KEY HUMAN FACTORS DEVELOPMENT PRINCIPLES FOR DMS ENHANCED COLLISION AVOIDANCE SYSTEM DEVELOPMENT

Mikael Ljung Aust

Emma Tivesten

Pär Gustavsson

Christian Applehult

Volvo Cars

Sweden

Paper number 23-0263

ABSTRACT

Driver monitoring systems (DMS) can enhance Collision Avoidance Systems (CAS) in numerous ways, for instance by adjusting warnings or interventions when drivers are inattentive or in other ways disengaged or impaired. However, the driver interaction principles applied when using DMS to enhance CAS must be based on State-of-the-Art Human Factors research and have a clear focus on understanding driver needs and in what way assistance should be provided to be appreciated by the driver. Otherwise, one risks implementing interactions that either do not make sense or are perceived as disturbing, both of which degrade the CAS's safety potential.

Some of these interaction principles may not be fully intuitive unless your background is in behavioral psychology. For example, it may be surprising that DMS is best used to delay certain collision avoidance warnings rather than supply them earlier. It may also not be fully intuitive that DMS is best used for detection of generic degradations in the behavioral patterns that define normal driving rather than for diagnosis of specific impaired states.

To use a CAS properly, you need to interact with it regularly to learn what its outputs mean. However, current accident and mileage statistics suggest that driving conflicts where a CAS could save you from an unrecoverable error that otherwise would have resulted in a high severity crash are rare; maybe as infrequent as once in a decade or lifetime depending on how one does the calculation.

From a design perspective, CAS are therefore best approached as lifetime driving companions. You may only need them once, but they still need to be interacted with regularly to work as intended. Hence, the conversation between driver and CAS should adhere to the same principles as applied between humans. For example, if your colleague is busy, you only interrupt for good reason, and if you interrupt regularly, both of you must agree its relevant and the message must be clear (though not necessarily loud) so the other person quickly can decide whether to interrupt the current task.

In this paper, first a general framework and corresponding design approach for CAS is formulated based on accident statistics, driving mileage and CAS interaction frequency analysis. Next, three specific development principles for DMS enhanced CAS are described to illustrate what the outcome is when the framework and design approach are applied in practice. These include how DMS enhancement can be used to avoid "cry wolf" effects in CAS interactions, how DMS enhancement can be used to get CAS timing right for both distracted and aware drivers and finally, how DMS offers a more efficient way than specific state diagnosis when tackling driver impairment.

By explicitly describing these fundamental Human Factors development principles for DMS enhanced CAS to the traffic safety engineering community, one may avoid unnecessary development pitfalls that could counteract DMS enhanced CAS deployment.

DEFINING A GENERAL FRAMEWORK

To define a general framework for the interaction principles for CAS, one has to consider three key facts. The first is that even though CAS come in many forms, they all have one thing in common which is that to use them properly, the driver has to know how they work. So, as long as a driver uses the same car, s/he needs to interact with the CAS regularly to first learn what the systems do, and then remember what their outputs mean.

Second, for the CAS to achieve its intended safety benefit, drivers need to trust their inputs and actions (i.e., the warnings and interventions provided). If driver trust in, and understanding of, the CAS can be established in a good way, chances are good that drivers will cooperate with the CAS in the intended way and perceive the CAS actions as beneficial. However, if what the CAS does is perceived as unintelligible, pointless or perhaps even scary, it follows that drivers will neither use the CAS as intended by the designer, nor spend money on CAS features in their next car purchase.

Third, when one studies how often people end up in traffic accidents and compare that to the distances travelled or hours driven without crashing, one can conclude that driving conflicts where an CAS would save you from unrecoverable error are rare. For example, in 2020 the US average was 1,33 deaths in 100 million vehicle miles travelled [1]. At the same time, the average mileage per car was about 13 500 miles per year [2]. This means that there are about 5500 years of successful driving per fatal crash.

Of course, non-fatal crashes are far more common, but the key point here is that crashes, and particularly severe crashes, are very infrequent compared to all the driving we do; maybe as infrequent as once in a decade or lifetime depending on which crash type and outcome one is looking at. This means that the average driver exposure to critical situations where the safety margins are so small that severe injury or death is imminent if not handled correctly is very low.

