Your Way Or My Way: Improving Human-Robot Co-Navigation Through Robot Intent and Pedestrian Prediction Visualisations

Xinyan Yu
xinyan.yu@sydney.edu.au
Design Lab, Sydney School of Architecture, Design and Planning
The University of Sydney
Sydney, NSW, Australia

Marius Hoggenmueller
marius.hoggenmueller@sydney.edu.au
Design Lab, Sydney School of Architecture, Design and Planning
The University of Sydney
Sydney, NSW, Australia

Martin Tomitsch
martin.tomitsch@sydney.edu.au
Design Lab, Sydney School of Architecture, Design and Planning
The University of Sydney
Sydney, NSW, Australia

ABSTRACT

As mobile robots enter shared urban spaces, operating in close proximity to people, this raises new challenges in terms of how these robots communicate with passers-by. Following an iterative process involving expert focus groups (n=8), we designed an augmented reality concept that visualises the robot's navigation intent and the pedestrian's predicted path. To understand the impact of path visualisations on trust, sense of agency, user experience, and robot understandability, we conducted a virtual reality evaluation (n=20). We compared visualising both robot intent and pedestrian path prediction against just visualising robot intent and a baseline without augmentation. The presence of path visualisations resulted in a significant improvement of trust. Triangulation of quantitative and qualitative results further highlights the impact of pedestrian path prediction visualisation on robot understandability as it allows for exploratory interaction.

CCS CONCEPTS

• **Human-centered computing → Empirical studies in interaction design**.

KEYWORDS

Mobile robots, understandable robots, robot intent, pedestrian path prediction, shared spaces, co-navigation, interface design

ACM Reference Format:

Xinyan Yu, Marius Hoggenmueller, and Martin Tomitsch. 2023. Your Way Or My Way: Improving Human-Robot Co-Navigation Through Robot Intent and Pedestrian Prediction Visualisations. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction (HRI '23), March 13–16, 2023, Stockholm, Sweden.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3568162.3576992

1 INTRODUCTION

The safe deployment of mobile service robots requires them to navigate autonomously in close human proximity, raising the need for socially acceptable robot navigation systems, particularly for shared urban spaces. To enable human-robot co-navigation [35], i.e. humans and robots safely navigating a shared space, robots have to predict human motion intention and plan their own trajectories correspondingly [32].

Apart from incorporating socially-aware navigation planning, these robots should also be designed in a way that humans understand how they think and act [18]. Failure to address this perspective might not only degrade user experience but also affect the operation efficiency of these robots [6]. Mobile service robots built in non-anthropomorphic forms (e.g. the majority of delivery robots) are not able to communicate their intention through natural means such as head orientation or gestures [44]. Thus, a growing body of human-robot interaction (HRI) research has investigated the use of additional visual cues, for example, through projections [9, 25], LED displays [31], and augmented reality (AR) interfaces [23, 37, 54, 70], to communicate the robot's intent. However, research has also found that a natural, efficient, and safe human-robot interaction requires not only people's understanding of what the robot does or is aiming to do next (i.e. its intent) but also why it acts the way it does [2, 26]. The HRI community has long recognised the need for understandable robots [15], which spurred a growing body of work designing interfaces providing information behind robot's decision-making process [24, 37, 56, 61, 73].

Human intention prediction is largely investigated from a technological standpoint [57]; however, there have been no investigations into how communicating this information to pedestrians affects co-navigation and their understanding of the robot's behaviour. Through an iterative process and consolidating feedback from two expert focus groups (n=8), we developed an AR design concept to visualise pedestrian's predicted path alongside the robot's navigation intent. We then prototyped and compared the concept in a virtual reality (VR) experiment (n=20) against two conditions, a baseline without visualisation and one with only the robot's navigation intent visualisation. The paper makes the following contributions to HRI: (1) Proposals for robot intent and pedestrian prediction path visualisations in a shared space. (2) Insights about the impact of path visualisations on trust, sense of agency, user experience, and robot understandability in unplanned human-robot co-navigation.

2 RELATED WORK

2.1 Mobile Robot Intention Communication

To overcome the shortcomings of robots' non-verbal communication abilities, a substantial body of work has investigated explicit communication approaches to convey a robot's state and motion intention to humans [3, 17, 25, 63, 71]. For example, projecting

HRI '23, March 13–16, 2023, Stockholm, Sweden

ACM ISBN 978-1-4503-9964-7/23/03…$15.00
https://doi.org/10.1145/3568162.3576992

the intended navigation trajectory and occupied space of a forklift robot was found to improve human response in a co-located environment [9]. Cleaver et al. [10] further investigated dynamic path visualisation to project a robot's motion intent at varying lengths depending on the complexity of the upcoming path, with results indicating that participants preferred longer path projections.

Apart from using projections onto the environment, other research also explored the use of on-robot communication modalities such as LED light patterns [31, 66] and on-screen displays [31, 62, 65]. In light of recent advancements in head mounted display (HMD) AR systems, the HRI community further recognised the potential of AR to enable improved communication and interactions between robots and people [23, 54, 64]. Walker et al. [70], for example, presented several HMD AR designs to convey an aerial robot's motion intent. Their study showed that the HMD AR designs helped users better understand the robot's intent, leading to significantly improved task efficiency compared to only relying on physically embodied orientation cues.

2.2 Understandable Robotics

Studies investigating user trust in intelligent systems have suggested that understandability is an essential factor determining trust, with the presence of explanations increasing perceived transparency and system understanding [28]. The HRI community has also acknowledged the need for understandable robots for a long time, for example, through the organisation of workshops on explainable and trustworthy robots [15, 49]. In response to the call for understandable robot-human interaction, Wortham et al. [73] developed a visualisation system of a robot execution planner (and later deployed it as a mobile AR application [56]) to provide real-time graphical visualisations of robot behaviour drives and their priorities. Selkowitz et al. [61] investigated designing for transparency and understandability in the context of human-robot collaborative teamwork. They compared participants' task performance for four interfaces containing different levels of system transparency information, including goals, reasoning, future projection, and uncertainty. Results indicate that combining the first three types increases subjective trust and facilitates people's comprehension of the robot's actions. Others explored enhancing user understanding through visualising robot sensor data [46] or navigation stack data (e.g., laser scans and environment maps) [37]. However, these studies exclusively focused on the technical evaluation of the visualisation system without involving user research, so the effect of such visualisations on people remains unclear.

According to the definition of robot understandability by Hellström and Bensch [26], humans understand a robot when they have sufficient knowledge about the robot's state of mind (SoM). Pedestrian movement intention prediction is a vital part of a mobile service robot's SoM as it determines its navigation planning. However, in current mobile robots, this information is hidden from pedestrians, hindering them from building a more accurate mental model of the robot's navigation and decision-making process.

2.3 Pedestrian Intention Prediction

Understanding and predicting human motion has become a critical ability for intelligent systems to coexist and interact with humans, especially in application domains such as autonomous vehicles and service robots [57]. State-of-the-art deep learning models consider different aspects, including gestures [53], group behaviour [39], and context [34], to improve motion prediction accuracy. While there is continuing technological advancement in pedestrian intention recognition and prediction, to the best of our knowledge, there is no research on how to make this information transparent to pedestrians in co-navigation scenarios with mobile robots. However, studies from related contexts and domains offer a foundation for our research. For example, in the context of driver vehicle interaction, Kim et al. [33] presented a head-up display (HUD) AR system using a shadow overlay to indicate where pedestrians are predicted to cross [33]. Their evaluation proved the effectiveness of visualising pedestrian motion intention to help drivers build an accurate time-distance judgement of the vehicle and other moving obstacles. In the context of highly automated driving, Colley et al. [11] displayed recognised pedestrian intention to passengers inside the vehicle and compared two visualisation modalities, tablet display and HUD AR. Their study results demonstrated the effectiveness of visualising pedestrian intention in improving users' trust and a preference for AR-based visualisations.

