

Exploring Interaction Design Considerations for Trustworthy Language-Capable Robotic Wheelchairs in Virtual Reality

Nicholas Woodward
nwoodward@mines.edu
Colorado School of Mines
Golden, Colorado

Carter Fowler
carterjfowler@mines.edu
Colorado School of Mines
Golden, Colorado

Stephen Thoenmes
sthoennes@mines.edu
Colorado School of Mines
Golden, Colorado

Teresa Nguyen
teresamysinhnguyen@mines.edu
Colorado School of Mines
Golden, Colorado

Taewoo Kim
taewookim111@gmail.com
Colorado School of Mines
Golden, Colorado

Tom Williams
twilliams@mines.edu
Colorado School of Mines
Golden, Colorado

Lixiao Zhu
lizhu@mines.edu
Colorado School of Mines
Golden, Colorado

Stacia Near
stacia.near@gmail.com
Colorado School of Mines
Golden, Colorado

ABSTRACT

In previous work, researchers in Human-Robot Interaction (HRI) have demonstrated that user trust in robots depends on effective and transparent communication. This may be particularly true for robots used for transportation, due to user reliance on such robots for physical movement and safety. In this paper, we present the design of an experiment examining the importance of proactive communication by robotic wheelchairs, as compared to non-vehicular mobile robots, within a Virtual Reality (VR) environment. Furthermore, we describe the specific advantages – and limitations – of conducting this type of HRI experiment in VR.

KEYWORDS

virtual reality, trust, autonomous vehicles, robotic wheelchairs, human-robot interaction

1 INTRODUCTION

Over the past two decades, there has been a significant amount of research into developing autonomous robotic wheelchairs [16] to assist the growing elderly and disabled population [38]. Because many wheelchair users have difficulty using a traditional wheelchair joystick, many researchers have been investigating *natural language capable* robotic wheelchairs capable of flexibly executing high-level natural language directives given by users [59]. Accordingly, much of this research has focused on how language-capable robotic wheelchairs can understand and execute users' directives [58]. In contrast, however, little attention has been paid to how language-capable robotic wheelchairs should speak to their users.

We argue that human-robot dialogue with language capable robotic wheelchairs is subject to a fundamentally different set of requirements than dialogue with other types of mobile and social robots:

Wheelchairs are vehicles: Robotic wheelchairs, as vehicles, are (we believe) more likely to be viewed as tools designed to provide a specific functionality than as agents with whom the user can develop a social connection. This stance is likely

to be reinforced by wheelchairs' embodiment, which forces users to view their environment from the same perspective as their wheelchair at all times rather than facing each other [cp. 25]). As such, robotic wheelchairs may need to expect to be communicated with primarily using well-specified goal-driven language, as opposed to free-form social conversation. Similarly, this lack of perceived agency may lead to requirements for how often and verbosely the wheelchair communicates with its users of its own initiative.

Wheelchairs are constant companions: Unlike other robot platforms, vehicular or otherwise, wheelchairs are unique in that they are in constant use and physical contact with their users and typically answer only to their immediate user and not other agents in the environment.. As such, robotic wheelchairs that *do* exhibit socially agentic behaviors may stand to develop (or need to develop) significantly closer relationships with their users than would other types of robots.

Paired Autonomy: Unlike when interacting with other mobile robots, when using a robotic wheelchair, the user surrenders their own physical autonomy to the robot, resulting in a system of paired autonomy, where the user's actions directly impact the wheelchair, and the wheelchair's actions directly impact the user.

These factors create a complex dynamic in which constant co-location and paired autonomy may override the traditionally asocial nature of autonomous vehicles, resulting in increased need for establishment of trust and rapport between wheelchairs and their users. We expect that these dynamics will lead to unique communication requirements for robotic wheelchairs, in which they will need to take more trust-building communicative actions, such as checking in frequently and regularly [10] or sharing information and explaining unexpected actions proactively [52].

In our current research, we are specifically investigating the unique requirements language capable robotic wheelchairs may

face with respect to *proactive communication*. As an illustrative example, consider the following two scenarios:

The Handcart

Alice is following her language-capable robotic handcart along a path they've travelled many times. Every previous time traveling this route, her handcart has driven straight down a long hallway without taking any turns. This time, however, the robot abruptly and unexpectedly turns, veering down a hallway they've never been down before.

The Wheelchair

Alice is riding in her language-capable robotic wheelchair along a path they've travelled many times. Every previous time traveling this route, her wheelchair has driven straight down a long hallway without taking any turns. This time, however, the robot abruptly and unexpectedly turns, veering down a hallway she has never been down before.

