Chapter 8

Session 8: Understanding Humans II: Social Interactions

Simulating Multiple Road User Perspectives on Autonomous Vehicle Behaviors

JiHyun Jeong
Cornell University
Ithaca, New York, USA
jj639@cornell.edu

David Goedicke
Informatics
University of Duisburg-Essen
Essen, NRW, Germany
david.goedicke@uni-due.de

Wendy Ju
Information Science
Cornell Tech
New York, New York, USA
wendyju@cornell.edu

Guy Hoffman
Cornell University
Ithaca, New York, USA
hoffman@cornell.edu



Figure 1: An all-way stop intersection in virtual reality with a red human-driven vehicle intending to go straight, a white autonomous vehicle with a human passenger intending to steer left, and a pedestrian intending to cross the street.

Abstract

This paper presents a virtual reality (VR) study that examines how multiple road users jointly interact with an autonomous vehicle (AV) in complex traffic scenarios. Moving beyond dyadic studies (e.g., AV-pedestrian or AV-passenger), our multi-user setup simulates ambiguous all-way stop intersections involving a pedestrian, a human driver in a conventional vehicle, and a passenger in an AV, all interacting simultaneously with the AV. We investigated how users perceive and respond to two distinct types of AV behaviors: an *efficient* AV that proceeds as soon as it is safe to do so, and a *prosocial* AV that yields to others before entering the intersection. Sixteen groups of three participants (N=48) took part in the study, with each group interacting with a single AV type across four ambiguous traffic scenarios. Our findings show that even simple AV behavior logics can meaningfully shape crossing negotiation dynamics and highlight how trust and perception can vary across different user roles. We conclude by discussing how our methods and insights can inform the research and design of AV interactions in complex multi-agent traffic environments.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**; **Collaborative interaction**; *Interaction design process and methods*; *Systems and tools for interaction design*; *Virtual reality*.

Keywords

autonomous vehicles, virtual reality, multi-agent, trust, interaction

ACM Reference Format:
JiHyun Jeong, David Goedicke, Wendy Ju, and Guy Hoffman. 2025. Simulating Multiple Road User Perspectives on Autonomous Vehicle Behaviors. In *17th International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI '25), September 21–25, 2025, Brisbane, QLD, Australia.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3744333.3747817

This work is licensed under a Creative Commons Attribution 4.0 International License.
AutomotiveUI '25, Brisbane, QLD, Australia
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2013-0/25/09
https://doi.org/10.1145/3744333.3747817

1 Introduction

Interactions on the road involve coordination among multiple people in different roles, such as pedestrians, drivers, and passengers.

This coordination often involves complex implicit negotiations, as each person has their own intentions and perceptions of the situation. Autonomous vehicles (AVs) that share the road with human users must recognize these different perspectives to interact effectively in traffic.

Although prior research has investigated AV interactions with road users, most have focused on dyadic setups with a single user type and the AV. These approaches may fail to capture the complexity of real-world traffic scenarios, which often involve multiple stakeholders with intersecting paths. To address this gap, we developed a multi-user virtual reality (VR) setup that enables three participants in different roles to interact in real time with a simulated AV. The participants in our setup took on the roles of an AV passenger, a pedestrian, and a human driver of a conventional vehicle. We focused on all-way stop intersection scenarios where all three road users arrived at the same time, creating uncertainty over who should proceed first.

We prototyped two distinct AV behavior strategies and compared how they affected multi-agent interactions in this ambiguous setting. Half of the participant groups interacted with an efficient AV that proceeds whenever it is safely able to do so, while the other half interacted with a prosocial AV that waits for both the pedestrian and the human-driven vehicle before proceeding. We then compared how different AV behaviors influenced joint decisions, trust, and perception across different road user roles.

This work contributes to AV interaction research in two key ways. First, it introduces a novel triadic interaction setup within a VR environment that supports a synchronized and repeatable study of complex traffic scenarios with multiple human road users. Second, it provides empirical evidence that road users in different roles can form diverging perceptions and levels of trust toward the same AV behavior, influencing how they negotiate and make joint decisions. Together, our work offers methodological and empirical insights for advancing the study of AV behaviors in complex traffic scenarios with multiple human stakeholders.

2 Related Work

Our work builds on a growing body of research that considers human factors in the design of AV behaviors. Prior studies have shown that passengers tend to prefer and trust AVs that exhibit defensive driving characteristics [2, 22], such as maintaining a longer distance from the lead car, use of full stops, and gradual acceleration [10, 35]. However, overly conservative behaviors can impede an AV's ability to navigate traffic efficiently and can lead to unnecessary confusion or congestion [36]. Additionally, road users may exploit the AV's defensiveness to violate road conventions, assuming that the AV will always yield to them [5, 25]. Therefore, in certain contexts, it may be necessary for AVs to appear less defensive for joint interaction. For example, an AV may activate its turn signal to initiate a lane change, prompting surrounding drivers to slow down and yield to accommodate the merge [29, 30].

Much work has explored how AVs coordinate, negotiate, and communicate with human road users [27, 34]. Explicit and implicit interaction between road users are particularly important, as solely relying on formal traffic rules without accounting for social behaviors can result in misunderstandings and even accidents [11]. To explore human interactions with AV behaviors, researchers have used a variety of prototyping methods. Some studies simulated AV behavior using Wizard-of-Oz techniques, where a hidden human driver controlled the vehicle while the passengers or pedestrians believed it was fully autonomous [6, 18, 28, 35]. Others have used driving simulators for users to compare autonomous driving policies that interact with static and dynamic characters [9, 15]. In recent years, virtual and mixed reality simulators, in particular, have been used to study road interactions in immersive, yet controlled environments. These methods have been used to investigate interaction between drivers [14, 23], drivers and pedestrians [7, 8, 21], AVs and pedestrians [1, 19, 31, 33], AVs and cyclists [17], and passengers in AVs [13, 37].

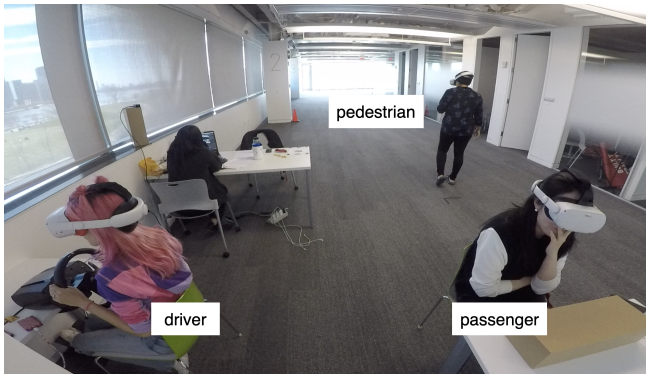


Figure 2: Three co-located participants are connected to the networked VR environment. The driver (bottom left) operates a virtual vehicle using a steering wheel and pedals. The passenger (bottom right) sits in the virtual AV, and the pedestrian (upper right) walks across the empty space to cross the road in the virtual world while a camera tracks them.