2.4 Summary

To sum up, there is an extensive body of work in HRI on mobile robot intention communication and a few explorations to make internal decision-making processes or robot data visible to users. However, a majority of work focused on indoor settings (e.g. shared floor space [9], corridors [63]) and addressed anticipated human-robot collaboration scenarios, in which people and robots operate jointly as a team [23, 54, 70]. As a consequence, previous work in this domain mainly evaluated the effects of robots' additional information cues on task performance [9, 10, 16, 25] or perceived communication efficiency [17, 31], while rarely paying attention to how they affect trust during occasional and unplanned encounters in public settings. Responding to the call of improving understanding of autonomous systems for casual bystanders [7], our work extends the information the robot communicates to pedestrians beyond intent and raw sensor data to the internal processes influencing its decision-making. Furthermore, we broaden the empirical investigation of such visualisations regarding their impact on people's perceived trust toward robots and sense of agency (SoA).

3 DESIGN PROCESS

Research through Design (RtD) [77] is a well-established research methodology that is known for its capability to facilitate design explorations and generate new knowledge throughout the process. In recent years, RtD has also been applied in HRI [40, 41]. We chose RtD to guide the design of our path prediction visualisation as it has been found to support reevaluating underlying assumptions and reframing design problems in exploratory research [78]. In this section, we outline the assumptions that we made based on related studies, describe the methods used as part of the design process, and discuss how the findings led to a reframing of our research questions and informed the proposed path prediction visualisation.

3.1 Methods

We used an iterative approach to developing design proposals and seeking expert feedback through focus groups. The initial design proposals were based on previous findings in related work and involved the first author developing and refining proposals based on feedback from the other two authors. Given the complexity and domain-specificity of our design investigation, we invited experts of various relevant backgrounds as our focus group participants in order to effectively identify key issues. We first conducted an expert focus group with four participants, including one robotic expert (E1), one HRI researcher (E2), and two interaction designers (E3, E4). Based on our findings, we reframed the goal of the path prediction visualisation to support participants in developing a mental model of the robot navigation. We revised the design proposals accordingly and conducted a second expert focus group with four participants, including two HCI researchers (E5, E6), one HCI/HRI researcher (E7), and one interaction designer (E8). A different set of participants was recruited for the second focus group in order to reduce potential effects of preconceptions or biases while also collecting feedback from a larger sample size. Based on the findings from both focus groups, we then identified the design parameters for the path visualisation and formulated the research questions for the subsequent evaluation study.

3.2 Initial Design Proposals

Previous findings suggest that the presence of mobile robots in shared spaces influences pedestrians' navigation behaviours in several ways. First, it makes some pedestrians deviate from their original trajectory to keep an unnecessarily large distance [36, 69]. Second, pedestrians might walk in a reckless way [52] or even perform risky behaviours towards robots [43], causing failures in robot navigation [58]. Studies investigating human-robot teamwork suggested that providing robot transparency information can improve people's performance in a collaborative task [24, 38].

Building on previous work that suggests AR as a strong candidate to enhance human-robot interactions [64] and the importance of human motion prediction in influencing socially acceptable robot navigation [70], we decided on an initial design concept of using HMD AR to visualise the predicted paths. The upcoming motion information of pedestrians predicted by the robot will be transmitted from the robot to the HMD that pedestrians wear so that they can see the future trajectories as predicted by the robot. As all data is sensed and calculated centralised by the robot unit, every pedestrian wearing an HMD receives path prediction information for all detected pedestrians, even those who don't wear a gear themselves. We hypothesised that if pedestrians could see their upcoming trajectories as predicted by the robot, they could follow this hint and cooperate, improving the pedestrian's experience as well as the robot's operating efficiency. Drawing on findings from in-vehicle path prediction visualisations [11], we also hypothesised that visualising the pedestrian's predicted path could positively influence pedestrians' trust towards robots.

To evaluate these preliminary assumptions, we probed the initial design idea in a focus group. We used images and video mock-ups as representations of the initial concept (see Fig. 1, left) to foster discussions with participants on different aspects, including their willingness to comply with the predicted path, subjective feelings and trust. We also discussed additional feedback they would like to receive from the robot, e.g. to encourage them to follow the predicted path.

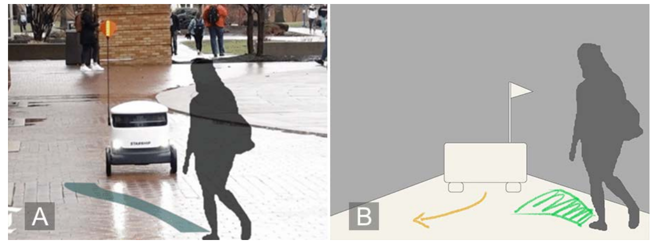

Figure 1: Representations of initial concept showing the pedestrian's predicted path (A) and refined concept featuring both the pedestrian's predicted path and robot intent (B).

Unexpectedly, we found that participants' interpretations of the predicted path visualisation and reactions to it vastly varied depending on their prior knowledge of robotics. The robotics expert (E1) instantly understood the predicted path visualisation and expressed a willingness to follow it. In contrast, the other three participants (E2-4) generally had inaccurate interpretations of the visualisation, assuming it was the *'recommended path'* or a *'command'* from the robot. This misinterpretation reduced participants' willingness to follow the path, refusing to follow *'what a robot tells me to do'*. In terms of trust, all participants agreed that the acquisition of more knowledge behind the robot's decision-making could improve their trust. However, the visualisation did not immediately increase trust for E2-4 who failed to understand the visualisation when first exposed to it. After being introduced to path prediction, they expressed their interest in testing the prediction, indicating their trust may be enhanced after assessing how it functions. Furthermore, apart from our original investigations, we also discovered a potentially negative impact visualising the predicted path could have on people's sense of agency as E2-4 expressed their feeling of *'being controlled'* by the robot.

3.3 Design Refinements

Based on the insights from the first focus group, we made adaptations to our initial design concept. Unlike in the case of visualising pedestrian intent recognition for passengers inside AVs [11], we found that additional information is needed for pedestrians in co-navigation scenarios to make use of the information. This might be due to the fact that the origin and addressee of the path prediction information are identical. Furthermore, the first focus group evaluation showed that a more accurate mental model of the robot's path planning is needed to make the predicted path visualisation understandable to non-expert users. To cater for this, we created several variations based on our initial design, such as adding the robot's navigation intent (see Fig. 1, right) and indicating the robot's field of view (as suggested in [42]) through outgoing radar waves and virtual eye designs. In the second focus group, we investigated the effectiveness of such additional information cues to contextualise the path prediction visualisation, thus facilitating people's understanding. We used images and video mock-ups of our adapted design variations to facilitate discussion of how additional information could improve the initial design proposal.

The predicted path visualisation – manifesting the 'reason' for the robot's behaviour – was more understandable for participants when accompanied by its 'result' – the robot's intended path. All participants agreed that the robot's intended path was the most crucial information to contextualise the predicted path visualisation as it also showed how their behaviour would influence the robot's path planning. Three participants further indicated that they would be more willing to follow the path when they could see the robot's navigation path. The findings confirmed the importance of indicating the robot's field of view. All four participants preferred eyes as a metaphor for the robot's perception, serving as an association that the predicted path is the calculation result of the data collected by the robot's sensor. Based on these findings, we further refined our design decisions: (1) visualising the predicted path using a green-coloured area with arrows; transparency gradually decreases from the pedestrian's position onward to indicate that prediction accuracy decreases with distance, (2) visualising the robot's intended path through yellow arrows, (3) visualising other pedestrians' predicted paths within robot detecting areas in cyan (to support understanding of the robot's path planning through interactions between the robot and other pedestrians), (4) overlaying the robot's front with virtual eyes to gaze at pedestrians when their predicted path starts to influence the robot's path planning.

The colours chosen for the visualisations follow the standards of road sign design in the Manual on Uniform Traffic Control Devices (MUTCD) published by the US Department of Transportation [47]. We used yellow for the robot's intended path to convey warning of potential interference and green for the predicted path visualisation to convey permitted traffic movements or directional guidance. The paths of other pedestrians are visualised using the same graphics but with cyan as a more neutral colour, allowing pedestrians to distinguish between their paths and those of others.

