Exploring the Relationship between False Alarms and Driver Acceptance of a Pedestrian Alert System during Simulated Driving

David R Large[1], Catherine Harvey[1], Gary Burnett[1]
Coleman Merenda[2], Samantha Leong[2], Joseph Gabbard[2]
[1]Human Factors Research Group, University of Nottingham, Nottingham. UK. {david.r.large; catherine.harvey; gary.burnett}@nottingham.ac.uk
[2]COGENT Lab, Virginia Tech University, Blacksburg. USA. {cjm120; sam94; jgabbard}@vt.edu.

Abstract
In-vehicle pedestrian-alert-systems (PASs) can be prone to 'false positive' declarations, with the likelihood of false interventions increasing as time-to-collision (TTC) extends. A high number of false alarms can annoy drivers and lead to poor acceptance and low trust in the technology. To explore this relationship, 24 experienced drivers negotiated a single-lane urban high-street – moderately populated with pedestrians – during 12 five-minute drives in a medium-fidelity driving simulator. PAS warnings were presented in response to pedestrians who approached the roadside, either as a static visual alert icon presented on a HUD and/or auditory icon. The number of accurately detected pedestrians (i.e. those who entered the roadway rather than waiting at the kerbside) decreased with increasing TTC, giving rise to 'false positive alarms'. Subjectively, participants associated the highest level of trust, confidence and desirability, and lower levels of annoyance, when warnings were presented at intermediate TTCs (3.0 and 4.0-seconds, corresponding to false-alarm rates of 40% and 60%, respectively); trust and confidence reduced significantly with both increasing and decreasing TTC. Driving performance data show that earlier warnings encouraged drivers to begin braking sooner and apply braking force more gradually, ultimately stopping further from the pedestrian – on average 18.0m following 5.0-second warnings compared to 6.2m with 2.0-second warnings. Nevertheless, evidence suggests that some drivers may have disregarded the system at longer TTCs, choosing to rely on their own judgement. The results have implications for the design, evaluation and acceptance of PASs.

Keywords
False Alarms; Driver Acceptance; Pedestrian Alert Systems; Vulnerable Road Users; Simulation.

1. Introduction

Accidents involving vulnerable road users (VRUs) remain a major issue for road safety, accounting for almost 40% of road fatalities in Europe, and almost 50% worldwide [1]. Pedestrians are one of the most vulnerable road user groups, both in terms of the likelihood of being involved in a near-miss or collision, and the potential ramifications should an incident occur – pedestrians represent 11-13% of those killed in collisions and are 1.5 times more at risk than vehicle passengers to be fatally injured [2, 3]. Analysis of accident data shows that a high percentage of pedestrian-related incidents occur in urban areas [4], where the density of pedestrians is likely to be higher, with many of these attributed to unexpected pedestrian behaviour. For example, in the UK in 2013 over 80% of pedestrian road-related fatalities occurred outside of demarcated pedestrian crossings [5]. Data indicate that the most common accident scenario involves adult pedestrians (15-64 years old) and occurs in situations when the car moved forward and the pedestrian crossed from the right (34.7%) or left (21.4%), and in particular where there were no obstacles obstructing drivers' vision [6]. Data also indicate that 85% of accidents involving pedestrians occur at vehicle speeds of 40km/h or below [6].

Automated in-car, pedestrian alert systems (PASs) have potential to mitigate the risk to pedestrians by warning drivers if a pedestrian is present within their vehicle's path and/or a collision is likely based on the current speed and trajectory of each party. However, the success of such technology is predicated on its ability to accurately detect pedestrians – and make a precise estimation of their current and future positions with respect to the moving vehicle – while providing drivers with the capacity (time/space) to respond to warnings. These are evidently conflicting goals: by warning drivers sooner, they have more time and space to respond and take evasive action (e.g. reduce speed/brake), but earlier warnings are likely to be offered at the expense of accuracy, due to limitations in sensors (range and capability). Current state-of-the-art systems make varying claims, e.g. the ability to detect pedestrians within 3 seconds of the time to collision (TTC) [7] or within an 80m range of the host vehicle [8].

The performance of PAS is also hindered by the highly dynamic behaviour of pedestrians, who can change their walking direction/speed, or abruptly start or stop moving at a moment's notice – e.g. a pedestrian may approach the road, apparently intending to cross, but then stops at the roadside kerb. This could result in a false declaration if the PAS has already identified the pedestrian as a potential hazard and highlighted their presence to the driver. Consequently, even state-of-the-art systems can be prone to declaring high numbers of false detections, with performance data suggesting that current systems miss approximately 20-30% of all pedestrians (‘false negatives’), and issue approximately one false alarm (‘false positives’), every ten images [9] (e.g. the system mistakenly detects a pedestrian that is not actually present, or falsely predicts their intention to enter the host vehicle’s trajectory). Developments in technology and detection algorithms are likely to assist in improving the capability, accuracy and range of systems, but due to the dynamic nature of pedestrians, even a perfect system may still ‘fail’ (i.e. issue a false alarm) if deployed too soon (e.g. a pedestrian walks towards the roadside, clearly intending to cross the road – and is consequently highlighted by the PAS – but then stops).

A high incidence of false alarms is likely to annoy drivers. This will influence their willingness to accept genuine alerts and may result in poor overall acceptance of the technology [10, 11]. For example, a PAS that alarms drivers every few seconds as drivers pass through an urban road network – with adjoining pathways densely populated with pedestrians – is likely to annoy drivers and may encourage them to neglect the system, find creative ways to bypass it, or deactivate the system completely [12]. It has therefore been suggested that drivers are not provided with the capacity to ‘switch off’ active safety systems but are instead given the ability to modify the sensitivity thresholds, auditory intensity and visual luminance, thereby maintaining a sense of control over the system [13].