Despite a rich body of road interaction studies, relatively few studies have investigated scenarios involving simultaneous interactions with multiple human road users [3, 32]. The most closely related work features a coupled VR simulator in which a human driver and a pedestrian interact with an AV in the same virtual environment [4, 26]. To the best of our knowledge, this study is the first to simulate triadic AV interactions where three participants, each in different roles as a driver of a conventional vehicle, a pedestrian, and a passenger of the AV, jointly navigate ambiguous intersections. This approach offers both methodological and empirical contributions to understanding how trust, perceptions, and coordination can vary across user roles in response to different AV behaviors in complex, ambiguous scenarios.

3 System Setup

We create an immersive virtual environment in which three participants can interact simultaneously in the virtual world, assuming the roles of a *pedestrian*, a *passenger of an AV*, and a *driver of a conventional vehicle* (Figure 2). Our system builds on StrangeLand [14][1],

[1]StrangeLand has been open-sourced and we invite others with similar research interests to explore and contribute to the project at https://github.com/Strange-Land.

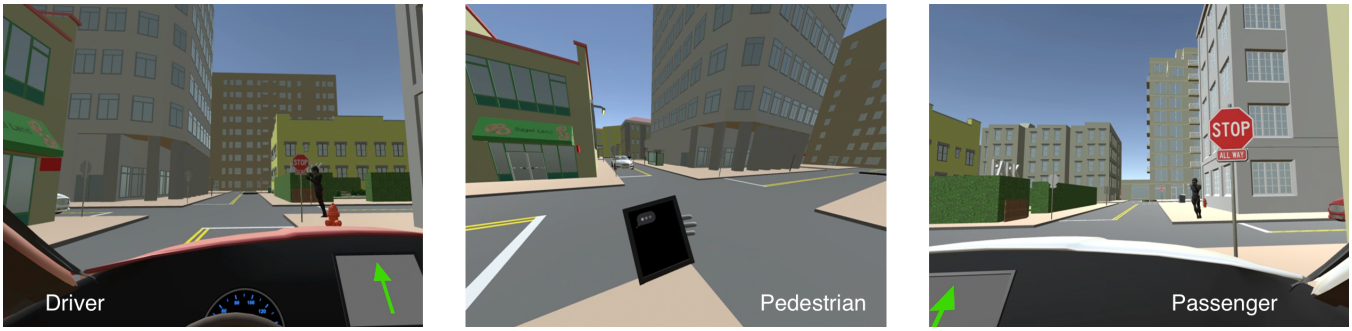


Figure 3: Each image depicts a viewpoint from each of the human's assigned roles. The driver's and passenger's head and hand movements, as well as the pedestrian's body movements are visible to one another.

an experimental VR platform for multi-participant driving simulation, by integrating modules for autonomous vehicle control and full-body pedestrian tracking (Figure 4).

Each participant wears a Meta Quest 2 headset and connects to a shared virtual environment through a networked Unity server. At the start of each scenario, participants are spawned at their designated positions and orientations in the virtual world. They interact in a shared virtual environment with the same roads, buildings, stop signs, moving vehicles, and with each other's avatars rendered in real time (Figure 3).

The driver's and passenger's head and hand movements are represented inside their respective vehicles. Each VR headset tracks their coordinates and sends them to the server to be shared with other clients. As a result, nonverbal cues such as gaze direction and hand gestures are visible to all participants.

The pedestrian's full-body movements are represented with a humanoid avatar. We used a Zed2i camera, Stereolab's AI model (v3.7), and their plugin for Unity to track the pedestrian's body pose. The pedestrian's avatar is rendered on the server using coordinates derived by aligning the neck position and orientation from the external camera to the headset's tracking data. Therefore, body movements such as walking and other emergent behaviors during traffic negotiations are visible to the driver and passenger.

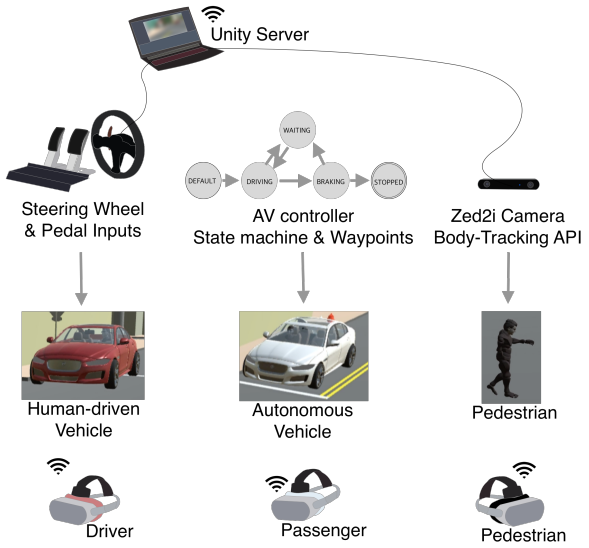


Figure 4: The system set up for three road user roles. The driver controls their virtual vehicle via steering wheel and pedal inputs. The AV is controlled via a state machine and waypoints. An external camera tracks the body movements of the pedestrian and projects it on the virtual server. The driver, AV passenger, and pedestrian all connect to the same Unity server through their headsets.

3.1 Participant Roles

Each participant engages with the virtual environment according to their assigned role: driver, AV passenger, or pedestrian.

Driver. The human-driven vehicle is red and visually and audibly mimics a conventional four door sedan. The driver controls the virtual vehicle by physically manipulating a Fanatec CSL steering wheel and pedal setup. The driver receives navigation guidance via an arrow displayed on the vehicle's dashboard. The driving experience is similar to that of a conventional vehicle (e.g. the car will slowly move forward if the participant takes their foot off the brake). To prevent cyber-sickness, we limited the vehicle's acceleration past the local speed limit. This also enabled the human-driven car and the AV to approach the intersection simultaneously, creating intentional ambiguity around right-of-way.

AV passenger. The autonomous vehicle visually and audibly resembles the overall look of the human-driven vehicle except its white exterior. The AV follows a pre-programmed path, with its intended direction displayed on the vehicle's dashboard. It drives below the local speed limit with gradual acceleration and deceleration. The vehicle's start is synchronized with the human-driven vehicle's approach to ensure that they arrive at the intersection at the same time. The passenger sits in the front-right seat and observes the AV's behavior.