3.4 Research Questions

We revisited our initial assumptions and underlying research questions in light of several insights gained from the findings across both focus groups. First, the findings validate the potential of predicted path visualisation in enhancing pedestrian trust. Second, findings demonstrate an inaccurate understanding of the path prediction when displayed in isolation. Third, feedback from the robotic expert (E1) challenged that pedestrian compliance with the predicted path has little effect on the operational efficiency of the robot itself because of the high algorithm update frequency. We therefore turned to investigate how the visualisation of robot intent and pedestrian's predicted path can improve understanding of the robot's navigation. Furthermore, participant's assessment that the predicted path visualisation may negatively affect their agency in co-navigation scenarios motivated us to further examine the impact on sense of agency [19] in addition to the earlier identified factors of trust and user experience.

We thus formulated the following research questions to guide the evaluation of our path visualisation concept:

- To what extent, if any, does the visualisation of robot intent and pedestrian path prediction enhance people's understanding of robot navigation?
- How will the robot intent and pedestrian path prediction visualisation affect people on (1) trust, (2) sense of agency, and (3) user experience?

4 EVALUATION IN VIRTUAL REALITY

Given that our design concept includes both the robot's intent and the pedestrian's predicted path, a comparison to a conventional robot communication method only conveying intent would yield clear insights into how visualising the robot's decision-making process affects pedestrians. We conducted a within-subject study with three experimental conditions (see Fig. 2): a robot without external communication visualisation (*baseline*), a robot with AR visualisation of its intended navigation path (*intent*), and a robot with AR visualisation of both the robot's intended path and the pedestrian's upcoming path predicted by the robot (*intent + predict*).

We decided on a VR simulation that allows participants to experience co-navigation with a delivery robot in a shared space. VR simulations have been widely used for prototyping and evaluating interactions with autonomous vehicles (e.g. [11–13, 21]). They lower development costs and potential risk to participants while at the same time offering high ecological validity and interaction fidelity, for example, in comparison to video-based simulations [27].

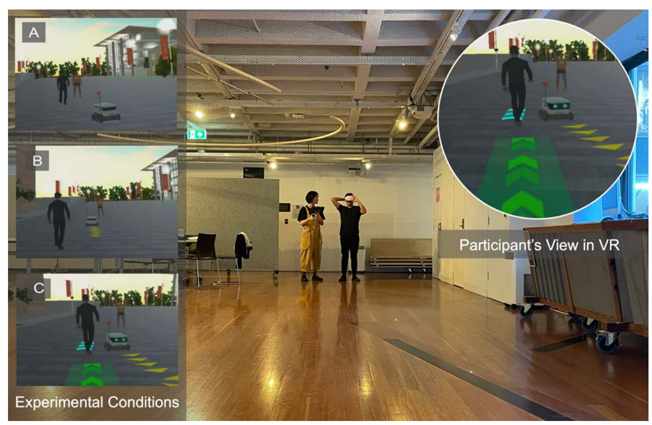

Figure 2: Study setup, the participant's perspective in VR, and experimental conditions: baseline (A), visualising the robot's intent (B), and visualising the robot's intent and the pedestrian's predicted path (C).

4.1 Study Apparatus and Implementation

We implemented the VR simulation using Unity3D[1] and deployed it to an HMD headset, the Oculus Quest 2 providing a fully untethered 6DOF experience[2]. The experiment was conducted in a 12x5-meter open floor space, where participants were able to physically walk through the simulated scenarios. We chose a 3D model of our university's campus avenue as the surrounding environment, since it represented a typical example of a shared space. Given that our participants were recruited from the university community, they were likely familiar with the environment, making it easier to relate to such a speculative scenario [21, 27]. For the mobile

[1] https://unity.com/
[2] https://www.oculus.com/quest-2/

robot, we used a 3D model of the delivery robot Starship, obtained from the modeling platform Sketch Fab[3]. The robot simulated the essential path planning functions. We predicted participants' upcoming trajectory using exponential smoothing model [20], which is a widely used statistical model in time series forecasting. We took eight points (equivalent to 4 seconds) to predict the estimated subsequent coordinate and repeated the process until the following eight points were calculated. The predicted points were used to render the real-time visualisation in the virtual environment. As the technical implementation was not the focus of our study, we decided on a simplified path prediction model. Furthermore, a state-of-the-art deep learning algorithm would have required extensive computing power not provided by standalone HMD hardware. It is worth mentioning, however, that the feedback provided by four participants during pilot testing validated that our approach offered a realistic sense of their upcoming path being predicted.

4.2 Participants

We recruited a total of 20 participants in the age range of 21 to 37 years (M=27.1/SD=4.43). Eight of them self-identified as male, twelve as female. Participants were recruited from our university's mailing lists, flyers and social networks. All participants voluntarily took part in the experiment and initial contact had to be made by them, following the study protocol approved by our university's human research ethics committee. Of our participants, seven had previously encountered robots in public spaces and five had seen cleaning robots in domestic spaces.

4.3 Procedure

After participants arrived at the study site, they were first given a brief introduction to the study background and procedure. They were then asked to sign a consent form, followed by a demographic questionnaire that obtained their basic information, including age group, gender, occupation, nationality, and previous experience with AR/VR and mobile robots. Following this, we briefly introduced the VR headset and its basic operations, and notified participants that the HMD VR was used to simulate the experience of seeing the visualisation with a wearable AR device. Prior to the conditions, each participant went through a familiarisation session to practise walking in the VR environment, which also ensured that participants did not experience motion sickness and were willing to proceed with the study.

Each participant experienced three experimental conditions in different orders (using balanced latin square design to minimise carryover effects). To ensure sufficient time for participants to get familiar with the visualisations, each participant experienced four different scenarios per experimental condition. In doing so, we also aimed to cater for different co-navigation situations that could occur in a shared space. The direction in which participants and the robot would move relative to each other (e.g., facing each other or passing in front of the robot) and the placement of other virtual pedestrians differed in each scenario. Virtual pedestrians were added to simulate a more realistic shared space experience, and to help participants to understand the visualisation based on how the robot responded to the predicted path of the virtual pedestrians.

To motivate participants to walk through the virtual scene (thus, encountering the robot while walking), we placed a virtual avatar that waved to the participant at the 'destination point'. Prior to the experiment, we briefed participants that this was a 'friend' that they were planning to meet with. After participants walked virtually (and physically) to the destination, they could enter the subsequent scenario by confirming a user interface overlay via the hand controller. Before continuing with the following scenario, the experimenter ensured that the participants would turn around by 180 degrees, thus making optimal use of the limited physical space and avoiding potential collisions. After each condition, participants removed the headset and were asked to fill out standardised questionnaires on trust, sense of agency, and user experience. After participants completed all three conditions, we conducted a post-study semi-structured interview to gain in-depth insights into their experiences. The whole study took approximately 45 minutes.

4.4 Data Collection

We collected both quantitative and qualitative data through questionnaires, observations and interviews, following a mix-methods approach [14]. To assess participants' trust towards the robot, we used the Trust in Automation Scale [30] with all items corresponding to 7-point Likert scales. This is a standardised questionnaire that is designed to measure people's trust in autonomous systems and has been widely used in HRI [74–76]. In considering the feedback obtained from the focus groups (i.e. that the predicted path visualisation might affect pedestrians' autonomy), we measured the participant's sense of agency using the Sense of Agency Scale (SoAS) [67] on a 7-point Likert scale. It measures people's perceived control over their own mind, body, and the environment and has been previously used to assess perceived agency when interacting with social robots [29]. To measure participants' user experience, we used the short version of the widely adopted User Experience Questionnaire (UEQ-S) [60] on a 7-stage scale from -3 to +3.

We took observation notes of participants' reactions towards the robot by observing their physical behaviour and monitoring interactions in the virtual environment (streamed in real-time to the experimenter's iPad). We additionally screen-recorded the VR interactions for later verification of observation notes and to contextualise participants' statements from the interviews. Through semi-structured post-study interviews, we asked questions about their preference between the three experienced conditions, their perceived trust towards the robot, their interpretation of the AR visualisations, and their perceived sense of agency.