Driver Acceptance

The ‘acceptance’ of in-vehicle driver support systems can be considered as, “the degree to which an individual incorporates the system in his/her driving, or if the system is not available, intends to use it” [14]. The determinants of drivers’ acceptance are complex and derive from various factors. These include how individuals use the technology, their understanding of its limits and the context in which it is implemented. Acceptance can also be affected by factors such as the level of trust, confidence, intrusion, or annoyance that users associate with the technology. For a hazard warning alert system, acceptance is considered to be closely related to the driver’s perception of risk – not necessarily the actual risk of collision [15]. Drivers are therefore more likely to accept alerts presented in situations that they perceive to be ‘alarming’ (i.e. they are aware of the hazard and/or potential risk of collision), even if they would have been able to avoid the incident without the warning [16].

A common ground in acceptance research is the fact that human behaviour is not primarily determined by objective factors, but by subjective perceptions [17]. This means that acceptance is based on individual attitudes, expectations and experience (including factors such as trust, confidence, annoyance, etc.), as well as their subjective evaluation of expected benefits [18]. In fact, it has been shown that the degree of technological innovation has a lesser effect on the acceptance and use of new systems than the personal importance for users [19].

For a pedestrian alert system (assuming technical competence), understanding when best to present alarms (i.e. at what TTC) appears key to ensuring driver acceptance – if alarms are presented too early, accuracy may be compromised or drivers may be unaware of the potential risk. Conversely, alarms that are ‘too late’ are of limited practical use and can in fact be detrimental to the braking process [20]; late alarms may also annoy drivers if they are already aware of the hazard and instigated an evasive manoeuvre. Nevertheless, systems must still accurately detect genuine situations. The concern is that by maximising detection rates, drivers are likely to be flooded with false-alarm warnings, with the inevitable consequence that drivers will miss or ignore genuine alerts in safety-critical situations (i.e. the system ‘cries wolf’) [21]. In contrast, systems designed to minimise false alarms may miss genuine safety-critical situations. Moreover, in the absence of false alarms, genuine alerts may be so rare as to be utterly unfamiliar and consequently drivers’ reactions will be unpredictable. It is therefore suggested that false alarms have pragmatic and genuine utility in design, and rather than aiming to eliminate false alarms, systems should provide an *appropriate level* of false alarms – this will also eliminate an ‘irony of automation’ (i.e. efficient recall of how to react depends on the frequency of use [22]) – and also helps to ensure that drivers are able to calibrate their trust in system [23]. It is therefore important to understand the relationship between driver acceptance and the safety implications of different activation thresholds, and base the design of future systems on driver expectations for when the system should issue an alert (i.e. design for acceptance of false alarms).

Overview and Aims of Study

In order to explore the relationship between false alarms and driver acceptance of a PAS, participants were asked to negotiate a single lane urban high-street (with pathways moderately populated with pedestrians) on multiple occasions using a medium-fidelity driving simulator. On five occasions during each drive (presented at random throughout the scenario), a pedestrian approached the roadside, as if intending to cross the road. Some of the pedestrians then continued to cross the road while others remained at the roadside – all five pedestrians were identified by the PAS system, thereby giving rise to ‘false positive alarms’ in the latter situation (i.e. the system accurately detected the presence of a pedestrian but falsely predicted their intention to enter the roadway). This scenario was chosen to correspond with current accident data which show that the most prominent crash types involving pedestrians occur in an urban setting where the pedestrian crosses a straight road (from either the nearside or off-side) outside of demarcated pedestrian crossings [5], is not obstructed from view, and the vehicle is travelling at speeds of 30-50km/h [24].

Warnings were presented at different times-to-collision (TTCs), with the number of *accurately* detected pedestrians (i.e. those who continued to cross the road) decreasing with increasing TTC, thus simulating the increasing false alarm rate associated with the provision of earlier warnings [24]. The approach therefore assumes that the technology performed correctly – in so far as it accurately detected pedestrians (and predicted that they were at risk based on their current speed/trajectory), in all situations, i.e. there were no 'false negative' alarms whereby the system mistakenly detects a pedestrian that is not actually present. Consequently, any false alarms were due to changes in the behaviour of pedestrians (which might be reasonably expected in a real-world system, particularly at elevated TTCs), rather than any technical limitations of the system. The aim of the study was therefore to explore the isolated effects of 'false positive' interventions only, i.e. situations when the system falsely predicts a pedestrian's intention to enter the host vehicle's trajectory.

2. Method

Twenty-four people took part in the study: 15 male, 9 female, with ages ranging from 19 to 55 years. All participants held a valid driving licence and were experienced and active drivers (mean number of years with licence, 10.7). Participants were self-selecting volunteers who responded to advertisements placed around the University of Nottingham campus and were reimbursed with £15 (GBP) of shopping vouchers as compensation for their time. All participants provided written informed consent.

The study took place in a medium-fidelity, fixed-based driving simulator at the University of Nottingham (Figure 1). The simulator comprises an Audi TT car located within a curved screen, affording a 270 degrees forward and side image of the driving scene via three overhead HD projectors, together with rear and side mirror displays. A Thrustmaster 500RS force feedback steering wheel and pedal set are integrated with the existing Audi steering wheel and pedals. The dashboard is created using a bespoke Java application and presented on a 7-inch LCD screen, which replaces the original instrument cluster.

