

Human-Robot Trust Assessment Using Motion Tracking & Galvanic Skin Response

Kasper Hald¹, Matthias Rehm¹ and Thomas B. Moeslund¹

Abstract—In this study we set out to design a computer vision-based system to assess human-robot trust in real time during close-proximity human-robot collaboration. This paper presents the setup and hardware for an augmented reality-enabled human-robot collaboration cell as well as a method of measuring operator proximity using an infrared camera. We tested this setup as a tool for assessing trust through physical apprehension signals in a collaborative drawing task, where participants hold a piece of paper on a table while the robot draws between their hands. Midway through the test we attempt to induce a decrease in trust with an unexpected change in robot speed and evaluate subject motions along with self-reported trust and emotional arousal through galvanic skin response. After performing the experiment with forty participants, we found that reported trust was significantly affected when robot movement speed was increased. The galvanic skin response measurement were not significantly different between the test conditions. The motion tracking method used in this study did not suggest that subjects' motions were significantly affected by the decrease in trust.

I. INTRODUCTION

Human-robot collaboration (HRC) is becoming increasingly relevant as we are seeing the long-term consequences of repetitive manual production work [1], such as industrial meat production, in which the staff are especially affected [2]. Musculoskeletal diseases not only affects the quality of life for those suffering, but also makes production work less attractive for potential staff. We are working towards enabling close-proximity HRC, allowing the robot to relieve the collaborating worker, from here referred to as the operator, of heavy and repetitive actions while keeping the operator safe and feeling secure. To this end we are researching methods of assessing the operator's level of trust towards the robot partner in a non-obtrusive way in real-time.

We use the same definition of HRC as Herrmann and Leonards [3], where the robot and the operator works on the same component at the same time. In an HRC task where the operator has to use both hands for manipulation of objects, the system needs hands-free and non-obstructive methods of human-robot interaction (HRI). The goal is to develop a HRC cell, a collaborative human-robot work space, that enables both implicit and explicit communication from the operator to the robot using computer vision while displaying task-relevant information to the user with augmented reality (AR). Specifically, in this report we propose and evaluate a method of measuring and recording the operator's proximity



Fig. 1. Full HRC cell setup with Sawyer robot, projectors and infrared camera.

to the robot with a depth camera setup with a very small footprint, allowing the robot system to adapt accordingly. We aim to develop methods for inferring the changes in operator trust through reactive body postures as physical apprehension signals from changes in the robot's behaviours as a violation of the operator's explicit expectations. The full prototype and test setup is shown in Figure 1. The long-term goal is to develop non-obstructive solution that will fit into a production context, allowing robot system to interpret the operator's trust towards the robot based on proximity tracking. Using cross-referencing with the current shared objective, the aim is to have the system use the information to adapt to the operator by adjusting movement patterns or secondary communications methods, such as AR.

II. BACKGROUND

While body and posture tracking have been used and assessed in enabling safe HRC, it is rarely utilized for real-time trust assessment. Morato et al. [4] used a setup of multiple Kinect sensors for ensuring safe HRC in an environment for standing work while Tan & Arai [5] used a triple-camera setup for sedentary HRC. Both used skeleton tracking algorithms. Similarly, Hald et al. [6] used skeleton tracking for proximity tracking and trust assessment, showing correlation between user proximity and attitude towards the robot. In order to limit the physical footprint of our setup to allow for close rows of HRC cells we use a single depth camera pointed downwards, not allowing for effective skeleton tracking. Our long-term goal is to develop methods for tracking and for both standing and sedentary HRC.

¹All author are with The Department of Architecture, Design & Media Technology, Aalborg University, Rendsburggade 14, 9000 Aalborg, Denmark kh@create.aau.dk

*This work is funded in part by Innovation Fund Denmark

Lee & See [7] defined trust in automation as the expectation that the agent will achieve their goal in a situation characterized by uncertainty and vulnerability. This definition requires elaboration in order to include whether the trust is appropriate, which is derived from the relationship between the capabilities of the agent and the level of trust. Additionally, we have to consider the influence of the automation context as well as the goal-related characteristics of the agent. Lee & See proposed that trust in automation is created through a combination of analytic, analog and affective processes of external information and internal believe. We use these characteristics of trust in automation to define human-robot trust in HRC. We focus on trust as a the attitude towards the robot in the moment, rather than necessarily as a results of long-term interaction.

