

Keep it Real: Investigating Driver-Cyclist Interaction in Real-World Traffic

Ammar Al-Taie
a.al-taie.1@research.gla.ac.uk
Glasgow Interactive Systems Section,
School of Computing Science,
University of Glasgow
United Kingdom

Yasmeen Abdrabou
yasmeen.essam@unibw.de
University of the Bundeswehr
Munich, Germany and
University of Glasgow, UK

Shaun Macdonald
shaun.macdonald@glasgow.ac.uk
Glasgow Interactive Systems Section,
School of Computing Science,
University of Glasgow
United Kingdom

Frank Pollick
frank.pollick@glasgow.ac.uk
School of Psychology, University of
Glasgow
United Kingdom

Stephen Brewster
stephen.brewster@glasgow.ac.uk
Glasgow Interactive Systems Section,
School of Computing Science,
University of Glasgow
United Kingdom



Figure 1: Cyclist's gaze fixation during an interaction with a driver in real-world traffic.

ABSTRACT

Cyclists encounter drivers in many traffic scenarios; good communication is key to avoiding collisions. Little is known about everyday driver-cyclist interaction and communication. This is important in designing Automated Vehicles (AVs) that must drive safely around cyclists. We explored driver-cyclist interaction across diverse scenarios through in-the-wild observations ($N = 414$) and a naturalistic study involving cyclists wearing eye-trackers ($N = 12$). Results showed cyclists attended to road markings and traffic signs in controlled traffic scenarios but to vehicle sides and windows in uncontrolled encounters. Interactions were unlikely at controlled intersections, but various techniques were used to negotiate right-of-way in uncontrolled scenarios, e.g. cyclists used arm gestures and

shoulder checks to communicate their intent and awareness when lane merging. Drivers communicated these through on-vehicle signals and head movements at roundabouts. We discuss the implications of driver-cyclist interaction behaviour on AV interaction design and offer insights into system requirements to support cyclists riding in traffic.

CCS CONCEPTS

• **Human-centered computing** → Empirical studies in HCI.

KEYWORDS

Cyclists, Vulnerable Road Users, Autonomous Vehicle-Cyclist Interaction, Observations, Field Study, Naturalistic Study, Eye-Tracking

ACM Reference Format:

Ammar Al-Taie, Yasmeen Abdrabou, Shaun Macdonald, Frank Pollick, and Stephen Brewster. 2023. Keep it Real: Investigating Driver-Cyclist Interaction in Real-World Traffic. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3544548.3581049>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

CHI '23, April 23–28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9421-5/23/04...\$15.00
<https://doi.org/10.1145/3544548.3581049>

1 INTRODUCTION

In many countries, cyclists must share the road with motorised vehicles. This often exposes riders to dangerous situations. For example, 84% of fatal cycling accidents between 2015 and 2020 in the UK involved a motorised vehicle, with over 11,000 vehicle-cyclist collisions [1]. For this reason, cyclists rely on social *interactions* with other vehicles to reduce the chance of collision and resolve *space-sharing conflicts* that happen when cyclists and drivers “are intending to occupy the same region of space at the same time in the near future” [20]. Researchers define road *interactions* as “situations where the behaviour of at least two road users can be interpreted as being influenced by a space-sharing conflict between the road users” [20]. This behaviour can be expressed explicitly through social cues such as hand gestures and facial expressions, for example, a driver waving their hand to signal a cyclist to proceed at an intersection or a cyclist using an arm gesture to indicate a turn, or implicitly through driving or cycling behaviour, for example, a driver decelerating as a response to a cyclist’s right of way at an intersection, or a cyclist accelerating to support their intent to merge lanes with a driver [9, 17, 20].

As automated vehicles (AVs) become more common [4], cyclists (and other road users) will no longer have these social interactions with drivers to resolve potential conflicts. Cues such as eye contact and facial expressions will be lost, and resolving ambiguous situations, such as negotiating right-of-way, will be more difficult [16]. This could significantly impact rider safety. Previous work has primarily focused on the challenges of AV interaction with pedestrians [26]. However, cyclists have more complex needs: they move at higher velocities, can be in different locations around a vehicle, and their encounters with AVs are not limited to crossings [14]. Therefore, AV-cyclist interaction must be *versatile*; AVs must be able to interact with cyclists in diverse and complex traffic scenarios, including dynamic manoeuvres such as lane merging and stationary road infrastructure such as roundabouts.

There is a lack of research on the types and methods of communication that take place in diverse real traffic scenarios. AV-cyclist researchers commonly use design sessions, vehicle-cyclist collision reports and AV-pedestrian interaction research to construct solutions [15, 22]. AV-pedestrian researchers have used empirical evaluations of pedestrian behaviour in real crossing scenarios [26], allowing them to understand pedestrians’ unique needs in real settings. Uncovering the nuances of everyday driver-cyclist interaction behaviour in real-world traffic would help designers to build AVs that extend this natural behaviour and drive safely around cyclists in all circumstances. Specifically, identifying the scenarios that prompt interactions between cyclists and drivers will allow designers to evaluate their concepts in critical situations relevant to cyclists. Moreover, understanding how interactions differ between scenarios will help researchers develop versatile interfaces that recognise and respond appropriately to cyclists. Finally, specifying the factors, for example, on-vehicle direction indicators and traffic signs, that play a role in preventing space-sharing conflicts could help researchers establish a design space for AV-cyclist interfaces based on current cycling behaviour. The aim of our work is to investigate current driver-cyclist interactions to develop AV-cyclist

interfaces that extend current everyday interaction behaviour and do not require a large learning curve for cyclists.

In this paper, we report the findings of two studies done in real-world traffic. First, we observed driver-cyclist encounters ($N = 414$) to identify how interaction behaviour differs across five traffic scenarios. Second, we conducted a naturalistic cycling study ($N = 12$) by instrumenting commuter cyclists with eye-tracking glasses and a bike computer. The studies allowed us to gain a first-person perspective on cycling in real traffic and understand the nuances of everyday driver-cyclist encounters. The results of the studies show for the first time the rich interactions between drivers and cyclists in the real world and the types of communications that should be supported when AVs and cyclists share the road. We contribute the following:

- An understanding of how driver-cyclist interaction behaviour (exchanged messages, social cues and implicit cues) changes between common traffic scenarios that result in space-sharing conflicts;
- Common vehicle and traffic control features that appear in a cyclist’s field of view and cyclists’ gaze behaviours toward them;
- A discussion of the implications of our findings on AV interaction design.

2 RELATED WORK

There is little existing work on driver-cyclist interactions to inform the design of AVs that can drive safely around cyclists. In this section, we review the literature on the problems that can occur in vehicle-cyclist interactions, existing work on AV-cyclist interfaces and related work done with pedestrians and AVs.

2.1 The Problem

Despite researchers recognising the need for AVs to compensate for the disappearance of human drivers, AV-cyclist interaction remains largely unexplored. Dey et al. [8] reviewed existing AV interaction solutions and found that most target pedestrians, not cyclists. There is a clear need for a greater understanding of AV-cyclist interaction as highlighted by Pokorny et al. [25] and Pelikan [23], who observed automated shuttle-cyclist encounters in real traffic scenarios. The very cautious driving style of the shuttles and the hesitation of cyclists to pass due to the shuttles’ intentions being unclear caused many issues and hard stops from both road users. The encounters resulted in ambiguous and potentially dangerous situations that cyclists traditionally relied on interactions with the vehicle to resolve. For example, some cyclists steered away from the shuttles and were exposed to oncoming traffic, impacting their safety. Hagenzieker et al. [12] established that AVs must compensate for lost social cues to be successfully integrated into traffic. Participants judged photos of car-bicycle encounters from a cyclist’s perspective, with the car (AV or conventional vehicle) as an independent variable. Participants were more confident about being noticed by drivers than AVs, suggesting that AVs must communicate their awareness of cyclists effectively. The works showed a need for AV-cyclist interaction, but it is critical to follow up on these results to understand how such interactions should be facilitated.

Existing research has differentiated cyclists from other road users. Holländer et al. [14] developed a taxonomy of vulnerable road users (VRUs) that highlighted some unique cyclist characteristics. They rely, at least partially, on muscle power to move but travel at higher velocities than pedestrians. Despite their vulnerability, cyclists will likely encounter AVs in many different traffic scenarios, not just crossings. Trefzger et al. [29] equipped cyclists and pedestrians with eye-tracking glasses and instructed them to follow a predetermined route. Cyclists conducted shorter and less frequent shoulder checks than pedestrians and were more likely to focus on their path than the surrounding environment. Both works motivate further explorations of cyclists' unique road behaviours through the lens of AV-cyclist interaction design to understand the requirements of novel interfaces. Interfaces catered toward the needs of other road users, such as pedestrians, may not be suitable for cyclists.