4.5 Data Analysis

Questionnaires: We first calculated Cronbach's alpha to assess the internal reliability of all scales used. Internal reliability was excellent for the *trust* subscale (α=0.945) and good for *distrust* (α=0.812). For the SoAS, internal reliability was good for *positive agency* (α=0.883), but poor for *negative agency* (α=0.627). Following advice on Cronbach's alpha in [68], we removed items with a correlation lower than 0.25 (item 3, item 6), which made the internal reliability for the remaining items acceptable (α=0.717). For the UEQ-S questionnaire, item reliability was good for *pragmatic quality* (α=0.889) and excellent for *hedonic quality* (α=0.916).

We conducted and report on descriptive and inferential analysis of questionnaire data. As data was non-parametric, we used Friedman test to determine statistically significant differences. We

[3]https://sketchfab.com/

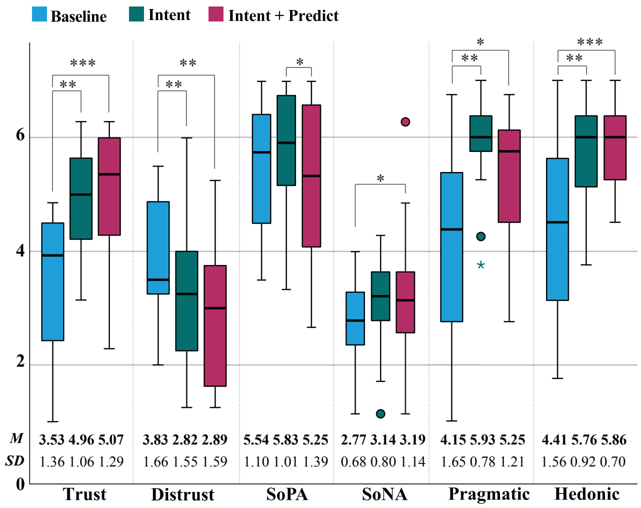

Figure 3: Results of Trust Scale, SoA Scale, and UEQ-S (scale adjusted to increase comparability). M: Mean, SD: Standard Deviations, *: p < .05, **: p < .01, ***: p < .001

further performed pairwise comparisons with Bonferroni corrections in case of significant differences. We considered an effect significant if the p-value was less than 0.05.

Interviews: All interviews were transcribed by the first author. The interview data was analysed following an inductive thematic analysis approach [8]. We first coded all quotes, sorted initial codes into sub-themes, and grouped them into final themes. Observation notes were used to supplemented the interview data analysis.

5 RESULTS

5.1 Trust

5.1.1 Trust Scale. Following descriptive data analysis (see Fig. 3), participants' trust was highest in *intent + predict*, while their distrust was lowest in *intent*. Friedman's ANOVA showed significant differences in both trust ($\chi^2(3)$=22.605, p<0.001) and distrust ($\chi^2(3)$=15.041, p<0.001) subscales. Post-hoc tests revealed that both *intent* ($p = 0.002$) and *intent + predict* ($p < 0.001$) received significantly higher trust ratings compared to *baseline*. Conversely, participants' distrust ratings in *baseline* was significantly higher than in *intent* ($p = 0.006$) and *intent + predict* ($p = 0.003$). There were no significant differences between *intent* and *intent + predict* in either trust and distrust subscales.

5.1.2 Qualitative Feedback on Trust. Even though statistical significance was not found between *intent + predict* and *intent*, during the interviews, over half of the participants (n=14) indicated that *intent + predict* improved their trust towards the robot. Below, we first discuss influencing factors for participants' improved trust discovered in the post-study interviews, followed by reasons that could explain the comparable quantitative results.

Firstly, *intent + predict* provided participants with a confirmation that they had been *'seen'* by the robot (n=7). In contrast, participants (n=5) indicated that their insecurity in *baseline* was mainly due to the uncertainty of whether the robot had detected them, e.g. *'I feel unsafe as you don't even know if the robot sees you or not'* (P6). Secondly, the dynamic response of the robot's intended path according to the predicted path made participants understand that the robot was actively calculating its paths to avoid collisions with pedestrians, which made some participants feel *'safer'* when navigating around the robot (n=9). When asked about the experience in *baseline*, some participants expressed scepticism about the robot's intelligence, suggesting that the robot *'looks silly'* (n=3), and even felt that the robot was just a *'toy'* (P16). In contrast, the robot in *intent + predict* was perceived by almost half of the participants (n=9) as *'more intelligent'* and 'smarter', which also contributed to their trust. Two participants even suspected that what they encountered in the three experimental conditions was *'not the same robot'*.

Although *intent + predict* provided additional motivating factors for trust and improved perceived trust, some participants (n=3) indicated a preference for *intent* rather than *intent + predict*: they pointed out a decrease in trust caused by potentially inaccurate prediction results of the pedestrian's upcoming trajectory - an issue that they did not find for only visualising the robot's navigation intent. Another participant, P17, also raised trust issues in relation to data security, which only surfaced for them through making the path prediction visible: *'What if it will send my data to somewhere I don't know?'*. Diverse perspectives and the variety of factors at play influencing participant's trust may account for the non-statistically significant difference between *intent* and *intent + predict*.

5.2 Sense of Agency

5.2.1 Sense of Agency Scale. Descriptive data analysis of the SoA scale (Fig. 3) showed that sense of positive agency (SoPA) was rated highest in *intent*, while unexpectedly sense of negative agency (SoNA) was rated lowest in *baseline*. This contradicting result might be due to the only moderate correlation between the two subscales [67]. *intent + predict* received the lowest ratings in SoPA and highest in SoNA, indicating a negative impact on participants' sense of agency. The Friedman's ANOVA showed significant differences in both the SoPA ($\chi^2(3)$=7.690, p=0.021) and SoNA ($\chi^2(3)$=8.247, p=0.016) subscales. Post-hoc tests revealed a significant difference between *intent* and *intent + predict* in SoPA subscale ($p = 0.027$), while no significant differences were found between *baseline* and the other two conditions. The SoNA rating in *intent + predict* was significantly higher than in *baseline* ($p = 0.043$). No significant difference was found between *intent* and the other two conditions.

5.2.2 Qualitative Feedback on SoA. Even though in the interviews only four participants explicitly expressed their feeling of *'being controlled'*, the quantitative data revealed the negative impact of predicted path visualisation on participants' SoA. Below we discuss interview data relating to SoA in order to gain further insight into which factors contributed to this assessment.

Most participants indicated a high degree of autonomy during the co-navigation task with the robot, e.g. determining their path *'based on own judgment'* (P11) instead of simply *'follow[ing] where the robot was telling to go'* (P15). Nonetheless, the predicted path visualisation still influenced participants' behaviour for two main reasons: Firstly, the concern that not following the predicted path would malfunction the robot made some participants feel their actions were *'restricted'* (P6, P11). As P6 stated: *'I think if I did something weird, for example, turning right or left out of sudden, I felt like the machine might not detect it.'* Secondly, some participants indicated that they followed the path *'unconsciously'* (n=6), which might lead to a decrease in their sense of agency.

5.3 User Experience

5.3.1 UEQ-S Questionnaire. Descriptive data analysis (see Fig. 3) of UEQ ratings shows that *intent* and *intent + predict* were the highest in pragmatic quality and hedonic quality, respectively. In contrast, *baseline* remained the lowest in both subscales. Friedman's ANOVA showed significant differences in both pragmatic ($\chi^2(3)$=15.53, p<0.001) and hedonic quality ($\chi^2(3)$=20.94, p<0.001). Post-hoc tests revealed that both *intent* ($p = 0.002$) and *intent + predict* ($p = 0.010$) received significantly higher ratings for pragmatic quality than the *baseline*. At the same time, the hedonic ratings were also significantly higher for *intent* ($p = 0.002$) and *intent + predict* ($p < 0.001$) in comparison to the *baseline* condition. Yet, no statistically significant difference was found between *intent* and *intent + predict* for both subscales.

5.3.2 Qualitative Feedback on User Experience. Despite the majority of participants (n=13) being in favour of additional information in helping them to understand the robot's navigation decision-making, the opinions on the necessity of the displayed information in an occasional encounter scenario varied. Some participants (n=8) suggested that the predicted path visualisation provided *'convenience'* (n=3) as it made the upcoming co-navigation with the robot *'predictable'* (P15) and *'so it's more clear that I don't have to give way to [the robot]'* (P10). Others stated that explanatory information such as the predicted path visualisation would add to their cognitive load and could be *'distracting'* (n=3) or even *'confusing'* (n=2). This was further underpinned by our observations, indicating that some participants (n=7) were more hesitant in *intent + predict*, i.e. walking slower than in the other two conditions or even stopping after a few steps (n=3).

