Directional Light-based External Human-Machine Interface with Onboarding

Existing eHMIs for automated vehicles face a design tradeoff: text-based interfaces communicate explicitly but are language-dependent and difficult to scale, whereas light-based designs are easier to deploy yet often convey only coarse go/no-go information. We present an informative light-based windshield eHMI that adds richer yielding cues while maintaining a minimal visual footprint. In a virtual-reality study ($N = 30$), participants encountered a yielding AV with and without eHMI; half received a short onboarding video, allowing us to compare trained and untrained pedestrians in repeated trials. The eHMI reduced crossing initiation time, with larger early benefits for trained participants. For perceived safety and trust, trained participants improved immediately, whereas untrained participants improved only after repeated exposure. These findings suggest that minimal directional eHMIs can reduce ambiguity in AV–pedestrian yielding interactions, and that brief onboarding can accelerate learnability of non-textual signals.

Additional Key Words and Phrases: External Human-Machine Interfaces, Pedestrian Crossing Behaviours, Training Effect, Learnability, Trust, Perceived Safety

ACM Reference Format:
. 2026. Directional Light-based External Human-Machine Interface with Onboarding. 1, 1 (April 2026), 13 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

As automated vehicles (AVs) assume more of the driving task, pedestrians lose access to familiar communicative cues such as eye contact, gaze direction, and conventional hand and head gestures that help them judge whether a vehicle has noticed them and intends to yield [12]. In human-driven traffic, these cues shape pedestrian crossing decisions and can increase perceived safety [21, 22, 24]. Therefore, their absence creates uncertainty in AV–pedestrian encounters, with implications not only for safety, but also for pedestrian trust in the vehicle and the fluency of the interaction [16]. Designing AV signals that support clear and timely pedestrian interpretation is, therefore, a central challenge for human-centred vehicle design.

Existing external human-machine interface (eHMI) designs and evaluations span a wide range of communication modalities, including text, symbols, projections, and lights, with light- and text-based concepts representing two prominent directions in the literature [2, 6]. However, despite promising results, an important design problem remains unresolved: text-based eHMIs are often more explicit but harder to broadly deploy due to language, legibility, and liability concerns, whereas light-based eHMIs are easier to deploy yet often provide less explicit information and can be harder to interpret without prior exposure or learning [2, 11]. This challenge is particularly important because pedestrians are among the most vulnerable road users (VRUs), and VRUs as a whole account for more than half of global road traffic deaths [26].

Light-based eHMIs have attracted sustained attention because they are visually lightweight, language-independent, and potentially easier to integrate into real vehicles than more explicit visual interfaces such as text-based displays [15].

Author's Contact Information:

Manuscript submitted to ACM

Their practical appeal has made them a prominent direction in AV–pedestrian communication research, and prior work has reported several benefits for pedestrian interaction. For example, front brake light concepts have been associated with greater perceived safety and, in simulation-based analyses, with potential reductions in certain intersection crashes, while dynamic light patterns have been found to increase the willingness of pedestrians to cross [5, 25]. Some light-based designs have also gone beyond general yielding cues by adding contextual or recipient-specific information. For example, Dey et al.'s situational-awareness eHMI used the windshield to indicate the pedestrian for whom the AV was yielding and improved ratings such as dependability, attractiveness, and efficiency, as well as willingness to cross [7]. At the same time, much of this work suggests that light-based signals often communicate only coarse yielding intent and may be harder to interpret at first encounter than more explicit alternatives, particularly when pedestrians do not receive prior explanation [2, 11, 13]. Thus, while light-based eHMIs are attractive from a deployment perspective, it is unclear how to make them more informative without sacrificing their practical advantages.

Text-based eHMIs represent the other side of this design space. By communicating through explicit verbal messages, they can make vehicle intent easier to understand at first encounter [23]. Therefore, previous work has often assessed text-based eHMIs in terms of clarity, comprehension, persuasiveness, and perceived safety. Survey-based studies found that textual concepts were generally regarded the clearest among the eHMI designs presented to pedestrians, while experiments on message phrasing showed that egocentric messages such as "WALK" or "DON'T WALK" were interpreted more easily than more ambiguous formulations [2, 9]. These findings make text-based eHMIs attractive from a communication perspective. However, their practical deployment remains challenging: text-based signals are language-dependent and raise concerns related to legibility, liability, and broad standardisation across traffic contexts [2, 11]. As a result, although text-based eHMIs may increase explicitness, they do not resolve the broader challenge of developing AV signals that are both sufficiently informative for immediate interpretation and practically deployable at scale.

1.1 Directionality of External Human-Machine Interfaces

Despite these advances, an important gap remains in the information quality of existing eHMIs. Many current concepts primarily communicate a coarse go/no-go message without clarifying how that signal relates to a specific pedestrian in context [2, 11]. This can make pedestrians uncertain of whether the vehicle's behaviour is directed at them, even when the vehicle is clearly slowing or turning. This ambiguity matters not only for perceived safety and trust, but also for how quickly pedestrians translate the interpreted intent into a crossing decision. In AV–pedestrian encounters, delays in this decision prolong uncertainty and reduce interaction fluency; in future traffic environments with multiple AVs and pedestrians, rapid interpretation of yielding signals may become even more important. Prior work has often focused on willingness to cross, clarity, perceived safety, and trust, but less is known about how richer light-based cues may support faster, action-proximal pedestrian decisions without sacrificing deployability.