Despite cyclists' unique characteristics, there is little foundational knowledge providing AV interaction designers with real-world requirements for AV-cyclist interfaces. Hegna Berge et al. [13] reviewed existing solutions for cyclists. They found no consensus on the placements and modalities that interfaces should use and the messages AVs should exchange with cyclists. This suggests a need for research that collects the requirements of AV-cyclist interfaces. Requirements were collected for AV-pedestrian interfaces; when Rasouli and Tsotsos [26] conducted a literature review of AV-pedestrian interaction research, they found that studies of driver-pedestrian interactions allowed for the creation of more effective solutions. The studies provided designers with the cues interfaces must recognise from pedestrians, messages they should respond with, and the optimal placements of interfaces. The studies even informed how AV-pedestrian interfaces should be evaluated. This strong foundation kickstarted AV-pedestrian research, and the AV-cyclist interaction domain would benefit from similar knowledge. Al-Taie et al. [3] provided a starting point by surveying cyclists about their interactions with drivers. Results showed that AVs must recognise eye contact and hand gestures from cyclists and communicate their awareness and intent to them. The authors identified cyclist-specific interaction scenarios not considered in previous work, such as cycle lanes merging into traffic.

These findings provide a general perspective on AV-cyclist interactions but not on real encounters in everyday traffic. While designers could build on real-world studies of cycling behaviour to construct their solutions, this could be challenging if these studies are not focused on informing AV-cyclist interface design. For example, Pokorny and Pitera [24] observed cyclists and human-driven trucks to inform safe infrastructure design for cyclists. While they found that some encounters resulted in social interactions, for example, cyclists waving to thank the drivers for yielding, the exchanged social cues were only recorded for the cyclists. Therefore, the results only highlight the cues that AVs should recognise, not those given in the response, which AVs would need to complete the communication.

2.2 User Requirements for AV-Cyclist Interfaces

Cyclists will encounter AVs in many traffic scenarios, but most existing solutions have only been evaluated in a single scenario, such as lane merging [15] or uncontrolled intersections [32]. Therefore,

it is unknown if these interface designs are *versatile* enough to cope with a wide range of traffic settings. Designers would benefit from knowledge of the scenarios that prompt interactions with cyclists and how interaction behaviour differs between them. Interaction behaviour includes *exchanged messages*: the content one road user communicates to the other. For example, thanking a road user or confirming their awareness of them; *social cues*: using social signals, such as hand gestures or facial expressions, as a mechanism to explicitly communicate a message to another road user; *implicit cues*: using driving or cycling behaviour as a mechanism to communicate a message. For example, braking to indicate intent to let the other road user proceed or accelerating to maintain the right of way. AV-pedestrian researchers are particularly interested in how implicit cues and social cues are used by drivers to negotiate right of way and may be transferred to AV interactions. Both Lee et al. [17] and Dey and Terken [9] observed driver-pedestrian encounters at crossings to understand how AVs should behave. They found that pedestrians and drivers rarely use social cues to interact and rely mostly on implicit cues. Whether cyclists rely on social or implicit cues when interacting with drivers is unknown. There is no work in the domain of vehicle-cyclist interaction where the detailed communication and social and implicit cues exchanged between riders and drivers have been recorded across a range of traffic scenarios. This is the key focus of our paper. Our first study takes an observation approach focusing on this to contribute new knowledge to the AV-cyclist design space.

Hou et al. [15] designed six prototype AV-cyclist interfaces, including on-road projections from the car and audio feedback from a helmet, which were evaluated in a simulated lane merging scenario. They found that cyclists preferred having an interface to facilitate their interactions with AVs. A design based on pedestrian needs (a display on the windscreen) received a low usability score as it prompted many shoulder checks from cyclists. While pedestrians are mostly exposed to the car's front in crossing scenarios [8], cyclists must interact with AVs behind, alongside and ahead of them. The results stress the need to inform AV-cyclist interface design with the optimal placements for these interfaces. Potential interface placements were investigated for pedestrians, where Dey et al. [11] conducted a controlled eye-tracking study in an outdoor setting to understand the kinds of information pedestrians seek from a car in a crossing scenario and optimal placements for AV-pedestrian interfaces. They collected gaze behaviour from participants wearing eye-tracking glasses when a car (driven by a collaborator) approached. Gaze patterns were distance-dependent, suggesting that AV-pedestrian interfaces should behave differently depending on the distance between the car and the pedestrian. In this paper, we conducted a naturalistic study with cyclists wearing eye trackers to provide AV-cyclist interaction designers with a similarly strong foundation to the one available to AV-pedestrian researchers.

2.3 Summary and Research Questions

Current work on AV-cyclist communication highlighted the need to facilitate interactions and stressed that the placement of these interfaces was essential to their usability. However, existing research has only considered individual scenarios, such as lane merging, as

use cases to evaluate these interfaces. Cyclists have unique needs compared to other road users and frequently encounter motorised vehicles in diverse traffic scenarios. Little is known about the requirements for AV-cyclist interfaces in real traffic scenarios. We constructed the following questions to address the gap in research:

- RQ1** How do the messages, social cues, implicit cues and on-vehicle signals that cyclists and drivers exchange differ between traffic scenarios?
- RQ2** How does cyclist gaze behaviour toward vehicle and traffic control features differ between traffic scenarios?

3 STUDY 1: IN-THE-WILD OBSERVATIONS

This study investigated how drivers and cyclists communicate during on-road manoeuvres and when negotiating in stationary infrastructure. Little is known about how behaviour differs between these encounters; capturing this knowledge is necessary to design versatile AVs that can safely drive around cyclists in all road scenarios. Previous research showed the suitability of conducting observations to record road user encounters in real-world traffic [9, 17, 24]. We conducted in-the-wild observations ($N = 414$) of driver-cyclist encounters in five different traffic scenarios.

3.1 Study Design

An observation approach was used to answer RQ1 and investigate the impact of traffic scenarios on driver-cyclist interaction behaviour. The method allowed us to collect data on many real encounters in Glasgow, UK, a city with dense urban areas and diverse road infrastructure. The city does not have much dedicated cycling infrastructure; cyclists often encounter drivers on their trips. An initial pool of traffic scenarios with varying traffic control levels was considered for observation. Work from Al-Taie et al. [3] was revisited to only include scenarios that cyclists perceive to be hazardous, allowing us to contribute findings that address cyclists' unique needs. We checked UK vehicle-cyclist collision reports [1] to observe scenarios critical to cyclists' safety. The following scenarios were selected: (1) controlled intersection, (2) roundabout, (3) uncontrolled intersection, (4) cycle lane that merges into traffic and (5) lane merging.

We then explored potential sites to observe. Prior work found that lower speed encounters resulted in more interactions and are easier to observe [28, 31], so we only considered sites with a speed limit of 30mph, the UK standard for urban areas [2]. We used Google Maps to determine the locations of possible sites within a reachable distance from the university and Google Street View to see the observer's field of view at each site. We used the Strava Heat Map¹ to ensure that the site had high cycling traffic. We visited each potential site to check the observer's field of view and record any traffic signs or road works that may impact interaction behaviour. Five sites were observed (see figure 2). An online form hosted on Qualtrics² was used along with an iPhone 12 mini to record the observations, which the lead author conducted. Our independent variable was the traffic scenario. Our dependent variables were (1) number of interactions/no interactions in a scenario, (2) messages exchanged by drivers and cyclists, (3) drivers' social cues, implicit

cues and on-vehicle signals, and (4) cyclists' social cues and implicit cues.

3.2 Procedure

Each site was observed for one day over two two-hour sessions during rush hour (session 1: 08:00-10:00, session 2: 16:00-18:00) on weekdays to maximise the number of observations. For scenarios with multiple cyclists, we only considered the first cyclist to avoid situations where a preceding cyclist could influence the observed cyclist's behaviour. Only the vehicle closest to the cyclist was observed to ensure it was the most likely to be in a space-sharing conflict. Upon arriving at each site, the observer first stood at the predetermined spot and filled in an online form to specify the site's conditions, including weather, road works and other factors that may impact interaction behaviour. They photographed their field of view and then filled in the observation form for each encounter.

Using the online form, the observer first specified whether an *interaction* occurred, allowing us to quantify the frequency of interaction in each scenario. If there was an interaction, the observer selected the exchanged messages, social cues, and implicit cues from predetermined lists informed by previous research [3, 8, 20]. For the messages, they selected *awareness* if a road user explicitly acknowledged the other's presence, *intent* if a road user indicated their next manoeuvre, e.g. a turn, *right of way negotiation*, if both road users negotiated who should proceed, e.g. at an intersection, and *positive* or *negative feedback* when a road user communicated their perception of the encounter's outcome.

Following this, the observer selected any social cues used by each road user to communicate a message. Choices included *hand/arm gestures*, *head movements* such as nodding, *facial expressions* and *vocal cues*. We did not record eye contact as this is difficult to observe. Instead, we recorded the head movements *toward a road user*. This could indicate eye contact being established [31]. The observer also selected any implicit cues used by each road user. Options included *deceleration*, *coming to a complete stop*, e.g. to indicate that the other should proceed, *acceleration* or *maintenance of speed*, e.g. to enforce right of way. The driver or cyclist could also be *stationary*, for example, if a driver is stationary behind a stop line at a roundabout. The observer then selected any on-vehicle signals used, including *hazard lights*, *flashing headlights*, *direction indicators* and *sirens*. The signals were identified from the UK Highway Code. This did not include implicitly triggered signals such as brake lights. The University's ethics committee approved this study.

