MOVE_UK

Data Analysis Report Phase 3

August 2019

TSB Reference No.	102584
Project Title	MOVE_UK - Accelerating automated driving by connected validation and big data analysis
Deliverable No.	D7.5
Principal Author	Simon Morley (Bosch)
Contributing Authors	Kia Hafezi (Bosch), Istvan Bajnok (Bosch), Tobias Mathony (Bosch), Michaela Semmler (Bosch), Chris Ndava (Bosch), Faisal Khan (Bosch), Jolyon Carroll (TRL), Arun Kalaiyarasan (TRL), Matthias Seidl (TRL), Jonathan Kent (TRL), Mark Bell (TRL), Tanya Robinson (TRL), Ian Barlow (Jaguar Land Rover), Sam Chapman (The Floow), Mark Burke (The Floow), Dan Freedman (Direct Line Group), Heather Yedigaroff (Royal Borough of Greenwich)
Date	1st August 2019
Issue/Revision No.	1.0

Data Analysis Report – MOVE_UK Phase 3 (Deliverable D7.5)

Table of Contents

Executive summary

Introduction

MOVE_UK is a project that has contributed to the progression towards automated driving. The particular focus for the MOVE_UK contribution has been connected systems validation and analysis of ‘big data’. MOVE_UK is a collaboration of industry partners led by Bosch and supported by TRL, Jaguar Land Rover (JLR), Direct Line Group, The Floow and the Royal Borough of Greenwich. The project ran from August 2016 to July 2019, with the final phase of the project (Phase 3) running from October 2018 until June 2019.

The primary focus of this report is on three new use cases, which relied on the Phase 3 corner radar capabilities: Surround sensing of Cut-In situations (SUR-CIN), Surround sensing of Cross-traffic situations (SUR-CROSS), and further telematics (Telematics 3). However, updates from the use cases defined in Phases 1 and 2 are also provided.

Update on Phase 1 and 2 Use Cases

The use of event-based data collection, where only relevant data is recorded and stored has been shown to save a lot of data storage space and processing time compared with conventional data collection. This new approach will enable shorter development cycles for Automated Driving Systems (ADS), as only relevant data needs to be processed. This reduces the time associated with moving, checking, reviewing and annotating the data prior to reprocessing the data with newer software, for instance in a Hardware in the Loop (HiL) simulation system.

The results of re-simulation of the captured Advanced Emergency Braking (AEB) events with production camera software were far better than expected and highlight the relevance of the data collected from the MOVE_UK trials. Critical AEB events, such as those collected in MOVE_UK, are very difficult to observe in the real world and are extremely useful for the development and validation of future ADS features. The main reason for their usefulness is that they often cover driving situations which the engineers/developers did not foresee and also situations which testing bodies (like Euro NCAP) do not currently assess. Using these events/sequences to test ADS software in a simulation environment, therefore, leads to more broadly effective and safer ADS features.

Through the Traffic Sign Recognition (TSR) use case, MOVE_UK developed an ‘evidence-based methodology’ that can determine the detection rates of speed signs and their change in detection likelihood over time. The methodology can be used to identify locations where the recognition system is performing poorly. These location(s) can then be communicated to concerned parties for restorative interventions which would improve detection going forward (via a feedback loop system).

Phase 3 Methodology and Use Cases

For Phase 3 of MOVE_UK, the Phase 1 and 2 vehicle modifications were extended in two of the five Land Rover vehicles with the addition of four corner radar sensors and the corresponding adaptations to hardware and software in each of the two vehicles. The corner radars were added to allow full 360-degree surround vehicle sensing. This enabled the following observations to be made.

The Time-To-Collision (TTC) for the following vehicle in a Cut-In scenario is generally between 1 and 10 seconds (s), around the point of Cut-In, and rarely goes below the 1s mark. This shows us that the drivers of following vehicles are, for the most part, leaving a sufficient gap to the ego-vehicle in-front of them to avoid a collision should the ego-vehicle brake sharply when another vehicle cuts-in in front of it; this is a behaviour that automated driving vehicles of the future should emulate. However, there are still a significant number of occasions when the TTC for the following vehicle is below 2s, which shows that some drivers leave smaller time gaps than the 2 seconds recommended by the Highway Code, and therefore could become frustrated if automated vehicles are designed to maintain time gaps of 2s and above only. Further investigation into the gaps which automated vehicles should maintain is required to ensure that ADS do not cause users to become frustrated. It is also important to note that if the time-gap between vehicles is less than 1s, an automated vehicle will have difficulty merging into a lane with other traffic. Again, further investigation of this kind of scenario is required.

The Geo-Fence based analysis approach was valuable in identifying sites for evaluation of traffic flow and driver behaviour. This approach has potential for being adapted for other use cases, as demonstrated in the determination of hotspots where TTC was observed to be low.

From the investigation of behaviours around specific locations, it is clear that vehicles fitted with surround sensors present a huge opportunity to learn about the decisions taken by human drivers. Data that can help the understanding of drivers' behaviours when interacting with other road users will be of value to those developing automated vehicles, to insurers for the purposes of identifying when customers participate in higher-risk manoeuvres, and to local authorities for providing an evidence-base to inform on improvements to infrastructure.

The high value of the forward radar data to both Event Data Recording (EDR) and risk assessment was demonstrated during Phase 2, and the analysis from the Phase 3 use cases strongly suggests that the surround data can add significant value to our understanding of both.

<u>Post Project</u>

Event-based data collection was demonstrated to be a valuable method in reducing time and resources required for future ADAS and ADS testing and development. It is the intention of JLR and Bosch to use this method to help develop its next generation of ADAS features. It is also the intention of Bosch to use the 23 critical AEB events captured within MOVE_UK to help test and develop its next generation of production AEB software.

The Traffic Sign Recognition (TSR) methodology is ready to be transferred to another setting where it could be used to identify locations where restorative interventions could improve visibility of signs and system performance. In the future, this methodology could also be adapted to work with other vehicle systems too. Potentially it could be used in pothole detection or with other road signs, road markings and infrastructure detection systems that support automated driving.

1 Introduction

MOVE_UK is a project contributing to the progression towards automated driving through connected systems validation and analysis of 'big data'. Specifically, MOVE_UK has trialled a new method of validating Automated Driving Systems (ADS) using a fleet of five Land Rover vehicles fitted with a range of sensors and data recording/connectivity equipment. The fleet has completed over 120,000 miles on public roads, predominantly in and around Greenwich, London. MOVE_UK is a collaboration of industry partners led by Bosch and supported by TRL, Jaguar Land Rover, Direct Line Group, The Floow and the Royal Borough of Greenwich. The project benefits from a grant of £3.4 million from the UK Government's £100 million Intelligent Mobility Fund, which is administered by the Centre for Connected and Autonomous Vehicles (CCAV) and delivered by the Innovate UK agency. The project began in August 2016 and concluded in July 2019. Phase 3 of the project ran from October 2018 until June 2019.

This report follows on from the Phase 1 report published in January 2018, *MOVE_UK Data Analysis Report - Phase 1 (D7.3)*[1] and the Phase 2 report published in November 2018, *MOVE_UK Data Analysis Report - Phase 2 (D7.4)*[1]. In D7.3, the project was introduced and the design and initial data analysis results of four project 'use cases' were described: camera based Autonomous Emergency Braking (AEB), Driver Harsh Braking (DHB), Traffic Sign Recognition (TSR) and Telematics 1. In D7.4, an update was given on the Phase 1 results and four additional use cases were introduced: radar based Autonomous Emergency Braking (ARB), Cut-In Scenarios (CIN), Lead Vehicle Statistics (LVS), and Telematics 2; all of these uses cases rely on additional data being received from a front radar sensor fitted to the trial vehicles at the beginning of Phase 2. Updates on the data analysis results related to the Phase 1 and Phase 2 use cases are contained in Section 3 of this report.

The defining feature of MOVE_UK Phase 3 was the addition of corner radar sensors to two of the trial vehicles. This additional hardware extended these vehicles' sensor perception yet further and allowed experimentation with new use cases which rely on a full 360-degree understanding of surrounding objects; something that is vital for the successful development and validation of more complex automated driving features. The project's data recording and analysis tools were updated and extended to allow recording and visualisation of the additional corner radar data. These improvements are described in Section 2.

The primary focus of this report is on three new use cases, which make use of the Phase 3 corner radar capabilities: Surround sensing of Cut-In situations (SUR-CIN), Surround sensing of Cross-traffic situations (SUR-CROSS), and further telematics (Telematics 3). Figure 1 illustrates the role of these use cases within the previously defined framework of capabilities which MOVE_UK aims to demonstrate, and the applications that are envisaged for the developed methodologies. The Phase 3 use cases and related capabilities are described in more detail in Section 4.

Finally, Section 5 captures our final conclusions from all Phases of MOVE_UK.

[1] http://www.move-uk.com/publications.html

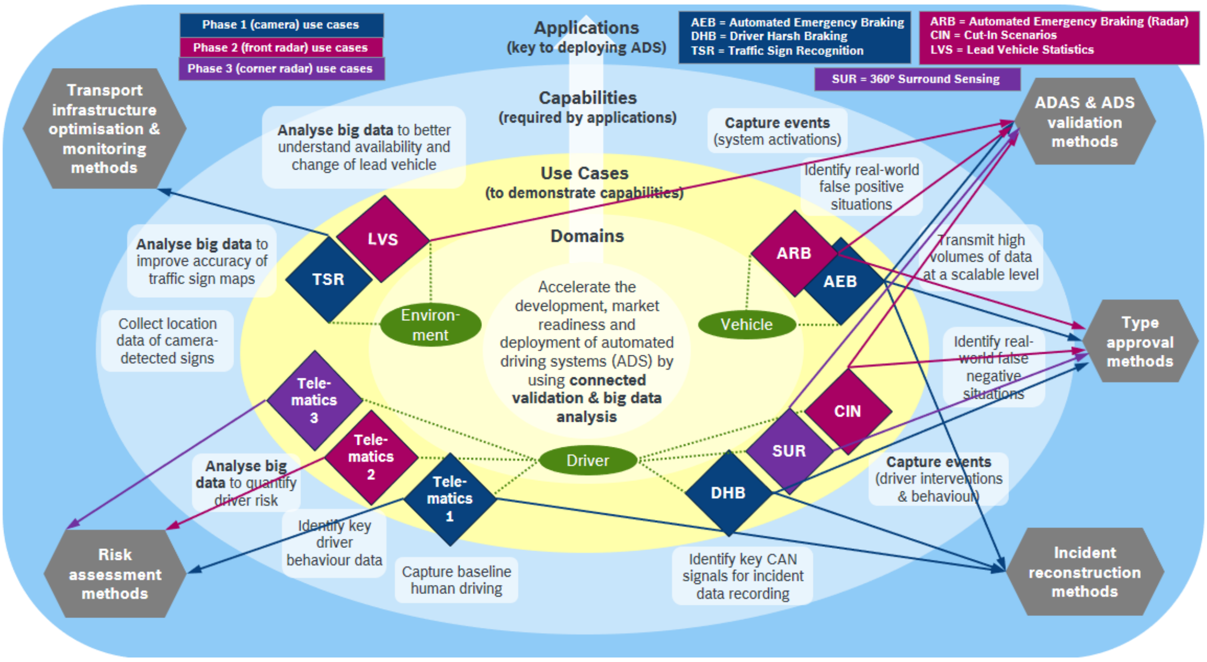

Figure 1 – MOVE_UK objectives and project strategy for Phases 1, 2 and 3

2 Data collection and analysis methodology for Phase 3

2.1 Additional hardware and software adaptions for Phase 3

For Phase 3, the Phase 1 and 2 vehicle modifications were extended in two of the five Land Rover vehicles with the addition of four corner radar sensors and the corresponding adaptions to hardware and software in each of the two vehicles. The corner radars were added to allow full 360-degree surround vehicle sensing.

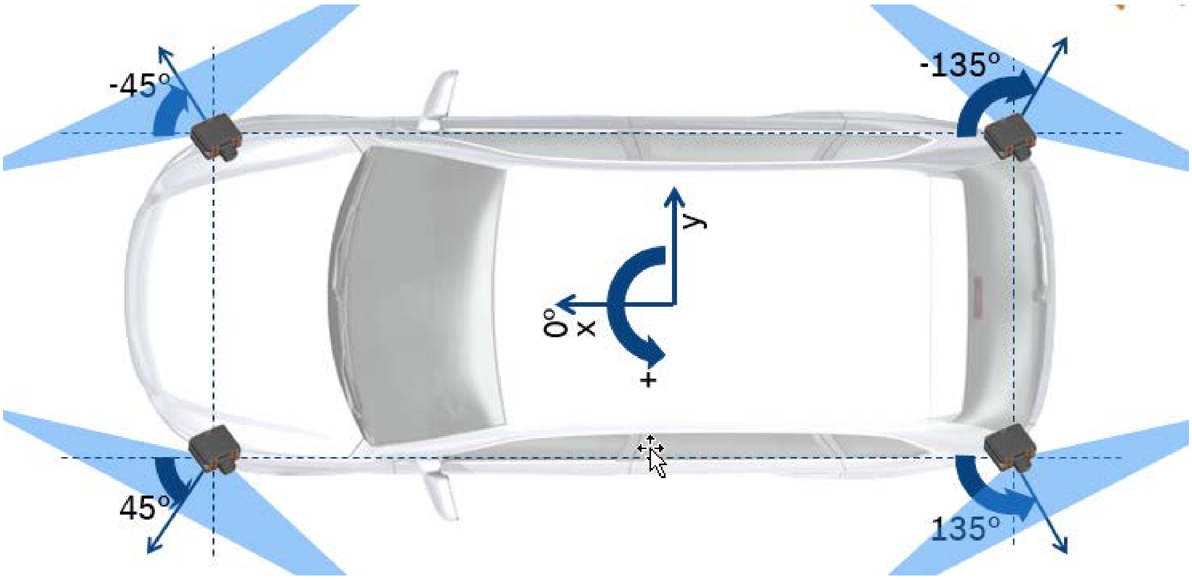

Figure 2 – Corner Radar integration positions for Phase 3 hardware setup

Table 1 – Technical specification of corner radar sensor

Technical data	
Frequency range	76...77 GHz
Detection range	0...~90 m
Field of view (horizontal)	±75° (for all antenna ranges)
Measuring accuracy **Distance** **Speed** **Angle**	 0.12 m 0.14 m/s ±0.8°
Object separation capability **Distance** **Speed** **Angle**	 0.72 m 0.78 m/s 7°
Cycle time	~60 ms
Modulation	Frequency modulation (FMCW)
Max. number of tracked objects	32
Dimensions (W x H x D) in mm	70 x 60 x 30 (without connector) 70 x 82 x 30 (with connector)
Weight	~190 g
Power consumption	4.5 W

For both Phase 2 and Phase 3, a Mid-Range-Radar (MRR) was chosen for use in this project because it is more suitable for urban environments. Table 1 displays the technical specification of the radar sensor selected for the Phase 3 surround- sensing application. This sensor, which is designed primarily for rear sensing, was added to all four corners of the vehicles. In each case, two of the four sensors were reconfigured to face forward. The front sensors are positioned at ±45° to the longitudinal axis of the vehicle. The rear sensors are positioned at ±135° to the longitudinal axis of the vehicle, as shown in Figure 2.

To make the radar integration possible for Phase 3, some external modifications were made to the two vehicles. In the front, fog lamps were removed and 3D printed brackets, designed specifically to hold the front corner radar sensors in the correct position, were installed in the fog lamp area. This approach allowed the 'radome' (a structural, weatherproof enclosure made of a special radar transparent material, protecting the radar antennas, at the very front of the sensor), to not be covered by the bumper. This is important for making the integration of the sensors easier, because normal bumper materials, especially painted bumpers, in front of the antennas create reflections. These reflections have to be compensated for in the software, otherwise ghost objects or loss of performance is visible. Front integration of one of the corner radar sensors is show in Figure 3.

CAD | Component | Integrated in vehicle

Figure 3 - 3D printed front sensor brackets

Rear integration was very similar to the front solution, except that rectangular holes needed to be cut in the rear corners of the vehicles to accommodate the rear sensor brackets, as shown in Figure 4. Both front and rear brackets for the sensors were designed using CAD, then manufactured using rapid prototyping techniques, before being finally integrated into the vehicles.

CAD | Component | Integrated in vehicle

Figure 4 – 3D printed rear sensor brackets

To avoid any adverse impact on the Phase 2 vehicle recording system (which was already working to near full capacity) and to support the newly added corner radar sensors, it was decided to install an additional/parallel vehicle recording system for Phase 3. This involved adding a further Flea3 data storage and connectivity device and a further CAN Gateway to each of the two Phase 3 vehicles. A forward-facing IP (web) Camera was also mounted behind each of the vehicles' front windscreens (on the passenger side) and connected to the new recording system. The video footage from this camera, which was recorded only during Phase 3 events, supplemented the corner radar sensor data and allowed a better understanding of the driving situation. The full setup including the additional Flea3 device and CAN Gateway device, as well as the IP Camera, are shown in Figure 5.

Additional Flea3 and Gateway IP Camera

Figure 5 – additional equipment

Figure 6 – Two vehicles after Phase 3 equipment integration at Bosch Denham site, near Uxbridge

2.2 Field data collection

2.2.1 Changes to continuous CAN data collected in Phase 3

The changes made in Phase 2 to the list of CAN signals collected continuously at 1Hz on the first Flea3 device proved successful. As a result, no further changes were made in Phase 3 to the continuous CAN data recording.

2.2.2 Event-triggered CAN data collected in Phase 3

For the Phase 3 use cases, where an understanding of the objects surrounding the vehicles was required, a range of CAN signals, mainly corner radar related signals, were recorded on the second Flea3 device when a Phase 3 event was triggered. These signals were recorded at a higher frequency (10 Hz) than the 1 Hz continuous CAN data. The recording duration was initially set at 30 seconds (s); 15 s before and 15 s after the event. In addition to these signals, the event-triggered data captured in Phase 2 continued to be recorded in Phase 3.

The second Flea3 device was not capable of recording all the CAN signals associated with all of the objects detected by the corner radar sensors. Each corner radar sensor can recognize and track 32 objects, so in total there was the potential to track 128 objects surrounding the MOVE_UK ("ego") vehicle. For each object there are a few very important and many less important signals. These signals describe the position and speed of the object relative to the ego vehicle, whilst also providing information as to whether the object is moving or stationary, measured or estimated from previous measurements, as well as the existence probability of the object and so on.

In order to reduce the amount of corner radar data captured and so also reduce the load on the second Flea 3 device, data filtering before recording was used. This filtering was implemented on the additional CAN gateway installed into the vehicles for Phase 3. A solution that provided the amount of data that fitted within the system limitations but still allowed an understanding of how the surrounding objects were moving around the ego vehicle was developed by applying different combinations of filtering to test measurements. This solution focussed on capturing the most relevant objects.

As a result of the data reduction activity, a maximum of seven objects from each corner radar sensor were recorded, resulting in recording data from a total of 28 objects. Static objects were disregarded because in most cases they are less critical and recording all the objects was not an option.

For the selected objects, five important CAN signals were collected when recording was triggered by an event. These signals were:

- Distance to the object;
- The horizontal (azimuth) angle to the object;
- Relative speed of the object to the Ego vehicle;
- A moving state signal which differentiates between moving and previously moving objects and was used to filter objects; and
- An ID signal for tracking objects.

In addition to these 140 (28 x 5) corner radar related CAN signals, a further 57 CAN signals were recorded on the second Flea3 device to support the Phase 3 use cases. These signals contained:

- Dynamic information from the vehicle, for example ego speed, steering wheel angle, brake pressure and longitudinal acceleration;
- Lane and first object information signals from the stereo camera installed in Phase 1;
- Object information from the MRR plus forward-facing radar sensor installed in Phase 2;
- Trigger signals which were used in different use cases (for example, in the AEB, ARB, CIN and Phase 3 use cases);
- GPS signals for location identification; and
- Some status signals of the Flea3 device.

2.2.3 Event-based IP-camera video sequences

For Phase 3, the two vehicles equipped with corner radar sensors were also fitted with an IP camera. The image stream from each of these cameras was grabbed, compressed and transferred by the second Flea 3 device, via mobile 3G network, to the cloud server along with the event-triggered CAN data mentioned above. Continuous measurement was not possible because of the mobile network bandwidth and system load limitations, so only event triggered capture of the IP-camera and corner radar data was implemented in Phase 3.

The frequency of the image and CAN recording was 10 samples per second and the resolution of the image recording was 352 x 240 pixels. This gave enough insight into the event in question and was a good compromise taking into consideration the amount of data to be transferred. This frequency and resolution also allowed the image stream to be transferred and viewed almost instantly, unlike the higher resolution images collected from the Stereo Video camera (used to support the Phase 1 use cases) which had to be transferred via Wi-Fi once the vehicles had returned to the Wi-Fi hotspots, leading to some delay. The recording of relatively low-resolution images also meant that the images required no further anonymisation to comply with data privacy regulation, such as GDPR, as it was not possible to recognise other vehicles' number plates or pedestrian's faces from them.

The goal of the IP camera recording was to make it easier for the MOVE_UK engineers and data scientists to follow and interpret the recorded event and associated radar data. It also served as a reference sensor, especially when an interesting situation arose, allowing further investigation. The two streams of data, IP camera data and CAN data, were not fully synchronised, meaning that there was a slight delay visible between the two. This was a result of the system load and implementation on the Flea 3 device. Further improvements were outside the scope of the MOVE_UK project.

2.2.4 Telematics data collected in Phase 3

The telematics devices remained in the vehicles throughout the lifetime of the project. The process for collecting 1 Hz GPS data was not changed during Phase 3. The telematics data was processed, the journeys graded for driving quality, and then stored as part of The Floow's data archive. The Floow developed a data fusion process to combine their telematics pool with data collected from the CAN bus.

2.3 Data storage and analysis tools

2.3.1 Improvements to data visualisation and analysis tools for Phase 3

Since publication of the Phase 2 report (D7.4), many improvements to the data visualisation and analysis tools have been made. This is particularly the case for matching the new requirements which have originated from the new surround radar-based use cases.

At the start of Phase 3, an offline visualization tool was developed by the data analytics team setup during the project. This tool, as depicted in Figure 7, allowed analysis of surround radar data prior to the deployment of the sFDE SUR-View (see details below). It also allowed more detailed analysis of surround events than was possible in sFDE.