To address this gap, we propose an informative light-based windshield eHMI designed to provide richer yielding cues while preserving the practical advantages of light-based signalling. The eHMI is a minimal light strip positioned at the bottom of the windshield. When the AV yields, a brightened segment of the strip shifts along the horizontal axis according to the relative position of the pedestrian, allowing the interface to convey more contextual information than the coarse go/no-go signals. This design aims to reduce ambiguity while maintaining language independence and a small visual footprint. The windshield was chosen because previous work has identified it as a promising location for pedestrian-facing eHMIs and one that aligns more closely with where pedestrians expect driver-related signals [7, 10].

By combining directional information with a minimal and deployable form factor, the eHMI is intended to support earlier pedestrian decisions and improve perceived safety and trust during AV–pedestrian interaction.

1.2 Onboarding for External Human-Machine Interfaces

A second challenge concerns the onboarding (learnability) of non-textual eHMIs. Even if a light-based signal is practical to deploy, pedestrians must still learn how to interpret it, especially when it departs from familiar traffic conventions. Previous work suggests that familiarity and repeated exposure shape how pedestrians use eHMIs: familiar signals can lead to earlier crossings, repeated encounters can increase trust and reduce crossing times, and learnt reliance on an eHMI can even contribute to misuse when the signal is interpreted too confidently in ambiguous situations [4, 18–20]. These findings indicate that the effectiveness of an eHMI depends not only on what it communicates, but also on how its meaning is acquired over time. This issue is particularly relevant for light-based interfaces, which can be harder to interpret than text-based alternatives without prior explanation or exposure [2, 11]. However, compared with research on interface form and message content, brief onboarding as an explicit design factor remains comparatively underexplored, especially for minimal light-based signals and in studies that compare trained and untrained pedestrians across repeated encounters. This motivates the inclusion of a brief onboarding component in the present study.

1.3 Aim of Study

To address these gaps, we conducted a virtual-reality study in which pedestrians encountered a yielding AV with and without eHMI, while half of the participants received a brief onboarding video before the experiment. This design allowed us to examine not only whether the eHMI supports a more confident and earlier pedestrian interpretation of the AV yield intent, but also how its interpretation differs between trained and untrained pedestrians in repeated encounters. We focus on two complementary measures: crossing initiation time as an action-proximal behavioural indicator of how quickly pedestrians commit to crossing once they interpret yielding intent, and perceived safety and trust as subjective indicators of interaction quality. In this way, the present study contributes an informative yet minimal directional light-based eHMI, empirical evidence of its effects on pedestrian decision-making, and a learnability perspective on how brief onboarding shapes early interpretation of non-textual AV signals. While the current evaluation isolates these effects in a controlled single-pedestrian setting, it establishes a foundation for studying richer light-based communication in more realistic multi-agent traffic.

2 Methods

The study was conducted in a virtual-reality (VR) environment developed in Unity 2022.3.5f1 (see supplementary material) and presented using a Meta Quest 3 headset. Adapted from previous research on AV-pedestrian interaction by Alam et al. [1], the platform provided a controlled yet immersive setting for evaluating the eHMI across repeated roadside encounters. VR was used as it enabled systematic manipulation of AV behaviour and signalling while maintaining an interactive pedestrian perspective.

2.1 Study Design

We used a mixed experimental design to examine how the eHMI and brief onboarding influenced pedestrian responses to AV yielding intent. Onboarding was manipulated between participants, whereas the eHMI presence and vehicle yielding behaviour were manipulated within participants. This resulted in four within-participant conditions: the eHMI present versus absent, crossed with yielding versus non-yielding trials.

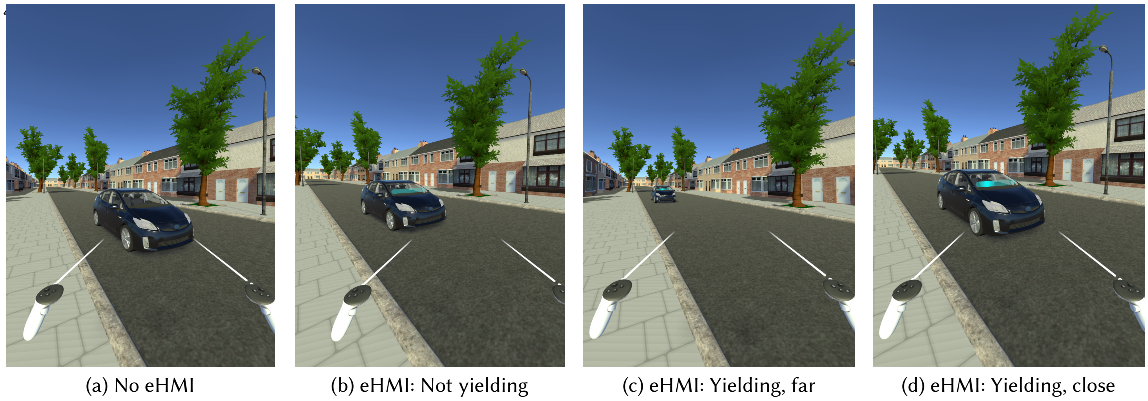
(a) No eHMI (b) eHMI: Not yielding (c) eHMI: Yielding, far (d) eHMI: Yielding, close

Fig. 1. Interface Conditions: Baseline (a), with eHMI and AV not yielding (b), with eHMI and AV yielding when far away from the pedestrian (c), with eHMI AV yielding when closer to the pedestrian on the left side of the pavement (d).

