Symmetric Telepresence using Robotic Humanoid Surrogates

Abstract

Telepresence involves the use of virtual reality technology to facilitate participation in distant events, including potentially performing tasks, while creating a sense of being in that location. Traditionally, such systems are asymmetric in nature where only one side (participant) is "teleported" to the remote location. In this manuscript, we explored the possibility of symmetric 3D telepresence where both sides (participants) are "teleported" simultaneously to each other's location; the overarching concept of *symmetric telepresence* in virtual environments is extended to telepresence robots in physical environments. Two identical physical humanoid robots located in UK and the United States serve as *surrogates* while performing a transcontinental shared collaborative task. The actions of these *surrogate* robots are driven by capturing the intent of the participants controlling them in either location. Participants could communicate verbally but could not see the other person or the remote location while performing the task. The effectiveness of gesturing along with other observations during this preliminary experiment are presented. Results reveal that the symmetric robotic telepresence allowed participants to use and understand gestures in cases where they would otherwise have to describe their actions verbally.

Keywords: 3D telepresence, telepresence robots, symmetric telepresence, robotic surrogates

1 Introduction

Telepresence is a concept that has been widely studied by researchers for several years. Marvin Minsky, in 1980, pioneered the concept of mechanical telepresence where each motion of a person's hand, arm and fingers was reproduced in a different room, city, country or planet using mobile mechanical hands [1]. The idea was to provide an ability to work in distant environments while allowing a user to see and feel what was happening - in other words, providing a sensation of "being there". The process of enabling telepresence is sometimes referred to as "teleportation". With this came several applications of the technology including the idea of remote surgery and applications related to space exploration. Since then, researchers in robotics and virtual reality have identified several elements that can enhance the telepresence experience. Their focus has traditionally been on unidirectional asymmetric telepresence. The research questions answered historically can generally be categorized into one of the following:

(i) What are the factors that influence people to believe that they are in a different location (Presence)?

(ii) How can we improve the ability of people to perform tasks in remote locations (Teleoperation)?

(iii) How can we combine (i) and (ii) so a person in the remote location and the person being "teleported" can effectively communicate or work with each other?

In this manuscript, the concept of unidirectional telepresence is extended by teleporting two users in different locations simultaneously to each other's locations. This is referred to as bi-directional symmetric telepresence. While the concept of symmetric telepresence has been explored in virtual environments (via virtual embodiments), the form of

telepresence described in this manuscript is mechanical in nature (and hence 3-D) since identical humanoid robots are used at both locations. This work is comparable to using bi-directional avatars but is a pilot demonstration of bi-directional working with robots. These humanoid robots are referred to as *surrogates* (Figure 1). *Surrogates* are defined as context-specific stand-ins for real humans. Traditionally, manifestations of *surrogates* are referred to as avatars or agents depending on the entity controlling them - avatars are controlled by humans (traditionally referred to as *Inhabiters*) while agents are controlled by computer programs. The term *Surrogate* avoids having to explicitly differentiate between avatars and agents thereby allowing hybrid versions of control i.e. a *surrogate* may be an avatar at one instant and an agent in the next. Physical manifestations of avatars or agents provide the ability to manipulate things in the remote environment. In this exploratory work (pilot), the focus is on providing users at either end, the ability to gesture using their *surrogates*. A shared collaborative task which involves solving tic-tac-toe puzzles is chosen to encourage *inhabiters* to gesture through their *surrogates*.

The remainder of this manuscript is organized as follows. In section 2, important background literature and previous work in the area of telepresence is covered. Section 3 contains a description of the system architecture used to control the robotic humanoid *surrogates*. The experimental setup and design are covered in Section 4. Section 5 is a discussion of the observations during this task for a small number of participants. Conclusions and future work form the last section of this manuscript.