Pedestrian. The pedestrian moves through the virtual world by physically walking in the real world. The pedestrian receives navigation instructions (e.g., to a bagel shop or park) via a message

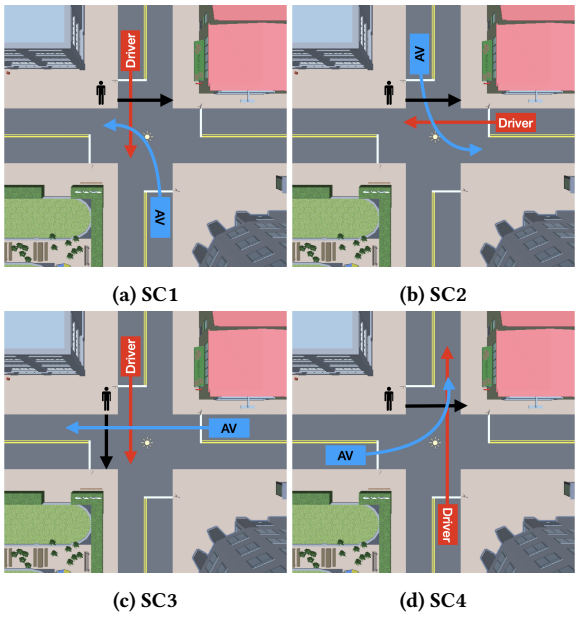
displayed on a virtual phone held in their left hand. Each destination requires the pedestrian to cross the street once. The timing of the message is synchronized to the AV and the human-driven vehicle's approach towards the stop sign.

3.2 Traffic Scenarios

To simulate complex interactions, we designed four scenarios at an all-way stop intersection, each involving different configurations of path conflicts among the three road users.

In all scenarios, the timing of each participant's approach was synchronized so they would arrive at the intersection simultaneously. According to state traffic regulations, the vehicle that arrives first should proceed first. If two vehicles arrive simultaneously, the vehicle to the right has the right-of-way. However, drivers are encouraged to use their judgment and communicate with others to navigate ambiguous situations.

To further the ambiguity, the path conflicts between the pedestrian and the vehicles are not immediately apparent. Although vehicles are expected to yield to pedestrians, the pedestrian's crossing directions are unknown until they receive a message on their virtual phone. In addition, a pedestrian's path can intersect with a vehicle's either immediately, crossing directly in front, or later, as the vehicle exits the intersection. Pedestrian zebra crossings are also intentionally omitted from the virtual environments. This uncertainty complicates negotiations between road users.



(a) SC1 (b) SC2

(c) SC3 (d) SC4

Figure 5: Four different traffic scenarios at an all-way stop intersection and the intended paths of the human-driven vehicle (red), the AV (blue) and the pedestrian (black)

We designed four different scenarios (SC1–SC4) to explore different configurations of path conflicts and create varied contexts for negotiation. Each scenario differed in which users had intersecting paths and the timing of those conflicts. The specific path configurations are shown in Figure 5, where directional arrows illustrate the intended paths of the pedestrian (black), the human-driven vehicle (red), and the AV (blue). For example, in SC1, the driver's path intersects first with the pedestrian before entering the intersection and then with the AV at the intersection, while the AV's path does not conflict with the pedestrian's. In contrast, in SC4, the vehicles' paths converge, and the pedestrian's path intersects with both vehicles as they exit the intersection.

3.3 Autonomous Vehicles

To enable controlled comparison of AV behaviors in multi-user scenarios, we implement a simplified driving logic using state machines, rather than using advanced motion planning algorithms. This approach allowed us to isolate and examine how different behavior strategies can influence negotiation, trust, and perceptions across user roles.

The AV follows a fixed set of waypoints for each scenario and transitions among five internal states: Default, Driving, Braking, Waiting, and Stopped. It begins in the Default state and transitions to Driving once the human-driven vehicle approaches a predefined distance from the intersection. When in motion, the AV calculates its steering and acceleration using linear interpolation between its current and target waypoint and speed. If the next waypoint is located at a stop sign, the AV transitions to the Braking state and decelerates with a target speed of zero. Once it comes to a full stop, it enters the Waiting state for three seconds before switching back to the Driving state. If, in the Driving state, another vehicle or a pedestrian is detected within its path, the AV re-enters the Waiting state for another three second interval to avoid collision. It continues cycling between Waiting and Driving until the intersection is clear, at which point it resumes Driving.

To detect a pedestrian or a human-driven car, we set two invisible boundaries for trigger detection zones in each scene. The first zone covers the area where the pedestrian's path intersects with the AV's path, and the second zone covers where the driver's path intersects with the AV's path. If a pedestrian or driver enters their respective zones before the AV, the AV will yield by remaining outside the detection boundaries until they pass. If the AV reaches the zone first, it will proceed unless another user unexpectedly enters in its path, in which case it will stop to avoid a collision.

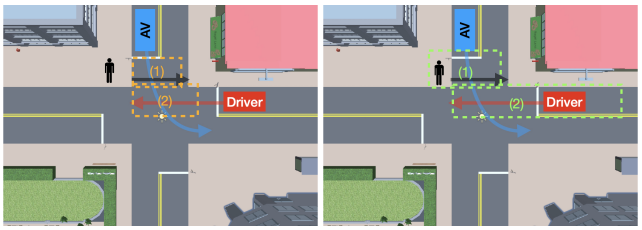


Figure 6: Zones in SC2 for detecting pedestrians (1) and human-driven vehicles (2) for efficient (left) and prosocial (right) AVs.

We designed and tested two AV behavior strategies: efficient and prosocial. Both share the same underlying driving logic. The difference lies entirely in the boundaries for the aforementioned detection zones (Figure 6). The efficient AV uses minimal boundaries

at path intersections to avoid collisions, so that it can proceed as soon as no other user is in its way. In contrast, the prosocial AV uses extended boundaries to detect pedestrians and drivers earlier as they approach the intersection. This results in the AV yielding and waiting until both have passed.

4 Study

We recruited participants using convenience sampling at a metropolitan campus in the United States. Sixteen groups each consisting of three people (N=48) completed our study. Eight groups interacted with efficient AVs and the other eight groups interacted with prosocial AVs.

Drivers, passengers, and most pedestrians had a U.S. driver's license or experience driving in the U.S. with an international license. The average age of the participants was 27.6 years (SD=6.2). Forty-eight percent identified as female, two percent identified as non-binary and the rest identified as male. The study took up to 45 minutes and all participants were compensated with a $15 gift card for their time. The study was approved by the university's Internal Review Board under protocol #IRB0146989.

4.1 Procedure

Participants gave informed consent and were assigned to one of three roles: driver of a conventional vehicle, a passenger of an AV, and a pedestrian. Before the study, participants received role-specific instructions, including how to operate the virtual vehicle or navigate the virtual space safely without walking into physical walls. In the simulation, participants were asked not to speak directly to one another, as road users typically cannot communicate verbally through vehicles. Instead, they were encouraged to behave as they would in normal traffic situations, knowing that their non-verbal actions such as hand gestures and body movements were visible to others through their avatars.

Participants first completed a practice session to become familiar with the system and their assigned roles. The practice scenario featured an all-way stop intersection with no conflicting paths among users. Each group then experienced four virtual traffic scenarios in sequence, interacting with one type of AV throughout. After each scenario, participants completed a situational trust questionnaire evaluating the AV's behavior. At the end of all four scenarios, each participant completed an individual post-study questionnaire.