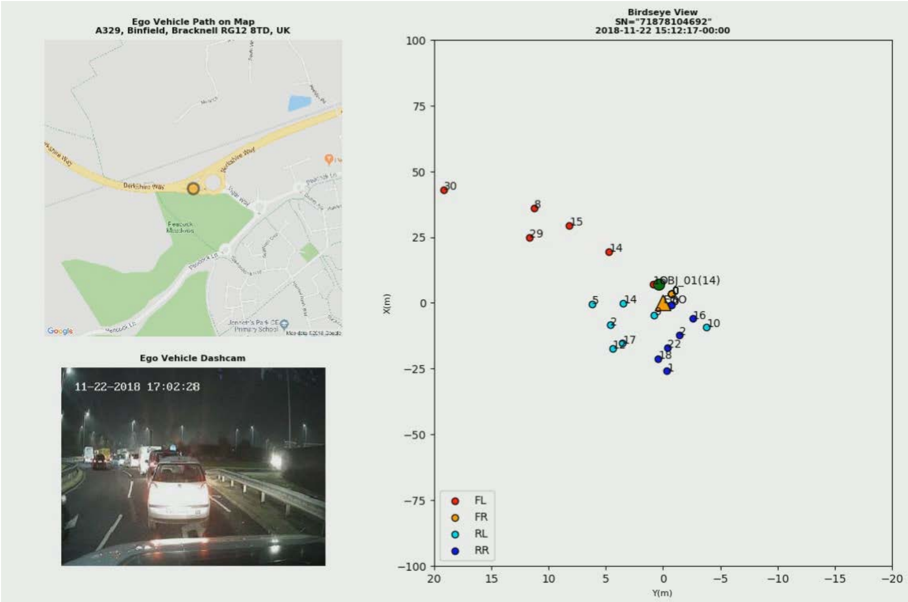

Figure 7 – Offline data visualisation tool for surround data

2.3.2 Systematic Field Data Exploration (sFDE)

As described in the Phase 1 and 2 reports (*MOVE_UK Data Analysis Report - Phase 1 (D7.3)* and *MOVE_UK Data Analysis Report - Phase 2 (D7.4)*), sFDE is a system, developed by Bosch Software Innovations GmbH (INST), used to gather data from any source and to make it available to a data management system. For the MOVE_UK project, a web user interface (UI) was provided for accessing and analysing data stored in sFDE. Besides providing a view for streamed continuous CAN data from each individual trip carried out by any one of the MOVE_UK vehicles, views were implemented in MOVE_UK Phases 1 and 2 to display the use case data for Traffic Sign Recognition (TSR), Autonomous Emergency Braking (AEB), Driver Harsh Braking (DHB), Lead Vehicle Statistics (LVS) and Cut-In (CIN).

For Phase 3, sFDE was extended with the addition of a new "SUR-View" to display the data from the new Phase 3 use cases (SUR-CIN and SUR-CROSS), as shown in Figure 8.

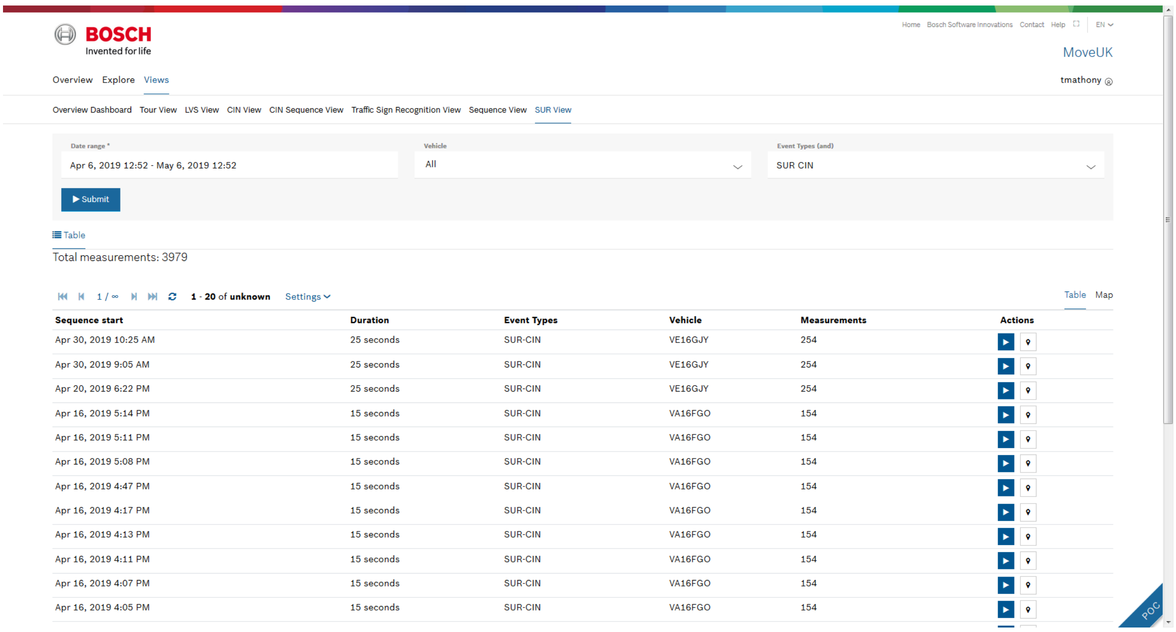

Figure 8 – New SUR-view for Phase 3

When clicking on a certain event, a detailed event view is displayed showing IP camera footage of the event, an event map displaying GPS data, a “Bird’s Eye View” displaying the front radar objects recorded, and a new “Surround Bird’s Eye View” displaying the surround radar objects recorded during SUR-CIN or SUR-CROSS event (see Figure 9).

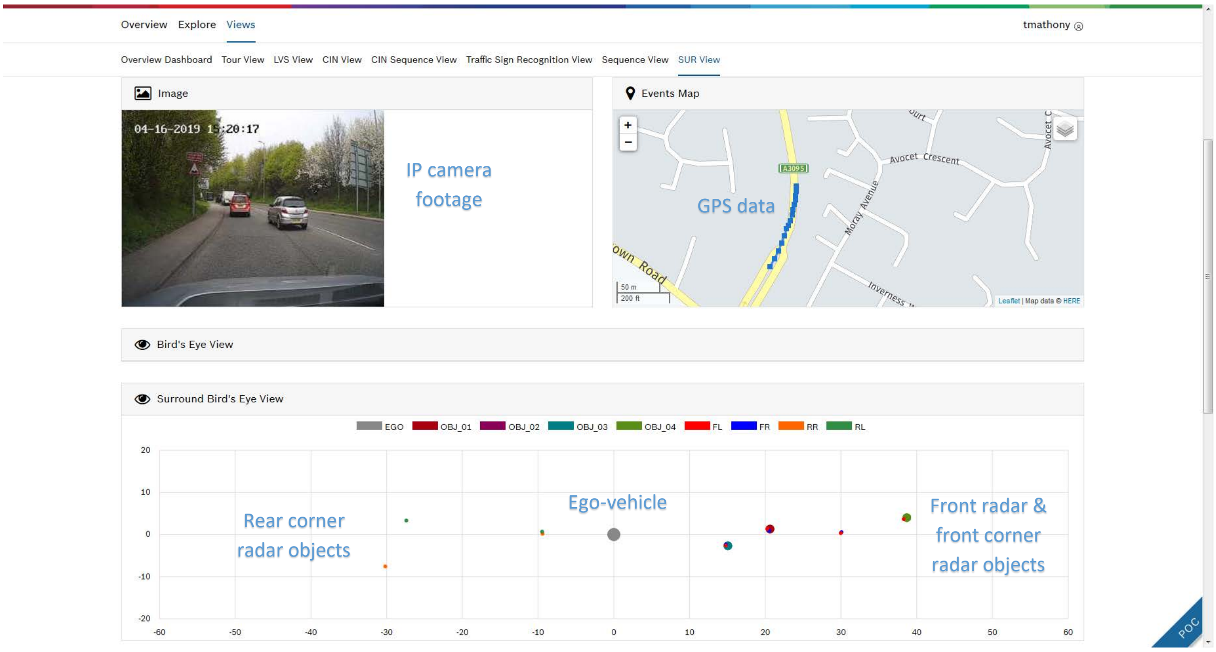

Figure 9 – Detailed event view for a SUR-CIN event

The objects detected by the four corner radar sensors were differentiated by using different colours for each (i.e. red – front left radar, blue – front right radar, orange – rear right radar, and green – rear left radar).

2.3.3 Enterprise Automotive Data Management (EADM)

As described more fully in the Phase 1 and 2 reports (*MOVE_UK Data Analysis Report - Phase 1 (D7.3)* and *MOVE_UK Data Analysis Report - Phase 2 (D7.4)*), EADM is a solution for the management and analysis of automotive measurement data and metadata, developed by ETAS (a Bosch subsidiary company).

In Phase 3 of MOVE_UK, EADM continued to be used to search and analyse data using the techniques developed in Phase 1 and 2. It was also used extensively by The Floow, ETAS and Bosch to carry out further and more advanced data analytics using R Language, MATLAB and Python, respectively.

No notable changes were made to EADM since the release of the phase 1 report (*MOVE_UK Data Analysis Report - Phase 1 (D7.3)*.

3 Update on Phase 1 and 2 use cases

3.1 Overview

As described in the Phase 1 report, *MOVE_UK Data Analysis Report - Phase 1 (D7.3)*, use cases were defined within the project which demonstrated different capabilities and related to the various different applications (which were envisaged for the developed methodologies). A use case consists of a description of the events considered to be relevant to a particular capability (the description being used to define the trigger condition for video sequence or dense CAN data capture) and a method for data analysis to fulfil the purpose of the use case.

Eight use cases were developed and implemented previously, in MOVE_UK Phases 1 and 2:

Phase 1:

1. Subcritical camera based Autonomous Emergency Braking (AEB);
2. Driver Harsh Braking (DHB);
3. Traffic Sign Recognition (TSR); and
4. Telematics 1.

Phase 2:

5. Subcritical Radar AEB Activation (ARB);
6. Cut-In scenarios (CIN);
7. Lead Vehicle Statistic (LVS); and
8. Telematics 2.

For more detail on the purpose and design of these use cases, please refer to the *MOVE_UK Data Analysis Report - Phase 1 (D7.3)**Error! Bookmark not defined.*** and *MOVE_UK Data Analysis Report - Phase 2 (D7.4).*

This section provides an update on the number of events recorded for these use cases during Phase 3, and during the entire project. In addition:

- For AEB, the two most-relevant sequences captured in Phase 3 are reviewed, and the results of the final sequence re-simulation work are presented;
- For DHB, the final distribution of reasons for harsh braking manoeuvres is revealed and results of a time to collision evaluation and an AEB false negative analysis (at detection level) are shown;
- For TSR, the results of further statistical analysis are presented, a new “big data” approach to processing data is introduced, and the results of a signage correction exercise carried out in Phase 3 revealed;
- For both ARB and CIN, a number of interesting events collected in Phase 3 are reviewed and the results from both Phases 2 and 3 are summarised;
- For LVS, a summary of the results and conclusions from Phases 2 and 3 is provided; and
- For Telematics 1 and 2, a final update on the event data recorder (EDR) analysis and risk method is provided.

3.2 Use case Subcritical Camera Based Autonomous Emergency Braking (AEB)

3.2.1 AEB events captured in Phase 3

In Phase 3 of MOVE_UK 38 events were captured within the Subcritical AEB use case during the course of an eight-month period (1 Oct 18 – 31 May 19). In total during all three phases of the project, 80 events were captured. The following still frames (Figure 10 and Figure 11), which were captured during Phase 3 of the trials, show different examples of the camera view at the moment of a Subcritical AEB trigger. The red arrow highlights the object which caused the trigger (the colour itself is not significant in this context). The blue boxes highlighting the vehicles, road users and roadside objects that have been recognised are created 'live' by the camera system (are not added during post-processing). The graph below the camera view indicates the moment of the Subcritical AEB trigger.

Figure 10 – Example 1 (Subcritical AEB Sequence no. 56): Braking car in front due to approaching a junction

In Figure 10, a MOVE_UK vehicle is driving behind another vehicle whilst approaching a junction. First the leading vehicle brakes slowly while it still has a relatively large distance to the junction, but then it suddenly brakes strongly, surprising the driver of the MOVE_UK vehicle. Due to the leading vehicle braking harshly, the driver of the MOVE_UK vehicle has to brake harshly too. During this sudden braking manoeuvre, when the time-to-collision drops, the stereo video camera detected and correctly classified the leading vehicle and activated AEB (in silent mode). The rear view of the vehicle in front was detected and caused the trigger.

Figure 11 displays another exemplary AEB activation from Phase 3. A MOVE_UK vehicle is driving when another vehicle (shown) cuts in front of it from a side road/junction. Immediately after cutting in, the now leading vehicle brakes harshly to turn right. The driver of the MOVE_UK vehicle therefore also brakes harshly. During this sudden braking manoeuvre, when the time-to-collision drops, the stereo video camera detected and correctly classified the car and activated AEB (in silent mode). The rear view of the vehicle in front was detected and caused the trigger.

Figure 11 – Example 2 (Subcritical AEB Sequence no. 67): Car immediately brakes after cutting in

3.2.2 Re-simulation of captured events

The re-simulation of all AEB events captured during the MOVE_UK Phase 1, 2 and 3 trials showed that a production version of the video camera based AEB system would have triggered in a large number of the situations captured. Out of the 80 AEB events captured, 23 were classified as critical AEB events by the latest Bosch production camera software, while one potential false positive was detected and correctly suppressed by the mono vehicle rear classification.

3.2.3 Conclusions

The results of re-simulation of the captured AEB events with production camera software were far better than expected and highlight the relevance of the data collected from the MOVE_UK trials. Critical AEB events, such as those collected in MOVE_UK, are very difficult to capture in the real world and are extremely useful for the development and validation of future ADS features. The main reason for this is that they often cover driving situations which the engineers/developers did not foresee and also a variety of situations which testing bodies (like Euro NCAP) do not include in their test protocols. Using these events/sequences to test ADS software in a simulation environment, therefore, leads to more effective and safer ADS features. To highlight the relevance of the AEB events captured in MOVE_UK, it is the intention of Bosch to use the 23 critical events captured to help test and develop its next generation of production AEB software.

3.3 Use case Driver Harsh Braking (DHB)

3.3.1 DHB events captured in Phase 3

In Phase 3 of MOVE_UK, 89 DHB events were captured between 1 October 2018 and 30 April 2019, which brings the total number of DHB events up until end April 2019 to 207. In addition, 8 events were captured by the AEB or ARB triggers, which also fulfilled the DHB criteria (the first trigger that occurs determines the type of recording). A monthly breakdown of events recorded throughout the duration of the project is given in Table 2.

Table 3 relates the number of harsh braking events observed to the miles driven during each month. This allowed the consortium to determine how frequently harsh braking was observed during real-world driving under the trial conditions.

Table 2 – Number of harsh braking events captured in each month

Month, Year	Phase	Number of DHB events captured		
		Primary trigger: DHB	Primary trigger: AEB	Primary trigger: ARB
August 2017	Phase 1	7	–	–
September 2017	Phase 1	8	–	–
October 2017	Phase 1	14	–	–
November 2017	Phase 1	5	–	–
December 2017	Phase 2	1	–	–
January 2018	Phase 2	3	–	1
February 2018	Phase 2	2	–	–
March 2018	Phase 2	2	–	–
April 2018	Phase 2	13	2	–
May 2018	Phase 2	15	–	–
June 2018	Phase 2	20	–	–
July 2018	Phase 2	18	–	–
August 2018	Phase 2	3	–	–
September 2018	Phase 2	7	–	–
October 2018	Phase 3	15	1	1
November 2018	Phase 3	7	–	–
December 2018	Phase 3	10	–	1
January 2019	Phase 3	14	–	2
February 2019	Phase 3	8	–	–
March 2019	Phase 3	17	–	–
April 2019	Phase 3	18	–	–
Total number	**215**			

Table 3 – Number of harsh braking events recorded and miles driven

Month, Year	Total number of events involving harsh braking	Miles driven	Number of harsh braking events per 1,000 miles
August 2017	7	4,516	1.6
September 2017	8	2,988	2.7
October 2017	14	3,266	4.3
November 2017	5	1,773	2.8
December 2017	1	2,565	0.4
January 2018	4	4,061	1.0
February 2018	2	3,902	0.5
March 2018	2	4,411	0.5
April 2018	15	5,352	2.8
May 2018	15	4,161	3.6
June 2018	20	3,746	5.3
July 2018	18	3,646	4.9
August 2018	3	2,336	1.3
September 2018	7	4,697	1.5
October 2018	17	5,796	2.9
November 2018	7	4,538	1.5
December 2018	11	4,011	2.7
January 2019	16	6,577	2.4
February 2019	8	4,289	1.9
March 2019	17	5,994	2.8
April 2019	18	5,595	3.2
Total	**215**	**88,220**	**2.4**

It can be seen in Figure 12 that the frequency of harsh braking events throughout the project varied widely between 0.4 and 5.3 events per 1,000 miles driven. The month where certain events were missed due to a software misconfiguration (between November 2017 and March 2018; see MOVE_UK *Data Analysis Report - Phase 2 (D7.4)*) can be identified clearly in the graph. Without these outliers taken into account the range of frequency observed was between 1.3 and 5.3 harsh braking events per 1,000 miles driven.

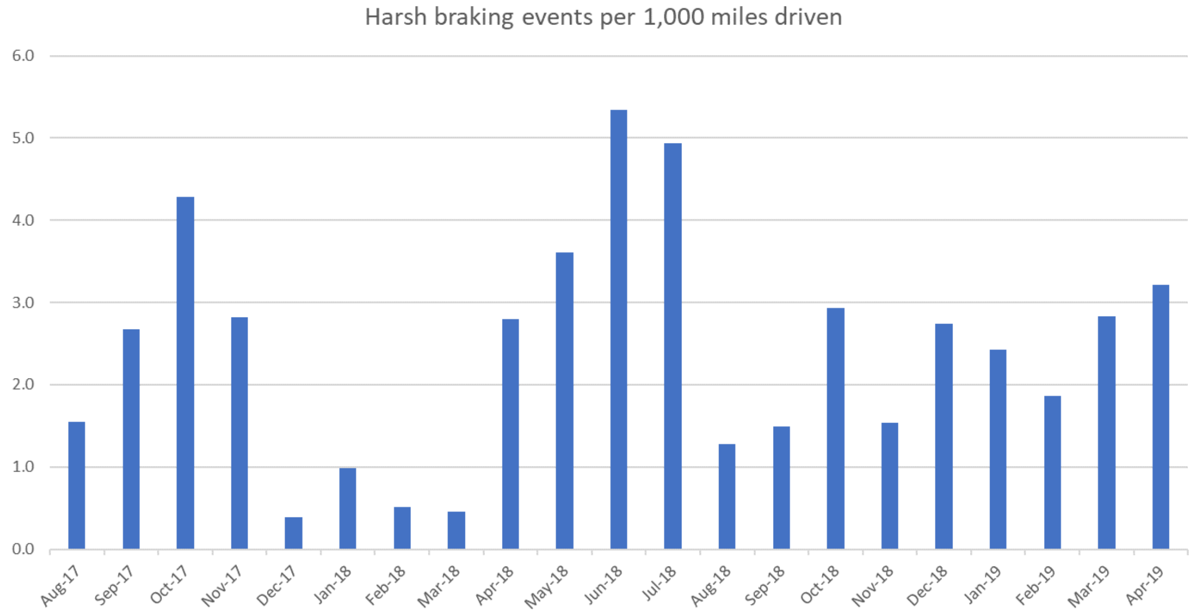

Figure 12 – Frequency of harsh braking events (number of events observed per 1,000 miles driven)

3.3.2 Reasons for drivers performing harsh braking manoeuvres – update

3.3.2.1 Method

As in the previous project phases a review of the captured DHB videos was undertaken to determine the likely reasons for drivers braking harshly in each of the events. The method for classifying events was developed further during Phase 3 and retrospectively applied to sequences from all phases. The final method applied is described below.

A systematic classification of DHB events (and AEB/ARB events that also fulfilled DHB trigger criteria) was undertaken to understand the *reasons* and *potential consequences* of such events (e.g. impact with a vehicle, vulnerable road user or other object). This analysis was undertaken to develop a better understanding of how critical events can be avoided by auto-braking or other driver assistance technologies.

The systematic classification aims to distinguish between two main types of events:

- TYPE 1: Events where harsh braking is used to avoid imminent impact with another vehicle, VRU or other object; and
- TYPE 2: Events where harsh braking is used as dictated by Highways Code rules without there being an imminent risk of colliding with another vehicle, a vulnerable road user (VRU; such as a pedestrian or cyclist) or other object.

Additionally, the classification distinguishes between different types of road layout where such events take place.

For a full description of TYPE 1 and TYPE 2 DHB events, please refer to Appendix 1.

3.3.2.2 Results

A total of 215 relevant events (207 DHB events plus 8 AEB and ARB sequences that also matched DHB trigger criteria) were classified using the above method throughout the duration of the project.

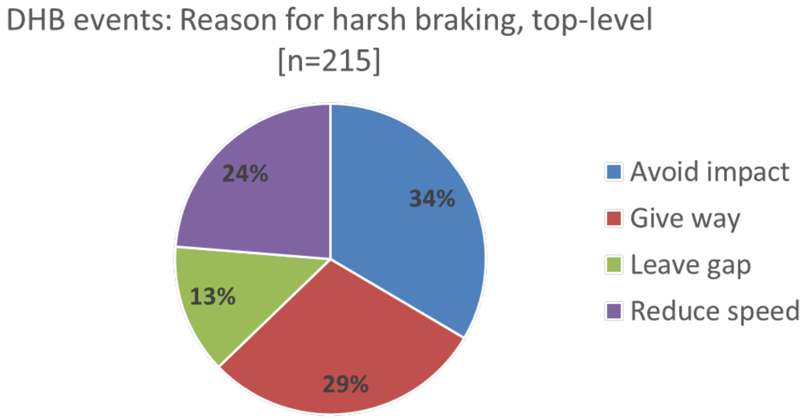

Figure 13 – Breakdown of 215 relevant events by reason for harsh braking

From the breakdown provided in Figure 13 it can be seen that 'avoid impact' (Type 1 events) represents the largest group with about a third of the events observed (72 of the total 215).

This group is followed closely by 'give way' situations with pedestrians or other vehicles, which were not deemed as avoidance of an imminent collision. 63 of the total 215 events fell into this category.

General 'reduce speed' situations, for instance for upcoming turns or speed humps, represented 51 of the recorded events.

Leaving a gap for another vehicle, mostly to pass in the opposite direction through a narrow road section, was found to be the reason in 29 of the events.

Figure 14 provides a further breakdown of all 'avoid impact' events to visualise the type of target for which the driver braked harshly. Approximately two thirds of events (48 of the total 72) involved a conflict with another car, followed by vans (7 events), HGVs (6 events) and pedestrians (4 events).

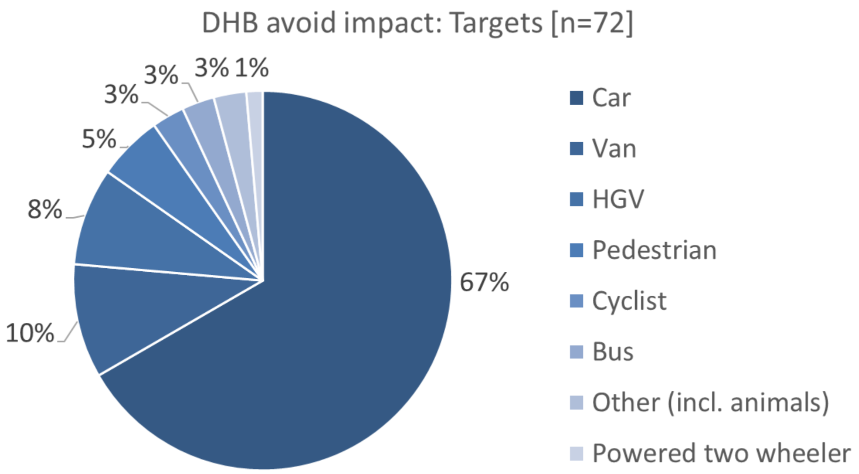

Figure 14 – Breakdown of the 72 Type 1 events (avoid impact) by the type of target that was avoided by harsh braking

When all motorised vehicle targets are grouped, these account for 89% (64 events) of all avoid impact events recorded, whereas pedestrian, cyclists and animals represent only 11% (8 events).

The review of events highlighted another location (in addition to that identified in *MOVE_UK Data Analysis Report - Phase 2 (D7.4)*) where two relevant events occurred. Figure 15 shows the view ahead during these events. The AEB sequence was caused by a car in front abruptly braking for the upcoming crossing; the DHB event was triggered by the MOVE_UK vehicle braking harshly in front of the crossing to let pedestrians pass.

