The Six Hug Commandments: Design and Evaluation of a Human-Sized Hugging Robot with Visual and Haptic Perception

Alexis E. Block
MPI-IS and ETH Zürich
Stuttgart, Germany

Sammy Christen
ETH Zürich
Zürich, Switzerland

Roger Gassert
ETH Zürich
Zürich, Switzerland

Otmar Hilliges
ETH Zürich
Zürich, Switzerland

Katherine J. Kuchenbecker
MPI for Intelligent Systems
Stuttgart, Germany

ABSTRACT

Receiving a hug is one of the best ways to feel socially supported, and the lack of social touch can have severe negative effects on an individual's well-being. Based on previous research both within and outside of HRI, we propose six tenets ("commandments") of natural and enjoyable robotic hugging: a hugging robot should be soft, be warm, be human sized, visually perceive its user, adjust its embrace to the user's size and position, and reliably release when the user wants to end the hug. Prior work validated the first two tenets, and the final four are new. We followed all six tenets to create a new robotic platform, HuggieBot 2.0, that has a soft, warm, inflated body (HuggieChest) and uses visual and haptic sensing to deliver closed-loop hugging. We first verified the outward appeal of this platform in comparison to the previous PR2-based HuggieBot 1.0 via an online video-watching study involving 117 users. We then conducted an in-person experiment in which 32 users each exchanged eight hugs with HuggieBot 2.0, experiencing all combinations of visual hug initiation, haptic sizing, and haptic releasing. The results show that adding haptic reactivity definitively improves user perception a hugging robot, largely verifying our four new tenets and illuminating several interesting opportunities for further improvement.

CCS CONCEPTS

• **Computer systems organization** → **Robotics**.

1 INTRODUCTION

Hugging has significant social and physical health benefits for humans. Not only does a hug help lower blood pressure, alleviate stress

Figure 1: A user hugging HuggieBot 2.0.

and anxiety, and increase the body's levels of oxytocin, but it also provides social support, increases trust, and fosters a sense of community and belonging [6]. Social touch in a broader sense is also vital for maintaining many kinds of relationships among humans and primates alike [32]; hugs seem to be a basic evolutionary need. They are therefore highly popular! An online study conducted in 2020 polled 1,204,986 people to find out "what is the best thing?" Hugs earned fifth place out of 8,850 things, behind only sleep, electricity, the Earth's magnetic field, and gravity [27]. The absence of social touch can have detrimental effects on child development [4]. Unfortunately, ever more interactions are happening remotely and online, especially during this unprecedented time of physical distancing due to COVID-19. An increasing number of people are suffering from loneliness and depression due to increased workload and population aging [21, 22]. Our long-term research goal is to determine the extent to which we can close the gap between the virtual and physical worlds via *hugging robots that provide high-quality social touch.*

Making a good hugging robot is difficult because it must understand the user's nonverbal cues, realistically replicate a human hug, and ensure user safety. We believe that such robots need multimodal perception to satisfy all three of these goals, a target no previous system has reached. Some approaches focus primarily on safety, providing the user with the sensation of being hugged without being able to actively reciprocate the hugging motion [11, 33, 36]. Conversely, other researchers focus on providing the user with an

item to hug, but that item cannot hug the user back [10, 30, 31]. Other robotic solutions safely replicate a hug, but they are teleoperated, meaning they have no perception of their user and require an additional person to control the robot any time a user wants a hug [16, 28, 42]. Finally, some robots have basic levels of perception but are not fully autonomous or comfortable [3, 20]. Section 2 further details prior research in this domain.

To tackle the aforementioned goal of safely delivering pleasant hugs, we propose the *six tenets ("commandments") of robotic hugging*: a hugging robot should be soft, warm, and sized similar to an adult human, and it should see and react to an approaching user, adjust automatically to that user's size and position while hugging, and reliably respond when the user releases the embrace. After presenting these tenets and our accompanying hypotheses in Section 3, we use the tenets to inform the creation of HuggieBot 2.0, a novel humanoid robot for close social-physical interaction, as seen in Fig. 1 and described in Section 4. HuggieBot 2.0 uses computer vision to detect an approaching user and automatically initiate a hug based on their distance to the robot. It also uniquely models hugging after robot grasping, using two slender padded arms, an inflated torso, and haptic sensing to automatically adjust to the user's body and detect user hug initiation and termination. HuggieBot 2.0 is the *first human-sized hugging robot with visual and haptic perception for closed-loop hugging*.

We then seek to validate the four new tenets by conducting two experiments with HuggieBot 2.0. First, we confirmed user preference for the created platform's physical size, visual appearance, and movements through a comparative online study, as described in Section 5. We then conducted an in-person study (Section 6) to understand how HuggieBot 2.0 and its three new perceptual capabilities (vision, sizing, and release detection) affect user opinions. Section 7 discusses the study results, which show that the six tenets significantly improve user perception of hugging robots. Section 8 discusses the limitations of our approach and concludes the paper.

2 RELATED WORK

Using Vision for Person Detection. One challenge of accurate and safe robotic hugging is detecting a user's desire for a hug. Many researchers solve this problem by using a remote operator to activate the hug [10, 31, 42, 43]. This form of telehug is not a universal approach because it requires a hugging partner to be available at the exact moment a user would like the comfort of a hug. Having the user press a button is a simpler alternative but differs greatly from human-human hugging. One method that could allow robots to provide hugs autonomously is to detect an approaching user via computer vision. Human detection has long been of interest in many research fields, including autonomous driving [44], surveillance and security [25], computer vision [38], and human-robot interaction [40]. Early works focus mostly on finding a representative feature set that distinguishes the humans in the scene from other objects. Different methods for the feature extraction have been proposed, such as using Haar wavelets [24], histograms of oriented gradients (HOG) [8], and covariance matrices [37]. Several approaches try to combine multiple cues for person detection; for example, Choi et al. [5] and Vo et al. [40] combine the Viola-Jones face detector [39] and an upper-body detector based on HOG features [8]. In computer

vision, deep-learning-based systems relying on convolutional neural networks (CNNs) are often used for person detection. However, the computational cost of these detection pipelines is very high, so they are not suited for use on a real-time human-robot interaction platform. The recent decrease of the computational cost of improved models, e.g., [18], has facilitated their adaptation to robot platforms. Hence, we integrate such a model into our pipeline so that HuggieBot 2.0 can recognize an approaching user to initiate a hug with minimal on-board computational load.