Putting these three key facts together, it becomes clear that we cannot rely on exposure to a particular set of truly critical situations with potential high crash severity for drivers to understand and learn how a particular CAS works. This insight provides the foundation for the general framework proposed in this paper. This framework states that each CAS must be capable of providing a frequent enough driver interaction also outside of truly critical situations to provide a sustained, acceptable and trustworthy learning experience. The reason for this is that it's that experience which in turn secures an adequate driver response on those rare occasions when a truly critical situation occurs, and the CAS can provide a real safety benefit. This general framework forms the basis for the rest of the paper.

INCREASING MARGINAL TRUST - A FRAMEWORK COMPATIBLE APPROACH TOWARD CAS DESIGN

As shown in the general framework discussion above, to maintain an interaction frequency that is sufficiently high to create and maintain learning in the driver, most CAS interactions need to happen in situations that are not truly critical, since the latter are too infrequent to provide a robust exposure rate for learning.

Still, for the interaction with the CAS to make sense and provide that learning experience, there must be a clear connection between the CAS output and the type of critical situation which the CAS is meant to provide protection from or create awareness of. If that connection is unclear, it will be very hard for the driver to understand what the purpose of the CAS is.

From this follows an interesting design constraint for CAS interactions. Since most CAS are intended to act on, or make drivers aware of, critically small safety margins, the learning experience needs to happen in situations where the safety margins either are shrinking in an obvious manner toward a point where they would be considered too small (things would have gone bad if this state had been allowed to continue), or where they are sufficiently small[1] to create an intuitive learning experience once the driver becomes aware of them.

---

[1] Note though that they cannot be too small, since if they were, the CAS would not provide the relevant safety benefit in the situations where it is actually needed.

Moreover, since the designer of the CAS interaction will not be riding along to explain what s/he meant when shaping the system, the driver is the sole interpreter of both meaning and relevance. The connection between an CAS's outputs and the truly critical situations it intends to prevent or mitigate therefore has to be obvious to all drivers without additional explanatory material being present or read/watched in advance (note that it has been shown that at least 50% of customers learn about CAS through trial and error [3,4,5].)

Now, as many studies in pedagogy will testify, people are different and have different ways of understanding things. However, since most vehicle manufacturers make hundreds of thousands or millions of vehicles, the possibilities for tuning the shape of the CAS connection to suit individual preferences or explanatory models are limited. Also, even though more computational power and intelligent algorithms are expected to make their way into future cars, deploying CAS where each system gets tuned to each (current) driver is likely a too large and complex undertaking to warrant the effort required.

Thus, to secure an intuitive and relevant CAS interpretation in all drivers, one instead has to find approaches that speak to all drivers, regardless of who they are and where they drive. A successful CAS interaction pedagogy thus needs to be anchored in basic human factors design principles, i.e. principles that one has good reason to believe applies to all drivers in general rather than just a few, to meet with the capabilities of the technology.

Fortunately, this is an area where extensive research has been done over the years and lots of good advice and guidelines exist, both in the scientific community and in the more popular science literature, such as the highly entertaining books "The design of everyday things" [6] and "Things that make us Smart" [7]. Still, the advice on offer here covers many different applications and contexts. It is therefore worthwhile to define an approach that is specifically suited to the automotive domain in general and CAS interaction designs in particular.

Here, we propose that the concept of *Marginal Trust*, i.e. the study of how trust is built or lost based on the acquisition of new information or input [8] provides the approach we need to design CAS interactions. As described above, the success of a CAS in preventing or mitigating a crash depends on the meaning of its outputs being intuitive (what does it want with me?) and perceived as relevant for driving (does that make sense?).

Since every interaction a driver has with a CAS provides information on what the CAS intends to achieve, following the model in [8] one can say that these two questions are revisited in each such interaction, which in turn influences marginal trust in the CAS (i.e. it either increases, decreases or remains unchanged). Aiming to design CAS such that each interaction has a positive influence on marginal trust, or not a detrimental one if trust already is high, can therefore be stated as main design goal if the full safety potential of an CAS is to be realized.