Here, the only difference between the two scenarios is whether Alice is following or riding in the mobile robot in question. And yet, we argue that Alice's response should be radically different in the two scenarios. In "The Handcart", we would expect Alice to be surprised and curious when her handcart turns down an unfamiliar hallway unexpectedly. On the other hand, in "The Wheelchair", we would expect Alice to be frightened and anxious, because her own autonomy and well being are being put at risk.

In both of these cases, Alice's emotional response could have been allayed if the robot had generated, in the period proceeding its unexpected turn, a *proactive response*, e.g. "Just to let you know, I've been informed that the hallway on our normal route is inaccessible, so I need to make a detour down a hallway on the left up ahead." Because the negative consequences of taking this turn unannounced are expected to be so much more negative in the case of the wheelchair than the handcart (i.e., fear and anxiety versus surprise and curiosity), we thus believe it will be much more important and impactful for robotic wheelchairs to generate these sorts of explanations than it will be for traditional non-vehicular mobile robots.

If this hypothesis is experimentally supported, it would represent a key insight to inform the dialogue policies of robotic wheelchairs. Moreover, we might expect this finding to carry over to inform the design of other types of autonomous vehicles, such as autonomous cars. Unfortunately, assessing this hypothesis is difficult in a traditional experimental context, for several reasons: (1) no autonomous robotic wheelchairs currently exist on the market, and the construction of an autonomous language-capable autonomous wheelchair is a significant engineering challenge¹. (2) Autonomous vehicles, including wheelchairs, are incredibly dangerous, being able to move at high speeds and cause significant damage to both walls and people². This is also obviously the case for other types of autonomous vehicles as well, such as autonomous cars, which is why other research groups have sought to investigate trust in autonomous vehicles within very limited carefully controlled contexts such as

autonomous parking [51]. (3) Testing this hypothesis in a *controlled* real-world environment would require maintaining full control over a network of hallways, preventing any of the building's occupants from entering those hallways, and the ability to artificially orchestrate realistic³ excuses for detours such as blockades and spills.

We argue that these challenges can be addressed by moving experiments into Virtual Reality environments: VR enables presentation of environments that are carefully controllable and manipulable, and in which users can be placed into situations in which they can be led to feel nervous or anxious despite being placed at no actual risk [32]. Accordingly, in this paper we present the design of a human-robot interaction experiment conducted in Virtual Reality in which we explore the benefits of proactive communication for participants who are either seated in a robotic wheelchair or who are working alongside a non-vehicular mobile robot. We address in particular the decisions made in designing this experiment and VR environment, and the strengths and limitations of this approach.

2 PREVIOUS WORK

2.1 Trust in Machines, Robots and Vehicles

Trust plays an important role in facilitating task success for both human-human and human-machine teams, especially in time critical contexts [4]. As shown by Robinette et al. [44], poor decisions made by autonomous systems significantly lowers human trust, and mistakes must be promptly addressed and corrected to repair trust [43]. As such, there has been significant research seeking to quantify or estimate human-robot trust [5, 17, 26, 40, 46, 47, 54].

These efforts have revealed the wide variety of *types* of trust that can be developed between humans and robots. In recent work, Ullman and Malle [54] have argued for a multi-dimensional measure of trust with two broad components: reliability/capability, and ethicality/sincerity. The first dimension, reliability/capability, aligns with insights from the more general human-machine teaming literature that emphasizes machine predicatability [34]. The second dimension, ethicality/sincerity, aligns instead with insights from the interpersonal trust literature [22]. This dichotomy also echoes recent efforts to bring together philosophical and psychological definitions of trust by scholars such as Danks [11], who differentiates between behavioral trust (grounded in predictability) and understanding trust (grounded in beliefs of suitable dispositions). Ultimately, we believe that interpersonal, ethical/sincerity-based trust, and understanding trust are all more likely to be of importance and more likely to be developed for robots with greater perceived social agency [24, 53] and for robots that more frequently prompt interactants to take the intentional stance [12].

In this paper, we focus specifically on requirements for trust between humans and language-capable intelligent wheelchairs, and how these differ from requirements for trust between humans and non-vehicular mobile robots. Trust in autonomous vehicles in particular has also been an area of significant recent study [9, 13, 15, 51], as autonomous vehicle companies attempt to convince consumers and investors of the promise of the technology in the wake of high-profile Tesla accidents. Autonomous wheelchairs,

¹The last author can attest to this based on his own experiences in graduate school.

²The last author can *also* attest to this based on his experiences in graduate school.