4.2 Data and Measures

We collected log, video, and survey data for analysis. Log data included the head position and orientation of each participant, as well as the position and rotation of both vehicles in the virtual environment. It was used to analyze and visualize participant trajectories for each scenario (Figure 7). We also captured recordings of the virtual environment from each participant's perspective using ReRun [12] (Figure 3) and recorded video of the physical environment with participants' consent (Figure 2). Survey data included a post-scenario and post-study questionnaire. Following each scenario, participants completed a situational trust questionnaire [16] adapted to their specific roles (Table 1). At the end of the study, participants reflected on their impressions of the AV, including what they liked or disliked, how clear they found the right-of-way, and how they felt about the crossing decisions (Table 2).

Table 1: Situational Trust Scale Items for Each Role

Driver (Cronbach's α=0.88)
I trust the other vehicle in this situation. I performed better than the other vehicle in this situation.* In this situation, the other vehicle performed well enough for me. The situation was risky.* The other vehicle made an unsafe judgement in this situation.* The other vehicle reacted appropriately to the environment.
Passenger (α=0.89)
I trust the autonomous vehicle in this situation. If I were driving, I would have performed better than the autonomous vehicle in this situation.* In this situation, the autonomous vehicle performed well enough for me to engage in other activities (such as reading). The situation was risky.* The autonomous vehicle made an unsafe judgement in this situation.* The autonomous vehicle reacted appropriately to the environment.
Pedestrian (α=0.87)
I trust the white autonomous vehicle in this situation. If I were driving, I would have performed better than the white autonomous vehicle in this situation.* In this situation, the white autonomous vehicle performed well enough for me. The situation was risky.* The white autonomous vehicle made an unsafe judgement in this situation.* The white autonomous vehicle reacted appropriately to the environment.

*Asterisk indicates reverse-scored items

Table 2: Post-study Questionnaire

Driver
1. What was your overall impression of the other vehicle? Was there anything you liked or disliked about the way it was driving? 2. Was it clear who had the right of way? How did you decide when to enter the intersection?
Passenger
1. What was your overall impression of the autonomous vehicle? Was there anything you liked or disliked about the way it was driving? 2. Was it clear who had the right of way? How do you feel about the decisions that the autonomous vehicle made on when to enter the intersection?
Pedestrian
1. What was your overall impression of the white autonomous vehicle? Was there anything you liked or disliked about the way it was driving? 2. Was it clear who had the right of way? How did you decide when to cross?

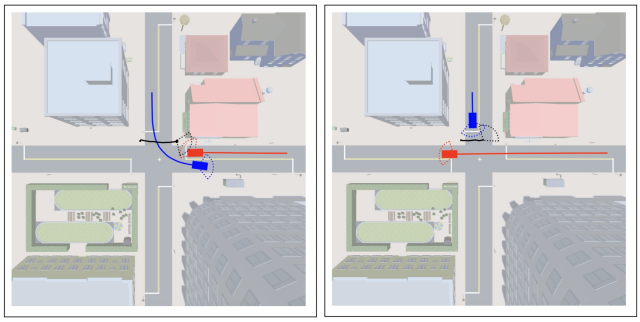


Figure 7: Trajectories of a human-driven vehicle (red), an AV (blue), and a pedestrian (black) in SC2. Solid lines indicate paths, dotted arcs represent horizontal fields of view. The efficient AV proceeds before the human-driven vehicle (left); the prosocial AV waits until others have passed (right).

5 Findings

We analyzed data from 16 groups of three people in four different scenarios, for a total of 64 runs. Two runs were excluded: one in SC1 due to missing data and the other in SC2 due to an outlier caused by the pedestrian deviating from the instructed path. The remaining data included 62 runs (31 with the efficient AV and 31 with a prosocial AV) and $3 \times 62 = 186$ participant data points.

5.1 Manipulation Check

5.1.1 Ambiguity in Perceived Right-of-way. Participant responses to our post-study questionnaire suggested that the study successfully created ambiguous traffic situations for most road users. A majority of the participants expressed doubt about right-of-ways and when they should proceed. Drivers and passengers mentioned being unsure who arrived first, or noted that they arrived at the intersection simultaneously. Pedestrians were likewise unsure about the vehicles' intentions and when it was safe for them to cross.

> *"Since me and the autonomous vehicle both arrived at around the same time, it was difficult to tell. It was also difficult to tell when the pedestrian was about to cross the road."* (E4-Driver[2])
>
> *"It was very ambiguous. I didn't notice who arrived at the stop signs first"* (P1-Passenger)
>
> *"Wasn't clear. I walked as soon as possible"* (E6-Pedestrian)

5.1.2 Difference in AV Behaviors. We derived AV Wait Time from our log files to illustrate the difference in AV behaviors. AV Wait Time measures the time that it spent waiting at the intersection, from its initial stop at the stop line until it fully committed to entering the intersection.

We ran an ANOVA test with Wait Time as the dependent variable and AV Type and Scenario as fixed factors. Results showed that there was a main effect for AV Type and an interaction effect between AV Type and Scenario (Table 3[3]). We further analyzed the interaction effect by performing simple main effects tests for AV Type with Scenario as the moderator (Table 4) and plotted the mean and standard error of wait times for each pair (Figure 8). We found that mean wait times were significantly shorter for efficient AVs compared to prosocial AVs across all scenarios except SC2. These results suggest that efficient AVs were able to reduce delays when navigating ambiguous multi-user intersections.

Table 3: ANOVA Results for AV Wait Time

Cases	SS	df	MS	F	p
AV Type	1325.811	1	1325.811	52.234	< .001
Scenario	114.778	3	38.259	1.507	0.223
AV Type * Scenario	282.518	3	94.173	3.710	0.017

Table 4: AV Type on Wait Time with Scenario as Moderator

Scenario	SS	df	MS	F	p
SC1	782.001	1	782.001	30.809	< .001
SC2	57.384	1	57.384	2.261	0.139
SC3	157.647	1	157.647	6.211	0.016
SC4	613.541	1	613.541	24.172	< .001

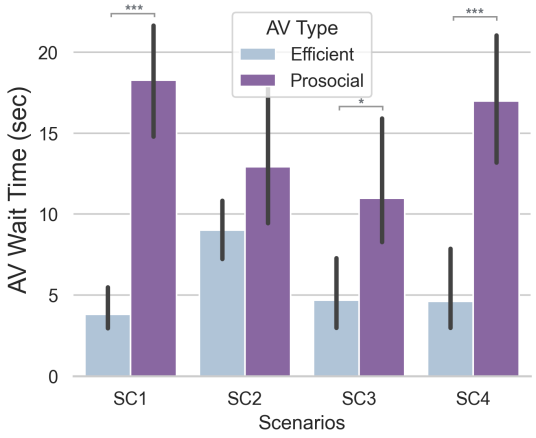


Figure 8: Barplot of average wait time for AV types in each scenario. Error bars represent 95% confidence intervals, and statistical significance is indicated by asterisks above the bars. (*: p<0.05, **: p<0.01 ***: p<.001)

5.2 Joint Decision in Negotiations

To better understand how differences in AV behavior influenced joint decision-making and negotiation, we analyzed log data capturing the AV's relative entry order at intersections. The intersection entry order refers to the AV's relative rank (first, second, last) of

[2]Quotes are attributed to participants as "<Study ID>-<Role>". Study ID indicates AV type and participant group within that condition.