3.3 Inter-Observer Reliability

Two authors independently observed the cycle lane that merges into traffic to eliminate bias. An inter-observer reliability analysis using the Kappa statistic was performed to determine consistency among observers. The reliability for the observers was found to be: Kappa = 0.857 ($P < .001$), suggesting a strong agreement between them.

3.4 Results

A Chi-Square test of independence was used to explore relationships between our independent and dependent variables. *Post hoc*

¹The Strava Global Heatmap: www.strava.com/heatmap

²Qualtrics online survey platform: www.qualtrics.com



Figure 2: Observed sites in the city of Glasgow.

analyses with the Chi-Square test of independence with Bonferroni corrections were conducted pairwise across our independent variable; 10 comparisons between each pair of traffic scenarios. We investigated the interplay of social and implicit cues by grouping the observations into three *cue categories* for each road user. We calculated the number of observations where the driver/cyclist used: (1) only social cues, (2) only implicit cues and (3) a combination of both cues.

3.4.1 Interaction Frequency. We observed 414 driver-cyclist encounters; 231 (55.8%) resulted in an interaction. Table 1 shows the number of observations and frequency of interaction in each scenario. We found a significant relationship between traffic scenario and frequency of interaction ($\chi^2(4, 414) = 71.96, P < .001$): the likelihood of an interaction is not the same for all traffic scenarios. *Post hoc* tests comparing the frequency of interaction between each pair of traffic scenarios revealed that the road users were less likely to interact at controlled intersections compared to uncontrolled intersections ($\chi^2(1, 244) = 42.664, P < .001$), lane merging ($\chi^2(1, 218) = 31.573, P < .001$), cycle lanes that merge into traffic ($\chi^2(1, 148) = 27.768, P < .001$) and roundabouts ($\chi^2(1, 176) = 36.177, P < .001$).

Exchanged Messages. Figure 3 shows the exchanged messages in the observed interactions. The relationship between traffic scenario and the type/content of exchanged messages was significant ($\chi^2(32, 351) = 176.23, P < .001$): road users exchange different messages in different traffic scenarios. We found significant effects when comparing exchanged messages in uncontrolled intersections with controlled intersections ($\chi^2(8, 183) = 42.56, P < .001$), roundabouts ($\chi^2(8, 200) = 33.78, P < .001$) and lane merging ($\chi^2(8, 205) = 127.25, P < .001$). Road users were more likely to negotiate their right of way and provide positive/negative feedback in uncontrolled intersections. Cyclists communicated their intent and awareness more often when lane merging and in controlled intersections than in uncontrolled intersections, while drivers were likely to communicate their intent and awareness at roundabouts.




Figure 3: Number and content of exchanged messages in driver-cyclist interactions per traffic scenario.

3.4.2 Driver Interaction Behaviour. Figure 4 illustrates drivers' social cues and implicit cues used in each traffic scenario. We found a significant relationship between cue categories and traffic scenarios ($\chi^2(8, 179) = 58.43, P < .001$). *Post hoc* tests revealed significant results when comparing cue categories in lane merging scenarios with uncontrolled intersections ($\chi^2(2, 107) = 40.05, P < .001$), cycle lanes that merge into traffic ($\chi^2(1, 51) = 23.04, P < .001$) and roundabouts ($\chi^2(1, 66) = 19.02, P < .001$): drivers did not use social cues when cyclists merged lanes with them, they often relied on social cues or a combination of social cues and implicit cues in uncontrolled intersections. Drivers only used combinations of both cues at roundabouts and cycle lanes that merge into traffic. *Post hoc* comparisons were also significant for cue categories in controlled intersections compared to uncontrolled intersections ($\chi^2(2, 94) = 19.41, P < .001$) and cycle lanes that merge into traffic ($\chi^2(1, 38) = 9.16, P = .0247$). Similar to our findings with lane merging, drivers were more likely to use implicit cues in controlled intersections than in other scenarios.

Traffic Scenario	N No Interaction	N Interactions	Total N Observations	Interaction Frequency (%)
Controlled Intersection	93	31	124	25
Roundabout	13	39	52	75
Uncontrolled Intersection	39	81	120	67.5
Cycle Lane Merging into Traffic	4	20	24	83.3
Lane Merging	34	60	94	63.8
Total	183	231	414	55.8

Table 1: Number of observed driver-cyclist encounters and interactions in each traffic scenario (N = Number of).




Figure 4: Drivers' social and implicit cues in each traffic scenario.

Social Cues. We did not find a significant association between social cues and traffic scenarios ($\chi^2(9, 89) = 7.68, P = .567$): drivers' social cues did not significantly differ when interacting across the scenarios.

Implicit Cues. We found a significant relationship between scenarios and implicit cues ($\chi^2(16, 184) = 113.83, P < .001$). *Post hoc* comparisons were significant when comparing controlled intersections with roundabouts ($\chi^2(4, 53) = 19.76, P = .0056$) and uncontrolled intersections ($\chi^2(4, 93) = 58.18, P < .001$). Drivers were more likely to maintain their speed in the cycle lane scenario. The tests also revealed significant differences between cycle lanes that merge into traffic and uncontrolled intersections ($\chi^2(4, 99) = 15.82, P = .0033$). Drivers were more likely to accelerate or maintain their speed in the cycle lane scenario. Comparisons of implicit cues in lane merging with roundabouts ($\chi^2(4, 66) = 23.49, P = .001$) and uncontrolled intersections ($\chi^2(4, 106) = 73.85, P < .001$) were also significant. Drivers rarely came to a complete stop when cyclists merged lanes. Drivers also used different implicit cues in roundabouts than uncontrolled intersections ($\chi^2(4, 108) = 22.29, P = .0018$). They were more likely to decelerate when encountering cyclists at roundabouts but accelerated or came to a complete stop at uncontrolled intersections.

On-Vehicle Signals. Drivers used direction indicators 7 times at controlled intersections, 14 times at roundabouts, 15 times at uncontrolled intersections and 11 times at cycle lanes that merge into traffic. They did not use direction indicators in lane merging, but they used hazard lights once and flashed their headlights once

when cyclists performed lane merging manoeuvres. Sirens were used once by an emergency vehicle at uncontrolled intersections.

There was a significant association between on-vehicle signals and traffic scenarios ($\chi^2(12, 50) = 52.13, P < .001$). *Post hoc* tests were significant between on-vehicle signals in lane merging scenarios with cycle lanes merging into traffic ($\chi^2(2, 13) = 13, P = .015$), roundabouts ($\chi^2(2, 16) = 16, P = .0034$) and uncontrolled intersections ($\chi^2(2, 18) = 18, P = .0044$). Drivers were unlikely to use direction indicators when cyclists merged lanes compared to other scenarios.

3.4.3 Cyclist Interaction Behaviour. Figure 5 illustrates cyclists' use of social cues and implicit cues. We found a significant association between traffic scenarios and cyclists' cue categories ($\chi^2(8, 228) = 36.04, P < .001$). *Post hoc* tests comparing cue categories in controlled intersections with cycle lanes that merge into traffic ($\chi^2(2, 48) = 17.524, P = .0016$), lane merging ($\chi^2(2, 85) = 8.39, P = .015$) and roundabouts ($\chi^2(2, 68) = 6.87, P = .032$) were significant. Cyclists were more likely to use only social cues in controlled intersections than in other scenarios where they used a combination of social cues and implicit cues. The pairwise comparisons were also significant for lane merging scenarios with roundabouts ($\chi^2(2, 99) = 13.19, P = .0137$), cycle lanes that merge into traffic ($\chi^2(2, 79) = 14.15, P = .0085$) and uncontrolled intersections ($\chi^2(2, 139) = 15.39, P = .0045$). Cyclists were more likely to use social cues in lane merging scenarios but used implicit cues to support the messages they communicated in the other scenarios.

Social Cues. The relationship between cyclists' social cues and traffic scenarios was significant ($\chi^2(16, 241) = 38.38, P = .0013$).




Figure 5: Cyclists' social and implicit cues used in each traffic scenario.

Post hoc tests showed statistically significant results when comparing riders' social cues in uncontrolled intersections with lane merging ($\chi^2(4, 151) = 26.89, P < .001$). Cyclists were likely to use a wide variety of social cues, including facial expressions, vocal cues and hand gestures in uncontrolled intersections, but mostly used hand gestures and head movements toward the driver when lane merging.

Implicit Cues. There was a significant relationship between traffic scenarios and cyclists' implicit cues ($\chi^2(12, 183) = 42.27, P < .001$). *Post hoc* comparisons were significant when comparing implicit cues in cycle lanes that merge into traffic with lane merging ($\chi^2(3, 57) = 19.19, P = .025$) and roundabouts ($\chi^2(3, 61) = 20.14, P = .016$). Riders often came to a complete stop in the cycle lane scenario but were more likely to accelerate or maintain their speed when lane merging and in roundabouts. We also found a significant difference between lane merging and uncontrolled intersections ($\chi^2(3, 106) = 18.14, P = .0041$). Cyclists were more likely to decelerate when lane merging than in uncontrolled intersections, where they more frequently accelerated or came to a complete stop.

