Simulation-Based Autonomous Systems Testing – From Automotive to Defence

Timothy Coley, Dave Fulker
XPI Simulation
Crawley, West Sussex, UK
timothy.coley@xpisimulation.com,
dave.fulker@xpisimulation.com

Rob McConachie
Thales UK
Crawley, West Sussex, UK
rob.mcconachie@uk.thalesgroup.com

ABSTRACT

The introduction of autonomous systems into safety-critical domains poses many challenges – from system safety assurance to public acceptance of the use of such systems. Recent autonomous vehicle incidents have underlined the need to establish clear regulations for the introduction of such systems into public environments, while the challenge of conducting adequate testing in the real world makes simulation-based testing an attractive proposition – reducing time, cost and risk associated with development and evaluation.

This paper elaborates on an XPI-led feasibility study examining the use of simulation for the certification of autonomous vehicles (CAVinSE), co-funded by the UK Centre for Connected and Autonomous Vehicles, and draw parallels with accrediting autonomous systems for defence applications. This includes consideration of the emerging standards, regulations and methodologies in the automotive sector with regard to autonomous systems. Moreover, it highlights the importance of establishing a trusted environment for autonomous systems testing activity – with a root of trust that can be relied upon by certifying bodies, system developers and the general public. Potential approaches for accrediting simulation tools associated with such autonomous systems testing are also addressed.

The paper considers how the different operating environment and capability of defence systems, as well as different regulatory frameworks, could nonetheless benefit from the outputs of CAVinSE and similar endeavours in the automotive domain. As well as XPI's work on CAVinSE, a use case for simulation-based autonomous systems testing in the maritime domain is covered.

In concluding, some of the key technical challenges that remain in this domain are identified – with a particular focus upon sensor models, representation of the physical environment and adequately representing human behaviours in simulation.

ABOUT THE AUTHORS

Timothy Coley has held the role of Product Specialist at XPI Simulation since May 2017. He is responsible for defining the company's product roadmap and ensuring its alignment to customer requirements. Timothy plays a key role in determining and executing company strategy; he is also responsible for coordinating XPI's research activities in autonomous transport solutions

Dave Fulker is the technical director of XPI Simulation and is responsible for the technical oversight of all projects within the company. Dave has worked in the simulation industry for over 25 years, starting as a systems engineer with GEC Marconi, moving up to R&D manager with Primary Image Ltd. and then co-founding XPI simulation in 2005.

Rob McConachie is Product Line Manager for UK Land training and simulation at Thales UK. He has previously enjoyed a number of systems- and software- engineering roles in the real-time simulation domain for Thales UK, McLaren Applied Technologies, and Sonda Aviation Enterprises.

Simulation-Based Autonomous Systems Testing – From Automotive to Defence

Timothy Coley, Dave Fulker
XPI Simulation
Crawley, West Sussex, UK
timothy.coley@xpisimulation.com, dave.fulker@xpisimulation.com

Rob McConachie
Thales UK
Crawley, West Sussex, UK
rob.mcconachie@uk.thalesgroup.com

INTRODUCTION

The development of autonomous systems in both civil and military domains is a growing matter of concern to policymakers, enterprises and the public alike. The technical hurdles associated with such systems notwithstanding, a considerable regulatory challenge lies in the ability to demonstrate and certify that autonomous platforms meet acceptable standards of safety.

In one domain – that of connected autonomous vehicles (CAV), also referred to as self-driving cars – the RAND Corporation has published research stating that, in order to demonstrate the same level of safety as human drivers – i.e. a fatality rate of 1.09 deaths per 100 million miles – that the vehicles would have to be driven 8.8 billion miles (assuming a 95% confidence level and a precision of 20% - see Figure 1).