Now, various types of trust have been identified in the literature [9,10]. Often, one makes a distinction between cognition- and affect-based trust [11]. Cognition-based trust is based on a rational evaluation of competence, responsibility, and dependability [10] and rests on logical and rational calculation of likely behavior and outcomes of future collaboration. Affect-based trust on the other hand happens when emotional bonds are created, and this bond can work as a replacement or surrogate for cognition-based trust, enabling people to take a "leap of faith" that trust will be honored [12].

Affective trust is based on beliefs that the one you place your trust in cares about your welfare, will act positively towards it and take care to avoid harming it. However, while clearly a strong influencing factor in the build of marginal trust, affective trust is rather hard to exploit through system interaction design. In the automotive domain, one might rather say that this type of trust instead should be built through brand communication and how one describes the stance of the company towards its customers in various contexts.

Cognition-based trust however fits the bill for CAS design. If viewing an CAS as a lifetime driving companion which you may only really need once but still need to interact with regularly, then the conversation between driver and CAS should strengthen the driver's view of the CAS as a competent, responsible, and dependable agent. This is very similar to conversations between humans. For example, if your colleague is busy (or in this case, if the driver is busy driving) you should only interrupt for good reason, and if you interrupt regularly, both must agree that its relevant, the message must be clear, etc.

Approaching CAS interactions as a design space where one needs to leverage cognition-based trust mechanisms to positively influence marginal trust thus seems to be a both viable and promising way forward. This brings us to the next question: how can DMS facilitate the design of CAS?

SPECIFIC DEVELOPMENT PRINCIPLES WHEN USING DMS TO ENHANCE CAS

So, how can one make use of DMS inputs to when one needs to leverage cognition-based trust mechanisms to positively influence marginal trust in CAS interactions? In the following, three specific principles will be described that illustrate how DMS can be used to support the designers in this endeavor. These principles are not meant to be exhaustive nor definitive; rather they serve to illustrate which knowledge must be acquired and which decisions need to be made when specifically using a DMS system to enhance CAS interactions.

Using DMS to avoid "Cry wolf" effects in CAS interactions

When applying collision warnings to alert the driver to external threats, the CAS doesn't have to be right every time in the sense that to the driver, a perceivable external threat always must exist to match the warning and hence make it appear relevant. However, the CAS cannot be wrong most of the time either, because then it will be perceived as a system which, in the words of Aesop [13], cries wolf all the time. Or in the context of the framework and approach described above, each false warning decreases marginal trust, so you can't have too many of those or all trust will be lost.

Some basic rules for how often you need to be right can actually be established. From a purely theoretically standpoint, one can make the argument that the system needs to be right more than half the time, otherwise the ADAS outputs literally appear stochastic in the eyes of the driver (being right every other time implies also being wrong every other time). Simply put, it has to perform better than chance.

Still, the CAS does not have to be right all the time. In an interesting study on behavioral adaptation to Lane Departure Warnings (LDW), Le Noy and Rudin-Brown [14] showed that drivers reported almost as high levels of trust in a flawed LDW system that was programmed to miss one in three true positives and also intermittently added false positives, as in a fully accurate LDW.

What the exact ratio of true to false positives should be to build sufficient driver trust in a CAS remains to be empirically determined. However, it is clear that a key aspect of CAS design is to secure a positive balance toward true positives, that means that the warning is warranted both from a situational and a driver perception perspective and avoid false positives where the warning appears to be given for no real reason.

Here, adding a DMS to the equation opens up for significant improvements of a CAS true to false positive ratio in two ways; one regarding the opportunities to be right and the other regarding the possibility to get the timing right.

To understand the being right part, we must first look at today's systems. These are often designed around a main conflict scenario with a number of exceptions added. For example, you might design your forward collision warning system (FCW) to warn when Time-to-Collision (TTC) is less than 1.7 s, based on the assumption that drivers would not voluntarily place themselves in a situation with such a small safety margin. However, once that general rule is in place, the developers immediately start adding exceptions for situations where the safety margins can become smaller, but where they believe the driver is in control and does not need a warning. An example of the latter would be overtaking. Here, TTC values can get very low while the driver is still in full control.

Still, even if the list of exceptions and their associated detection criteria is well thought out and tested in development and through customer feedback, it is very difficult to get it right all the time. Developers therefore generally take a conservative approach and avoid letting warnings through when they are unsure about whether the driver really needs the warning, an approach that potentially leads to under- or disuse of the CAS.