3.5 Discussion

RQ1, which asked how interaction behaviour differs between traffic scenarios, was answered by observing 414 encounters in real-world traffic. The likelihood and nature of interactions were different between the scenarios, suggesting that AV interfaces need to be versatile and exhibit different behaviours depending on the type of traffic scenario.

Road users adopted various social and implicit cues at uncontrolled intersections, such as decelerating and using vocal cues, facial expressions and hand gestures to signal the other road user to proceed. They are a complex scenario for road users. Special attention should be given to supporting communication at them to enable safe interactions. Both Von Sawitzky et al. [32] and Matvienko et al. [22] proposed using Head-Up Displays (HUDs) to facilitate interactions at uncontrolled intersections, for example, to allow cyclists to "see" the AV through obstacles. However, these interfaces are most useful if the AV is not yet at the intersection, but we found that negotiations are complex when both the cyclist

and driver are near each other. AVs must recognise cyclists' social and implicit cues and respond by enforcing their right of way or yielding using techniques cyclists will understand and recognise.

Interaction behaviours were similar between uncontrolled intersections and cycle lanes that merge into traffic. Road users often had to negotiate their right of way in these scenarios. However, more drastic speed changes were noticeable in the cycle lane scenario, suggesting that road users were unsure how to proceed. Cycle lane scenarios had the highest proportion of interactions, aligning with previous work which showed that cyclists perceive these scenarios as high-risk [3, 27], which identified them as interaction triggers. These are cyclist-specific scenarios yet to be explored by AV interaction designers. While previous work has explored segregated cycle lanes as a solution to preventing space-sharing conflicts [7], this is not practical for all parts of a journey; designers should explore how vehicles can assist cyclists in transitioning from segregated to mixed traffic.

Roundabouts have similar physical properties to uncontrolled intersections: stationary infrastructure with give way lines that make road users decelerate or stop. The UK Highway Code recommends that drivers use direction indicators to communicate their intent at roundabouts. This resulted in different interaction behaviours from the ones at uncontrolled intersections. Drivers communicated their awareness through head movements toward the cyclist and intent through on-vehicle signals. Cyclists often communicated their intent through arm gestures, suggesting that mutual understanding and negotiation of the right of way was more straightforward than at uncontrolled intersections. These interactions also resulted in less drastic speed changes from cyclists, who were most likely to maintain their speed or accelerate at roundabouts. Our findings align with previous work highlighting the positive effect of communicating awareness and intent on the road [8, 21]. AVs can already communicate their intent through direction indicators at roundabouts, and designers should focus on compensating for the loss of eye contact and communication of awareness often utilised by drivers. Future work should explore whether communicating intent and awareness more explicitly in other scenarios, e.g. cycle lanes

that merge into traffic could help cyclists navigate such scenarios better.

Controlled intersections reduced the frequency and complexity of interactions. Road users were unlikely to interact at intersections with traffic lights determining the right of way, and the observed interactions were more straightforward. Cyclists used arm gestures and head movements to communicate intent and awareness. Drivers rarely responded through social cues and often maintained their speed to preserve their right of way, as determined by traffic lights. This finding shows potential and indicates a place for interfaces to facilitate social interactions in future traffic.

Lane merging is a popular scenario in the AV interaction domain. Al-Taie et al. [3] found that cyclists perceive them as high risk. In a simulated environment, Hou et al. [15] showed that riders preferred having an interface when lane merging around AVs. We found that cyclists used arm gestures and shoulder checks to communicate their intent and awareness to drivers. Drivers did not respond through social cues, suggesting that even though cyclists prefer having an interface in these scenarios, they may not be essential to resolving ambiguities here and could be optional devices, such as Augmented Reality (AR) headsets. AVs must detect a cyclist's social cues, such as arm gestures and shoulder checks, rather than just their overall presence, suggesting more sophisticated sensing may be needed. Designers could develop and assess the value of interfaces that explicitly respond to cyclists' messages when lane merging. However, they should avoid overwhelming the rider, as both the cyclist and AV are moving and navigating a complex traffic scenario. We observed some instances where drivers flashed their headlights or used hazard lights to respond to cyclists. Future work should explore whether traditional on-vehicle signals are enough for AVs to use in these scenarios.

We showed the range of messages exchanged by road users to provide designers with information about the cues that AVs should recognise from cyclists and highlight the appropriate responses. Having AVs communicate these messages could simplify their integration into mixed traffic with a minimal transition from current social norms. As to how these messages should be communicated, Pokorný et al. [25] and Pelikan's [23] observations suggest that AVs should mimic human drivers for implicit communication. For example, they should decelerate when they intend to yield or accelerate to preserve their right of way. We also identified the social cues drivers use. AVs may not need to mimic how these are currently presented; for example, external Human-Machine Interfaces (eHMIs) on the AV do not need to be in the form of a waving hand to mimic hand gestures. However, they should communicate appropriate messages clearly and predictably, which could be through light signals on the vehicle. Future work could compare the effects on cycling behaviour by mimicking human social cues or using novel communication techniques. Designers can use our findings to prioritise the scenarios that require AVs to explicitly communicate their messages and how implicit cues can be used alongside these explicit communications. We identified how interactions currently happen, but we also need to explore how these cues and messages may be translated into future interaction behaviour; we explored this in Study 2.




Figure 6: Defined Areas of Interest (AOI) illustrated on a 3D car model, traffic signs and road markings.

4 STUDY 2: NATURALISTIC CYCLING STUDY

Study 1 showed how interaction behaviour changed with infrastructure. However, there was a need to understand the detailed effects of external factors, e.g. the environment, that cyclists were exposed to and their reliance on them to resolve traffic ambiguities. This could contextualise the earlier findings on, for example, why interactions were unlikely in controlled intersections and inform the placements for AV-cyclist interfaces. Previous work showed the rich data that comes with collecting road user gaze behaviour [11, 27, 29]. Therefore, we followed up our observations with a naturalistic cycling study to gain a first-hand perspective of cycling behaviour in the real world.

4.1 Study Design

We recruited commuter cyclists and equipped them with eye-tracking glasses and a bike computer to record two home/office commutes in real-world traffic. We collected: (1) route GPS data, (2) first-person video footage, and (3) cyclists' gaze data. Participants took their usual commuting routes and were asked to behave as they normally would, preserving ecological validity. Routes were at least 5km and included at least three types of road infrastructure to increase the likelihood of recording diverse driver-cyclist encounters.

We defined Areas of Interest (AOIs) on vehicle and traffic control features (see figure 6) to map and analyse participants' gaze fixations. We used Dey et al.'s [11] vehicle AOIs as a basis and extended them to a 3D vehicle model as our study was conducted in the wild where cyclists could fixate on any part of a vehicle. For traffic control AOIs, we considered the UK highway code to specify traffic signs and road markings that organise the right of way in space-sharing conflicts. These included (1) traffic lights, (2) road works signs, (3) signs giving orders, e.g. stop signs, (4) across the carriageway road markings, e.g. stop lines and (5) warning of 'give way' ahead road markings.

We analysed *visit counts* to investigate riders' gaze behaviour in different traffic scenarios. A *visit* is defined as "the period of time when a participant first focuses on a region until the person looks away from that region. A visit consists of one or more fixations." [30]. All recordings were in the city of Glasgow. Our independent variable was the traffic scenario. Our dependent variables were: (1) the number of times an AOI was in a cyclist's field of view (regardless of whether the rider fixated on it) and (2) the number of visits riders had to an AOI (visit count).




Figure 7: A commuter cyclist equipped with (A) Tobii Pro Glasses 2, (B) recording unit attached to the eye-tracker, (C) Garmin Edge 530 bike computer and (D) Dell XPS 13 9300 laptop. Participants kept the laptop in their backpacks or bike bags.

4.2 Apparatus

Participants signed up through a survey hosted on Qualtrics. They included their commuting routes by plotting them on a map using RideWithGPS³. We used the Tobii Pro Glasses 2⁴ eye-tracker to record gaze fixations and video footage at a frame-rate of 100Hz with a resolution of 1920 x 1080 pixels. The glasses, which feature a recording unit, are portable and easy to calibrate and wear when cycling [27]. Participants were equipped with a Dell XPS 13 9300 laptop with Tobii Pro Glasses Controller⁵ installed to calibrate the eye-tracker and start/end a recording.

Participants were also instrumented with a Garmin Edge 530 bike computer, which collected GPS data, allowing us to visualise and confirm the cyclists' routes. Participants used their own bikes in the study. We equipped riders with a bicycle helmet and front/rear bicycle lights if they did not have them. All participants were provided with a copy of the UK Highway Code's rules for cyclists⁶. Figure 7 shows an instrumented commuter cyclist.

