Technical Paper Series - Intoxication: Part 3

Turning integrated awareness into appropriate action

Brook Shiferaw, PhD • Kyle Wilson, PhD • Mike Lenne, PhD

Introduction

When Driver Monitoring Systems (DMS) were first introduced, their primary target was a well-understood safety problem: drivers falling asleep at the wheel. This was followed by the similarly well-defined problem of visual distraction.

Early implementations of DMS solutions therefore relied on specific and readily interpretable behavioural markers, such as head position (e.g. involuntary head drops associated with fatigue), ocular state (e.g. eye closure characteristics), and gaze direction. Given the specificity and clarity of such behaviours, intervention strategies could be equally targeted, making it easier for drivers to understand the cause of an alert and recognise how to self-correct[1].

While this approach simplifies both detection and intervention, such behaviours provide only a partial view of many safety-relevant driver states, which often present in more nuanced patterns of behaviour.

As interior sensing capabilities continue to evolve, they are increasingly enabling inference of driver state from more subtle and complex patterns of behaviour. As a result, assessment is no longer limited to specific behavioural indicators such as eye closure and glance behaviour, but can increasingly infer broader driver states including drowsiness, fatigue, cognitive distraction, and intoxication-related impairment[2,3].

Similarly, other Advanced Driver Assistance Systems (ADAS) address relatively discrete aspects of driving, with each function detecting a specific condition and applying a correspondingly specific response. Lane-departure warning, for example, detects when the vehicle crossed a lane boundary and alerts the driver, while automatic emergency braking responds to an imminent collision by applying the brakes.

As sensing and vehicle-control capabilities have advanced, ADAS equipped with visual sensing can draw on information about lane geometry, surrounding road users, vehicle trajectories, and closing distances to build a more complete picture of the driving situation[4].

Together, these developments provide increasingly rich and complementary forms of perception, allowing simultaneous observation of driver state and external context.

While the importance of interpreting driver state in the context of the driving situation has been recognised for well over a decade[5,6], only recently have advances in sensing technologies and embedded algorithms begun to make such assessments practical in real-world systems. Such assessments may be especially valuable for managing risks associated with complex impaired states such as intoxication.

While this approach simplifies both detection and intervention, such behaviours provide only a partial view of many safety-relevant driver states, which often present in more nuanced patterns of behaviour.

Rethinking intervention strategy for complex impaired states

Previous papers in this series provide an important basis for understanding why such integrated assessment may be necessary.

- **Part 1** argued that Blood Alcohol Concentration (BAC)is insufficient as a real-time ground truth because measured concentration and functional impairment do not necessarily follow the same temporal course.
- **Part 2** further distinguished intoxication from impairment, arguing that detecting functional impairment is more closely aligned with safety-relevant changes in capability and associated risk than detecting whether a chemical concentration has crossed a static threshold.

Together, these arguments suggest that neither measured concentration nor presumed cause can individually determine the most appropriate intervention: drivers with the same measured BAC may present different degrees or patterns of impairment, while similar behavioural manifestations may arise from different or interacting causes.

One reason intoxication presents a particularly difficult intervention problem is that it does not necessarily present as a single, coherent driver state.

Depending on the intoxicating substance (e.g. alcohol, sedative or stimulative drugs), the individual, and the circumstances, observable behaviour may resemble drowsiness, distraction, impaired coordination, heightened activity, unstable attention, or some combination of these and other effects.

These behavioural manifestations may also change across the course of a journey and in response to differing driving demands. Recent naturalistic driving research similarly suggests that alcohol- and cannabis-induced impairment may present through context-specific behavioural patterns rather than as a uniform decline in driving ability[7].

Such variability means that the significance of an observed behavioural change cannot always be understood in isolation, reinforcing the value of interpreting driver state in relation to the surrounding driving situation.

For intervention purposes, the more immediate question may therefore be not simply *what has caused the driver's state?*, but *what does the observed state mean for the driver's ability to manage the situation, and what system response is most appropriate?*

Identifying a likely intoxicant may help explain why behaviour has changed, but it does not necessarily establish the extent of functional impairment or the most appropriate action under the prevailing circumstances. A more useful basis for intervention may lie in how impairment presents, how it evolves over time, and the risks it creates within the current driving context.

One reason intoxication presents a particularly difficult intervention problem is that it does not necessarily present as a single, coherent driver state.