Hugging as a Form of Grasping. Once the user arrives, safely delivering a hug is challenging for robots because users come in widely varying shapes and sizes and have different preferences for hug duration and tightness. No existing hugging robots are equipped to hug people adaptively. We propose looking to the robotics research community to find a solution. Grasping objects of varying shape, size, and mechanical properties is a common and well-studied problem, e.g., [1, 7, 19, 26]. Therefore, we look at hugging as a large-scale two-fingered grasping problem, where the item to be grasped is a human body. For example, the BarrettHand, a commercially available three-fingered gripper, automatically adjusts to securely grasp widely varying objects by closing all finger joints simultaneously and then stopping each joint individually when that joint's torque exceeds a threshold [34].The robot arms used for HuggieBot 2.0 have torque sensors at every joint, making this torque-thresholded grasping method an ideal way to achieve a secure embrace that neither leaves air gaps nor applies excessive pressure to the user's body. Torque sensors can also enable the robot to feel when the user wishes to leave the embrace.

Previous Hugging Robots. In recent years, there has been a dramatic increase in the number of researchers developing and studying hugging robots. High press coverage shows this topic is of great interest to the general public. Interestingly, researchers are taking many different approaches to create robotic systems that can receive and/or give hugs. Smaller hugging robots have been created to provide comfort, but they typically cannot actively hug the user back. The Huggable is a small teddy bear to accompany children during long stays in the hospital [30]. The Hug is a pillow whose shape mimics a child wrapping his or her limbs around an adult [10]. Hugvie is also a pillow users can hug; a cellphone inside lets the user talk to a partner while he or she hugs the pillow [31, 43]. Their small size makes these huggable systems inherently safer than larger devices, but it also prevents them from providing the benefits of social touch to the user because they cannot administer deep touch pressure therapy [12].

Teleoperated hugging robots have also been created, and they can be closer to human size. Some research groups focus on non-anthropomorphic solutions like using a large panda or teddy bear stuffed animal to hide the mechanical components [16, 28, 29]. These robots all require that either an operator or partner is available at the time any user wants a hug. They also may not be very comfortable for the user because the robots are unable to stand on their own; the user must crouch or crawl to get a hug from the robot. Disney also patented a teleoperated hugging robot to be similar to its famous character Baymax [42], though a physical version has not been reported. Negatively, none of these robots seem to have the ability to adjust their embrace to the size or location of the user.

In addition to creating an appropriately sized hugging robot, all of these researchers covered the robot's rigid components with soft materials to create an enjoyable contact experience for their users. In contrast, Miyashita and Ishiguro [20] previously created a robot that measures the distance between the robot and user with range sensors to initiate the hug sequence; this robot has a hard surface that may be uncomfortable to touch, and its inverted pendulum design appears to use the human to balance during the hug.

Block and Kuchenbecker added padding, heating pads, and a soft tactile sensor to a Willow Garage Personal Robot 2 (PR2) to enable it to exchange hugs with human users [2, 3]. The experimenter manually adjusted the robot to match the height and size of each user, a process that takes time and prevents spontaneous hugging. She also initiated every hug for the user. The PR2 was successful in adapting to the user's desired hug duration through the use of the tactile sensor, but the user had to place his or her hand in a specific location on the PR2's back, which was not natural for all users, and the user had to press this sensor to tell the robot to release them. Some users also criticized the size and shape of this robot as being awkward to hug. On the positive side, this study showed that both softness and warmth are important for a robot to deliver good hugging experiences; we therefore incorporate these already validated elements as our first two tenets. From Trovato et al. [35] we learned that softness alone is not enough for a hugging robot; people are more receptive to hugging a robot that is wearing clothing, so we placed suitable clothes on HuggieBot 2.0.

3 HUGGING TENETS AND HYPOTHESES

We propose six tenets to guide the creation of future hugging robots. The first two were validated by Block and Kuchenbecker [3], and the other four are proposed and validated in this paper. We believe that a hugging robot should:

T1. be soft,
T2. be warm,
T3. be sized similar to an adult human,
T4. visually perceive and react to an approaching user, rather than requiring a triggering action such as a button press by the user, an experimenter, or a remote hugging partner,
T5. autonomously adapt its embrace to the size and position of the user's body, rather than hug in a constant manner, and
T6. reliably detect and react to a user's desire to be released from a hug regardless of his or her arm positions.

Building off the previously described research, this project seeks to evaluate the extent to which the six tenets benefit user perception of hugging robots. Specifically, we aim to test the following three hypotheses:

H1. When viewing from a distance, potential users will prefer the embodiment and movement of a hugging robot that incorporates our four new tenets over a state-of-the-art hugging robot that violates the tenets.
H2. Obeying the four new tenets during physical user interactions will significantly increase the perceived safety, naturalness, enjoyability, intelligence, and friendliness of a hugging robot.
H3. Repeated hugs with a robot that follows all six tenets will improve user opinions about robots in general.

4 SYSTEM DESIGN AND ENGINEERING

We introduce a new human-sized hugging robot with visual and haptic perception. This platform is designed to have a friendlier and more comfortable appearance than previous state-of-the-art hugging robots. Building off feedback from users of prior robots (Section 2), we focused on six areas for the design of this new, self-contained robot: the frame, arms, inflated sensing torso, head and face, visual person detection, and software architecture.

Frame. HuggieBot 2.0's core consists of a custom stainless steel frame with a v-shaped base. The robot's height can be manually adjusted if needed, and the shape of the base allows users to come very close to the robot, as seen in Fig. 1. The user does not need to lean over a large base to receive a hug, unlike the PR2-based hugging robot [3]. The v-shaped base also increases the safety and stability of the robot by acting to counteract any leaning force imparted by a user. This large base and counterweight ensure that even a large user approaching at a high speed intending to make an incorrect impact with the robot will not flip it over, inflict injury upon themselves, or cause damage to the robot.

Arms. Two 6-DOF Kinova JACO arms are horizontally mounted to a custom stainless steel bracket attached to the top of the metal frame. To create a more approachable appearance, the grippers of the JACO arms were removed, and large padded mittens were placed over the wrist joints that terminate each arm. The arms are controlled by commanding target joint angles; movement is quiet, and the joints are not easily backdrivable when powered. The torque sensors at each joint are continually monitored so that hugs can be automatically adjusted to each user's size and position. The second joint (shoulder pan) and third joint (elbow flex) on each arm stop closing individually when they surpass a torque threshold, which we empirically set to 10 Nm and 5 Nm, respectively. The joint torques are also used to detect when a user is pushing back against the arms with a torque higher than 20 Nm, indicating his/her desire to be released from a hug. To create a comfortable and enjoyable tactile experience for the user, we covered the arms in soft foam and a sweatshirt.