4.3 Participants

We recruited 12 commuter cyclists (*Male* = 8, *Female* = 4) (20-29 years old = 5, 30-39 = 5, 40-49 = 1, 50-59 = 1) through flyers and social media advertising. All participants were affiliated with the university to achieve ease of communication, primarily due to the requirement of handling expensive hardware. Participants had normal to corrected vision. None wore eyeglasses during the study. One participant had contact lenses on.

4.4 Procedure

Participants attended a briefing session where we explained the study's purpose and showed how to use the eye-tracker and associated software. We demonstrated what the collected data looked like by showing them sample footage from a pilot study. We then

³Ride With GPS: www.ridewithgps.com

⁴Tobii Pro Glasses 2: www.tobii.com/product-listing/tobii-pro-glasses-2

⁵Tobii Pro Glasses Controller Software: www.tobii.com/learn-and-support/learn/steps-in-an-eye-tracking-study/setup/installing-tobii-glasses-controller

⁶The UK Highway Code's Rules for Cyclists: www.gov.uk/guidance/the-highway-code/rules-for-cyclists-59-to-82



Figure 8: A video frame showing labelled AOIs at an uncontrolled intersection.

showed participants how to use the bike computer and informed them about pressing the *lap button* if they encountered something relevant to the study; this would record a timestamp in the event logfile, allowing us to revisit the event in our analysis. Participants had the opportunity to ask any questions during the session. After that, participants were tasked with recording two commutes - from their workplace to their residence and vice versa. Each participant kept the equipment overnight and returned it after their commute the following day. They were compensated with a £10 Amazon voucher. The University's ethics committee approved this study.

4.5 Data Validation and Pre-Processing

Tobii Pro Lab⁷ was used to label traffic scenarios within the footage and analyse gaze data. There was no predetermined list of scenarios to label, giving us the advantage of identifying spontaneous ones. Like Study 1, scenarios were labelled if there was a potential driver-cyclist space-sharing conflict, regardless of whether there was an interaction.

We ensured that all data were collected appropriately, i.e. the hardware recorded the entire commute. We visualised the bike computer route in case there were any changes to the RideWithGPS one, e.g. due to road works, and played the video footage overlayed with gaze samples to be familiar with the trip. We replayed the footage to label scenarios with potential space-sharing conflicts. Labels were *times of interest* (time range); starting at the first video frame a scenario appears in a participant's field of view and ending at the first frame where the scenario is no longer visible.

We used Tobii Pro Lab's AOI tool to manually (frame-by-frame) map fixations to AOIs in each scenario time range (see Figure 8). Like Study 1, we labelled the vehicle closest to the cyclist and the closest road marking/traffic sign where relevant. Three authors completed the process independently to eliminate bias. First, one author labelled all 24 videos, and a second labelled 6 of these chosen at random. Comparing the results showed no discrepancies. A third author labelled 3 of the 6 videos selected randomly. Comparing their results again showed no discrepancies.

4.6 Results

We collected 24 video clips with a total duration of 8 hours, 50 minutes and 25 seconds. We identified 171 scenarios featuring driver-cyclist encounters. Scenarios had different AOI distributions (see table 2), which could impact gaze behaviours. Therefore, we report

⁷Tobii Pro Lab eye tracking analysis software: www.tobii.com/products/software/data-analysis-tools/tobii-pro-lab

Traffic Scenario	N	AOI Distribution (%)														
		Side Window	Side Body	Roof	Bonnet	Windscreen	Front Bumper	Back Body	Back Bumper	Direction Indicator	Hazard Light	BrakeLight	Reverse Light	Traffic Light	Traffic Sign	Road Marking
Controlled Intersection	48	13.3	13.3	10	9.2	8.8	10.4	4.6	5	2.1	N/A	2.9	N/A	13.3	N/A	7.1
Controlled Crossing	19	15.9	15.9	7.9	15.9	12.7	12.7	3.2	3.2	1.6	N/A	N/A	N/A	9.5	N/A	1.6
Road Works	7	9.1	18.2	N/A	9.1	N/A	9.1	N/A	N/A	N/A	N/A	N/A	N/A	27.3	N/A	N/A
Uncontrolled Intersection	48	15.3	15.8	11.7	11.7	9.9	12.6	3.6	4.5	3.2	N/A	1.4	N/A	N/A	1.4	9
Roundabout	10	13.3	15.6	4.4	11.1	11.1	13.3	4.4	2.2	2.2	N/A	N/A	N/A	N/A	11.1	11.1
Cycle Lane into Traffic	2	14.3	14.3	7.1	14.3	14.3	14.3	N/A	N/A	N/A	N/A	N/A	N/A	N/A	7.1	14.3
Lane Merging (Cyclist)	5	14.3	14.3	4.8	9.5	4.8	4.8	9.5	14.3	NA	4.8	4.8	N/A	N/A	4.8	9.5
Overtaking (Driver)	18	19.2	20.5	11.5	5.1	1.3	7.7	16.7	15.4	N/A	N/A	2.6	N/A	N/A	N/A	N/A
Parking Manoeuvre	9	18.6	18.6	14	2.3	2.3	2.3	14	14	2.3	N/A	2.3	9.3	N/A	N/A	N/A
Bottleneck	7	9.1	15.2	12.1	18.2	18.2	15.2	3	3	N/A	3	N/A	3	N/A	N/A	N/A
Total	171	14.8	15.6	10.1	10.3	8.7	6.1	5.8	6.1	1.9	0.3	1.8	0.7	5.3	1.7	6.1

Table 2: Labelled traffic scenarios and the proportion of AOIs in them. N = Frequency of Scenario Appearance, N/A= AOI did not appear in cyclists' field of view. Bottlenecks happen when the driver and cyclist move in the same lane in opposite directions.




Figure 9: AOI visit distributions in scenarios. Refer to Table 2 for the frequency of each scenario and the AOIs that appeared in them.

general patterns in our data before comparing visit counts in scenarios with similar AOI combinations to minimise their variation.

4.6.1 Visit Count Patterns. We grouped Direction Indicators, Hazard, Brake and Reverse Lights into one category: *On-Vehicle Signals* to simplify our visualisations (see Figure 9) and conduct further analysis. Cyclists visited different vehicle sides throughout their commutes. For example, 27.6% of visits in uncontrolled intersections and 58.3% in parking manoeuvres were to side bodies, with 40% in overtaking manoeuvres to back bodies. The vehicle's front was mostly visited in bottlenecks (e.g. the bonnet was 34.5% of visits). On-vehicle signals were our smallest AOI, but cyclists still visited them in multiple scenarios. They formed 2.6% of visits in uncontrolled intersections, 1.1% in controlled intersections, and 20% in overtaking manoeuvres.

Windscreens or windows are often visited with the expectation of social interaction with the driver [11]. Cyclists mostly visited vehicle windscreens or windows in scenarios with little traffic control. For example, 27.6% at bottlenecks and 24.9% at roundabouts. In comparison, only 13.7% of visits in controlled intersections were to these AOIs, despite only a 2.3% difference in their appearance between controlled intersections and roundabouts. Road markings were the most common traffic control feature and were often visited when they were present. For example, 25% in roundabouts and 16% in controlled intersections. Traffic lights impacted cyclist gaze behaviour, with 17.2% in controlled intersections and 29.6% of road work visits to traffic lights. Although, cyclists visited traffic signs (48.1%) more frequently than traffic lights in road works.

4.6.2 Visits in Scenario Categories. We grouped the traffic scenarios into three categories: (1) *controlled scenarios*: controlled intersections, crossings and road works; (2) *uncontrolled infrastructure*: uncontrolled intersections, roundabouts and cycle lanes merging

into traffic; (3) *dynamic manoeuvres*: bottlenecks, lane merging, parking and overtaking. We used a Chi-Square test of independence to explore the relationship between scenarios and visit counts for each category. *Post hoc* tests were performed using a Chi-Square test of independence with a Bonferroni correction.

We found a significant association between traffic scenarios and AOI visit counts for *controlled scenarios* ($\chi^2(22, 104) = 54.55$, $P < .001$). *Post hoc* comparisons showed significant results between controlled intersections and road works ($\chi^2(11, 102) = 50.34$, $P < .001$); cyclists were more likely to visit vehicle features in controlled intersections compared to road works, where they mostly visited traffic control AOIs. No statistically significant association was found between the variables for *uncontrolled infrastructure* ($\chi^2(20, 90) = 17.2$, $P = .64$). Finally, results were significant for *dynamic manoeuvres* ($\chi^2(24, 56) = 50.49$, $P = .0012$). *Post hoc* comparisons show cyclists were likely to visit the vehicle's windscreen and front (e.g. bonnet and front bumper) in bottleneck scenarios, but the side and back of the vehicle in parking manoeuvres ($\chi^2(7, 41) = 24.82$, $P = .0049$).

4.7 Discussion

We answered RQ2 by investigating cyclist gaze behaviours in diverse traffic scenarios and showed that different AOI combinations within scenarios impacted gaze behaviour. We discuss the influence of traffic control, vehicle features and overall AOI distributions on gaze behaviours with special consideration to AV-cyclist interaction design.