By adding a DMS to the vehicle, one can replace that exception list with real time analysis of where the driver's attention actually is directed when a conflict arises. In a first step, the system can determine whether the driver is looking in the direction of the threat at all. In a second step (though a bit more challenging on the detection side), the system can determine whether someone with his or her visual attention in the right place also is ready to act.

Adding these enhancements to the situation analysis can provide a huge step forward in terms of equipping CAS designers with the confidence they need to let their system warn the driver. They no longer need to make that decision based on assumptions and predictions about whether the driver is attentive or not. Instead, they can use the driver's actual, real time, direction of attention to help arbitrate whether CAS inputs are needed or not.

This opens up for a much more forward, less conservative, approach in terms of getting the CAS dialogue right with the driver. The achieved safety benefit of these systems could thus get a significant boost, since being able to increase the true to false positive ratio would build driver trust and confidence in the CAS actions.

Using DMS to get the timing right for both distracted and aware drivers

Even if a DMS can detect where the driver's attention is directed in a given moment and hence whether a CAS input might be warranted or not, the designer still has to decide on the input timing. Conceptually speaking, a key element in drivers' judgement of situational relevance will depend on what the safety margins are when the CAS activates. For example, consider a situation where you are approaching a lead vehicle and the DMS has determined that your visual attention is directed through the side window. Also, let's say the car is set to warn drivers of a potential lead vehicle collision extremely early, e.g., when TTC is 5 seconds. At 50 kph, this translates to being about 65 meters away from a stationary lead vehicle, which means most drivers would say there is no immediate danger present. Chances are therefore quite high that you as a driver would consider this a false warning when you get it. On the other hand, if the warning is set to come extremely late, e.g., when TTC is 0.5 s, there will not be enough time to react. Today's CAS systems have been developed to strike a balance between these end points. For example, the European New Car Assessment Programme (EuroNCAP9 requires that forward collision warnings be given at TTC > 1.2 s [15].

As described above, a basic assumption when deciding on a warning timing threshold is that the driver would not be in this position voluntarily, which means that there also is an underlying assumption that driver must be distracted or unaware for some other reason. An inherent challenge for this approach is the side effect it has on drivers who have their visual attention in the right place when the warning is given. Looking at what happens a bit more in detail, the time required for visually distracted drivers to move their gaze from e.g. a secondary task to the road scene ahead is typically around 500-700 ms [16]. If we start from a warning timing threshold of say 1.7 s TTC, the critical event will thus have developed to a point where TTC is 1.0-1.2 s before a visually distracted driver is able to assess the forward road scene.

Visually attentive drivers on the other hand do not need additional time to get their eyes back on the road. They will therefore assess the scene at the same time as the warning is given, i.e. when TTC = 1.7 s. The same warning timing thus presents a considerably less critical situation to the visually attentive driver, who therefore less likely to consider the safety margins small and the CAS input relevant. The same warning timing can thus lead to a true positive perception by the visually distracted driver but a false positive perception by the visually attentive driver.

One simple resolution to this problem would be to suppress all CAS inputs for visually attentive drivers. However, that would disregard the possibility that the driver, although visually attentive, for some reason does not realize the need to act and hence still needs the CAS warning. The latter is sometimes referred to as cognitively distracted drivers, who would be labelled as 'failed to look properly' or 'looked but failed to see' cases in British crash causation analysis [17].

Here, the DMS provides a very interesting alternative, since it can be used to *delay* warnings for the visually attentive drivers. Since they don't need those 500-700 ms to get their eyes back on the road, a delay of the warning by 500-700 ms will move the perception of the situation to the same TTC value as that for the visually distracted drivers (1.0-1.2 s in this example). This means that the visually attentive driver will have as much time to respond as the visually distracted driver. At the same time, the CAS is bringing attention to a driving situation that is more critical, in the sense that safety margins are perceivably smaller when the warning is issued. So, adding a delay gives visually inattentive and 'looked but failed to see' drivers equal time to react, but also increases the chance that a driver who is both looking in the right place and aware of the need to act still will judge the CAS input as relevant and meaningful, and hence as a true positive.