2 Related Work

Telepresence robots, including mobile telepresence robots and humanoid robots have been studied by several researchers over the years. They provide a connection between a user and a distant participant or a remote environment to perform social interactions or specific tasks. Mobile telepresence robots such as MeBot V4 [2], PRoP [3], Anybots'QB and the VGo [4], allow a remote user to control the robot's movement around a space while the user converses with other users in that space. Using these telepresence robots, remote co-workers can wander the hallways and engage in impromptu interactions, increasing opportunities for connection in the workplace [5]. While such mobile robots have been introduced to support telepresence, the anthropomorphic nature of humanoid robots may allow for better conveyance of a person's remote physical presence. In addition, these humanoid robots could allow for manipulation of objects in the remote environment, thereby increasing the feeling of "presence" for a user. Among such humanoid robots is the Geminoid HI-1 [6] developed to closely resemble a specific human. While not capable of manipulating objects in the environment, it was evaluated as being highly human-like but uncanny [7]. Related research includes the concept of Animatronic Shader Lamp Avatars (SLA) [8]. Here, researchers use the technique where an image of an object is projected onto a screen whose shape physically matches the object. Cameras and projectors are used to capture and map the dynamic motion and appearance of a real person onto a humanoid animatronic model. These avatars can potentially be used to represent specific visitors at a destination but are limited in their flexibility to gesture in the remote environment. Another related concept is that of tele-existence [9] where a user is given the sense that they are inside the robot itself. Not all telepresence systems however support tele-existence and this particular concept although relevant, is in-fact not explored in this manuscript.

When using robotic systems for telepresence, understanding the psychology of human-robot interactions is critical. In [10] and [11], the interaction between humans and robots is studied. In this research, the robots are agents, not inhabited by humans, but capable of automatically processing a person's physical motion and verbal communication. The authors of these papers were primarily concerned with giving the robot the "gestural awareness" and eye contact required in a natural interaction. Several researchers have investigated the effect of human comfort and trust when interacting with robots. It must be noted that these robots are typically not *surrogates* i.e. no human in the loop controls them. In a meta-analysis of literature in the

Figure 1: Two identical humanoid robots located in UK and USA were used as *Surrogates* to explore the concept of symmetric bi-directional telepresence.

area [12], features such as anthropomorphism, co-location, robot personality, behavior, predictivity and level of automation were all important factors in establishing trust between a human and a robot during interactions. For the purposes of telepresence, one can envision that these factors are not only predictable (as a result of the human inhabiting the robot) but also important to facilitate the interaction.

Most of the previous work in telepresence has focused on unidirectional robotic systems. Two-way symmetric telepresence via robotic *surrogates* has not been investigated yet. In [13] and [14], a unified framework for generic avatar control is presented. With this control paradigm, an *inhabiter* is able to manipulate a single avatar while automated "behaviors" power the remaining characters in the environment. In this paper we present a similar control strategy for controlling robotic avatar manifestations in order to facilitate telepresence. In addition, our framework lends itself to controlling multiple robotic manifestations simultaneously and thereby facilitates symmetric bi-directional telepresence.

3 System Architecture

One of the central components of this manuscript is the ability to remotely inhabit a *surrogate*. For the purposes of this study, a commercial off-the-shelf robotic humanoid called the *Robothespian™* was used as the *surrogate*. The Robothespian features a hybrid actuation mechanism consisting of fluidic muscles and DC motors along with passive compliance

elements and offers 24 degrees of freedom. To support telepresence using this *Surrogate*, an *inhabiter's* intent is realized and transmitted accurately in real-time, while the closed-loop response of the Robothespian's hybrid actuation mechanism is adapted to faithfully represent this action. The control aspects of this paradigm are not covered in detail here since the focus of this manuscript is on presenting the concept of symmetric telepresence.

To support teleoperation and telepresence, a Master-Slave architecture is employed. The Master uses virtual characters that can be controlled using generic input devices, one of which is a magnetic tracking device called the Razer Hydra. A calibration routine on the Master allows users to map their motion to corresponding actions of the master virtual characters. This is a gestural interface and not a literal interface i.e. the motions of the *inhabiter* do not have to explicitly match the desired motions of the virtual character. The person controlling the virtual characters (or avatars) is referred to as an *inhabiter*. The *inhabiter's* intent is transmitted via a lightweight networking protocol to a Slave (client) program. The Slave has the same continuous avatar state representations as that of the Master. A subroutine on the Slave maps the motions of the active virtual characters onto any secondary hardware manifestations such as a humanoid robotic *surrogate*. In this case, the active avatar's motions are mapped to those of

the *Robothespian™*. The mapping is achieved via a custom routine that identifies the number of degrees available on the specific robotic *surrogate*, extracts the relevant data from the active avatar and applies it in joint-space using a traditional PID controller (positionally or via velocity control). An illustration of the general architecture is provided in Figure 2. We refer to this as the "Teleoperation Paradigm" in the remainder of this manuscript.