AEB event 35 – April 2018 | DHB event 74 – June 2018

Figure 15 – AEB sequence 35 and DHB sequence 74. Still frames extracted from the event videos.

The road layout ahead of the crossing is shown in Figure 16. It can be speculated that the cause for both events is the fact that the pedestrians waiting to use the crossing itself were spotted late by drivers. It is unclear whether the road layout with a bent road and an adjoining side road before the crossing have contributed to these events, but the events observed might warrant further observation of this location.

Figure 16 – View of the road layout ahead of the pedestrian crossing, approximately three second before DHB event 72

3.3.3 Time to Collision (TTC) evaluation

3.3.3.1 Method

The following section details the method that has been applied to calculate the minimum TTC for DHB, AEB and ARB sequences. The TTC between two vehicles represents the time until a collision would occur if both vehicles continued moving at their current speeds on their current trajectory (i.e. the time remaining before a collision if none of the vehicles applied their brakes or steering to avoid a collision). TTC is calculated from the distance between two vehicles and their closing speed.

For the purpose of this analysis, sequences were only considered where the following criteria were satisfied:

1. The braking event occurred to avoid an impact between the ego vehicle and the vehicle driving in front of it (lead vehicle);
2. The lead vehicle was detected by the ego vehicle's CAN system as the most relevant object at the key point in the sequence, namely the point where the driver applied the maximum level of brake pressure; and
3. The point at which the maximum level of brake pressure was applied occurred either after or in the second before the event trigger (15 seconds into the sequence).

Sequences identified as not meeting these criteria were excluded from the analysis. These included:

1. DHB sequences where the driver had to brake for reasons other than to avoid impact with a vehicle, such as traffic lights, giving way or a pedestrian crossing;
2. False positive AEB or ARB sequences where there was no target vehicle in front of the ego vehicle;
3. Sequences where the maximum level of brake pressure occurred more than one second before the event trigger; and
4. Sequences where at the point of maximum brake pressure, an object other than the vehicle in front was detected by the CAN system as OBJ_01 (i.e. as the target object of most relevance; see *MOVE_UK Data Analysis Report - Phase 2 (D7.4) for definition of target objects*). Typically, this was an object at the side of the road, such as a parked car or bollard.

One important limitation of this analysis was identifying sequences that were covered by point 4. In this case, the analysis was done by reviewing videos of sequences where the angle from the ego vehicle to the centre of OBJ_01, at the time of maximum braking, was comparatively large (of magnitude greater than 11.5 degrees). This suggested that an object on the side of the road had mistakenly been detected as OBJ_01. Applying this rationale excluded 14 sequences from further analysis. It should be noted, however, that this approach is not automated and therefore would not be suitable for analysing data from a larger vehicle fleet. In addition, it does not cover any potential cases where OBJ_01 was not a lead vehicle but was still in front of the ego vehicle, such as a pedestrian crossing the road, rather than at an angle to it. Therefore, without performing a full manual video review, it cannot be completely guaranteed that OBJ_01 has been correctly detected in every remaining sequence.

More information about defining the key time period and calculating the time to collision can be found in Appendix 1.

3.3.3.2 Distribution of Minimum TTC

Figure 17 presents histograms showing the distribution of the minimum TTC for the different classifications of sequence; All; DHB only, AEB only and ARB only. As explained in section A1.2.2, the minimum TTC was defined as the minimum TTC between time points 1 and 2, only considering the frames where OBJ_01 was detected correctly.

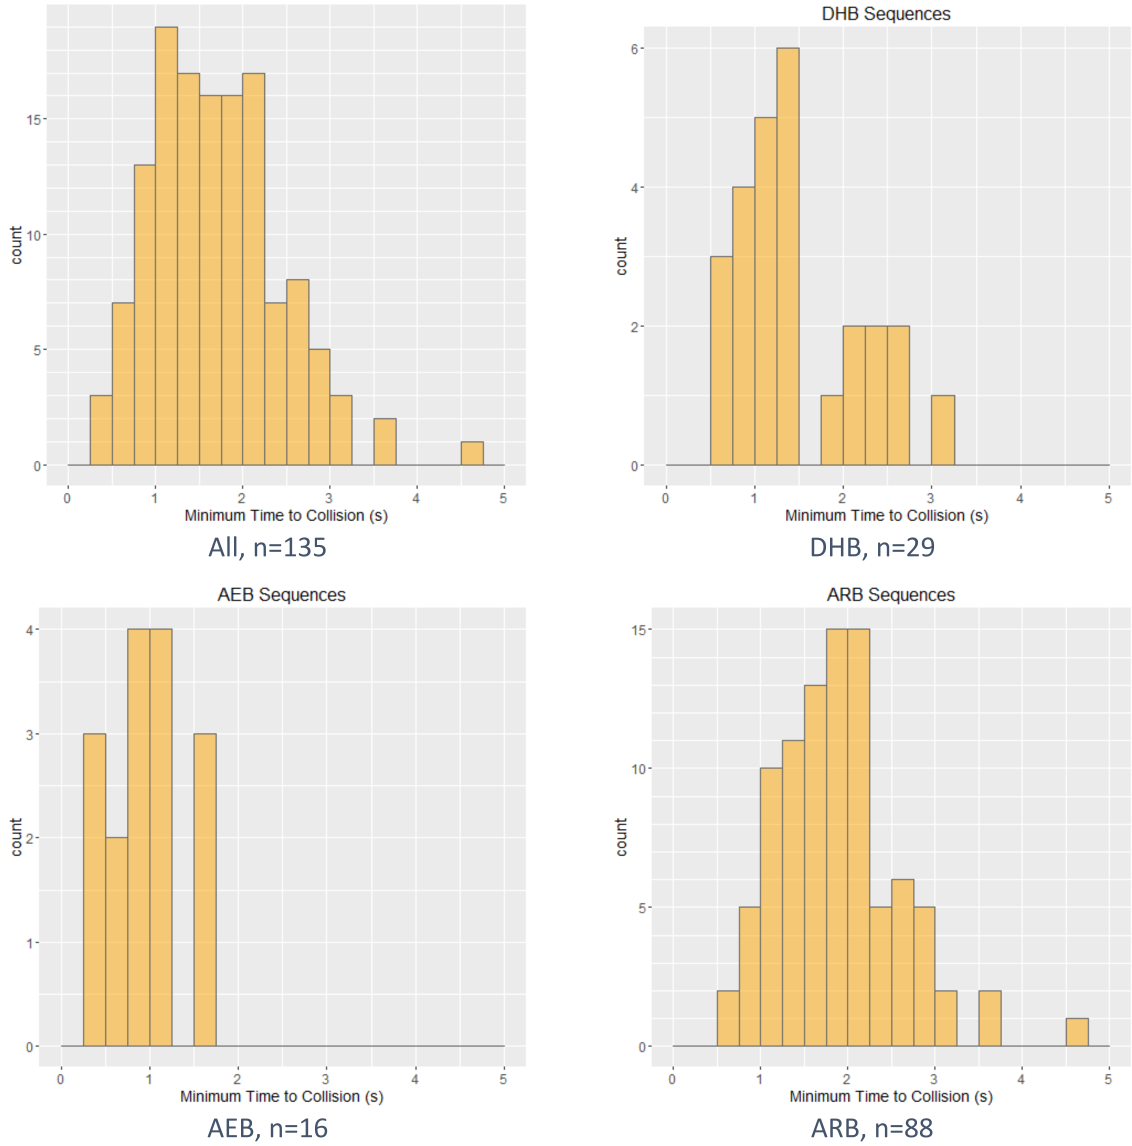

Figure 17 – Distribution of minimum TTC

The histogram for all sequences shows that most of the sequences had a minimum TTC of between 1-2.5 seconds. However, there are a small number of more critical sequences with a minimum TTC of less than a second.

Most of the DHB and AEB sequences had a minimum TTC of between 0.5-1.5 seconds, and were therefore more critical than ARB sequences, where the minimum TTC was mainly between 1-2.5 seconds. This suggests that the ARB trigger is more sensitive than the AEB trigger. In addition, statistical tests[2] showed that there were significant differences in minimum TTC between AEB and ARB sequences[3].

[2] Two-sample t-tests

[3] t=7.358, p<0.001

3.3.3.3 Observations

The top panel of Figure 17 shows that there were seven DHB sequences with a minimum TTC of less than a second. However, none of these sequences were classified as being subcritical AEB or ARB events. It should be noted that in four of six of these sequences, there was some part of the key time period where the ego vehicle was not travelling directly behind the target vehicle (OBJ_01). AEB systems are not designed to activate in such a scenario where there is little or no lateral overlap between the two vehicles (i.e. the vehicles are not travelling directly in line with one another). One of the other two sequences also had a low amount of lateral overlap (33%). In the remaining sequence, the ego vehicle braked behind an HGV at a junction, but the HGV was not detected as a vehicle during the key time period.

3.3.3.4 Conclusion

The methodology described here gives a means of calculating TTC for braking events which are of relevance to AEB systems. The distribution of values presented corresponds with what would be expected to be seen in real world driving. In addition, there are justifiable reasons why any critical DHB events might not trigger a subcritical AEB activation, with the most apparent one being an absence or low level of overlap at least during parts of the key time period.

However, a significant limitation with the outlined approach is the regularity with which the CAN system did not identify the most relevant object as OBJ_01 at the point of maximum brake pressure, or at some point during the key time period. The first case in particular is not something that can be identified using an automated procedure, but only by reviewing the videos from each individual sequence. In the future, if there was a desire to calculate TTC for a larger sample of events, this issue would be a major barrier to overcome, unless ADAS systems such as AEB can be adapted to reduce the likelihood of identifying objects with little or no relevance to the braking event, in particular those on the side of or off the road.

3.3.4 Event cluster analysis – update

In Phase 2, an event cluster analysis was performed on the DHB events to investigate whether CAN signal data associated with harsh braking events could tell us anything about human braking behaviour. The method and results can be found in section 3.3.3 of the Phase 2 report. The analysis identified four distinct clusters of DHB events; however, it was difficult to associate these clusters with particular driving situations, due to the small sample size (35 events in total). In Phase 3, the cluster analysis was repeated, using DHB events from across the whole of the project, to see if increasing the sample size would lead to some clearer patterns emerging in the results.

3.3.4.1 Variables of interest

Table 4 shows the variables which were considered for inclusion in the analysis, as well as a description of how each variable was measured, using data from the CAN signals on the MOVE_UK vehicles. It should be noted that *Steering Wheel Angle* was measured in absolute terms, meaning that no distinction was made between steering to the left or right. Instead, the focus here was on the magnitude of steering.

Table 4 – Initial list of variables considered for inclusion in the event cluster analysis

Variable	Description
Maximum Brake Pressure	Maximum recorded brake pressure during braking event
Brake Time to Max	Time between start of braking event and point of maximum brake pressure
Rate of Change in Brake Pressure	$\frac{\text{Maximum Brake Pressure}}{\text{Brake Time to Max}}$
Speed	Speed of car at start of braking event
Change-over Period	Time between the driver releasing pressure from the accelerator pedal and the start of the braking event
Steering Wheel Angle	Absolute steering wheel angle at start of braking event
Maximum Deceleration	Maximum deceleration during braking event
Hour	Hour of the day during which event took place

There were some differences between the variables considered in this analysis and those considered in Phase 2. Firstly, the total braking time and the distance to the object in front were no longer included. It was decided that the time period measured by *Brake Time to Max* was more relevant than the total braking time, which would include, for example, time after the point of maximum brake pressure when the driver might be travelling very slowly or even be stationary. The distance to the object in front was removed, as more in-depth analysis of the CAN data (see Work Package 8 final report - deliverable D8.3) showed that the object that was detected was not always relevant, and in particular was not always in the driver's line of sight (e.g. a parked car or bollard at the side of the road). Therefore, data from this variable was not considered to be sufficiently reliable.

Secondly, it was decided that *Maximum Deceleration* and time of day, represented by *Hour*, might add some further insight to the variables already included, although it was acknowledged that *Maximum Deceleration* is closely related to a combination of *Maximum Brake Pressure, Rate of Change in Brake Pressure* and *Speed*, and therefore there would be a need to reduce the set of variables in Table 4.

Figure 18 shows the level of correlation between each pair of variables (higher levels of correlation are shown by darker colours). A high level of correlation (>0.7) was found between *Maximum Brake Pressure* and *Maximum Deceleration*. It was decided to keep *Maximum Brake Pressure* and remove *Maximum Deceleration* from the analysis, given that this variable is a consequence of *Maximum Brake Pressure, Rate of Change in Brake Pressure* and *Speed*. These three variables would provide more information on the details of driver behaviour for a particular event. In addition, a moderately high correlation (>0.6) was found between the *Rate of Change in Brake Pressure* and *Brake Time to Max*. It was decided to keep *Rate of Change in Brake Pressure* in the analysis, as this was considered to be of more closely associated with driver behaviour, and remove *Brake Time to Max*.

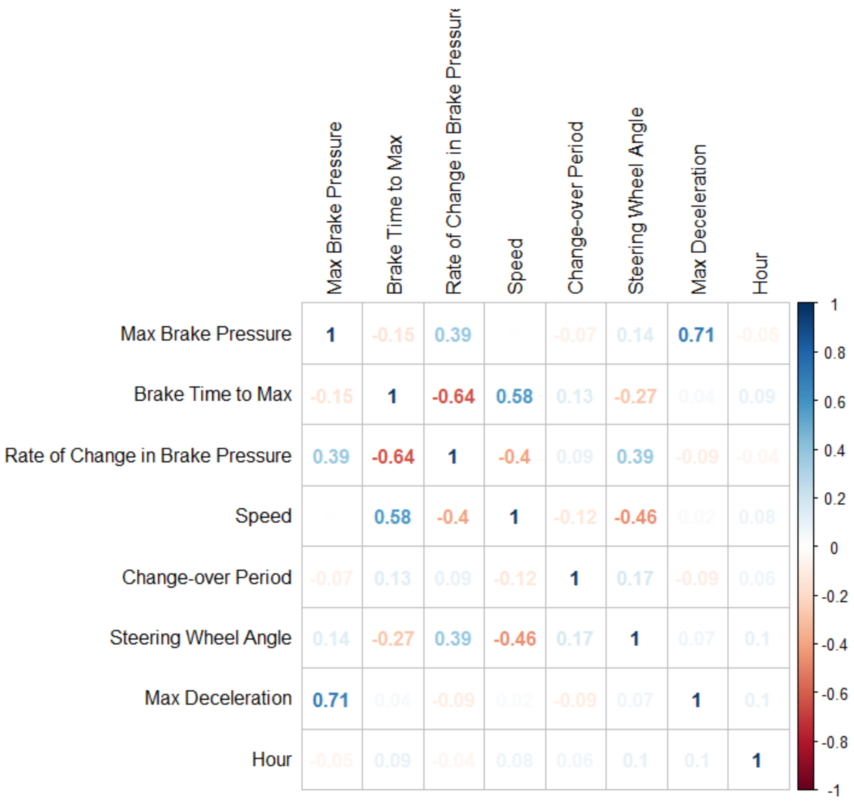

Figure 18 – Plot of correlations between variables

Table 5 shows the final list of variables that progressed into the cluster analysis.

Table 5 – Final list of variables included in the event cluster analysis

Variable	Description
Maximum Brake Pressure	Maximum recorded brake pressure during braking event
Rate of Change in Brake Pressure	$\frac{\text{Maximum Brake Pressure}}{\text{Brake Time to Max}}$
Speed	Speed of car at start of braking event
Change-over Period	Time between the driver coming off the accelerator pedal and the start of the braking event
Steering Wheel Angle	Steering wheel angle at start of braking event
Hour	Hour of the day during which event took place

3.3.4.2 Comparing clusters with reasons for braking

Figure 19 shows the distribution of the reasons for braking within the two clusters. For simplicity, the reasons have been grouped into four categories; avoiding an impact, giving way, leaving a gap and reducing speed. The results show a couple of differences:

1. A higher proportion of events in Cluster 2 occurred in order to avoid an impact, compared to those in Cluster 1, and;
2. A lower proportion of events in Cluster 2 occurred because the driver was leaving a gap, compared to Cluster 1

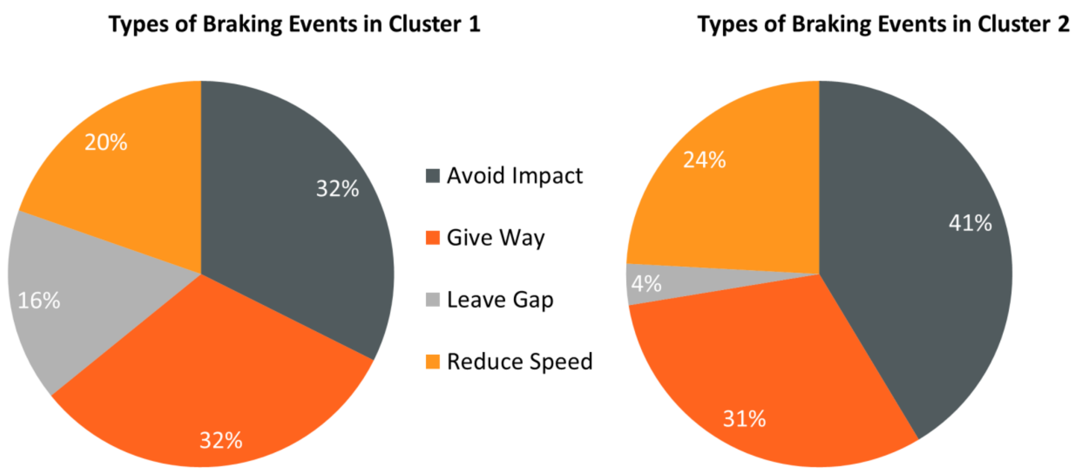

Figure 19 – Reasons for braking within the two final clusters

3.3.5 Conclusions

During four months of Phase 2, a software misconfiguration of the Flea3 box resulted in a number of DHB events not triggering a recording. Considerable effort was then involved in identifying that an incorrect software configuration was causing the issue. If silent connected validation trials were applied in an approval setting, it would be important for the technical service to be able to determine if no events being recorded in a certain phase is indeed an indicator of there being no events or if a failure in the recording equipment is responsible. The frequency of DHB events (i.e. recorded events per 1,000 miles driven; calculated on a monthly basis) was seen to drop to exceptionally low levels in this period (0.4 events per 1,000 miles driven). Monitoring this figure for sudden drops throughout a trial may alert to potential recording issues. However, also in phases with correct configuration the frequency varied widely between 1.3 and 5.3 events per 1,000 miles driven, which makes the frequency not suitable as a sole indicator of any failures or misconfigurations throughout the trial. Therefore, thorough definition and oversight of the experimental setup and review of technical records (at the initial trial setup and for any configuration or hardware changes) by the technical service will still be required to ensure consistency of approvals.

The analysis of reasons for harsh braking manoeuvres showed that 34% of DHB events recorded fell into the ‘avoid impact’ category, meaning that these events would have resulted in an impact with another road user or object had the driver not undertaken the manoeuvre. These events, together with AEB events, form the basis for a comparison between the test cases covered in regulatory type approval tests for AEB systems and those encountered in the real world. This analysis was carried out as part of project Work Package 8 (see WP8 final report - deliverable D8.3) and demonstrates that recording of DHB events can be used to highlight potential gaps in regulatory tests and devise future test setups for AEB systems based on real-world data.

An analysis of the minimum time to collision (TTC) of DHB events was undertaken, which is an indicator of how critical the observed events were. This allowed a comparison between DHB events and those events triggering an AEB or ARB activation. There were seven DHB sequences with a minimum TTC of less than a second, which indicates highly critical events for which an AEB system activation might be desired. However, none of these sequences were classified as being subcritical AEB or ARB events. It should be noted that in four of six of these sequences, there was some part of the key time period where there was no lateral overlap between the ego vehicle and OBJ_01 (the ego vehicle was not travelling directly behind OBJ_01), a scenario that AEB systems are not designed to activate in. One of the other two sequences also had a low minimum overlap (33%). In the remaining sequence, the ego vehicle braked behind an HGV at a junction, but the HGV was not detected as a vehicle during the key time period. Observations like these can inform the AEB test scenario analysis in WP8.

The event cluster analysis, initially performed in Phase 2 of MOVE_UK, was repeated with a refined method on a larger set of events from all three project phases. Two clusters sharing similar characteristics emerged based on five event variables: Maximum brake pressure, rate of change in brake pressure, speed at start of braking event, change-over period and steering wheel angle. Associating these clusters with the reasons for braking showed that braking events in order to leave a gap for another road user to pass were mostly found in cluster 1 and braking to avoid an impact was more likely associated with the characteristics of cluster 2 although avoid impact events were common in both clusters. However, the results do not allow a clear association of the reason for braking with the braking characteristics, which means that a prediction of the reason for braking using this analysis method alone does not appear feasible based on the data set gathered.

3.4 Use case Traffic Sign Recognition (TSR)

The purpose of the Traffic Sign Recognition (TSR) use case was to develop big data analysis methodologies relevant to ADS with a focus on location-based data.

In Phase 1 of the project, the following initial steps were taken:

- The TSR use case was commissioned by selecting the vehicle and sensor signals that were relevant to the use case;
- An interactive user interface (sFDE – Systematic Field Data Exchange) containing a TSR view was developed to display the spread of traffic signs detected by the MOVE_UK fleet vehicles on a map; and
- A clustering methodology was developed within the sFDE TSR view to calculate the actual and missed detections.

Then, in Phase 2, more focussed investigation and analysis was carried out, as follows:

- Statistical analysis was performed on the data collected in sFDE to understand the effect of external and vehicle related factors on both the probability of detecting a speed limit sign, and the probability of detecting the correct speed; and
- Data was clustered based on the location of traffic signs, in order to identify detection rates at specific signs. From the analysis, a number of hotspot locations were identified where the TSR system had a poor detection rate. A site visit was conducted at each of these locations to uncover the issues associated with poor detection of a traffic sign.

Phase 3 included the following activities, which are described in more detail in the following sections:

- Further statistical analysis was performed on a larger dataset;
- The clustering methodology from Phase 2 was improved;
- Rectification actions were performed at one of the seven hotspot locations identified in Phase 2; and
- An interactive dashboard was produced to provide an interesting and interactive means of visualising the outputs of the analysis.

3.4.1 Further statistical analysis

In Phase 2 of the project, a statistical methodology was developed to identify factors affecting the probability of detections (the system detecting a road sign when one is present) and the probability of true detections (the system detecting the correct speed on a detected speed limit sign). More details of the methodology can be found in section 3.4.1 of the Phase 2 report.

In Phase 3, this methodology was applied to all the TSR data from both the Greenwich and Bracknell (TRL) areas over the length of the project, incorporating some improvements from Phase 2. The two key improvements were:

1. The inclusion of direction of travel in the clustering step; and
2. The inclusion of cross-validation techniques to prevent overfitting (including too many factors in the model), and to improve the predictive ability and robustness of the final statistical models

3.4.2 Data processing

The data was cleaned in four stages before beginning the analysis:

1. Any uncertain or missing data was removed. This included data where the GPS uncertainty error was above 100 metres; the direction of movement was undeterminable; the road type was unknown or the speed of the vehicle was less than 6 km/h (as the traffic sign recognition system used in MOVE_UK does not work below this speed).
2. Data from the two 'depots' in Greenwich and Bracknell (TRL) was removed, as the layouts of these locations were substantially different to that of standard roads and included a number of traffic signs within a small area.
3. A number of sanity checks were performed on the clusters in order to ensure that pairs of clusters were not too close together - which could indicate that these clusters might correspond to the same traffic sign. Clusters were combined where the following conditions held:
    - Their most common detected speeds were the same.
    - Their most common directions of travel were less than 45 degrees apart.
    - Their centre points (based on latitude and longitude coordinates) were less than 200m apart.
    - The shortest distance between points in the two clusters was less than 50m.
4. A second level of sanity checks was performed on the missed detections (where the MOVE_UK vehicle drove through a cluster location but did not detect a sign). Data was removed for missed detections where the direction of travel was more than 45 degrees away from the average direction of detections in the corresponding cluster. This was to ensure that when a missed detection was recorded, the vehicle did actually drive past a sign, and was not travelling in the opposite direction or turning down a different road.