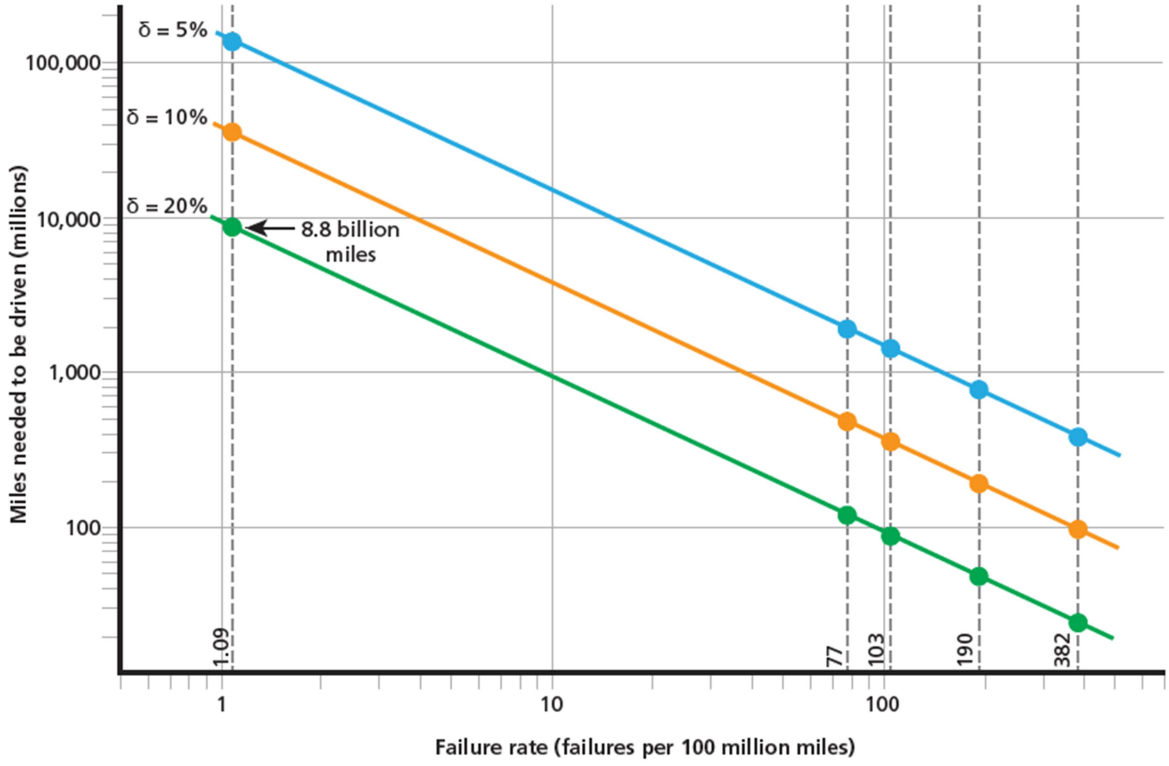

SOURCE: Authors' analysis.
NOTE: These results use a 95% CI. The three colored lines show results for different levels of precision δ, defined as the size of the CI as a percent of the failure rate estimate. The five dashed vertical reference lines indicate the failure rates of human drivers in terms of fatalities (1.09), reported injuries (77), estimated total injuries (103), reported crashes (190), and estimated total crashes (382).
RAND *RR1478-2*

Figure 1: Miles Needed to Demonstrate Failure Rate to a Particular Degree of Precision (Kalra and Paddock, 2016)

It is self-evident that such a requirement to accumulate CAV mileage would place an impossible burden upon technology developers in the domain – even if large fleets were deployed for 24 hours a day, 365 days a year in pursuit of this goal. Moreover, the crude yardstick of miles travelled provides no consideration for the type of miles travelled – or the road layouts, signage, markings, junctions, roundabouts, vehicles, pedestrians, cyclists, animals and inanimate objects that are encountered in the course of such mileage. Driving along 8.8 billion miles of well-maintained, clearly marked and spacious motorway will not demonstrate that a CAV is equipped to navigate the cluttered, congested, chaotic and confounding road environments in developed cities – let alone the bedlam of developing metropolises.

Moreover, failures in complex systems are typically manifested through the incidence of a confluence of factors – such as a combination of weather, external entity behaviour and a sensor at the edge of its operating envelope – rather than a single failure. These corner cases (i.e. the points at which multiple edge cases meet) are rarely encountered in real-world testing, or require such control over the test conditions as to render their repeated replication infeasible.

One proposed solution that could help to resolve this is the application of simulation to the accreditation of autonomous systems. By creating a synthetic environment in which to exercise the capabilities of the autonomous system and observe its behaviour, it could be possible to run millions of simulations in parallel in order to achieve a fuller characterisation of an autonomous system's capability than would be possible through real-world trialling.

The ability to expose such corner cases in synthetic environments will necessitate simulation capabilities that go beyond many of the solutions that have been developed today for engineering or training. Adequately representing complex scenarios for autonomous systems to navigate will require high-fidelity models of the environment, external entities (traffic, pedestrians etc.), weather, connectivity and so on.

Furthermore, in order to enable evidence gathered in the synthetic environment to be part of an overall system certification process, one can argue that the synthetic environment itself – as well as the data and models used – will require a level of accreditation to demonstrate that it is an appropriate test tool that produces valid results. Accredited simulators and approved models already exist in both civil and military domains – for example, full-flight simulators are approved by aviation authorities in order to be used for civil pilot training. As far back as 1985, the North Atlantic Treaty Organisation (NATO) issued a report proposing methodologies to validate missile system simulation, having found procedures for simulation model validation to be generally lacking (AGARD, 1985).