The participants completed two experimental blocks separated by a short break. The conditions were presented repeatedly in random order, allowing us to analyse both immediate responses and changes over repeated encounters. Throughout the experiment, each participant encountered each condition eight times, resulting in 32 experimental trials in total, in addition to two familiarisation trials before the experiments started. This repeated design allowed us to study both the effects of the eHMI and its learnability over time.

2.2 External Human-Machine Interface and Onboarding

We implemented a directional light-based eHMI as a cyan light strip positioned at the bottom of the AV's windscreen. When the vehicle was not yielding, they drove through at a steady speed of $50km/h$, with the strip remaining faintly illuminated across its full width (Figure 1b), indicating that the interface was active without signalling yielding intent. When the vehicle yielded, a segment of the strip brightened and shifted along the horizontal axis according to the pedestrian's relative position. As shown in Figure 1, when the car moved towards the pedestrian who stood by the pavement on the left, the light strip moved from the centre, as illustrated in Figure 1c, to the left side of the windshield (Figure 1d). From the pedestrian's perspective, this created a directional cue that allowed the interface to convey more contextual information than a coarse go/no-go signal. In the experimental implementation, the brightened segment maintained a constant width and moved in discrete steps as the vehicle approached and came to a stop.

To examine whether prior explanation influenced the interpretation of this non-textual signal, we included a brief onboarding manipulation. Participants in the onboarding condition watched a pre-recorded video of approximately three minutes before the experiment (see supplementary material). The video introduced the eHMI using slides showing images of AVs in their non-yielding and yielding states, audio explaining the meaning of these states, and illustrations of how the brightened segment moved relative to the pedestrian. Participants in the no-onboarding condition did not receive this explanation.

2.3 Participants and Measures

Thirty participants took part in the study. The sample comprised 19 males and 11 females, aged 20–60 years ($M = 27.6$, $SD = 10.0$, median = 25). Participants were recruited through the researchers' networks and a group database and received monetary compensation for their participation. The study was approved by the Ethics Review Board of [redacted]. A sensitivity analysis indicated that the sample size was sufficient to detect approximately medium within-participant effects, but only relatively large between-participant effects; onboarding-related contrasts should therefore be interpreted with appropriate caution.

We focused on one primary behavioural measure and two secondary subjective measures. The primary measure was the crossing initiation time, operationalised as the moment when participants pressed the controller button to indicate that they felt safe to cross. The participants were instructed to press the button when they believed that the vehicle was yielding to them and to keep pressing until the vehicle started moving again. For each trial, we extracted the timestamp of the first full button press, providing an action-proximal indicator of how quickly the participants committed to cross.

The two secondary measures were perceived safety and trust. After each trial, participants rated how safe they felt, and how much they trusted the vehicle during this interaction on 7-point Likert scales presented in VR. In addition, participants completed a 7-item trust questionnaire at baseline, during the break, and after the second block. Trust was assessed using an adapted version of Jian et al.'s trust in automation scale [17], with items reworded to refer to automated vehicles and to the two vehicle conditions used in the experiment (with and without eHMI).

2.4 Procedure

After giving their informed consent, the participants received a standardised explanation of the study and completed the baseline trust questionnaire. Participants in the onboarding condition then watched the pre-recorded explanation video, whereas those in the no-onboarding condition proceeded directly to the VR task. Before the main experiment, all participants completed two familiarisation trials without the eHMI: one non-yielding trial and one yielding trial.

The participants then completed two experimental blocks separated by a five-minute break. After the first block, they completed the vehicle-specific trust questionnaires, which were administered again after the second block. The full session lasted approximately 45–60 minutes.

2.5 Data Analysis

Data were pre-processed in Python and analysed using linear mixed-effects models to account for repeated observations nested within participants. For the initiation time, six observations were excluded because the participants either accidentally pressed the button or reported responding only after the vehicle had nearly left; all other observations were retained. In the main models, the eHMI presence, onboarding condition, and trial number were included as fixed effects, together with relevant interaction terms, and participant was included as a random intercept. Separate models were fitted for crossing initiation time, perceived safety, and trust. For the questionnaire-based trust measure, negatively worded items were reverse-coded and averaged to form a composite score after confirming acceptable internal consistency. Statistical significance was evaluated at $\alpha = .05$ and 80% power, the design is powered to detect an approximately medium within-participants effect and a large between-participants effect.

3 Results

3.1 Crossing Initiation Time

Figure 2 shows the crossing initiation times were lower when the eHMI was present than when no eHMI was shown, indicating earlier pedestrian decisions in the eHMI condition. This reduction was more pronounced for participants who received onboarding. In the no-onboarding group, the difference between eHMI and no-eHMI was modest, whereas in the onboarding group, the eHMI was associated with a clearer reduction in crossing initiation time.

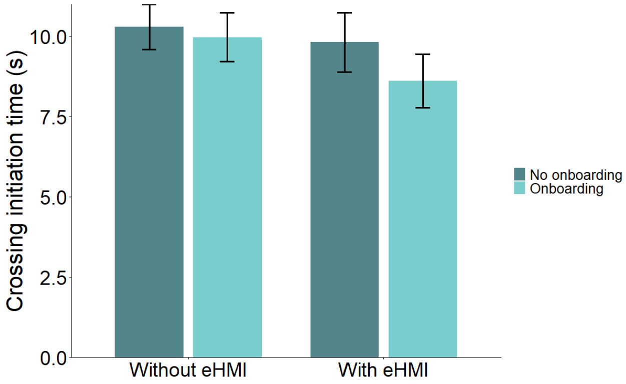

Fig. 2. Mean crossing initiation time by eHMI condition and onboarding group. Lower values indicate earlier pedestrian decisions. Error bars represent 95% confidence intervals.