This research is relevant as enabling implicit communication has been shown to improve human-robot interaction [8] [9]. Also, Rani et al. [10] showed that human-aware motion planning systems improve the feeling of safety and comfort when used to adapt to the operator.

In a meta-analysis, Hancock et al. [11] found that trust factors in HRI were mainly influenced by the robot’s characteristics, in particular its performance, while environmental factors had moderate effect. Dragan et al. [12] tested the effect of a robot’s motion pattern on human trust using Hoffman’s metrics for fluency in human-robot collaboration [13] and found that predictable motions were more accepted by the operator than purely functional motions. Hoffman’s metrics were collected using a post-test questionnaire after the participants had been through all three conditions which includes questions pertaining to trust. In order to avoid the potential effects of delaying the questionnaire, we derived a shorter version of the trust metrics to be administered throughout the experiment.

Rani et al. [10] successfully used physiological measurement in an affect recognition system in the context of interacting with a remote robot. While physiological measurement might prove intrusive during daily operations, for our experiment we will use them to help verify our assessment methods.

III. PROXIMITY TRACKING

The setup for the human-robot collaboration cell is shown in Figure 1 and consists of a roughly two meter by two meter aluminum rig equipped with two projectors and an infrared (IR) camera. A seat, a work surface and the robot are arranged in the center of the rig. The dual-projection setup, with projectors positioned at either side and at a roughly forty-five degree angular offset from the work surface, enables projection-based AR as an output modality and is hard to fully occlude when reaching across the surface, as long as the projections are calibrated to match. The IR camera, an Orbbec Astra, specifically, is mounted at the top-center of the rig and pointed downwards toward the user. These types of camera allow for skeleton tracking, but this requires a lot of distance while facing the front of the user. This is problematic in this setup with the user seated and facing the

robot, so we have designed the cell to rely on the IR images only. The camera has a resolution of 320 by 240, but the outer edges never receive light, leaving the useful resolution at around 277 by 213.

A. Data Processing

In order to infer the proximity and posture of the user we make an aggregate of the IR camera frame to see how they reflect the light along the vertical axis of the frame. Examples of the infrared frames are shown in Figure 2 where a user is shown sitting upright and leaning back.

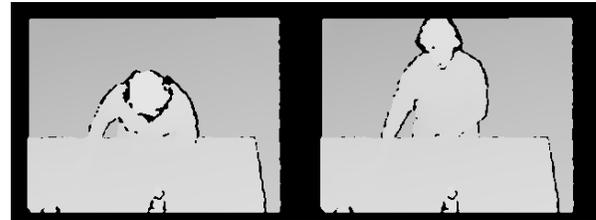


Fig. 2. Images from the top-down IR camera. Left: User sitting upright. Right: User leaning back.

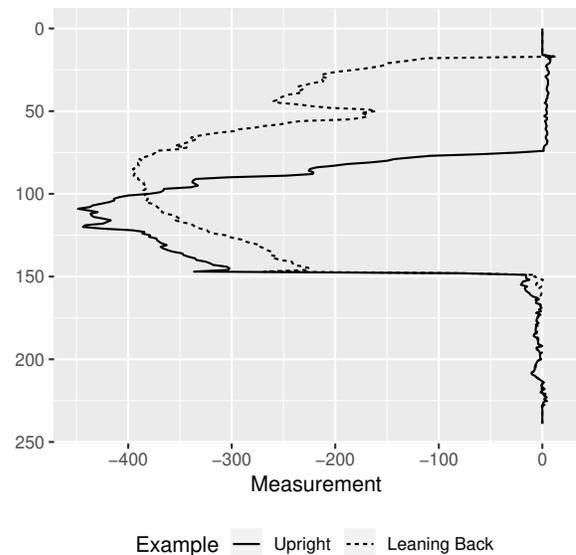


Fig. 3. Depth measurements aggregated from the top-down IR camera with background subtraction of a user, sitting upright and leaning backwards.

The first step of processing the frame is saving an empty background average of the work environment. The per-line average is from all the non-zero pixels in the line, as black pixels are areas with no reflected light and are considered noise. Subtracting the background averages from the per-line averages with user in the frame shows how much closer the user is to the camera than the background. Examples of these measurements are shown in Figure 3. In the example there is a visible difference in overall magnitude and spread, which can be used to infer the proximity and posture of the user by looking at the bounds and peaks of the curve. However, as can be seen in the examples, having the work

surface so close to the camera compared to the floor makes the measurements less sensitive in that part of the frame as the difference, such as from the user's arms, is lowered.