[3]We report results in a table using sum-of-squares (SS), degrees of freedom (df), mean squares (MS), F ratios and p values.

when it entered the all-way stop intersection in relation to the human-driven car and the pedestrian.

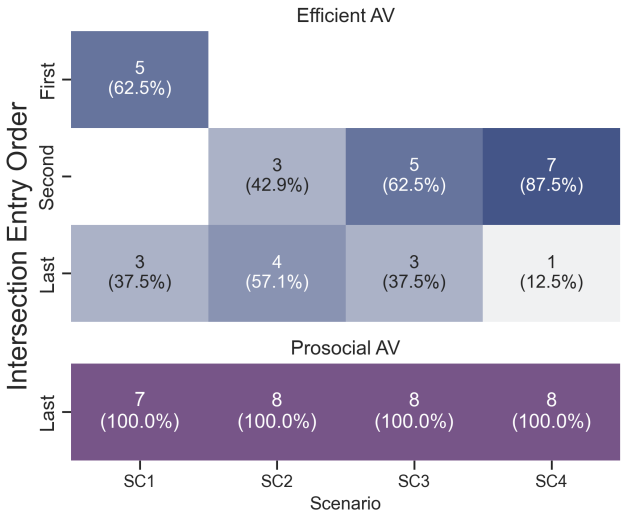


Figure 9: The relative order that efficient (top) and prosocial (bottom) AVs entered the intersection for each scenario, in relation to the human driver and pedestrian.

Figure 9 visualizes these entry orders across scenarios for both efficient and prosocial AV types. By design, the prosocial AV entered the intersection last, yielding to both the human-driven car and the pedestrian. In contrast, the efficient AV entered earlier in most cases, though its entry order varied across scenarios depending on the specific configuration of path conflicts.

We also supplement our analysis of AV entry order with qualitative observations from VR recordings, noting behaviors, movements, or other contextual factors that influenced crossing negotiations. Particularly in SC2, the AV's path intersected first with the pedestrian's and then with the driver's, while the driver and the pedestrian had parallel paths. Drivers interacting with both AV types typically proceeded to the intersection when they realized that their path did not conflict with the crossing pedestrian. However, in three cases involving the efficient AV, drivers were moving slowly when they noticed the AV subtly inch forward. In response, they yielded and allowed the AV to enter the intersection ahead of them.

Similarly in SC1, SC3, and SC4, the efficient AV slowly inched forward from the stop line after their initial stop, implicitly signaling its intent to proceed. Human drivers similarly made subtle advances forward but would brake repeatedly in response to the pedestrian's movements or the AV's behavior. In most of these cases, the efficient AV took advantage of the driver's hesitation and entered the interaction before them. In contrast, the prosocial AV always yielded to both the pedestrian and the driver. Drivers would proceed except to stop for the pedestrian, likely learning over time that the AV would always yield to them.

5.3 Situational Trust in the Autonomous Vehicle

We ran an ANOVA test with situational trust in the AV as the dependent variable and AV Type, Role, and Scenario as fixed factors. Results showed a main effect for AV Type and Role, as well as an interaction effect between AV Type and Role (Table 5). We performed simple main effects tests for AV Type with Role as the moderator (Table 6) to investigate the interaction between the two variables. We found that drivers and passengers trusted prosocial AVs more than efficient AVs ($p<.01$). Pedestrians, in contrast, showed no significant difference in trust levels between efficient and prosocial AVs (Figure 10).

Table 5: ANOVA Results for Situational Trust in AV

Cases	SS	df	MS	F	p
AV Type	11.108	1	11.108	7.142	0.008
Role	29.595	2	14.798	9.514	< .001
AV Type * Role	21.189	2	10.595	6.811	0.001
Scenario	4.774	3	1.591	1.023	0.384
AV Type * Scenario	0.247	3	0.082	0.053	0.984
Scenario * Role	5.644	6	0.941	0.605	0.726
AV Type * Scenario * Role	3.542	6	0.590	0.380	0.891

Table 6: AV Type on Situational Trust with Role as Moderator

Role	SS	df	MS	F	p
Driver	13.406	1	13.406	9.357	0.003
Passenger	15.540	1	15.540	10.847	0.002
Pedestrian	3.351	1	3.351	2.339	0.132

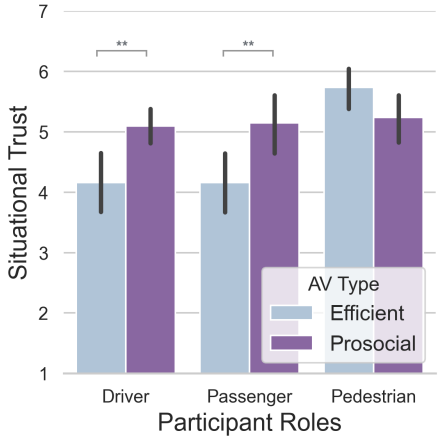


Figure 10: Barplot for average situational trust rating for AV types given by participant roles.

5.4 Qualitative Analysis of User Perceptions

In addition to our quantitative metrics, we collected descriptive impressions that participants had of the AV they interacted with. To synthesize the participants' open-ended text responses from the survey, we used affinity mapping to cluster the quotes from each role into positive and negative impressions, and extracted representative quotes for role-specific findings.

5.4.1 Drivers. Drivers believed that efficient AVs were responding to their actions but were too hasty at times.

> *"The other vehicle was great in 3/4 runs, patient, cognizant of the person crossing the road, seemed aware of my driving. In one of the runs it started moving a bit too soon when I was coming through the intersection."* (E5-Driver)

The predictability of the prosocial AV was appreciated by other drivers, but they also noted that the AV was conservative and yielding excessively.

> *"I like that the other car was driving in a stable and predictable manner."* (P8-Driver)

> *"I didn't like how it was too conservative. I felt in a couple of the trials they had the right of way but still let me go first."* (P2-Driver)

One common sentiment expressed by drivers, regardless of the type of AV they were interacting with, was wanting explicit communication on who should go first.

> *"It was very unnerving to not be able to gesture and decide who should go first."* (E6-Driver)

5.4.2 Passengers. Passengers perceived the efficient AV's behaviors as intelligent and responsive, but at times felt it was too risky or unsafe.

> *"I liked that it is very intelligent. I think it has its own observation of what's happening and is making responses accordingly. I disliked that sometimes it is taking the risks. Like when the other car stops at the stop sign at the same time, it moves first instead of waiting for the other."* (E3-Passenger)

> *"It did occasionally get quite close to the pedestrian or the other car, although there was no collision, but that made me feel a bit less safe."* (E5-Passenger)

Conversely, passengers in the prosocial AV noticed that the AV always yielded to everyone else. Passengers liked the cautious and defensive driving, but reported feeling impatient about it at times.