Table 1. Linear mixed-effects model for crossing initiation time.

Effect	*b*	*SE*	95% CI	*z*	*p*
eHMI	-0.48	0.11	[-0.70, -0.26]	-4.28	< .001
Onboarding	-0.33	0.49	[-1.30, 0.64]	-0.67	.506
Trial number	-0.11	0.02	[-0.14, -0.07]	-6.22	< .001
eHMI × Onboarding	-0.87	0.16	[-1.18, -0.56]	-5.51	< .001

A linear mixed-effects model supported this descriptive pattern. The proposed eHMI significantly reduced crossing initiation time ($b = -0.48$, $SE = 0.11$, 95% CI $[-0.70, -0.26]$, $z = -4.28$, $p < .001$). Trial number also had a significant effect ($b = -0.11$, $SE = 0.02$, 95% CI $[-0.14, -0.07]$, $z = -6.22$, $p < .001$), indicating that participants initiated crossing earlier over repeated encounters. The main effect of onboarding was not significant ($b = -0.33$, $SE = 0.49$, 95% CI $[-1.30, 0.64]$, $z = -0.67$, $p = .506$). Importantly, there was a significant interaction between the eHMI and onboarding ($b = -0.87$, $SE = 0.16$, 95% CI $[-1.18, -0.56]$, $z = -5.51$, $p < .001$), indicating that the effect of the eHMI was stronger for participants who received prior explanation of the interface. Pairwise comparisons further clarified this interaction. In the no-onboarding group, eHMI reduced crossing initiation time by 0.48 s ($SE = 0.11$, 95% CI $[-0.70, -0.26]$, $z = -4.28$, $p < .001$). In the onboarding group, the corresponding reduction was 1.35 s ($SE = 0.11$, 95% CI $[-1.57, -1.13]$, $z = -12.12$, $p < .001$). Onboarding significantly reduced crossing initiation time only when eHMI was present ($b = -1.20$, $SE = 0.49$, 95% CI $[-2.17, -0.23]$, $z = -2.43$, $p = .015$), but not when it was absent ($b = -0.33$, $SE = 0.49$, 95% CI $[-1.30, 0.64]$, $z = -0.67$, $p = .506$). Together, these results indicate that our eHMI supported earlier pedestrian crossing decisions overall, while onboarding substantially strengthened this behavioural benefit.

3.2 Perceived Safety

Figure 3 shows that the effect of the proposed eHMI on perceived safety differed markedly between the onboarding and no-onboarding groups. In the onboarding group, the eHMI produced a clear and stable positive effect on perceived safety from the first trial onward. In contrast, in the no-onboarding group, the effect of the eHMI was initially close to zero, but increased steadily over repeated encounters. By later trials, the estimated safety benefit of eHMI in the no-onboarding

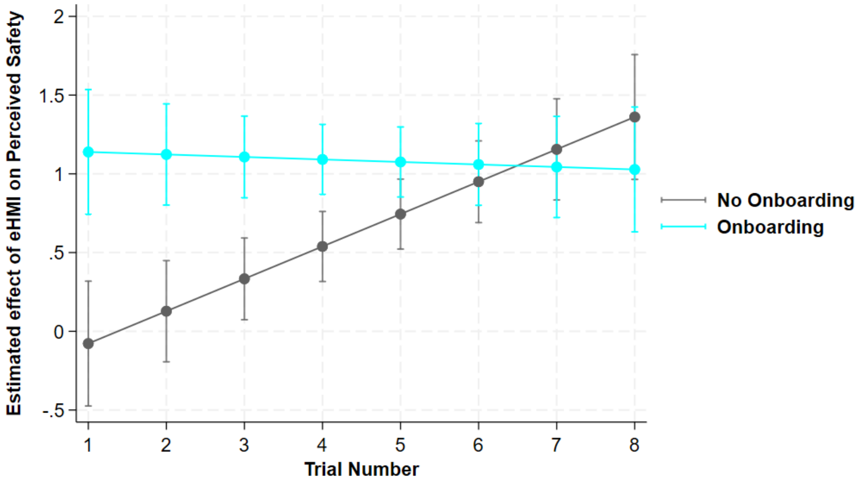

Fig. 3. Estimated marginal effect of eHMI on perceived safety across repeated trials, separately for the no-onboarding and onboarding groups. Values above zero indicate that eHMI increased perceived safety relative to the no-eHMI condition. Error bars represent 95% confidence intervals.

Table 2. Linear mixed-effects model for perceived safety.

Effect	*b*	*SE*	95% CI	*z*	*p*
eHMI	-0.28	0.24	[-0.76, 0.20]	-1.16	.246
Onboarding	-0.86	0.41	[-1.66, -0.06]	-2.10	.036
Trial number	-0.13	0.03	[-0.20, -0.07]	-3.92	< .001
eHMI × Onboarding	1.44	0.35	[0.76, 2.11]	4.17	< .001
eHMI × Trial number	0.21	0.05	[0.11, 0.30]	4.25	< .001
Onboarding × Trial number	0.12	0.05	[0.02, 0.21]	2.45	.014
eHMI × Onboarding × Trial number	-0.22	0.07	[-0.36, -0.09]	-3.24	.001

group approached and slightly exceeded that observed in the onboarding group. This pattern suggests that onboarding enabled participants to benefit from eHMI immediately, while participants without onboarding gradually learned to interpret the signal through repeated exposure.