HuggieChest: Inflatable Torso. We developed a simple and inherently soft inflatable haptic sensor to serve as the torso of our hugging robot, as pictured in Fig. 2. The torso was constructed by both heat sealing and gluing (with HH-66 vinyl cement) two sheets of PVC vinyl to create an airtight seal. This chest has one chamber located in the front and another in the back. There is no airflow between the two chambers. Each chamber has a valve from an inflatable swim armband to inflate, seal, and deflate the chamber. Inside the chamber located on the back of the robot is an Adafruit BME680 barometric pressure sensor and an Adafruit electret microphone amplifier MAX4466 with adjustable gain. Both sensors were secured in the center of the chest on the inner wall of the chamber. The two sensors are connected to a single Arduino Uno micro-controller outside the chamber. The microphone and pressure sensor are sampled at 45 Hz, and the readings are sent over serial to the HuggieBot 2.0 computer for real-time processing. We originally tested with the same sensing capabilities in both chambers but did not find the information from the front chamber to be very useful; thus, no data are collected from the front chamber. This novel inflatable haptic sensor is called the HuggieChest.

Figure 2: The inflated HuggieChest when lying flat. The two air chambers form the front and back of the robot's torso.

The HuggieChest's shape was created by following a pattern of a padded vest that goes over the wearer's head and is secured with a belt around the waist. Because the HuggieChest is heat-sealed at the shoulders to stop airflow between the chambers and allow the chest to bend once inflated, the robot is unable to feel contacts in these locations. The HuggieChest is placed directly on top of the metal frame of HuggieBot 2.0. On top of the inflatable torso, we put two Thermophore MaxHeat Deep-Heat Therapy heating pads (35.6 cm × 68.6 cm), which are attached together at one short edge with two shoulder straps. The sweatshirt is placed on top of the heating pads to create the final robot torso.

Head and Face. We designed and 3D-printed a head to house a Dell OptiPlex 7050 minicomputer that controls the robot, a face screen, an RGB-D camera, the Arduino from the HuggieChest, a wireless JBL speaker, and cables. The head splits into two halves with a rectangular plate on each side that can be removed to access the inside. The final piece of the head is the front-facing frame, which secures the face screen and camera. The face screen is an LG LP101WH1 Display 10.1" LCD screen with a 1366 × 768 resolution in portrait orientation. The screen displays faces based on designs created and validated for the Baxter robot [13].

Vision and Person Detection. We use a commercially available Intel Realsense RGB-D camera with custom software to recognize an approaching person and initiate a hug. To this end, we integrate a deep-learning-based person detection module into our pipeline. The module consists of two parts. First, our software recognizes an approaching person using an open-source Robot Operating System (ROS) integration [23] of Tensorflow's object detection library, which is based on the SSD mobilenet model [18]. In the next step, we utilize the camera's depth sensor to estimate the distance of the person to the robot. We use a sliding window to ensure a person is actively approaching the robot; we observe the distance measured from the depth sensor, which can be noisy, and check whether the mean distance decreases. This strategy prevents undesired hugs in case a person walks away from the robot. Once the person is actively approaching the robot, we initiate a hug as soon as a tuned distance threshold of 2.45 m is passed. This threshold was selected as it informs the robot the person is attempting to move from the social space into the robot's personal space [15].

Robot Software Architecture. The robot is controlled via ROS Kinetic. Each robot arm joint has both angle and torque sensors. A PID controller is used to control each joint angle over time. The robot arms begin by moving to a home position. The camera module starts and waits for an approaching user. Upon detection, the robot

asks the user, "Can I have a hug, please?" as in [2], while the robot's face changes to an opening and closing mouth. The specific hug it is supposed to run (with or without haptic sizing and release) is executed by commanding each joint to move at a fixed angular velocity toward a predetermined goal pose. For hugs without haptic sizing, the robot hugs in a one-size-fits-most manner, where the robot's second and third joints each close by 20°. This pose was large enough such that it did not apply high forces to the bodies of any of our users; it was not adjusted for different subjects. For hugs with haptic sizing, the robot arms move toward a pose sized for a small user; we continually monitor each joint torque and stop a joint's movement if it exceeds the pre-set torque threshold. This method leads to automatic adjustment to the user's size (T5).

The Arduino communicates the microphone and pressure sensor data from inside the back chamber of the HuggieChest to ROS over serial at 45 Hz. This data stream is analyzed in real time. The program first determines the ambient pressure and noise in the chamber by averaging the first 20 samples to create a baseline that accounts for different levels of inflation and noise. The chamber detects the user beginning to hug when the chamber's pressure increases by 50 kPa above the baseline pressure. Contact is determined to be broken, thus indicating that the user wants to be released, when the pressure returns to 10 kPa above the baseline pressure. The measured torques from the robot's shoulder pan and elbow flex joints are monitored continually during a hug. A haptic release is also triggered when any of these torques surpasses a threshold limit of 20 Nm. For a timed hug, rather than detecting the instant at which the user wants to be released, the robot waits 1 second after the arms fully close before releasing the user, so it is apparent to the user they are not in control of the duration of the hug. Overall, our proposed method of closed-loop hugging works on a higher level of abstraction than the low-level control, i.e., including both visual and haptic perception in the loop of the hugging process. The robot's haptic perception is two-fold: adjusting to the size of the user and sensing when he/she wants to be released.

5 ONLINE USER STUDY

We ran an online study to get feedback from a broad audience on the embodiment and movement of our robot as part of our user-centered design process, and to compare it to the PR2-based HuggieBot 1.0 [3]. The main stimuli were two videos from [3] along with two newly recorded videos of people hugging HuggieBot 2.0 with matched gender, enthusiasm, and timing; these videos are included as supplementary material for this paper. This study was approved by the Max Planck Society ethics council under the HI framework.

Participants. All participants for the online study were non-compensated English-speaking volunteers recruited via emails and social media. A total of 117 subjects took part in the online survey: 42.7% male, 56.4% female, and 0.9% who identify as other. The participants ranged in age from 20 to 86 (M = 37.5, SD = 16.75). The majority of respondents indicated they had little (30.7%) or no experience (43.6%) interacting physically with robots.

Procedures. After someone reached out to the experimenter and indicated interest in participating in the online study, the experimenter sent the subject an informed consent document by email.