Figure 2: The system architecture used for asymmetric uni-directional teleoperation or telepresence using robotic manifestations of avatars. A synchronously updated slave avatar instance is used to drive the actions of the robotic *surrogate* based on an *inhabiter's* intent.

The advantage of such an architecture is the support offered for multiple configurations involving several hardware devices, each of which can inherit different actions based on their parent avatar's characteristics. While not the focus of this manuscript, these robotic *surrogates* also support appearance changing via rear-projected faces.

The Teleoperation Paradigm shown in Figure 2 can be extended to work bi-directionally. When instantiated simultaneously in two locations, it is possible for 2 masters and 2 slaves to function in parallel. This results in an architecture that supports symmetric telepresence as seen in Figure 3. The components of the "Teleoperation Paradigm" can be inferred from Figure 2.

One of the key features of such a paradigm is the closed-loop nature of the approach as seen in Figure 3. In specific, the actions of "*Inhabiter* 1" drive those of "*surrogate* 1". This in turn causes "*Inhabiter* 2" to respond potentially with both gestures and verbally. The actions of "*Inhabiter* 2" now drive those of "*surrogate* 2". The process continues and this results in each *inhabiter*


Figure 3: The teleoperation paradigm described in Figure 2 is extended to support symmetric bi-directional telepresence. The architecture results in a closed-loop scenario where the actions of each *inhabiter* continuously drive their corresponding *surrogates* resulting in collaboration on each side simultaneously.

collaborating indirectly with the other via their respective *surrogates*.

4 Experimental Design

To evaluate the system, we created a set of tasks that required the collaboration of two participants located at transcontinental sites (USA, UK). Specifically, the tasks involved solving a series of twenty tic-tac-toe puzzles since this was likely to promote discussion via gestures. For each puzzle, the participants were presented with a partially completed tic-tac-toe game, and instructed to come to an agreement on the next best move for either "X" or "O". An example of two such puzzles is shown in Figure 4.

As mentioned, our experimental design relied on the notion that the tic-tac-toe puzzle designs would encourage gesturing at the board. Puzzles were chosen such that all squares on the board would be potential solutions, and thus participants would likely gesture towards all nine squares over the course of the twenty puzzles. Of course, certain squares can be expected to be favored by the participants because they are known to be good moves (e.g. the center square) in general. We note that participants may not always identify or converge on the optimal or best possible move on the board. The study did not test subjects on this aspect as the emphasis for this trial was on encouraging discussion with gestures. The full set of puzzles will be available at *http://anonurl* for those interested in explor-
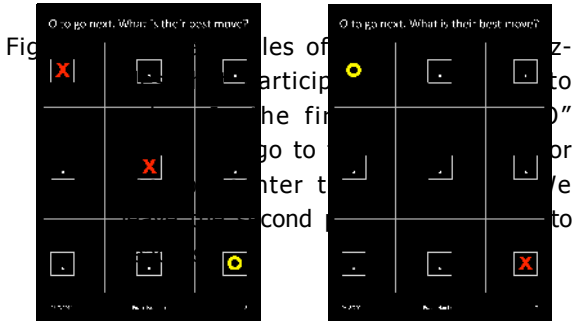
Fig... les of ...z-...articip... to ...he fir... 0" ...go to ... or ...nter t... ...e ...cond ... to ing similar setups.