A linear mixed-effects model supported this descriptive pattern. The main effect of the eHMI was not significant ($b = -0.28$, $SE = 0.24$, 95% CI $[-0.76, 0.20]$, $z = -1.16$, $p = .246$), whereas onboarding showed a significant main effect ($b = -0.86$, $SE = 0.41$, 95% CI $[-1.66, -0.06]$, $z = -2.10$, $p = .036$). Importantly, there was a significant interaction between the eHMI and onboarding ($b = 1.44$, $SE = 0.35$, 95% CI $[0.76, 2.11]$, $z = 4.17$, $p < .001$). Trial number also had a significant effect ($b = -0.13$, $SE = 0.03$, 95% CI $[-0.20, -0.07]$, $z = -3.92$, $p < .001$). In addition, both the eHMI × trial number interaction ($b = 0.21$, $SE = 0.05$, 95% CI $[0.11, 0.30]$, $z = 4.25$, $p < .001$) and the onboarding × trial number interaction ($b = 0.12$, $SE = 0.05$, 95% CI $[0.02, 0.21]$, $z = 2.45$, $p = .014$) were significant. These effects were further qualified by a significant three-way interaction between the eHMI, onboarding, and trial number ($b = -0.22$, $SE = 0.07$, 95% CI $[-0.36, -0.09]$, $z = -3.24$, $p = .001$).

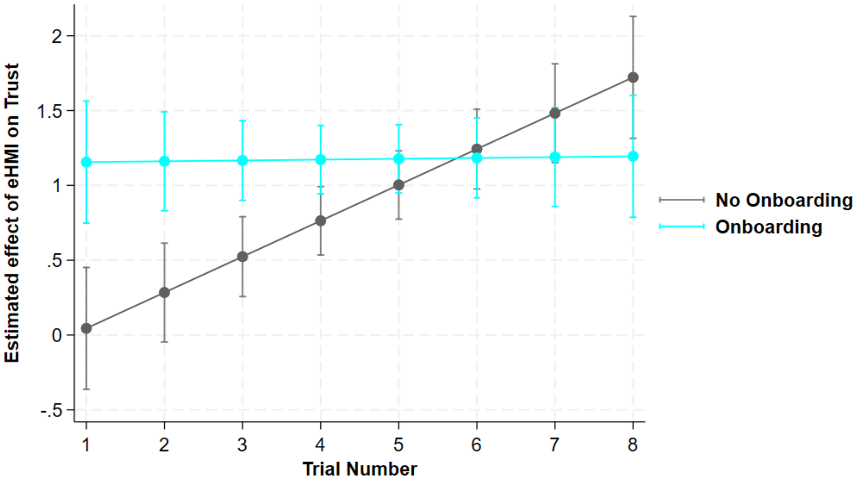

Fig. 4. Estimated marginal effect of eHMI on perceived trust in the AV across repeated trials, separately for the no-onboarding and onboarding groups. Values above zero indicate that eHMI increased perceived trust relative to the no-eHMI condition. Error bars represent 95% confidence intervals.

Pairwise comparisons clarified this interaction. In the onboarding group, eHMI significantly increased perceived safety from the first trial onward, with an estimated effect of 1.14 in Trial 1 that remained relatively stable across repeated encounters. In the no-onboarding group, eHMI did not significantly affect perceived safety in Trials 1 and 2, but became significant from Trial 3 onward and increased steadily through Trial 8. Together, these findings indicate that eHMI improved perceived safety overall, while onboarding primarily influenced how quickly this benefit emerged.

3.3 Trust

Trust was assessed in two complementary ways: a 7-item questionnaire administered at baseline, during the break, and after the second block, and a trial-level trust rating collected after each encounter. The questionnaire-based measure captures participants' more stable overall impressions of the AV, whereas the trial-level measure reflects how trust developed over repeated interactions. For the 7-item trust questionnaire, a linear mixed-effects model that included eHMI, onboarding, and survey timing showed a strong positive effect of eHMI on trust ($b = 1.58$, $SE = 0.13$, $z = 12.54$, $p < .001$). In contrast, neither onboarding ($b = 0.09$, $SE = 0.24$, $z = 0.38$, $p = .706$) nor survey timing (break vs. end of the experiments) ($b = 0.01$, $SE = 0.13$, $z = 0.08$, $p = .940$) had a significant effect. These results indicate that the eHMI substantially increased the overall trust of the participants in the AV, whereas repeated exposure and prior explanation did not significantly alter these impressions of more global trust. A different pattern emerged for trial-level trust ratings. Figure 4 shows that, in the onboarding group, eHMI produced a clear and stable positive effect on trust from the first trial onward. In contrast, in the no-onboarding group, the effect of eHMI was initially small and non-significant, but increased steadily over repeated encounters. From Trial 3 onwards, eHMI significantly increased trust in the no-onboarding group, and by later trials the estimated trust benefit approached and exceeded that observed in the onboarding group. This pattern suggests that onboarding accelerated and stabilised the trust benefit of eHMI, while participants without onboarding required repeated exposure before eHMI significantly increased trust.

Table 3. Linear mixed-effects model for the 7-item trust questionnaire.