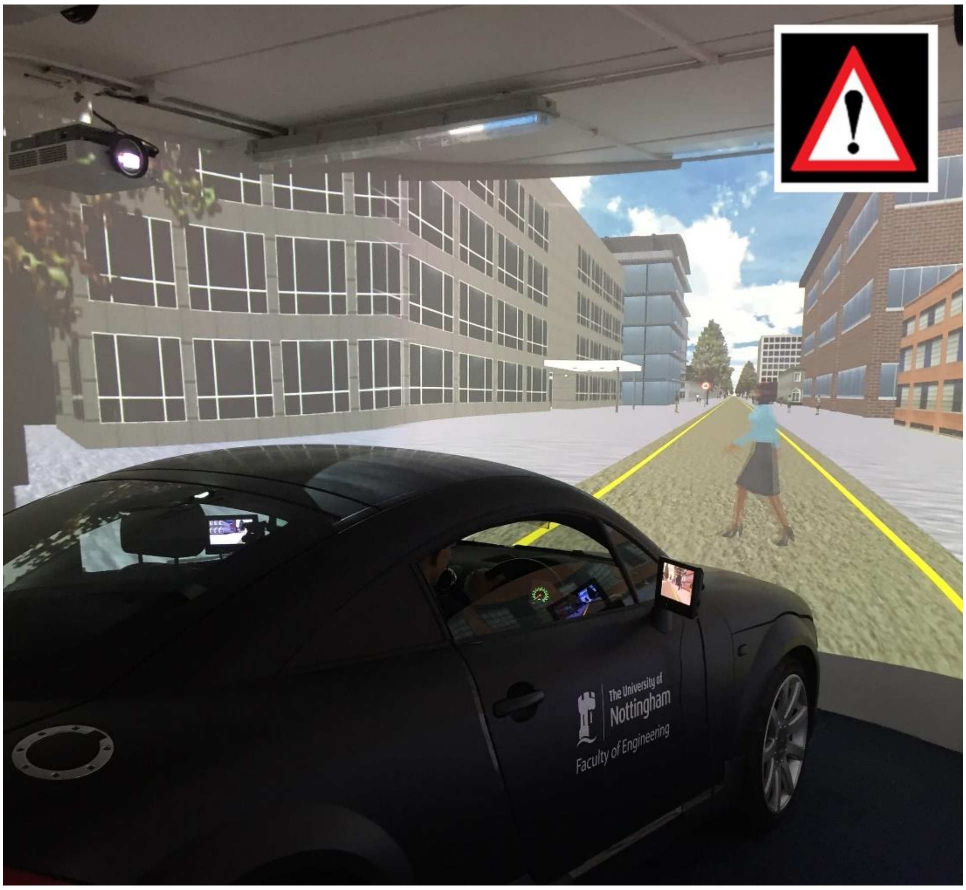

Figure 1: Driving simulator showing urban high-street scenario and pedestrian crossing road ahead, with visual alert icon that was presented on the HUD (inset).

The driving scenario was created using STISIM Drive (version 3) to replicate an urban 'high-street' scenario, with pavements ('side-walks') moderately populated with walking pedestrians. Participants were asked to drive along the high street and aim to achieve a constant speed of 25mph (40kph). No specific instruction were provided regarding how drivers should respond to warnings and/or the behaviour of pedestrians.

PAS warnings were presented within the vehicle as a static visual alert icon and/auditory icon presented at 75dB (based on design guidelines [25]). To add a further novel element, visual warnings were provided using a head-up display (HUD), and both visual and auditory warnings were spatially congruent (i.e. corresponded with the side of the road from which the pedestrian approached). A generic high-contrast alert icon was selected to aid ease of recognition and saliency to visual attention (Figure 1). The icon was presented for a period of 4.0-seconds, based on design recommendations [26].

Warnings were presented at different TTCs, with the number of accurately detected pedestrians (i.e. those who continued to cross the road) decreasing with increasing TTC, thus simulating the increasing false alarm rate associated with the provision of earlier warnings [24]. The TTCs (and associated false alarm rates) were selected based on a compromise between external validity (i.e. to correspond with the capability of existing detection technology/algorithms [24]) and internal validity (i.e. to ensure than an array of TTCs to which drivers were able to respond, were examined). TTCs therefore ranged from 5.0-seconds (with a false alarm rate of 80%, i.e. 4 out of 5 pedestrians did not cross the road) to 2.0-seconds (with a false alarm rate of 20% or 1 in 5). The locations of the pedestrians who approached the roadside, and those who ultimately entered the roadway, were selected at random throughout each scenario. For comparison, participants also repeated each drive with no warnings. Thus, participants competed 12 drives (Table 1). Conditions were counterbalanced, but participants completed all conditions for each warning type together.

After each scenario (corresponding to a different warning type and time, i.e. TTC), participants rated their attitudes towards the PAS using bespoke Likert-type scales (from 1 to 10, where 1 = ‘not at all’ and 10 = ‘completely’) exploring the constructs of trust, confidence, annoyance, desirability. Scales were constructed based on previous usability assessment recommendations [27], with ratings captured in response to specific questions following each drive. For example, ‘confidence’ was determined by inviting responses to the question: *“How confident are you that the system will be able to cope with all situations in the future?”*, and ‘desirability’ with the question *“How likely would you be to use the system if it was available in your own car?”*

Driving performance data were captured from the simulator software. These were analysed to extract the headway distance to the pedestrian when drivers responded to the warning/hazard, i.e. when they lifted their foot from the accelerator, began braking (applied foot to brake) and brought the car to a stop (or reached minimum speed).