The initial dataset included 100,447 journeys and 1,756 clusters, and the cleaning and filtering steps above reduced it to 48,254 journeys and 893 clusters. The most common reasons for data points being removed were because the road type was unknown (32%); the location was at one of the two 'depots' (23%); the direction of movement was undeterminable (15%) or the direction of travel of a missed detection meant it was removed according to point 4 above (14%).

Two new datasets were generated from the final cleaned dataset:

1. Detection/missed detection dataset (46,071 data points including 40,823 detections and 5,248 missed detections from 893 clusters); and

2. True Detection/False Detection dataset (43,006 data points including 40,823 true detections and 2,183 false detections).

3.4.2.1 Factors influencing detection/missed detection probability

In Phase 2, the detection/missed detection dataset was used to build a statistical model to determine which factors influenced the probability of detecting a speed sign, and by how much. This identified four fixed factors as being influential: vehicle speed, distance from the vehicle/object in front (OBJ_01 in the CAN data), windscreen wipers (on/off) and road type. In addition, cluster identity also had a significant influence on the probability of detection. This was classified as a random effect in a "mixed-effects" model. Such a model can separately consider the variability caused solely by the different cluster locations. This variability is understood and then taken into account. This ensures that results from the model – the extent to which each fixed factor influences the probability of detection - are not inappropriately influenced by the variability across the different clusters.

In Phase 3, the same methodology was applied to the new dataset, which was now considerably larger in size (46,071 data points compared to 6,122 in Phase 2). However, the distance to the object in front (OBJ_01) was not considered as a potential influential factor. This was because more in-depth analysis of the CAN data (see Work Package 8 final report - deliverable D8.3) showed that the object that was detected was not always relevant, and in particular was not always in the driver's line of sight (e.g. a parked car or bollard at the side of the road). Therefore, data from this variable was not considered to be sufficiently reliable.

Figure 20 shows that the updated statistical model identified five factors as being influential on the probability of detection. Two of these factors, speed and road type, were also identified as being influential in Phase 2. However, the other three factors, acceleration, time of day and the low beam headlight (on/off), were not. Windscreen wipers were no longer identified as an influential factor. Average detection rates across the dataset were high at 89%, although slightly lower than in the smaller dataset in Phase 2.

Once again, speed was the most influential factor. A 10km/h drop in speed (relative to the average) reduced the probability of detection by 0.36%. Road type was the next most influential factor, with road types 4 (connective roads) and 5 (slow urban roads) associated with reductions in detection probability of 0.31% and 0.10% respectively, compared with journeys made on road type 3 (major roads, excluding motorways and trunk roads). These results indicate that cars travelling on lower speed roads, where signage tends to be smaller, were less likely to detect speed signs than those on higher speed, more major roads with more prominent signage.

Acceleration, the time of day and whether or not the low beam headlight was on also had a small but statistically significant influence on detection probability. A 1 m/s^2 increase in longitudinal acceleration (relative to the average) was associated with a 0.09% reduction in detection probability. Each increase of 3 hours later in the day (from a baseline of 11am) was associated with a 0.04% reduction and having the low beam headlight on with a 0.03% reduction (compared with the headlight being off), suggesting that lower light levels reduce the probability of detection.

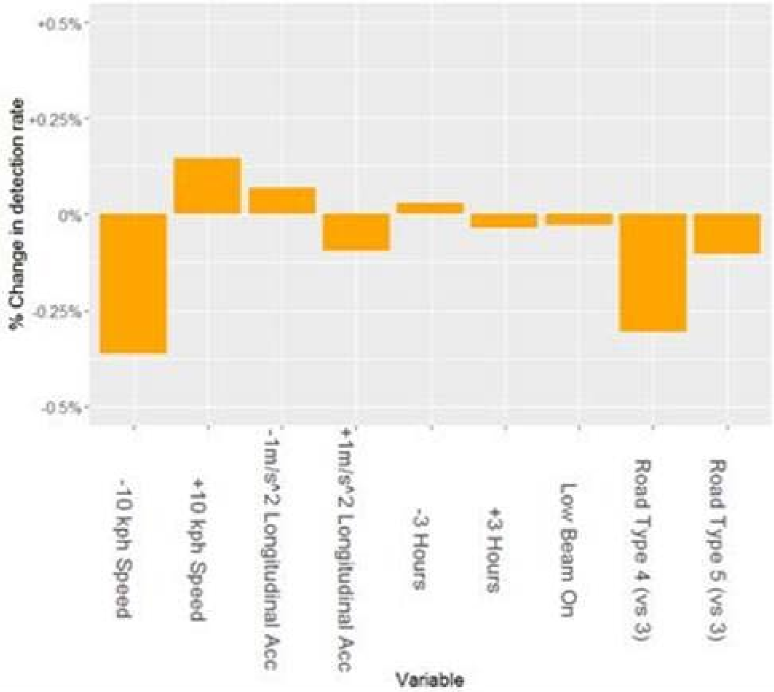

Figure 20 – Influence of fixed factors on probability of detecting a speed sign

As well as the effects already discussed, cluster identity also had a significant influence on the probability of detection. Out of the 893 clusters in the data, 671 (75.1%) had a detection rate of more than 99%, indicating that the TSR system was performing well in these locations. However, there were 47 clusters (5.3%) with a detection rate of between 50% and 80%, and for 25 clusters (2.8%) it was less than 50%, indicating that the TSR system was not performing well in these locations.

Overall, it can be seen that none of the fixed factors had a particular large impact on the detection probability (maximum change of 0.36% associated with a 10kph reduction in speed), whereas there were substantial variations in detection rates between different clusters. This indicates that the layout of individual locations and speed signs has a much larger influence on the probability of detection than driver behaviour or other external factors such as road layout or weather conditions.

In addition, the factors identified as influencing the probability of detection, as well as their level of influence, were substantially different from the results in Phase 2 (see section 3.4.1.5 of the Phase 2 report), indicating that the 15 locations used in that analysis were not wholly representative of the Greenwich and Bracknell areas in terms of the factors affecting detection probability.

3.4.2.2 Factors influencing true/false detection probability

A similar modeling approach was taken to identify which factors influenced the likelihood of true or false detection. The Phase 2 analysis identified three fixed factors as being influential: vehicle speed, vehicle lateral acceleration and road type. In addition, cluster identity also had a significant influence on probability of a true detection. This was classified as a random effect.

Applying the same approach to the new, larger dataset in Phase 3 identified four influential fixed factors: lateral and longitudinal vehicle acceleration, the low beam headlight (on/off) and road type. Figure 21 shows the magnitude and direction of the fixed effects in the model. Average true detection rates across the dataset were high at 95%.

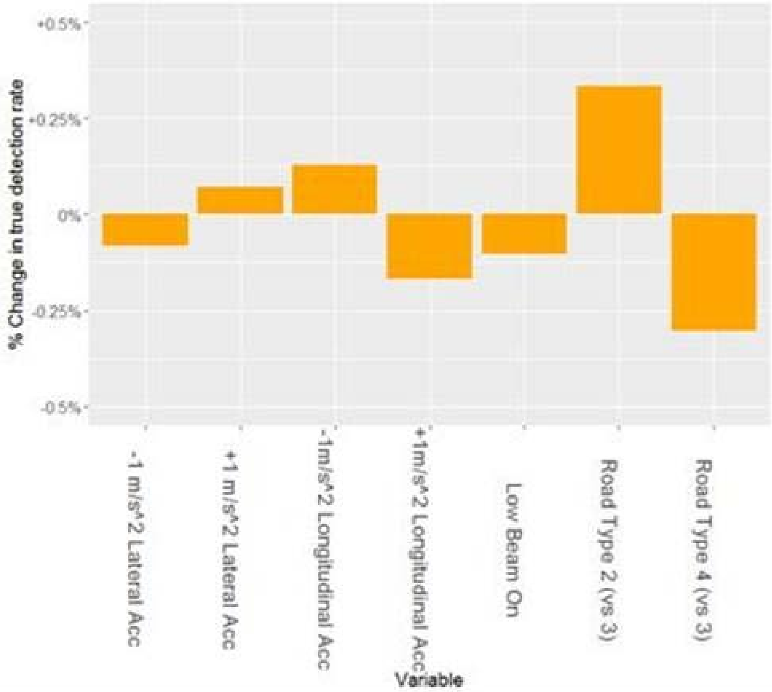

Figure 21 – Influence of fixed factors on probability of true detection

Road type was the most influential factor on the probability of a true detection. Compared to type 3 roads, the probability of true detection was 0.33% higher on type 2 roads and 0.3% lower on type 4 roads. This indicates that cars travelling on lower speed roads, where signage tends to be smaller, are less likely to detect speed signs correctly than those on higher speed, more major roads.

Positive lateral acceleration (turning or bearing left) was associated with a higher probability of true detection than a lateral acceleration at or below zero (travelling in a straight line or turning/bearing right). On the other hand, positive longitudinal (forward) acceleration was associated with a lower probability of true detection than a longitudinal acceleration at or below zero.

Having the low beam headlight on was associated with a lower probability than when it was turned off, suggesting that lower lighting levels reduced the probability of true detection.

Cluster identity also had a significant influence on the probability of true detection. Out of the 893 clusters in the data, 654 (73.2%) had a true detection rate of more than 99%, indicating that the TSR system was detecting the correct speed in these locations. However, there were 65 clusters (7.3%) with a true detection rate of between 50% and 80%, and for 7 clusters (0.7%) it was less than 50%, indicating that the TSR system was not performing well at detecting the correct speed in these locations.

As with the results presented in 3.4.2.1, cluster identity had a much greater influence on the probability of true detection than the fixed effects, indicating that the layout of individual locations and speed signs has a much larger influence on the correct speed being detected than driver behaviour or other external factors.

In addition, the factors identified as influencing the probability of true detection, as well as their level of influence, are substantially different from the results in Phase 2 (see section 3.4.1.6 of the Phase 2 report), indicating that the 15 locations used in that analysis were not wholly representative of the Greenwich and Bracknell areas in terms of the factors impacting true detection probability.

3.4.3 Sign correction exercise

The site visits in phase 2 uncovered issues with both speed sign placement and configuration, and issues with the methodology itself. More details on the information from the ground truth exercise can be found in section 3.4.1.7 of the phase 2 report.

In phase 3, rectification actions were performed at one of the seven locations with the help of the local council. The left-hand panel of Figure 22 shows the location that was identified with two speed signs (20mph and 30mph) in close proximity to each other, both facing traffic travelling west on the main road. Up until the end of January 2019, one of the two signs (either 20mph or 30mph) was detected for the vast majority of journeys (82%) by the MOVE_UK fleet. However, both speed signs were supposed to be for vehicles on the side road to the right and were not intended to be for vehicles staying on the main road. Therefore, on 5th February 2019, action was taken to rotate the signs so that they were only detected by vehicles turning into/out of the side road. The right-hand panel of Figure 22 shows the image of the speed limit signs after this action was taken.

Figure 22 – Location where speed signs were rectified, before (left) and after (right) rectification

Figure 23 shows the probability of detecting a traffic sign by the MOVE_UK TSR system at this location from December 2016 to April 2019. It can be seen that since action was taken, there have been no further detections of any sign, although there have only been 5 journeys made at that location. However, it should be noted that in the period December 2018 to January 2019, there were eleven missed detections at this location from eleven journeys, suggesting further data would be required to fully understand the effect of actions taken.

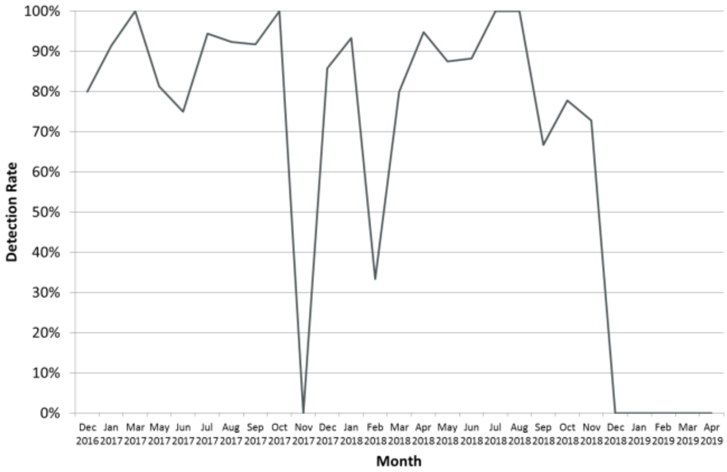

Figure 23 – Detection rate over time

3.4.4 Interactive dashboard

In Phase 3, an interactive dashboard was produced with the purpose of providing an interactive means of visualising the results of the TSR analysis. It was created using R-Shiny software, which enables interactive web applications to be built directly from the R programming language.

The dashboard was produced with the aim of giving the user the capability to identify specific locations with consistently poor detection / true detection rates, as well as focus in on particular regions or locations of interest. In addition, it includes the results of the statistical analysis presented in 3.4.2.1 and 3.4.2.2.

The main feature of the dashboard is an interactive OpenStreetView map showing the locations of the 893 clusters included in the analysis. This is shown in Figure 24. For each cluster, a point has been plotted at the mean latitude/longitude location of the detections in that cluster, and the icon at that point corresponds to the most common speed detected at that location. The user has the ability to zoom in and out, enabling them to look more closely at the clusters in certain regions.

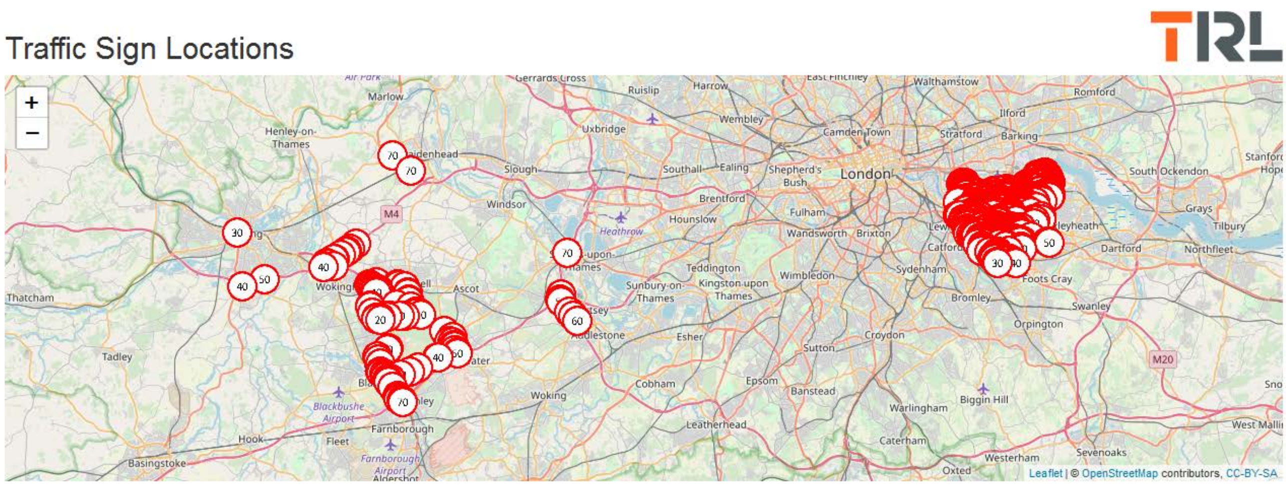

Figure 24 – Interactive map of locations in the Greenwich and Bracknell (TRL) areas where speed signs were detected by the MOVE_UK vehicles

The user can interact with the dashboard in a number of different ways, allowing them to focus on what they are particularly interested in. These include:

1. A dropdown menu to select whether to look at detection rates or true detection rates;
2. Sliding scales to restrict the map to only showing clusters where the detection / true detection rate was between the selected minimum and maximum;
3. Sliding scales to restrict the map to only showing clusters within a specified coordinate system, defined by the minimum and maximum latitude and longitude values that are selected;
4. A check box which, when ticked, adds to the map a plot of the detection / true detection rate over time at each cluster; and
5. A free text box to enable the user to type in the name of a particular location/cluster of interest. The map will then only show that one cluster.

For example, these filters could be used to focus on clusters with low detection rates, as these are the places where intervention may be necessary. Figure 25 shows how the user can identify clusters with a detection rate of less than 50%.

Figure 25 – Locations with an overall detection rate of less than 50%

Alternatively, it might be of interest to focus on a particular region or stretch of road. Figure 26 gives an example of where the filters on the dashboard have been used to do this. Adding in the plots of detection rate over time at each cluster shows that one cluster had a 12 month period where there were no detections at all. This was due to roadworks taking place, as identified in section 4.4.3 of the Phase 1 report. The 40mph sign was covered up as the speed limit had been changed to 20mph.

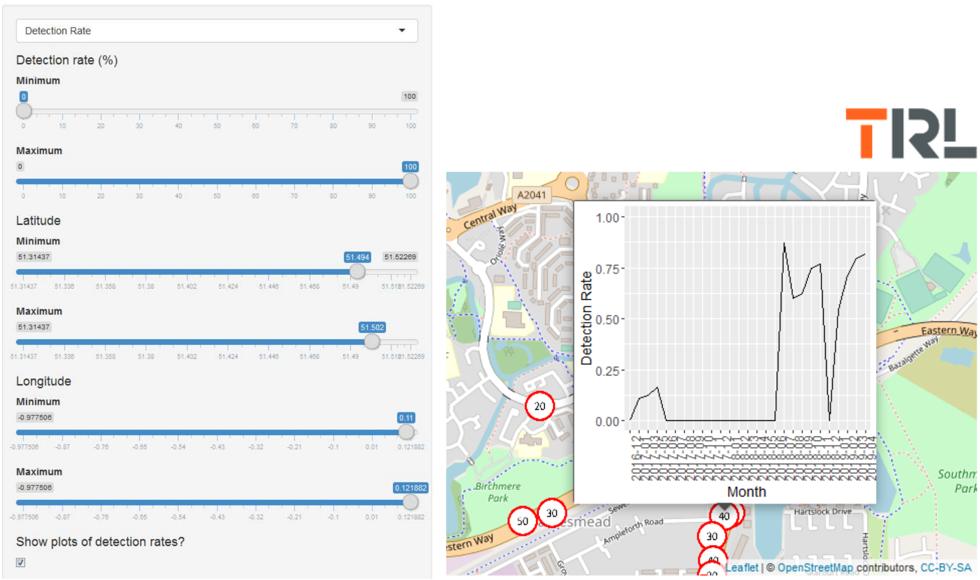

Figure 26 – Example of location where detection rate was 0% for a period of time

There are a number of ways in which the dashboard could be further developed. These include:

- Flagging locations/clusters when a step change occurs in the detection/true detection rate, indicating that something has changed at that location;
- Adding in a filter to only include data from certain months or periods of time; and
- As and when data becomes available, producing similar dashboards for other locations/regions, particularly where there is interest from local authorities to identify speed signs that are not being detected correctly.

3.4.5 Exploratory "big data" approach to processing of TSR data

The current MOVE_UK data collection setup involves sending of data packets from vehicles to a cloud storage server. The data collected until the end of the project is over 450 GB of sensor signal which was then analysed to perform work package related activities. The general schema of analysis interaction pipeline is shown in Figure 27.

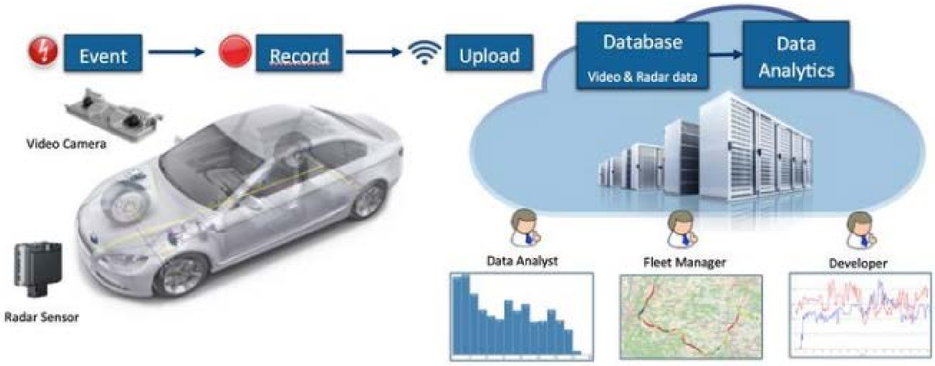

Figure 27 – Current analysis interaction pipeline in MOVE_UK

This activity gave a good experience in trying to handle relatively larger datasets. There was a realization that conventional techniques will not be sufficient to carry out analysis practically with datasets of such growing size. This raises the question of potential scalability of analysis techniques beyond this size of data.

Two approaches for performing the TSR analysis over this large dataset were trialed.

The first approach can be seen in Figure 28. This involved downloading the TSR relevant CAN signals through the SFDe application from the storage server and running the analysis scripts in a Python library to identify detection/missed detections and subsequent calculations. This approach is memory intensive because the amount of data downloaded is large.

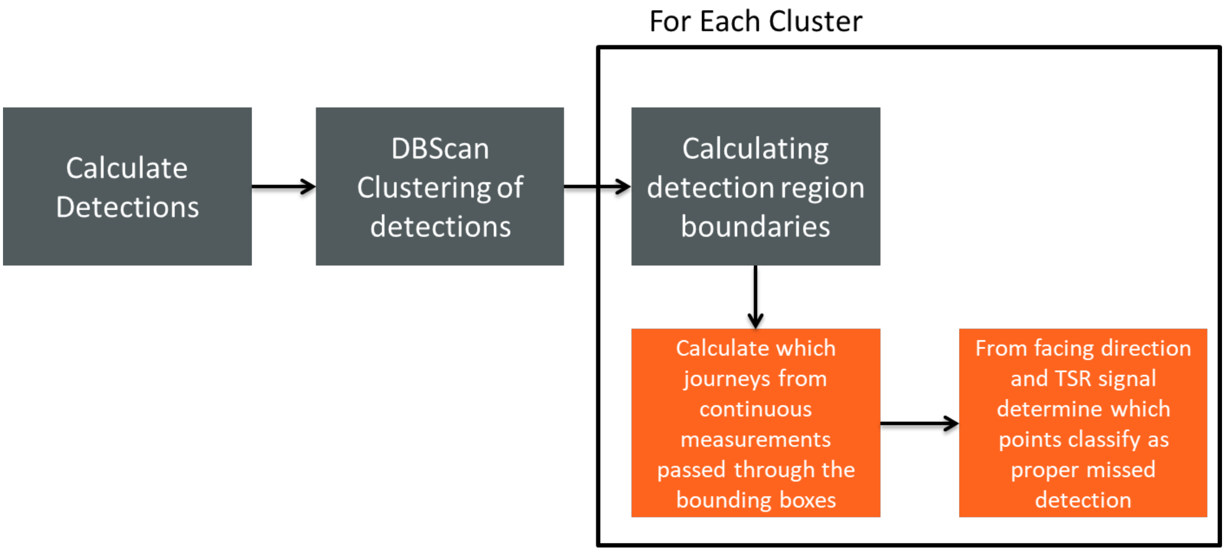

Figure 28 – Pipeline for performing TSR analysis on a larger dataset

An alternative approach was explored for analysing the TSR data and it performed the analysis using the 'Hadoop' cloud system which was already established for the data collection/storage. The techniques developed with this approach would deal with data scalability and allocation of necessary compute power for performing the analysis on larger datasets. Another advantage of this approach is that the data would not need to be moved over the internet in order to be analysed at different locations. Analysis could be performed using 'PySpark' or other similar arrangements in 'R' within the server and only the outputs needed to be downloaded.