IV. EXPERIMENT

The objective of the experiments is to evaluate the top-down proximity tracking system as a tool to infer operator trust during close-proximity collaboration. To do this we have designed a collaborative drawing task in which the operator has to position a piece of A4 paper in a space on a table in front of the robot which is marked using the AR rig. The operator's role is to hold the paper down to the table as the robot, equipped with a felt pen in a 3D printed mount, moves in and draws a square on the paper, between the operator's hands. Midway through the experiment the robot's movement speed is changed without warning, changing pattern in order to provoke a decrease in trust. The aim is to determine if the operators proximity correlates with their trust towards the robot, which is assessed from participant arousal inferred from galvanic skin response (GSR) and with self-reporting using a questionnaire. In addition, we investigate whether the measures are affected by the operator's ability to hear the robot, as different movement speeds produce different motor noise, which may affect the operator's perception of the robot. For assessing the proximity tracking we test the following eight hypotheses:

- **H1:** The participants' reported trust toward the robot is significantly affected by changing velocity of the robot arm's movements and whether they are wearing ear protection.
- **H2:** The participants' movements are significantly affected by changing velocity of the robot arm's movements and whether they are wearing ear protection.
- **H3:** The participants' GSR response is significantly affected by the velocity of the robot arm's movements and whether they are wearing headphones.
- **H4:** The participants' movements and proximity to the robot correlate with their reported levels of trust toward the robot.

A. Procedure

At the beginning of the test the participants are presented with a printed consent form and description of the experiment. After signing the consent form the participants are sat at a table with the Sawyer robot facing them. They are then introduced to the task procedure: First they have to take a piece of paper from a stack on their left, which they position in a marked space on the table. Once the paper is positioned and they hold it down, the robot moves from its resting position and draws a square on the paper. This motion is activated manually by the test conductor. This is demonstrated during the introduction using the robot speed the participants starts with. This is to help the participant feel comfortable at the beginning of the test. In addition to marking the area for the paper, the AR rig is also used to

show the lines that the robot is going to draw. This is done with projected red lines, which the robot will draw over. This is to help inform the participants of the current status of the robot and when they can safely let go of the paper. Figure 4 shows a participants sitting in front of the robot, holding the paper as the robot draws along the projected red line.



Fig. 4. A participant sitting in front of the robot, holding the paper down. The paper is held in the area marked with a projected white rectangle, fitting an A4 sheet of paper. The robots holds the felt pen in a custom-designed 3D printed mount that is fitted with foam. The robot draws along the red line which is projected on the paper. To the participant's left is the stack of paper they can grab from and the tablet used to answer the questionnaire throughout the test.

Once the drawing is complete, the participant has to put the paper off to their right, after which they have to report their attitude and trust towards the robot using a questionnaire on a tablet to their left. The participants are instructed to state agreement to three statement on a scale between strongly disagree or strongly agree on sliding scales, yielding scores between 0 and 1, respectively. Based on the Hoffman's metrics on human-robot trust [13] the three statements are:

- I trusted the robot to do the right thing at the right time.
- I felt safe working next to the robot.
- The robot's reaching motion was surprising.

The participants are told to grab a new piece of paper and repeat the task until the test is over. Before the test, electrodes are attached to the back of the participants shoulder, opposite their dominant hand to limit disturbances, for measuring GSR with a Bluetooth-enabled device strapped to their upper arm. The GRS device infers the level of arousal in the participant by measuring the electric conductivity across the skin between the attached electrodes.

The task is repeated a total of twenty times, and after the first ten repetitions the robot movements speed is changed. In order to determine if the participants reacts to an increase in speed or rather to just a change in speed, half the participants start at a slow speed while the other start at high. The low speed is at a ratio of 0.2 of the robot's highest speed and the high speed is a ratio of 0.4. The test conditions, whether the participant is wearing ear and the robot's beginning speed, are counter-balanced with ten participants for each combination of conditions, leading to a total of forty participants.

V. RESULTS

A total of forty subjects participated in the experiment. 34 were male and 6 were female, 34 were right-handed, 6 were left-handed, and ages ranged from 21 to 28 age with and average age of 23.5 years.