Interview data showed that the predicted path visualisation raised participants' *'interest'* (n=16), describing their experience in the *intent + predict* condition as *'interesting'* (n=6), *'cool'* (n=3) and *'fun'* (n=2). This again is in line with our observations that revealed how participants increasingly engaged in exploratory interactions with the robot in the *predicted path* condition (n=11). Post-study, nine participants explained their exploratory behaviours with a feeling of *'curiosity'* invoked by the predicted path visualisation, which even made the interaction with the robot feel like *'play[ing] a game'* (n=2). AR, on the other hand, was regarded as an unfamiliar application for most participants (n=12), which also contributed to the perceived novelty of the experience.

5.4 Robot Understandability

When asked about their interpretation of the predicted path visualisation, most participants (n=17) stated that it improved their understanding of how the robot *'thinks'*. Three of them further indicated that they did not fully understand the robot when first seeing it but started to do so after seeing the effect that their actions had on the robot's behaviour: *'At first, I didn't realise that my path was linked to [the robot's] path. The moment I realised this was when I walked past it and I saw that its path had changed and veered off.'* (P18). Some participants indicated that their understanding of how the robot worked was improved after actively testing it (n=6): *'I tried to block it and to see its reaction [...], so I got it.'* (P9).

The predicted path visualisation also caused inaccurate interpretations of the robot for a few participants (n=3). For example, P7 incorrectly interpreted the path prediction visualisation as 'mind-reading': *'I was thinking to turn right (without actually performing the turning action), but [the robot] didn't notice that'*. This misinterpretation of the robot's prediction capabilities caused suspicion in P7, considering the robot as *'not reliable'*.

6 DISCUSSION

In this section we discuss the findings from across the design process and the evaluation study's quantitative and qualitative data sources and contrast them to prior work. Within each subsection we describe recommendations for designing robot explanatory visualisation and opportunities for future work.

6.1 Reflections on Visualisation Preferences

Expanding study results on the effect of robots' motion intent on improving collaborative task efficiency [10, 16, 70], our results show that communicating such information can also enhance people's trust towards robots in unplanned encounters (e.g., co-navigation scenarios in public spaces). Although no significant difference of trust was discovered between the two visualisations, qualitative data revealed positive aspects that explanatory information might have on people's trust, such as making the robot's capability of safely navigating explicit and enhancing people's perception of the robot as an intelligent agent. A robot's perceived intelligence was found to correlate with its animacy [4], while our study indicates the possibility of increasing non-humanoid robots' perceived intelligence by providing explanatory information. Yet, it is worth noting that the theme of perceived intelligence emerged from qualitative interview data. Thus, a systematic evaluation using scale measures for perceived intelligence [5] would be able to shed further light on these findings.

Despite the additional incentives for trust and the improvement of robot understandability, interview data showed comparable participant preferences between *intent* (n=11) and *intent+predict* (n=8), which are likely linked to the evenly matched user experience scores, and the higher SoA score in *intent*. While the information conveyed by *intent + predict* strengthened participant's understanding of the robot's behaviour, intent was effective enough for a bit less than one-third of participants (n=6) as co-navigation guidance during casual encounters, which indicates different informational needs across individuals. In addition, people's needs for explanations may also vary depending on the situational context, as suggested by two participants that the predicted path visualisation was especially helpful in a scenario where they could not decide how to give way to prevent the robot from bumping into a virtual pedestrian.

In comparison, it was *'unnecessary'* in scenarios where the co-navigation went smoothly. Building on these findings, future work could therefore investigate how to balance the benefits of conveying prediction information versus causing information overload, for example, through alternative information visualisation techniques. At the same time, in line with similar findings within the explainable AI research domain (e.g. [55, 72]), we also suggest that the design of explanatory visualisations for mobile robots should consider offering pedestrians flexibility to customise explanations according to different user needs and situations.

6.2 The Relation Between Explainability and Sense of Agency

During the focus groups, we identified people's willingness to dominate interactions with robots in co-navigation situations and the potential influence of exploratory information on people's SoA [45]. Previous work in HRI has investigated the impact of the robot's presence on people's SoA, suggesting a decrease in SoA in human-robot joint action [22]. Thus, researchers are looking for solutions to restore SoA for operators of autonomous intelligent systems [50, 59]. Our study results show that people's SoA could potentially increase when the robot communicates its intention, which is in line with Pagliari et al. [50], who argue that communicating an AI system's intention to its operators can improve their self-agency. Although pedestrians are not direct operators of mobile service robots, communicating the robot's motion intention still improves their confidence in co-navigating around robots, thus positively influencing their user experience during interactions.

In contrast, our study results also showed that visualising path prediction reduces pedestrians' SoA and increases hesitancy in co-navigation decisions. The interviews revealed that this negative impact on SoA is partly because participants may unconsciously take the predicted path visualisation as an instruction, causing the feeling of being 'controlled' by the robot. These findings combined have important implications for the design of transparent and explainable autonomous physical systems: while some information regarding a robot's behaviour and action may increase people's SoA, in some cases information can even have a reverse effect. Future research should therefore investigate what explanatory information to provide, through which message and communication strategy, and how to visualise, thereby considering the situational context of both robot and addressee.

6.3 Interactive Approach to Establishing Understanding

From the VR observation results, we found that the predicted path visualisation raises participants' interest in engaging in exploratory interaction with mobile robots. These interactions can support the understandability of the mobile robot's navigation behaviours: users can test and explore how the robot responds to their changing behaviour on the spot, which in turn helps them understand the capability and reliability of the robot's navigation. This finding echoes recent calls within the community to engage users in 'open-ended' explorations of an AI system's behaviour through interactive explanations [1, 64]. Currently, research investigating robot understandability in human-robot collaboration settings often provides interfaces with text-heavy descriptions (e.g. by highlighting labels in a static decision tree [73]). As such interfaces require a longer and more focused attention span, they are less suited for public spaces, in which people are likely to be in a rush and encounter robots as casual bystanders in an unplanned manner [48]. Our findings thus suggest that highly visual and dynamic explanatory cues such as path prediction visualisations can foster exploratory engagements that better fit casual encounters with robots.

Nevertheless, as our study showed, exploratory interactions stimulated by explanations could also increase the likelihood of passers-by interfering with service robots and thus negatively impacting their operational efficiency. In previous research, it has also been observed that due to curiosity, people disturb the robots' task performance [51] and even treat robots aggressively, causing damage [58]. Future research could explore solutions to counterbalance these issues. At the same time, as suggested by some of our participants, predicted path visualisation may be more suitable for educational purposes in order to train the general public about the capabilities and behaviour of robotics.

6.4 Limitation

First, the findings of our evaluation study drew on the experiences of a small number of mostly university students and young professionals. Although we anticipate comparable outcomes, a larger representative sample would be beneficial, particularly in resolving some borderline quantitative results. Second, the ecological validity of this work is limited by the use of a VR simulation. The virtual environment could not entirely recreate the complex multi-sensory experiences in the real world. The novelty of the VR experience and the fact that the robot in VR does not threaten participants' safety may also increase the likelihood of engaging in exploratory and risky behaviours. Third, because exposure to such internal information of a robot was a novel experience for all participants, causing *'unreal'* feelings for some (n=4), results may differ after long-term deployment and as familiarity increases, which need to be further validated in future studies.

7 CONCLUSION

In this paper, we described the design process and evaluation study of an AR predicted path visualisation to promote robots' understandability for pedestrians in an unplanned encounter scenario. The evaluation study results showed that both the *intent* and *intent + predict* visualisations can significantly improve pedestrian's trust and user experience, with *intent + predict* additionally reinforcing people's understanding of how such systems work by engaging them in exploratory interactions, thus providing extra incentives for further trust improvement. We also discovered a positive influence of robot intention communication on people's sense of agency and that adding the pedestrian's predicted path visualisation had a negative effect, which is an issue raised during focus groups. Participants' divergent informational requirements highlighted a need for flexible, customised and context-dependent explanatory visualisation. Our study unfolds the discussion of 'opening the black box' by helping people develop a mental model of how mobile robots operate through dynamic path visualisations. Future studies could more systematically study what other information about a robot's navigation decision-making process could benefit its understandability.

