ZENZIC
SELF-DRIVING REVOLUTION

A HORIBA COMPANY MIRA

United Kingdom

Embedding Safety Case: Next Steps for Trialling

A report by Zenzic in association with HORIBA MIRA

Centre for Connected & Autonomous Vehicles

zenzic.io

Executive summary

This report builds upon previously released guidance, examples and templates relating to the safety assurance of connected and automated mobility (CAM) trials taking place upon CAM Testbed UK that is aimed at supporting both creators and reviewers of trial safety cases. As part of this project, a series of workshops were held that had a dual aim: to disseminate information from the detailed safety case guidance, and to collect feedback from CAM stakeholders on what could be done to further enhance safety, interoperability or convenience.

The workshop findings were distilled into a set of recommendations for future work and a set of recommendations that should not be taken forward. A brief summary of these is as follows:

Recommendations for Future Work

Sharing of Incident and Near Miss Data: whilst road collision statistics have been used extensively to inform safety developments relating to manually driven traffic, there is limited data available relating to autonomous vehicles. This should be addressed in two phases:

- A project to provide a 'significant hazards log' based on the best data currently available, including use of expert opinion to extrapolate from data for human drivers;
- A project to plan and implement a scheme to collect incident and near miss data so that, as the mileage driven by automated vehicles increases, the significant hazards log can be updated using data from CAM trials and deployments.

Development of a 'Community of Practice' for Reviewers: this would allow stakeholders who have responsibility for reviewing safety cases to align their expectations and discuss areas where more clarity is needed, improving consistency and therefore interoperability.

Research on the Controllability Provided by Safety Drivers: many CAM trials rely upon a safety driver positioned in the vehicle being able to override system errors using traditional driver controls. However, despite this being key to the safety of trials, little information is available on what expectations can reasonably be placed upon a safety driver. Research should therefore be undertaken to investigate alertness, ability to react to unsafe control inputs, human factors of the controls, and how this learning could be adapted for 'advanced trials' without a conventional safety driver.

Database of Key Features Within Each Testbed: this would help customers, or potential customers, of CAM Testbed UK by making it easy to identify suitable locations to perform the tests they wish to undertake. Such a facility would highlight the variety of scenes that CAM Testbed UK can provide, thereby providing a commercial advantage.

Validating Correlation between Test Environments: CAM Testbed UK incorporates public roads, controlled facilities and virtual test environments. To make the most efficient use of the strengths and weaknesses of each test environment, it is necessary develop a means to conduct an integrated test programme, including establishing processes for comparing results between environments for validation purposes.

Suggestions not Recommended for Future Work

Definition of Terms: whilst the value of an aligned vocabulary is recognised, this is provided by the BSI CAV Vocabulary. Rather than duplicating this work, CAM Testbed UK stakeholders are encouraged to make use of the BSI initiative.

Guidance on the Use of Non-Type-Approved Vehicles: similarly, this overlaps with work that a consortium led by HORIBA MIRA, TRL and WMG are currently undertaking for the UK Department for Transport to develop a regulatory process for low-speed automated vehicles.

Library of Previous Safety Cases: it was determined that commercial and legal sensitivities would make this impractical to implement. Therefore, whilst the value of knowledge transfer is recognised, stakeholders should instead make use of other mechanisms such as the safety case guidance and the proposed initiative to collate and share information on hazards.

Contents

For further information on this report, please contact the Zenzic team at info@zenzic.io

Author
Richard Hillman – HORIBA MIRA

1 Introduction

1.1 CAM Testbed UK

CAM Testbed UK is the centre for the innovation and development of connected and automated mobility (CAM) technologies. It is the only place worldwide with the capability to take ideas from concept to development both virtually and physically, all within a 3-hour drive. CAM Testbed UK is made up of 5 physical testbeds and one data exchange. The UK's comprehensive and integrated facilities are world-leading, with a unique ability to cross-share data and a collaborative way of working.

1.2 Guidance Documents

Zenzic, in conjunction with CAM Testbed UK partners, have produced two separate guidance documents relating to safety cases: one targeted at safety case 'creators', and the other at safety case 'reviewers' (Zenzic, 2021). The aim of these is to support best practice by outlining the key evidence that needs to be considered and providing examples of how this might be achieved. This guidance has been developed to be used in conjunction with existing regulations and standards, as illustrated in Figure 1.

Figure 1: Illustration showing how the Zenzic Safety Case Guidance supports the existing landscape of regulations and standards

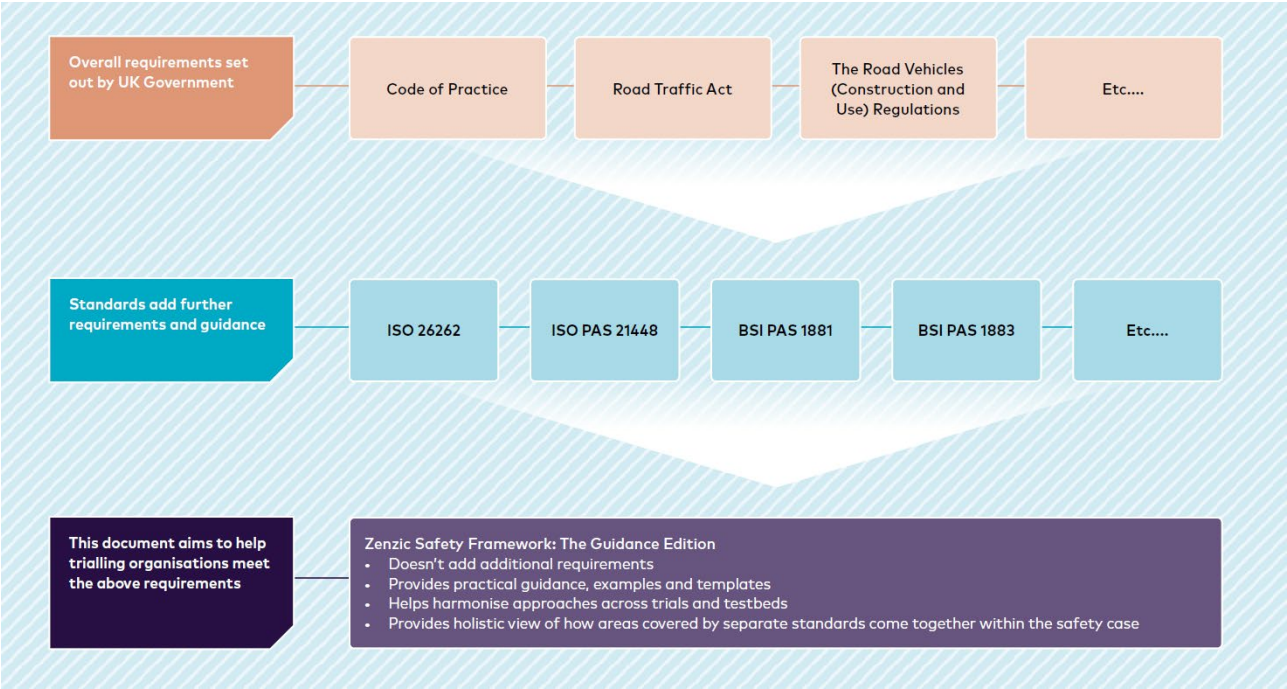

The safety case guidance documents were authored by HORIBA MIRA, TRL and WMG, with HORIBA MIRA holding the lead technical author role. In addition to leveraging the expertise resulting from these organisations being responsible for the safety management in many CAM trials, the project also involved three workshops attended by a wide spectrum of stakeholders,

together with regular meetings with an advisory group. This, combined with an extensive use of existing literature within the domain, ensures that the documents represent a consensus view upon good practice for safe trialling.