Table 1: Twelve scenarios completed by all participants, showing false alarm rates (i.e. number of pedestrians who did not step into the road but were nevertheless highlighted by the PAS) associated with each TTC

TTC	Audio Only	Combined (HUD/audio)	None
2.0-sec	20% (1 in 5)	20% (1 in 5)	20% (1 in 5)
3.0-sec	40% (2 in 5)	40% (2 in 5)	40% (2 in 5)
4.0-sec	60% (3 in 5)	60% (3 in 5)	60% (3 in 5)
5.0-sec	80% (4 in 5)	80% (4 in 5)	80% (4 in 5)

3. Results and Analysis

Subjective Ratings

Paired-samples t-tests revealed that subjective ratings were comparable across most scales for the combined HUD/audio and audio-only warnings. There was, however, a significant difference in ratings of Confidence (t(23) = 2.14, p = .043), with participants indicating greater confidence in combined HUD/audio warnings (M=6.65, SD=1.85), compared to audio only (M=6.20, SD=1.83). Nevertheless, the primary aim of the investigation was to explore differences associated with different TTCs. Thus, data were combined for both warning types (Audio and HUD/audio), and compared across the different TTCs (2.0, 3.0, 4.0 and 5.0-seconds) using a repeated-measures ANOVA.

Trust

There was a significant difference in ratings of Trust (*“Overall how much do you trust the system?”*) associated with different warning times ($F(3,69) = 5.14$, $p = .003$). Pairwise comparisons showing that drivers placed significantly less trust in 2.0-second warnings compared to 3.0-seconds and 4.0-seconds (p = .017 and .003, respectively). Warnings provided at 5.0-seconds were rated as significantly less trustworthy than 4.0-second warnings but tended to be rated as more trustworthy than 2.0-second warnings (p = .095) (Figure 2).

Confidence

There was a significant difference in ratings of Confidence (*"How confident are you that the system will be able to cope with all situations in the future?*") ($F(3,69) = 4.05$, $p = .010$), with pairwise comparisons revealing that drivers placed less confidence in 2.0-second warnings, compared to 3.0-seconds and 4.0-seconds (*p* = .032 and .005, respectively). There was also a trend for participants to place less confidence in 5.0-second warnings compared to 4.0-second warnings ($p = .092$), although these tended to inspire higher ratings of confidence than 2.0-second warnings ($p = .101$) (Figure 2).

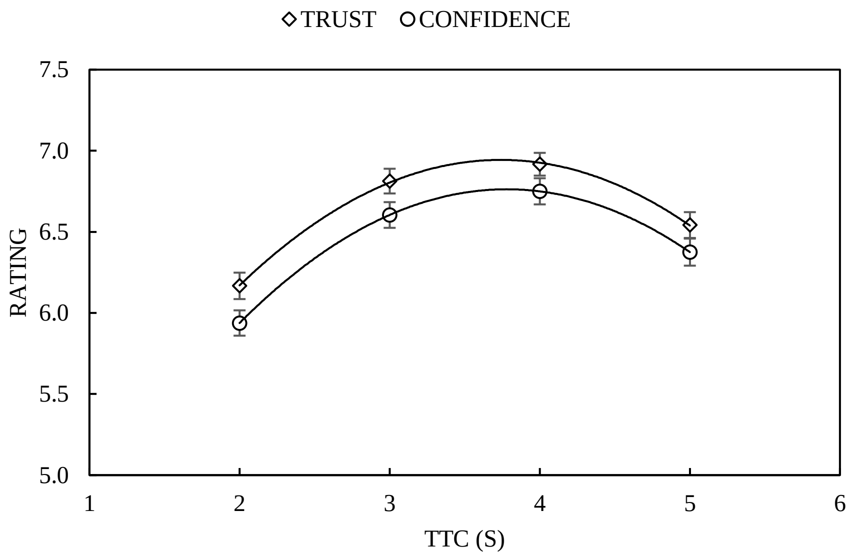

Figure 2: Subjective ratings of Trust and Confidence, with standard error bars

Annoyance

There was no significant difference overall in ratings of Annoyance ("*How annoying was the system?*") as TTC varied ($F(3,69) = .69$, $p = .564$), although there was a tendency for 5.0-second warnings to be rated as more annoying than 4.0-second warnings ($p = .136$) (Figure 3).

Desirability

There was a significant difference in Desirability (*"How likely would you be to use the system if it was available in your own car?"*) ($F(3,69) = 2.75$, $p = .049$). Pairwise comparisons revealed significant differences between 2.0-second and 4.0-second warnings ($p = .050$) – participants indicated that they would be most likely to use the system based on their experience of 4.0-second warnings. There was also a tendency for participants to rate 3.0-second warnings as more desirable than 2.0-second warnings ($p = .070$) (Figure 3).

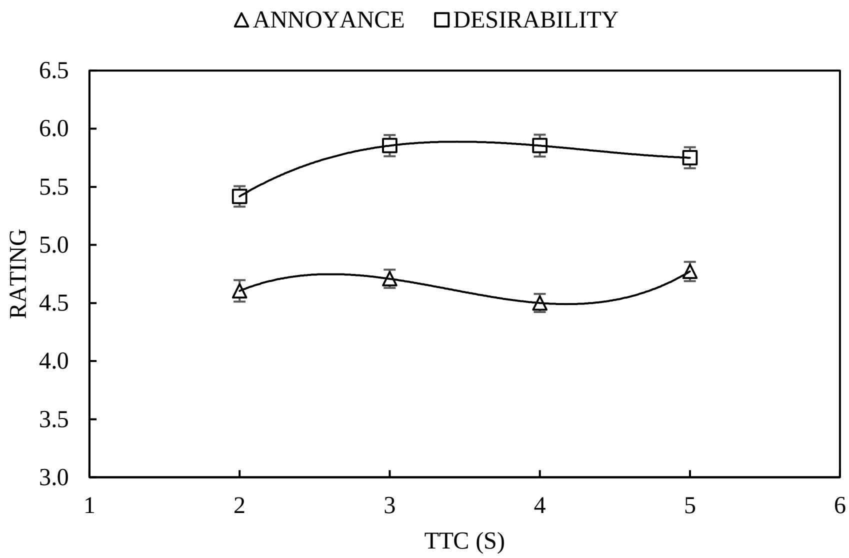

Figure 3: Subjective ratings of Annoyance and Desirability, with standard error bars

Driving Performance

Given that drivers also experienced the scenario without warnings, driving performance measures were also compared between warning types. A two-way, repeated-measures ANOVA was conducted to compare hazard-response indicators (*foot off accelerator*, *foot on brake* and *car stopped*) between different warning types (HUD/audio, audio only or no warning) and TTCs (2.0, 3.0, 4.0 and 5.0-seconds). During the calculation of these metrics, no distinction has been drawn between participants moving from the nearside or the off-side, nor between accurate detections and false positive alarms, with mean values generated for all five incidents in each scenario.