Table 1: The questions participants answered after viewing or experiencing robot hugs.

This hug made the robot seem (unfriendly – friendly)
This robot behavior seemed (unsafe – safe)
This hug made the robot seem socially (stupid – intelligent)
This hug interaction felt (awkward – enjoyable)
This robot behavior seemed (unnatural – natural)
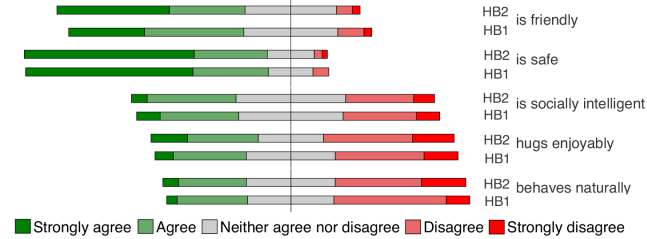
Figure 3: The responses to the five questions asked after users watched two videos of people hugging HuggieBot 2.0 (HB2) and HuggieBot 1.0 (HB1).

The participant read it thoroughly, asked any questions, and signed it and sent it back to the experimenter only if they wanted to participate. At this point, the user was assigned a subject number and sent a unique link to the online survey.

First, participants filled out their demographic information. Next, they were shown two videos of an adult (one male and one female) hugging a robot labeled "Robot A". Robot A was HuggieBot 2.0 for half of the participants and Block and Kuchenbecker's HuggieBot 1.0 [3] for the other half. Users could watch these videos as many times as they liked before answering several questions. Users described their first impressions of the robot they saw. Then, they answered the questions shown in Table 1 on a 5-point Likert scale. Afterwards, there was an optional space for additional comments. Next, they were shown two videos of people hugging the other robot labeled "Robot B" under the same conditions as the first videos. Users answered the same questions for Robot B as they did for Robot A. Finally, since the videos were shot from behind the robot, users were shown frontal images of both robots posed in a similar manner. Participants were then asked "In what ways is Robot A better than Robot B?" and "In what ways is Robot B better than Robot A?" Then, they were asked to select which robot they would prefer to hug, Robot A or B, and why.

After the first 40 participants, we noticed that several users were commenting on the purple fuzzy appearance of the PR2, rather than on the robots themselves. Therefore, we added a new section at the end of the survey, which was completed by the remaining 77 participants. This new section showed photographs of Robot A and Robot B plus two additional photographs showing HuggieBot 2.0 in different clothing. In one of the new photographs, HuggieBot 2.0 wore a fuzzy purple robe similar to the fabric cover on the PR2, and in the other it wore its gray sweatshirt over this same fuzzy purple robe. Participants were asked which of these four robots they would prefer to hug and why.

Results. The responses to the five Likert-style questions from Table 1 can be seen in Fig. 3. For all statistical analyses, we applied a Bonferroni alpha correction to $\alpha = 0.05$ to determine significance
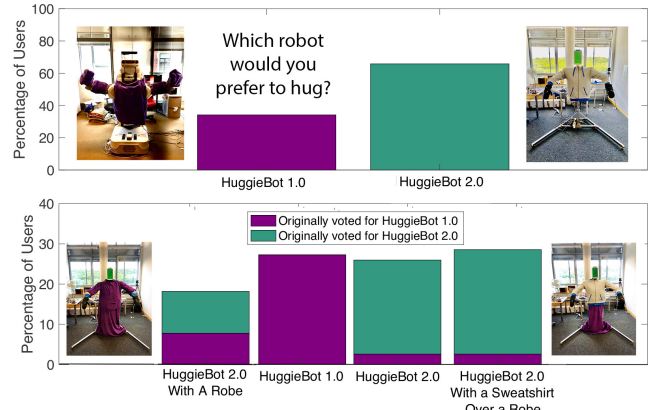
Figure 4: The breakdown of preferences when users had two choices (top) and four choices (bottom), with the associated images. The colors of the second plot show which robot the user preferred in the first selection round.

and to account for the multiple comparisons. Because the data from these questions were non-parametric, we conducted a Wilcoxon signed-rank test. No statistically significant differences were found between the responses to any of these questions for the two robots.

The responses to the first and second rounds of voting for which robot users would prefer to hug can be found in Fig. 4. We ran a Wilcoxon signed-rank test for the first round of voting and determined users would significantly prefer to hug our new robot over HuggieBot 1.0 ($p < 0.001$). In the second voting round, no significant preference was found between any of the four options, indicating HuggieBot 2.0 was preferred over HuggieBot 1.0 approximately 3:1. We ran several one-way analyses of variance (ANOVA) to see if participant gender, robot presentation order, or participant level of extroversion had a significant effect on which robot the user selected. No significance was found in any of these cases.

Changes to Platform. After analyzing the results of the online study and reviewing the user feedback, we found several areas to improve our hugging robot. Some users found the initial HuggieBot 2.0 voice off-putting, so we changed the robot's voice to sound less robotic. We made the purple robe and the sweatshirt the robot's permanent outfit, as it had the highest number of votes and received many positive comments. We changed the color of the robot's face from its initial green to purple to match the robe and create a more polished look. Several users commented on the slow speed of HuggieBot 2.0's arms. Since the arm joints cannot move faster than the maximum angular velocity specified by the manufacturer, we instead moved their starting position closer to the goal position to reduce the time they need to close.

6 IN-PERSON USER STUDY

The goal of the in-person study was to evaluate our updated robotic platform and the four new hugging tenets that drove its design. This study was also approved by the Max Planck Society ethics council under the HI framework.

Participants. The recruitment methods for the in-person study were the same as for the online study. Participants not employed by the Max Planck Society were compensated 12 euros. A total

Table 2: The fifteen questions asked in the opening and closing questionnaires.

I feel understood by the robot
I trust the robot
Robots would be nice to hug
I like the presence of the robot
I think using the robot is a good idea
I am afraid to break something while using the robot
People would be impressed if I had such a robot
I could cooperate with the robot
I think the robot is easy to use
I could do activities with this robot
I feel threatened by the robot
This robot would be useful for me
This robot could help me
This robot could support me
I consider this robot to be a social agent

of 32 subjects participated in the in-person study: 37.5% male and 62.5% female. Our participants ranged in age from 21 to 60 (M = 30, SD = 7) and came from 13 different countries. We took significant precautions beyond government regulations to protect participant health when running this study during the COVID-19 pandemic.