1.3 Project Background

Following the release of the guidance documents in March 2021, HORIBA MIRA were commissioned by Zenzic to undertake a follow up project to help embed the guidance within the CAM ecosystem. This involved the creation of further materials to support the safety case guidance documents, in the form of downloadable templates that can optionally be used within CAM trial safety cases, case studies describing how the guidance has been used successfully in the past, a brief 'Non-Technical Explainer' document to provide an overview, and a set of explanatory slides. The latter were derived from a series of four workshops (held in November 2021) delivered as part of the project, which had a dual purpose:

- To explain key concepts relating to the guidance and provide practical activities to aid understanding
- To gain feedback from the participants in order to inform future work relating to safety assurance on CAM Testbed UK

As such, there was a two-way flow of information within the workshops, with feedback being gained through interactive activities involving group discussions and the collation of virtual sticky notes. This report, the final phase within the project, makes use of the stakeholder feedback in order to identify opportunities for future work that could result in improvements to CAM Testbed UK from the perspective of safety, interoperability or convenience.

A summary of the raw findings from these workshops is presented within the Appendix. From this, the strongest themes were extracted and discussed within the project team to identify which ones should be taken forward as recommendations. As a result, a total of five recommendations for future work are set out in Section 2, with a further three suggestions that are not recommended for future work (e.g. due to the risk of duplicating other work that is already underway) set out in Section 3.

Each recommendation commences with a suggested organisation that it is primarily targeted at; this is because some are appropriate to be actioned by Zenzic as part of the development of CAM Testbed UK, whereas others are suitable for government departments such as CCAV or DfT due to a reach beyond CAM Testbed UK.

2 | Recommendations Future Work

2.1 Sharing of Incident and Near Miss Data and Information

Recommendation aimed at: CCAV and/or DfT

Best practice in safety requires the monitoring of systems when in use in order to validate the data, models and assumptions that were used to assure safety prior to the system being deployed. Such in-use data has the potential to highlight many flaws in a safety case, including hazards that have not been identified, risk levels that have been unrealistically scored, failure modes that occur in unpredicted ways, or probability distributions of situations the system will be exposed to varying over time (UL4600, 2020).

Statistics have been collected for existing manually driven traffic via a number of means, such as STATS-19 (DfT, 2021) and RAIDS (DfT, 2013), and indeed many other industries have established processes that aggregate reports or perform investigations relating to safety-related incidents and near-misses (RAIB, 2022; AAIB, 2022), which the RAC Foundation (2021) has looked to incorporate into road collision investigation processes. Monitoring data supports ongoing optimisation of safety within a transport system, such as highlighting the need for improvements to a specific item or piece of infrastructure, or highlighting the need to update regulations, standards or guidance. This allows a transport system to become safer over time.

The problem CAM trials, and future commercial CAM deployments, face is that there is limited data on incidents and near missed relating to CAM technology due to the proportionately very low mileage of driving completed. Furthermore, where such data is available, there is no established means to share it such that trials using CAM Testbed UK, or indeed other public roads and private facilities within the UK, can draw upon past learning gained by other organisations.

To overcome this, it is recommended that a process be developed to facilitate the collation and dissemination of data, and the corresponding knowledge gained from the data, within the UK CAM ecosystem. There are two possible ways that such a process could be implemented: sharing raw data and sharing processed data. These options are examined in the subsequent sections.

Raw Data Sharing Approach

CAM trials typically collect significant volumes of data to allow analysis and optimisation of the system performance; this may include data directly from sensors monitoring the external environment (e.g. camera footage), data from sensors monitoring the vehicle (e.g. wheel speeds, steering wheel angle) and data logs from the system itself (e.g. logs of messages sent over the CAN bus or of system status). The Code of Practice for Automated Vehicle Trialling (CCAV, 2019a) sets out the need for data recording, and some minimum requirements for parameters that should be logged but does not give a detailed specification or attempt to harmonise the data formats used by trials. There are efforts to standardise event data recorders (EDRs) for regular vehicles and data storage systems for automated driving (DSSAD) within the EU and UNECE (Interregs, 2021), but currently there remains wide variation in the data collected.

This highlights a key challenge if raw data is to be shared: the proprietary formats used by each trial will not allow data from one trial to be readily compared with another and will prove problematic to analyse and interpret for anyone outside the organisation responsible for collecting the data. Such large volumes of data will also be problematic to store and share.

Furthermore, there would be significant concerns relating to the intellectual property (IP) of trialling organisations, who have invested large amounts of resource into such technology and do not want information entering the public domain that could benefit their competitors. This could be expected to result in significant opposition to sharing of raw data on incidents becoming mandatory, or in little or no data being shared if it is voluntary.

For these reasons, it is concluded that the mechanism for sharing incident and near miss data should not attempt to share raw data from vehicle systems and sensors.

Processed Information Sharing Approach

The alternative approach is for a centralised body to collate data, including conducting investigations where appropriate, and disseminate this in the format of readily intelligible information rather than as raw data. Alignment on a standardised list of hazards for a trial or deployment type would facilitate mutual understanding between safety case reviewers and developers, supporting interoperability and providing a benchmark for reviewers' expectations. It could also support system developers by highlighting early in the development lifecycle the capabilities that the system must have and ensuring that the test programme provides good coverage of scenarios that present a high risk. CAM Testbed UK could play a significant role in facilitating the collection and dissemination of data and in supporting trials of the processes developed within the projects.

The focus should be upon high-level operational hazards, including hazards resulting from other road user behaviour, hazards resulting from failure of non-ADS aspects of the vehicle (e.g. tyre blowout or lighting failure) and hazards relating to the errors that the ADS could make. However, the latter should be restricted to vehicle-level errors (e.g. system fails to brake when it should, system provides inappropriate steering input) and should not attempt to capture subsystem or component level hazards, these being bespoke to the technology employed within each system and therefore not lending themselves to aggregation.

Despite the paucity of incident data currently available relating to CAM trials and deployments, there is a need to ensure safety in the short term, prior to data from CAM vehicles becoming more widely available in the future. As a result, two separate phases have been identified, one to support CAM operations in the short term by making use of the data that is already available, and a second to support long term safety through a process to acquire and disseminate data that is specific to CAM. It is envisaged that these two phases would be procured separately due to the differing sources, skills and timescales involved.

Phase 1: Significant Hazards Log Using Best Available Current Data

The ultimate output of this work would be a prioritised list of the most significant operational hazards that are presented by CAM trials and early deployments. The identification and subsequent prioritisation of hazards should be based upon two key methods:

- A backward-looking analysis of non-CAM road traffic collision data from a range of relevant sources, such as STATS 19 (DfT, 2021) or RAIDS (DfT, 2013);
- A forward-looking analysis by experts to predict how this baseline data for existing traffic is likely to be affected by CAM technology, such that the hazard log can be adapted as far as possible given the available sources.

The expert analysis is important because the non-CAM data alone would not reflect how the distribution of collision modes may be expected to shift when CAM technology is introduced, due to it 'failing' in a different way to humans, and also because the level of exposure to hazards caused by other road users will be different for limited scale trials or deployments. For example, the hazard of another vehicle reversing on a slip road after taking the wrong exit is far more likely to occur within the lifecycle of the whole transport system than it is to occur in the vicinity of CAM trials operating a limited number of vehicles for a limited duration.