Experimental Setup: Figure 5 shows the setup of the experiment at the USA site. This setup was mirrored at the site in UK. The robotic humanoid *surrogate* was capable of all nine different gestures required to point at each square of the tic-tac-toe board. When the participant at Location A did not perform a gesture, their robotic *surrogate* in Location B would default to returning to a neutral stance while observing the participant in its location. Participants in either location used a magnetic tracking device called the Razer Hydra to inhabit their robotic *surrogate*. A Kinect device was positioned appropriately in the experimental area to collect data for analysis including video and audio streams. Participants were asked to answer a post-interaction questionnaire (Table 1) to correlate the qualitative and quantitative metrics collected during the study.

Priming: Before each experiment, the participants watched an instructional video detailing their task and the usage of the system. Participants were not made aware of the symmetric control system or told that their Hydra gestures were mapped onto another robot. Instead, they were simply informed that they would have to use the Hydra device in order for the robot in their location to understand their intent (pointing gestures). Participants were unable to view the other side since no video was used to support telepresence. The robotic *surrogates* were a di-

No.	Question
1.	How well do you feel the collaboration with your partner went?
2.	How much did you feel that your collaborator was here in the lab?
3.	Did you find it easy to communicate with your collaborator?
4.	How confident were you that the meaning of your gestures was conveyed to your collaborator?
5.	Did you feel that you did more of the puzzle solving or them?
6.	Do you feel that the robotic *surrogate* acted naturally?
7.	Could you understand what your collaborator was attempting to convey through the robotic *surrogate*?
8.	Did you feel comfortable around the robotic *surrogate* (did it invade your space / did you feel unsafe)?

Table 1: Table showing the post-interaction questionnaire

rect means of telepresence and as a result, they were forced to observe the robotic *surrogate* (of the other collaborator) in their own space to understand visual pointing cues. The participants were allowed to verbally communicate with the *surrogate*. The speakers on both sides were positioned in such a way to make the sound appear as if it was coming from the robotic *surrogate* itself.

5 Results and Discussion

Several qualitative and quantitative metrics were collected during the interactions at both sites. Since this was an exploratory pilot study, a video analysis of all the participant interactions was performed to gain insight into the effectiveness and usefulness of the physical bi-directional robotic telepresence system. In this section, we discuss some of observations during the interactions including computing the latency times, the number of times participants pointed to a square during the interaction, general notes from observation of audio and video data streams and interviews with the participants.

A total of seven pairs of participants tested the system (labeled as P1-P7 below). Due to software failure, data for one set (P2) was incom-

Figure 5: The experimental setup at the site in USA showing the display surface, the control device, the robotic *surrogate* (Robothespian), and the Kinect device used to collect data for analysis of the interaction. This setup was mirrored at the site in UK with the only exception being the display surface which was replaced by a traditional flip chart.

plete, and thus the quantitative metrics for those are not available. Qualitative responses are still included.

5.1 Latency

As a part of the communication architecture, "heartbeats" were sent every 10 seconds between the server and the client. The client recorded the local time at which a heartbeat was received from the server and also recorded the time at which the heartbeat originated from the server. Clock time difference and latency between the two locations were estimated from this data. To do this, a latency test was run between the two networks (USA and UK) situated reasonably close to backbones. This value was noted as $est_{latency}$. The heartbeats were analyzed to identify a value with minimum clock difference. This was the most representative first estimate of the $ClOCkshift_{meas}$ since anything greater can be attributed to latency or program delays. $ClOCkshift_{meas} - est_{latency}$ is now the best estimate for the actual $ClOCkshift_{act}$ between the two locations. This $ClOCkshift_{act}$ was then subtracted from the vector of heartbeats received. The resultant is a vector of latencies corresponding to each heartbeat sent during a session. Using this data, the mean latency per

session (interaction between pairs of users) was calculated.


Figure 6: The computed mean latency from the heartbeat data during each of the six sessions seemed to vary at the two sites. It was fairly consistent at the UK site, varying roughly around the 1.0s mark. At the site in the USA, the observed latency varied between 0.2s and 1.0s

Since the heartbeats were implemented only one-way, it was possible to differentiate the observed latencies at the two remote sites. The graphs show that the latency was quite variable, indicating, potentially, that route changes were occurring or there was significant load. It should be noted that an instance of TeamViewer (Remote Desktop) was running on the machine in UK during the experiment, though this machine was not CPU bound, and the bandwidth used was well within the local network capabilities. Thus it could be inferred that it took approximately 1s between the user gesturing and their robotic *surrogate* at the other end moving. The *surrogate* then took a small amount of time to reach its final destination as a result of its inherent actuation mechanism involving fluidic muscles, typically characterized by smooth and non-jerky responses.