Procedures. After confirming their eligibility given the exclusion criteria, users scheduled an appointment for a 1.5-hour-long session with HuggieBot 2.0. Upon arrival, the experimenter explained the study, and the potential subject read the informed consent document and asked questions. If he/she still wanted to participate, the subject signed the consent form and the video release form, at which point the video cameras were turned on to record the experiment.

Users began by filling out a demographics survey on a computer. Next, the investigator introduced the robot as the personality "HuggieBot" and explained its key features, including the emergency stop. The experimenter explained how the trials would work and how the subject should be prepared to move. She also explained the two different ways to initiate a hug (walking, key press) and the three different ways to be released from a hug (release hands, lean back, wait until robot releases). At this point the subject filled out an opening survey to document his/her initial impressions of the robot; participants rated on a sliding scale from 0 (disagree) to 10 (agree) how much they agreed with the statements found in Table 2. Note that these questions were asked before the user had any experience physically hugging the robot. Next, the user performed practice hugs with the robot to acclimate to the hug initiation methods and the timing of the robot's arms closing. A participant was allowed to perform as many practice hugs as desired, verbally indicating to the experimenter when they were ready to begin the experiment. On average users did 2 or 3 practice hugs, but users taller than the robot (1.75 m) averaged 5 or 6 practice hugs because it took more time for them to find the most comfortable arm positions. The eight hugging conditions that made up this experiment are all three possible pairwise combinations of our three binary factors (vision, sizing, and release detection). The video associated with this paper shows a hugging trial from each of the eight conditions.

We used an 8 × 8 Latin square to counter-balance any effects of presentation order [14] and recruited 32 participants to have complete Latin squares. After each hug, the participant returned

to the computer and answered six questions. The first question was a free-response asking for the user's "first impressions of this interaction." Then, the participant used a sliding scale from 0 to 10 to answer the five questions found in Table 1, which were the same questions as in the online study. A subject could request to experience the same hug again if needed. After experiencing all eight hug conditions, the participants experienced an average of 16 more hugs, during which they contacted the robot's back in different ways and received various robot responses. Data were collected for these additional hugs, but they will not be analyzed in this paper due to space constraints. At the end of the experiment, the subject answered the same questions from the beginning of the study (Table 2). Finally, users could provide additional comments at the end.

All slider-type questions in the survey were based on previous surveys in HRI research and typical Unified Theory of Acceptance and Use of Technology (UTAUT) questionnaires [17, 41]. The free-response questions were designed to give the investigators any other information the participant wanted to share about the experience. A within-subjects study was selected for this experiment because we were most interested in the differences between the conditions, rather than the overall response levels to a robot hug. We also preferred this design for its higher statistical power given the same number of participants compared to a between-subjects study [9].

Results. This in-person study was the first robustness test of the fully integrated HuggieBot 2.0 system. Each subject experienced a minimum of 24 hugs during the study, plus practice hugs. With 32 total participants, the robot executed more than 850 successful hugs over the course of the entire study, sometimes giving 200 hugs in one day.

For all statistical analyses, we applied a Bonferroni alpha correction (to account for 15 multiple comparisons) to $\alpha = 0.05$ to determine significance. We use Pearson's linear correlation coefficient, ρ, to report effect size. Box plots of the responses to the opening and closing survey questions from Table 2 are shown in Fig. 5. In this study, answers were submitted on a continuous sliding scale, so a paired t-test comparison of the opening and closing survey was conducted. We found that users felt understood by (p = 0.0025, ρ = 0.57) and trusted the robot more (p < 0.001, ρ = 0.70) after participating in the experiment. Users also felt that robots were nicer to hug (p < 0.001, ρ = 0.76)

The responses to the five questions asked after each hug can be seen grouped by the presence and absence of each of the three tested factors (vision, sizing, and release) in Fig. 6. These responses were analyzed using three-way repeated measures analysis of variance via the built-in MATLAB function ranova; our data satisfy all assumptions of this analytical approach. No significant improvements were noticed in the perceived safety of the robot in any of the tested conditions, as the robot was consistently rated highly safe. The automatic size adjustment significantly increased users' impressions of the naturalness of the robot's movement (F(1,31) = 25.192, p < 0.001, ρ = 0.4158). Users found the robot's hug significantly more enjoyable when it adjusted to their size (F(1,31) = 70.553, p < 0.001, ρ = 0.3610). Automatic size adjustment to the user caused a significant increase (F(1,31) = 25.102, p < 0.001, ρ =0.4258) in the perceived intelligence of the robot. Finally, the robot was considered significantly friendlier when it adjusted to the size of
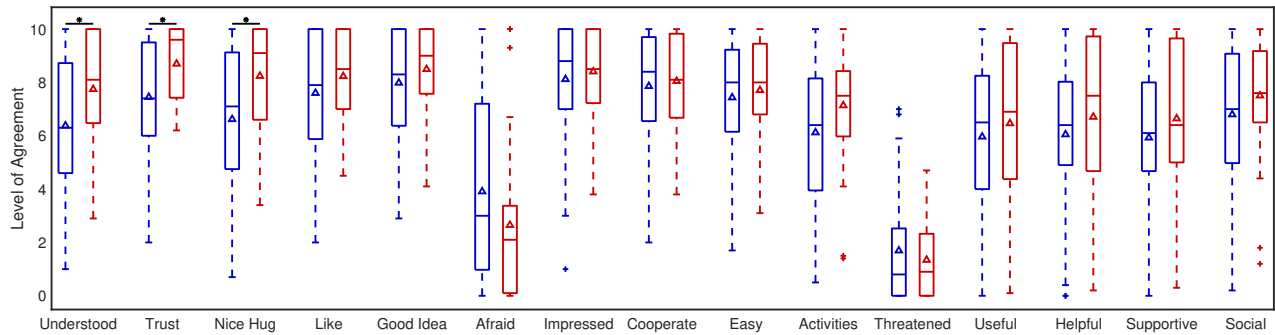
Figure 5: A comparison of the responses to the opening (blue) and closing (red) surveys. The top and bottom of the box represent the 25th and 75th percentile responses, respectively, while the line in the center represents the median, and the triangle indicates the mean. The lines extending past the boxes show the farthest data points not considered outliers. The + marks indicate outliers. The black lines with stars at the top of the graph indicate statistically significant differences.
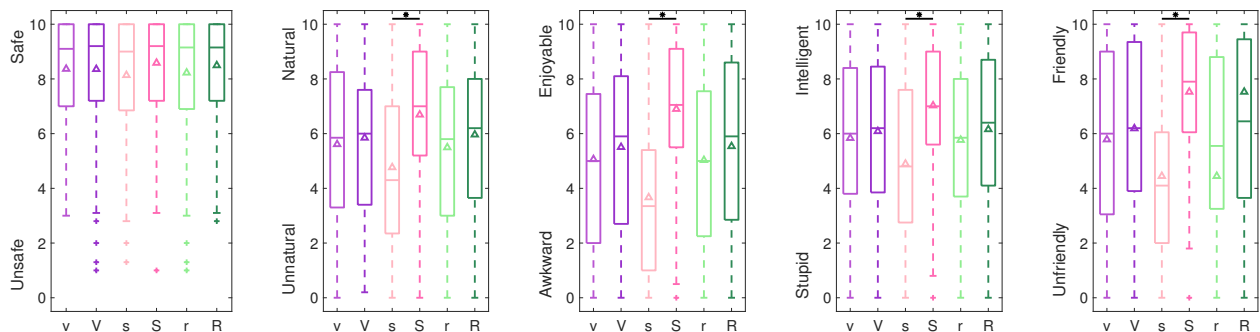
Figure 6: A comparison of the responses to the survey questions after each hug, grouped by factors. The purple colors represent without vision (v) and with vision (V), the pink colors represent without sizing (s) and with sizing (S), and the green colors represent timed release (r) and haptic release (R).