> *"It was cautious. It let others pass. I think it was appropriate for this situation. I myself might have eased in a bit earlier."* (P3-Passenger)

> *"I liked that it was cautious and seemed to drive only when it was certain no collision would occur. [...] In one scenario, I did think it was slightly too cautious, and that we could have continued through the intersection before the other car, and I got the impression that the other car was irritated with our delay (honking)."* (P7-Passenger)

> *"The autonomous vehicle always prefer driving the last. I might be anxious about it if I'm in a hurry"* (P6-Passenger)

Similar to drivers, passengers in both AVs wanted cues on whether it planned to yield or enter the intersection. Some participants even noted that they wanted to signal on behalf of the AV to coordinate with other road users.

> *"I did not like not knowing whether the AV was planning on going before or after the red car. By the end, I think we learned that the AV would wait for the red car no matter what, but the first few trials I did not like not knowing whether to signal the other car that they could go or not."* (P1-Passenger)

5.4.3 Pedestrians. The pedestrians thought the efficient AV behaved like human driven vehicles, but some noted that it was too close to them at times.

> *"Couldn't feel any difference between the autonomous car and the red car in terms of their behavior."* (E6-Pedestrian)

> *"The white one looks normal to me. I would not recognize it is a self driving car actually."* (E8-Pedestrian)

> *"At one time I did feel like the car was too close to me."* (E7-Pedestrian)

Conversely, pedestrians noticed that the prosocial AV was waiting until they finished crossing and appreciated it.

> *"Driving was pretty cautious and always waited for pedestrian to finish walking."* (P7-Pedestrian)

With both AVs, pedestrians were not always confident to cross the streets and wanted explicit acknowledgment from AVs that it was waiting for them to cross, especially when they were crossing directly in front of the AV.

> *"I don't know whether the car sensed me or not, and this made me hesitate to cross the streets. this is as opposed to the red car with actual human driver."* (P1-Pedestrian)

6 Discussion

Our work investigates multi-user interactions with autonomous vehicles (AVs) in a controlled yet immersive environment. By simulating ambiguous all-way stop intersections and capturing real-time negotiation through embodied avatars, we noted how different road users perceived and responded to different AV behaviors.

Even a simple logic for efficient and prosocial AVs produced meaningful differences in interaction dynamics and user perceptions. The efficient AV, designed to proceed as soon as it was safe, minimized wait times. This behavior was perceived by many drivers and passengers as competent and responsive but at times risky. As a result, they were less likely to trust an AV that appeared to take risks. In contrast, the prosocial AV, designed to yield to all other road users, earned higher trust and was perceived as stable and predictable. Even passengers who might be directly affected by delays rated the prosocial AV higher. This finding aligns with prior work on passenger preferences for defensive driving styles [2, 10, 22, 35].

Still, perceptions varied as some drivers and passengers found the prosocial AV overly cautious and impractical. Passengers, in

particular, expressed that they may grow impatient with excessive yielding if in a hurry, or felt uneasy when the AV's overly conservative behavior seemed to frustrate other road users. These reactions suggest that perceptions of the AV may vary depending on the context, especially for passengers. Future research could explore additional contextual factors such as situational urgency or traffic density to evaluate how different AV behaviors can influence perception and trust across various road user roles.

Interestingly, pedestrian outcomes diverged from those of the other user roles. Despite arguably being the most vulnerable road user, they showed no significant differences in trust between AV types. Many accepted the efficient AV's behavior as "normal" noting it was similar to how human drivers typically act, such as inching forward from the stop line while looking for opportunities to proceed. In fact, a few pedestrians explicitly compared the AV's movements to those of the human-driven car in the scenario. Although some pedestrians noticed that the prosocial AV always waited and did not move until they had fully crossed, this behavior did not appear to increase trust.

Pedestrian perceptions may be better understood in light of prior work showing that pedestrians often do not realize a vehicle is autonomous when it follows conventional driving norms [28]. Our findings suggest that AVs could maintain pedestrian trust while driving efficiently, so long as their behavior aligns with familiar social norms. This points to an interesting direction for future research in identifying the specific driving conventions that pedestrians expect, to guide when and how AVs can adopt efficient strategies without compromising social acceptability.

Finally, road users across all roles expressed a need for more explicit communication to resolve ambiguity during traffic negotiations. Drivers wanted a way to signal or gesture to the AV to go first. Passengers wanted the AV to communicate its intentions, both for their own situational awareness and to communicate the AV's intentions to others. Pedestrians sought clear recognition that the AV was waiting for them rather than simply stopping at the sign, especially in close-proximity scenarios like SC2. These findings align with broader conversations in HRI about the value of expressive and communicative interfaces on AVs [24] and highlight role-specific communication needs.

6.1 Implications for Multi-role Simulators in AV Research

Our work demonstrates the methodological benefits of studying multi-user interactions with an AV. Unlike dyadic AV studies, our multi-role simulation explores complex, ambiguous scenarios involving pedestrians, human drivers, and AV passengers. Such configurations more closely reflect real-world intersections where joint decisions unfold dynamically among multiple actors.

Furthermore, the multi-agent setup allows AV behavior to be understood in context, particularly in relation to the human-driven vehicle. Such comparison between the AV and the human-driven car can reveal trust dynamics that single-role studies might miss. For example, pedestrian perceptions of the efficient AV's behavior appeared to be influenced by simultaneously observing similar actions from a human-driven vehicle. Seeing these actions together may have helped normalize the efficient AV's behavior as typical and socially acceptable.

The setup also supports experimental control through manipulation of synchronized variables, such as timing and detection boundaries, while preserving the unscripted and emergent qualities of social negotiation in ambiguous traffic scenarios. Subtle or unexpected behaviors and naturalistic reactions that influence negotiations are often overlooked in purely quantitative experiments. While such behavioral variability is sometimes treated as noise, we argue that this degree of behavioral freedom offers a valuable source of qualitative insight into the nuances of road user interactions with AVs.

Overall, our methodological approach supports synchronized yet naturalistic interaction among pedestrians, drivers, and AV passengers, providing a rich context to examine how different AV strategies are perceived and responded to across roles. By identifying the role-specific perceptions that impact their negotiation with the AVs, we can inform the design of driving strategies that consider travel efficiency and social expectations in complex multi-agent interactions.

6.2 Limitations

While our VR-based multi-role simulation study provided valuable insights into AV interactions, several limitations affect the scalability of our methods and the generalizability of our findings, pointing towards opportunities for future investigations.

First, technical constraints influenced the study design. The system currently supports tracking for a single pedestrian within a limited walking range, which restricted the complexity of pedestrian trajectories. In addition, aligning coordinates between the external camera and the headset introduced minor latency and imprecision. While the setup was sufficient for our scenarios, studies involving multiple pedestrians or more dynamic movements would require expanded tracking infrastructure to support larger spaces, additional users, and precise movements reliably.

Moreover, the risk of cyber-sickness for drivers and passengers caused by simulated motion without corresponding physical feedback constrained the duration and the complexity of driving maneuvers in our scenarios. To mitigate this risk, we used a series of relatively short and synchronized scenarios to induce ambiguity in joint interaction without prolonging each session. Future setups could incorporate motion platforms or haptic feedback systems to support longer, more immersive experiences while maintaining comfort.