This paper will explore the challenge of autonomous systems testing in synthetic environments through two lenses. The first will focus on activity carried out by XPI Simulation, Warwick Manufacturing Group (WMG) at the University of Warwick, and Thales under the Certification of Autonomous Vehicles in Synthetic Environments (CAVinSE) feasibility study, which considers the automotive domain. The second will consider the defence domain in general, with a maritime autonomous system (MAS) certification approach providing a particular use case. While highlighting the contrasts between autonomous system certification requirements in the automotive and defence domains, common elements and scope for complementary methodologies and tools will be identified. Outstanding technical challenges and areas for future research will also be elaborated upon.

AUTOMOTIVE AUTONOMY

The dream of self-driving cars has been pursued, mirage-like, by technicians for decades. Projects have abounded, but the launch by the US Defence Advanced Research Project Agency (DARPA) of the 'Grand Challenge' in 2002 (the first event taking place in 2004) can be seen as a milestone in the development of modern CAV technology. While none of the competitors in the first 'Grand Challenge' met the target of completing a 150 mile (240km) course in the Mojave Desert (the farthest distance covered being 7.32 miles), the following year saw five vehicles completing the 132 mile course. The intervening period has seen rapid development of the underlying technology to enable self-driving cars – including sensors (ultrasound, camera, radar and LiDAR), image processing, computer vision and machine learning.

At the time of writing, much of this enabling technology is being proven through the development and fielding of advanced driver assistance systems (ADAS) – such as adaptive cruise control (ACC), automated emergency braking (AEB) and automated parking features. In accordance with the SAE International (formerly known as the Society for Automotive Engineers) publication J3016™, such features are typically categorised as Level 0 (warnings and momentary assistance), Level 1 (automation of a longitudinal or lateral driving function) or Level 2 (automation of both longitudinal and lateral driving functions) (see Figure 2). However, the move to higher levels of autonomy will inevitably require a far greater level of technical capability and more stringent safety requirements, as Level 4 and Level 5 automation does not assume that the human driver is able to take over in the event of the system failing to operate safely.