³In the case of the robotic handcart, there's no reason why Alice could not continue the traditional route herself, even after receiving a proactive explanation.

unlike traditional autonomous vehicles, stand to be much closer to constant companions than mere vehicles, and this is especially true for future language-capable robotic wheelchairs. But while there has been some work on developing language-capable robotic wheelchairs with high level cognitive capabilities [58], the vast majority of language-capable wheelchairs only respond to low-level metric commands such as “go forward”, “turn left”, and “stop” [59]. As such, we believe that most robotic wheelchairs will ultimately be perceived more as vehicles or machines than as intentional agents, and as such, reliability/capability trust will be more relevant to examine in this work.

2.2 Proactive and Reactive Explanations

Foundational work in human factors has clearly demonstrated the importance of *transparency* for human-machine trust [30, 31]: appropriate trust necessitates understanding of past performance history, current process, and purpose towards future action [7]. However, it is also well understood that transparency can come at a cost of increased operator workload [21]. And to make this calculus more complex, the relationship between transparency and trust is nonlinear: when a system is trustworthy, transparency can of course lead to increased trust, but when a system is already highly trusted – especially when that high level of trust is not warranted – transparency can instead decrease trust [1]. These factors conspire to create a difficult engineering challenge in which transparency must be carefully employed to jointly optimize trust and workload.

One of the most effective ways that interactive robots can carefully and intentionally manage transparency is by selectively generating and tailoring the content of natural language explanations [56]. For example, while humans may distrust a robot that makes mistakes, giving an explanation regarding *why* the robot might make such mistakes can help to increase trust to an appropriate level [14].

Most previous work on robot explanation generation, including work in autonomous vehicles, has focused on explanations generated on demand for actions already executed; what we refer to as *reactive* explanations. In contrast, recent work by Koo et al. [27] focuses on *proactive* explanations generated before actions are taken, especially when actions are expected to be surprising or alarming, and analyzes the effects that different kinds of explanations can have on human-vehicle trust. However, this previous work has not explored the difference in importance of proactive communicative actions between autonomous vehicles and non-vehicular mobile robots. And moreover, to the best of our knowledge, there has been no previous exploration of the importance of proactive communication for robotic wheelchairs specifically, which as we have described occupy a unique niche within the landscape of interactive robots.

2.3 VR for Human-Robot Interaction

Researchers have been recently investigating several different uses for VR technologies to facilitate HRI research [60, 61]. Most relevant are approaches that use VR as a testbed for exploring interactions with vehicular robots, both robotic wheelchairs specifically [55] as well as traditional autonomous vehicles [18, 45, 48, 49]. Also relevant are approaches using VR to study interaction with other sorts of dangerous robots, such as robots in meat-packing plants [19].

There have also been approaches using VR for other purposes, e.g. providing training environments for robots [23, 50, 57], visualizing robots in large-scale maritime or aerial environments [20, 29, 36], and providing immersive teleoperation interfaces [2, 3, 6, 37, 41, 62].

2.4 VR Design Considerations

A key difference between robotic wheelchairs and non-vehicular mobile robots in virtual reality is user motion control, as users accompanying non-vehicular mobile robots in virtual environments need some way to propel themselves through the environment. Previous studies have investigated a variety of movement forms in virtual reality, taking into account factors including realism, possibility to act, quality of the interface, possibility to examine the environment, and self-evaluation of performance [39]. Users accompanying non-vehicular mobile robots must be able to freely move around a virtual reality environment, with minimized potential for motion-sickness and maximized potential for immersion. There are several possible options for satisfying these criteria:

Manual translation requires the participant to use a controller and physical buttons to move throughout a scene. This movement system typically emulates those utilized in non-virtual reality video games. Due to the lack of physical movement paired with the digital translation of the participant, manual translation detracts from the immersion of a virtual reality simulation. Motion-sickness is caused by dissonance between a user’s mental understanding of how they are moving and the reality of their movement; in other words, manual translation does little to pair physical actions with digital movement, resulting in a greater potential for motion sickness [8].

Teleportation requires the participant to use an HTC Vive controller to point at a location, after which they are immediately moved to that target location. Teleportation is a more commonly-used locomotion mechanism in virtual reality simulations, using a point-and-click mechanism to move the user. While studies have shown that teleportation reduces the potential for motion sickness and requires little physical demand [33], this method has low capacity for immersion; using teleportation to follow a robot allows users to move extremely rapidly, which not only fails to accurately emulate the experience of walking behind a robot, but also may lead to reduced spatial awareness [35] and situational awareness [28].