4.7.1 Traffic Control. Road markings are a shared component between controlled intersections and scenarios in the uncontrolled infrastructure category. Riders often visited these AOIs. This is similar to Trefzger et al.'s [29] finding that cyclists focused more on their path than the surrounding environment, which could be due




Figure 10: Overall fixation counts of all participants in each traffic scenario. These data were the source for calculating visit counts. The dots represent the number of fixations. Green represents smaller numbers and red represents larger numbers.

to ergonomics, i.e. riders not having to lift their heads to see AOIs. We offer new opportunities for interface placements; cyclists could respond well to interfaces placed beyond the vehicle. While this is similar to Hou et al.'s [15] finding with cyclists reacting positively to road projections from the vehicle when lane merging, our results showed that this effect could extend to other scenarios, such as uncontrolled intersections and roundabouts, where the cyclist or AV is likely stationary.

As in Study 1, cyclists responded well to traffic lights with minimal visits to vehicle AOIs in controlled scenarios, suggesting that novel interfaces may be less needed here. As long as AVs correctly adhere to traffic rules, making them predictable for cyclists, designers can prioritise scenarios with more complex social interactions. This also shows that cyclists acknowledge on-environment interfaces, motivating future work to explore interfaces placed on the environment to support interaction in other scenarios, such as cycle lanes merging into traffic. This was investigated with pedestrians (who were accepting of these interfaces) [18], but not cyclists. Participants were more likely to visit traffic control AOIs than vehicle ones in road works compared to controlled intersections, which could be due to road works being described as *dynamic intersections* [5]. They borrow components such as traffic signs from stationary infrastructures but are temporary, and cyclists may not expect them. Road works can also be challenging for AVs to navigate as they may not be in the map database [5]. These factors make them an interesting scenario to explore, and we showed the available AOIs that would make road work simulations more accurate.

While road works may have traffic lights to determine the right of way, interaction may sometimes be needed. Vehicles could use road projections to communicate with riders; this may bridge the gap between them and controlled intersections. Our findings with cyclists having more visits to traffic signs than lights at road works highlight the efficiency and effectiveness of abstractions in traffic lights, a minimal light-based design with no writing. Future work

should further compare traffic lights' effects against traffic signs, as this could directly influence the modalities and abstractions taken by AV-cyclist interfaces.

4.7.2 Vehicle Features. Previous work proposed eHMIs as solutions for AV interactions [8]. Our results support this consensus; eHMIs would be a practical solution as cyclists encounter vehicles in highly diverse traffic scenarios, including dynamic manoeuvres, where interfaces may not be feasible to place on the road or environment. Our findings directly impact how eHMIs should be designed and operated around cyclists. First, cyclists could be anywhere around the vehicle, and we found that they visit different vehicle sides throughout their commutes. Placing an interface on a single part of the AV could result in ambiguities due to cyclists not seeing the AV's messages. Second, cyclists were most frequently exposed to and visited the vehicle's side body (its biggest component) in various scenarios, giving designers a large surface to deploy more unrestrained displays.

Third, exploring dynamic manoeuvres showed that both the vehicle and cyclist could be moving, impacting gaze behaviours. Riders prioritised certain vehicle AOIs to visit; for example, the windscreen in bottlenecks with 27.6% of visits. This could be due to the higher velocities and cyclists not having much time to decide on their next manoeuvre compared to uncontrolled infrastructures where one of the road users is usually stationary, giving cyclists time to visit different AOIs. More complex eHMIs with multiple components could accommodate interactions at dynamic manoeuvres by only having a specific component working at a time. Designers can refer to our findings to understand which components should be active in a scenario. Finally, cyclists visited on-vehicle signals in multiple scenarios, suggesting that riders rely on them to navigate traffic. Replacing signals could increase the learning curve for cyclists riding around AVs, so eHMIs must co-exist with these light-based

communication mechanisms. Future work can further explore on-vehicle signals to identify animation and colour patterns that can be extended to novel interfaces.

4.7.3 AOI Distributions. The study offered a first-person perspective of riding in mixed urban traffic. With 171 driver-cyclist encounters in diverse traffic scenarios, we quantified the frequency of each scenario, which could help designers prioritise scenarios to evaluate concepts. We identified the different AOI combinations and distributions between scenarios, which could result in more accurate simulations and ecologically valid interface evaluations. The varying AOIs between scenarios impacted gaze behaviours, motivating using a larger interface design space (beyond the vehicle) to address cyclists' natural gaze behaviours and ensure interface messages are always in cyclists' fields of view.

5 LIMITATIONS AND FUTURE WORK

There are some limitations that could reduce the generalisability of our results. Both were conducted in the city of Glasgow. While representative of many urban areas in the UK, it is unknown how our findings would generalise to other countries with different road infrastructure. Future work should conduct similar studies in different regions worldwide, similar to Lee et al. [17], who conducted observations of driver-pedestrian interactions in different European cities. Both studies were conducted in urban areas, but cyclists also encounter drivers on rural roads. We investigated encounters on urban roads as most cycling traffic and vehicle-cyclist collisions happen there [1]. Both our studies considered only one-to-one driver-cyclist interactions. However, multiple drivers or cyclists may be present in an encounter, with future work needed to explore *scalability*. This is also a challenge yet to be resolved in the AV-pedestrian interaction domain [10]. We chose single interactions as our starting point as it gives baseline knowledge that others could extend to more complex environments. We focused on *versatility* to cover a wider range of traffic scenarios. Both of these issues must be resolved for AVs to work safely in diverse and complex traffic scenarios.

We observed only one site per traffic scenario in the daytime during the summer (June-August) for Study 1. We ensured that the site had high cycling traffic and a 30mph speed limit. Future work could observe scenarios with different settings, for example, during winter, nighttime or at higher speeds, as these factors may impact interaction behaviour. Our observations were complemented by the naturalistic study in which participants took their own routes, allowing us to capture cycling behaviour in a wider set of scenarios with different features.

While we could have further differentiated the AOIs in Study 2, e.g. by splitting the windscreen between the driver and passenger sides, the study's naturalistic nature (where road users were at varying speeds, angles and distances, with diverse lighting/weather conditions and locations) required some AOI simplification due to the precision of the Tobii eye-tracker. Therefore, we adopted elements of Dey et al.'s [11] approach, where they used the windscreen, for example, as a single entity. Future work could explore more specific and detailed AOIs, with our work guiding where more detailed AOI analysis would be appropriate. A limitation of Study 2 was the small number of riders sampled. This was due

to the complexity of the study (participants kept the equipment overnight) and the requirement of recruiting commuter cyclists affiliated with the University (for ethics and insurance purposes). However, participants took different routes across the city, allowing us to capture over 8 hours of footage in many traffic scenarios. Future researchers should consider longer-duration rides to extend the analysis period. Despite this limitation, we found participants had similar gaze patterns throughout their commutes, suggesting that our data is representative of city cycling in the UK.

6 OVERALL DISCUSSION AND IMPLICATIONS FOR AV RESEARCH

We used the findings from both studies to form design guidelines for versatile AV-cyclist interfaces. Study 2 showed potential in using the road, vehicle and environment as part of a holistic design space. There are three interpretations to addressing this design space alongside our guidelines. First, designers may develop a versatile interface placed on a single entity, e.g. the road or vehicle. This would not use the larger design space and could challenge designers who must ensure that the interface is always in a cyclist's field of view. Study 1 showed that road users exchange a wide range of messages, and designers should ensure that cyclists can differentiate between these messages; this would require establishing a clear design language that classifies factors such as light colour for light-based displays. Another challenge is allowing a smooth transition between the interface states in traffic scenarios close to each other, for example, an uncontrolled intersection leading to a roundabout. Nevertheless, this approach could be useful to cyclists, who would not need to shift their attention beyond a single interface.

In the second interpretation, designers can use the guidelines to develop interfaces with different placements which work independently from each other. For example, vehicles could have eHMLs, cyclists may be wearing AR glasses, and there may be an on-road interface, with all interfaces working individually to facilitate interaction. Designers do not need to consider connectivity between individual interfaces, which is useful in remote areas where that is not possible. It also protects cyclists' privacy, as their location may not be revealed to external parties [6]. The approach could preserve cyclists' safety; if one interface goes down, others could still facilitate interaction. Still, this approach may challenge cyclists, who must learn to interact with different interfaces and potentially conflicting messages; having the interfaces work independently may send contradictory messages that could impact cyclists' safety.

