SHOULDN'T CARS REACT AS DRIVERS EXPECT?

Jan-Erik Källhammer,[1,2] Kip Smith,[2] Johan Karlsson,[1] Erik Hollnagel[2]
[1]Autoliv Research, Vårgårda, Sweden
E-mail: jan-erik.kallhammer@autoliv.com
[2]Department of Computer and Information Science
Linköping University, Sweden

Summary: The objective of this project is to develop and test a multi-method empirical approach for predicting drivers' assessments of the level of acceptability of a warning issued in response to accidents, near-accidents, and other incidents. The role of humans (drivers) in the pre-crash phase means that systems that protect occupants and pedestrians must be seen as distributed, cognitive systems. Driver acceptance therefore has to be an important design goal. One obstacle to acceptance is the human dislike for false alarms. An approach to overcoming driver dislike for false alarms is to focus on driver expectations and to design systems to issue alarms when and only when the driver is likely to accept them. In this paper we discuss one such approach.

INTRODUCTION

A wide variety of active safety systems currently under development promise to reduce significantly the frequency and severity of injuries sustained in collisions. To make progress toward the goal of reducing fatalities and accidents on our roads, designers of safety systems have focused on the prevention or mitigation of accidents. Issues of automation, system accuracy, and user reactions will therefore become major stumbling blocks to the introduction of systems for collision warning and avoidance. Since humans (drivers) play an important role in the pre-crash phase, systems that protect occupants and pedestrians must be seen as distributed, cognitive systems rather than as purely technical systems. Driver acceptance therefore becomes an important design goal. Accordingly, knowledge of how drivers and vehicle systems function together is critical in achieving a successful design. One hallmark of this interaction is the human dislike for false alarms. Indeed, the rate of false alarms may be a key factor for driver acceptance of novel active safety systems. Unfortunately, because even highly accurate systems are likely to have a high rate of false alarms, they are likely to encounter driver resistance. This generates a dilemma for designers of active safety systems. An approach to overcoming the dilemma posed by driver dislike for false alarms is to focus on driver expectations and to design systems to issue alarms only for conditions where the driver is likely to accept them. In this paper we discuss one such approach.

Historically, the goal when designing and developing active safety systems has usually been to achieve the highest rate of system accuracy from an engineering perspective. Unfortunately, false alarms are the rule when the base rate of the event to be detected is low (Parasuraman, Hancock, & Olofinboba, 1997). When false alarms are the rule, people either ignore them or disable the system. Because the active safety systems currently under development will not become 100% accurate within the foreseeable future, designers and developers face a dilemma: How to design

the system so that drivers and society can reap their benefits while minimizing the likelihood of rejection due to the innately human dislike for false alarms?

Farber and Paley (1993, cited in Parasuraman et al., 1997) speculate that the ideal detection algorithm may provide alarms in conditions that may lead to collision, even though the driver will probably avoid the accident. A warning only for situations leading to a traffic accident will be very rare. In the US during 2004 there were 38,253 fatal, police-reported motor vehicle traffic crashes, 1,862,000 crashes with injuries and 4,281,000 crashes with property damage only (NHTSA, 2005). Given that the total vehicle miles traveled was 2,962 billion with almost 199 million licensed drivers, a fatal motor vehicle traffic crash would be expected once every 5,200 driver years. Similarly, a crash resulting in an injury would be expected once every 107 driver years and property damage crashes once every 46 driver years. A warning *only* in situations leading to any of these crashes would therefore be so rare that it would likely exacerbate the already critical situation and make the driver's reactions hard to predict. Thus, we need to accept warnings also in situations that do not always lead to a crash.

Instead of a traditional engineering performance criteria, the goal may equally well be to maximize the driver's expectations of when the system should become active. By conforming to driver expectations, the system should be able to achieve a relatively high level of user acceptance, and become an effective partner in the driver-vehicle system. Following this approach we therefore need to identify in which situations drivers would expect alerts and therefore also accept them as relevant and appropriate. In the next section we describe our empirical approach to identifying and classifying the types of situations in which drivers are more likely to welcome an alarm or intervention from an active safety system. We then discuss the first set of results from our study and their implications for designing active safety systems that respond when, and only when, drivers do.