Foot Off Accelerator

There was a significant main effect of Warning Type ($F(2,46)= 44.78$ $p < .0005$) and TTC ($F(3,69)= 1352.7$ $p < .0005$) and a significant interaction between Warning Type and TTC ($F(6,138)= 1797$ $p < .0005$). Pairwise comparisons with Bonferroni corrections show that drivers first responded to the pedestrian hazard (i.e. lifted foot from the accelerator) much earlier when either combined HUD/audio or audio-only warnings were provided (both $p_{max} < .0005$), compared to situations where no warnings were given; there was no significant difference between HUD/audio or audio-only warning types.

Additionally, the distance at which drivers lifted their foot from the accelerator increased with increasing TTC, with drivers lifting their foot furthest from the pedestrian (M=150.9ft) when 5.0-second warnings were provided and closest (M=56.8ft) when 2.0-second warnings were given; all distances were significantly different from one another ($p_{max} < .0005$) (Figure 4).

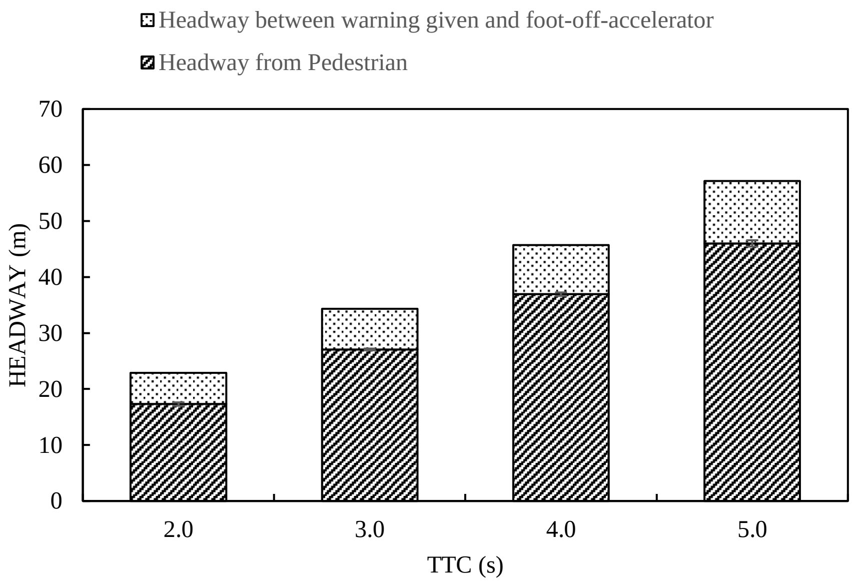

Figure 4: Mean headway from pedestrian when participants lifted their foot from the accelerator pedal, split by TTC

Foot On Brake

The distance at which participants first applied the brake in response to warnings, differed significantly based on TTC ($F(3,69)=86.93$, $p < .0005$), but was insignificant (@ $p < .05$) for warning type ($F(2,46)= 2.61$, $p = .084$). There was, however, a significant interaction between warning type and TTC ($F(6,138)=3.01$ $p = .009$). Pairwise comparisons with Bonferroni corrections indicated that drivers first applied the brake soonest when 5.0-second warnings were provided (M=32.8m, $p < .0005$); drivers were latest to apply braking following 2.0-second warnings (M=12.7m, $p < .0005$). Brake actuation times/distances were similar for 3.0 and 4.0-second warnings (Figure 5).

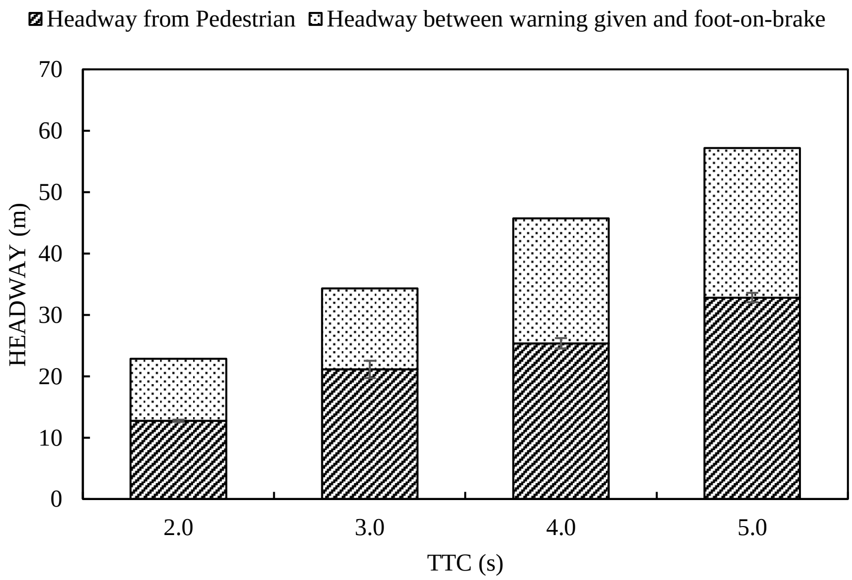

Figure 5: Mean headway from pedestrian when participants first applied their foot to the brake pedal, split by TTC

Car Stopped

There was a significant main effect of TTC (F(3,69)=19.60 $p < .0005$) and near-significant main effect of warning type (F(2,46)=2.86 $p = .067$). Pairwise comparisons with Bonferroni corrections showed that drivers stopped the car furthest from the pedestrian when audio-only warnings were presented (M=15.0m), compared to situations where either HUD (M=13.9m) or no warnings (M=11.7m) were provided (both $p < .0005$).