Effect	b	SE	z	p
eHMI	1.58	0.13	12.54	$< .001$
Onboarding	0.09	0.24	0.38	.706
Survey timing	0.01	0.13	0.08	.940

A linear mixed-effects model supported this descriptive pattern. The main effect of eHMI was not significant ($b = -0.20$, $SE = 0.25$, $z = -0.78$, $p = .437$), while the registration showed a significant negative main effect ($b = -1.10$, $SE = 0.44$, $z = -2.49$, $p = .013$). Importantly, there was a significant interaction between the eHMI and onboarding ($b = 1.35$, $SE = 0.35$, $z = 3.79$, $p < .001$). Trial number also had a significant effect ($b = -0.14$, $SE = 0.04$, $z = -4.11$, $p < .001$). In addition, both the eHMI × trial number interaction ($b = 0.24$, $SE = 0.05$, $z = 4.82$, $p < .001$) and the onboarding × trial number interaction ($b = 0.15$, $SE = 0.05$, $z = 2.95$, $p = .003$) were significant, and these effects were qualified by a significant three-way interaction between eHMI, onboarding, and trial number ($b = -0.23$, $SE = 0.07$, $z = -3.33$, $p = .001$).

Pairwise comparisons further clarified this interaction. In the onboarding group, eHMI significantly increased trust from the first trial onward, with an effect of 1.16 in Trial 1 that remained relatively stable across repeated encounters. In the no-onboarding group, eHMI did not significantly affect trust in Trials 1 and 2, but became significant from Trial 3 onward and increased steadily through Trial 8. Together, these findings indicate that the eHMI increased trust overall, while onboarding primarily influenced how quickly this benefit emerged.

4 Discussion

This study examined how an informative light-based eHMI, and a brief onboarding intervention influenced pedestrian crossing behaviour, perceived safety, and trust during interactions with a simulated SAE Level 5 AV. Overall, the findings suggest that our eHMI can support earlier pedestrian decisions and improve subjective interaction quality, while onboarding primarily affects how quickly these benefits emerge. Together, these results contribute to a more nuanced understanding of how pedestrians interpret non-textual AV signals and how such interpretations develop over repeated encounters.

4.1 Effects of the eHMI on pedestrian decision-making and subjective evaluations

The clearest behavioural effect of the eHMI was on crossing initiation time. Participants initiated crossing earlier when the eHMI was present than when no eHMI was shown, indicating that the signal helped pedestrians decide more quickly that the vehicle was yielding. This finding is consistent with prior work showing that light-based eHMIs can facilitate pedestrian crossing decisions and willingness to cross [5, 7, 19]. In the context of the present paper, this result is particularly important because it supports the broader design goal of moving beyond coarse go/no-go light signals toward a minimal but more informative interface. Rather than merely signalling that the AV is yielding, our eHMI appears to provide sufficient directional information to reduce ambiguity and support earlier action.

The eHMI also positively affected perceived safety. However, unlike crossing initiation time, this effect was not equally immediate across groups. For participants who received onboarding, the eHMI produced a stable safety benefit from the first trial onward. For participants without onboarding, the safety benefit emerged only after repeated encounters. This pattern suggests that the meaning of the signal was not fully transparent at first encounter, but became interpretable with experience. Such a finding is consistent with prior work indicating that light-based eHMIs may be less immediately intuitive than more explicit alternatives, while still being learnable over repeated interactions [5, 13]. In other words, our eHMI appears effective, but its effectiveness depends partly on familiarity. A similar pattern was observed for trust. The 7-item trust questionnaire showed a strong overall positive effect of the eHMI, suggesting that the presence of the interface substantially improved participants' broader impressions of the AV.

By contrast, the trial-level trust measure revealed a more dynamic pattern: in the onboarding group, the eHMI increased trust from the first encounter, whereas in the no-onboarding group this benefit became significant only after repeated exposure. This difference between the global questionnaire and the trial-level ratings is informative rather than problematic. It suggests that the questionnaire captured a more stable, learned evaluation of the vehicle, whereas the trial-level ratings were more sensitive to the unfolding interpretation of the signal during individual encounters. This interpretation aligns with prior distinctions between more stable trust in automation and more situational, interaction-dependent trust [3, 14].

Overall, these findings suggest that the eHMI had a positive effect across behavioural and subjective outcomes, but that these outcomes differed in how quickly the benefit emerged. Crossing behaviour changed relatively quickly, whereas perceived safety and trust more clearly reflected the gradual process by which pedestrians learned to interpret the interface. This strengthens the argument that evaluating eHMIs should not rely on a single outcome alone: behavioural measures may show whether pedestrians act differently, while subjective measures help explain how that behavioural change is experienced.

4.2 Onboarding as a learnability intervention

Beyond the effects of the eHMI itself, the results highlight the importance of learnability for non-textual eHMIs. Onboarding did not produce a broad stand-alone benefit across all outcomes. Instead, its role was more specific: it accelerated and stabilised the positive effects of the eHMI. Participants who received onboarding crossed earlier when eHMI was present, and they also showed immediate gains in perceived safety and trial-level trust. When eHMI was absent, onboarding did not improve crossing behaviour. This pattern is theoretically important because it suggests that onboarding should not be understood as an independent driver of pedestrian confidence, but rather as a mechanism that helps pedestrians acquire the meaning of a novel signal more quickly. Prior work has shown that familiarity and repeated exposure shape how pedestrians interpret eHMIs, influencing crossing behaviour, trust development, and even misuse [4, 18, 19]. The present findings extend this literature by showing that even a brief explanatory intervention can shape the early phase of interaction with a minimal light-based eHMI. At the same time, the results also show that onboarding was not strictly necessary for eHMI to become useful: participants without onboarding eventually showed similar improvements in perceived safety and trust after only a few encounters. This implies that the interface was learnable through interaction alone, but that onboarding reduced the initial ambiguity and accelerated the transition from uncertainty to confident interpretation. One notable trust-related finding was that onboarding did not significantly affect the 7-item trust questionnaire, even though it influenced trial-level trust in early encounters. This reinforces the view that onboarding primarily affects short-term interpretability rather than stable overall impressions. It also suggests that, for non-textual eHMIs, learnability may be especially important during first exposure, when pedestrians have not yet formed a working understanding of the signal.