The hazards should be prioritised in terms of the risk that they present such that resources utilised for mitigation are targeted where they will make the most difference. Therefore, the proposed 'significant hazards log' for CAM should include data on prevalence of the hazardous scenarios and severity of the likely outcomes.

Because the prioritisation of hazards will change according to characteristics of the trial or deployment, such as the types of roads used, it is important that the information provided to safety case creators is customised to their needs. It is therefore recommended that a tool should be created that allows users to input key characteristics of their trial, which then influences which hazards they are presented with and what risk prioritisation they have. For example, a hazard relating to identifying a gap in oncoming traffic before overtaking would not apply on a motorway, and a hazard relating to failing to vacate a closed lane upon a managed motorway would not apply to a country lane. Hazard prioritisation could potentially even be linked to the location of trialling, i.e. be specific to a testbed or an area within a testbed.

As such, an initial stage of the project, before data collection starts, should be to identify what parameters should be available to users in order to optimise the data efficiently, and what attributes need to be recorded in the database to support this. It would be possible to implement the tool within software such as Microsoft Excel or Access, but a more interactive solution would be to develop a bespoke web application that allows the user to input the attributes of their trial or deployment in order to receive tailored results. The project should also include the development and demonstration of a suitable methodology by which the significant hazards log can be applied within risk assessments for CAM trials and deployments, thereby supporting industry in making the most effective use of the available information.

It should be noted that National Highways have a 'Top-Level Hazards' log for manually-driven traffic upon their Strategic Road Network (SRN), which has been valuable within CAM trials to support the identification of hazards which would otherwise have been missed (Hillman, 2021). Furthermore, they have recently completed a project to convert this log into one that is specific to CAM trials, thereby reflecting the expected shift in accident mode distributions and in exposure to hazard types described previously. The proposed work should therefore take the National Highways log into consideration, but would still be expected to add significant value over and above it because:

- A wider pool of available data can be used within the 'backward looking' analysis such that the tool covers routes and road types outside the SRN;
- A far larger and more diverse pool of expert opinion can be used in the 'forward-looking' analysis to adapt the hazards according to the characteristics of CAM;
- The data would be publicly available;
- The development and sharing of best-practice methodologies to make effective use of the data would support safety case creators.

Phase 2: Planning and Implementation of a Process to Collect and Share CAM Incident and Near Miss Data

Whilst phase 1 provides a short-term stop gap to support safety within trials, nonetheless it is limited by the data currently available. Whilst this limitation is tolerable given the low mileage being completed by such vehicles, and the lack of a superior alternative, it will cease to be tolerable as and when more significant rollout of CAM technology occurs.

Therefore, a process should be developed to collect data on hazards that are identified for CAM technologies in the field, thereby allowing the significant hazards log to evolve so that the evidence it is based upon becomes progressively more biased towards data on CAM, rather than non-CAM, over time. This project needs to identify:

- What hazard data is desirable to include in the log
- What data it is practicable to collect in the field
- How this data should be collected
- What is the threshold for recording a hazard? (e.g. how "near" does a "near miss" need to be?
- In what format should it be presented
- How data should be classified into categories and/ or assigned metadata to allow users to access a list of the most significant hazards that is customised to the attributes of their trial or deployment?

Whilst this should take into account the data captured within phase 1, it may be determined that deviation from this is desirable due to richer information being available and/or needed. The methodology should consider learnings from existing processes such as the aforesaid ones used by the RAIB and AAIB, and also existing work in the CAM domain such as the DSSAD regulations being developed by UNECE, BSI PAS 1882 (2021), recommendations by the Law Commissions (2022), and the CAV PASS (CCAV, 2019b) programme.

Consideration should be given to how to anonymise the data as far as is practicable, both to protect individuals' privacy in line with the General Data Protection Regulation (GDPR), and also to protect organisations from exposure of their IP and from negative publicity. It should seek to find ways to aggregate data in such a way that trends can be observed and should highlight key learnings in a clear and constructive way in order to ensure lessons are learnt without blame or negative publicity.

By following such an approach, it will be possible for safety mitigations and safety cases to have an evidential basis for the assumptions and decisions that are made, with the level of safety progressively increasing over time as further data is aggregated and disseminated. Furthermore, the data will be able to support the wider rollout of regulations, standards, codes of practice and guidance, providing a feedback loop such that requirements can be better calibrated to practical experience of CAM technology and the particular hazards it can present.

The scope of the project should include working with a range of stakeholders to identify how the process could work, taking consideration of what is desirable and also what is technically, commercially and legally feasible. Once stakeholder agreement, including relevant government departments, has been reached, the project should then move on to implementation of the scheme, including the creation of a suitable database and communication channels, and also engagement with industry to ensure widespread awareness and support.

2.2 Development of a 'Community of Practice' for Reviewers

Recommendation aimed at: Zenzic & CAM Testbed UK

Issues Relating to Consistency of Safety Case Reviews

Within the workshops to support the creation of the safety case guidance (Sept-Oct 2020) and the workshops that formed part of the subsequent project to embed the guidance within the CAM ecosystem (Nov 2021), the need for a means of standardisation and sharing of good practice between testbeds was raised. This would present the following advantages:

- Enhanced consistency in the level of safety and safety oversight expected by the testbeds and other stakeholders, supporting interoperability by ensuring that a safety case for a trial on one testbed can be adapted to suit another testbed without having to make fundamental changes to the style, methodology or level of detail;
- Enhanced sharing of knowledge from previous trials, such as where safety processes have failed in practice or where difficult decisions have been required on "grey areas". This would allow CAM Testbed UK to learn from past experience and develop processes in a collective manner, making better use of the available data to reach informed decisions.

An ambition was raised within some of the workshops to have a central approver acting on behalf of testbeds such that they review and make decisions upon safety evidence. This would, in theory, allow trials to move seamlessly between testbeds, with safety evidence that has been approved once being able to be used on all testbeds. However, following further discussions upon the practicalities within the workshops, this concept is not proposed for further development, due to the following limitations:

- Each testbed has its own safety procedures that must be followed, and have a duty to provide an appropriate level of safety oversight themselves, e.g. to adhere to corporate governance policies or the requirements of insurers. It is therefore not envisaged that they would be able to accept a trial as safe purely on the basis of a prior acceptance by a 3rd party. This will also be the case for many other stakeholders within trials, e.g. highway authorities, local authorities, land owners.

- Each trial has its own unique characteristics in terms of the operational hazards that are presented, the risk those hazards pose, and the resulting mitigations that are put in place. As such, it is important that safety cases are bespoke to the particular trial, and therefore are subject to a bespoke review. This precludes the possibility of a single safety case, or documents that form part of the safety case, being 'carried over' without further review.

- No model has yet been proposed for how a centralised review role would be funded or how a person with the appropriate knowledge, but also the necessary impartiality, would be resourced.

Proposed Solution

Given the impracticability of using a centralised reviewer role, an alternative approach that was suggested, and favoured by many stakeholders, is to set up a 'community of practice' (Wenger-Trayner and Wenger-Trayner, 2015), whereby stakeholders in CAM Testbed UK trials have regular meetings that act as a forum for sharing information.