Figure 5 shows the aggregated questionnaire for each task, grouped by condition along with confidence interval. The aggregates are the average responses to the three questions, each answered on a scale between 0 and 1, where the weight of the last question is inverted, so that that participant disagreement with the statement that the robot's motion was surprising will counts positively towards trust. The vertical lines marks the midway through the tests where the speed ratio was either increased or decreased. For the groups that started with the slow speed ratio we see that reported trust started high throughout the first half of the test, followed by a dip in trust when the speed was increased, as we would expect. The reported trust is then gradually recovered throughout the later half of the test. These effects are recognizable, though less pronounced, for the participants who started with high speed with a speed decrease in the middle. Still, this groups starts with lower trust towards the robot, which gradually builds up towards the halfway point.

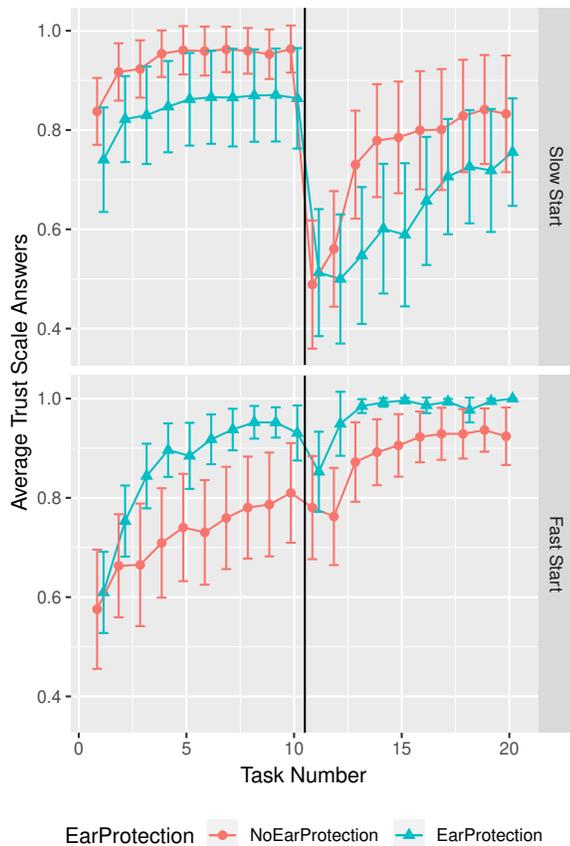


Fig. 5. Average reported trust as aggregated from the questionnaire answers throughout the test along with confidence interval.

Running the Shapiro-Wilk test on the reported trust scores, the tracking data and the GSR revealed that most data

groupings are not normally distributed. As such, all the data is treated as non-parametric. To evaluate hypothesis one we run a Wilcoxon rank sum test on the reported trust scores immediately before and after the change in robot movement speed. The test shows significant difference for participants for whom speed was increase, both with ($W = 88, p < .01$) and without ear protection ($W = 100, p < .01$), while not significant for participants who experienced a decrease in speed.

To compare the robot start speed and ear protection conditions we compare the pairs of trust scores for the tenth and eleventh tasks, separately. The Wilcoxon rank sum test showed significant difference in trust scores between the speed conditions after the speed change, both with ($W = 11, p < .01$) and without ear protection ($W = 12, p < .01$), while there were no significant differences before the speed change. There were no significant effects from the ear protection in any condition. From these results we can retain hypothesis one in that unforeseen changes in robot movement speed affects reported trust, but only for increases in speed, and the ability to hear the robot motors has insignificant effect.

The participants' movement and proximity between the conditions are shown in Figure 6 and Figure 7. The proximity is measured as the position of the participant's highest point, usually the top of the head, along the vertical axis of the depth image, measured in pixels. We measure the participants' movement reactions to the robot as the delta changes in proximity within the first second of the robot moving to draw. Figure 6 shows the average delta movement among participants for each task, showing whether the participants overall moved away or towards the robot.

Figure 7 shows the average absolute. The difference is that this is the total movement of the participant, regardless of whether they are moving away or towards the robot, showing how much the participants move in general. Looking at both Figures, we see no obvious tendencies among the conditions, whether it being in the first or later half or right before and after the changes in robot speed ratio.

When comparing both the delta and absolute movement between conditions similarly to how we did with the reported trust scores, the Wilcoxon rank sum test yield no significant differences, regardless of data groupings. Due to the lack of significant difference and the inconsistencies we cannot reject the null hypotheses for hypothesis two.