4.3 Implications for eHMI design

From a design perspective, the findings support the broader claim that a light-based eHMI does not need to rely on text in order to be useful, but that minimal designs must still communicate enough information to reduce ambiguity. Our eHMI appears to move in that direction by combining a small visual footprint with richer directional information. This is important because much of the existing design space reflects a tension between explicitness and deployability: text-based eHMIs can be clearer, but raise concerns regarding language dependence, legibility, and scalability, whereas light-based eHMIs are easier to deploy but often communicate only coarse intent. The present results suggest that this tradeoff may not be fixed. Instead, a minimal light-based design can still support meaningful behavioural and subjective benefits if it conveys more contextual information and if pedestrians are given either a brief explanation or repeated exposure. At the same time, the present study does not yet show that our eHMI is superior to other dynamic light-based eHMIs in more complex settings. What it does show is that the design is promising in a controlled single-pedestrian interaction and that its benefits are especially visible when interpreted as a learnable signal rather than as a universally intuitive one.

4.4 Limitations and future work

Several limitations should be considered when interpreting these findings. First, the study was conducted in a controlled VR environment featuring a single pedestrian and a single AV on a straight road without additional traffic or competing cues. This level of control was useful for isolating the effects of the eHMI and onboarding, but it limits ecological validity. Real traffic environments often involve multiple vehicles, multiple pedestrians, and competing sources of information. Future work should therefore examine whether the observed benefits persist in more complex multi-agent environments. Second, the current study cannot fully evaluate one of eHMI's most interesting design premises, namely that directional information may be particularly valuable when multiple pedestrians are present. Because only one pedestrian was included at a time, the study cannot show whether our eHMI performs better than other dynamic light-based eHMIs in situations where the AV's yielding target is genuinely ambiguous. This should be a central focus of future research. Third, the sample size was relatively small ($N = 30$), and the between-participant onboarding manipulation resulted in two groups of only 15 participants each. As noted earlier, this limits sensitivity to between-group effects unless they are relatively large. The fact that several onboarding-related effects nevertheless emerged suggests that onboarding may play a meaningful role, but larger studies are needed to assess its effects more robustly. Fourth, the vehicle kinematics were intentionally held constant across conditions: whenever the AV yielded, its approach speed, braking onset, deceleration, and stopping position were identical. This consistency was important for experimental control, but it may also have allowed participants to anticipate the vehicle's behaviour over time. As a result, some of the observed changes across repeated encounters may reflect not only learning of the eHMI, but also increasing familiarity with the vehicle's motion pattern. Future studies should therefore introduce greater variability in vehicle behaviour to better approximate naturalistic AV–pedestrian interactions. Finally, the present sample consisted primarily of participants familiar with Dutch traffic environments, which may limit the generalisability of the findings to other traffic cultures and pedestrian norms. Since eHMI interpretation is likely shaped by learned expectations about road-user communication, cross-cultural validation would be a valuable next step.

5 Conclusion

In summary, the present study suggests that our eHMI can support earlier pedestrian crossing decisions and improve perceived safety and trust during AV–pedestrian interaction. These benefits were strongest immediately when pedestrians received onboarding, but they also emerged through repeated exposure in the no-onboarding group. The findings therefore support two broader conclusions: first, that minimal light-based eHMIs can be made more informative without abandoning deployability; and second, that the value of such interfaces depends not only on their design, but also on how quickly pedestrians can learn to interpret them.

6 Supplementary Material

Following open science practices and transparency recommendations in automotive user research [8], the authors openly share study artefacts to support reproducibility and collaboration. The simulator and analysis code, as well as the materials used in the experiment, are available at https://www.dropbox.com/scl/fo/wp4hsxv4jknzmpnos7q9c/AD2qe_I69lrc0Qje09nyg3A?rlkey=zrne0tckku91g6q9xy3o86azv.

References

[1] Md Shadab Alam, Debargha Dey, Marieke H. Martens, and Pavlo Bazilinskyy. 2026. You'll never walk alone: Inter-pedestrian distance, eHMIs, and crossing decisions in virtual reality. *Under Review* (2026).

[2] Pavlo Bazilinskyy, Dimitra Dodou, and Joost De Winter. 2019. Survey on eHMI concepts: The effect of text, color, and perspective. *Transportation research part F: traffic psychology and behaviour* 67 (2019), 175–194.

[3] Thierry Bellet, Sébastien Laurent, Jean-Charles Bornard, Isabelle Hoang, and Bertrand Richard. 2022. Interaction between pedestrians and automated vehicles: Perceived safety of yielding behaviors and benefits of an external human–machine interface for elderly people. *Frontiers in psychology* 13 (2022), 1021656.

[4] 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*. 1–11.

[5] Koen De Clercq, Andre Dietrich, Juan Pablo Núñez Velasco, Joost De Winter, and Riender Happee. 2019. External human-machine interfaces on automated vehicles: Effects on pedestrian crossing decisions. *Human factors* 61, 8 (2019), 1353–1370.

[6] 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 (2020), 100174.

[7] Debargha Dey, Kai Holländer, Melanie Berger, Berry Eggen, Marieke Martens, Bastian Pfleging, and Jacques Terken. 2020. Distance-dependent eHMIs for the interaction between automated vehicles and pedestrians. In *12th international conference on automotive user interfaces and interactive vehicular applications*. 192–204.