METHOD, EXPERIMENT 1

The study was designed to expose a driver to ordinary driving situations that might (not) pose the opportunity for an active safety system to issue a warning. After the drive, the driver rated how welcome a warning would have been to each of those situations.

Participants and task

Twelve drivers, 3 female and 9 males (age: M 28 yr, SD 2.0, range 23 - 30; driving experience: M 9.7 yr, SD 2.2, range 5 - 12) participated in our study. All reported driving relatively old cars (M 9.7 yr, SD 4.1) with manual transmissions and without modern safety systems like stability control. All but one reported driving at least 10,000 km annually. Four participants had corrected-to-normal vision. The volunteers were recruited at the Autoliv plant in Vårgårda Sweden, in a rural setting about 60 km from Gothenburg, Sweden.

Task

Each volunteer individually drove an instrumented car alone for a total of approximately three hours in normal traffic. Their assigned route presented the driver with a mix of rural, industrial,

highway, and urban driving. The route description was provided by a Tom-Tom GO 910 nomadic navigator with a pre-programmed route. The navigation system's guidance to the driver was provided on a 4" LCD display screen and by a synthesized voice played through the vehicle's audio system. The driver could adjust the audio level with the normal audio controls. A GPS-based vehicle tracking device enabled the researchers to track the progress of the vehicle over the internet. A cell phone was also provided in case the driver wanted to contact the researcher.

The driver was instructed to drive safely and normally and to follow the route indicated by the navigator. In addition, they were instructed to push a large red button on the dashboard to flag any and all incidents perceived as potentially risky, threatening, or otherwise worthy of retrospective analysis. No other instructions were given about when and in which situations to push the button.

The instrumented vehicle

The vehicle (a Volvo V70) was equipped with a four-camera Smart-Eye eye tracking system that continuously recorded the driver's gaze direction. Three roof-mounted cameras with a 180 degree field of view captured the environment and traffic dynamics. Continuous recording of these data sets enables an off-line visualization tool to generate a movie of the entire drive with a superposed circle representing the driver's gaze direction. A snapshot is shown in Figure 1. In addition, the car's location (from GPS) and all data sent to the CAN bus were continuously recorded.

Procedure and dependent variables

Data collection was split on two days, with the first day consisting of going through the written instructions to participants, informed consent, calibration of the eye-tracking device, and general familiarization of the vehicle and instrumentation.

**Figure 1. Screenshot from offline visualization tool. Gaze direction
from eye tracker visualized as circle**

The second day started immediately with the driving session. Driving was conducted outside rush-hour from about 9 AM to noon. The drive was designed to expose the driver to ordinary driving situations that might or might not lead an active safety system to issue a warning. Three

triggers were used to define those situations. The first was the driver's pushing the button on the dashboard. Two others were automatic triggers generated by longitudinal deceleration (XAcc) and fast brake preparedness (Pedal). These triggers flagged additional situations to complement those identified by the driver using the push button. Brake preparedness was derived by measuring the time between the release of the accelerator and the activation of a proximity sensor attached to the brake pedal. The threshold values of both automated triggers were adjusted to produce about 30-40 situations for each driver.

After the drive (and the inevitable Swedish lunch), the driver met with the researchers for approximately two hours to conduct a post-drive review of the situations flagged by the three sets of triggers (Button, XAcc, Pedal). The review session was conducted as a structured interview. The driver watched a video of the drive (showing the image from the roof-top cameras and the eye-tracker, as shown in Figure 1) of the 20 seconds before and 10 seconds after each trigger. After viewing a video-clip, the driver was asked to rate the level of acceptance of an alarm in the presented situation. The values were entered on the computer using a slide bar with the continuous scale from 0 (totally unacceptable) to 100 (totally acceptable) shown in Figure 2. We have analyzed these ratings to assess the likelihood of driver acceptance of an active safety system to actual driving situations.
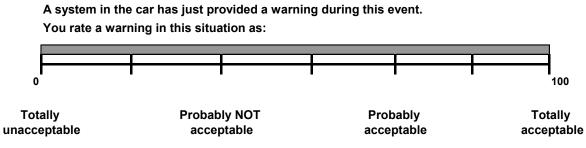
A system in the car has just provided a warning during this event.
You rate a warning in this situation as:

0						100

Totally unacceptable	**Probably NOT acceptable**	**Probably acceptable**	**Totally acceptable**

Figure 2. Scale to rate the level of acceptance of an assumed warning

RESULTS

Relative acceptance of warnings by triggers

Half of our volunteer drivers (6 of 12) pushed the button fewer than 4 times during the 3-hour drive. This half of our sample apparently did not perceive many situations as worthy of alerts or of retrospective evaluation. In contrast, the other half pushed the button more than a dozen times and one 34 times. This difference suggests that many in the driving public may be reluctant to accept the introduction of automated warning systems. This resistance to systems poses a severe challenge to the future of active safety systems and underscores the need for designers to understand the (rare) situations in which a warning might be welcome.

The 12 drives recorded a total of 99 button pushes, 189 rapid pedal shifts, and 111 rapid decelerations. A between ANOVA for trigger conditions indicates that driver ratings differ across the three triggers, $F(2, 396) = 78$, $p < .001$. Post-hoc comparisons indicate that all three differ from each other, with the degree of acceptance greatest for button pushes and least for

rapid pedal shifts. It appears that it may be tricky to develop automatic triggers based purely on these data.

Categorization of situations

We iterated several times to develop a parsimonious categorization of the 399 situations captured by the three types of triggers. Our original scoring scheme included 19 categories (e.g., animal in road, pedestrian, speed bump, intersection). This scheme has been consolidated to the five categories shown in Figure 3. This scheme is admittedly arbitrary but is manageable and replicable.

Relative acceptance of warnings by situation type

The ratings given to each incident were sorted by situation type (Figure 3) and driver. Each driver's ratings for each situation type were averaged and analyzed using a repeated measures ANOVA, where categories are the repeated measure and drivers the blocking variable. This procedure controls for differential scale use across drivers. The ANOVA rejects the null hypothesis of no difference across categories, $F(4, 43) = 4.08$, $p < .01$. A post-hoc comparison shows that the ratings given to alerts to pedestrians are significantly different ($p < .05$) than ratings to all other categories except vehicles ahead, and that ratings given to traffic differs significantly from ratings to vehicles ahead ($p < .05$).
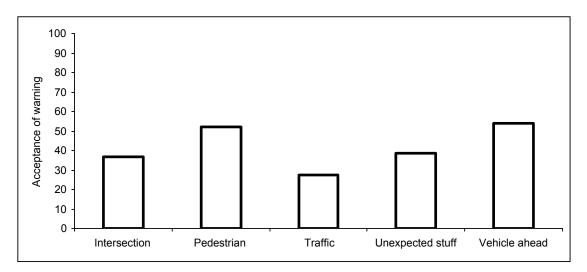

Figure 3. Categories of incident types with mean ratings
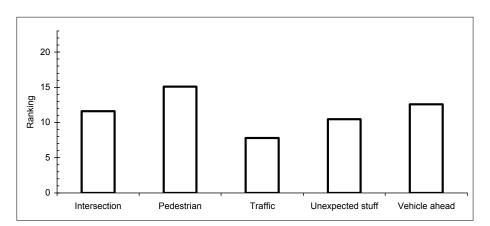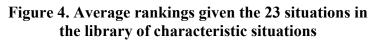
EXPERIMENT 2

Method

To generate additional ratings of the relative acceptability of a warning to actual driving situations, we assembled a library of 23 representative video-clips from Experiment 1 and conducted a table-top test. An independent group of eight volunteers watched all 23 clips and rated them using the replay procedure of Experiment 1 and the scale shown in Figure 2.