ArmSwinger movement requires the participant to swing two HTC Vive controllers back and forth in a motion similar to that of walking. ArmSwinger was designed for the purpose of pairing a physical movement with movement in VR. Each ‘swing’ of the arm constitutes one stride forward in the simulation, with the speed and distance governed by the speed and distance of the swing. The participant moves forward in the direction that their whole body is facing. This means that the participant must be more physically present in order to move, fully turning their body and swinging their arms to move around the scene, as opposed to the relatively passive movement styles of manual translation and teleportation. As such, ArmSwinger offers a movement method that

maintains a relatively high level of immersion, while moderately reducing the potential for motion sickness by utilizing physical movement to function [39].

3 METHOD

3.1 Hypotheses

In this experiment, we are particularly interested in the *Reliability* and *Capability* dimensions of human-robot trust rather than *Ethicality* and *Sincerity* dimensions (cf. [54]). Accordingly, we constructed the following hypotheses:

- H1a** Participants will perceive the robot as more reliable if it offers proactive explanations for its actions...
- H1b** ...and any drop in perceived reliability after failure to proactively explain a decision will be greater for wheelchairs than for non-vehicular mobile robots.
- H2a** Participants will perceive the robot as more capable if it offers proactive explanations for its actions...
- H2b** ...and any drop in perceived capability after failure to proactively explain a decision will be greater for wheelchairs than for non-vehicular mobile robots.
- H3a** Participants will perceive the robot as more trustworthy across both of the above dimensions if it offers proactive explanations for its actions...
- H3b** ...and any drop in perceived trustworthiness after failure to proactively explain a decision will be greater for wheelchairs than for non-vehicular mobile robots.
- H4** For robotic wheelchairs specifically, participants will perceive the robot more positively in terms of usability if it offers proactive explanations for its actions.

3.2 Measures

To assess H1 and H2 we used a subset of items from Ullman and Malle [54]’s Reliability and Capability trust scales [Forthcoming] (with reliability items used to assess H1, capability items used to assess H2, and all items in aggregate used to assess H3). In addition, two additional trust-related survey items were used to differentiate between predictability-based and understanding-based trust [11]. Finally, to assess H4, we used two simple wheelchair-oriented usability questions. All survey items were delivered as Likert items on a scale from 1 to 5 (1 = disagree, 5 = agree) with participants asked to answer the questions verbally at set points during the experiment, as described below. All survey items are listed in the Appendix.

3.3 Experimental Design

Our study uses a between-subjects Wizard-of-Oz design [42] in which participants interacted with one of two types of robots, which communicated using one of two communication styles.

The two communication styles were *reactive* and *proactive*:

Reactive Condition: In the *reactive* condition, the robot would only generate an explanation for a behavior if explicitly asked for one by the participant.


Proactive Condition: In the *proactive* condition, the robot generated explanations for upcoming actions likely to be surprising

approximately 3-4 seconds before execution.


The two virtual robots used were a *robotic wheelchair* (Fig. 1a) and a *non-vehicular mobile robot* (Fig. 1b).

Wheelchair Condition: In the wheelchair condition, the participant’s point of view is as though they were seated in a wheelchair. The participant is able to look around by moving their head, but is unable to directly control the movement of the wheelchair. Instead, the participant can only move themselves by issuing voice commands to the wheelchair, telling it where to go.

Non-vehicular mobile robot Condition: In the non-vehicular mobile robot condition, the participant works alongside a mobile robot. While the user must still use voice commands to tell this mobile robot where to go, their own position is not tied to that of the robot, and they are instead able to travel freely using ArmSwinger locomotion.



(a) Wheelchair



(b) Non-vehicular Mobile Robot

Figure 1: VR Design for the (a) Wheelchair and (b) Non-vehicular Mobile Robot

3.4 VR Design

The VR environment was designed as a two-player game using the Unity game engine. The first player, the participant, was seated at a desktop computer in a large experiment room, and fitted with an HTC Vive head-mounted display (HMD) and two HTC Vive hand controllers. The second player, the Wizard, was unbeknownst to participants, seated in an adjacent Wizard Control room, at another desktop computer, where they could control the robot’s movements through the Unity game interface on that desktop. User audio was recorded through the Vive headset and streamed to the Wizard using Zoom, which was then also used to record audio and video of the participant’s session.

3.5 Environment Design

A diagram depicting the VR environment is shown in Fig. 2. The environment consisted of three main areas: a front desk, and two sets of offices (1A-1E and 2A-2E), with the first set of offices accessible either through a direct route (which runs past the second set of offices) or an indirect, roundabout route (which does not).