[8] Patrick Ebel, Pavlo Bazilinskyy, Mark Colley, Courtney Michael Goodridge, Philipp Hock, Christian P. Janssen, Hauke Sandhaus, Aravinda Ramakrishnan Srinivasan, and Philipp Wintersberger. 2024. Changing lanes toward open science: Openness and transparency in automotive user research. In *Proceedings of the 16th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Stanford, CA, USA) *(AutomotiveUI '24)*. Association for Computing Machinery, New York, NY, USA, 94–105. doi:10.1145/3640792.3675730

[9] Yke Bauke Eisma, Anna Reiff, Lars Kooijman, Dimitra Dodou, and Joost CF de Winter. 2021. External human-machine interfaces: Effects of message perspective. *Transportation research part F: traffic psychology and behaviour* 78 (2021), 30–41.

[10] Yke Bauke Eisma, Steven van Bergen, SM Ter Brake, MTT Hensen, Willem Jan Tempelaar, and Joost CF de Winter. 2019. External human–machine interfaces: The effect of display location on crossing intentions and eye movements. *Information* 11, 1 (2019), 13.

[11] Jiawen Guo, Quan Yuan, Jingrui Yu, Xizheng Chen, Wenlin Yu, Qian Cheng, Wuhong Wang, Wenhui Luo, and Xiaobei Jiang. 2022. External human–machine interfaces for autonomous vehicles from pedestrians' perspective: A survey study. *Sensors* 22, 9 (2022), 3339.

[12] Azra Habibovic, Victor Malmsten Lundgren, Jonas Andersson, Maria Klingegård, Tomas Lagström, Anna Sirkka, Jonas Folkesson, Anders Kullgren, Annika Larsson, and David Saluäär. 2018. Communicating intent of automated vehicles to pedestrians. *Frontiers in Psychology* 9 (2018), 1336. doi:10.3389/fpsyg.2018.01336

[13] Ann-Christin Hensch, Isabel Neumann, Matthias Beggiato, Josephine Halama, and Josef F Krems. 2019. Effects of a light-based communication approach as an external HMI for Automated Vehicles–a Wizard-of-Oz Study. *Transactions on Transport Sciences* 10, 2 (2019).

[14] Kevin Anthony Hoff and Masooda Bashir. 2015. Trust in automation: Integrating empirical evidence on factors that influence trust. *Human factors* 57, 3 (2015), 407–434.

[15] Kai Holländer, Mark Colley, Christian Mai, and Enrico Rukzio. 2019. Investigating the influence of external communication of autonomous vehicles on pedestrian crossing behavior and trust. In *Proceedings of the 11th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. 292–301. doi:10.1145/3342197.3344516

[16] 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 (2019), 117. doi:10.3389/frobt.2019.00117

[17] Jiun-Yin Jian, Ann M Bisantz, and Colin G Drury. 2000. Foundations for an empirically determined scale of trust in automated systems. *International journal of cognitive ergonomics* 4, 1 (2000), 53–71.

[18] Anees Ahamed Kaleefathullah, Natasha Merat, Yee Mun Lee, Yke Bauke Eisma, Ruth Madigan, Jorge Garcia, and Joost de Winter. 2022. External human–machine interfaces can be misleading: An examination of trust development and misuse in a CAVE-based pedestrian simulation environment. *Human factors* 64, 6 (2022), 1070–1085.

[19] Yee Mun Lee, Ruth Madigan, Chinebuli Uzondu, Jorge Garcia, Richard Romano, Gustav Markkula, and Natasha Merat. 2022. Learning to interpret novel eHMI: The effect of vehicle kinematics and eHMI familiarity on pedestrian'crossing behavior. *Journal of safety research* 80 (2022), 270–280.

[20] Yee Mun Lee, Vladislav Sidorov, Ruth Madigan, Jorge Garcia de Pedro, Gustav Markkula, and Natasha Merat. 2025. Hello, is it me you're stopping for? The effect of external human machine interface familiarity on pedestrians' crossing behaviour in an ambiguous situation. *Human Factors* 67, 3 (2025), 264–279.

[21] Vishal Onkhar, Pavlo Bazilinskyy, Dimitra Dodou, and JCF De Winter. 2022. The effect of drivers' eye contact on pedestrians' perceived safety. *Transportation research part F: traffic psychology and behaviour* 84 (2022), 194–210.

[22] Amir Rasouli, Iuliia Kotseruba, and John K Tsotsos. 2017. Agreeing to cross: How drivers and pedestrians communicate. In *2017 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 264–269.

[23] Alexandros Rouchitsas and Håkan Alm. 2019. External human–machine interfaces for autonomous vehicle-to-pedestrian communication: A review of empirical work. *Frontiers in psychology* 10 (2019), 2757. doi:10.3389/fpsyg.2019.02757

[24] Aïsha Sahaï, Elodie Labeye, Loïc Caroux, and Céline Lemercier. 2022. Crossing the street in front of an autonomous vehicle: An investigation of eye contact between drivengers and vulnerable road users. *Frontiers in psychology* 13 (2022), 981666.

[25] Ernst Tomasch, Bernhard Kirschbaum, and Wolfgang Schubert. 2025. Assessment of the potential of a front brake light to prevent crashes and mitigate the consequences of crashes at junctions. *Vehicles* 7, 2 (2025), 40.

[26] World Health Organization. 2023. Road traffic injuries. https://www.who.int/news-room/fact-sheets/detail/road-traffic-injuries. Fact sheet, published 13 December 2023, accessed 7 April 2026.