ACKNOWLEDGMENTS

This study is funded by the Australian Research Council through the ARC Discovery Project DP220102019 Shared-Space Interactions Between People and Autonomous Vehicles. We thank Dr. Stewart Worrall and Kunming Li from the Australian Centre for Field Robotics for their support with the algorithm implementation.

REFERENCES

[1] Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y. Lim, and Mohan Kankanhalli. 2018. Trends and Trajectories for Explainable, Accountable and Intelligible Systems: An HCI Research Agenda. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–18. https://doi.org/10.1145/3173574.3174156

[2] Sule Anjomshoae, Amro Najjar, Davide Calvaresi, and Kary Främling. 2019. Explainable Agents and Robots : Results from a Systematic Literature Review. In *AAMAS '19: Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems : (Proceedings)*. International Foundation for Autonomous Agents and MultiAgent Systems, 1078–1088. http://www.ifaamas.org/Proceedings/aamas2019/pdfs/p1078.pdf

[3] Kim Baraka, Stephanie Rosenthal, and Manuela Veloso. 2016. Enhancing human understanding of a mobile robot's state and actions using expressive lights. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 652–657. https://doi.org/10.1109/ROMAN.2016.7745187

[4] Christoph Bartneck, Takayuki Kanda, Omar Mubin, and Abdullah Al Mahmud. 2009. Does the design of a robot influence its animacy and perceived intelligence? *International Journal of Social Robotics* 1, 2 (2009), 195–204. https://doi.org/10.1007/s12369-009-0013-7

[5] Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. 2009. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International journal of social robotics* 1, 1 (2009), 71–81. https://doi.org/10.1007/s12369-008-0001-3

[6] Suna Bensch, Aleksandar Jevtic, and Thomas Hellström. 2017. On interaction quality in human-robot interaction. In *ICAART 2017 Proceedings of the 9th International Conference on Agents and Artificial Intelligence, vol. 1*. SciTePress, 182–189. https://doi.org/10.5220/0006191601820189

[7] Susanne Boll, Marion Koelle, and Jessica Cauchard. 2019. Understanding the Socio-Technical Impact of Automated (Aerial) Vehicles on Casual Bystanders. In *1st International Workshop on Human-Drone Interaction*. Ecole Nationale de l'Aviation Civile [ENAC], Glasgow, United Kingdom. https://hal.archives-ouvertes.fr/hal-02128379

[8] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. https://doi.org/10.1191/1478088706qp063oa arXiv:https://www.tandfonline.com/doi/pdf/10.1191/1478088706qp063oa

[9] Ravi Teja Chadalavada, Henrik Andreasson, Robert Krug, and Achim J. Lilienthal. 2015. That's on my mind! robot to human intention communication through on-board projection on shared floor space. In *2015 European Conference on Mobile Robots (ECMR)*. 1–6. https://doi.org/10.1109/ECMR.2015.7403771

[10] Andre Cleaver, Darren Vincent Tang, Victoria Chen, Elaine Schaertl Short, and Jivko Sinapov. 2021. Dynamic Path Visualization for Human-Robot Collaboration. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction* (Boulder, CO, USA) *(HRI '21 Companion)*. Association for Computing Machinery, New York, NY, USA, 339–343. https://doi.org/10.1145/3434074.3447188

[11] Mark Colley, Christian Bräuner, Mirjam Lanzer, Marcel Walch, Martin Baumann, and Enrico Rukzio. 2020. Effect of Visualization of Pedestrian Intention Recognition on Trust and Cognitive Load. In *12th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Virtual Event, DC, USA) *(AutomotiveUI '20)*. Association for Computing Machinery, New York, NY, USA, 181–191. https://doi.org/10.1145/3409120.3410648

[12] Mark Colley, Benjamin Eder, Jan Ole Rixen, and Enrico Rukzio. 2021. Effects of Semantic Segmentation Visualization on Trust, Situation Awareness, and Cognitive Load in Highly Automated Vehicles. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 155, 11 pages. https://doi.org/10.1145/3411764.3445351

[13] Mark Colley, Max Rädler, Jonas Glimmann, and Enrico Rukzio. 2022. Effects of Scene Detection, Scene Prediction, and Maneuver Planning Visualizations on Trust, Situation Awareness, and Cognitive Load in Highly Automated Vehicles. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 6, 2, Article 49 (jul 2022), 21 pages. https://doi.org/10.1145/3534609

[14] John W. Creswell. 2014. *A concise introduction to mixed methods research*. Sage, Thousand Oaks, California, USA.

[15] Maartje M.A. de Graaf, Bertram F. Malle, Anca Dragan, and Tom Ziemke. 2018. Explainable Robotic Systems. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction* (Chicago, IL, USA) *(HRI '18)*. Association for Computing Machinery, New York, NY, USA, 387–388. https://doi.org/10.1145/3173386.3173568

[16] Anca D. Dragan, Shira Bauman, Jodi Forlizzi, and Siddhartha S. Srinivasa. 2015. Effects of Robot Motion on Human-Robot Collaboration. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction* (Portland, Oregon, USA) *(HRI '15)*. Association for Computing Machinery, New York, NY, USA, 51–58. https://doi.org/10.1145/2696454.2696473

[17] Rolando Fernandez, Nathan John, Sean Kirmani, Justin Hart, Jivko Sinapov, and Peter Stone. 2018. Passive Demonstrations of Light-Based Robot Signals for Improved Human Interpretability. In *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 234–239. https://doi.org/10.1109/ROMAN.2018.8525728

[18] Terrence Fong, Illah Nourbakhsh, and Kerstin Dautenhahn. 2002. A survey of socially interactive robots: Concepts, Design and Applications. (2002).

[19] Shaun Gallagher. 2012. Multiple aspects in the sense of agency1. *New Ideas in Psychology* 30, 1 (2012), 15–31. https://doi.org/10.1016/j.newideapsych.2010.03.003

[20] Everette S. Gardner Jr. 1985. Exponential smoothing: The state of the art. *Journal of Forecasting* 4, 1 (1985), 1–28. https://doi.org/10.1002/for.3980040103 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/for.3980040103

[21] Michael A. Gerber, Ronald Schroeter, and Julia Vehns. 2019. A Video-Based Automated Driving Simulator for Automotive UI Prototyping, UX and Behaviour Research. In *Proceedings of the 11th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Utrecht, Netherlands) *(AutomotiveUI '19)*. Association for Computing Machinery, New York, NY, USA, 14–23. https://doi.org/10.1145/3342197.3344533

[22] Ouriel Grynszpan, Aïsha Sahaï, Nasmeh Hamidi, Elisabeth Pacherie, Bruno Berberian, Lucas Roche, and Ludovic Saint-Bauzel. 2019. The sense of agency in human-human vs human-robot joint action. *Consciousness and Cognition* 75 (2019), 102820. https://doi.org/10.1016/j.concog.2019.102820

[23] Morris Gu, Akansel Cosgun, Wesley P Chan, Tom Drummond, and Elizabeth Croft. 2021. Seeing thru walls: Visualizing mobile robots in augmented reality. In *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*. IEEE, 406–411. https://doi.org/10.1109/RO-MAN50785.2021.9515322

[24] S. Guznov, J. Lyons, M. Pfahler, A. Heironimus, M. Woolley, J. Friedman, and A. Neimeier. 2020. Robot Transparency and Team Orientation Effects on Human–Robot Teaming. *International Journal of Human–Computer Interaction* 36, 7 (2020), 650–660. https://doi.org/10.1080/10447318.2019.1676519 arXiv:https://doi.org/10.1080/10447318.2019.1676519

[25] Zhao Han, Jenna Parrillo, Alexander Wilkinson, Holly A. Yanco, and Tom Williams. 2022. Projecting Robot Navigation Paths: Hardware and Software for Projected AR. (2022), 623–628. https://doi.org/10.48550/arXiv.2112.05172