It is therefore suggested that further work should be undertaken to identify a suitable format for a community of practice and to subsequently implement it. This work should consider:

- What meeting cadence would be appropriate. Excessively frequent meetings could compromise attendance, whereas excessive gaps between meetings could result in delays in responding to problems. One option to overcome these difficulties would be to implement a process whereby significant events trigger special/ extraordinary meetings, thereby allowing the cadence of the regular meetings to be longer.

- What stakeholders should be involved. It could be limited those responsible for safety assurance within the testbeds, or could be extended to include others such as insurers or technology developers.

- What the scope of the meetings should be.

- How the findings should be stored and shared to support future safety case reviews.

- How the findings could feed into wider CAM Testbed UK, Zenzic, BSI, DfT or CCAV documents, processes or projects; for example, by informing future updates to the safety case guidance available on the Zenzic website, or by triggering updates to operational procedures used by testbeds.

- What lessons can be learned from current best practice within the field of knowledge management with regards to forming effective communities of practice

The Zenzic Interoperability Working Group would be a suitable format in which to agree a process and cadence for meetings, but there would be additional project work required to investigate the available options and to put into effect an agreed solution.

2.3 Research on the Controllability Provided by Safety Drivers

Recommendation aimed at: CCAV and/or DfT

Although some trials have used methods such as remote safety operators, or safety operators within the vehicle who only have access to simple controls such as an emergency stop button, trials of CAM technology taking place upon public roads typically make use of a safety driver, i.e. a safety operator who is positioned within the vehicle and has access to the same controls that a regular driver has (Zenzic, 2021).

This can be an effective mitigation where the level of maturity of a prototype autonomous driving system is insufficient for it to be relied upon to perform the task without an additional layer of protection. However, if the safety driver is to be relied upon to be able to avoid or mitigate collisions when the system makes an error or finds itself in a scenario it is not designed to react to, it must be ensured that the safety driver is able make such interventions consistently, correctly and rapidly. Trialling organisations should not delegate responsibility, or indeed blame, to safety drivers unless it can be shown that it is reasonable to expect them to be able to perform the task that is required of them. In order to do this, evidence is needed.

Although the ultimate goal of the industry is to get to a point where systems can be deployed without a safety driver, the need for safety drivers will not cease; it should be borne in mind that new CAM developers, or existing developers working on new iterations of their technology, will still need to test systems whilst they remain in an immature state, prior to being ready for commercial deployment. It is therefore anticipated that the use of safety drivers within development testing will remain an important safety measure in the longer term.

Whilst research has been conducted on driver behaviour when a member of the public is required to take control by an advanced driver assistance system (ADAS) in a production vehicle, equivalent research has not been conducted into the very different demands of a professional safety driver testing prototype technology within a risk-managed project. This report therefore sets out the following areas for future work relating to the ability of a safety driver to control hazards:

Alertness

Within the UK CAM Standards landscape, BSI PAS 1881 (2019) sets high level requirements for ensuring that any reliance upon a safety driver or safety operator results in appropriate safety, and BSI PAS 1884 (2021) takes this further by focusing exclusively on the role undertaken by safety drivers and safety operators.

Clause 4.3 of the latter includes an absolute maximum duration for trialling without a break (2.5 hours) and an absolute minimum length of break (a 15-minute break after a duration of not more than 2 hours; otherwise, a 30-minute break). However, these requirements only provide for a basic level, and are effectively caveated by the requirement for each trial to consider the balance of risk and select appropriate policies for break periods on a case-by-case basis. No data or references are provided to support the trialling organisation in making this judgement –

understandably so, as this is an aspect that has had minimal research with regards to safety drivers within research trials.

Therefore, it is recommended that consideration is given to undertaking research to analyse how safety driver attentiveness can be best assured. This could include secondary research using sources such as studies on members of the public supervising ADAS functions (TRL, 2021) or sources from other industries (e.g. research on train driver attentiveness), but could also include primary research in the form of user trials in a simulator or upon a proving ground. The research should look at what expectations are reasonable (e.g. the duration for which concentration can be adequately maintained), what parameters affect safety driver performance (e.g. whether a system that makes regular mistakes helps the safety driver maintain focus), and what measures can be put in place to detect or reduce inattentiveness (e.g. driver monitoring and alerts, or vigilance devices requiring regular input from train drivers).

Ability to Mitigate Against Unsafe Control Inputs

When an automated vehicle (AV) detects a fault or a situation it is unable to manage safely, the safety driver can be provided with a warning message such that they can take control as soon as possible.

However, if the AV makes an error due to an inherent limitation in its design (or machine learning training), such as mistaking a tar line for a lane boundary, the system will have no knowledge that it is making an error, and therefore no warning will be given to the driver. As such, the driver will have no awareness of the hazardous behaviour until the vehicle has already adopted a perceptibly incorrect path or speed (HumanDrive, 2019).

Each system should have a maximum level at which inputs can be made, whether limited by algorithms that have been through an appropriate safety assurance process, or by the physical capabilities of the actuators. It is therefore possible to assess the safety driver responses to worst-case erroneous inputs such that a model can be built up of what is and is not reasonable to expect of a safety driver. This can then be compared to the characteristics of the trial route, e.g.

- If the available space to the side of the vehicle before a collision would occur is less than the identified deviation before a safety driver manages to correct worst-case steering inputs, it would be unsafe to operate autonomously in this region (See Figure 2);
- If the gap between where a vehicle pulls up at a side road of a junction and the path of traffic passing in front of the vehicle is less than that which would be required to correct a worst-case acceleration input, the safety driver should take manual control pre-emptively when traffic is crossing (e.g. put foot on brake pedal).

Figure 2: Illustration of how the available space to the side of a vehicle, when compared to data on how much the vehicle is expected to deviate before the safety driver corrects a fault, helps determine safety

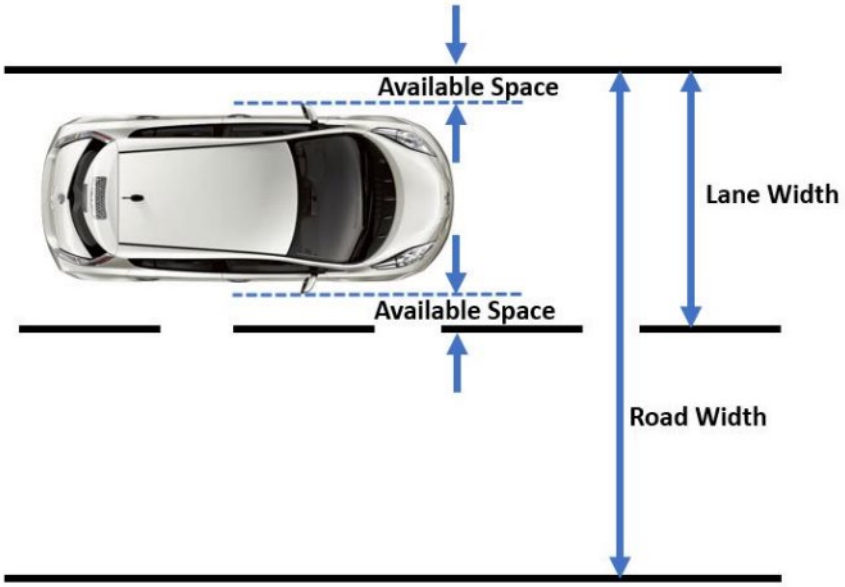

Source: HumanDrive (2019)

This highlights the importance of having accurate and unbiased data available on what can be expected of a safety driver; however, it is an area where there is very little research available in the public domain. Whilst case-by-case consideration of the needs of individual trials would always be necessary, research to better understand how deviations from safe paths and speeds vary as a function of various parameters (such as vehicle speed, steering rate limits or required override force) would provide both safety case creators and reviewers with far more information to make decisions. This information would not only improve safety through analysing a key layer of protection within trials, but would also result in economies of scale and a reduction in the time and cost required for safety assurance through the data being freely available.