the user (F(1,31) = 84.925, p < 0.001, ρ = 0.3205). In summary, haptic hug sizing significantly affected every aspect except safety. Trials that included visual perception trended slightly positive but were not significantly different for any five of the investigated questions; small positive trends for haptic release were also not significant.

Eight users (25%) verbally stated and wrote about their preference for not having to push a button to activate a hug. Out of the 256 distinct hug response surveys (32 users, 8 surveys per user), the physical warmth of the robot was positively mentioned 100 times (39%), further validating that physical warmth is critical to pleasant robot hugs (T2) [3]. These positive comments were most commonly seen in the conditions with automatic hug sizing, presumably because the increased contact with the robot torso made the heat more apparent. Additionally, we observed that our participants used a mixture of pressure release and torque release to indicate their desire to end the hugs in the study. 17 users (53%) voiced their preference for the haptic release hugs, saying when the robot released before they were ready (in hugs with a timed release), "he didn't want to hug me!" or that the hug was "too short!" Interestingly, the hug condition where all three perceptual factors were present had the highest number of positive comments. 31 out of 32 users (96.8%) commented that this condition was the most "pleasant interaction," "natural," "friendly," or "fun."

7 DISCUSSION

Our three hypotheses were largely supported by the results. First, H1 hypothesized that when viewing from a distance, potential users will prefer the embodiment and movement of a hugging robot that incorporates our four new tenets over a state-of-the-art hugging robot that violates these tenets. The online study found that users significantly preferred HuggieBot 2.0 over HuggieBot 1.0; we believe our new robot was preferred because it obeys all six tenets. Written comments from the online community mention the "large," "hulking," and "over-powering" PR2 robot as unnerving when compared to the size of the user. In comparison, our robot is considered "nice" and "friendlier." Several users also wrote comments on how the people in the PR2 videos had to push a button on the robot's back, which seemed "unnatural", whereas the HuggieBot 2.0 release seemed more "intuitive". We did our best to match the videos of our new robot to the pre-existing videos of the PR2 so as not to bias the online viewers. These videos included but could not showcase visual hug initiation and haptic size adjustment. Based on the strong preference for our hugging robot in our carefully controlled online study, we conclude that users do prefer the embodiment and movement of a hugging robot that obeys our four new tenets over a state-of-the-art hugging robot that violates most of them.

H2 conjectured that obeying the four new tenets during physical interactions with users will significantly increase the perceived safety, naturalness, enjoyability, intelligence, and friendliness of a

hugging robot. We found that the haptic perception tenets had the greatest effects on these aspects of the robot, with haptic sizing positively affecting many responses and both haptic sizing and haptic release garnering positive comments. The lack of significant effects of visual initiation does not match the comments that users prefer the interaction when the robot recognizes their approach, rather than them having to push a button to initiate the hug. It is possible that users might not have included the button pushing in their rating of the hug as we simply asked users to "rate their experience with this hug" and did not explicitly tell them to include the hug initiation. Users might also have been confused that they had to walk towards the robot to initiate a hug, and then the robot would ask "Can I have a hug, please?" We chose to have the robot say the same phrase for both initiation methods to minimize variables, but as the user was initiating the hug, it may have made more sense for the robot to say something else or not speak.

We also believe there is room for improvement of the visual perception of our robot, which could contribute to higher ratings of the five questions. Currently, our perception of the user is based solely on his/her approach. To take perception even further, we believe hugs would be more comfortable if the robot could adjust its arm poses to match the approaching user's height and arm positions. Our taller users found the robot hugged them too low, and our shorter users found the opposite. Adjusting to user height would more fully obey T4 (visual perception) and therefore should be more acceptable to users. While the torso of the robot and dual release methods ensure our robot follows T6 of reliably releasing the user, a robot that could adjust its arm positions to the reciprocating pose of the user could greatly strengthen the user's impression of the robot's visual perception and improve the user opinion of the robot. We concede that our robot's rudimentary visual perception contributed to our lack of finding significant differences in the areas we investigated when testing with and without this factor.

Finally, H3 hypothesized that repeated hugs with a robot that follows all six tenets will improve user opinions about robots in general. We asked the same opening and closing survey questions as Block and Kuchenbecker [3], with similar results. Both studies' users felt more understood by, trusted, and thought that robots were nicer to hug after participating. HuggieBot 1.0's users also liked the presence of the robot more afterwards, found the robot easier to use, and viewed it as more of a social agent after the experiment, although these findings were reported without any statistical correction for multiple comparisons. The PR2 robot used in that experiment is significantly larger than an adult human, which violates T3. This domineering physical presence, therefore, contributed to lower initial ratings for users liking the presence of the robot, their perceived ease of use of the robot, and viewing the robot as a social agent. Our new robot, whose physical stature obeys the first three tenets, received higher initial ratings in these categories. HuggieBot 2.0 appeared as a friendly social agent from the beginning, and prolonged interaction with it confirmed these high initial impressions, which is why we did not find any significant differences for these questions in our study. As first impressions are often critical to determine whether a user will interact with a robot, here we see that it is important to obey the tenet prescribing the physical size of a robot. Therefore, we conclude that a robot that follows the six tenets does indeed improve user opinions about robots in general. A positive impression of the robot is crucial because it will make users more willing to receive a robot hug, and thus more likely to receive these health benefits when they cannot receive them from other people.