Second, the use of simplified driving logic, while intentional for experimental control, may limit real-world applicability. The AV behaviors in our study do not capture the full capabilities of modern motion planners used in real autonomous vehicles. To scale to longer, more complex, and dynamic traffic scenarios, future studies could integrate intelligent and adaptive driving models.

Finally, the virtual environment lacked cultural specificity and fidelity. We used generic urban buildings and streets without ambient traffic or bystanders. Additionally, virtual avatars did not convey facial expressions. Participant recruitment also relied on convenience sampling and may not generalize to broader populations, particularly those from different driving cultures. Future studies

could incorporate more representative or cross-cultural samples and use richer avatar representations to better capture culturally specific social interactions with greater fidelity.

Acknowledgments

This work includes materials gathered and presented in the first author's dissertation [20]. This research was supported by a Cornell Center for Social Sciences Grant, Defense Advanced Research Projects Agency (DARPA) Grant No. W911NF2010004, and National Science Foundation (NSF) Grant No. IIS-2212431. Opinions, findings, conclusions, or recommendations in this work do not reflect the views of these agencies.

References

[1] Raul Almeida, Emanuel Sousa, Dário Machado, Frederico Pereira, Susana Faria, and Elisabete Freitas. 2024. Analysis of the Interaction between Humans and Autonomous Vehicles Equipped with External Human–Machine Interfaces: The Effect of an Experimental Reward Mechanism on Pedestrian Crossing Behavior in a Virtual Environment. *Sustainability* 16, 8 (April 2024), 3236. doi:10.3390/su16083236

[2] Chandrayee Basu, Qian Yang, David Hungerman, Mukesh Singhal, and Anca D. Dragan. 2017. Do You Want Your Autonomous Car To Drive Like You?. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction* (Vienna, Austria) *(HRI '17)*. ACM, New York, NY, USA, 417–425. doi:10.1145/2909824.3020250

[3] Pavlo Bazilinskyy, Patrick Ebel, Francesco Walker, Debargha Dey, and Tram Thi Minh Tran. 2024. It Is Not Always Just One Road User: Workshop on Multi-Agent Automotive Research. In *Adjunct Proceedings of the 16th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Stanford, CA, USA) *(AutomotiveUI '24 Adjunct)*. ACM, New York, NY, USA, 268–272. doi:10.1145/3641308.3677400

[4] Pavlo Bazilinskyy, Lars Kooijman, Dimitra Dodou, Kirsten Mallant, Victor Roosens, Marloes Middelweerd, Lucas Overbeek, and Joost de Winter. 2022. Get Out of The Way! Examining eHMIs in Critical Driver-Pedestrian Encounters in a Coupled Simulator. In *Proceedings of the 14th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Seoul, Republic of Korea) *(AutomotiveUI '22)*. ACM, New York, NY, USA, 360–371. doi:10.1145/3543174.3546849

[5] Mark Colley, Elvedin Bajrovic, and Enrico Rukzio. 2022. Effects of Pedestrian Behavior, Time Pressure, and Repeated Exposure on Crossing Decisions in Front of Automated Vehicles Equipped with External Communication. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. ACM, New York, NY, USA, Article 367, 11 pages. doi:10.1145/3491102.3517571

[6] Rebecca Currano, So Yeon Park, Lawrence Domingo, Jesus Garcia-Mancilla, Pedro C. Santana-Mancilla, Victor M. Gonzalez, and Wendy Ju. 2018. !'Vamos! Observations of Pedestrian Interactions with Driverless Cars in Mexico. In *Proceedings of the 10th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Toronto, ON, Canada) *(AutomotiveUI '18)*. ACM, New York, NY, USA, 210–220. doi:10.1145/3239060.3241680

[7] Meiting Dang, Yan Jin, Peng Hang, Luca Crosato, Yuzhu Sun, and Chongfeng Wei. 2024. Coupling intention and actions of vehicle–pedestrian interaction: A virtual reality experiment study. *Accident Analysis & Prevention* 203 (Aug. 2024), 107639. doi:10.1016/j.aap.2024.107639

[8] Debargha Dey, David Goedicke, Chishang Yang, David Sirkin, Rebecca Currano, and Wendy Ju. 2024. GrokWalks: A Portable Virtual Reality Platform to Facilitate Studying Driver-Pedestrian Interactions. In *Adjunct Proceedings of the 16th International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI '24)*. ACM, 284–288. doi:10.1145/3641308.3685013

[9] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. CARLA: An Open Urban Driving Simulator. In *Proceedings of the 1st Annual Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 78)*, Sergey Levine, Vincent Vanhoucke, and Ken Goldberg (Eds.). PMLR, 1–16.

[10] Fredrick Ekman, Mikael Johansson, Lars-Ola Bligård, MariAnne Karlsson, and Helena Strömberg. 2019. Exploring automated vehicle driving styles as a source of trust information. *Transportation Research Part F: Traffic Psychology and Behaviour* 65 (Aug. 2019), 268–279. doi:10.1016/j.trf.2019.07.026

[11] Roni Factor, David Mahalel, and Gad Yair. 2007. The social accident: A theoretical model and a research agenda for studying the influence of social and cultural characteristics on motor vehicle accidents. *Accident Analysis & Prevention* 39, 5 (Sept. 2007), 914–921. doi:10.1016/j.aap.2006.12.015

[12] David Goedicke, Harald Haraldsson, Navit Klein, Lunshi Zhou, Avi Parush, and Wendy Ju. 2023. ReRun: Enabling Multi-Perspective Analysis of Driving Interaction in VR. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction* (Stockholm, Sweden) *(HRI '23)*. ACM, New York, NY, USA, 889–890. doi:10.1145/3568294.3580211

[13] David Goedicke, Jamy 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* (Montreal QC, Canada) *(CHI '18)*. ACM, New York, NY, USA, 1–11. doi:10.1145/3173574.3173739

[14] David Goedicke, Carmel Zolkov, Natalie Friedman, Talia Wise, Avi Parush, and Wendy Ju. 2022. Strangers in a Strange Land: New Experimental System for Understanding Driving Culture Using VR. *IEEE Transactions on Vehicular Technology* 71, 4 (2022), 3399–3413. doi:10.1109/TVT.2022.3152611

[15] Bassam Helou, Aditya Dusi, Anne Collin, Noushin Mehdipour, Zhiliang Chen, Cristhian Lizarazo, Calin Belta, Tichakorn Wongpiromsarn, Radboud Duintjer Tebbens, and Oscar Beijbom. 2021. The Reasonable Crowd: Towards evidence-based and interpretable models of driving behavior. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 6708–6715. doi:10.1109/IROS51168.2021.9635960

[16] Brittany E. Holthausen, Philipp Wintersberger, Bruce N. Walker, and Andreas Riener. 2020. Situational Trust Scale for Automated Driving (STS-AD): Development and Initial Validation. In *12th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Virtual Event, DC, USA) *(AutomotiveUI '20)*. ACM, New York, NY, USA, 40–47. doi:10.1145/3409120.3410637