[26] Thomas Hellström and Suna Bensch. 2018. Understandable robots - What, Why, and How. *Paladyn, Journal of Behavioral Robotics* 9, 1 (2018), 110–123. https://doi.org/doi:10.1515/pjbr-2018-0009

[27] Marius Hoggenmüller, Martin Tomitsch, Luke Hespanhol, Tram Thi Minh Tran, Stewart Worrall, and Eduardo Nebot. 2021. Context-Based Interface Prototyping: Understanding the Effect of Prototype Representation on User Feedback. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 370, 14 pages. https://doi.org/10.1145/3411764.3445159

[28] Daniel Holliday, Stephanie Wilson, and Simone Stumpf. 2016. User Trust in Intelligent Systems: A Journey Over Time. In *Proceedings of the 21st International Conference on Intelligent User Interfaces* (Sonoma, California, USA) *(IUI '16)*. Association for Computing Machinery, New York, NY, USA, 164–168. https://doi.org/10.1145/2856767.2856811

[29] Aike Horstmann and Nicole Krämer. 2022. The Fundamental Attribution Error in Human-Robot Interaction: An Experimental Investigation on Attributing Responsibility to a Social Robot for Its Pre-Programmed Behavior. *International Journal of Social Robotics* 14 (07 2022), 1–17. https://doi.org/10.1007/s12369-021-00856-9

[30] Jiun-Yin Jian, Ann M. Bisantz, and Colin G. Drury. 2000. Foundations for an Empirically Determined Scale of Trust in Automated Systems. *International Journal of Cognitive Ergonomics* 4, 1 (2000), 53–71. https://doi.org/10.1207/S15327566IJCE0401_04

[31] Shyam Sundar Kannan, Ahreum Lee, and Byung-Cheol Min. 2021. External human-machine interface on delivery robots: Expression of navigation intent of the robot. In *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*. IEEE, 1305–1312. https://doi.org/10.1109/RO-MAN50785.2021.9515408

[32] Harmish Khambhaita and Rachid Alami. 2020. Viewing Robot Navigation in Human Environment as a Cooperative Activity. In *Robotics Research*, Nancy M. Amato, Greg Hager, Shawna Thomas, and Miguel Torres-Torriti (Eds.). Springer International Publishing, Cham, 285–300. https://doi.org/10.1007/978-3-030-28619-4_25

[33] Hyungil Kim, Jessica D Isleib, and Joseph L Gabbard. 2016. Virtual shadow: making cross traffic dynamics visible through augmented reality head up display. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 60. SAGE Publications Sage CA: Los Angeles, CA, 2093–2097. https://doi.org/10.1177/1541931213601474

[34] Julian Francisco Pieter Kooij, Nicolas Schneider, Fabian Flohr, and Dariu M Gavrila. 2014. Context-based pedestrian path prediction. In *European Conference on Computer Vision*. Springer, 618–633. https://doi.org/10.1007/978-3-319-10599-4_40

[35] Markus Kuderer, Henrik Kretzschmar, Christoph Sprunk, and Wolfram Burgard. 2012. Feature-based prediction of trajectories for socially compliant navigation. In *Robotics: science and systems*. https://doi.org/10.7551/mitpress/9816.001.0001

[36] Rainer Kümmerle, Michael Ruhnke, Bastian Steder, Cyrill Stachniss, and Wolfram Burgard. 2015. Autonomous robot navigation in highly populated pedestrian zones. *Journal of Field Robotics* 32, 4 (2015), 565–589. https://doi.org/10.1002/rob.21534

[37] Linh Kästner and Jens Lambrecht. 2019. Augmented-Reality-Based Visualization of Navigation Data of Mobile Robots on the Microsoft Hololens - Possibilities and Limitations. In *2019 IEEE International Conference on Cybernetics and Intelligent Systems (CIS) and IEEE Conference on Robotics, Automation and Mechatronics (RAM)*. 344–349. https://doi.org/10.1109/CIS-RAM47153.2019.9095836

[38] Shan G. Lakhmani, Julia L. Wright, Michael R. Schwartz, and Daniel Barber. 2019. Exploring the Effect of Communication Patterns and Transparency on Performance in a Human-Robot Team. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 63, 1, 160–164. https://doi.org/10.1177/1071181319631054 arXiv:https://doi.org/10.1177/1071181319631054

[39] Kunming Li, Mao Shan, Karan Narula, Stewart Worrall, and Eduardo Nebot. 2020. Socially Aware Crowd Navigation with Multimodal Pedestrian Trajectory Prediction for Autonomous Vehicles. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*. 1–8. https://doi.org/10.1109/ITSC45102.2020.9294304

[40] Michal Luria, Marius Hoggenmüller, Wen-Ying Lee, Luke Hespanhol, Malte Jung, and Jodi Forlizzi. 2021. Research through Design Approaches in Human-Robot Interaction. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction* (Boulder, CO, USA) *(HRI '21 Companion)*. Association for Computing Machinery, New York, NY, USA, 685–687. https://doi.org/10.1145/3434074.3444868

[41] Michal Luria, John Zimmerman, and Jodi Forlizzi. 2019. Championing Research Through Design in HRI. (2019). https://doi.org/10.48550/ARXIV.1908.07572

[42] Greg Lynn. 2021. *Robot Etiquette*. https://medium.com/@formgreglynn/robot-etiquette-6509abc92e32

[43] Ruth Madigan, Sina Nordhoff, Charles Fox, Roja Ezzati Amini, Tyron Louw, Marc Wilbrink, Anna Schieben, and Natasha Merat. 2019. Understanding interactions between Automated Road Transport Systems and other road users: A video analysis. *Transportation Research Part F: Traffic Psychology and Behaviour* 66 (2019), 196–213. https://doi.org/10.1016/j.trf.2019.09.006

[44] Conor McGinn. 2020. Why do robots need a head? The role of social interfaces on service robots. *International Journal of Social Robotics* 12, 1 (2020), 281–295. https://doi.org/10.1007/s12369-019-00564-5

[45] James W. Moore. 2016. What Is the Sense of Agency and Why Does it Matter? *Frontiers in Psychology* 7 (2016). https://doi.org/10.3389/fpsyg.2016.01272

[46] Faizan Muhammad, Amel Hassan, Andre Cleaver, and Jivko Sinapov. 2020. Creating a Shared Reality with Robots. In *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction* (Daegu, Republic of Korea) *(HRI '19)*. IEEE Press, 614–615. https://doi.org/10.1109/HRI.2019.8673191

[47] United States Road Symbol Signs FHWA MUTCD. 2002. *Manual on Uniform Traffic Control Devices*. https://mutcd.fhwa.dot.gov/services/publications/fhwaop02084/index.htm

[48] Mohammad Naiseh, Deniz Cemiloglu, Dena Al Thani, Nan Jiang, and Raian Ali. 2021. Explainable recommendations and calibrated trust: two systematic user errors. *Computer* 54, 10 (2021), 28–37. https://doi.org/10.1109/MC.2021.3076131

[49] Daniel Omeiza, Sule Anjomshoae, Konrad Kollnig, Oana-Maria Camburu, Kary Främling, and Lars Kunze. 2021. Towards Explainable and Trustworthy Autonomous Physical Systems. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI EA '21)*. Association for Computing Machinery, New York, NY, USA, Article 88, 3 pages. https://doi.org/10.1145/3411763.3441338

[50] Marine Pagliari, Valérian Chambon, and Bruno Berberian. 2022. What is new with Artificial Intelligence? Human–agent interactions through the lens of social agency. *Frontiers in Psychology* 13 (Sept. 2022). https://doi.org/10.3389/fpsyg.2022.954444

[51] Kelian van Pernis. 2022. *Investigating pedestrian-robot interaction in a context manipulation experiment*. Master's thesis. https://studenttheses.uu.nl/handle/20.500.12932/41450

[52] Manon Predhumeau, Anne Spalanzani, and Julie Dugdale. 2021. Pedestrian Behavior in Shared Spaces with Autonomous Vehicles: An Integrated Framework and Review. *IEEE Transactions on Intelligent Vehicles* (2021), 1–1. https://doi.org/10.1109/TIV.2021.3116436

[53] Photchara Ratsamee, Yasushi Mae, Kazuto Kamiyama, Mitsuhiro Horade, Masaru Kojima, and Tatsuo Arai. 2015. Social interactive robot navigation based on human intention analysis from face orientation and human path prediction. *ROBOMECH Journal* 2 (2015), 1–18. https://doi.org/10.1186/s40648-015-0033-z

[54] Christopher Reardon, Kevin Lee, and Jonathan Fink. 2018. Come see this! augmented reality to enable human-robot cooperative search. In *2018 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE, 1–7. https://doi.org/10.1109/SSRR.2018.8468622

[55] Mireia Ribera and Agata Lapedriza. 2019. Can we do better explanations? A proposal of user-centered explainable AI.. In *IUI Workshops*, Vol. 2327. 38.