8 LIMITATIONS AND CONCLUSION

This study represents an important step in understanding intimate social-physical human-robot interactions, but it certainly has limitations. Due to COVID-19, the first study relied solely on videos and images, rather than the participants physically interacting with robots. Since we do not have access to a PR2, this online study enabled a fair comparison between our new platform and a previously well-rated hugging robot. Watching other people hug a robot is how users will decide if they also want to interact with a hugging robot in the wild.

One weakness of our new platform, HuggieBot 2.0, was the slow speed of the Kinova JACO arms. We selected these arms because of their inherent safety features; however, the distance the arms had to travel made their slow speed obvious and caused a long delay after hug initiation. When users started the hug with a button press, they could wait before walking to the robot for better timing. With visual hug initiation, the users were required to walk before the arms began moving, which resulted in many awkwardly waiting in front of the robot for the arms to close. Related to this limitation is the size of the room in which we conducted the experiment. A larger room would have let us set the threshold distance farther back to accommodate the speed of the arms. Both of these limitations could have contributed to the lack of significant differences between the hugs with and without visual perception.

Another limitation is the self-selection bias of our participants. For transparency, we advertised the experiment as a hugging robot study. While we succeeded in recruiting a diverse and largely non-technical audience to make our results as applicable to the general public as possible, we nevertheless acknowledge that users who chose to participate in the study were interested in robots. Because we did not hide the nature of our study, we did not have any participants who refused to hug the robot, as might occur in a more natural in-the-wild study design.

This project took a critical look at state-of-the-art hugging robots, improved upon their flaws, built upon their successes, and created a new hugging robot, HuggieBot 2.0. We also propose to the HRI community six tenets of hugging that future designers should consider to improve user acceptance of hugging robots. During times of social distancing, the consequences of lack of physical contact with others can be more damaging and prevalent than ever. If we cannot seek comfort from other people due to physical distance or health or safety concerns, it is important that we seek other opportunities to reap the benefits of this helpful interaction.

REFERENCES

[1] James Barber, Richard A. Volz, Rajiv Desai, Ronitt Rubinfeld, Brian Schipper, and Jan Wolter. 1986. Automatic two-fingered grip selection. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Vol. 3. 890–896.

[2] Alexis E. Block and Katherine J. Kuchenbecker. 2018. Emotionally supporting humans through robot hugs. In *Companion of the ACM/IEEE International Conference on Human-Robot Interaction*. Chicago, USA, 293–294.

[3] Alexis E. Block and Katherine J. Kuchenbecker. 2019. Softness, warmth, and responsiveness improve robot hugs. *International Journal of Social Robotics* 11, 1 (2019), 49–64. https://doi.org/10.1007/s12369-018-0495-2

[4] Carissa J. Cascio, David Moore, and Francis McGlone. 2019. Social touch and human development. *Developmental Cognitive Neuroscience* 35 (2019), 5–11. https://doi.org/10.1016/j.dcn.2018.04.009

[5] Wongun Choi, Caroline Pantofaru, and Silvio Savarese. 2011. Detecting and tracking people using an RGB-D camera via multiple detector fusion. In *Proceedings of the IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*. 1076–1083.

[6] Sheldon Cohen, Denise Janicki-Deverts, Ronald B. Turner, and William J. Doyle. 2015. Does hugging provide stress-buffering social support? a study of susceptibility to upper respiratory infection and illness. *Psychological Science* 26, 2 (2015), 135–147. https://doi.org/10.1177/0956797614559284

[7] Marco Costanzo, Giuseppe De Maria, and Ciro Natale. 2020. Two-fingered in-hand object handling based on force/tactile feedback. *IEEE Transactions on Robotics* 36, 1 (2020), 157–173.

[8] Navneet Dalal, Bill Triggs, and Cordelia Schmid. 2006. Human detection using oriented histograms of flow and appearance. In *Proceedings of the European Conference on Computer Vision (ECCV) - Volume Part II* (Graz, Austria). Springer-Verlag, Berlin, Heidelberg, 428–441. https://doi.org/10.1007/11744047_33

[9] Joost C. F. de Winter and Dimitra Dodou. 2017. Experimental design in: human subject research for engineers. *SpringerBriefs in Applied Sciences and Technology* (2017).

[10] Carl DiSalvo, Francine Gemperle, Jodi Forlizzi, and Elliott Montgomery. 2003. The Hug: an exploration of robotic form for intimate communication. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 403–408.

[11] Julia C. Duvall, Lucy E. Dunne, Nicholas Schleif, and Brad Holschuh. 2016. Active hugging vest for deep touch pressure therapy. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct*. 458–463.

[12] Stephen M. Edelson, Meredyth Goldberg, David C. R. Kerr, and Temple Grandin. 1999. Behavioral and Physiological Effects of Deep Pressure on Children With Autism: A Pilot Study Evaluating the Efficacy of Grandin's Hug Machine. 53, 1979 (1999), 145–152.

[13] Naomi T. Fitter and Katherine J. Kuchenbecker. 2016. Designing and assessing expressive open-source faces for the Baxter robot. In *Proceedings of the International Conference on Social Robotics*. Springer, 340–350.

[14] David A. Grant. 1948. The Latin square principle in the design and analysis of psychological experiments. *Psychological Bulletin* 45, 5 (1948), 427–442.

[15] Edward T. Hall, Ray L. Birdwhistell, Bernhard Bock, Paul Bohannan, A. Richard Diebold Jr., Marshall Durbin, Munro S. Edmonson, J. L. Fischer, Dell Hymes, Solon T. Kimball, Weston La Barre, Franck Lynch, S. J, J. E. McClellan, Donald S. Marshall, G. B. Milner, Harvey B. Sarles, George L. Trager, and Andrew P. Vayda. 1968. Proxemics [and Comments and Replies]. *Current Anthropology* 9, 2/3 (1968), 83–108.

[16] Hooman Hedayati, Srinjita Bhaduri, Tamara Sumner, Daniel Szafir, and Mark D. Gross. 2019. HugBot: a soft robot designed to give human-like hugs. In *Proceedings of the ACM International Conference on Interaction Design and Children*. 556–561.

[17] Marcel Heerink, Ben Krose, Vanessa Evers, and Bob J. Wielinga. 2009. Measuring acceptance of an assistive social robot: a suggested toolkit. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 528–533.