5.2 Gesturing

During the interaction, all gestures towards the tic-tac-toe board performed by both the participant as well as the robotic *surrogate* were recorded in each location. Once all interactions were complete, the log files showing the interactor's intent (gesture) and the corresponding robot's gesture (pose obtained) were verified. The intent of the *inhabiter* was found to be transmitted to their robotic *surrogate* via the master-slave architecture on all occasions. The data revealing the robotic *surrogate's* pose is viewed as a heatmap in Figure 7. The heatmap reveals that participants pointed to every square on the board at least once. In addition to this, Table 2 and Table 3 show the mean and standard deviation of the total number of times each participant pointed to each square on the board. This demonstrates that subjects in both locations were using gestures to communicate during the interaction. As an aside, the increased pointing to the middle center and bottom center squares was a result of the particular state of the Tic-Tac-Toe board. The uneven distribution simply indicates that the set of puzzles was not "rotation symmetric" and the bottom center square was a reasonable choice more often. If the puzzles were indeed designed to be "rotation symmetric", the uneven distribution would be indicative of a mechanical or control problem with the robotic *surrogate* (e.g., the actuators would not have sufficient power to point to the upper squares). The robotic *surrogate* systems at both ends were checked and tuned before the experiment to mitigate this risk.

Participant	Mean	Standard Deviation
1	3.0000	3.0000
3	3.2222	3.1535
4	3.8889	2.4721
5	4.8889	4.1667
6	5.6667	2.7386
7	4.7778	4.0552

Table 2: The mean and standard deviation for the number of times a participant pointed to each square in UK

Participant	Mean	Standard Deviation
1	0.6667	1.0000
3	4.7778	3.8333
4	3.2222	2.0480
5	0.7778	1.0929
6	3.1111	2.3688
7	3.4444	3.2582

Table 3: The mean and standard deviation for the number of times a participant pointed to each square in the USA


Figure 7: The heatmap showing the number of times all participants pointed to each square on the board in UK (Top) and USA (Bottom)

5.3 Body language

The results from the trials show that participants were successful in using gestures to communicate. If they successfully agreed upon the best possible position for the next "O" or "X" without verbalizing their location, it indicated that the robotic *surrogate's* movements conveyed their intentions correctly.

In analysis of videos of the participants performing the task we saw a variety of collaborative strategies come to light: one participant pointing and the other just agreeing; one participant pointing and the other participant pointing at the same square to confirm (see Figure 8 top); or a more complex exchange where pointing was used to express differences of opinion. We also saw several failed communication attempts: including pointing with the un-tracked hand, or gesturing at the board in a more complex way (e.g. painting lines on the board) that was not captured by the system. This was attributed

Figure 8: Examples from trials showing various communicative acts. Top: Participant and robotic *surrogate* pointing at the same square at the same time. Middle: Participant glancing at robot while gesturing. Bottom: Participant glancing at robotic *surrogate's* hand even though it is not pointing.

to the fact that our *surrogate* control paradigm did not involve full body motion. Instead, we only tracked gestures that were considered important for the task - in this case, this involved pointing towards the board (i.e. reaching out). These observations may suggest a need to fully interpret all user gestures since they could subconsciously use hand and arm gestures without knowing or considering how the system will be able to convey these.

We observed the participants making several other communicative actions such as smiling at the robotic *surrogate*, waving good bye, nodding and shrugging. None of these were captured and relayed via the robotic *surrogates* to the other participant. The participants did not gaze frequently at the *surrogate's* face, but tended to focus on the board and the *surrogate's* gestures near the board. We hypothesize that there may be more glancing towards the face if the robotic *surrogate* had a human-like appear-

ance via the rear-projection display. A couple of participants noted a lack of information from the face in their interviews as discussed below. We do see participants looking towards the robotic *surrogate* as they gesture, and also occasionally looking at the *surrogate's* hands as if they are about to move although they do not. Please refer to Figure 8 middle and bottom.