3.6 Task Design

Our experiment requires participants to execute a series of delivery tasks in a VR environment. In this task, participants start at the front desk, and are given (through a pop-up window) a target office




Figure 2: Paths Taken by Participant During Experiment

to which a document needs to be delivered. Participants then need to command the robot to travel to that location using a command like “Go to room 1B”. Upon receiving this command, the robot travels to this room. In the Non-vehicular mobile robot condition, the participant must manually move to follow the robot as it makes its way to this location. Upon arriving at the target room, the user is again given (through a pop-up window) an office to which a new report must be delivered, and they must again give a command, like “Go to room 2C”. Upon arriving at the target room, the user is instructed to return to the front desk for a new assignment.

This sequence is repeated three times. During the first two sequences, the robot travels according to the direct route, indicated using the red and blue lines in Fig. 2. On the third sequence, the robot instead takes the detour depicted using the orange line in Fig. 2. The explanation for this detour, given reactively or proactively according to condition, is that there was a “water spill” on the original path (in the hallway junction between 1B and 1C) that prevented the normal route from being taken. This water spill was visualized on the ground at this location so that if participants in the non-vehicular mobile robot condition went to investigate the spill, they would observe it for themselves. This detour served as a “trust reducer” similar to those employed in [48]’s work on human-vehicle trust in VR.

The first nine (trust-related) survey items described in Sec. 3.2 and listed in the Appendix were administered each time the participant returned to the front desk (i.e., after each of the three delivery sequences). This allowed us to establish both baseline levels of trust in the robot, as well as the level of trust after the detour was taken, to allow us to assess our hypotheses. The additional usability questions were also administered on the last of these three surveys.

4 DISCUSSION AND CONCLUSION

In this paper, we have described how VR technologies can be used to examine interactions with autonomous vehicles such as robotic wheelchairs, and presented the design of a human-subject experiment intended to assess how communication policies may need to differ between robotic wheelchairs and non-vehicular mobile robots, in order to maintain human-robot trust.

Our use of VR for this study presents a number of advantages. VR makes for a convenient method for setting up experiments due to increased experimental control and decreased hardware costs. Moreover, VR allows for investigation of high-risk contexts that could not otherwise be investigated due to risk of bodily harm. This being said, questions may of course be raised as to how closely findings from studies such as that which we have proposed will carry over into reality, especially in environments such as that used in our experiment, which are not photorealistic and thus not highly immersive. Moreover, the use of VR environments does of course present other physical risks, including motion sickness for some participants.

Upon completion of this experiment, our next step will be to analyze our collected data in order to assess our hypotheses. In future work, we also hope to investigate how interaction design requirements for robotic wheelchairs will change along with human perceptions thereof; i.e., we are interested in determining how robotic wheelchairs should frame their explanations, and the importance of those explanations for human-robot trust, depending on whether they are perceived as more socially agentic, viewed more frequently from the intentional stance, and considered more as companions than as vehicles.

ACKNOWLEDGMENTS

This work was supported in part by an Early Career Faculty grant from NASA’s Space Technology Research Grants Program and by NSF Grant CNS-1810631.

REFERENCES

- [1] Kumar Akash, Tahira Reid, and Neera Jain. 2019. Improving Human-Machine Collaboration Through Transparency-based Feedback—Part II: Control Design and Synthesis. *IFAC-PapersOnLine* 51, 34 (2019), 322–328.
- [2] Jordan Allspaw, Lilia Heimold, and Holly A Yanco. 2019. Design of Virtual Reality for Humanoid Robots with Inspiration from Video Games. In *International Conference on Human-Computer Interaction*. Springer, 3–18.
- [3] Maxwell Bennett, Tom Williams, Daria Thames, and Matthias Scheutz. 2017. Differences in interaction patterns and perception for teleoperated and autonomous humanoid robots. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 6589–6594.
- [4] Deborah R Billings, Kristin E Schaefer, Jessie YC Chen, and Peter A Hancock. 2012. Human-robot interaction: developing trust in robots. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. 109–110.
- [5] George Charalambous, Sarah Fletcher, and Philip Webb. 2016. The development of a scale to evaluate trust in industrial human-robot collaboration. *International Journal of Social Robotics* 8, 2 (2016), 193–209.
- [6] Junshen Chen, Marc Glover, Chenguang Yang, Chunxu Li, Zhijun Li, and Angelo Cangelosi. 2017. Development of an immersive interface for robot teleoperation. In *Annual Conference Towards Autonomous Robotic Systems*. Springer, 1–15.
- [7] Jessie YC Chen, Shan G Lakhmani, Kimberly Stowers, Anthony R Sellkowitz, Julia L Wright, and Michael Barnes. 2018. Situation awareness-based agent transparency and human-autonomy teaming effectiveness. *Theoretical issues in ergonomics science* 19, 3 (2018), 259–282.
- [8] W. Chen, J.Z. Chen, and R.H.Y. So. 2011. *Visually Induced Motion Sickness: Effects of Translational Visual Motion Along Different Axes*. CRC Press, 281–287.
- [9] Jong Kyu Choi and Yong Gu Ji. 2015. Investigating the importance of trust on adopting an autonomous vehicle. *International Journal of Human-Computer*