[18] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. 2016. SSD: Single Shot MultiBox Detector. *Lecture Notes in Computer Science* (2016), 21–37. https://doi.org/10.1007/978-3-319-46448-0_2

[19] Raymond R. Ma, Adam Spiers, and Aaron M. Dollar. 2016. M2 gripper: extending the dexterity of a simple, underactuated gripper. *Advances in Reconfigurable Mechanisms and Robots II. Mechanisms and Machine Science* 36 (2016). https://doi.org/10.1007/978-3-319-23327-7_68

[20] Takahiro Miyashita and Hiroshi Ishiguro. 2004. Human-like natural behavior generation based on involuntary motions for humanoid robots. *Robotics and Autonomous Systems* 48, 4 (2004), 203–212.

[21] Catriona M. Morrison and Helen Gore. 2010. The relationship between excessive Internet use and depression: a questionnaire-based study of 1,319 young people and adults. *Psychopathology* 43, 2 (2010), 121–126. https://doi.org/10.1159/000277001

[22] Corey J. Blomfield Neira and Bonnie L. Barber. 2014. Social networking site use: linked to adolescents' social self-concept, self-esteem, and depressed mood. *Australian Journal of Psychology* 66, 1 (2014), 56–64. https://doi.org/10.1111/ajpy.12034

[23] Cagatay Odabasi. 2017. ROS people object detection and action recognition Tensorflow. https://github.com/cagbal/ros_people_object_detection_tensorflow

[24] Mike Oren, Constantine Papageorgiou, P. Sinha, E. Osuna, and Tomaso Poggio. 1997. Pedestrian detection using wavelet templates. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. 193–199.

[25] Manoranjan Paul, Shah Haque, and Subrata Chakraborty. 2013. Human detection in surveillance videos and its applications: a review. *Eurasip Journal on Applied Signal Processing* 176 (2013), 1–16. https://doi.org/10.1186/1687-6180-2013-176

[26] Joseph M. Romano, Kaijen Hsiao, Günter Niemeyer, Sachin Chitta, and Katherine J. Kuchenbecker. 2011. Human-inspired robotic grasp control with tactile sensing. *IEEE Transactions on Robotics* 27, 6 (December 2011), 1067–1079.

[27] Tom Scott. 2020. 1,204,986 Votes Decided: What Is The Best Thing? *YouTube* (2020). https://www.youtube.com/watch?v=ALy6e7GbDR

[28] Masahiro Shiomi, Aya Nakata, Masayuki Kanbara, and Norihiro Hagita. 2017. A hug from a robot encourages prosocial behavior. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 418–423.

[29] Masahiro Shiomi, Aya Nakata, Masayuki Kanbara, and Norihiro Hagita. 2017. A robot that encourages self-disclosure by hug. In *Proceedings of the International Conference on Social Robotics (ICSR)*. Springer, 324–333.

[30] Walter Dan Stiehl, Jeff Lieberman, Cynthia Breazeal, Louis Basel, Levi Lalla, and Michael Wolf. 2005. Design of a therapeutic robotic companion for relational, affective touch. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 408–415.

[31] Hidenobu Sumioka, Aya Nakae, Ryota Kanai, and Hiroshi Ishiguro. 2013. Huggable communication medium decreases cortisol levels. *Scientific Reports* 3, 1 (2013), 3034. https://doi.org/10.1038/srep03034

[32] Juulia T. Suvilehto, Enrico Glerean, Robin I. M. Dunbar, Riitta Hari, and Lauri Nummenmaa. 2015. Topography of social touching depends on emotional bonds between humans. *Proceedings of the National Academy of Sciences* 112, 45 (2015), 13811–13816. https://doi.org/10.1073/pnas.1519231112

[33] James Keng Soon Teh, Adrian David Cheok, Roshan L. Peiris, Yongsoon Choi, Vuong Thuong, and Sha Lai. 2008. Huggy Pajama: a mobile parent and child hugging communication system. In *Proceedings of the ACM International Conference on Interaction Design and Children*. 250–257.

[34] William Townsend. 2000. The BarrettHand grasper: a programmably flexible part handling and assembly. *Industrial Robot: An International Journal* 27, 3 (2000), 181–188.

[35] Gabriele Trovato, Martin Do, Ömer Terlemez, Christian Mandery, Hiroyuki Ishii, Nadia Bianchi-Berthouze, Tamim Asfour, and Atsuo Takanishi. 2016. Is hugging a robot weird? Investigating the influence of robot appearance on users' perception of hugging. In *Proceedings of the IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. 318–323.

[36] Dzmitry Tsetserukou. 2010. HaptiHug: a novel haptic display for communication of hug over a distance. In *Proceedings of the International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*. Springer, 340–347.

[37] Oncel Tuzel, Fatih Porikli, and Peter Meer. 2007. Human detection via classification on Riemannian manifolds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1–8.

[38] Paul Viola and Michael Jones. 2004. Robust real-time object detection. *International Journal of Computer Vision* 57, 2 (2004), 137–154.

[39] Paul Viola and Michael J. Jones. 2001. Robust real-time face detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Vol. 2. 747–747.

[40] Duc My Vo, Lixing Jiang, and Andreas Zell. 2014. Real time person detection and tracking by mobile robots using RGB-D images. In *Proceedings of the IEEE International Conference on Robotics and Biomimetics (ROBIO)*. 689–694.

[41] Astrid Weiss, Regina Bernhaupt, Manfred Tscheligi, Dirk Wollherr, Kolja Kuhnlenz, and Martin Buss. 2008. A methodological variation for acceptance evaluation of human-robot interaction in public places. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 713–718.

[42] Katsu Yamane, Joohyung Kim, and Alexander Nicholas Alspach. 2017. Soft Body Robot for Physical Interaction. US Patent Application 20,170,095,925.

[43] Ryuji Yamazaki, Louise Christensen, Kate Skov, Chi-Chih Chang, Malene F. Damholdt, Hidenobu Sumioka, Shuichi Nishio, and Hiroshi Ishiguro. 2016. Intimacy in phone conversations: anxiety reduction for Danish seniors with Hugvie. *Frontiers in Psychology* 7 (2016), 537. https://doi.org/10.3389/fpsyg.2016.00537

[44] Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and Kazuya Takeda. 2020. A survey of autonomous driving: common practices and emerging technologies. *IEEE Access* 8 (2020), 58443–58469. https://doi.org/10.1109/ACCESS.2020.2983149