Magnitude and temporal evolution

The observable presentation of impairment can be characterised along two dimensions:

- the degree of change in driving-relevant behavioural functions at a given moment (its **magnitude**), and
- how those changes evolve over time (its **temporal evolution**), whether they recover, persist, or worsen.

In this framing, the transient and non-transient states introduced in Part 2 may be understood as points along a broader continuum of persistence and recovery.

Repeated lapses in attentional engagement, increasing duration, incomplete recovery, or failure to respond to an initial prompt may signal a transition towards a non-transient state. Recognising this transition is particularly important for intervention, as the driver's capacity for effective self-correction may be diminishing before the impairment becomes severe.

The influence of magnitude and temporal evolution on intervention depends on both the presentation of impairment and the risks posed by the driving situation. For instance, safety-critical situations such as sudden and complete unresponsiveness may justify immediate protective intervention without first establishing temporal progression of impairment.

At lower magnitudes, however, recognising when the observed state transitions from transient to persistent may help determine whether driver-led recovery remains viable and whether a warning alone is sufficient.

Contextual inference derived from external sensing and vehicle dynamics can then inform the urgency of escalation and how protective action can be carried out safely.

Recognising this transition is particularly important for intervention, as the driver's capacity for effective self-correction may be diminishing before the impairment becomes severe.

The opportunity for driver response

Earlier recognition of behavioural change may allow intervention to begin while the driver retains sufficient capability and time to respond.

Where no immediate external hazard is present, the initial response should generally inform the driver, communicate the expected action, and provide an opportunity for self-correction. This is particularly important in systems where the driver remains legally responsible for operating the vehicle.

The driver's response then provides further evidence: stable recovery may support de-escalation, while incomplete or absent recovery may indicate that a stronger response is required.

When driver state information influences how other ADAS functions behave, the need for driver communication depends on the nature of the adaptation.

Some changes may occur entirely within the system and have no direct effect on the driver's experience. Others may materially alter vehicle behaviour through earlier braking, stronger steering support, more conservative following behaviour, or other protective interventions.

While such adaptations may improve safety, unannounced changes in vehicle behaviour could leave the driver unaware that the vehicle is compensating for apparent impairment. This may make improved vehicle performance appear to reflect driver recovery when it instead reflects increased system assistance.

Not every internal parameter adjustment requires explicit communication, and urgent protective action should not be delayed while awaiting acknowledgement. However, where circumstances permit, drivers should ordinarily be informed before DMS-informed adaptation materially changes the timing, assertiveness, or authority of vehicle intervention.

Conclusion

In this three-part series on intoxication, we have illustrated that chemical concentration alone cannot fully characterise real-time driving risk.

The same measured level of intoxication can produce markedly different patterns of impairment and risk, while similar impairments can arise from different substances, combinations of substances, and individual circumstances. Chemical thresholds therefore give only a partial indication of the threat posed by an impaired driver.

A more meaningful assessment comes from how impairment is expressed in the present moment and what risk it creates under the current driving conditions, and this reframing changes how intervention decisions should be made.

The continued development of DMS and ADAS is providing vehicles with significantly enhanced forms of awareness: DMS observes changes in the driver's behavioural state, while ADAS establishes the demands and risks of the driving environment. Considered together, these capabilities provide a richer understanding of both the driver and the situation in which impairment occurs.

The significance of an observed impairment depends not only on how it presents, but also on the surrounding driving context. Early signs may warrant a gentle response on an empty rural highway yet a more assertive one in dense urban traffic.

Such awareness can help determine when driver-led recovery remains viable, when additional support may be beneficial, and when protective intervention becomes necessary. This allows responses to be matched more closely to the level of risk being encountered, helping ensure intervention remains aligned with the demands of the situation.

The result is a more cooperative relationship between driver and vehicle. The driver remains informed and engaged for as long as meaningful participation in managing risk is possible; as impairment progresses, the vehicle assumes greater responsibility and, when necessary, acts protectively on the driver's behalf.

Intervention then becomes not a single reaction to detected impairment, but a coordinated response informed by continuous awareness of both the driver and the driving environment, turning integrated awareness into action that is timely, proportionate, and appropriate to the situation.

The significance of an observed impairment depends not only on how it presents, but also on the surrounding driving context.