Rankings of characteristic situations

The eight participants in Experiment 2 gave lower ratings to situations triggered by the Button than the driver in Experiment 1 who pushed the Button. This difference is likely explained by the driver's recall of the full contextual significance of the event. No such difference was found for situations triggered braking and brake preparedness.

Because all eight volunteers rated the same situations, it is possible to rank their ratings (within subjects) and to ascertain the degree of their consensus about the relative acceptability of alarms across event categories. We found that consensus to be statistically significant (Kendell's $W = 0.210$, $\underline{p} < .01$). The data are shown in Figure 4. The average ranking of the situations in this library reveals a consensus that alerts issued in response to threats posed by pedestrians are more welcome than alerts to the other four types.

**Figure 4. Average rankings given the 23 situations in
the library of characteristic situations**

DISCUSSION

The objective of this project is to develop and test a multi-method empirical approach for predicting drivers' assessments of the level of acceptability of a warning issued in response to accidents, near-accidents, and other incidents. The argument for this is that the acceptance of an active safety system is likely to improve if the system activates only when the driver finds it reasonable. The ultimate objective is to make use of this predictive capacity in an active safety system.

Minimizing false alarms is a key to acceptance of active safety systems. The low base rates of an accident imply that the posterior probability of a collision given an alarm will also be low, even if the sensing system will have a high probability of detection (Parasuraman et al., 1997). Also, as most drivers will experience at most a few real accidents in their lifetimes, such low frequency of alarms would make drivers' reactions hard to predict. Driver acceptance has to be an important design goal for any collision warning and avoidance systems. The rate of false alarms will be key factor for driver acceptance.

The ratings by drivers, shown in Figure 3, and the rankings by participants in the table-top study, shown in Figure 4, are markedly similar. The correlation between the two data sets is .89. This impressively high correlation has auspicious implications. First, it indicates that we have converging evidence for the relative acceptability of warnings issued in response to situations involving pedestrians and vehicles ahead. Ratings and rankings given to warnings to other categories of situations (intersections, traffic, and unexpected stuff) were consistently lower.

Second, the high correlation between the retrospective ratings by drivers who triggered events and the rankings by independent volunteers of video-clips of selected events has implications for future empirical studies. It implies that table-top tests using video clips of events would likely generate responses that would correlate positively with the responses of drivers in cars on the road. If this finding receives additional support, it may be possible to simplify considerably the testing of active safety warning strategies. Table-top reply sessions could readily compare driver responses to simulations of a wide variety of warning strategies and could be conducted with a large number of participants. The naturalistic setting afforded by the instrumented vehicle could then be used to confirm or disconfirm findings from the table-top sessions.

CONCLUSION

We have found converging evidence for the relative acceptability of alerts issued in response to situations involving pedestrians and vehicles ahead. These findings can be used to guide the development of active safety systems. The increased understanding of alarm acceptance by drivers in various situations will be used to develop warning strategies that will likely lead to higher driver acceptance and trust in the safety systems. System performance may be optimized based on driver acceptance rather than objective performance criteria that do not necessarily match the expectations of the drivers. After all, the driver expects the car to react in an acceptable way.

ACKNOWLEDGMENTS

REFERENCES

Farber, E. & Paley, M. (1993). Using freeway traffic data to estimate the effectiveness of rear end collision countermeasures. Paper presented at the *Third Annual IVHS America Meeting, IVHS America*, Washington, DC.

NHTSA [National Highway Traffic Safety Administration]. (2005). *Traffic Safety Facts 2004.* Washington, DC: U.S. DOT.

Parasuraman, R., Hancock, P.A., & Olufinboba, O. (1997). Alarm effectiveness in driver centered collision-warning systems. *Ergonomics 40*, 3, 390 - 399.

Sorkin, R.D. & Woods, D.D. (1985). Systems with human monitors: A signal detection analysis. *Human Computer Interaction 1*, 49-75.