The third interpretation bridges the previous two, with multiple interconnected interfaces working as part of a single holistic interface. Individual interfaces could have different placements to suit natural gaze behaviour. Designers can split their designs into smaller interfaces to avoid being restricted by a single placement or modality, allowing cyclists to differentiate between messages easily. Having interfaces work together also means they could use the same *design language*, requiring a lower mental effort from cyclists. The approach could help designers address *scalability*; AVs could send direct messages to cyclists wearing interfaces such as AR glasses. Still, designers should avoid sending unnecessary information causing visual clutter and overwhelming cyclists when taking this approach. Figure 11 illustrates an early sketch of a prototype




Figure 11: A prototype AV-cyclist interface implemented using the third interpretation of the guidelines. A light-based eHMI uses a three-colour design language inspired by traffic lights. Red = cyclist cannot proceed (or negative feedback when active all around the vehicle), yellow = awareness and right-of-way is up for negotiation when appropriate, green = cyclist can proceed (or positive feedback when active all around the vehicle). The eHMI uses different animations and colour combinations to interact with cyclists. Cyclists can use multimodal AR glasses that can communicate with the AV for additional support. The glasses use the same design language as the eHMI by extending its colours. Five initial example states are proposed. States 1-3 could be useful when interacting at uncontrolled infrastructure. States 4 and 5 could be useful in controlled scenarios and dynamic manoeuvres.

holistic AV-cyclist interface implemented using this interpretation. It involves a light-based eHMI that uses different animations and colour combinations for interaction between drivers and cyclists. The eHMI includes a sensor to detect cyclist gestures. Cyclists can also use AR glasses connected to the eHMI for extra support.

We showed that versatility is an important challenge that must be resolved before AVs are integrated into traffic. Study 1 revealed that road users interact differently between the scenarios, and Study 2 demonstrated how different AOI combinations in traffic scenarios influence gaze behaviours. These results could impact the design of versatile AV-cyclist interfaces, and we contribute the following guidelines to help designers address this. The guidelines were developed for SAE level 5 AVs (no human driver in any traffic scenario), but some may apply to AVs with other automation levels.

Driver-cyclist communication is two-way. Study 1 showed that interactions are not limited to drivers communicating messages to cyclists; riders also share messages with drivers, for example, when negotiating right of way at uncontrolled intersections. Designers must accommodate the back-and-forth between drivers and cyclists. AVs should not only communicate messages to riders but also recognise their social and implicit cues and respond appropriately. For example, AVs must recognise cyclists' arm gestures in lane merging manoeuvres and decelerate if they intend to let riders proceed. This may require more precise sensing than is common in most AVs, which can only detect the presence of riders and not more subtle body movements.

The road as a design space for interfaces. Study 2 showed that road markings occur in most stationary road infrastructure, and riders regularly visit them. The road could be used as an interaction space, for example, using projections from vehicles or AR glasses worn by cyclists to facilitate interactions. Designers should use this to their advantage. The road offers a larger display space than a vehicle, so AVs can communicate more comprehensive messages, e.g. right of way negotiations, in scenarios where they are likely stationary, giving riders the time and space to read the messages.

Interfaces should not overwhelm cyclists. Road users may exchange multiple messages when interacting, e.g. right of way negotiation followed by feedback or the same message through social and implicit cues. Designers should avoid overwhelming cyclists with unnecessary information, especially in scenarios where road users are moving. AVs could communicate messages sequentially to avoid overwhelming cyclists. For example, an AV could communicate its awareness at a bottleneck scenario, followed by its intent to yield once the cyclist knows the vehicle is aware of them.

eHMI messages should be perceivable anywhere around the vehicle. Different vehicle parts (front, sides and back) were in cyclists' fields of view throughout their commutes in Study 2. Designers should ensure that cyclists can receive the AV's messages from anywhere around the vehicle. To avoid overwhelming cyclists (and potentially communicating information to the wrong rider), only the parts of the eHMI in a cyclist's field of view should be used. Again, additional sensing on the vehicle may be required to detect this.

Messages should be communicated in the right place at the right time. Timing is critical in AV-cyclist interaction. Study 2 showed cyclists and drivers could both be moving when interacting, so AVs must communicate messages without delay in such scenarios, e.g. overtaking. The vehicle may be behind the cyclist or out of view. Designers should ensure cyclists can always receive the AV's messages. For example, the vehicle should not use hazard lights when the cyclist is not facing the vehicle in lane merging scenarios; other solutions, such as directional audio, may be more appropriate.

New interfaces must co-exist with current on-vehicle signals. Study 1 showed that drivers frequently used existing on-vehicle signals, e.g. direction indicators, and Study 2 showed that cyclists visited these signals along with implicitly triggered ones such as reversing lights. Designers should develop interfaces that work concurrently with these signals and do not obstruct or conflict with them, minimising the learning curve for cyclists. It could also make it easier to develop more effective interfaces, as designers can primarily focus on replacing lost social cues.

Interfaces must reflect implicit cues. Drivers often use implicit cues to support explicitly communicated messages, e.g. decelerating and using hand gestures to signal a cyclist to proceed. AV driving behaviours should reflect the interface messages and not contradict them, as this could cause ambiguities and confusion. AV implicit cues should be predictable for cyclists, as miscommunications can arise from scenarios where, for example, an AV decelerates once a cyclist indicates their intent to merge lanes but the AV does not intend to allow for the cyclist to complete the manoeuvre.

Interfaces do not need to be restricted to communicating a single message. Study 2 showed that some scenarios have similar AOIs, for example, uncontrolled intersections and roundabouts. However, Study 1 revealed that interactions differ between scenarios, and designers should accommodate these differences. Therefore, even though interfaces such as eHMIs may have the same placements, they need to communicate different messages to operate in different scenarios.

Positive/negative feedback should not be overlooked. Study 1 showed that positive and negative feedback are components of natural interaction behaviour. Designers should consider facilitating the exchange of feedback between AVs and cyclists to gain feedback on the AV's driving performance in real traffic or create expressive AVs that do not obstruct the current social traffic paradigm.

These interfaces must be comprehensible and predictable. The guidelines do not specify whether interface behaviours should mimic those of human drivers, but we discuss the benefits and drawbacks of utilising current interaction behaviour. For example, we showed the benefits of embracing traditional on-vehicle signals when designing future interfaces. We stress that designers should take care of how AVs use implicit cues. Given the interactions observed in scenarios such as lane merging, having AVs exhibit different implicit cues to human drivers would disrupt the current traffic paradigm and other road users' expectations. Designers can use our guidelines to establish a novel design language that replaces social cues. However, this language must be comprehensible and avoid visual clutter to not overwhelm cyclists. Alternatively, future

work could explore mimicking drivers' social cues; this was done successfully with pedestrians [19]. Still, going beyond drivers' social cues may result in a more usable interaction experience. For example, cyclists naturally fixate on the road, motivating its use as a design space for interfaces that benefit riders.

7 CONCLUSION

We investigated driver-cyclist interaction in real-world traffic to inform AV-cyclist interface design. We conducted a set of in-the-wild observations ($N = 414$) of driver-cyclist encounters in five traffic scenarios. Road users exchanged different messages, social cues and implicit cues between the scenarios, suggesting that AVs must exhibit different behaviours depending on the traffic scenario. We then conducted a naturalistic study ($N = 12$) with cyclists wearing eye-trackers to gain a first-person perspective of real-world commutes. We provided empirical evidence showing that cyclists could be anywhere around vehicles on the road, and their gaze behaviour changes between traffic scenarios. Our findings inform the potential placements and design of AV-cyclist interfaces. We concluded the paper with guidelines to assist AV-cyclist interface designers in developing versatile interfaces that work in diverse traffic scenarios. Our results establish a foundation for AV-cyclist interaction research and provide knowledge enabling designers to create AVs that can effectively communicate with cyclists and drive safely around them.

ACKNOWLEDGMENTS

This publication was supported by the Excellence Bursary Award from the University of Glasgow and the Royal Society of Edinburgh (RSE) grant number 1931.

REFERENCES

- [1] 2020. Reported road casualties in Great Britain: pedal cycle factsheet, 2020 - GOV.UK. <https://www.gov.uk/government/statistics/reported-road-casualties-great-britain-pedal-cyclist-factsheet-2020/reported-road-casualties-in-great-britain-pedal-cycle-factsheet-2020>
- [2] 2022. Speed limits - GOV.UK. <https://www.gov.uk/speed-limits>
- [3] Ammar Al-Taie, Frank Pollick, and Stephen Brewster. 2022. Tour de Interaction: Understanding Cyclist-Driver Interaction with Self-Reported Cyclist Behaviour. In *14th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. ACM, New York, NY, USA, Seoul, Republic of Korea.
- [4] James Anderson, Nidhi Kalra, Karlyn Stanley, Paul Sorensen, Constantine Samaras, and Oluwatobi Oluwatola. 2016. *Autonomous Vehicle Technology: A Guide for Policymakers*. RAND Corporation. <https://doi.org/10.7249/RR443-2>
- [5] Shunsuke Aoki and Ragunathan Rajkumar. 2018. Dynamic Intersections and Self-Driving Vehicles. *Proceedings - 9th ACM/IEEE International Conference on Cyber-Physical Systems, ICCPS 2018* (8 2018), 320–330. <https://doi.org/10.1109/ICCPS.2018.00038>
- [6] Siri Hegna Berge, Marjan Hagenzieker, Haneen Farah, and Joost de Winter. 2022. Do cyclists need HMIs in future automated traffic? An interview study. *Transportation Research Part F: Traffic Psychology and Behaviour* 84 (1 2022), 33–52. <https://doi.org/10.1016/j.trf.2021.11.013>
- [7] Bryan Botello, Ralph Buehler, Steve Hankey, Andrew Mondschein, and Zhiqiu Jiang. 2019. Planning for walking and cycling in an autonomous-vehicle future. *Transportation Research Interdisciplinary Perspectives* 1 (6 2019), 100012. <https://doi.org/10.1016/J.TRIP.2019.100012>
- [8] Debargha Dey, Azra Habibovic, Andreas Löcken, Philipp Wintersberger, Bastian Pfleging, Andreas Riener, Marieke Martens, and Jacques Terken. 2020. Taming the eHMI jungle: A classification taxonomy to guide, compare, and assess the design principles of automated vehicles' external human-machine interfaces. *Transportation Research Interdisciplinary Perspectives* 7 (9 2020), 100174. <https://doi.org/10.1016/J.TRIP.2020.100174>