[56] Alexandros Rotsidis, Andreas Theodorou, Joanna J. Bryson, and Robert H. Wortham. 2019. Improving Robot Transparency: An Investigation With Mobile Augmented Reality. In *2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. 1–8. https://doi.org/10.1109/RO-MAN46459.2019.8956390

[57] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Dariu M Gavrila, and Kai O Arras. 2020. Human motion trajectory prediction: a survey. *The International Journal of Robotics Research* 39, 8 (2020), 895–935. https://doi.org/10.1177/0278364920917446 arXiv:https://doi.org/10.1177/0278364920917446

[58] P. Salvini, G. Ciaravella, W. Yu, G. Ferri, A. Manzi, B. Mazzolai, C. Laschi, S.R. Oh, and P. Dario. 2010. How safe are service robots in urban environments? Bullying a robot. In *19th International Symposium in Robot and Human Interactive Communication*. 1–7. https://doi.org/10.1109/ROMAN.2010.5654677

[59] Supraja Sankaran, Chao Zhang, Marisela Gutierrez Lopez, and Kaisa Väänänen. 2020. Respecting Human Autonomy through Human-Centered AI. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society* (Tallinn, Estonia) *(NordiCHI '20)*. Association for Computing Machinery, New York, NY, USA, Article 134, 3 pages. https://doi.org/10.1145/3419249.3420098

[60] Martin Schrepp, Andreas Hinderks, and Jörg Thomaschewski. 2017. Design and evaluation of a short version of the User Experience Questionnaire (UEQ-S). *International Journal of Interactive Multimedia and Artificial Intelligence* 4, 6 (2017), 103. https://doi.org/10.9781/ijimai.2017.09.001

[61] Anthony R. Selkowitz, Cintya A. Larios, Shan G. Lakhmani, and Jessie Y.C. Chen. 2017. Displaying Information to Support Transparency for Autonomous Platforms. In *Advances in Human Factors in Robots and Unmanned Systems*, Pamela Savage-Knepshield and Jessie Chen (Eds.). Springer International Publishing, Cham, 161–173. https://doi.org/10.1007/978-3-319-41959-6_14

[62] Moondeep Chandra Shrestha, Ayano Kobayashi, Tomoya Onishi, Erika Uno, Hayato Yanagawa, Yuta Yokoyama, Mitsuhiro Kamezaki, Alexander Schmitz, and Shigeki Sugano. 2016. Intent Communication in Navigation through the Use of Light and Screen Indicators. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction* (Christchurch, New Zealand) *(HRI '16)*. IEEE Press, 523–524. https://doi.org/10.1109/HRI.2016.7451837

[63] Moondeep C. Shrestha, Tomoya Onishi, Ayano Kobayashi, Mitsuhiro Kamezaki, and Shigeki Sugano. 2018. Communicating Directional Intent in Robot Navigation using Projection Indicators. In *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 746–751. https://doi.org/10.1109/ROMAN.2018.8525528

[64] Ryo Suzuki, Adnan Karim, Tian Xia, Hooman Hedayati, and Nicolai Marquardt. 2022. Augmented Reality and Robotics: A Survey and Taxonomy for AR-Enhanced Human-Robot Interaction and Robotic Interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 553, 33 pages. https://doi.org/10.1145/3491102.3517719

[65] Daniel Szafir, Bilge Mutlu, and Terry Fong. 2015. Communicating Directionality in Flying Robots. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction* (Portland, Oregon, USA) *(HRI '15)*. Association for Computing Machinery, New York, NY, USA, 19–26. https://doi.org/10.1145/2696454.2696475

[66] Yenet Deribe Tafesse, Maggie Wigness, and Jeff Twigg. 2018. *Analysis techniques for displaying robot intent with led patterns*. Technical Report. US Army Research Laboratory Adelphi.

[67] Adam Tapal, Ela Oren, Reuven Dar, and Baruch Eitam. 2017. The Sense of Agency Scale: A Measure of Consciously Perceived Control over One's Mind, Body, and the Immediate Environment. *Frontiers in Psychology* 8 (2017). https://doi.org/10.3389/fpsyg.2017.01552

[68] Mohsen Tavakol and Reg Dennick. 2011. Making sense of Cronbach's alpha. *International journal of medical education* 2 (2011), 53. https://doi.org/10.5116/ijme.4dfb.8dfd

[69] Christian Vassallo, Anne-Hélène Olivier, Philippe Souères, Armel Crétual, Olivier Stasse, and Julien Pettré. 2018. How do walkers behave when crossing the way of a mobile robot that replicates human interaction rules? *Gait & posture* 60 (2018), 188–193. https://doi.org/10.1016/j.gaitpost.2017.12.002

[70] Michael Walker, Hooman Hedayati, Jennifer Lee, and Daniel Szafir. 2018. Communicating Robot Motion Intent with Augmented Reality. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction* (Chicago, IL, USA) *(HRI '18)*. Association for Computing Machinery, New York, NY, USA, 316–324. https://doi.org/10.1145/3171221.3171253

[71] Atsushi Watanabe, Tetsushi Ikeda, Yoichi Morales, Kazuhiko Shinozawa, Takahiro Miyashita, and Norihiro Hagita. 2015. Communicating robotic navigational intentions. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 5763–5769. https://doi.org/10.1109/IROS.2015.7354195

[72] Gesa Wiegand, Malin Eiband, Maximilian Haubelt, and Heinrich Hussmann. 2020. "I'd like an Explanation for That!"Exploring Reactions to Unexpected Autonomous Driving. In *22nd International Conference on Human-Computer Interaction with Mobile Devices and Services* (Oldenburg, Germany) *(MobileHCI '20)*. Association for Computing Machinery, New York, NY, USA, Article 36, 11 pages. https://doi.org/10.1145/3379503.3403554

[73] Robert H. Wortham, Andreas Theodorou, and Joanna J. Bryson. 2017. Improving robot transparency: Real-time visualisation of robot AI substantially improves understanding in naive observers. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 1424–1431. https://doi.org/10.1109/ROMAN.2017.8172491

[74] Julia L. Wright, Jessie Y. C. Chen, and Shan G. Lakhmani. 2020. Agent Transparency and Reliability in Human–Robot Interaction: The Influence on User Confidence and Perceived Reliability. *IEEE Transactions on Human-Machine Systems* 50, 3 (2020), 254–263. https://doi.org/10.1109/THMS.2019.2925717

[75] Yaqi Xie, Indu P Bodala, Desmond C. Ong, David Hsu, and Harold Soh. 2019. Robot Capability and Intention in Trust-Based Decisions Across Tasks. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 39–47. https://doi.org/10.1109/HRI.2019.8673084

[76] Jin Xu, De’Aira G. Bryant, and Ayanna Howard. 2018. Would You Trust a Robot Therapist? Validating the Equivalency of Trust in Human-Robot Healthcare Scenarios. In *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 442–447. https://doi.org/10.1109/ROMAN.2018.8525782

[77] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through Design as a Method for Interaction Design Research in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI ’07)*. Association for Computing Machinery, New York, NY, USA, 493–502. https://doi.org/10.1145/1240624.1240704

[78] John Zimmerman, Aaron Steinfeld, Anthony Tomasic, and Oscar J. Romero. 2022. Recentering Reframing as an RtD Contribution: The Case of Pivoting from Accessible Web Tables to a Conversational Internet. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI ’22)*. Association for Computing Machinery, New York, NY, USA, Article 541, 14 pages. https://doi.org/10.1145/3491102.3517789