Hazardous situations can arise from safety drivers failing to correctly identify when a system has left its operational design domain (ODD); it is therefore recommended that the research into safety driver responses could also include an investigation into how accurate it is reasonable to expect a safety driver's perception to be of the operating conditions that the vehicle is in and how successfully they compare this to the ODD specification in order to make appropriate decisions. This would help inform how ODDs should be specified for trials and what elements of ODD monitoring are appropriate for the driver to assume responsibility for.

The research could involve simulator trials, but it is envisaged that physical trials would be needed at the very least to validate the simulation accuracy and may be determined to be the optimal way to conduct all of the testing. Driver performance could be assessed through fault injection (e.g. inputting maximum braking without warning, and recording how quickly it is overridden) or by deploying hazards (e.g. pedestrian dummy crossing path without warning).

Suitable mitigations should also be considered; for example, an engineer on-board the vehicle observing data on what trajectory the system plans to take or what obstacles have been detected

in the scene could allow them to identify errors and provide early warning to the driver, or a reduction in the maximum operating speed of the vehicle within a particular area of concern could reduce both the likelihood and severity of a collision. It would not be possible to provide precise data that would be universal to all trials – for example, the vehicle dynamics or the level of steering torque the system can provide are factors that would affect controllability – but it would be possible to provide a reasonable benchmark to allow safety case creators and reviewers to make informed decisions.

Suitability of Overrides

One key parameter in the above consideration of what can reasonably be expected of safety drivers is the human factors of the interface that they use. Problems could occur, for example, as a result of:

- The torque required to override steering being too high, making it challenging for the safety driver to push past the resistance;
- The position of override button being difficult to access such that time is lost finding or accessing it;
- Controls being counter-intuitive such that the wrong input could be made by mistake;
- Mode confusion resulting from unclear feedback from the system as to what mode it is in.

It is therefore important that research is undertaken on how to optimise the human-machine interface (HMI) for maximum controllability by the safety driver. It is proposed that this would take the form of a literature study on the human factors for both automotive driver controls and controls used in other industries, and user trials in a simulator or in a controlled environment to gain quantitative data on safety driver performances as parameters relating to the HMI are varied.

Progressing Beyond Safety Drivers

Although safety drivers positioned at regular driver controls within an AV remain a significant means of safety assurance within trials, and indeed may be expected to remain indefinitely as a key safety measure for newly developed production systems in the early stages of accumulating test evidence, nonetheless it must be recognised that there is a strong desire to move beyond this such that advanced trials without a safety driver can take place. Such trials may rely upon other solutions to allow a safety operator to monitor system performance and intervene where necessary, such as remote supervision or the use of simplified in-vehicle controls that can be operated from positions other than a traditional driver's seat.

To support this, it is recommended that the above work to investigate safety driver controllability should also consider what other solutions could be used, what strengths and weaknesses they have with regards to safety and practicability, and whether any of the data collected or synthesised with regard to traditional safety drivers (for example, relating to maintaining alertness) can be extrapolated to other safety operator roles. The work should also consider inherent system limitations that may affect safety such as latency and robustness of wireless communications links, and should examine how MRMs performed by the system may overlap or clash with manual interventions performed by a safety operator.

Where there is insufficient data available to determine the safety of a solution, further recommendations for future work should be made, such that the industry can iteratively improve its understanding of how safety operators can support safety in CAM trials.

2.4 Database of Key Features Within Each Testbed

Recommendation aimed at: Zenzic & CAM Testbed UK

Whereas for traditional automotive systems, a vehicle prototype that has negotiated the same test case 999 times could reasonably still be expected to fail on the 1000th run (e.g. as a result of fatigue), repeating the same test case will provide little value within the safety assurance of advanced driver assistance systems (ADAS) or AVs; instead, the test effort would be more efficiently used sampling other scenarios from the vast range of permutations that the system could experience in service (HumanDrive, 2020). And whereas a human who has performed a right turn across traffic successfully within a driving test could reasonably be expected to use their general intelligence to achieve a similar result upon a different geometry of junction with different roadside furniture present, autonomous systems do not possess equivalent general intelligence; therefore, it cannot be assumed that an autonomous vehicle will adapt to new situations safely.

As such, CAM technology presents an unprecedented challenge in terms of the sheer variety of scenario permutations a vehicle needs to be tested in to provide appropriate coverage of the range of permutations that could be encountered in service.

CAM Testbed UK incorporates a variety of road types, including a wide range of different features such as traffic lights and roundabouts in a wide range of geometries. This variety is a key means by which the testbeds can provide value for customers, allowing them to expose the system to the required diversity of test scenarios. However, whilst the scale of CAM Testbed UK presents an opportunity in terms of the variety it encompasses, it also poses a challenge: out of so many locations, how do trialling organisations find the ideal locations for their testing needs?

Whilst not raised in other workshops, this was a strong theme within the workshop aimed at creators of safety cases, suggesting that a means to support the selection of test locations would be of interest to many technology developers who are customers or potential customers of CAM Testbed UK.

Proposed Solution

It is therefore recommended that a project is undertaken to scope, plan and implement a database that allows stakeholders to input key scene attributes that they are looking for in order to retrieve information on what locations would suit their needs. Such a database would need to capture a significant level of detail to allow meaningful searches to be undertaken. For example, merely identifying a feature as a roundabout would be of minimal use, with the search returning a large number of roundabouts that would appear to just be duplicating each other, whereas in practice there would be subtle differences in each roundabout that would result in it adding unique value to the range of scenes available.

It would therefore be necessary to capture data on attributes such as the geometry (e.g. radius, number of lanes, number of entries/exits), the prevailing speed limit, what range of angles and speeds a test vehicle or other actor would be able to approach from or whether there are obstacles that obscure line of sight. Other supporting information could also be captured such as any safety concerns identified that may restrict the scenarios that could be performed at the location, whether it is possible to use robotic soft target actors at the location, whether there is a digital twin available, or whether there are any facilities provided to support V2X communication or data collection.

The project could consist of the following phases:

- Stakeholder consultation, particularly with customers and potential customers, to understand what attributes of a scene should be captured within the database. This needs a careful balance, as increased data adds utility, but the resources required to collect and process it should not be underestimated. It is therefore important to prioritise the attributes that are investigated.
- A study of the available routes, in collaboration with the testbeds, to identify what features can be found at what geolocation, and what attributes they possess.
- Development of a database that provides an intuitive and accessible front-end such that trialling organisations can retrieve results on appropriate locations as efficiently as possible.

It may prove possible to leverage prior work done on scenario databases within Midlands Future Mobility (WMG's Safety Pool and HORIBA MIRA's Assured CAV) and Smart Mobility Living Lab, but the project should expand beyond individual testbeds such that it covers the entire CAM Testbed UK ecosystem.

Such work will make CAM Testbed UK more attractive to customers as it will show them scenarios that meet their needs, thereby providing a significant market advantage over testing in other locations, whether dedicated test facilities or public roads, that are not a part of CAM Testbed UK. It will also aid interoperability, as the data will allow safety case creators and reviewers to gain a clearer understanding of what the similarities and differences when a trial moves from one testbed to another. Finally, it will aid safety by making hazards (e.g. obscured line of sight) and accident trends for the location visible to all stakeholders.