Drivers stopped closest to pedestrians when the warning time was 2.0-seconds (compared to all other warnings) (M=6.2m, $p_{max} = .001$). Additionally, the headway to pedestrians associated with 4.0-second (M=16.8m) warnings was larger than both 3.0-second and 2.0-second warnings (M=6.2m and 13.0m, respectively; both $p_{max} < .0005$). Headway from the pedestrian was greatest when warnings were provided at 5.0-second TTC (18.0m) (Figure 6).

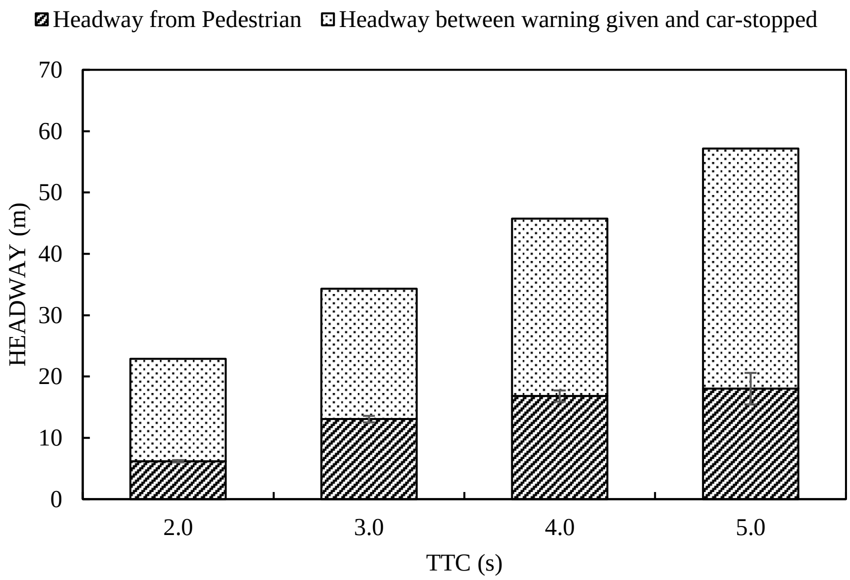

Figure 6: Mean headway from pedestrian when car stopped or reached minimum speed, split by TTC

Total Stopping Distance

To consider total stopping distance, driving performance measures were combined (Figure 7). Total stopping distance increased significantly with TTC, with participants taking the longest distance to stop when warnings were provided at 5.0-seconds (M=39.1m). The shortest stopping distances were associated with 2.0-second warnings (M=16.7m).

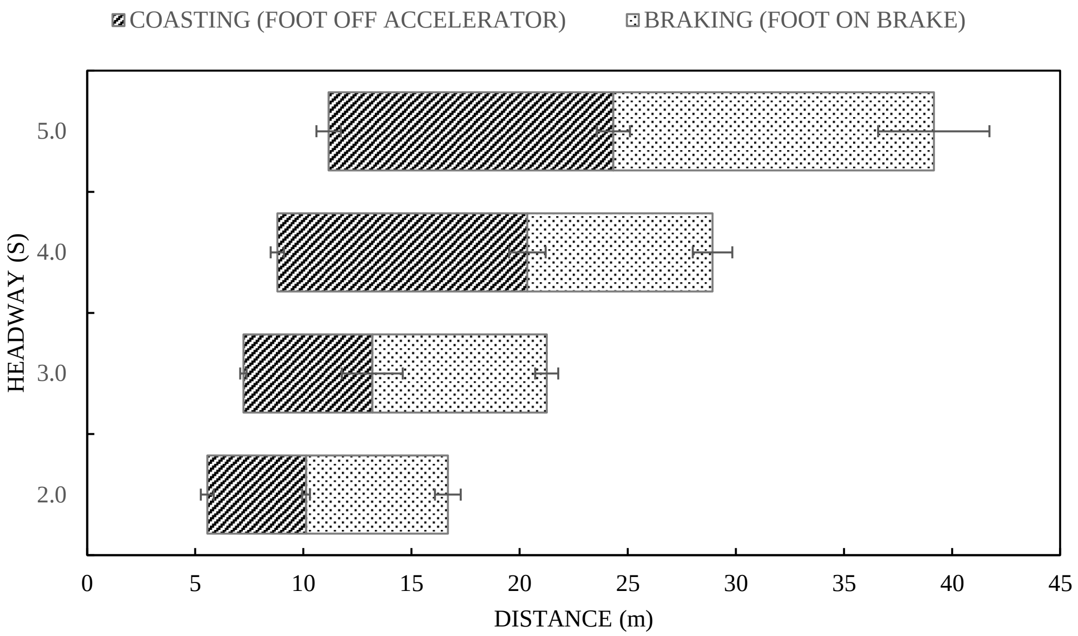

Figure 7: Total stopping distance from initial delivery of warning (at '0'-feet), split by TTC

4. Discussion

The study explored the relationship between false alarms and driver acceptance of a pedestrian alert system by exposing drivers to different warning types and TTCs/headways using visual and auditory alerts. It is evident that participants placed the highest level of trust and confidence in the system when warnings were presented at intermediate TTCs (3.0 and 4.0 seconds). These corresponded to false alarm rates of 40% and 60%, respectively, suggesting that drivers were apparently willing to accept relatively high false alarm rates without affecting their trust or confidence in the system. It is noteworthy that ratings of trust and confidence reduced significantly with both increasing *and* decreasing TTC, with the lowest levels evident at the shortest TTC of 2.0 seconds. This is likely to be because drivers felt that these warnings were provided too late for comfort rather than due to concerns about the technical capability of the system. Participants also indicated that they were most likely to use the system (revealed through the construct of 'Desirability') based on their experiences of 3.0 and 4.0 second warnings.