The working flow diagram of the new method is shown in Figure 29 using the TSR use case as an example. This method reads the stored parquet data, processes it and stores the results in a results file. This operation is summarized with the acronym Extract (Read), Transform (Analyze), Load (Output) or ETL.

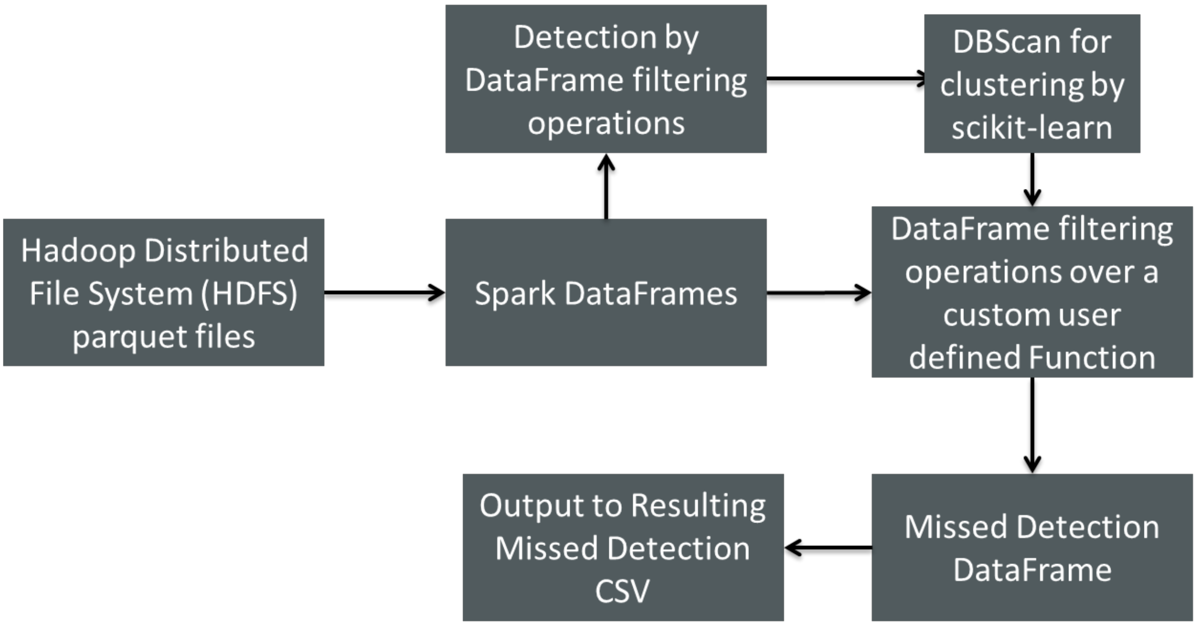

Figure 29 – ETL pipeline for performing TSR methodology on a larger dataset

This is a much more scalable approach because it works with distributed datasets and does not need to load all the data in memory. This would be proposed as the new way forward for such scalable analysis.

Table 6 presents a summary comparison of the two approaches for the particular application to the TSR clustering analysis.

Table 6 – Comparison of two methods for performing TSR clustering analysis

Local analysis with SFDe download	Hadoop ETL pipeline
Data downloaded locally	Analysis Jobs submitted and performed where the data is situated
Local machine a limitation	Can be scaled with the server processing power and scale
Good for prototyping and performing complex/ quick queries	Good for handling large scales of data

3.4.6 Conclusions

The key outcome of the TSR use case is an **'evidence-based methodology'** that can effectively determine the detection rates of speed signs and their change in detection overtime. The methodology can be used to identify locations where the recognition system (not just speed recognition) is performing poorly and notify the location(s) to concerned parties for restorative interventions which would improve detection going forward (feedback loop system).

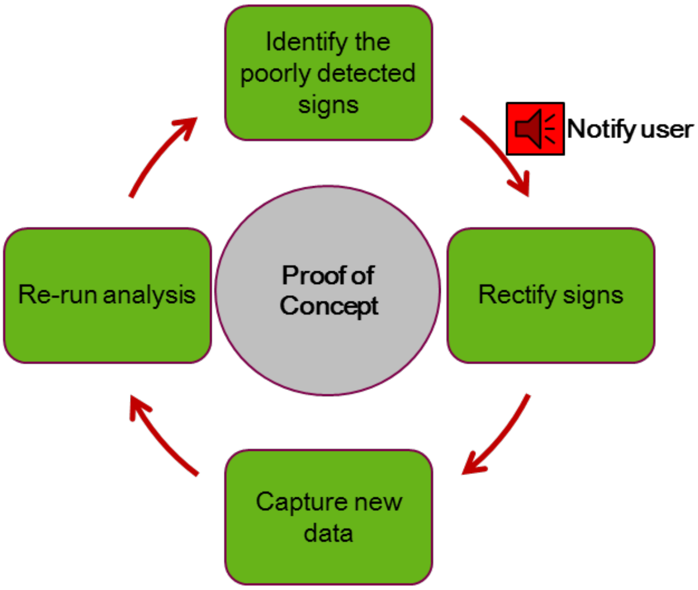

Figure 30 – 'Evidence-based methodology' for TSR systems

The results presented in 3.4.3 give a good example of how the desired 'evidence-based methodology' has been used to do the following:

- Identify a location where the TSR system was not performing well;
- Visit the location and take necessary action to improve system performance; and
- Examine data after the intervention to ensure that it has been successful.

Since the intervention took place, only 5 journeys have been recorded past the location, so in practice, more time/data would be needed to be able to definitively say that the TSR system is now performing well. However, the results here give an initial indication that it is performing as expected.

This methodology can be adapted to work with other systems such as pothole detection, road marking and infrastructure detection that support Automated Driving.

3.5 Use case Subcritical Radar Based Autonomous Emergency Braking (ARB)

As in *MOVE_UK Data Analysis Report - Phase 2 (D7.4)*, for the purposes of the MOVE_UK project, radar based AEB functionality is referred to as ARB in order to distinguish it from the video based AEB use case.

3.5.1 ARB events captured in Phase 3

Between 1st June 2018 and 31st May 2019, 98 ARB events (sequences) were collected. More precisely, the distribution of ARB function activations over the time period was as follows: **Jun '18:** 11, **Jul '18:** 16, **Aug '18:** 2, **Sep '18:** 10, **Oct '18:** 13, **Nov '18:** 8, **Dec '18:** 6, **Jan '19:** 4, **Feb '19:** 1, **Mar '19:** 11, **Apr '19:** 11, **May '19:** 5.

A further, final adjustment of the ARB parameter settings was made on 22nd of August 2018, prior to the start of Phase 3. The final setting was very close to the production radar based AEB settings which meant that mainly near-critical situations were captured. The continuing presence of some sequences that, from the analysis of the video footage, were perceived as possessing low criticality can be explained by the fact that the front radar sensor software is capable of several escalation steps, from warning, to low partial brake, to maximum activation. When re-simulated with production software, most of these sequences resulted in warnings or, at the maximum, only brief partial activations to support the driver.

Two sequences were chosen for a more detailed analysis. The images below show the point in time of ARB function activation.

Figure 31 shows an activation which occurred because the car in front was stopping due to another car reversing out of its parking space. With the sub-critical parameter setting, a partial brake request was sent out by the radar sensor. Re-simulation showed that the production software would have given out a warning. This situation is a good example of where lighting conditions can hamper visibility and thus affect image analysis for camera systems. However, the lighting conditions did not influence the outcome from the radar sensor since it is not influenced by the environmental lighting.

Figure 32 shows the moment of ARB function activation caused by the car in front stopping to let another car merge in. The radar sensor deems intervention necessary due to the short distance to the car in front (the leading object) in combination with the amount of remaining negative relative velocity of the leading object. This difference in speed was significant enough to trigger an ARB activation with production parameters, as shown by the re-simulation. It is interesting to note that the camera AEB system (with sub-critical parameter set) also reacted with an activation in this situation.

In both scenarios, the driver reacted by braking strongly – albeit not harshly (according to the consortium's definition) – in order to avoid collision.

Figure 31 – ARB activation 1

Figure 32 – ARB activation 2

3.5.2 Summary of results from Phases 2 and 3

Over the whole front radar/ARB trial duration, a total number of 151 ARB sequences (i.e. sequences where radar function activation on CAN was the trigger starting the recording) were collected. In order to gauge the relevance of these sequences for a production software release, all the sequences were re-simulated with a production version of the radar based AEB system (i.e. using production radar parameters).

The results of this re-simulation were as follows: of the total 151 sequences, 6 sequences were unusable due to data corruption during recording. Of the remaining 145, 24 (~16,5%) showed an AEB braking activation and were classified as true positives, and 85 (~58,6%) showed that the situation would have led to the activation of a warning. It is important to note that the 24 true positives are included in the 85, since the present radar system operates with escalation steps, from warning to partial activation until full activation, the level of which depends on the development of the situation.

In summary, for 58% of the re-simulated sequences, the production system would have shown a reaction.

After re-simulation, it was also possible to find a correlation between the parameter optimisation steps and the percentage of true positives in the triggered sequences. The three parameter updates were put into effect at three consecutive Quarterly Meetings of the consortium. The following numbers of sequences (Table 7) were counted in between updates and the end of May 2019, respectively.

Table 7 – Effect of parameter optimisation updates on number of sequences and outcomes

Time of update (mm/yy)	# months	Total # ARB	True positives	Warnings
Initial build-up, end 11/17	2,5	16	0 / 0%	5 / 31,3%
1. Update, 02/18	3	34	1 / 2,9%	9 / 26,5%
2. Update, 05/18	3	31	1 / 3,2%	16 / 51,6%
Final Update, 09/18	9	70	22 / 31,4%	55 / 78,6%

The low amount of sequences in the first period of collection can be explained by the duration of the Phase 2 system build-up and the holiday period at the end of 2017. Otherwise, the rate at which ARB sequences were collected has gone down with each parameter update, while the percentage of both true positives and warnings has gone up which shows the positive effect of the parameter optimisation.

A further analysis of all the ARB sequences collected was also carried out by re-simulating them using the production version of the video camera based AEB system (described in section 3.22). The result was that two of the ARB sequences collected were classified as true positives i.e. would have triggered the activation of the production version of the camera based AEB system. It is interesting to note that one of these sequences produced only a warning when re-simulated with production radar software, while the other one did not lead to a radar based AEB activation at all during re-simulation. Here, the difference in perception between radar sensor and video camera again becomes apparent.

3.5.3 Conclusions

Since the main difference between the principle of triggered collection of ARB data and triggered collection of camera-based AEB data is the sensor used, the conclusions about the usability of event triggered data recording in section 3.2.3 are also valid here and *vice versa*. More specifically, a multitude of situational problems in urban traffic and thus highly interesting “corner cases” could be found that fulfil criteria that are only slightly below the activation threshold of production ADS systems. Examples for these situations are cars in front that abruptly brake in order to turn a corner or to slow down for a speed bump. Such recordings can on the one hand show urban conditions that might prove challenging – or even a blocking item – for autonomous vehicles. On the other hand, the triggered data collection can form a library of real-world scenarios to help developers to configure their parameter sets for the optimal trade-off between early activation (which means higher performance in use case scenarios) and “annoying” regular system activations in urban traffic.

A future application of the connectivity-based validation approach could be operation in internal company car fleets. Here, a measurement setup – not present in standard production vehicles – in combination with regular Wi-Fi uploads at the office could be used to add to the above-mentioned library of real-world situations. Such a library could then be consulted upon, with a certain statistical significance, for the comparison with and analysis of possible incidents in the field where no high-density internal sensor data is available.

The additional benefit of trialling event triggered data recording with another sensor, more specifically a radar sensor, is that a 2-sensor-setup allows for comparison due to redundancy. Thus, it becomes possible to compare function activation and, in theory, use one to validate the other and potentially find false negatives that way. Moreover, a video camera and its output is very intuitive for humans since a big part of human perception is based on visual stimuli. A radar sensor, on the other hand, allows further information which is not accessible for humans which enabled the members of the consortium to explore a different perspective of perception.

Another potential use for the shown approach of real-world event triggered data collection is the improvement of the robustness of the system against activations caused by false detection of a -relevant object, as opposed to situational robustness. An example of this would be the radar system triggering due to the reflections of a manhole cover or a bridge. Throughout the period of the project where the ARB use case was implemented (i.e. Phases 2 and 3), no such activations were recorded which speaks for the robustness of even the sub-critical parameter setting of the radar sensor used. However, this point shows the significance of the differentiation between false positives due to false detection – caused by false perception of objects or the environment, i.e. a weak spot of the sensor itself – and situational false positives where there is no fault in object detection but the eventual criticality of the situation has been falsely classified as needing an emergency braking response. For the latter, activation is deemed necessary due to required assumptions about the future development of the situation even though it may turn out that the criticality of the situation eases by itself. An example of this is the situation in which a driver does not brake for stopped cars in front since he can see the traffic lights already turning green. The radar system does not have this situational awareness about the reason for the cars in front to be stopped or that they are due to start moving again. The evaluation of the validity of these situational decisions can be up to interpretation of criticality and the strength of response available to the system and driver. It also depends on the behaviour of the particular driver, as to how aggressive or predictive their behaviours are and how often such events are encountered.

In this context, it was observed that some drivers provoke significantly more activations than others. Similarly, today, with the possibility to select a driving profile for the distance warning (early – medium – late), the deduction could be to offer the learning of personalised driving profiles that influence the parametrisation of the driver assistance sensor systems to provide reassuring support to timid drivers while not bothering more “sporty” drivers.

3.6 Use case Cut-In (CIN) scenarios

The Cut-In (CIN) use case is defined in section 4.3 of the *MOVE_UK Data Analysis Report - Phase 2 (D7.4)*. Cut-Ins are interesting real-world scenarios due to the fact that they frequently show human behaviour which is in conflict with the expectations of the Highway Code, e.g. specifically the rule-of-thumb of maintaining a lead time of at least 2 seconds to the car in front.

Systematic analysis of the behaviour of both the driver of the vehicle cutting in and that of the ego driver can benefit the development of autonomous vehicles, which could potentially have trouble with Cut-In scenarios. For this, remarkable individual Cut-In events were investigated and statistics over all Cut-In events were analysed.

3.6.1 Removal of non-CIN (false positive) events

As documented in *MOVE_UK Data Analysis Report - Phase 2 (D7.4)*, it was observed that, due to the simplified nature of the Cut-In trigger implemented on the vehicles, a small proportion of the sequences collected were not real Cut-Ins (i.e. false positives). These could be subdivided into 2 classes which were labelled "Cut-Outs" and "Drift Spikes". In order to prevent these False Positives from distorting the (statistical) analysis of this use case, it was decided to eliminate them from the collected data. Below is an explanation of the two measures implemented to filter out "Cut-Outs" and "Drift Spikes" from the CIN data.

3.6.1.1 Cut-Out filter

Since the CIN trigger is formulated by detecting a change in the ID of the vehicle driving in front, the scenario of the ego vehicle itself "cutting out" from behind one vehicle to merge behind another vehicle can have the same effect and lead to the start of a CIN sequence recording. Thus, the Cut-Out filter checks whether or not the ego vehicle is changing lanes at the ID change point by observing whether the previous OBJ_01 (i.e. vehicle in front/lead vehicle) becomes either OBJ_02 or OBJ_03 (the vehicle on the left or right, respectively) simultaneously with the ID change of OBJ_01. If that is the case, the event is classified as a Cut-Out (COUT).

The Cut-Out filter was developed by Bosch, implemented by The Floow and applied to all CIN data in order obtain a clean data set (from dates 19 July 2018 to 31 May 2019).

3.6.1.2 Drift Spike filter

Another type of False Positive that can occur is the scenario dubbed "Drift Spike" where the ID change of OBJ_01 indicates that the label "lead vehicle" switches briefly to a nearby car – e.g. due to a curve in the road – and then returns to the original vehicle. Thus, the Drift Spike filter checks whether the ID of the new OBJ_01 matches the pre-previous ID of up to 5 seconds before the ID change point. If this is the case, the event is classified as a Drift Spike.

Again, the Drift Spike filter was developed by Bosch, implemented by The Floow and applied to all CIN data.

3.6.2 Summary and analysis of CIN events captured in Phases 2 and 3

Following the implementation of both filters, a delete list consisting of all false positives was compiled and subsequently, the false positives were deleted from the sFDE database. As of 30 April 2019, 397 CIN events were recorded, 125 of which were false positives which brings the number of real CIN events to 272.

The majority of CIN events exhibited a fairly uniform nature which includes driving characteristics considered to be normal. However, there were some CIN sequences that stood out in terms of unexpected driver behaviour. An example of such a sequence is described below.

The Cut-In trigger is activated by a motorcyclist seemingly appearing from out of the line of sight of the camera and suddenly cutting-in in front of the ego-vehicle, as illustrated in Figure 33. Please note that the image is provided by the webcam integrated in the surround radar vehicles in Phase 3.

Figure 33 – Example of a CIN event

The two Birds-Eye views in Figure 34 below show the target objects selected by the front-facing radar sensor, right before the Cut-In point (A) as well as right after (B). The car driving in front (ID 18) switches from the classification of "lead vehicle" (OBJ_01) to "target object in front of lead vehicle" (OBJ_04) – denoted in the figure by a colour change – in the same instance as the motorcycle (ID 1) becomes the new lead object, without previously being selected as a neighbouring target object. The described sequence of target object selection illustrates that the motorcyclist also seemingly appears out of nowhere, from the perspective of the front-facing radar sensor. For a detailed description of the convention of target object classification and corresponding IDs, please refer to section 4.3.2 in *MOVE_UK Data Analysis Report - Phase 2 (D7.4)*.

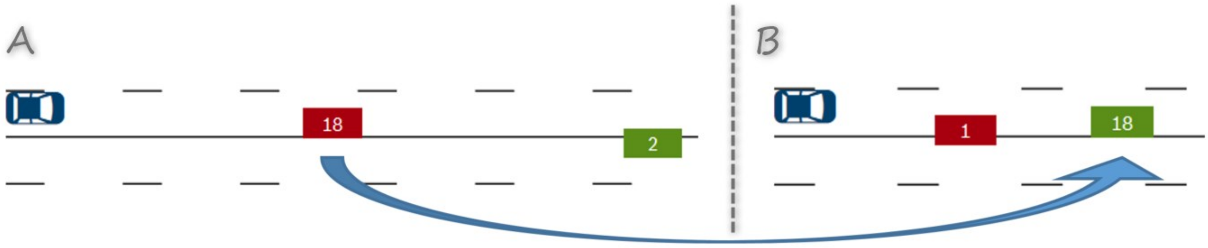

Figure 34 – Birds-eye view of the CIN example

The chosen example shows the kind of unexpected behaviour that future automated driving systems will have to take account of. It is interesting to note that the human driver hardly reacts to such a close Cut-In, presumably because the ego driver has already seen the motorcycle with their peripheral vision before it appears in front. The question of the consequence of an automated system reacting abruptly to such a Cut-In leads to the conclusion that analysis and understanding of these kinds of situations, which cannot be observed on a test track, is necessary for the development of comfortable future ADS system that can be embraced by the public.

3.6.3 Statistical analysis (using project-developed Web-UI)

As described in *MOVE_UK Data Analysis Report - Phase 2 (D7.4)*, a web-based User Interface (web-UI) for certain statistical analysis over all cleaned CIN data was developed in Phase 2 of the project. This web-UI was improved in Phase 3 with replacement of the scatter plots by heat maps for performance reasons, the introduction of fixed bin sizes, and bug fixes.

The analysis was divided into statistics for the vehicle cutting-in in front of the ego vehicle and for the actual behaviour of the ego driver. The user was able to select the desired time-frame, the vehicle they want to consider, as well as the road type (see chapter 4.4.2 of D7.4 for details). The data examples shown below are of evaluations resulting from the selection of all vehicles and both road types, for a date range from the 19th of July 2018 until the 28th of February 2019. Changing the road type to only road type 1 did not change the results in a significant manner.

The intention behind the statistics was to find the most prevalent human driving behaviour during Cut-Ins in order to be able to deduce parameters for the development of human-like automated driving systems. In-depth analysis done in this context which went beyond the analysis in the web-UI is described in chapter 3.6.4.

3.6.3.1 CIN vehicle statistics

The CIN vehicle statistics were computed from 100 measurement points taken at 10 Hz during 10 seconds after the CIN trigger event. The statistics shown in this report are taken from 142 events which means a pool of 14200 measurement points. The first visualisation takes all measurement points and plots the value pairs of "relative velocity" vs. "distance" in a heat map, shown in Figure 24.

The left-hand side of the heat map shows the data points for which the CIN vehicle was driving slower than the ego vehicle, while the right-hand side shows the data points for which the CIN vehicle was driving faster than the ego vehicle, which thus can be considered to show non-critical manoeuvres. The cluster of red colour in the middle and slightly to the left shows that in the time immediately after a CIN has occurred – i.e. our evaluation period – the ego drivers tended to either match speed with or be slightly **faster** than the new lead vehicle. The slight peak of 427 data points at -5m/s vs 15m shows that most drivers cutting in at slightly lower speeds will do so at around 15 metres.

The blank space at high relative speed and low distance – the bottom right-hand corner area of the map – can be explained by the fact that positive relative velocity means that the distance in the next time step will increase and will bring the CIN vehicle to a further distance. On the other hand, the blank space at bottom left shows what can be considered a critical area, since negative relative velocity means decreasing distance, i.e. for safety considerations this area has to be blank so that collisions can be avoided. It is interesting to note that slope on the left-hand side is flatter than the one on the right.

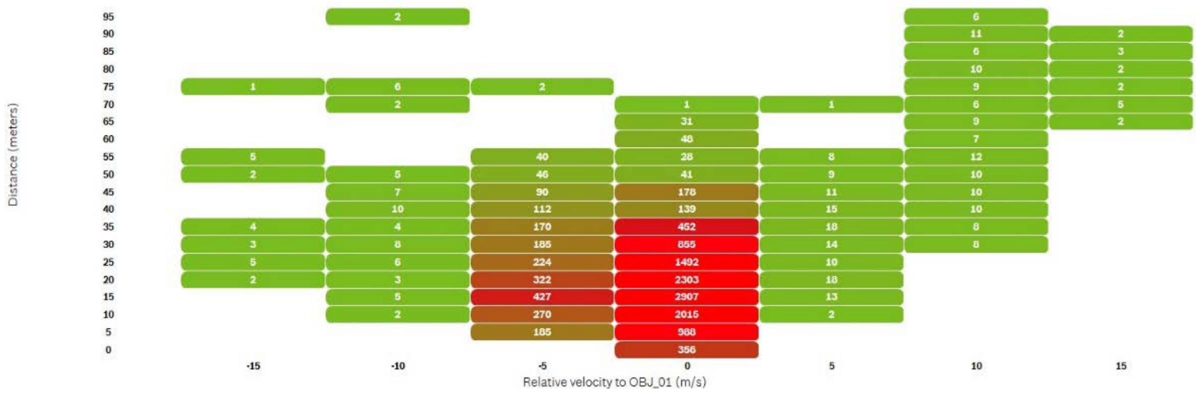

Figure 35 – Heat map - relative velocity of OBJ_01 vs distance

In order to more easily read out the criticality of the situation of the individual data point, another heat map (Figure 36) was provided in which the distance in metres was combined with the ego vehicle speed to calculate the time gap between ego vehicle and CIN vehicle. Here, the red cell in the centre bottom of the map, where the speed of the ego vehicle and CIN vehicle match, shows that the majority of drivers tend to maintain their original small time gaps, i.e. 0-5s, during CIN manoeuvres. Another cluster can be observed at -5 m/s showing that, in a significant number of cases, the time gap to the new vehicle is reduced. Possible reasons for this behaviour include slow or relaxed reaction of the ego driver or possibly an intentional closing of the gap in order to avoid another CIN.