2.5 Validating Correlation Between Test Environments

Recommendation aimed at: Zenzic/CAM Testbed UK. However, project should link to wider work funded by CCAV.

Advantages of Different Test Environments

The test modalities available within CAM Testbed UK can broadly be placed within three categories:

1. Real-world environments (public roads or private facilities accessible for purposes other than testing);

2. Controlled and semi-controlled environments (e.g. proving ground or other controlled facility);

3. Simulation using models of locations.

Although real-world environments allow perfect realism such that the results are representative of the system performance, the impracticality of collecting sufficient data to validate an AV purely through on-road mileage accumulation (RAND, 2016) means that test programmes would typically be expected to use other test modalities to augment real-world testing.

Controlled environments allow enhanced safety, privacy, the ability to have a high level of control over actors such as vehicles or pedestrians within scenarios, and the ability to collect detailed and precise data. As such, testing within a controlled environment provides significant value within a wider test programme; this is especially so for critical scenarios that are likely to result in collisions, where it is necessary to be able to explore the behaviour of the complete vehicle in a safe and scientific manner.

However, there are disadvantages in such testing, including the time and cost involved in setting up the scenarios and the limitations in the realism and variety that can be accommodated within the scenes; for example, whilst temporary white lines, facades to represent buildings or moveable roadside furniture can be used to replicate real scenes, there will inevitably be compromises in the detail and realism of such features, and other features such as gradients and cambers of the road will typically be restricted to those that are already present at the test location.

Simulation is widely seen as having a key role to play in the development and safety assurance of CAM technology (SaFAD, 2019; HumanDrive, 2020). This is because simulation allows:

- Test cases to be performed relatively quickly and cheaply in comparison to physical tests.
- Testing to be performed prior to a complete vehicle, or even subsystem hardware, being available. This allows control systems to be tested and optimised earlier, saving further time and cost.
- Scenarios that would be unsafe in the real world to be created and tested in perfect safety.
- A high level of control over variables, resulting in excellent repeatability and the flexibility to allow attributes to be parameterised to suit the sampling methodology/ experimental design.
- Rich data to be collected such that results can be analysed in detail, against a known ground truth.
- Previously performed scenarios to be repeated and extrapolated, e.g. to understand what would have happened if the safety driver had not intervened during a physical test, or what effect parameterising one or more variables would have.

Challenges Faced

Naturally, this makes it tempting to maximise the use of simulation within a CAM test programme. However, it must be considered that the value of the simulation results depends upon the

accuracy of the simulation itself, including the models used within it. This is problematic as modelling and simulation inevitably requires assumptions and simplifications to be made to manage the computational load and the time and cost involved in model creations. As such, there will always be imperfection within aspects such as:

- Dimensional and visual accuracy of models of physical locations;
- Physics models used to replicate sensor detection;
- Environmental conditions such as rainfall and lighting effects;
- Vehicle dynamics;
- The physics of actuators.

If the benefits of all three test modalities are to be realised, it is vital to have a means to compare results obtained in one modality to those obtained in another modality such that their similarity can be validated.

Opportunities

Such validation could, for example, be achieved by selecting samples from within a set of simulation test scenarios that have been performed and replicating some of the scenarios as closely as possible upon a controlled environment. Another approach would be to take real world test results, replicate them as closely as possible in simulation to validate the simulation's accuracy, then extrapolate around them by adjusting parameters in the simulation to expand the coverage of scenario permutations.

A summary of a method to integrate all three test environments is shown in Figure 3, where a proportion of the test cases in a controlled environment are validated against identical scenarios from real-world testing, but the controlled environment testing extrapolates beyond the real-world scenarios (e.g. to include rare edge cases or emergency scenarios). The simulation testing is then validated against some scenarios that also occur in the real world, some that occur in a controlled environment, and some that occur in both. Using the confidence gained from this validation, the simulation can also be extrapolated beyond the validated test cases to further expand test coverage.

Figure 3: Illustration of how validation can be performed between test environments, and how the range of test scenarios can be extrapolated beyond those used for validation

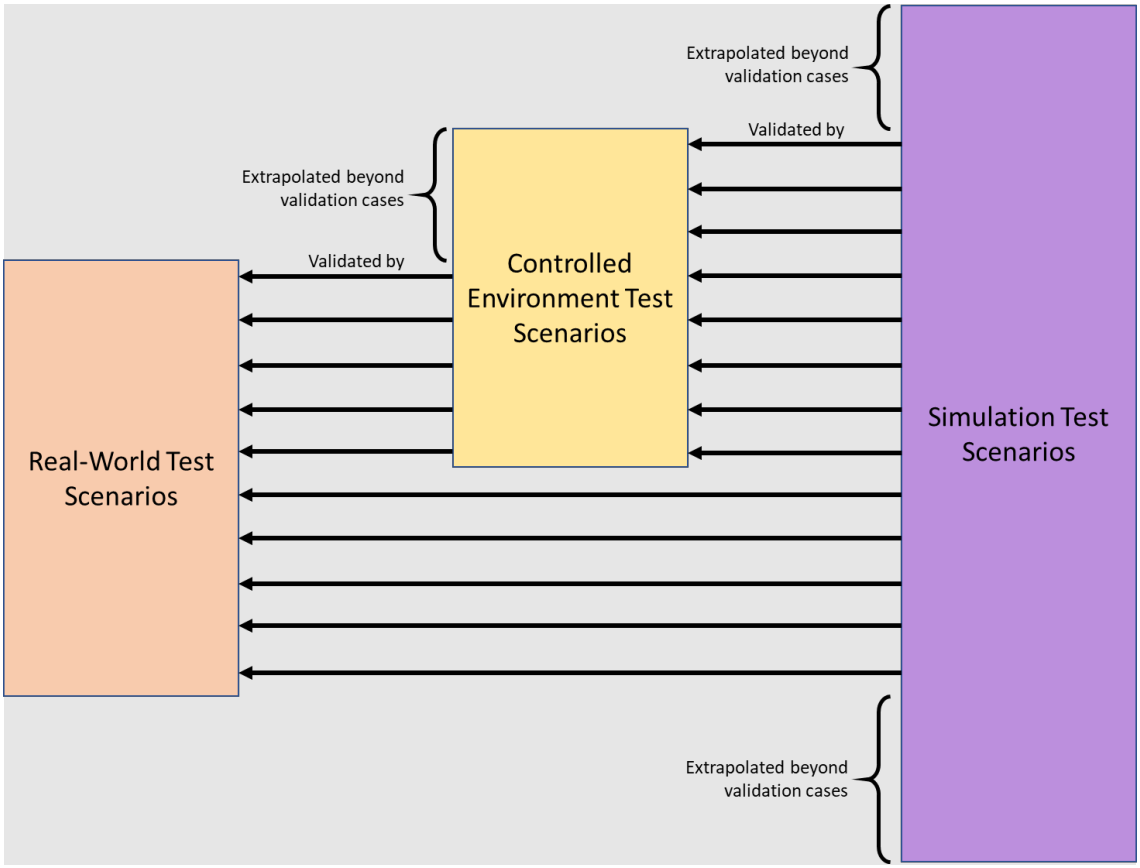

Source: Author generated

Although research projects have considered how these modalities could be combined within a test programme (HumanDrive, 2020) such that the strengths of each can be leveraged, further research is needed in order to develop and demonstrate a process that is practicable, robust and scaleable. CAM Testbed UK holds a particular advantage as it features all three test modalities, and therefore provides an ideal opportunity to undertake further research into incorporating multiple test environments within an assurance programme. It would also be possible to examine the correlation between test results within the same type of environment, such as comparing results from identical scenarios performed in different physical testbeds to identify how the subtle differences between two superficially similar road layouts can have a significant affect upon vehicle behaviour.