There were few differences highlighted specifically associated with ratings of Annoyance, although drivers tended to rate warnings delivered at the longest TTCs as most annoying. At this distance, it is possible that drivers were unaware of the potential threat posed by the pedestrian – indeed, earlier alerts can often be interpreted as 'false positives' [28]. Alternatively, drivers may have felt that they had sufficient time to respond without the assistance of a PAS.

These findings are also supported by the driving performance data, which show that earlier warnings encouraged drivers to respond sooner (i.e. reduce speed and begin braking), than in situations where there were no warnings or warnings were provided later. For example, drivers lifted their foot from the accelerator at 46.0m from the crossing pedestrian when provided with 5.0-seconds notification – 28.7m sooner than when 2.0-second warnings were provided. Similarly, drivers placed their foot on the brake 20.1m sooner following warnings at 5.0-seconds compared to 2.0-seconds and stopped the car 11.8m sooner. These data suggest that alerting drivers later is likely to delay the onset of braking. Whilst this may be considered inevitable, it also indicates that drivers were delegating the responsibility for hazard-detection to the technology, and waiting for the system to declare the hazard before taking evasive action. Although this is a further indicator of trust in the system, it also suggests disengagement from the environment. Further (ongoing) analysis of eye-tracking data (also collected as part of the study) is likely to elucidate this point – in particular, whether the driver was aware of the pedestrian hazard before it was highlighted by the system.

It is noteworthy that overall stopping distances (determined from the presentation of the warning to 'car stopped') were longer following earlier warnings, suggesting drivers took advantage of the longer headway and applied a more gradual braking force in these situations, rather than evoking emergency responses: this implies that informing drivers sooner should therefore be inherently safer. It is also of interest that some participants did not bring the vehicle to a complete stop in all situations (particularly at longer TTCs), but instead reduced speed (in anticipation of stopping) until they could determine whether the pedestrian was actually crossing or not. For the purpose of analysis in these situations, the distance at which minimum speed was achieved was recorded. However, this naturally tended to occur later than had the driver actually stopped the vehicle, leading to greater variability in this metric, particularly at a TTC of 5.0-seconds (Figure 7). This phenomenon was most evident with 5.0-second warnings, but also occurred – to a lesser extent – during the other TTCs, with the number of 'non-stop' events correlating with false alarm rates, i.e. drivers tended to continue moving in situations when it was evident that the pedestrian had stopped at the roadside. This suggests that drivers may begin to ignore the advice of the system and rely on their own judgement at longer headways, when the incidence of genuine alarms is much lower. This behaviour corresponds with existing literature, which demonstrates that human perception is more accurate and reliable at longer distances [28].

While one might conclude from the objective data alone that warnings should be provided as soon as possible, the subjective data tend to indicate an optimum level (between 3.0 and 4.0-seconds) corresponding with the highest ratings of trust, confidence and desirability associated with the system, and before drivers may begin to disregard the advice of the system.

Limitations of Study

For experimental convenience, the study design assumed that false alarm rate and TTC are closely correlated and increase at a linear rate. While there is strong evidence to support this (e.g. [24]), it is unlikely that this relationship is completely linear and correlations may be inconsistent: modelling system performance in this manner may therefore be an over-simplification and means that it is impossible to differentiate between TTC and false alarm rate. However, this was a necessary compromise to ensure internal validity. Moreover, the study assumed that the PAS was technically proficient (i.e. completely absent of false negative alarms), thus 'false alarms' were only created by the behaviour of deviant pedestrians. In reality, there are many other factors that are likely to influence the reliability and capability of PAS and contribute to false declarations. Nevertheless, our approach to evaluating system performance is predicated on the belief that system design should seek to maximise driver acceptance of the system.

It is also worth noting that in a real-world situation, subtle cues, such as head-movements, eye-contact and other postural changes, would likely indicate a pedestrian's intent to enter the roadway, or demonstrate that they had seen the approaching vehicle and were modifying their behaviour accordingly. Such behaviour was notably absent from the simulation (due to a lack of fidelity in the software), and this may have subsequently affected drivers' behaviour and responses.

5. Conclusions

The study revealed that, although drivers are likely to stop sooner and apply more gradual braking force, if warnings are provided earlier – thereby providing a greater safety margin between vehicle and pedestrian – warnings that are provided too early (or indeed, too late) are likely to annoy drivers and inspire lower levels of confidence in the technology. This can lead to drivers disregarding the advice and relying on their own judgement, especially at longer TTCs.

Based on the results, we could conclude that, in designing for maximum driver acceptance – and assuming that the system is technically proficient – warnings associated with PASs should be provided at intermediate TTCs (between 3.0 and 4.0-seconds TTC, based on the current data). However, this conclusion should be treated with caution – findings were based on repeated exposure over a short duration and in rapid succession. In a real-world situation, the system would likely intervene only very occasionally, and drivers' view of 'acceptance' may consequently differ. Therefore, results should not be applied verbatim but rather inform further long-term investigations and validation work. Further work should also consider other factors that are also likely to influence the reliability and capability of PASs and contribute to false alarms, such as variations in the behaviour and trajectories of pedestrians.

Acknowledgements

The study was funded by the European Community's Eighth Framework Program (Horizon2020) (agreement no. 634149) and conducted as part of the Proactive Safety for Pedestrians and Cyclists (PROSPECT) project.