Figure 8 shows the average normalized GSR measurement between participants and split between conditions. After noise removal the data is normalized by fitting the range of reading for each participant between zero and one. As such, the confidence intervals increase as the tests go on as most participants start the test with high resistance across the skin, placing them close to the value one, with the resistance decreasing at different rates throughout the tests.

To test the curves of the normalized GSR we perform a Wilcoxon rank sum test between the four conditions at the midpoint of the experiment. This yielded no significant differences. As such we cannot reject the null hypotheses for hypothesis three. Lastly, performing both Kendall's and

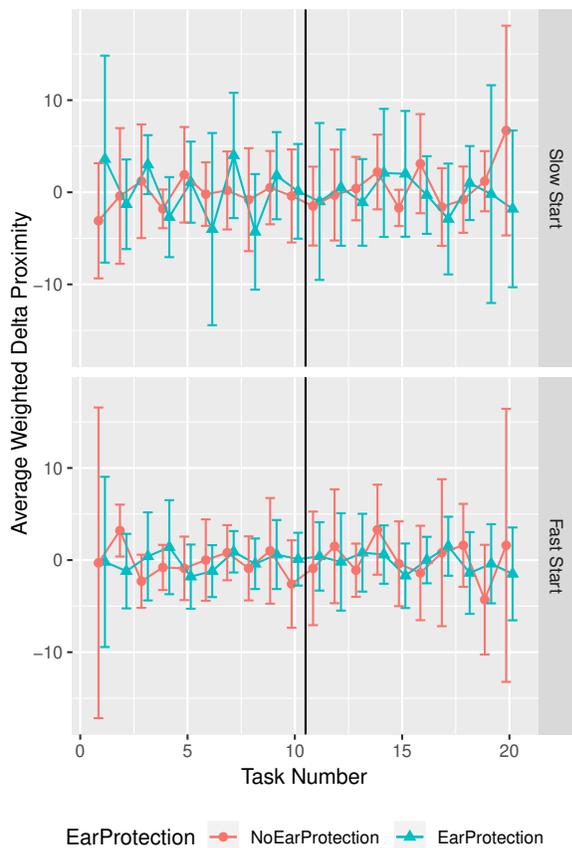


Fig. 6. Average participant motion for the first second of robot movement weighted by movement direction along with confidence interval. The vertical lines mark the midway point of the test where the robot speed ratio changes.

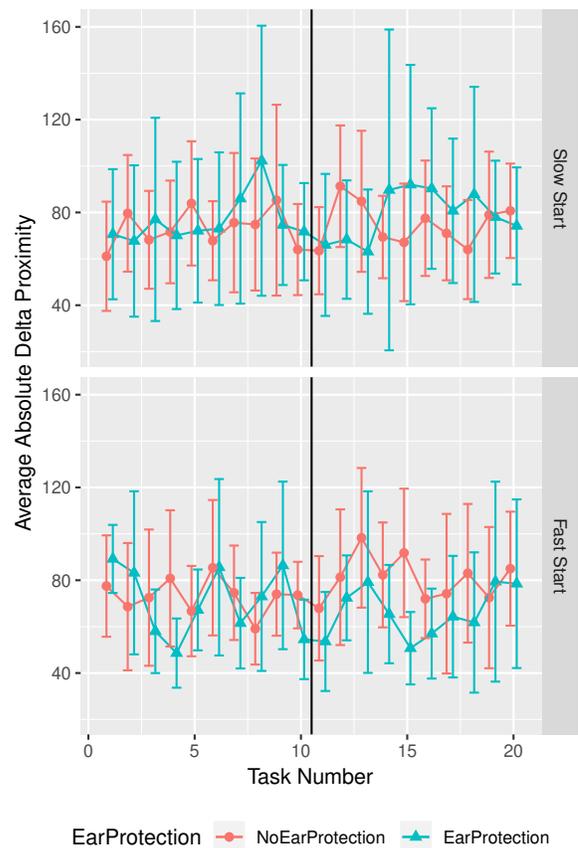


Fig. 7. Average absolute participant motion for the first second of robot movement along with confidence interval. The vertical lines mark the midway point of the test where the robot speed ratio changes.

Spearman's rank correlations, neither the weighted delta movements nor the absolute movements showed significant correlation with the reported trust scores, meaning we cannot reject the null hypothesis for hypothesis four.