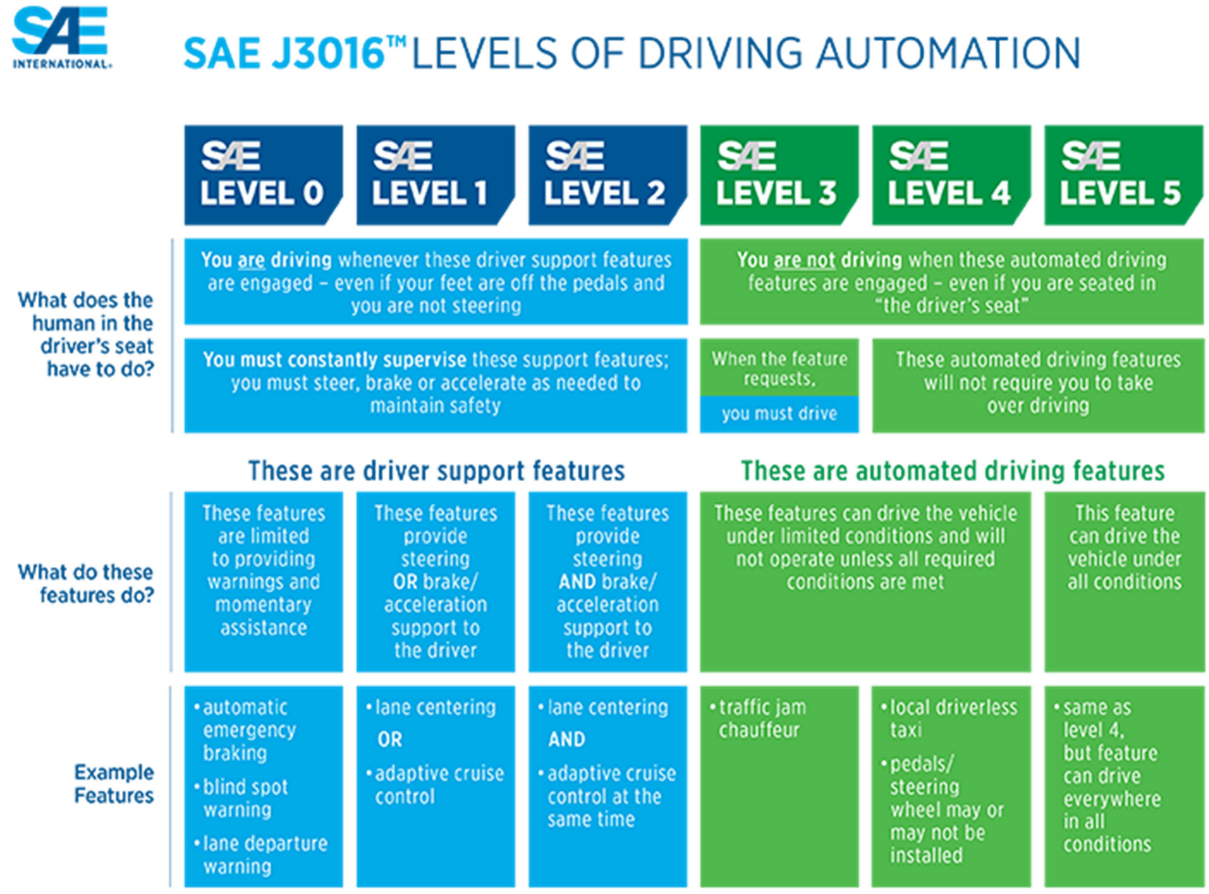

Figure 2: Levels of Driving Automation (SAE, 2019)

Participants in the race to autonomy include traditional automotive companies (commonly referred to as OEMs – original equipment manufacturers), Tier 1 suppliers to OEMs, mobility services companies and the technology sector. Although this paper will not seek to establish the leaders in the field, it is worth noting that Waymo, the self-driving car arm of Alphabet (the parent company of Google), claims to have achieved over 10 million miles of on-road testing in the US as of February 2019 (Waymo, 2019). This figure, although believed to be highest of any autonomous vehicle developer, pales in comparison to the 8.8 billion miles cited previously.

However, Waymo (among other developers) has deployed simulation in order to train and characterise their autonomous vehicle algorithms – presenting them with synthetic data and evaluating the resultant performance of the CAV. As of February 2019, the number of miles simulated by Waymo amounted to over 7 billion (Waymo, 2019).

This accumulation of mileage (whether in the real-world or in simulation) is primarily focused on the development of the underlying technologies that will enable autonomy – seeking to expose the CAV to as many different driving conditions as possible and gathering data about those that are safety navigated (or not) in order to refine the capabilities of the overall system. While evidence gathered in the course of such mileage may help to characterise the safety-related performance of the system, a structured process that will demonstrate the safety of the autonomous vehicle is clearly a requirement in regulated markets. Miles travelled are only of use if there are relevant to the operational design domain (ODD) of the CAV in question – the road type, the traffic conditions and the weather are all key features of the CAV operating envelope. Demonstrating safe behaviour in the scenarios that arise in these ODDs is therefore crucial to the demonstration of coverage – and these scenarios may not manifest even in a million (or a billion) miles of on-road driving. Identification and execution of the right scenarios – with adequate parameters to allow for variation in the real-world, is therefore a key part of a structured certification approach (PEGASUS, 2019).

Guidelines exist for the deployment of autonomous vehicles on public highways for testing (such as the UK's *Code of Practice: Autonomous vehicle trialling*). ISO:26262, the standard for development of safety-critical automotive software, does not seek to address autonomous driving, although new standards are emerging – such as the ISO/PAS (Publicly Available Standard):21448 covering Safety of the Intended Functionality (SOTIF) for road vehicles and UL 4600 (Standard for Safety for the Evaluation of Autonomous Products). Moreover, organisations with an interested in characterising the safety of vehicles – such as the European New Car Assessment Programme (Euro NCAP) – are developing testing standards and scenarios in order to enable assessments of features enabling lower-level autonomy. However, there are currently no formal regulations for CAV type approval – nor, indeed, for the accreditation of tools that form part of this approval process – which could include simulation software and models.

CERTIFICATION OF AUTONOMOUS VEHICLES IN SYNTHETIC ENVIRONMENTS

It is within this context that XPI Simulation, WMG and Thales embarked upon the CAVinSE feasibility study. This 16-month project, which started in July 2018, encompassed the following broad research goals:

- **Certification:** Develop a certification approach that would achieve the expected repeatability and robustness within the synthetic environment to test autonomous vehicles;
- **Security:** Understand the security vulnerability of a synthetic system being used in this manner to inform future simulator architectures
- **Standards:** Identify candidate architectures, formats and interfaces to drive and influence future standardisation and government policy;
- **Prototype:** Demonstrate a simulated environment incorporating a typical autonomous vehicle sensor as an early proof of concept.