[17] Ming Hou, Karthik Mahadevan, Sowmya Somanath, Ehud Sharlin, and Lora Oehlberg. 2020. Autonomous Vehicle-Cyclist Interaction: Peril and Promise. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. ACM, New York, NY, USA, 1–12. doi:10.1145/3313831.3376884

[18] Oliver Jarosch, Svenja Paradies, Daniel Feiner, and Klaus Bengler. 2019. Effects of non-driving related tasks in prolonged conditional automated driving – A Wizard of Oz on-road approach in real traffic environment. *Transportation Research Part F: Traffic Psychology and Behaviour* 65 (Aug. 2019), 292–305. doi:10.1016/j.trf.2019.07.023

[19] Suresh Kumaar Jayaraman, Chandler Creech, Dawn M. Tilbury, X. Jessie Yang, Anuj K. Pradhan, Katherine M. Tsui, and Lionel P. Robert. 2019. Pedestrian Trust in Automated Vehicles: Role of Traffic Signal and AV Driving Behavior. *Frontiers in Robotics and AI* 6 (Nov. 2019). doi:10.3389/frobt.2019.00117

[20] Ji Hyun Jeong. 2023. *Balancing Performance and Social Considerations for Autonomous Agents Interacting With Humans*. Cornell University.

[21] Amir Hossein Kalantari, Yue Yang, Jorge Garcia de Pedro, Yee Mun Lee, Anthony Horrobin, Albert Solernou, Christopher Holmes, Natasha Merat, and Gustav Markkula. 2023. Who goes first? A distributed simulator study of vehicle-pedestrian interaction. *Accident Analysis & Prevention* 186 (June 2023), 107050. doi:10.1016/j.aap.2023.107050

[22] Jesper Karlsson, Sanne van Waveren, Christian Pek, Ilaria Torre, Iolanda Leite, and Jana Tumova. 2021. Encoding Human Driving Styles in Motion Planning for Autonomous Vehicles. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. 1050–1056. doi:10.1109/ICRA48506.2021.9561777

[23] Navit Klein, Hauke Sandhaus, David Goedicke, Wendy Ju, and Avi Parush. 2024. Modeling Social Situation Awareness in Driving Interactions. In *Proceedings of the 16th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Stanford, CA, USA) *(AutomotiveUI '24)*. ACM, New York, NY, USA, 259–271. doi:10.1145/3640792.3675717

[24] Karthik Mahadevan, Sowmya Somanath, and Ehud Sharlin. 2018. Communicating Awareness and Intent in Autonomous Vehicle-Pedestrian Interaction. 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, 1–12. doi:10.1145/3173574.3174003

[25] Adam Millard-Ball. 2018. Pedestrians, Autonomous Vehicles, and Cities. *Journal of Planning Education and Research* 38, 1 (2018), 6–12. doi:10.1177/0739456X16675674

[26] Chun Sang Mok, Pavlo Bazilinskyy, and Joost de Winter. 2022. Stopping by looking: A driver-pedestrian interaction study in a coupled simulator using head-mounted displays with eye-tracking. *Applied Ergonomics* 105 (Nov. 2022), 103825. doi:10.1016/j.apergo.2022.103825

[27] Amir Rasouli and John K. Tsotsos. 2020. Autonomous Vehicles That Interact With Pedestrians: A Survey of Theory and Practice. *IEEE Transactions on Intelligent Transportation Systems* 21, 3 (2020), 900–918. doi:10.1109/TITS.2019.2901817

[28] Dirk Rothenbücher, Jamy Li, David Sirkin, Brian Mok, and Wendy Ju. 2016. Ghost driver: A field study investigating the interaction between pedestrians and driverless vehicles. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 795–802. doi:10.1109/ROMAN.2016.7745210

[29] Dorsa Sadigh, Nick Landolfi, Shankar S. Sastry, Sanjit A. Seshia, and Anca D. Dragan. 2018. Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state. *Autonomous Robots* 42, 7 (May 2018), 1405–1426. doi:10.1007/s10514-018-9746-1

[30] Dorsa Sadigh, Shankar Sastry, Sanjit A. Seshia, and Anca D. Dragan. 2016. Planning for Autonomous Cars that Leverage Effects on Human Actions. In *Proceedings of Robotics: Science and Systems*. Ann Arbor, Michigan. doi:10.15607/RSS.2016.XII.029

[31] Paul Schmitt, Nicholas Britten, JiHyun Jeong, Amelia Coffey, Kevin Clark, Shweta Sunil Kothawade, Elena Corina Grigore, Adam Khaw, Christopher Konopka, Linh Pham, Kim Ryan, Christopher Schmitt, and Emilio Frazzoli. 2022. Can Cars Gesture? A Case for Expressive Behavior Within Autonomous Vehicle and Pedestrian Interactions. *IEEE Robotics and Automation Letters* 7, 2 (2022), 1416–1423. doi:10.1109/LRA.2021.3138161

[32] Tram Thi Minh Tran and Callum Parker. 2024. Advancing VR simulators for autonomous vehicle–pedestrian interactions: A focus on multi-entity scenarios. *Transportation Research Part F: Traffic Psychology and Behaviour* 107 (Nov. 2024), 775–790. doi:10.1016/j.trf.2024.10.006

[33] Tram Thi Minh Tran, Callum Parker, Xinyan Yu, Debargha Dey, Marieke Martens, Pavlo Bazilinskyy, and Martin Tomitsch. 2024. Evaluating Autonomous Vehicle External Communication Using a Multi-Pedestrian VR Simulator. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8, 3, Article 130 (Sept. 2024), 26 pages. doi:10.1145/3678506

[34] Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and Kazuya Takeda. 2020. A Survey of Autonomous Driving: Common Practices and Emerging Technologies. *IEEE Access* 8 (2020), 58443–58469. doi:10.1109/ACCESS.2020.2983149

[35] Nidzamuddin Md. Yusof, Juffrizal Karjanto, Jacques Terken, Frank Delbressine, Muhammad Zahir Hassan, and Matthias Rauterberg. 2016. The Exploration of Autonomous Vehicle Driving Styles: Preferred Longitudinal, Lateral, and Vertical Accelerations. In *Proceedings of the 8th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Ann Arbor, MI, USA) *(Automotive'UI 16)*. ACM, New York, NY, USA, 245–252. doi:10.1145/3003715.3005455

[36] Wei Zhan, Changliu Liu, Ching-Yao Chan, and Masayoshi Tomizuka. 2016. A non-conservatively defensive strategy for urban autonomous driving. In *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*. 459–464. doi:10.1109/ITSC.2016.7795595

[37] Xin Zou, Steve O'Hern, Barrett Ens, Selby Coxon, Pascal Mater, Raymond Chow, Michael Neylan, and Hai L. Vu. 2021. On-road virtual reality autonomous vehicle (VRAV) simulator: An empirical study on user experience. *Transportation Research Part C: Emerging Technologies* 126 (May 2021), 103090. doi:10.1016/j.trc.2021.103090