- Interaction* 31, 10 (2015), 692–702.
- [10] CB Crisp and SL Jarvenpaa. 2000. Trust over time in global virtual teams. In *Academy of Management Meeting*.
 - [11] David Danks. 2019. The Value of Trustworthy AI. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 521–522.
 - [12] Daniel Clement Dennett. 1989. *The intentional stance*. MIT press.
 - [13] Murat Dikmen and Catherine Burns. 2017. Trust in autonomous vehicles: The case of tesla autopilot and summon. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 1093–1098.
 - [14] Mary T Dzindolet, Scott A Peterson, Regina A Pomranky, Linda G Pierce, and Hall P Beck. 2003. The role of trust in automation reliance. *International journal of human-computer studies* 58, 6 (2003), 697–718.
 - [15] Mica R Endsley. 2017. Autonomous driving systems: A preliminary naturalistic study of the Tesla Model S. *Journal of Cognitive Engineering and Decision Making* 11, 3 (2017), 225–238.
 - [16] Ahmetcan Erdogan and Brenna D. Argall. 2017. The effect of robotic wheelchair control paradigm and interface on user performance, effort and preference: An experimental assessment. *Robotics and Autonomous Systems* 94 (2017), 282 – 297. <https://doi.org/10.1016/j.robot.2017.04.013>
 - [17] Amos Freedy, Ewart DeVisser, Gershon Weltman, and Nicole Coeyman. 2007. Measurement of trust in human-robot collaboration. In *2007 International Symposium on Collaborative Technologies and Systems*. IEEE, 106–114.
 - [18] David Goedicke, Jany Li, Vanessa Evers, and Wendy Ju. 2018. VR-OOM: Virtual Reality On-road driving simulation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 165.
 - [19] Laerke I. N. Hansen, Niklas Vinther, Lukas Stranovsky, Mark P. Philipsen, Haiyan Wu, and Thomas B. Moeslund. 2018. Collaborative Meat Processing in Virtual Reality. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*.
 - [20] Kerstin S Haring, Victor Finomore, Dylan Muramoto, Nathan L. Tenhundfeld, J. Wen, and B. Tidball. 2018. Analysis of using virtual reality (vr) for command and control applications of multi-robot systems. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*.
 - [21] Tove Helldin. 2014. *Transparency for Future Semi-Automated Systems: Effects of transparency on operator performance, workload and trust*. Ph.D. Dissertation. Örebro Universitet.
 - [22] RW Hunt, PM Kohn, and CB Mallozzi. 1983. Factor analysis of the interpersonal trust scale with a noncollege population. *Journal of personality assessment* 47, 5 (1983), 507.
 - [23] Michael L Iuzzolino, Michael E Walker, and Daniel Szafrir. 2018. Virtual-to-Real-World Transfer Learning for Robots on Wilderness Trails. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 576–582.
 - [24] Ryan Blake Jackson and Tom Williams. 2019. On Perceived Social and Moral Agency in Natural Language Capable Robots. In *2019 HRI Workshop on The Dark Side of Human-Robot Interaction*. Jackson, RB, and Williams. 401–410.
 - [25] Peter H Kahn, Nathan G Freier, Takayuki Kanda, Hiroshi Ishiguro, Jolina H Ruckert, Rachel L Severson, and Shaun K Kane. 2008. Design patterns for sociality in human-robot interaction. In *Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction*. 97–104.
 - [26] Poornima Kaniarasu, Aaron Steinfeld, Munjal Desai, and Holly Yanco. 2012. Potential measures for detecting trust changes. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 241–242.
 - [27] Jeamin Koo, Jungsuk Kwac, Martin Steinert, Larry Leifer, and Clifford Nass. 2014. Why did my car just do that? Explaining semi-autonomous driving actions to improve driver understanding, trust, and performance. *International Journal on Interactive Design and Manufacturing (IJIDeM)* 9 (01 2014). <https://doi.org/10.1007/s12008-014-0227-2>
 - [28] Danielle Krukowski. 2019. *The Effect of Object Selection and Operator Movement on Situation Awareness in Virtual Reality*. Ph.D. Dissertation. Carleton University.
 - [29] Mårten Lager, Elin A. Topp, and Jacek Malec. 2018. Remote Operation of Unmanned Surface Vessel through Virtual Reality - a low cognitive load approach. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*.
 - [30] JD Lee. 2012. Trust, trustworthiness, and trustability. In *Presentation at the Workshop on Human Machine Trust for Robust Autonomous Systems*, Vol. 31.
 - [31] John D Lee and Katrina A See. 2004. Trust in automation: Designing for appropriate reliance. *Human factors* 46, 1 (2004), 50–80.
 - [32] Jih-Hsuan Tammy Lin. 2017. Fear in virtual reality (VR): Fear elements, coping reactions, immediate and next-day fright responses toward a survival horror zombie virtual reality game. *Computers in Human Behavior* 72 (2017), 350–361.
 - [33] Guillaume Loup and Emilie Loup-Escande. 2019. Effects of Travel Modes on Performances and User Comfort: A Comparison between ArmSwinger and Teleporting. *International Journal of Human-Computer Interaction* 35, 14 (2019), 1270–1278. <https://doi.org/10.1080/10447318.2018.1519164> arXiv:https://doi.org/10.1080/10447318.2018.1519164