In the interests of brevity, this paper will limit itself to description of the outputs from CAVinSE associated with certification and standards, although the security of any system used in certification is clearly a matter of considerable import in civil and defence domains alike (both for regulators and system manufacturers, and to engender trust among the wider public).

Certification

Given the lack of established regulations pertaining to type approval for CAVs, CAVinSE commenced with a review of extant certification approaches in the automotive, maritime, nuclear, rail and aerospace domains. The approach to tool accreditation in ISO:26262 was also considered, enabling a view to be taken of the overall development approach for the simulation-based certification tool.

In reviewing these approaches, the need for a modular and scalable approach to certification was identified – starting with system components and finishing with an entire system. Such a proposed approach is and is agnostic on the certification method – whether it be through simulation or real-world evaluation (see Figure 3).

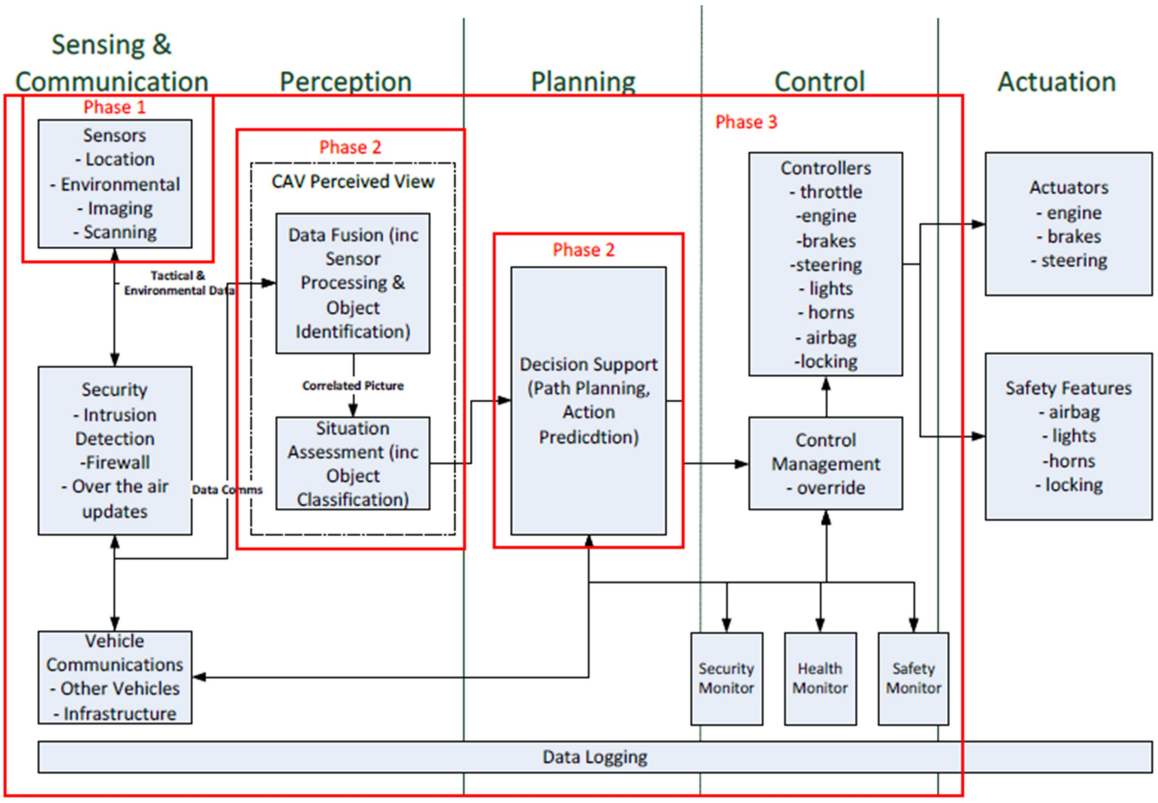

Figure 3: Candidate CAV Architecture and Certification Phases

It is imperative to note that this phased approach enables the overall system to flexibly meet the specific simulation needs of the components in question. For example, certifying a sensor in simulation would require high-fidelity models and a considerable amount of data about the environment model to be presented. Once the sensor model has been adequately characterised and accredited, a simplified version thereof can be used (with appropriate, but approximated, responses to noise, weather effects, interference and so on) for system-level simulation testing to evaluate the behaviour of the system under certain conditions (Koopman and Wagner, 2018).