This work would not only be of value to the industry as a whole; it would also increase the value that CAM Testbed UK is able to offer to potential customers. By being able to offer testing that uses all three modalities in an integrated manner, customers will be able to gain coverage of a wider range of scenarios whilst maintaining confidence that the results are representative, have the ability to manage the safety effectively, and benefit from enhanced efficiency in terms of time and cost. It is therefore recommended that research should be funded to use CAM Testbed UK to develop and demonstrate processes to integrate the different modalities. This would build directly upon recent Zenzic projects such as Interoperable Simulation, ALKS and the Safety Case

Framework, and could in turn feed into future work; for example, the findings could inform future updates to the physical testbeds by highlighting what features or capabilities are needed to optimise them for use in validating simulations.

3 | Suggestions Not Recommended for Future Work

3.1 Definition of Terms

Multiple participants within the workshops expressed a desire for a glossary providing a standardised definition of terms to be introduced, on the basis that this would facilitate a common understanding between stakeholders, thereby supporting discussions relating to safety. Whilst this is a valid idea, it is not recommended for further work as this need is already being provided for via a vocabulary document maintained on an ongoing, flexible basis by BSI (BSI Flex 1890, 2020).

Although there may be some instances where new terminology is needed that is not yet covered within the BSI document, there is a process to allow ongoing chances such that it can stay up to date within the rapidly changing landscape, and the document has wide involvement from stakeholders within the UK industry such that it is able to reflect the state of the art. It is therefore recommended that, rather than initiating a CAM Testbed UK equivalent to the BSI document, CAM Testbed UK should instead encourage testbed stakeholders to engage with the ongoing BSI work such that the industry can align upon a single set of terms and definitions.

3.2 Guidance on the Use of Non-Type-Approved Vehicles

Although many CAM trials make use of adapted production vehicles that have been through the type approval process, many other CAM trials use dedicated vehicles that have not been type approved, such as low speed 'pods'. Consequently, there were suggestions raised within the workshops that further work should be done to develop guidance, standards or regulations relating to the safety, security and emissions of bespoke CAM prototype vehicles in order to ensure that the physical vehicle has a level of safety that is in line with type approved production vehicles.

Whilst it is felt that this additional assurance for vehicles that have not been type approved would add value, it must be viewed in the context of a current project in which the author is currently closely involved. This project is led by the UK Department for Transport (DfT), who have established a consortium to develop regulatory requirements and processes for the assurance of safety and security of low-speed automated vehicles, with technical work led by HORIBA MIRA, TRL and WMG. In particular, work package 4 of this project is examining regulatory requirements relating to the non-ADS (autonomous driving system) aspects of the vehicle such as seating requirements or pedestrian protection.

As this work is aimed specifically at low-speed automated vehicles, the outcomes from the DfT project will be directly applicable to most, if not all, bespoke automated 'pod' vehicles. Furthermore, DfT plan for the regulations to be further developed in future such that they cover a wider range of CAM solutions. It is therefore felt that it would not be appropriate to initiate a CAM Testbed UK project related to the non-ADS safety of vehicles that have not been type approved, as this would involve significant duplication of the work being undertaken by DfT.

3.3 Library of Previous Safety Cases

There was some desire expressed for CAM Testbed UK to compile a library of previous safety cases that can be accessed by trialling organisations in order to provide exemplars that support the creation of new safety cases. This could also provide benefits for reviewers, who could use the exemplars as a benchmark to help determine the expectations they set for safety cases.

In practice, however, such a library will be difficult to compile due to commercial sensitivities relating to safety cases. Trialling organisations, and potentially testbeds who have reviewed safety cases and accepted that trials can proceed, are often very cautious about what information is made available due to the risk that acceptance of a certain level of risk could be presented in an unfavourable way within the media or could be used against them in court should a serious incident occur. Furthermore, many safety case stakeholders are, quite naturally, concerned about sharing their intellectual property; for example, CAM technology developers will not want proprietary information relating to their technical solutions or to their level of technological maturity being made available to their competitors, and similarly testbeds may wish to protect their commercial interests by not exposing information about their processes.

One option to avoid such sensitivity would be to create entirely fictitious safety cases to act as exemplars. However, it would be challenging to make such safety cases truly representative of the detail and complexity in a real one, resulting in the likelihood that such exemplars would be overly artificial, and therefore of minimal value. Furthermore, there could be liability concerns for those responsible for creating the exemplars should an incident occur that could be traced back to safety decisions in a real trial that were based upon the fictitious example.

It must also be borne in mind that there are other initiatives that will enable learnings from past trials to be shared, such as:

- The proposal presented in this report for further work to be undertaken to collate and share information relating to hazards and risks presented by CAM trials and commercial deployments (see Section 2.1).
- Template documents and high-level 'case studies', which have already been made available upon the Zenzic website (Zenzic, 2021). These are further augmented by the extensive use of examples within the safety case framework documents themselves.
- Publicly available safety case summaries are available for many trials, as per guidance within the Code of Practice for Automated Vehicle Trialling (CCAV, 2019a). Examples include HumanDrive (2019) and StreetWise (2019).
- BSI PAS 1881 (2020) includes normative guidance on what content should be included within a safety case, and also has appendices that provide informative guidance.

As a result, it is not recommended that creation of a library of past safety cases should be an area for future work.

4 | References

AAIB (2022) *Air Accidents Investigation Branch Website*. Last accessed 19th January 2022. Available at: https://www.gov.uk/government/organisations/air-accidents-investigation-branch

BSI Flex 1890 (2020) *BSI Connected and automated vehicles – Vocabulary*. Available at: https://www.bsigroup.com/en-GB/CAV/cav-vocabulary/

BSI PAS 1881 (2020) *PAS 1881:2020 Assuring the safety of automated vehicle trials and testing – Specification*, available at: https://www.bsigroup.com/en-GB/CAV/pas-1881/

BSI PAS 1882 (2021) *PAS 1882:2021 Data collection and management for automated vehicle trials for the purpose of incident investigation – Specification*, available at: https://www.bsigroup.com/en-GB/CAV/pas-1882/

BSI PAS 1884 (2021) *PAS 1881:2020 Safety operators in automated vehicle trials and testing – Guide*, available at: https://www.bsigroup.com/en-GB/CAV/pas-1881/

CCAV (2019a) *Code of Practice: Automated Vehicle Trialling*, Centre for Connected and Autonomous Vehicles, Published 6th February 2019. Available at: https://www.gov.uk/government/publications/trialling-automated-vehicle-technologies-in-public

CCAV (2019b) *New system to ensure safety of self-driving vehicles ahead of their sale*, accessible at: https://www.gov.uk/government/news/new-system-to-ensure-safety-of-self-driving-vehicles-ahead-of-their-sale

DfT (2013) *Road Accident In-Depth Studies (RAIDS)*. available at: https://www.gov.uk/government/publications/road-accident-investigation-road-accident-in-depth-studies/road-accident-in-depth-studies-raids