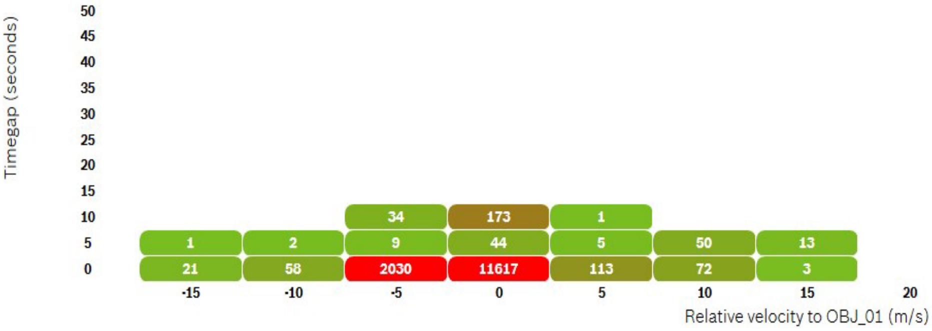

Figure 36 – Heat map- relative velocity of OBJ_01 vs time gap

Apart from the described analysis, it must be remarked that choosing appropriate data bins are vital for interpreting the results. In the present case, the area above 15s remains empty, which relates to the reach of the radar sensor. The most interesting area is between zero and four seconds of time gap which is predominantly chosen as the operating range of driver assistance systems. Therefore, the recommendation for future analysis is to decrease the bin size and limit the view to zero to 15 seconds.

Lastly, both the values for minimum time gap (or distance, if preferred) to the new lead vehicle and its maximum relative negative speed were counted for all individual CIN manoeuvres, see Figure 37. The results show that in most cases the minimum time gap is one second which again confirms the above stated conflict of human driving behaviour with the recommended two second safety gap. Furthermore, the majority of drivers cut-in with a maximum speed difference of -2.5m/s which agrees with the previously stated observation that the ego driver usually doesn't immediately match the speed of the new lead vehicle. Again, the reason could be trying to avoid another cut-in or preferring comfortable braking over abrupt reactions.

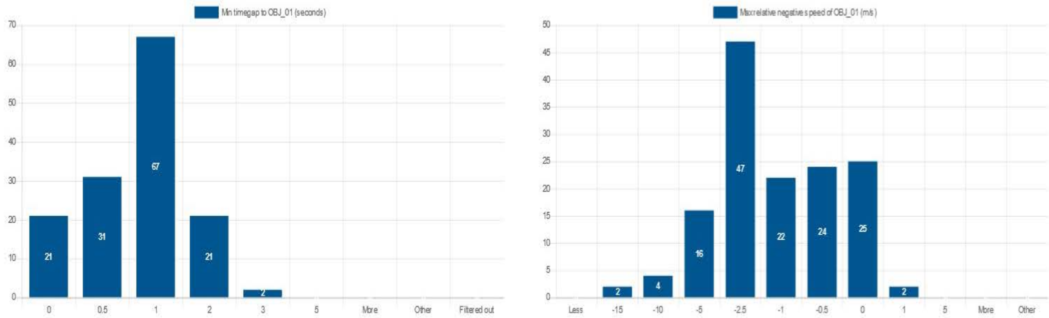

Figure 37 – CIN vehicle statistics

3.6.3.2 *Ego driver statistics during CIN*

The statistics pertaining to the behaviour of the ego driver during the CIN manoeuvre are calculated over a different evaluation period to that used for the CIN vehicle statistics. This is because the ego driver can already start reacting to the situation before the new vehicle has moved into the ego lane enough so that the radar sensor classifies it as the new lead vehicle. Thus, the ego driver statistics are accumulated maximum values, which are computed from the measurement points recorded in the period starting seven seconds before and finishing seven seconds after CIN trigger. The results are shown in Figure 38.

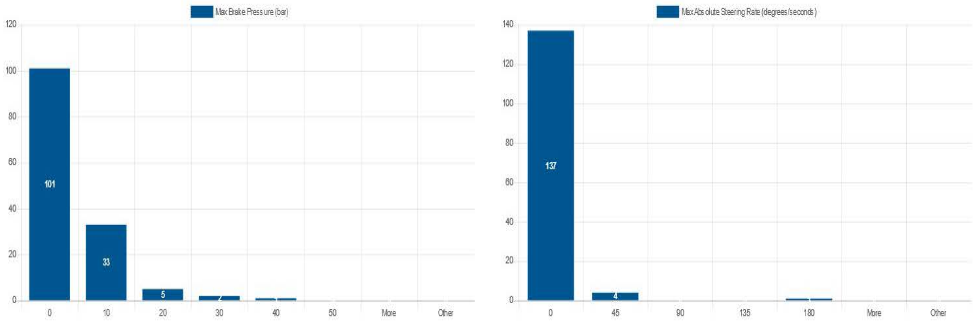

Figure 38 – Ego driver statistics during CIN

In most cases, the maximum brake pressure of the ego driver during a CIN manoeuvre remained between zero and ten bars, and it rarely exceeded 20 bars. Since the chosen threshold for the Driver Harsh Braking Use Case (see section 3.3) lies at 40 bars, it can be stated that the count of events where the driver is braking harshly before, during or after a CIN event is negligible. This observation of the drivers' tendency to brake smoothly in reaction to CIN fits with the above observation that the distance to the new lead vehicle tends to decrease. For optimal visualisation results, it is recommended for future analysis to limit the bins to the range of 0-40 bars and decrease the bin sizes to 5 bars per step.

In almost all cases, the maximum absolute steering rate was between 0 and 45 degrees/s, with only a few events with a maximum absolute steering rate of above 5 degrees/s. In the context of autonomous driving, steering angle rates above an absolute value of 5 degrees/s are considered harsh steering. Therefore, in order to gain a more differentiating distribution, it is recommended again to reduce the total bin range to 0 to 45 and split it up into a larger number of smaller bins (e.g. 10 bins).

3.6.4 Statistical analysis (beyond project-developed Web-UI)

As previously stated, anecdotal evidence suggests that disregarding the two second recommendation by the Highway Code is a common occurrence. It is also known that CINs can take place whenever the physical space between vehicles is large enough to accommodate an additional vehicle. Such CIN events, at close quarters, pose a tangible and credible risk to the safe passage of autonomous vehicles.

For example, at 80 km/h a two second gap equates to a physical distance of some 44m to the lead vehicle, or approximately 10 car lengths. If a human driver decides to perform a CIN at these speeds, then how should the autonomous vehicle respond? On the one hand, the vehicle could choose to maintain its current velocity, relying on other road users to mitigate risk and increasing the potential severity of any accident. Alternatively, the autonomous vehicle could decide to slow down, perhaps to strictly adhere to the recommended lead time of 2s.

Slowing down poses two immediate concerns, the first of which is the retardation of the vehicle increasing the probability of rear impact from following vehicles. Secondly, is slowing down only going to create a new space for further CIN events from adjacent vehicles? If so, is the motion of the autonomous vehicle going to be continuously impeded as a result of overly courteous behaviour leading to perpetual CIN events? What would the real impact on risk be to introduce to the road network a population of users who behave fundamentally differently to all other drivers?

A basic understanding of human behaviour during CINs has been derived by looking at the exact instant of the MOVE_UK CIN events and recording the relevant lead times. When available, the lead time to any lead vehicle that is subsequently replaced by the CIN vehicle is also recorded. Figure 39 shows the lead time of CIN vehicle, and the vehicle in front, as a function of Ego vehicle speed.

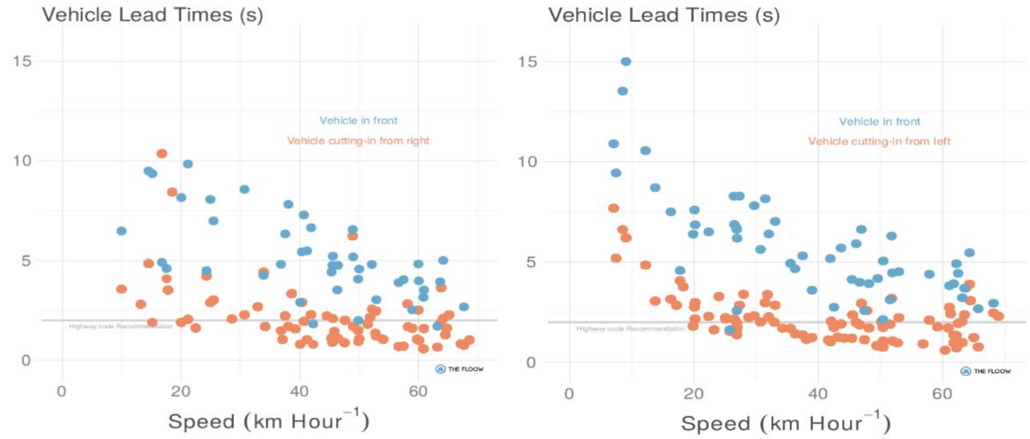

Figure 39 – Lead times of CIN vehicle

From these results, we can make the following conclusions:

1. The stipulation for a two second lead time is nearly always being obeyed when a CIN takes place and
2. CIN vehicles occupy a space with increasingly small lead times at higher speed. After 30 km/h this lead time is frequently < 1s.

In SUR-CIN (section 4.2) we expand on this result to show how drivers respond to the CIN vehicle; is there a following vehicle that means the Ego vehicle cannot safely slow down? How does the behaviour change when such a vehicle is present? This information is also important for future telematics projects (see section 4.4), as it gives a clear picture of how people perceive and react to emerging risks.

Finally, we examined the trajectories of the CIN vehicles in the few seconds immediately prior to becoming the lead vehicle (see Figure 40). We found that while the majority of CINs occur at constant velocity relative to the Ego vehicle, there is a higher probability of a vehicle accelerating to the middle lane from the left, and a higher probability of a vehicle retarding from the right. Such behaviour is to be expected for roads in the UK, where lanes to the right are intended for faster motion and overtaking manoeuvres.

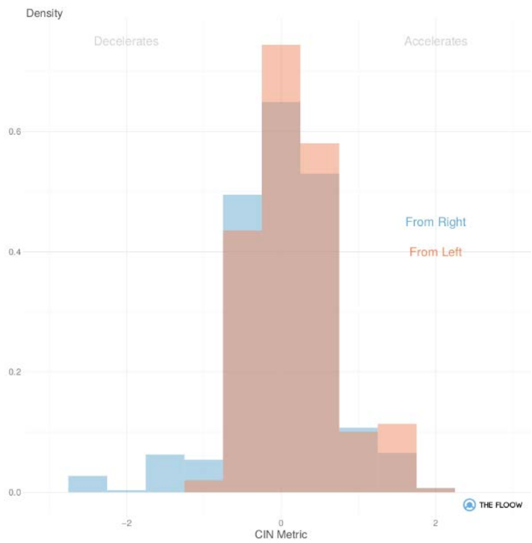

Figure 40 – Trajectories of CIN vehicles

3.7 Use case Lead Vehicle Statistics (LVS)

As stated in *MOVE_UK Data Analysis Report - Phase 2 (D7.4)*, the intention behind the LVS use case is to collect and display real-world statistical data from driving scenarios on major arterial roads (type 2) and motorways (type 1), to better understand the likely availability of level 3 ADS features and the behaviour of the driver and surrounding drivers in these situations (availability means the proportion of time that a typical motorway traffic jam related level 3 feature, which relies on the presence of a lead vehicle, would be available for the driver to use). With this increased understanding, the evidence can then help define parameters for safe and human-like level 3 and 4 ADS features.

For details on the design of the evaluation of the LVS Use Case, please refer to D7.4. The reader is reminded that unlike the CIN Use Case, where a new stream of dense CAN data was triggered and collected for a limited time at higher frequency, the LVS Use Case is intended as a proof of concept for deriving the desired information from the continuous 1 Hz CAN data.

The metrics analysed in the LVS Use Case are:

- The distribution of the distance to the lead vehicle vs the ego vehicle speed;
- The length of time for which all LVS conditions are continuously fulfilled, a.k.a. the Snippet Length; and
- The LVS condition that is violated first and thus breaks the continuous snippet length, a.k.a. the Exit Reason.

In the following sections, all the data collected since the implementation of the LVS Use Case at the start of Phase 2 until the end of May 2019 is selected for analysis, unless otherwise stated.

3.7.1 LVS analysis: distance vs speed distribution

The left-hand side of Figure 41 shows the resulting heat map for all data pairs of ego vehicle speed vs distance to lead vehicle for the entire evaluation period (1 Dec 2017 to 31 May 2019), where both available road types (types 1 & 2) are selected. In comparison, the graph on the right-hand side shows the data pairs where only road type 1, i.e. motorways, is selected.

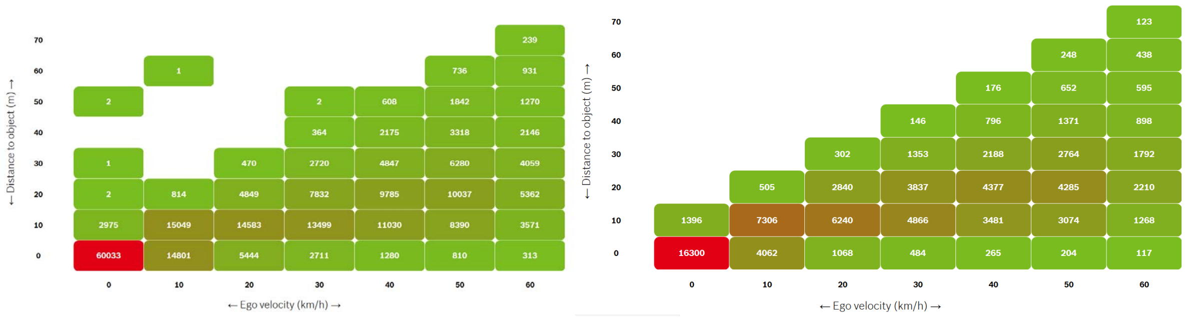

Figure 41 – Distance vs speed distributions 0-65 km/h

It is clear that for both road types some outliers can be observed in the lower speed ranges. Moreover, over a quarter of all data points are encompassed in the "0-10 km/h" vehicle speed bin. This can be explained by the fact that road type 2 refers to major arterial roads that can include traffic lights. In order to not skew the analysis results with non-motorway-type driving, only road type 1 is selected for the remainder of the LVS analysis. This is particularly relevant as consideration of human behaviour on motorways is currently of much interest for the development of level 3 ADS systems for the application of assisted or autonomous motorway driving.

Considering the right-hand side of Figure 41, it is possible to draw an approximate line through the maxima of sample numbers for each velocity bin, meaning that as speed increased, the distance to the lead vehicle increases. However, the slope of said line appears flatter than should be suggested by the advised safety distance. Also, although there is a clear maximum for each velocity bin, the distribution inside the velocity bins become less concentrated the higher the velocity. This agrees with the observation during the stop-and-go of traffic jams on the motorway that in low speeds (during the "stop" phase) the gaps tend to get closed and the vast majority of people drive bumper to bumper while during the "go" phase the cars pull apart more.

It is interesting to note that the clustering at the upper end of the selected velocity range described in chapter 4.4.2 of report D7.4 is not apparent anymore. A possible explanation for this observation is that with the relatively low amount of data available at the time of the Phase 2 analysis, outliers can tend to skew the data heavily in comparison with the amount of data that was at our disposal for the present report.

Figure 42 below shows the result for selecting road type 1 and the speed range of "65-120 km/h", which may be considered the typical cruising speeds on motorways while following other vehicles. Referring to the left-hand side of Figure 49, which shows the results for the entire evaluation period (1 Dec 2017 to 31 May 2019), it can be seen that the clear line of the lower speed ranges becomes even more dispersed. Also, a large number of data points are distributed across a broader range of distances in the "100-110 km/h" speed bin which corresponds to the legal speed limit on motorways in the UK. It is interesting to note that restricting the evaluation period from the beginning of 2019 until end of May 2019 (as shown in the right-hand side of Figure 49), results in the emergence of a spike at 70 km/h. This could be explained by a higher amount of travelling through temporary road works in operation around/during this period.

Figure 42 – Distance vs speed distribution 65-120 km/h

Lastly, the possibility of differing driving behaviour during winter and summer months was considered. For this, January and February were selected, as well as July and August. However, no significant difference in distances could be observed in either the higher or lower speed ranges. In order to put this observation into context, two points have to be kept in mind. Firstly, the meteorological winter does not necessarily mean different driving conditions – e.g. the UK winter in the beginning of 2019 was especially mild. Secondly, it has to be taken into account that restricting the data to only 2 months also means less data and thus less statistical significance. For the future, either a larger fleet has to be employed for considerations of this kind, or the fleet has to be specifically deployed for motorway driving, unlike the MOVE_UK fleet.

As a conclusion it can be said that the evaluation of ego vehicle speed versus distance to lead vehicle is useable for generic inferences, e.g. the typical braking distance at 60km/h being around 20m. With that, the proof of concept to further exploit existing low-frequency data for gaining statistical knowledge of driver behaviour was successful. However, in order to answer more specific questions changes in the deployment might be necessary, depending on how much or little these questions align with the typical driving behaviour of the existing fleet configuration.

3.7.2 LVS analysis: snippet length and exit reason

The distribution of the snippet length, i.e. the length of time for which all LVS conditions are continuously fulfilled, is found to steadily decline from its clear minimum at one second which corresponds to the sample time. For the whole available data range (1 Dec 2017 – 31 May 2019) and for road type 1, the longest snippet for the speed range “0-65km/h” is 1047 seconds (i.e. approx. 17 min) and for the speed range “65-120km/h” is 1812 seconds (i.e. approx. 30 minutes).

The reason for the unexpected amount of very short snippets can be surmised when also taking into account the evaluation of the exit reason, i.e. the distribution of LVS conditions that are violated first and thus break the continuous length of one individual snippet. The main exit reason differs depending on which speed bracket is chosen:

- For the “0-65 km/h” range, most snippets end because the speed of the ego driver exceeds the defined maximum, i.e. 65 km/h;
- The majority of the snippets in the “65-120 km/h” are terminated by the condition called “object loss”.

It is also interesting to note that there are considerably more snippets in the higher speed bracket than in the lower which can be explained by the fact that the MOVE_UK fleet is not typically driven in slow moving traffic on the motorway.

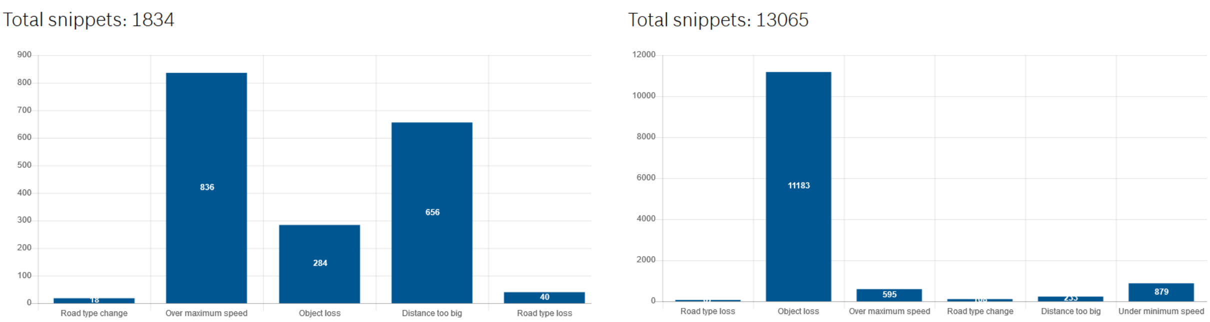

Figure 43 – Distribution of exit reasons

As already stated in chapter 4.4.4 of the D7.4 report, it becomes apparent that the definition of hard limits for all pre-conditions combined with the nature of the data is mostly responsible for the unexpected distributions of both snippet length and exit reason. Taking the example of the pre-condition based on vehicle speed, each slight crossing causes a new snippet. Considering that a prevalent speed limit at constriction zones in the UK is 40 mph which corresponds to 64.4 km/h, the maximum number of snippet endings due to “Over maximum speed” for the “0-65 km/h” range (see left-hand side of Figure 43) is easily explained. This also fits with the low amount of “Over maximum speed” in the “65-120 km/h” range (see right-hand side of Figure 43), because the maximum speed limit in the UK is at 70 mph or 112.7 km/h and drivers thus rarely cross 120 km/h. The range of “0-65 km/h” was chosen purposefully because it corresponds to the operational range of automated driving functions for slower moving motorway traffic and should thus not be increased or decreased. Therefore, a more computationally expensive method such as hysteresis is advisable for future analysis.

Similar considerations apply for the high amount of the exit reason called “Object loss”. This condition was chosen to gain an impression of how often, and in which time intervals, drivers choose to overtake other vehicles instead of following them. However, it has to be taken into account that the radar object information taken from the CAN corresponds to raw classification at lower computation levels where no plausibility check or tracking is employed as it is at higher-level functions. Thus, in order to smooth out momentary deselection of the target object in the raw data due to the probability-based calculation, especially at the larger distances that come with higher velocities, higher-level pro-processing such as tracking or again hysteresis is recommended before conducting statistical evaluation in the future. As a consequence, in terms of showing how long a driver follows a lead vehicle on a main artery road or motorway, no definite statement can be made at this juncture.

3.8 Use case Telematics – Phases 1 and 2

During the first two phases of the project, extensive analysis of signals from the MOVE_UK vehicles was carried out in order to understand the potential value of data to improve the understanding of risk and incidents.

During Phase 1, a framework for evaluation of data value, to both areas of interest, was developed. This framework was utilised again during Phase 2, when new data fields became available for analysis.

3.8.1 Update on work carried out in Phases 1 and 2

In order to assess the value of data obtained during earlier phases of the project, analysis was undertaken across all available data fields, aiming to understand each in the context of incidents (EDR, see project deliverable D10.3) and risk (see project deliverable D10.5).

In order to assess any potential new data fields, a process was developed to judge a field’s value for potential use as a discriminator. To understand risk, for example, data that is more volatile and occupies a broad range of values (such as those relating to the operation of the pedals, or directly to the physics of the vehicle’s motion) will be fundamentally more valuable than those data fields that are static or rarely triggered (e.g. those associated with ACC or AEB systems). On the other hand, for understanding incidents, rarely triggered parameters can have a much higher value, as those fields tend to change only in the event of a sudden emergency.

Positional information regarding other road users has a high value for understanding both incidents and risk. Detailed information about the behaviour and positioning of surrounding vehicles from camera and Radar systems has a revolutionary potential for incident understanding and risk estimation compared to observation of telemetry alone. The focus of study in Phase 2 was directed towards analysis of additional additive fields from the likes of the forward-mounted Radar system and enhanced camera data.