Standards

CAVinSE identified that a critical aspect of simulation-based certification will be the implementation of standards that will enable industry-wide solutions to emerge. These standards can be expected to cover a variety of aspects, including methodological approaches, data formats, interfaces and scenarios.

In exploring existing standards for CAV simulation, CAVinSE identified three *de facto* standards that covered various elements salient to autonomous vehicle certification – OpenDRIVE (which describes the road network), OpenSCENARIO (which describes the dynamic elements of a scenario) and Open Simulation Interface (OSI). Other standards outside of simulation – such as ISO:23150, which defines (real) sensor interfaces – are also likely to be of use in developing solutions that enable suitable simulation of real-world components.

Unexplored areas for standardisation could also include noise models for sensors, weather effects and a consideration of the sensor parameters that need to be represented in simulation. Moreover, the investigation of existing (real-world) scenarios for testing autonomous vehicles (and autonomous features) identified a considerable lack of existing content that could be adopted in simulation – indicating that the various efforts (both underway and completed) to develop standard scenario libraries (such as PEGASUS and StreetWise) are addressing a real need.

Recommendations to Regulators

The high-level recommendations from CAVinSE to regulators concern the development of a phased certification process and the adoption (and influencing) of standards to enable broad certification approaches to emerge – and are applicable whether using simulation-based testing or other methods.

Considering simulation-based certification specifically, it is clear that the tools and models used in this regulatory context must themselves be accredited in order to ensure the validity of the synthetic environment used for testing. While it was identified that simulation is already used in type approval (such as in large vehicle roll-over testing), the guidance on the use of tools and models does not extend to stipulating the use of accredited instances thereof – leaving industry to decide what is appropriate and submit a justification to the type approval authority. Given the safety-critical nature of CAVs it was considered that more stringent accreditation procedures would contribute to improved safety, reduced risk and greater public acceptance.

Such an accreditation procedure could be considered to be composed of two key activities. Firstly, tools and models should be developed in accordance with established procedures for safety-critical systems (for example, the 'V' model that is at the heart of systems engineering). Secondly, testing should be carried out to compare the outputs of the synthetic environment with real-world observations and demonstrate a suitable, statistically significant correlation thereof. Various methodologies already exist that are used to align modelled with observed data, and these should be adopted in order to satisfy regulatory authorities that tools and models produce adequate levels of fidelity compared to actual behaviour (Youngblood and Petty, 2018).

DEFENCE CONSIDERATIONS

It is clear that the principle of the use of simulation to test and certify autonomous systems is relevant to the challenges of approving such systems in the defence domain as much as in the automotive sector. The benefits of reduced real-world testing, with the associated demands upon budget and time, are similarly attractive to defence users. Indeed, defence platforms often have higher running costs owing to their greater mass and the specialised nature of their components and consumables. Moreover, there is a considerable additional benefit in the testing of weaponised systems, in that the costs of munitions are far greater than that of an autonomous vehicle's chief consumable – fuel.

A key differentiator in considering CAV and autonomous defence platforms (in any operating domain) is that, while the former is focused on autonomous navigation (transport being the primary purpose of the platform), defence platforms may have a wider scope of autonomous functions depending on the mission. Therefore, autonomous defence platforms must balance safe operation (primarily focused on navigation and collision avoidance, as with CAVs) and effective employment (targeting, probability of success, payload management).

While the options of a CAV can be distilled to three principal actions – accelerate, brake, steer – those available to a military autonomous system may be much greater, depending on the level of capability, including activation of specific sensors (such as fire-control radar) or weapon release. Use of autonomous weapon systems has created controversy in the defence domain regarding the possibility (and desirability) of devolving certain autonomous functions (such as weapon release), with countries such as the the UK committing to keep a human-in-the-loop for any weapon usage (UK MOD, 2017). The UK Ministry of Defence (MOD) also maintains a distinction between automated and autonomous weapons systems, noting that the latter do not yet exist and averring no intent to develop such.

The evaluation of autonomous systems' behaviour in a defence context therefore requires a nuanced understanding of the appropriateness of any actions in accordance with the completion of mission objectives (which are likely to be more complex than the navigational objectives of a CAV), rules of engagement (ROE) and interpretations of international law.

The defence domain poses certain challenges with regard to the different models that are required to adequately represent autonomous systems. Military platforms typically use different – and more performant – sensors than

those slated for deployment on CAVs. This will necessitate the development of correspondingly more complex sensor models – and the creation of those, like sonar, which do not have a use-case for autonomous vehicle testing.