 - [34] Joseph B. Lyons, Kevin T. Wynne, Sean Mahoney, and Mark A. Roebke. 2019. Chapter 6 - Trust and Human-Machine Teaming: A Qualitative Study. In *Artificial Intelligence for the Internet of Everything*, William Lawless, Ranjeev Mittu, Donald Sofge, Ira S. Moskowitz, and Stephen Russell (Eds.). Academic Press, 101 – 116. <https://doi.org/10.1016/B978-0-12-817636-8.00006-5>
 - [35] Kasra Rahimi Moghadam, Colin Banigan, and Eric D Ragan. 2018. Scene transitions and teleoperation in virtual reality and the implications for spatial awareness and sickness. *IEEE transactions on visualization and computer graphics* (2018).
 - [36] Michael Novitzky, Michael R. Benjamin, Paul Robinette, Hugh R.R. Dougherty, Caleigh Fitzgerald, and Henrik Schmidt. 2018. Virtual Reality for Immersive Simulated Experiments of Human-Robot Interactions in the Marine Environment. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*.
 - [37] Yeonju Oh, Ramviyas Parasuraman, Tim McGraw, and Byung-Cheol Min. 2018. 360 vr based robot teleoperation interface for virtual tour. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*.
 - [38] Jennifer M Ortman, Victoria A Velkoff, Howard Hogan, et al. 2014. *An aging nation: the older population in the United States*. United States Census Bureau, Economics and Statistics Administration, US
 - [39] Yun Suen Pai and Kai Kunze. 2017. Armswing: using arm swings for accessible and immersive navigation in AR/VR spaces. 189–198. <https://doi.org/10.1145/3152832.3152864>
 - [40] Eui Park, Quaneisha Jenkins, and Xiaochun Jiang. 2008. Measuring trust of human operators in new generation rescue robots. In *Proceedings of the JFPS International Symposium on Fluid power*, Vol. 2008. The Japan Fluid Power System Society, 489–492.
 - [41] Corey Pittman and Joseph J LaViola Jr. 2014. Exploring head tracked head mounted displays for first person robot teleoperation. In *Proceedings of the 19th international conference on Intelligent User Interfaces*. 323–328.
 - [42] Laurel D Riek. 2012. Wizard of oz studies in hri: a systematic review and new reporting guidelines. *Journal of Human-Robot Interaction* 1, 1 (2012), 119–136.
 - [43] Paul Robinette, Ayanna M Howard, and Alan R Wagner. 2015. Timing is key for robot trust repair. In *International Conference on Social Robotics*. Springer, 574–583.
 - [44] Paul Robinette, Ayanna M. Howard, and Alan R. Wagner. 2017. Effect of Robot Performance on Human–Robot Trust in Time-Critical Situations. *IEEE Transactions on Human-Machine Systems* 47, 4 (2017), 425–436. <https://doi.org/10.1109/THMS.2017.2648849>
 - [45] Sandro Ropelato, Fabio Zünd, Stéphane Magnenat, Marino Menozzi, and Robert Sumner. 2018. Adaptive tutoring on a virtual reality driving simulator. *International SERIES on Information Systems and Management in Creative EMedia (CreMedia)* 2017, 2 (2018), 12–17.
 - [46] Kristin Schaefer. 2013. The perception and measurement of human-robot trust. (2013).
 - [47] Kristin E. Schaefer. 2016. *Measuring Trust in Human Robot Interactions: Development of the “Trust Perception Scale-HRI”*. Springer US, Boston, MA, 191–218. https://doi.org/10.1007/978-1-4899-7668-0_10
 - [48] Shervin Shahrard, Corey Park, and Mehrdad Nojournian. 2019. Human Trust Measurement Using an Immersive Virtual Reality Autonomous Vehicle Simulator. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (Honolulu, HI, USA) (AIES ’19)*. Association for Computing Machinery, New York, NY, USA, 515–520. <https://doi.org/10.1145/3306618.3314264>
 - [49] Daniele Sportillo, Alexis Paljic, Mehdi Boukhris, Philippe Fuchs, Luciano Ojeda, and Vincent Roussarie. 2017. An immersive Virtual Reality system for semi-autonomous driving simulation: a comparison between realistic and 6-DoF controller-based interaction. In *Proceedings of the 9th International Conference on Computer and Automation Engineering*. ACM, 6–10.
 - [50] Francesca Stramandinoli, Kin Gwn Lore, Jeffrey R. Peters, Paul C. O’Neill, Binu M. Nair, Richa Varma, Julian C. Ryde, Jay T. Miller, and Kishore K. Reddy. 2018. Robot Learning from Human Demonstration in Virtual Reality. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*.
 - [51] Nathan L Tenhundfeld, Ewart J de Visser, Anthony J Ries, Victor S Finomore, and Chad C Tossell. 2019. Trust and distrust of automated parking in a Tesla Model X. *Human factors* (2019), 0018720819865412.
 - [52] James B Thomas and Linda Klebe Trevino. 1993. Information processing in strategic alliance building: A multiple-case approach. *Journal of Management Studies* 30, 5 (1993), 779–814.
 - [53] Daniel Ullman, Lolanda Leite, Jonathan Phillips, Julia Kim-Cohen, and Brian Scassellati. 2014. Smart human, smarter robot: How cheating affects perceptions of social agency. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 36.
 - [54] Daniel Ullman and Bertram F Malle. 2019. Measuring gains and losses in human-robot trust: evidence for differentiable components of trust. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 618–619.
 - [55] Guillaume Vailland, Yoren Gaffary, Louise Devigne, Valérie Gouranton, Bruno Arnaldi, and Marie Babel. 2020. Vestibular Feedback on a Virtual Reality Wheelchair Driving Simulator: A Pilot Study. In *ACM/IEEE International Conference on Human-Robot Interaction, HRI 2020*.