Using DMS to detect deviations from the normal driving

In addition to enhancing the precision of CAS warnings and interventions, DMS systems are also expected to have an application in the understanding of when drivers are fit to drive or not. The latter can result from many well-known risk states such as the driver being intoxicated or extremely sleepy [18,19]. While previous efforts largely focused on preventive work such as alcohol interlocks or “don’t drink and drive” campaigns, many new initiatives are being brought forward to promote detection and mitigation of impaired states while driving in both the legislation [20] and the consumer rating [21] arena.

A main, and at first sight reasonable, track in this endeavor to combat impaired driving is to try to turn new and existing vehicle sensors (including DMS systems) into diagnostic tools that can be used to precisely determine impaired states, such as intoxication above a certain level or a sudden drop in blood sugar to name a few. However, under closer scrutiny, this approach has several severe challenges. There are challenges coupled to the drivers’ privacy if the car would read and store medical diagnosis information. Also, there are severe technical gaps to be closed, since precise medical diagnosis that has to rely on non-invasive sensors and be performed in a moving vehicle is very hard. Additionally, if these two challenges are overcome, the result would be a system that most likely has to go through medical certification procedures for the information to be considered valid to act on. As medical certification procedures are, for a good reasons, both rigorous and slow, going down that route would make both development and updating of these features cumbersome and costly.

Now, this problem can be simplified greatly if one reverses the perspective. Instead of aiming to precisely diagnose specific driver impairment states such as a particular alcohol intoxication level, one can instead aim to detect generic degradations in the behavioral patterns that define normal driving and make the vehicle’s CAS act on those. Put differently; knowing that the driver has left normal driving behind is actually enough to give a CAS the freedom it needs to act on an impaired driver state. It does not have to be more precise than that.

The reason why this approach works in practice is because driving is a highly practiced and overlearned skill in most drivers. To travel those average 13500 miles a year [2], one has to spend several hundreds of hours driving, which also means that we train our speed, distance and lane keeping skills for hundreds of hours per year. Furthermore, the possibilities for variability (i.e. the possibility to drive in a very different style compared to how others drive) are limited if you’re to keep the car on the road and disrupt the traffic flow or pattern. And the latter you don’t want to do, since other drivers will give you feedback (basically tell you off) for doing so.

This means that the “normal driving box”, i.e. the control parameter space within which you normally constrain yourself while driving, is 1) small and 2) very similar between drivers. Hence, if we detect significant deviations from that box, we can deduce that you’re no longer driving like people normally do, which means there is now reason to think that you are impaired for some reason.

So, if the goal is to detect significant deviations from normal driving, having a DMS onboard offers the completely new opportunity to look at gaze patterns, in addition to vehicle control patterns. And as previous research has shown [22], gaze patterns can be a very good predictor of whether a driver has mentally checked out from driving, even though s/he is still in the driver’s seat and mostly looking at the road ahead.

Aiming to detect traffic relevant driver impairments by looking for significant deviations from normal driving offers a path toward combatting impaired driving that does not require medical grade detection procedures, which are also likely is more robust in the face of the natural variability that always is associated with large populations. Also, by leveraging DMS to study gaze patterns in additions to vehicle control patterns, one has new opportunity to give the CAS systems ‘license to intervene’ in situations where warnings and/or interventions might otherwise have been suppressed, for instance where pedal or steering wheel use studied by themselves might indicate an alert and aware driver.

CONCLUSIONS

To use an CAS properly, you need to interact with it regularly to learn what its outputs mean. However, current accident and mileage statistics suggest that driving conflicts where an CAS would save you from unrecoverable error are rare; maybe as infrequent as once in a decade or lifetime depending on crash type.

CAS are therefore best approached as lifetime driving companions. You may only need them once, but they still need to be interacted with regularly to work as intended. Hence, the conversation between driver and CAS should adhere to the same principles as applied between humans.

Based on this general approach, a few specific development principles for DMS enhanced CAS can be derived to maximize the benefit one can get by adding DMS assessments of driver state and attention to the CAS threat assessment platform. By explicitly describing these fundamental Human Factors development principles for DMS enhanced CAS to the traffic safety community, the designer may avoid unnecessary development pitfalls that could counteract the deployment of these systems.

REFERENCES

[1] National Center for Statistics and Analysis. (2022, May). Early estimates of motor vehicle traffic fatalities and fatality rate by sub-categories in 2021 (Crash Stats Brief Statistical Summary. Report No. DOT HS 813 298). National Highway Traffic Safety Administration.