In addition, the modelling of behaviour of the external entities within the synthetic environment poses challenges for defence. While the behaviour of drivers and pedestrians can be aggressive and erratic, CAVs can be considered to operate in a relatively benign environment. The modelling of opposing forces with representative tactical behaviours operating at a granular level introduces new areas of complexity – and the contested nature of the electromagnetic spectrum requires consideration of deliberate, sophisticated interference with communications and sensors.

Notwithstanding these differences, there are clearly areas of overlap between civil and military systems that could enable the defence sector to take benefit from the greater scale of developments for civil markets. The use of CAV simulation tools to ascertain the appropriate navigational behaviours of autonomous defence systems (particularly when operating during peacetime) presents a useful opportunity for re-use. Commercially available sensors could be integrated onto military vehicle platforms for non-critical manoeuvres – providing collision avoidance warnings at low-speeds, for instance. The re-use of sensor models across both domains may therefore be possible.

More fundamentally, the methodologies proposed for CAV safety certification – with a key role for simulation – can also be proposed for autonomous defence platforms. The structured approach posited by the likes of PEGASUS regarding the clear articulation of the system capability and the ODD in order to generate scenarios for testing (both in simulation and the real-world) enables a clear demonstration of coverage and generation of evidence. This core approach can be considered to be platform and domain agnostic. Furthermore, the issues of tool and model accreditation are equally salient in the defence domain as in the automotive domain.

Beyond these technical considerations, there is a differentiated path for the approval of defence systems compared to CAVs – and the specific regulatory agencies involved in certification of defence systems will impose their own requirements for the validity of evidence, simulated or otherwise, in support of an overarching safety-case. The following section on use case in the maritime defence domain considers some of these differences.

USE CASE – MARITIME AUTONOMOUS SYSTEMS

In considering the use of autonomous systems in a defence context, it is useful to examine a maritime autonomous system (MAS) use case in the defence domain. MAS solutions are being developed across the world and the defence domain can envisage multiple military and non-military means to which such systems can be directed. MAS may be employed in the waters of the nation itself, those of allied and partner nations, or in hostile contexts (see Figure 4).

Figure 4: Apollo Unmanned Surface Vehicle (USV)

As with other users of automated systems, defence use of MAS often focuses on tasks that are 'dull, dirty and dangerous'. However, they also offer opportunities to increase overall capabilities, with new modes of operations and features likely to emerge with time. Examples in defence span the full range of MAS, including:

- Environmental measurement using sensors mounted on long-endurance MAS (e.g. wave gliders) in order to inform the use of other equipment, such as sonar;
- Persistent monitoring of areas to protect valuable assets (e.g. ships), locations (e.g. harbours) and regions (e.g. channels) using a range of surveillance sensors mounted on MAS, for example, cameras, sonars;
- Detecting and clearing areas of mines using systems formed from MAS and unmanned underwater vehicles with the benefit of removing people from the main area of risk;
- Large MAS that are designed for long-duration and long-distance operations.

A key difference for the use of MAS in defence and security scenarios is in the level of risk that is acceptable to the operating organisation. Peacetime use of MAS might include training activities or regular monitoring of areas; in these situations it is to be expected that the MAS will be operated to the same standard as commercial or scientific operations with the risk managed to the same degree. Where the MAS is used in contexts with tension or conflict, the level of risk to the MAS and others that is acceptable to the operating organisation may change. For example, compliance to International Maritime Organisation (IMO) Convention on the International Regulations for Preventing Collisions at Sea (COLREGs) may be reduced to ensure the MAS mission is not disrupted by intentional obstruction, or the MAS may take a more covert approach with reduced lighting to minimise risk of detection.

The UK's Maritime and Coastguard Agency (MCA) is working closely with the IMO to ensure that the regulatory environment in the UK is suitable for the safe use of Maritime Autonomous Surface Ships. In addition, the UK is an example of a nation where there is a dedicated regulatory organisation responsible for maritime safety for Defence related activities: the Defence Maritime Regulator (DMR) which is part of the UK's Defence Safety Authority. The Maritime UK led MAS Regulatory Working Group is a good example of pan-organisation collaborative approach that has published the widely-adopted Maritime Autonomous Surface Ships Code of Practice.

While the regulators are proactively and positively addressing the need to enable the development of MAS through defining relevant regulations, a key challenge is in providing confidence and evidence of the safety of MAS against those regulations. An approach being taken by Thales in the development of a substantial system of systems in the maritime autonomy domain is to build confidence in compliance with COLREGs through multiple simulations of a range of scenarios involving the system of interest and other ships. The simulation is built on VR Forces, with an implementation of the COLREGs behaviour model. An analysis tool has also developed to provide assessment of compliance against COLREGs.