One pair of subjects (P4) both noted in interviews that they mostly used verbal communication for the task as they felt this was sufficient. However data revealed that they did gesture using the robotic *surrogate*.

5.4 Co-Presence

Appropriate body language is evidence of co-presence between participants. It is also evidenced by the responses of the participants to the questionnaires. Specifically, we enquired how participants felt with regards to whether or not their collaborator was with them. In addition we also looked at how easy they found it to communicate with their collaborator. When asked directly "How much did you feel that your collaborator was here in the lab?", seven participants reported yes to some extent, two could not say and five reported no. One of the most interesting responses was the following:

> *P3@SiteInUSA: I felt he was standing right next to me when I was not looking at the robotic surrogate. The physical presence of the surrogate disrupted me. If I was looking at the robotic surrogate's face then I felt he was not in the lab.*

This might indicate that when focused on the task, the participants only peripherally aware of the robotic *surrogate*. Three participants said that the audio was a distraction, with comments *"it was like a telephone call"*, *"the sound came from the whole room"* and *"... voice is far away ..."*. Producing an authentic sounding voice that originates from the *surrogate* is a top priority for future work.

From the observations of all trials it appears that participants became more comfortable with

the collaboration over time. One participant directly commented (when asked "Did you find it easy to communicate with your collaborator?")

> *P2@SiteInUSA: After we got started. I was not sure how to talk to the robotic surrogate at the beginning.*

Another commented (when asked "How confident were you that the meaning of your gestures was conveyed to your collaborator?")

> *P7@SiteInUSA: Very confident, other than the first one.*

This again suggests that participants did not have a good understanding of the capabilities of such robotic *surrogates* initially because they are still unusual.

Several participants mentioned that the robotic *surrogate's* movement was clunky and slow. This was expected because of the limitations of the hardware control loops and the distance between the two sites. Improving the responses of such robotic *surrogates* also forms a component of our planned future work in the mainstream area of robotics and control systems.

5.5 Safety

Most participants did not report feeling unsafe around the robotic *surrogate*. We did observe participants stepping back and participants reported that they felt the need to step back, but this did not make them feel unsafe. The robotic *surrogate* (and its *inhabiter*) did not have knowledge (situational awareness) of its local environment. As a result, the robot sometimes invaded the participant's space thereby triggering an avoidance instinct in them. This behavior would typically not occur if people collaborating closely in a physical space had knowledge about each other (including via their robotic *surrogates*). In the interview, one participant said that they felt unsafe because they felt the *surrogate* was not looking at what they were doing. This suggests that the robotic *surrogate* needs to appear to be continually aware of the participant's activity, even if it is not interacting with the participant. This ties back into the "situational awareness" discussion of the *surrogate*. We also refer back to Figure 8 where the participants frequently glance towards the

robotic *surrogate*, perhaps to gauge whether it will move; one could investigate if the robotic *surrogate* should do the same. Another participant said that they would have felt more comfortable around the robotic *surrogate* if they had known its capabilities. This is an interesting observation as it suggests that even as these robotic *surrogates* become more realistic in appearance, those interacting with them may not trust them because they understand that robots in general can have different capabilities than humans. We have covered some of the previous research regarding trust during human-robot interaction in the related work section of this manuscript.

6 Conclusion

Collaboration at a distance has long been a topic of interest to the virtual reality community. The system we have described shares many of the same software components as a collaborative virtual environment; the distinction is that the realization of the shared environment is done through physical manifestations: the robotic *surrogates*.

In the paper we have shown that two remote participants can collaborate on a shared task that involved voice and gestural communication. In this pilot trial, we found that participants would gesture to communicate spatial positions and did not have to resort to voice generally to complete the task. The trials highlighted several potential directions for research and development, such as having the robotic *surrogate* appear to monitor the participant, capturing more of the participants' behavior, improving audio reproduction and improving overall system latency.