[2] Highway Statistics (2019) . Available from https://www.fhwa.dot.gov/policyinformation/statistics/2019/

[3] Lubkowski, S. D., Lewis, B. A., Gawron, V. J., Gaydos, T. L., Campbell, K. C., Kirkpatrick, S. A., et al. (2021). Driver trust in and training for advanced driver assistance systems in real-world driving. Transportation Research Part F: Psychology and Behaviour, 81, 540–556.

[4] Abraham, H., Reimer, B., & Mehler, B. (2018, September). Learning to use in-vehicle technologies: consumer preferences and effects on understanding. In Proceedings of the human factors and ergonomics society annual meeting, 62(1), 1589-1593.

[5] Boelhouwer, A., Van den Beukel, A. P., Van der Voort, M. C., Hottentot, C., De Wit, R. Q., & Martens, M. H. (2020). How are car buyers and car sellers currently informed about ADAS? An investigation among drivers and car sellers in the Netherlands. Transportation Research Interdisciplinary Perspectives, 4, 100103.

[6] Norman, D. (2013). The design of everyday things: Revised and expanded edition. Basic books.

[7] Norman, D. (2014). Things that make us smart: Defending human attributes in the age of the machine. Diversion Books.

[8] White, M. P., & Eiser, J. R. (2006). Marginal trust in risk managers: Building and losing trust following decisions under uncertainty. Risk analysis, 26(5), 1187-1203.

[9] Young, L. (2006). Trust: looking forward and back. Journal of Business & Industrial Marketing. 21(7), 439-445.

[10] Huang, Y., & Wilkinson, I. F. (2013). The dynamics and evolution of trust in business relationships. Industrial Marketing Management, 42(3), 455-465.

[11] McAllister, D. J. (1995). Affect-and cognition-based trust as foundations for interpersonal cooperation in organizations. Academy of management journal, 38(1), 24-59.

[12] Lewis, J. D., & Weigert, A. (1985). Trust as a social reality. Social forces, 63(4), 967-985.

[13] Aesop. The Shepherd Boy & the Wolf. Available from https://read.gov/aesop/043.html

[14] Rudin-Brown, C. M., & Noy, Y. I. (2002). Investigation of Behavioral Adaptation to Lane Departure Warnings. Transportation Research Record, 1803(1), 30–37.

[15] European new car assessment programme (euroNCAP) assessment protocol – Safety assist. Collision avoidance - Implementation 2023. Available from: https://cdn.euroncap.com/media/75438/euro-ncap-assessment-protocol-sa-collision-avoidance-v102.pdf

[16] Aust, M. L. (2014). Effects of haptic versus visual modalities when combined with sound in forward collision warnings. In Proceedings of the Driving Simulation Conference 2014.

[17] Department for Transport (2014) Reported Road Casualties Great Britain: 2014 Annual Report. Available from: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/463043/rrcgb2014-02.pdf

[18] National Highway Traffic Safety Administration. (2020, October). Quantifying drowsy driving (Traffic Safety Facts Research Note. Report No. DOT HS 813 003).

[19] Thomas, F. D., Darrah, J., Graham, L., Berning, A., Blomberg, R., Finstad, K., Griggs, C., Crandall, M., Schulman, C. Kozar, R., Lai, J., Mohr, N., Chenoweth, J., Cunningham, K., Babu, K., Dorfman, J., Van Heukelom, J., Ehsani, J., Fell, J. & Moore, C. (2022, December). Drug prevalence among seriously or fatally injured road users (Report No. DOT HS 813 399). National Highway Traffic Safety Administration. NHTSA

[20] European Commission (2022) Vehicle General Safety Regulation. Available from: https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:32019R2144.

[21] European new car assessment programme (euroNCAP) assessment protocol – Safety assist. Safe Driving – Implementation 2023. Available from: https://cdn.euroncap.com/media/75554/euro-ncap-assessment-protocol-sa-safe-driving-v1011.pdf

[22] Tivesten, E., Victor, T. W., Gustavsson, P., Johansson, J., & Ljung Aust, M. (2019). Out-of-the-loop crash prediction: the automation expectation mismatch (AEM) algorithm. IET Intelligent Transport Systems, 13(8), 1231-1240.