VI. DISCUSSION

The trust score results indicate a decrease in trust towards the robot, as is to be expected from the experiment design, where trust were only significantly affected by increases in speed. However, our goal to verify it as a measurement using GSR as a measure of arousal was not met in this study. This may be due to the measurements not being sensitive enough for the arousal experienced, but it may also be a flaw in the procedure, as the electrodes may not have had enough time to warm up and level out before starting the tasks. Future experiments will be started with a warm-up period as well as a period for taking baseline measurements.

The method of motion tracking used in this experiment is not an effective indicator of trust through physical apprehension signals. This may be due to low sensitivity from the low camera resolution, but it may also be an issue with the nature of the task, in that having the participants be sedentary and holding down the paper may inhibit movement. Alternatively, only looking at the movement of the peak position of the

user is not enough to indicate movement or apprehension signals. In follow-up studies we will investigate alternative data processing methods that takes advantage of all the data we're collecting. In addition, aggregating the IR frame data along the horizontal axis may reveal more physical signals. We can also look into operator proximity to the robot throughout the test as they place the paper in order to infer trust. In addition, it could be beneficial to design a shared task where the participant is standing, allowing for more motion as they are not pinned down by having to sit in a chair.

For this experiment we focused on changes in robot speed for inducing a decrease in trust, but for real world application it would be valuable to look into more subtle signs of system error. One possibility with our current setup would be to either remove or alter the AR projection that shows what the robot is doing, which may induce more uncertainty and thereby mistrust in the system. A problem with changing the speed is that the participants may expect that something unexpected may happen due to the experiment context, so simulating subtle system errors may do less to immediately affirm their suspicion. With a longer experiment it would be interesting to look into the recovery in trust that we see in the later half of the test. As the participants know that changes

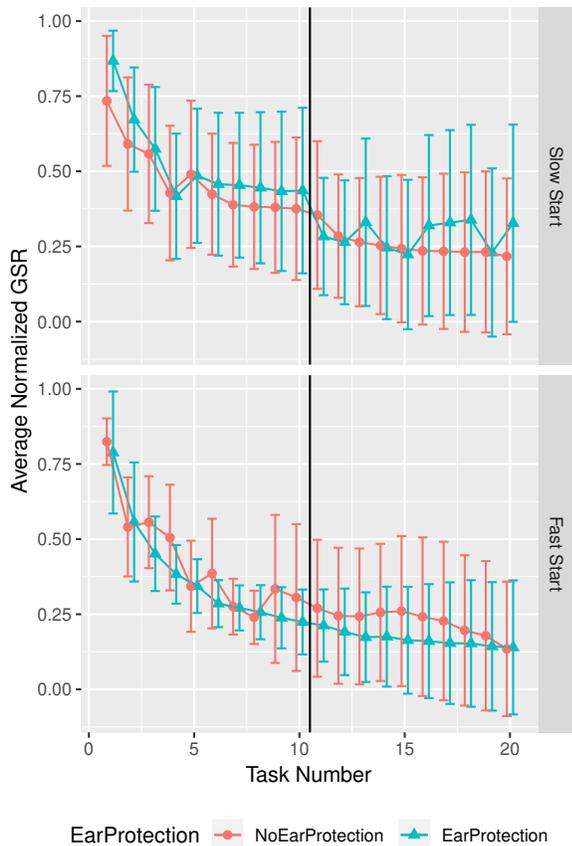


Fig. 8. Average normalized GSR at the beginning of each task along with confidence interval. The vertical lines mark the midway point of the test where the robot speed ratio changes.

can happen, they may expect a break in movement patterns to happen again, which may affect trust recovery.

VII. CONCLUSION

We aimed to develop a system to assess human-robot trust in real time during close-proximity HRC. Using a top-down IR depth camera we aggregated the frame data to measure the operator's proximity to the robot in order to infer trust towards the robot from physical apprehension signals.