The simulations, supported by at-sea trials, are a key element as they permit a significantly larger, and more varied, body of evidence to be generated in scenarios that would be unfeasible to test at-sea. For example, an autonomous vessel will need to respond safely in the event that another ship does not follow COLREGs expected behaviours. Similarly, in scenarios of tension or hostility the acceptable level of risk may be adjusted, with simulation enabling an assessment and characterisation of the impact on safety needs from such adjustments.

CONCLUSIONS AND FURTHER RESEARCH

This paper has demonstrated that the principles of using synthetic environments for the certification of autonomous (road) vehicles can also be applied to the defence domain. The benefits of reduced real-world testing, greater control over the environment, improved safety and the ability to massively scale simulation-based testing make it a compelling cross-domain proposition.

It is clear, however, that research challenges still pertain in the development of modelling and simulation tools for certification activity. The development of appropriate sensor models is a particular area where further investigation is required – not only to create higher-fidelity representations, but to determine the appropriate level of fidelity that is required at different stages of the certification process.

Moreover, the representation of the physical environment in order to provide data for these sensor models presents technical complexity. Capturing real-world environments – as has been carried out by XPI in the creation of LIDAR-scanned databases – is readily achievable, albeit requiring many hours of manual labour to correct imperfections in the data and remove undesired objects. Creating geo-typical environments with the same level of detail and richness as real-world areas is a further challenge.

The representation of human behaviours in simulation – whether as drivers, pedestrians, cyclists or opposing combatants – requires an acceptable degree of realism in order to truly evaluate an autonomous systems' performance and interaction with other entities in particular scenarios. The use of machine learning (ML) to analyse observed behaviours – which is being pursued by Latent Logic as part of OmniCAV, a follow up project to CAVinSE in which XPI is a key partner – presents a possible route to developing such smart actors.

Fundamentally, the technical challenges notwithstanding, the key enabler to the use of simulation for autonomous system certification requires a clear regulatory framework that defines the parameters associated with the use of synthetic environments – and the models employed therein – for this task. The creation of these regulatory frameworks – encompassing all aspects of autonomous systems testing and the evidence required to support safety cases – could enable the innovative potential of such systems across multiple domains, enabling mobility for the disabled, reducing risks to pedestrians and removing humans from all manner of dull, dirty and dangerous tasks.

ACKNOWLEDGEMENTS

CAVinSE and OmniCAV are both part-funded by the Centre for Connected and Autonomous Vehicles (CCAV), delivered in partnership with Innovate UK. It is part of the government's £100 million Intelligent Mobility Fund, supporting the Future of Mobility Grand Challenge. As a key part of the UK government's modern Industrial Strategy, the Future of Mobility Grand Challenge was announced in 2017 to encourage and support extraordinary innovation in UK engineering and technology, making the UK a world leader within the transport industries. This includes facilitating profound changes in transport technologies and business models, to make the movement of people, goods and services across the nation greener, safer, easier and more reliable.

REFERENCES

Advisory Group for Aerospace Research and Development (1985). *Final Report of the Flight Mechanics Working Group WG-12 on Validation of Missile System Simulation*.

Kalra, N. and Paddock, S. (2016). *Driving To Safety: How Many Miles Would it Take to Demonstrate Autonomous Vehicle Reliability*. RAND Corporation

Koopman, P, and Wagner, M. (2018). Towards a Framework for Highly Automated Vehicle Safety Validation. *PREPRINT: SAE World Congress*.

PEGASUS (2019), *The PEGAUS Method: An Overview*. Retrieved June 13, 2019, from https://www.pegasusprojekt.de/files/tmpl/Pegasus-Abschlussveranstaltung/PEGASUS-Gesamtmethode.pdf

SAE (2019). *SAE J3016™ Levels of Driving Automation.* Retrieved June 4, 2019, from https://www.sae.org/news/2019/01/sae-updates-j3016-automated-driving-graphic

UK Ministry of Defence (MOD) (2017), *Joint Doctrine Publication 0-30.2: Unmanned Aircraft Systems*. Retrieved June 13, 2019, from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/673940/doctrine_uk_uas_jdp_0_30_2.pdf

Waymo (2019). *An update on Waymo disengagements in California.* Retrieved June 4, 2019, from https://medium.com/waymo/an-update-on-waymo-disengagements-in-california-d671fd31c3e2

Youngblood, S. and Petty, M. (2018). *Assessing the Challenges of Rigorous Simulation Validation*. Tutorial given at I/ITSEC 2018 on November 26, 2018.