- [9] Debargha Dey and Jacques Terken. 2017. Pedestrian interaction with vehicles: Roles of explicit and implicit communication. *AutomotiveUI 2017 - 9th International ACM Conference on Automotive User Interfaces and Interactive Vehicular Applications, Proceedings* (9 2017), 109–113. <https://doi.org/10.1145/3122986.3123009>
- [10] Debargha Dey, Arjen van Vastenhoven, Raymond H. Cuijpers, Marieke Martens, and Bastian Pfleging. 2021. Towards Scalable eHMI: Designing for AV-VRU Communication Beyond One Pedestrian. In *13th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. ACM, New York, NY, USA, 274–286. <https://doi.org/10.1145/3409118.3475129>
- [11] Debargha Dey, Francesco Walker, Marieke Martens, and Jacques Terken. 2019. Gaze patterns in pedestrian interaction with vehicles: Towards effective design of external human-machine interfaces for automated vehicles. *Proceedings - 11th International ACM Conference on Automotive User Interfaces and Interactive Vehicular Applications, AutomotiveUI 2019* (9 2019), 369–378. <https://doi.org/10.1145/3342197.3344523>
- [12] Marjan P Hagenzieker, Sander Van Der Kint, Luuk Visser, Ingrid N L G Van Schagen, Jonathan De Bruin, Paul Van Gent, Jacques J F Commandeur, Å Marjan, and P Hagenzieker. 2019. Interactions between cyclists and automated vehicles: Results of a photo experiment Interactions between cyclists and automated vehicles: Results of a photo experiment. *Journal of Transportation Safety & Security* 12, 1 (2019), 94–115. <https://doi.org/10.1080/19439962.2019.1591556>
- [13] Siri Hegna Berge, Joost de Winter, and Marjan P Hagenzieker. 2022. Support systems for cyclists in automated traffic: A review and future outlook. (2022). <https://doi.org/10.13140/RG.2.2.29549.05606>
- [14] Kai Holländer, Mark Colley, Enrico Rukzio, and Andreas Butz. 2021. A Taxonomy of Vulnerable Road Users for HCI Based On A Systematic Literature Review. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–13. <https://doi.org/10.1145/3411764.3445480>
- [15] Ming Hou, Karthik Mahadevan, Sowmya Somanath, Ehud Sharlin, and Lora Oehlberg. 2020. Autonomous Vehicle-Cyclist Interaction: Peril and Promise. In *Conference on Human Factors in Computing Systems - Proceedings*. Association for Computing Machinery. <https://doi.org/10.1145/3313831.3376884>
- [16] Nóra Krizsik and Tibor Sipos. 2022. Social Perception of Autonomous Vehicles. *SSRN Electronic Journal* (2022). <https://doi.org/10.2139/ssrn.4060916>
- [17] Yee Mun Lee, Ruth Madigan, Oscar Giles, Laura Garach-Morcillo, Gustav Markkula, Charles Fox, Fanta Camara, Markus Rothmueller, · Signe, Alexandra Vendelbo-Larsen, · Pernille, Holm Rasmussen, Andre Dietrich, Dimitris Nathanael, Villy Portouli, Anna Schieben, and Natasha Merat. 2021. Road users rarely use explicit communication when interacting in today's traffic: implications for automated vehicles. *Cognition, Technology & Work* 23 (2021), 367–380. <https://doi.org/10.1007/s10111-020-00635-y>
- [18] Karthik Mahadevan, Sowmya Somanath, and Ehud Sharlin. 2018. Communicating Awareness and Intent in Autonomous Vehicle-Pedestrian Interaction. (2018). <https://doi.org/10.1145/3173574.3174003>
- [19] Karthik Mahadevan, Sowmya Somanath, and Ehud Sharlin. 2018. Communicating awareness and intent in autonomous vehicle-pedestrian interaction. In *Conference on Human Factors in Computing Systems - Proceedings*, Vol. 2018-April. Association for Computing Machinery. <https://doi.org/10.1145/3173574.3174003>
- [20] G Markkula, R Madigan, D Nathanael, E Portouli, Y M Lee, A Dietrich, J Billington, A Schieben, and N Merat. 2020. Theoretical Issues in Ergonomics Science Defining interactions: a conceptual framework for understanding interactive behaviour in human and automated road traffic. (2020). <https://doi.org/10.1080/1463922X.2020.1736686>
- [21] Milecia Matthews, Girish V Chowdhary, and Emily Kieson. 2017. Intent Communication between Autonomous Vehicles and Pedestrians. (8 2017). <http://arxiv.org/abs/1708.07123>
- [22] Andrii Matvienko, Florian Müller, Dominik Schön, Paul Seesemann, Sebastian Günther, and Max Mühlhäuser. 2022. BikeAR: Understanding Cyclists' Crossing Decision-Making at Uncontrolled Intersections using Augmented Reality. In *CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–15. <https://doi.org/10.1145/3491102.3517560>
- [23] Hannah R M Pelikan. 2021. Why Autonomous Driving Is So Hard: The Social Dimension of Traffic; Why Autonomous Driving Is So Hard: The Social Dimension of Traffic. (2021). <https://doi.org/10.1145/3434074.3447133>
- [24] Petr Pokorný and Kelly Pitera. 2019. Observations of truck-bicycle encounters: A case study of conflicts and behaviour in Trondheim, Norway. *Transportation Research Part F: Traffic Psychology and Behaviour* 60 (1 2019), 700–711. <https://doi.org/10.1016/j.trf.2018.11.018>
- [25] Petr Pokorný, Belma Skender, Torkel Bjørnskau, and Marjan P. Hagenzieker. 2021. Video observation of encounters between the automated shuttles and other traffic participants along an approach to right-hand priority T-intersection. *European Transport Research Review* 13, 1 (12 2021), 59. <https://doi.org/10.1186/s12544-021-00518-x>
- [26] 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 (3 2020), 900–918. <https://doi.org/10.1109/TITS.2019.2901817>
- [27] Megan S. Ryerson, Carrie S. Long, Michael Fichman, Joshua H. Davidson, Kristen N. Scudder, Michelle Kim, Radhika Katti, George Poon, and Matthew D. Harris. 2021. Evaluating cyclist biometrics to develop urban transportation safety metrics. *Accident Analysis & Prevention* 159 (9 2021), 106287. <https://doi.org/10.1016/J.AAP.2021.106287>
- [28] Katayoun Salamat, Bastian J. Schroeder, Duane R. Geruschat, and Nagui M. Rouphail. 2013. Event-Based Modeling of Driver Yielding Behavior to Pedestrians at Two-Lane Roundabout Approaches. *Transportation Research Record: Journal of the Transportation Research Board* 2389, 1 (1 2013), 1–11. <https://doi.org/10.3141/2389-01>
- [29] Mathias Trefzger, Tanja Blaschek, Michael Raschke, Sarah Hausmann, and Thomas Schlegel. 2018. A visual comparison of gaze behavior from pedestrians and cyclists. *Eye Tracking Research and Applications Symposium (ETRA)* (6 2018). <https://doi.org/10.1145/3204493.3204553>
- [30] Mathias Trefzger, Waldemar Titov, Karol Sgodzaj, and Thomas Schlegel. 2020. Analysis and Comparison of the Gaze Behavior of E-Scooter Drivers and Cyclists Depending on Road Surface Quality in a Real Test Environment. *Lecture Notes in Informatics (LNI), Proceedings - Series of the Gesellschaft für Informatik (GI) P-307* (2020), 1163–1180. https://doi.org/10.18420/LNF2020_{ }108
- [31] J. Uttley, Y.M. Lee, R. Madigan, and N. Merat. 2020. Road user interactions in a shared space setting: Priority and communication in a UK car park. *Transportation Research Part F: Traffic Psychology and Behaviour* 72 (7 2020), 32–46. <https://doi.org/10.1016/j.trf.2020.05.004>
- [32] Tamara Von Sawitzky, Philipp Wintersberger, Andreas Löcken, Anna-Katharina Frison, and Andreas Riener. 2020. Augmentation Concepts with HUDs for Cyclists to Improve Road Safety in Shared Spaces. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–9. <https://doi.org/10.1145/3334480.3383022>