We developed this scenario primarily to push the technical boundaries of what was possible with robotic *surrogate* representations. We found the use of physical robots for telepresence interesting because of issues with latency and timing that are perhaps not a major challenge with purely virtual avatars. In addition, the telepresence occurs in a physical environment at both locations allowing the trials to physical manipulation in the next stage of this work. We also note that the scenario has potential use in training or rehearsal scenarios where tactile and haptic cues are important. For example, a trainer and trainee could both have physical access to

an engineering piece, where it is important that both have "hands-on" access simultaneously to the piece. While the robotic *surrogates* we are using today are not able to manipulate objects, the next generation will be able to do so.

There are several routes for future research. Firstly, we would to explore other important scenarios and natural interaction (e.g., trust) by using a more sophisticated task. Secondly, we would like to use a more systematic user study with more subjects and quantitative report of the experimental results (e.g., the contribution of Physical manifestations with and without audio.) Thirdly, the autonomous capability augments to human control will be an important research area of robotics. We would also like to explore the intelligent capability of robot.

References

[1] Marvin Minsky. Telepresence. 1980.

[2] S.O. Adalgeirsson and C. Breazeal. Mebot: a robotic platform for socially embodied presence. In *Proceeding of the 5th ACM/IEEE international conference on Human-robot interaction*, pages 15–22. ACM, 2010.

[3] E. Paulos and J. Canny. Prop: personal roving presence. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 296–303. ACM Press/Addison-Wesley Publishing Co., 1998.

[4] K.M. Tsui, M. Desai, H.A. Yanco, and C. Uhlik. Exploring use cases for telepresence robots. In *Proceedings of the 6th international conference on Human-robot interaction*, pages 11–18. ACM, 2011.

[5] Min Kyung Lee and Leila Takayama. Now, i have a body: Uses and social norms for mobile remote presence in the workplace. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, pages 33–42. ACM, 2011.

[6] Daisuke Sakamoto, Takayuki Kanda, Tetsuo Ono, Hiroshi Ishiguro, and Norihiro Hagita. Android as a telecommunication medium with a human-like presence. In *Human-Robot Interaction (HRI), 2007 2nd ACM/IEEE International Conference on*, pages 193–200. IEEE, 2007.

[7] Masahiro Mori. The uncanny valley. *Energy*, 7(4):33–35, 1970.

[8] Peter Lincoln, Greg Welch, Andrew Nashel, Adrian Ilie, Henry Fuchs, et al. Animatronic shader lamps avatars. In

Mixed and Augmented Reality, 2009. IS-MAR 2009. 8th IEEE International Symposium on, pages 27–33. IEEE, 2009.

[9] Susumu Tachi, Kiyoshi Komoriya, Kazuya Sawada, Takashi Nishiyama, Toshiyuki Itoko, Masami Kobayashi, and Kozo Inoue. Telexistence cockpit for humanoid robot control. *Advanced Robotics*, 17(3):199–217, 2003.

[10] Yang Xiao, Zhijun Zhang, Aryel Beck, Junsong Yuan, and Daniel Thalmann. Human-robot interaction by understanding upper body gestures.

[11] Zerrin Yumak, Jianfeng Ren, Nadia Magnenat Thalmann, and Junsong Yuan. Modelling multi-party interactions among virtual characters, robots, and humans. *PRESENCE: Teleoperators and Virtual Environments*, 23(2):172–190, 2014.

[12] Peter A Hancock, Deborah R Billings, Kristin E Schaefer, Jessie YC Chen, Ewart J De Visser, and Raja Parasuraman. A meta-analysis of factors affecting trust in human-robot interaction. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 53(5):517–527, 2011.

[13] Arjun Nagendran, Remo Pillat, Adam Kavanaugh, Greg Welch, and Charles Hughes. Amities: avatar-mediated interactive training and individualized experience system. In *Proceedings of the 19th ACM Symposium on Virtual Reality Software and Technology*, pages 143–152. ACM, 2013.

[14] Arjun Nagendran, Remo Pillat, Adam Kavanaugh, Greg Welch, and Charles Hughes. A unified framework for individualized avatar-based interactions. *Presence: Teleoperators and Virtual Environments*, 23(2):109–132, 2014.