3.8.2 Extension of work in Phase 3

Unlike at the commencement of Phase 2, Phase 3 saw no major changes to the composition of the continuous data. Additional data was collected from the surround radar during a CIN (Section 4.2) or when one of the two vehicles passed through a geo-fenced location (Section 4.3). The intermittent capture of this data meant that we cannot assess the value of new parameters for EDR and risk assessment in a consistent way to those examined under the methodology summarised in *MOVE_UK Data Analysis Report - Phase 1 (D7.3)* and *MOVE_UK Data Analysis Report - Phase 2 (D7.4)*. However, the high value of the forward radar data to both EDR and risk assessment was demonstrated during Phase 2, and the analysis from the Phase 3 use cases (Section 4) strongly suggests that the surround data can add significant value to our understanding of both.

Section 4 describes more of the Telematics activities in Phase 3.

4 Phase 3 use cases and capabilities

4.1 Overview

In Phase 3, the main aim of the consortium was to use the additional corner radar (surround [SUR]) data which was collected to better understand human driver behaviour; something that is essential to the successful development and public acceptance of future ADS. For example, how aggressively do drivers merge into a busy motorway lane (in front of another vehicle) or enter a busy roundabout or intersection?

To answer these questions and others, the following three additional use cases were developed and implemented in Phase 3 of the project:

- Surround sensing of Cut-In situations (SUR-CIN);
- Surround sensing of Cross-traffic situations (SUR-CROSS); and
- Telematics 3.

The updated objectives diagram (Figure 1), presented in the Introduction, provides an overview of how each of these use cases relates to the applications and capabilities defined at the beginning of the project and described in more detail in the Phase 1 and Phase 2 reports (*MOVE_UK Data Analysis Report - Phase 1 (D7.3)* and *MOVE_UK Data Analysis Report - Phase 2 (D7.4)*. Note that the above-mentioned aim of Phase 3 is closely linked to the capability of "Capturing events (driver interventions & behaviour)" which can be seen in Figure 1 adjacent to "SUR" which is the collective term used for the SUR-CIN and SUR-CROSS use cases.

The focus of the SUR use cases was on analysing the surround data to better understand the behaviour of the ego driver and surrounding drivers during cut-in and cross traffic situations (e.g. at roundabouts and junctions). SUR-CIN used the same trigger condition as the CIN use case implemented in Phase 2, which helped reduce the complexity of implementing SUR-CIN (the first Phase 3 use case). In contrast, SUR-CROSS used a new location-based trigger.

Work towards the final telematics use case focused on using the continuous data to explore the potential for improving existing telematics solutions and developing metrics suitable for grading driver behaviour in the era of AVs.

Collection of corner radar data in the final format started at the beginning of November 2018 and finished towards the end of June 2019.

Details of the additional hardware required to support the Phase 3 use cases can be found in Section 2 of this report.

The Phase 3 use cases, and the results from each, will now be described in more detail in the following sections.

4.2 Use case Surround sensing of Cut-In situations (SUR-CIN)

4.2.1 Purpose

As with the Phase 2 CIN use case, the main purpose of the SUR-CIN use case was to understand the behaviour of the ego-driver and other drivers during cut-in situations. However, unlike Phase 2 CIN, where the only object data available to interpret the situation was from in front of the ego vehicle, with SUR-CIN, the consortium had the additional benefit of object data from all around the ego vehicle, as well as video footage from the forward-facing IP camera. This not only allowed a better understanding of the cut-in manoeuvre taking place in front (see Figure 44) but, for the first time, allowed an understanding of the behaviour of the vehicle(s) behind the ego vehicle at the point of cut-in (See Figure 45).

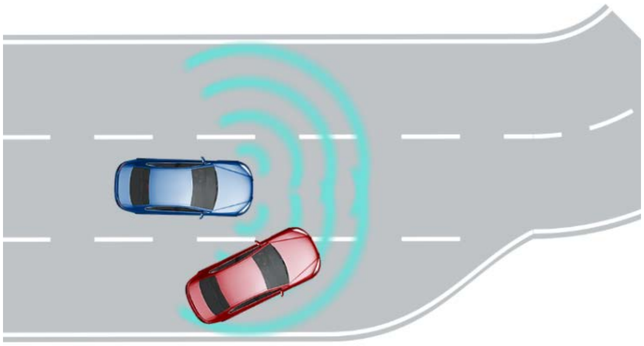

Figure 44 – Front corner radars provide improved view of vehicle cutting in, in front of ego vehicle

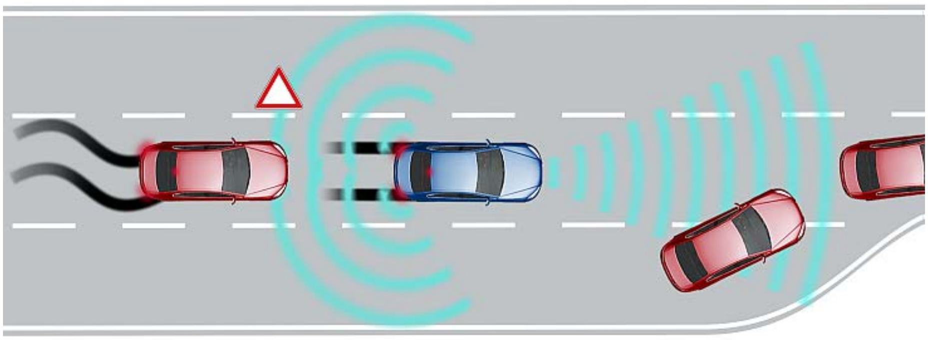

Figure 45 – Rear corner radars pick up vehicle braking harshly behind ego vehicle during a cut-in

4.2.2 Design

In Phase 2, the detection of cut-in scenarios was based on data from the forward-facing radar sensors fitted to the vehicles at the beginning of that phase. The trigger conditions were calculated by the original Flea3 device installed in the vehicles which was connected to the forward-facing radar sensor. A signal was transmitted on the CAN network when the trigger conditions were met. For a more detailed description, please refer to section 4.3 of *MOVE_UK Data Analysis Report - Phase 2 (D7.4).*

In Phase 3, this same trigger signal was also captured by the second Flea3 device which was connected to the corner radars and IP camera. So, in theory, the two Flea 3 devices were expected to capture the same trigger signal at the same time, and collect data from different sources, at the same time, around the trigger point.

Given that the trigger conditions were the same, it was expected that the number of CIN measurements recorded by the two Flea 3 devices would be the same. However, because the Flea3 devices used in MOVE_UK could only record one event at a time, the event which matched the trigger conditions first, within the designated recording timeframe, was the only one recorded and the next event falling within the recording timeframe was disregarded. This was significant because, during nearly all of Phase 3, triggered recording of two separate events (SUR-CIN and SUR-CROSS) was configured on the second Flea3 device. So, whenever a SUR-CIN and SUR-CROSS event was triggered at around the same time, and the follow-up processing, compressing and sending actions executed, there was a high probability that the second Flea3 would fail to record one of the events. The situation was not helped by the length of time it took to complete these follow up actions, which could be 2-3 minutes after the event had finished. The reason for this was largely down to system load and the fact that the Flea3 device was not developed to handle three CAN streams and an additional camera stream (refer to Section 2.2.3 of this document for more details on camera stream data capture process).

During Phase 3, three different SUR-CIN pre and post-trigger times were implemented for recording the event. With the first implementation, data was recorded 15s before the trigger and 15s after the trigger, so in total 30s. This time period was selected bearing in mind the known limitations of the Flea3 device (i.e. knowing that more data collected would result in higher system load). However, even with this relatively short time period, the system load was found to be too great and the data collection time had to be reduced as the Phase 3 trials progressed. The other two SUR-CIN trigger time implementations are described in the SUR-CROSS section (4.3).

4.2.3 Synthesis and analysis

4.2.3.1 Phase 2 (CIN) count vs Phase 3 (SUR-CIN) count

Initial analysis of the SUR-CIN data consisted of a like-for-like comparison of CIN and SUR-CIN events. During this analysis, it was discovered that around 50 % of the SUR-CIN measurements were missing. Follow-up investigations revealed that the reason for the discrepancy was largely due to high load on the second (SUR) Flea3 device. This caused CAN messages, including the actual trigger signal, to be missed by the second Flea3 device, which meant event recording was not started.

To solve the problem, different approaches were used. The first approach was to increase the priority for the CANLogger application (the application responsible for sampling the measurements) on the Flea3 device. This resulted in a slight improvement in the performance of the Flea3 and the correlation between Phase 2 (CIN) and Phase 3 (SUR-CIN) events increased to around 65 %. The second approach was to reduce the number of SUR-CROSS events recorded. During testing, through checking many combinations and configurations, it was discovered that the SUR-CROSS trigger was having a huge effect on SUR-CIN recordings and the decision was taken to reduce the number of SUR-CROSS geo-fence locations which, in turn, reduced the number of SUR-CROSS recordings. Unfortunately, this change, did not improve the situation, and so in a final attempt to solve the problem, it was decided to turn off SUR-CROSS recording, on the vehicle assigned to TRL, altogether.

At first, this approach seemed quite drastic, however, given that the vast majority of SUR events recorded by this vehicle were SUR-CIN events (around 85%), with only 15% being SUR-CROSS recordings, the impact on the total number of SUR-CROSS events recorded across the fleet was minimal. At the same time, the improvement in the number of SUR-CIN events recorded after SUR-CROSS recording on the TRL managed vehicle was turned off was significant, with the correlation of Phase 2 CIN events and Phase 3 SUR-CIN events increasing to around 95%. The decision was, therefore, taken by the consortium to turn off SUR-CROSS recording on the TRL vehicle for the remainder of the trials.

4.2.3.2 'Following' vehicle Time-To-Collision (TTC) during Cut-In events

Results of statistical analysis of TTC can be seen in Figure 46 below. These results were derived from a number of pre and post-trigger points taken from each SUR-CIN event recorded over the entire Phase 3 trial period.

The histograms in Figure 46 show the time-to-collision statistics for the vehicle following the ego-vehicle ('Following' vehicle) relative to the ego-vehicle around the point of cut-in (for all cut-in events captured in Phase 3).

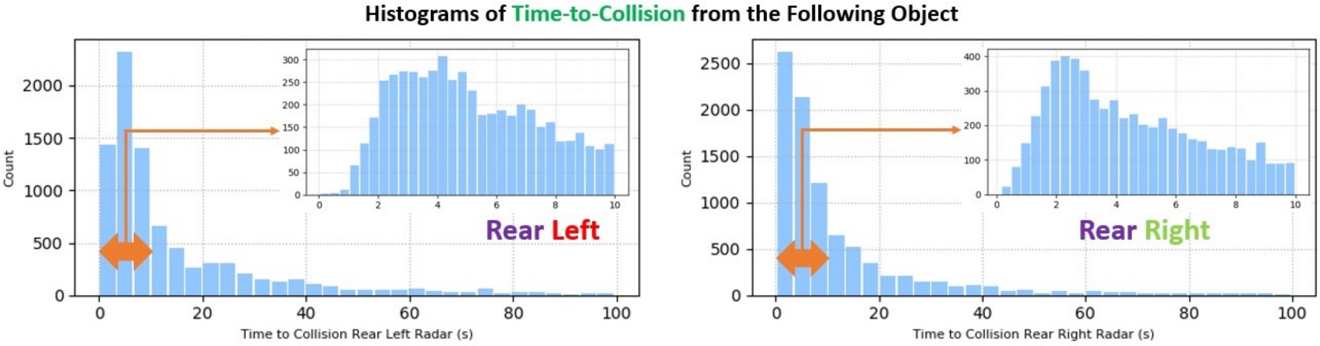

Figure 46 – Histograms of TTC for the following vehicle

It is clear from these histograms that the TTC for the following vehicle is generally between 1 and 10s around the point of cut-in, and rarely goes below the 1s mark.

This means that the drivers of following vehicles are, for the most part, leaving a sufficient time gap to the ego-vehicle in front of them to avoid a collision should the ego-vehicle brake sharply when another vehicle cuts-in in front of it. This is behaviour which automated vehicles of the future should emulate.

The above said, there are still a significant number of occasions when the TTC for the following vehicle is below 2s, which shows that some drivers leave smaller time gaps than the recommended 2s, and therefore could become frustrated if highly automated vehicles are designed to maintain time gaps of 2s and above only. This frustration could be made worse in a mixed environment of automated and non-automated vehicles, where drivers of non-automated vehicles could take advantage of the gaps left by automated vehicles and continually cut-in in front of them. Whilst safety should always take precedence, it is the MOVE_UK consortium's view that further work/investigation into the gaps which automated vehicles should maintain is required to ensure that automated driving systems do not cause users to become frustrated.

One thing is clear, however, if the time-gap between vehicles is less than 1s, an automated vehicle will have difficulty merging into a lane with other traffic. Again, further investigation of this kind of scenario is required.

4.3 Use case Surround sensing of Cross-traffic situations (SUR-CROSS)

Recording of SUR-CIN events is triggered using a situation-based event triggering process. In contrast, the SUR-CROSS use case (the second new use case developed in Phase 3) used a location-based event triggering approach. See Figure 47 below for examples of SUR-CROSS scenarios where location-based triggering was used.

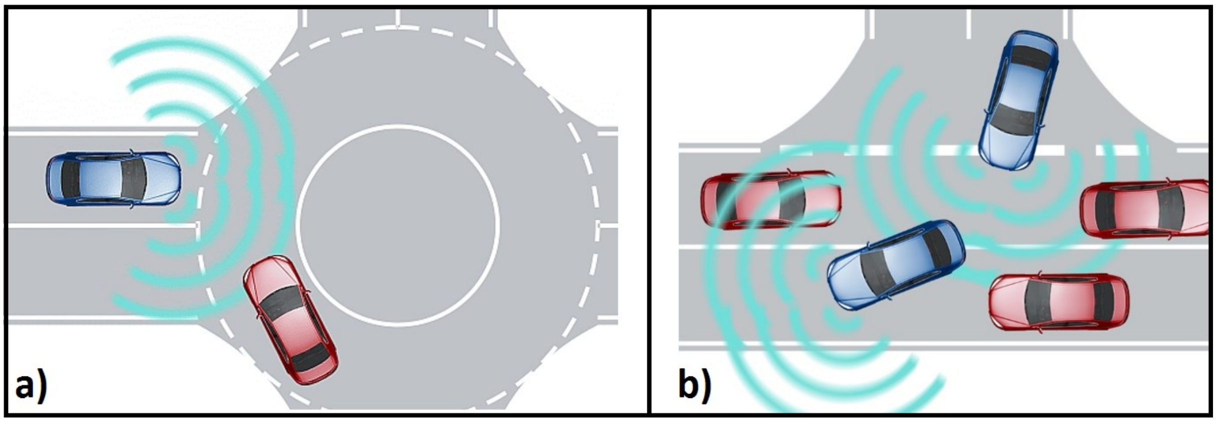

Figure 47 – SUR-CROSS example scenarios: a) roundabout b) T-junction

4.3.1 Purpose

The goal of this use case was to collect data which would allow the MOVE_UK consortium to better understand the behaviour of the ego-vehicle and surrounding vehicles when the ego-vehicle entered an intersection or roundabout and merged into traffic at these busy locations. This is a task which can be difficult for automated vehicles to complete without being overly cautious, something which could lead to driver frustration and potentially traffic congestion.

4.3.2 Design

A configurable functionality, known as a Geo-Fence-Monitor, was implemented on the second (Phase 3) Flea3 device, fitted to two of the vehicles. A number of locations were identified and added to a list. These locations each had specific GPS co-ordinate parameters and a defined radius which was 30 m. The second Flea3 was programmed to monitor the GPS co-ordinates along the route of the vehicle it was fitted to and compare them with those on the list. Whenever one of the two Phase 3 vehicles entered or exited the radius of the listed locations, an action was realised, the trigger for recording was then fired and the data collection process initiated.

The Floow used their analysis of junction complexity (refer to *MOVE_UK Data Analysis Report - Phase 2 (D7.4)* for more details) together with their telematics aggregation processing to identify the candidate geo-fence locations. This process identified locations which were both challenging from the perspective of autonomous driving and were visited frequently enough to ensure good statistical coverage for analytical purposes. The final list of geo-fence locations was manually refined to ensure as rich a variety of road topology as possible was covered, including roundabouts, T-junctions, Y-junctions and slip-roads.

Two versions of the SUR-CROSS locations list were created, at different stages of Phase 3. Initially, a list consisting of 20 locations around Greenwich and 15 around Wokingham was generated and deployed at the same time as the first version of the SUR-CROSS trigger. Figure 48 and Figure 49 show the initial distribution of the geo-fence locations (denoted by red dots) on maps of the Greenwich and Wokingham areas, respectively.

Figure 48 – First version of SUR-CROSS locations in Greenwich area

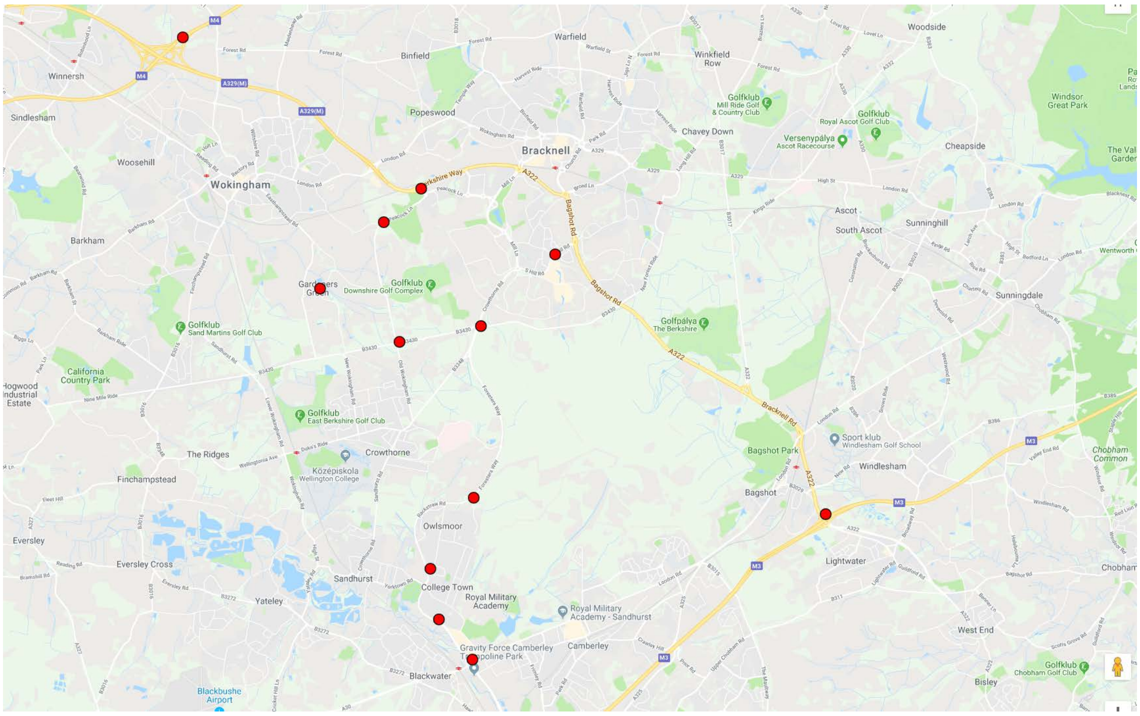

Figure 49 – First version of SUR-CROSS locations in Wokingham (TRL) area

After initial collection of data and subsequent analysis, it was concluded that some of the listed locations should be fine-tuned or deleted, and that new and more relevant locations should be introduced. As a result, the consortium created a second (final) list of locations, comprising 33 in the Greenwich area and 18 in the Wokingham area, thereby extending the list to 51 locations in total. Figure 50 and Figure 51 show the updated distribution of Geo-Fence locations in the two test areas.

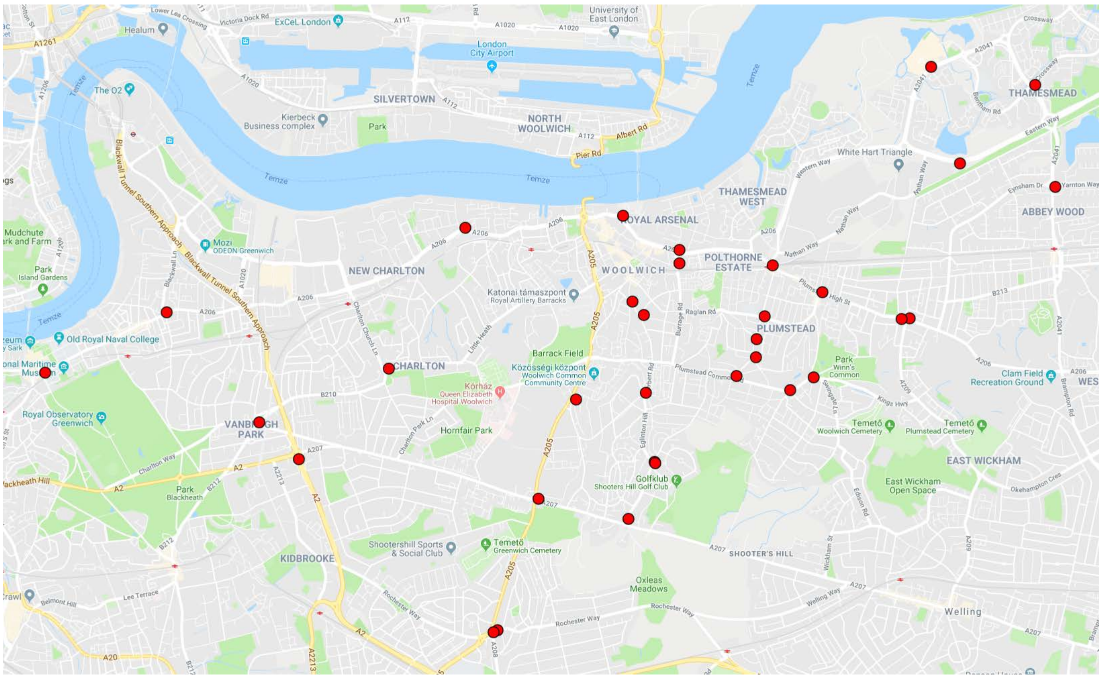

Figure 50 – Second version of SUR-CROSS locations in Greenwich area

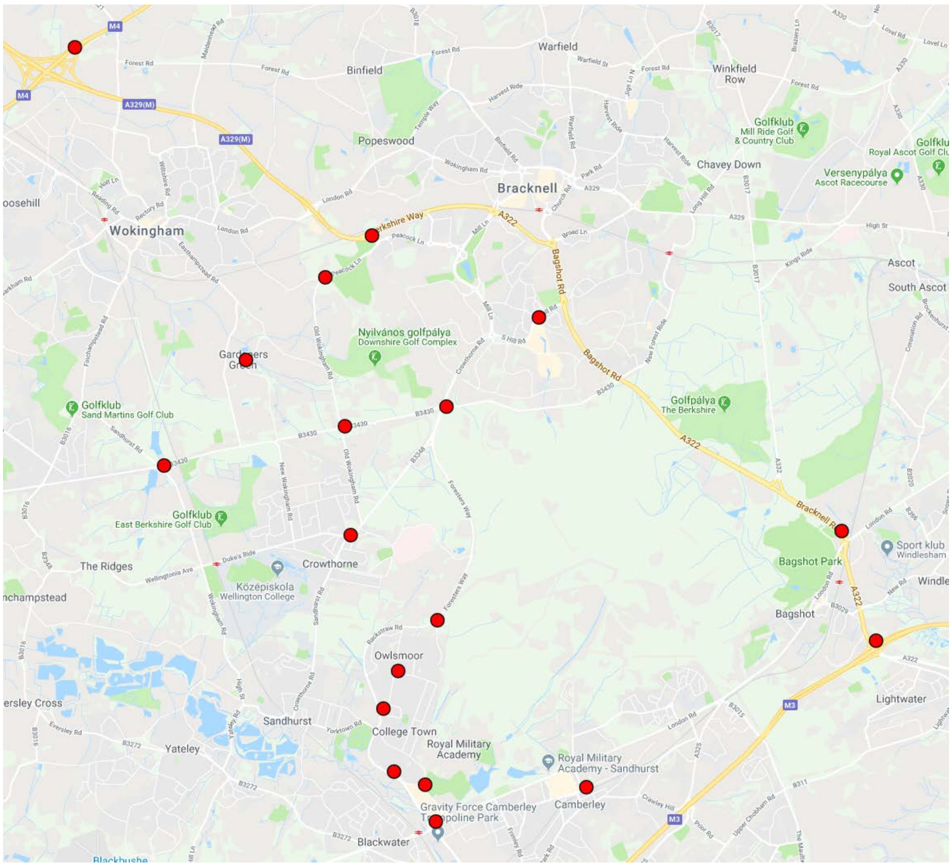

Figure 51 – Second version of SUR-CROSS locations in Wokingham (TRL) area

4.3.3 Synthesis and analysis

4.3.3.1 Analysis of initial results and trigger refinement

Soon after SUR-CROSS recording started, the consortium partners started to analyse the quality of the data that was being collected. The analysis revealed that, in a little less than 10% of the cases, the recording was found to be over before the real point of interest (i.e. the point at which the ego-vehicle was about to merge into traffic and/or turn) had been captured. The reason for this was found to be due to limitations in GPS accuracy, traffic jams leading up to the roundabout/intersection, and the chosen recording time (-15s to +15s).