We tested this setup as a tool for assessing operator trust based on reactions to a sudden change in robot movement speed in order to provoke a disruption of expectations. We looked at both increases and decreases in speed, as well as participants with and without ear protection to see if motor noise from the robot on speed changes have any effect. To determine the effects on operator trust we assessed the subjects attitudes towards the robot using self-reporting through questionnaires and emotional arousal from GSR. After performing the experiment with forty participants, we found that reported trust towards the robot was significantly affected when the robot's movement speed was unexpectedly increased. This was not the case for speed decreases. Wearing ear protection did not yield any significant difference, suggesting little effect from the motor noises. The GSR

measurement were not significantly different between the test conditions, which may be due to an insufficient warm-up period. The analyses for the motion tracking method used in this study did not suggest that the participants' motions were significantly affected by a decrease in trust. The method was based on tracking the position of the point of the participant closest to the camera. For future studies we will work with the data we collected to design a data processing method that takes better advantage of the amount of data collected in order to obtain a more sensitive motion measure. We will also re-design the shared task, allowing the participant to stand up, allowing for more motion.

REFERENCES

- [1] L. Punnett and D. H. Wegman, "Work-related musculoskeletal disorders: the epidemiologic evidence and the debate," *Journal of Electromyography and Kinesiology*, vol. 14, no. 1, pp. 13–23, feb 2004.
- [2] Arbejdstilsynet Denmark, "Anmeldte erhvervs sygdomme 2012-2017," 2017. [Online]. Available: <https://at.dk/media/4814/bilag-b-tabeller-med-anonymitet-final.pdf>
- [3] G. Herrmann and U. Leonards, *Human-Humanoid Interaction: Overview*. Dordrecht: Springer Netherlands, 2016, pp. 1–16.
- [4] C. Morato, K. N. Kaipa, B. Zhao, and S. K. Gupta, "Toward Safe Human Robot Collaboration by Using Multiple Kinects Based Real-Time Human Tracking," *Journal of Computing and Information Science in Engineering*, vol. 14, no. 1, p. 011006, jan 2014.
- [5] J. T. C. Tan and T. Arai, "Triple stereo vision system for safety monitoring of human-robot collaboration in cellular manufacturing," in *2011 IEEE International Symposium on Assembly and Manufacturing (ISAM)*. IEEE, may 2011, pp. 1–6. [Online]. Available: <http://ieeexplore.ieee.org/document/5942335/>
- [6] K. Hald, M. Rehm, and T. B. Moeslund, "Proposing human-robot trust assessment through tracking physical apprehension signals in close-proximity human-robot collaboration," in *The 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN) IEEE International Symposium on Robot and Human Interactive Communication*. IEEE Computer Society Press, 2019.
- [7] J. D. Lee and K. A. See, "Trust in Automation: Designing for Appropriate Reliance," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 46, no. 1, pp. 50–80, jan 2004. [Online]. Available: <http://hfs.sagepub.com/cgi/doi/10.1518/hfes.46.1.50.30392>
- [8] J. Greenstein and M. Revesman, "Two Simulation Studies Investigating Means of Human-Computer Communication for Dynamic Task Allocation," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 16, no. 5, pp. 726–730, sep 1986. [Online]. Available: <http://ieeexplore.ieee.org/document/4075638/>
- [9] P. A. Lasota and J. A. Shah, "Analyzing the Effects of Human-Aware Motion Planning on Close-Proximity Human–Robot Collaboration," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 57, no. 1, pp. 21–33, feb 2015. [Online]. Available: <http://journals.sagepub.com/doi/10.1177/0018720814565188>
- [10] P. Rani, N. Sarkar, C. A. Smith, and L. D. Kirby, "Anxiety detecting robotic system towards implicit human-robot collaboration," *Robotica*, vol. 22, no. 1, pp. 85–95, jan 2004. [Online]. Available: http://www.journals.cambridge.org/abstract_S0263574703005319
- [11] P. A. Hancock, D. R. Billings, K. E. Schaefer, J. Y. C. Chen, E. J. de Visser, and R. Parasuraman, "A Meta-Analysis of Factors Affecting Trust in Human-Robot Interaction," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 53, no. 5, pp. 517–527, oct 2011. [Online]. Available: <http://journals.sagepub.com/doi/10.1177/0018720811417254>
- [12] A. D. Dragan, S. Bauman, J. Forlizzi, and S. S. Srinivasa, "Effects of Robot Motion on Human-Robot Collaboration," in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI '15*. New York, New York, USA: ACM Press, 2015, pp. 51–58. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2696454.2696473>
- [13] G. Hoffman, "Evaluating Fluency in Human-Robot Collaboration," in *Robotics: Science and Systems*, 2013. [Online]. Available: <http://guyhoffman.com/publications/HoffmanRSS13Workshop.pdf>