DfT (2021) *STATS19 Forms and Guidance*. Available at: https://www.gov.uk/government/publications/stats19-forms-and-guidance

Hillman, R. (2021) *An approach to managing the operational safety of autonomous vehicle trials, Journal of Safety and Reliability*, Oxford: Taylor and Francis. Available at: https://www.tandfonline.com/doi/abs/10.1080/09617353.2021.1920300

HumanDrive (2019) *HumanDrive: Autonomous Vehicle Project Safety Management*. Available at: https://humandrive.co.uk/downloads/

HumanDrive (2020) *Test Methods for Interrogating Autonomous Vehicle Behaviour – Findings from the HumanDrive Project*. Available at: https://humandrive.co.uk/downloads/

Interregs (2021) *EU and UN ECE Develop New Regulations on Event Data Recorders*. Available at: https://www.interregs.com/articles/spotlight/eu-and-un-ece-develop-new-regulations-on-event-data-recorders-000235

Law Commissions (2022) Automated Vehicles: Joint Report, the Law Commission and the Scottish Law Commission. Available at: https://www.lawcom.gov.uk/project/automated-vehicles/

RAC Foundation (2021) *Road Collision Investigation Project*. Available at: https://www.racfoundation.org/collaborations/road-collision-investigation-project

RAIB (2022) *Rail Accident Investigation Branch Website*. Last Accessed 19th January 2022. Available at: https://www.gov.uk/government/organisations/rail-accident-investigation-branch

RAND (2016) *Driving to Safety – How Many Miles Would it Take to Demonstrate Autonomous Vehicle Reliability*, RAND Corporation. Available at: https://www.rand.org/content/dam/rand/pubs/research_reports/RR1400/RR1478/RAND_RR1478.pdf

SaFAD (2019) *Safety First for Automated Driving*. Available at: https://www.daimler.com/innovation/case/autonomous/safety-first-for-automated-driving-2.html

StreetWise (2019) *StreetWise - Abridged Safety Case for public road trials of Automated Vehicles in London*. Available at: https://www.trl.co.uk/Uploads/TRL/Documents/StreetWise---Abridged-Safety-Case-for-public-road-trials-of-Automated-Vehicles-in-London_2.pdf

TRL (2021) *Driver Availability Monitoring Systems*, Brussels: European Commission. Available at: https://trl.co.uk/uploads/trl/documents/MIS070-Technical-study-for-General-Safety-Regulation,-driver-assistance-monitoring-systems-DAMS.pdf

UL 4600 (2020) *Standard for Evaluation of Autonomous Products*, Underwriters Laboratories. Available at: https://www.shopulstandards.com/ProductDetail.aspx?productid=UL4600

Wenger-Trayner, B. and Wenger-Trayner, E. (2015) Communities of Practice: a Brief Introduction. Available at: https://wenger-trayner.com/introduction-to-communities-of-practice/

Zenzic (2021) *Safety Case Framework: The Guidance Edition*. Available at: https://zenzic.io/projects-and-resources/safety-case-framework/

5 | Appendix – Summary of Feedback from Workshops

- The use of operational databases within the rail industry to share best practice and information on accidents/ hazards was raised by some participants, and it was agreed that such an approach would be preferable to organisations operating in isolation. It was envisaged that this would include capturing of accidents and near misses such that the data could be anonymised and shared. This is potentially linked to suggestions of a centralised list of hazards for CAV trials to support the hazard identification process, and for a log of scenarios derived from previous incidents/ near misses. The need for sharing of data on incidents and near misses was raised more than any other, and across the different stakeholder groups that the workshops were divided up into, suggesting a strong desire for such a repository. This broadly feel into two categories: sharing of raw data, or processing of raw data to allow a targeted list of hazards to be put forth. The balance between using large amounts of data for manually driven vehicles (RAIDS, Stats-19) and smaller amounts of more relevant CAV data also needs to be considered.

- There were some suggestions that advice on transitioning between testbeds could be helpful, e.g. guidance on what the delta is between the challenges in one testbed and another.

- It was suggested that a reference set of previous safety cases could help new trialling organisations. This could include real and/ or hypothetical safety cases, although the former may be challenging as organisations may be wary of sharing sensitive data.

- One comment proposed shifting to thinking about interoperability on an international basis. This could help make Testbed UK more attractive to those who have tested in other countries.

- There were some concerns about how reliably a safety driver is able to assess whether a system is operating outside its ODD. This could be an area of further study, e.g. an experiment to assess how accurately safety drivers determine ODD compatibility.

- Given that trials typically progress to more complex environments as safety evidence is accumulated, it was suggested that more thought could be put in to processes to ensure the activities within a trial are providing the necessary safety evidence to support subsequent trials.

- The concept of an advisory board to ensure varied perspectives are included in safety case reviews was raised.

- It was suggested that training on the specific skills needed for reviewers could be provided.

- Legal status of testbed reviews – some comments suggested a need for more clarity around the language used relating to reviews (e.g. "endorse", "approve", "accept") and the implications that might exist surrounding legal liability in the event of an incident. This may be an area that justifies further work to understand the implications and agree suitable terminology.

- It was suggested that accident data from systems with lower levels of autonomy could potentially provide valuable data to support trials of higher levels of autonomy, e.g. statistical data relating to accident modes. NRPOI (National Roads Policing Operations and Intelligence) are aiming to start recording data on incidents with automated driving systems engaged, which could support such analysis.

- A centralised reviewer could support interoperability, although it was acknowledged within discussions that it may be difficult to resource someone with the appropriate skills and neutrality, that funding may not be available, and that testbeds would not be able to rely on an outside party and would have to conduct their own assurance for governance, legal and insurance reasons.

- An alternative approach was a 'community of practice' for safety case reviewers, providing an opportunity for relevant persons across the testbeds to meet and discuss cases that could prove/ have proved contentious such that a consistent approach can be adopted.

- It was suggested that further guidance on what level of safety analysis is proportionate would be valuable.

- Further information to support assertions that safety drivers are able to mitigate accidents was desired. BSI PAS 1884 has set out normative requirements for best practice, but there would be value in further research relating to the ability of safety drivers to maintain alertness, methods to ensure the controls available to safety drivers are suitably robust and ergonomic, how well safety drivers are able to react (e.g. quantitative measurement of vehicle deviations when undesired control inputs are injected) and research on approaches to safely test without a safety driver who has conventional driver controls. This could be linked to looking at what testbeds can do to support (e.g. testing with pedestrian dummies). There was concern that CAV trials place a lot of responsibility on safety drivers to mitigate hazards, in a way that wouldn't be done in an industrial safety case; evidence is therefore needed to support viability of safety drivers.

- There were multiple comments relating to providing a definition of terms used in order to support a common understanding within risk assessments and safety cases.

- Guidance on how system safety cases can feed into operational safety cases was suggested.

- Research and guidance was requested on how to make use of simulation within safety cases, including ensuring that results can be trusted.

- It was raised that local environmental effects could affect safety (weather, vegetation etc.). Guidance could be provided on how to identify and manage such issues.

- It was indicated that there would be an appetite amongst CAV developers for testbeds to record interesting operational situations within their routes (e.g. 3 way junctions).

- Guidance on what safety standards are needed for non-type approved vehicles (e.g. pods) would be appreciated by safety case 'creators'.

ZENZIC

SELF-DRIVING REVOLUTION

United Kingdom

zenzic.io