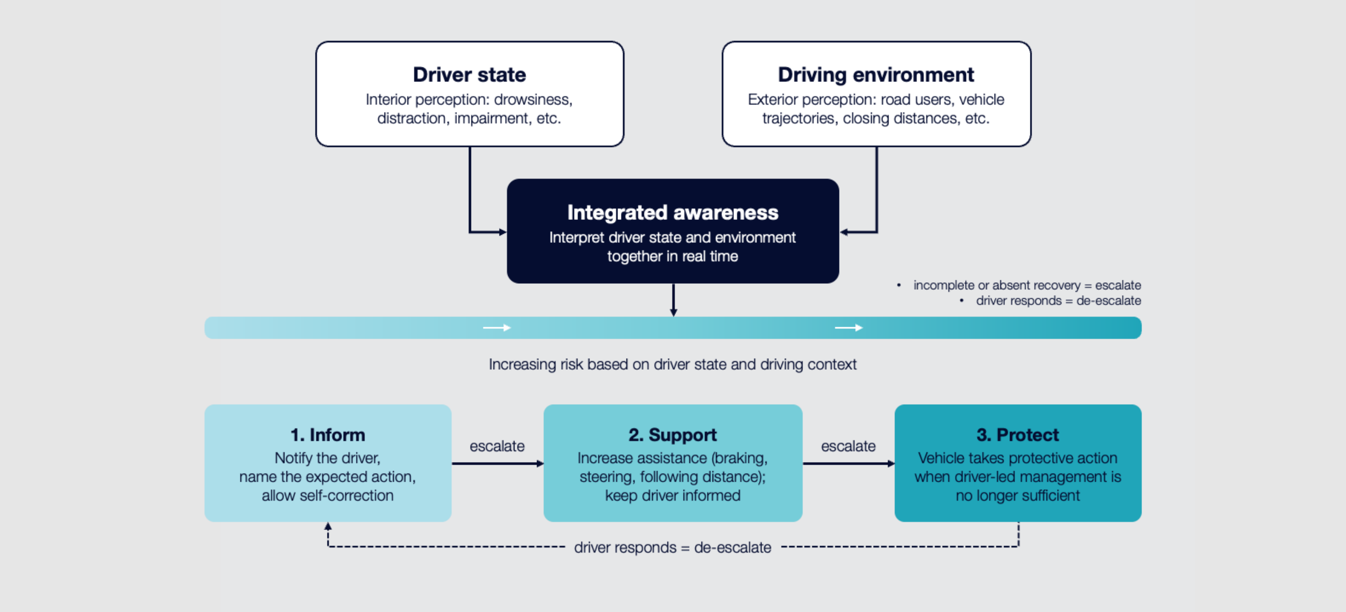

Figure 1: Matching system response to how and where impairment occurs.

References

1. McKerral, A., Mulhall, M. D., Gonçalves, R. C., Harms, I., Merat, N., Tuckwell, G., ... & Wilson, K. (2025). HMI design in the context of DMS and automation: emerging themes from an expert workshop. *Cognition, Technology & Work, **27***(1), 181-192.

2. Shiferaw, B. A., Crewther, D. P., & Downey, L. A. (2019). Gaze entropy measures detect alcohol-induced driver impairment. *Drug and alcohol dependence, **204***, 107519.

3. Manning, B., Hayley, A. C., Catchlove, S., Shiferaw, B., Stough, C., & Downey, L. A. (2023). Effect of CannEpil® on simulated driving performance and co-monitoring of ocular activity: A randomised controlled trial. *Journal of psychopharmacology, **37***(5), 472-483.

4. Yang, K., Al Haddad, C., Alam, R., Brijs, T., & Antoniou, C. (2024). Adaptive intervention algorithms for advanced driver assistance systems. *Safety, **10***(1), 10.

5. Coughlin, J. F., Reimer, B., & Mehler, B. (2011). Monitoring, managing, and motivating driver safety and well-being. *IEEE Pervasive Computing, **10***(3), 14-21.

6. Lenne, M. G., Fitzharris, M., Fredriksson, R., Palao, A., & Young, M. (2024). Occupant status monitoring: A framework to integrate human factors, technology & policy to realise injury reductions across the safety chain. In *International Conference on Driver Distraction and Inattention 2024* (pp. 112-115). Chalmers University of Technology.

7. Jain, S., Terranova, P., Bedwell, K., & Perez, M. (2026). Real-world driving behavior, performance, and risk following alcohol and cannabis use in habitual cannabis users. *Accident Analysis & Prevention*, ***236***, 108693.

seeingmachines