References

1. World Health Organization, "Global status report on road safety," 2015.
2. K. J. Clifton, C. V. Burnier and G. Akar, "Severity of injury resulting from pedestrian–vehicle crashes: What can we learn from examining the built environment?," *Transportation research part D: transport and environment,* vol. 14, pp. 425-436, 2009.
3. L. F. Beck, A. M. Dellinger and M. E. O'neil, "Motor vehicle crash injury rates by mode of travel, United States: using exposure-based methods to quantify differences," *American Journal of Epidemiology,* vol. 166, pp. 212-218, 2007.

4. F. Graves, D. Lloyd, D. Wilson, D. Mais, W. Deda and A. Bhagat, “Reported Road Casualties Great Britain: 2013 Annual Report,” 2014.
5. M. Yanagisawa, E. Swanson and W. G. Najm, “Target Crashes and Safety Benefits Estimation Methodology for Pedestrian Crash Avoidance/Mitigation Systems,” 2014.
6. M. Lindman, L. Jakobsson and S. Jonsson, “Pedestrians interacting with a passenger car a study of real world accidents,” in *IRCOBI*, Krakow, Poland, 2011.
7. Mobileye, “Mobileye Foward Collision Warning (FCW),” Online.. Available: http://www.mobileye.com/technology/applications/vehicle-detection/forward-colision-warning/. Accessed 31 May 2016..
8. A. Broggi, P. Cerri, S. Ghidoni, P. Grisleri and H. G. Jung, “A new approach to urban pedestrian detection for automatic braking,” *IEEE Transactions on Intelligent Transportation Systems,* vol. 10, no. 4, pp. 594-605, 2009.
9. S. Zhang, R. Benenson, M. Omran, J. Hosang and B. Schiele, “How Far are We from Solving Pedestrian Detection?,” *arXiv preprint arXiv:1602.01237,* 2016.
10. J.-E. Källhammer, K. Smith and E. Hollnagel, “An empirical method for quantifying drivers' level of acceptance of alerts issued by automotive active safety systems,” in *Driver acceptance of new technology: theory, measurement and optimisation,* M. A. Regan, T. Horberry and A. Stevens, Eds., Ashgate, 2014, pp. 121-136.
11. M. Saffarian, J. De Winter and R. Happee, “Automated driving: human-factors issues and design solutions,” in *Human Factors and Ergonomics Society Annual Meeting*, 2012.
12. R. Parasuraman and V. Riley, “Human and automation: use, misuse, disuse, abuse,” *Human Factors,* vol. 39, no. 2, pp. 230-253, 1997.
13. M. McCallum, J. Campbell, C. Richard and J. Brown, “Integrated vehicle-based safety system heavy truck driver vehicle interface (DVI) literature review (No. UMTRI-2008-22),” 2006.
14. E. Adell, Driver experience and acceptance of driver support systems: a case of speed adaptation, 2009.
15. T. Sheridan, “Risk, human error, and system resilience: fundamental ideas,” *Human Factors: The Journal of the Human Factors and Ergonomics Society,* vol. 50, no. 3, pp. 418-426, 2008.
16. J. Källhammer and K. Smith, “Assessing contextual factors that influence acceptance of pedestrian alerts by a night vision system,” *Human Factors: The Journal of the Human Factors and Ergonomics Society,* vol. 54, no. 4, pp. 654-662, 2012.
17. M. Ghazizadeh, J. Lee and L. Boyle, “Extending the Technology Acceptance Model to assess automation,” *Cognition, Technology & Work,* vol. 14, no. 1, pp. 39-49, 2012.
18. J. Schade and M. Baum, “Reactance or acceptance? Reactions towards the introduction of road pricing,” *Transportation Research Part A: Policy and Practice,* vol. 41, no. 1, pp. 41-48, 2007.
19. K. Ausserer and R. Risser, “Intelligent transport systems and services-chances and risks,” in *ICTCT-workshop*, 2005.
20. J. Lee, D. McGehee, T. Brown and M. Reyes, “Collision warning timing, driver distraction, and driver response to imminent rear-end collisions in a high-fidelity driving simulator,” *Human Factors: The Journal of the Human Factors and Ergonomics Society,* vol. 44, no. 2, pp. 314-334, 2002.
21. J. Lee and B. Kantowitz, “Perceptual and cognitive aspects of intelligent transportation systems,” *Human factors in intelligent transportation systems,* 1998.
22. L. Bainbridge, “Ironies of automation,” *Automatica,* vol. 19, no. 6, pp. 775-779, 1983.
23. J. D. Lee and K. A. See, “Trust in automation: designing for appropriate reliance,” *Human Factors,* vol. 46, no. 1, pp. 50-80, 2004.
24. PROSPECT (Proactive Safety for Pedestrians and Cyclists), “Deliverable D2.1: Accident Analysis, Naturalistic Observations and Project Implications,” 2016.
25. C. Keller and D. Gavrila, “Will the pedestrian cross? a study on pedestrian path prediction,” *IEEE Transactions on Intelligent Transportation Systems,* vol. 15, no. 2, pp. 494-506, 2014.
26. J. L. Campbell, C. M. Richard, J. L. Brown and M. McCallum, “Crash warning system interfaces: human factors insights and lessons learned,” US Department of Transportation, National Highway Traffic Safety Administration, 2007.
27. N. Reader, “Development and validation of functional definitions and evaluation procedures for collision warning/avoidance systems,” 1999.
28. J. Brooke, “ SUS-A quick and dirty usability scale,” *Usability evaluation in industry,* vol. 189, no. 194, pp. 4-7, 1996.
29. G. Abe and J. Richardson, “Alarm timing, trust and driver expectation for forward collision warning systems,” *Applied ergonomics,* vol. 37, no. 5, pp. 577-586, 2006.