- [56] Ning Wang, David V Pynadath, and Susan G Hill. 2016. Trust calibration within a human-robot team: Comparing automatically generated explanations. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 109–116.
- [57] David Whitney, Eric Rosen, and Stefanie Tellex. 2018. Learning from Crowdsourced Virtual Reality Demonstrations. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*.
- [58] Tom Williams, Collin Johnson, Matthias Scheutz, and Benjamin Kuipers. 2017. A Tale of Two Architectures: A Dual-Citizenship Integration of Natural Language and the Cognitive Map. In *Proceedings of the 16th International Conference on Autonomous Agents and Multi-Agent Systems*.
- [59] Tom Williams and Matthias Scheutz. 2017. The state-of-the-art in autonomous wheelchairs controlled through natural language: A survey. *Robotics and Autonomous Systems* 96 (2017), 171–183.
- [60] Tom Williams, Daniel Szafr, and Tathagata Chakraborti. 2019. The Reality-Virtuality Interaction Cube. In *Proceedings of the 2nd International Workshop on Virtual, Augmented, and Mixed Reality for HRI*.
- [61] Tom Williams, Daniel Szafr, Tathagata Chakraborti, and Heni Ben Amor. 2018. Virtual, augmented, and mixed reality for human-robot interaction. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. 403–404.
- [62] Tom Williams, Nhan Tran, Josh Rands, and Neil T Dantam. 2018. Augmented, mixed, and virtual reality enabling of robot deixis. In *International Conference on Virtual, Augmented and Mixed Reality*. Springer, 257–275.

APPENDIX: SURVEY QUESTIONS

- (1) Reliability:
 - The robot is reliable.
 - The robot is predictable.
 - The robot can be counted on.
 - The robot is consistent.
- (2) Capability:
 - The robot is capable.
 - The robot is skilled.
 - The robot is competent.
- (3) Additional Trust Items:
 - I can predict what the robot will do when I give it a command.
 - I understand why the robot is doing what it is doing.
- (4) Usability:
 - I would like to work with the robot again
 - I would choose this wheelchair over a manually driven one