In order to solve this issue, a new SUR-CROSS trigger was developed. During the development, six different options for a new trigger were analysed all of which were based on a combination of the existing geo-fence trigger signal and additional CAN signal conditions such as vehicle speed threshold, longitudinal acceleration or deceleration thresholds, and switch from brake pedal to accelerator pedal.

The analysis showed that if more conditions were introduced for triggering and/or more strict rules were created, important data would be lost and the real point of interest (i.e. merge into traffic and/or turn) would still probably be missed. To avoid this, a new approach was taken based on using the exit from the geo-fence location (where traffic build-up was generally much less than at entry) as the base for the trigger, rather than the entry to the location. At the same time, pre-trigger time was changed to 20s, in order to make sure that the main point of interest was not missed. Post-trigger time was also changed to 7s and then later on, to 5s, as it turned out that 5s was sufficient for analysis.

The result of all these changes was not only the successful capture of the main point of interest but also a reduction in size of each SUR-CROSS event/measurement by an average of 7.5% (due to shorter recording time). This led to less system load and greater stability.

Notwithstanding the above, due to the effect of SUR-CROSS event recording on SUR-CIN recording, it was decided to turn off SUR-CROSS recording on the vehicle assigned to TRL in the last few months of the Phase 3 trials. For further details of the reasons for this change, please see section 4.2.3.1 above.

4.3.3.2 GPS error

During further analysis of the captured SUR-CROSS events, a small flaw was identified in the second (Phase 3) Flea3 device. In very rare cases, around 1.5 % of all measurements, the Flea3 picked up the longitudinal co-ordinates of the location with the wrong sign, from the list of configured locations, and started recording in the wrong area. In Figure 52, which shows a location with identity G_12 in the Greenwich area, instead of recording measurements from co-ordinates 0.024 longitude and 51.474 latitude, the system recorded measurements from co-ordinates -0.024 longitude and 51.474 latitude. The solution for this issue was out of scope for this project and therefore was not investigated further.

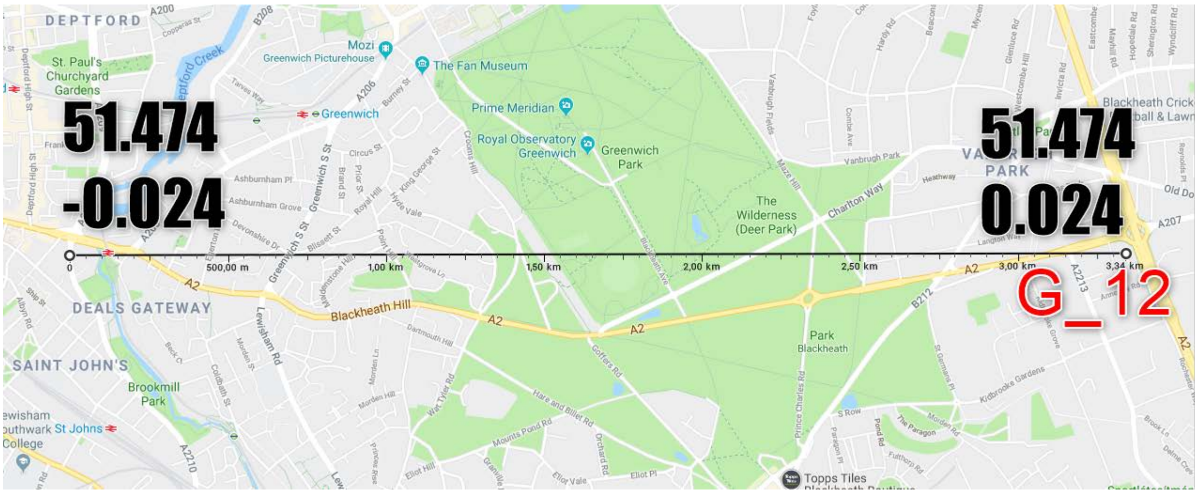

Figure 52 – Location picked up with the wrong longitudinal sign

4.3.3.3 TTC analysis

One of the tasks carried out by the projects data analytics team was to identify hotspots where the TTC was lower than usual, either from the front or to the rear of the ego-vehicle. This was done by generating heat-maps from front and rear corner radar data. TTC was plotted only when it was less than 2s either from the front or to the rear of the ego-vehicle.

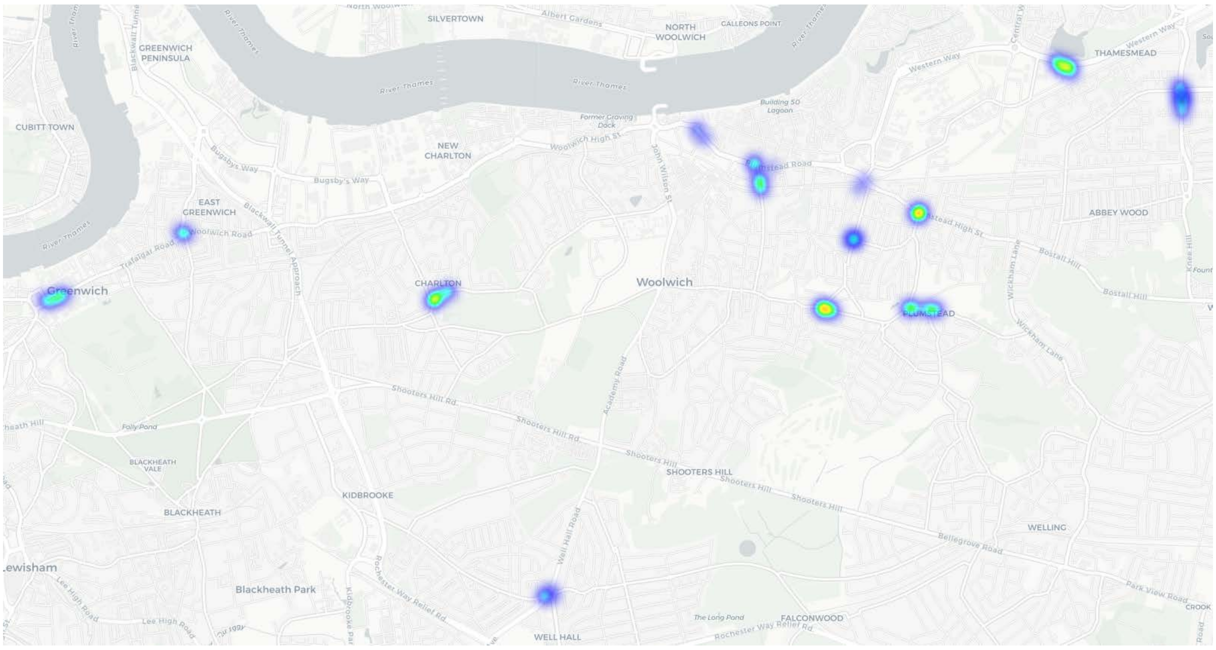

Figure 53 – Heat-map of TTC from the front of the ego-vehicle (brighter shading = lower TTC)

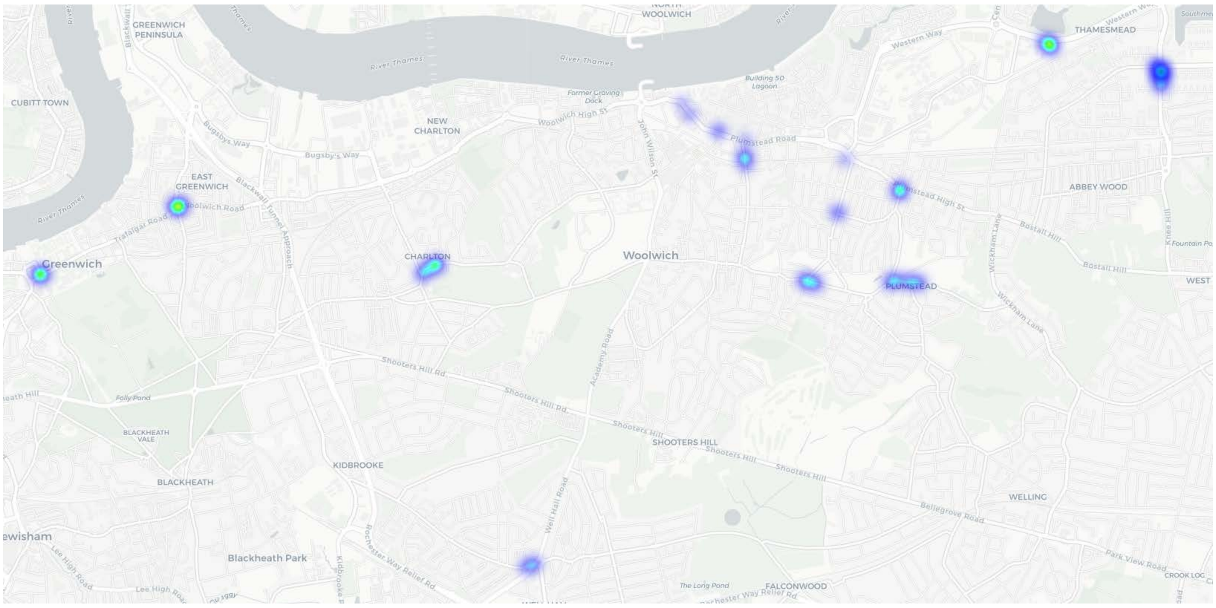

Figure 54 – Heat-map of TTC to the rear of the ego-vehicle (brighter shading = lower TTC)

From Figure 53 and Figure 54, it is clear that there were some locations where TTC was low to the rear but not from the front of the ego-vehicle, and vice-versa. There were also some junctions where the values seemed to be very similar on either end. On average TTC from the front was lower than to the rear.

4.3.3.4 Classification and clustering of junctions

Another task carried out by the projects data analytics team was to use the collected data to develop algorithms to classify the junctions into categories such as T-junctions, Y-junctions, 4-way junctions and roundabouts. The follow-up task to this was to discover the junctions used most by the vehicle fleet. The results of this analysis are shown below in Figure 55 (brighter shading = positions the ego-vehicle spent most time in, during the SUR-CROSS event, and vice-versa).

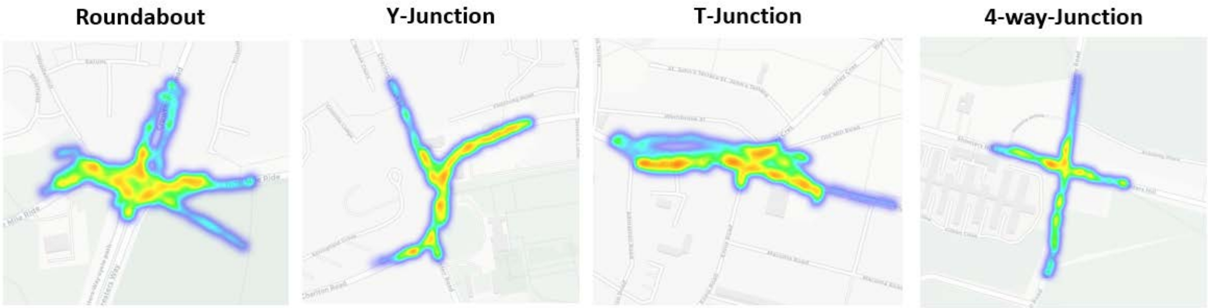

Figure 55 – Most frequently visited locations from each category of junction

4.3.3.5 Route-based analysis of SUR-CROSS data

SUR-CROSS sequences were categorised by 'route' to enable the clustering of similar sequences, in order to build up a deeper statistical understanding of SUR-CROSS events. Routes were defined, for example, to make a distinction between sequences involving left-turns and right-turns at a particular junction. The definition was done using the road segment map developed by The Floow, which is a global map of all roads segmented into short sections.

4.3.3.6 Bespoke route based analysis of SUR-CROSS data by location

The value of SUR-CROSS data stems from its ability to inform on situations where human drivers are required to make a judgement. For example, consider the geo-fence location at the entrance to the Greenwich council depot where four of the five MOVE_UK vehicles were based (this was the location with the highest number of SUR-CROSS triggered events). Figure 56 below shows several instances from SUR-CROSS sequences that passed through this geo-fence location, colour-coded by journey route (red points are those journeys that leave the depot, purple points are journeys returning to the depot, blue points are journeys that pass the depot entrance).

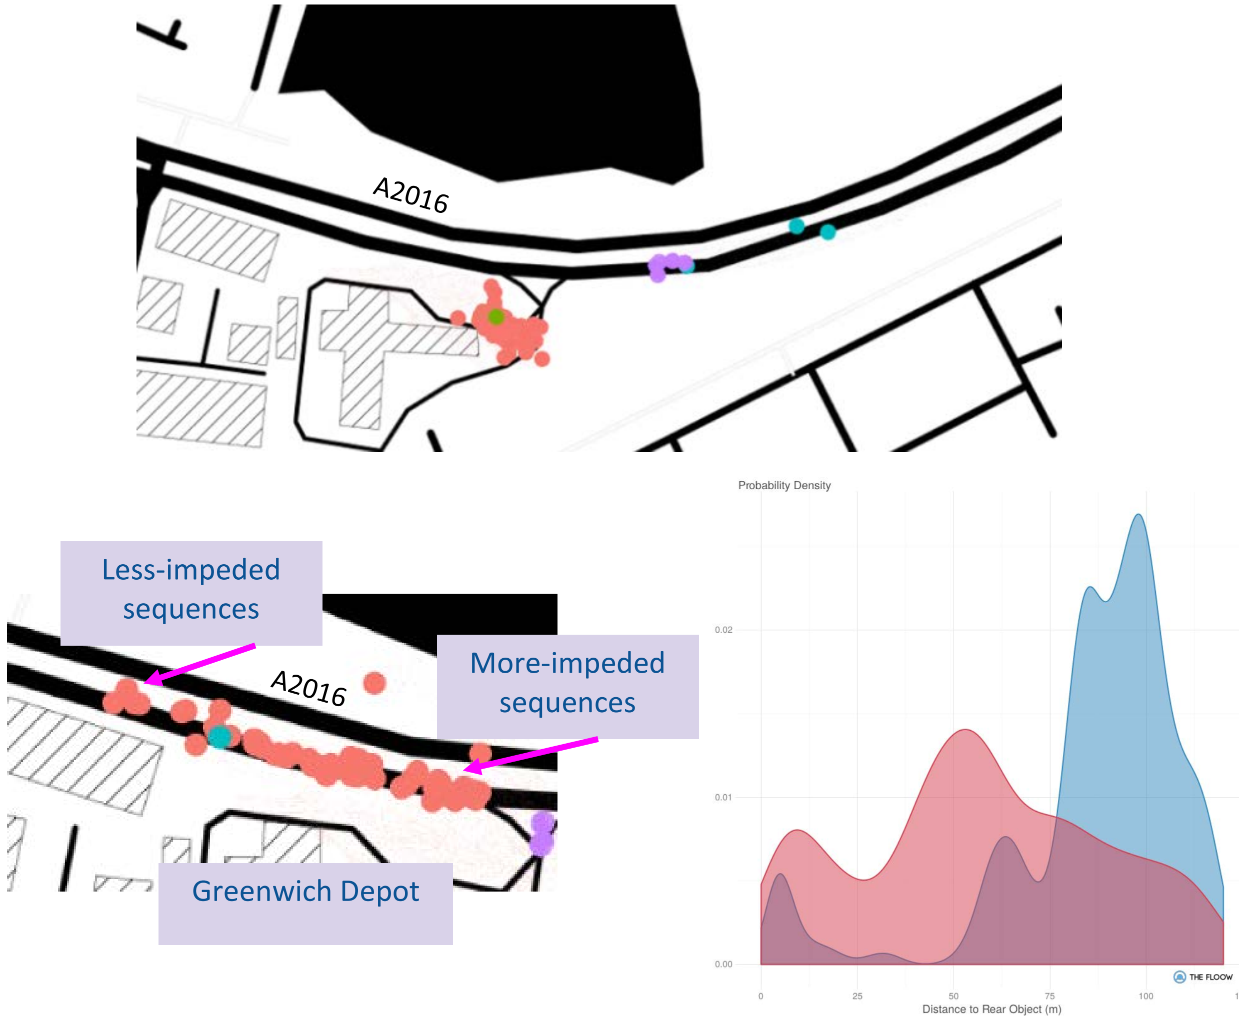

Figure 56 – Clustering of journeys at Greenwich council depot entrance

Exiting the depot car park involves taking a minor slip road onto the A2016 road, where the speed limit is 80 km/h (50 mph) but with strong excessive speeds observed. The slip road eventually runs parallel to the A2016, permitting vehicles leaving the depot to join the flow of traffic safely. The safe passage from the slip road to the A2016 road requires observations, judgements, and decisions to be made by the driver. First of all, the driver must observe if there are vehicles coming from behind, and if so, decide whether the following vehicles are at a distance such that joining their lane will not prove unsafe. The driver can then decide whether to join the A-road traffic or to defer entry until such point that a manoeuvre can be performed safely.

Using the data from the rear-facing corner radars, it is possible to begin to understand the specifics of what drivers consider to be safe when it comes to performing this manoeuvre. In the left panel of

Figure 56 comparison can be made between the SUR-CROSS sequences for which the ego-vehicle travels the least distance, during the ten seconds after leaving the depot, and those sequences that cover the furthest distances. The adjacent panel shows the corresponding detection probability densities for the rear corner radars as a function of distance, with the red curve corresponding to those sequences when the vehicle travelled shorter distances and the blue curve corresponding to those occasions when the vehicle progressed onto the A2016 road more rapidly. Clearly, there is a greater probability of the radar detecting an object closer than 70 m (corresponding to just over 3 seconds of travel time at 80 km/h) when the ego-vehicle takes longer to transition onto the A2016 road. This very simple statistical analysis provides a direct insight into the judgement of the driver and how they modify their behaviour in light of knowledge of the surrounding traffic.

4.3.4 Conclusions

As already mentioned, the geo-fence locations used for the SUR-CROSS analytics were defined using prior knowledge from Phase 1 and Phase 2 (1 Hz) data. However, following initial dense (10 Hz) data collection in the SUR-CROSS use case and related analysis, the list of locations was refined. This confirmed the already known fact that dense data is more insightful than lower resolution data.

The TTC analysis proved the surround sensors' capability of detecting areas where the traffic is denser or more aggressive driver behaviour is common. These areas may be more difficult areas for automated vehicles to navigate and need to be investigated further.

From the 'classification and clustering of junctions' analysis, the possibility to identify the busiest junctions using surround radar data was established. From this, it is also possible to study these junctions further or earmark them for the testing of features fitted to automated vehicles of the future.

The extent to which there is 'hand-over' between the forward and rear corner radar sensors was examined, and it was found that, passing objects were not detected by both forward and rear sensors simultaneously (i.e. no field of view overlap), with detections registering in the rear corner radars a couple of seconds after disappearing from the forward corner radars. In practice, this result means that vehicles with this configuration of radar have similar 'blind spots' to human drivers. It should be noted, however, that in production systems there would be fusion of the radar sensor data to eliminate such blind spots.

The geo-fence based analysis approach has proven to be useful in terms of being adaptable for other use cases, as demonstrated in its application to determining TTC hotspots. It can also be applied in other use cases for example in the Traffic Sign Recognition use case, which is covered in section 3.4.

Finally, section 4.3.3.6 showed that vehicles fitted with surround sensors present a huge opportunity to learn about the decisions taken by human drivers. Data that can help the understanding of drivers' behaviours when interacting with other road users will be of value to those developing automated vehicles, to insurers for the purposes of identifying when customers participate in higher-risk manoeuvres, and to local authorities for providing an evidence-base to inform on improvements to infrastructure.

4.4 Use case Telematics – Phase 3

4.4.1 Insight for existing telematics

At present, telematics analysis overwhelmingly focuses on processing GPS data that is collected from black boxes, smartphones, or devices that plug into on-board diagnostic (OBD) ports. MOVE_UK has created a rich dataset consisting of continuous 1 Hz measurements over hundreds of parameters for thousands of journeys. These richer data present the opportunity to enhance existing telematics propositions in two important ways. First, the data can be examined to determine the absolute accuracy of GPS data and its reliability to inform on behaviours-of-interest. Secondly, there is the potential to expand our knowledge of which behaviours can be observed using GPS data alone; using direct measurements of those behaviours to identify any corresponding characteristic patterns in the GPS data.

The primary concern for motor insurance providers is the assessment of risk and the extent to which risk can be understood using telematics. It is well established that drivers who frequently indulge in excessive speeding will be more likely to be involved in incidents with greater frequency and severity than those drivers who behave more conservatively. GPS data contains a direct measurement of speed that can be used to help assess the extent to which the driver is operating safely. Additional quantities, such as the longitudinal acceleration of the vehicle, might not be measured directly by some devices but can be estimated from these GPS speed measurements. Knowledge of a vehicle's acceleration provides another valuable tool with which risk can be parameterised, serving as an effective proxy for driver aggression and preparedness. For example, one can infer a driver's 'smoothness', the extent to which a vehicle will slow down gradually to approach a junction or will tend towards an abrupt halt. The latter could suggest that the driver is either overly aggressive or not applying an appropriate level of foresight, both of which increase the risk of collision with other road users. Clearly, understanding the consistency and accuracy of GPS data is crucial to providing reliable telematics solutions using current technology.

In Phase 3, a series of experiments were carried out to understand the measured GPS speed as a function of the speed of the vehicle as measured at the wheels. Figure 57 shows a direct comparison between the wheel speed and the speed measured by GPS. While the figure shows excellent agreement between the two, subsequent analysis focused on identifying a set of parameters for which we suspect the relationship draws from broader distributions (e.g. during manoeuvres-in-the-road at low speed) or will break down all together (e.g. such as during the approach of a vehicle to a tunnel, when the quality of GPS data drops).

The Floow have explored the enriched MOVE_UK data to help validate and enhance their ability to detect risk-linked behaviours, such as crashes and extreme cornering. These investigations pave the way to improving traditional telematics technology.

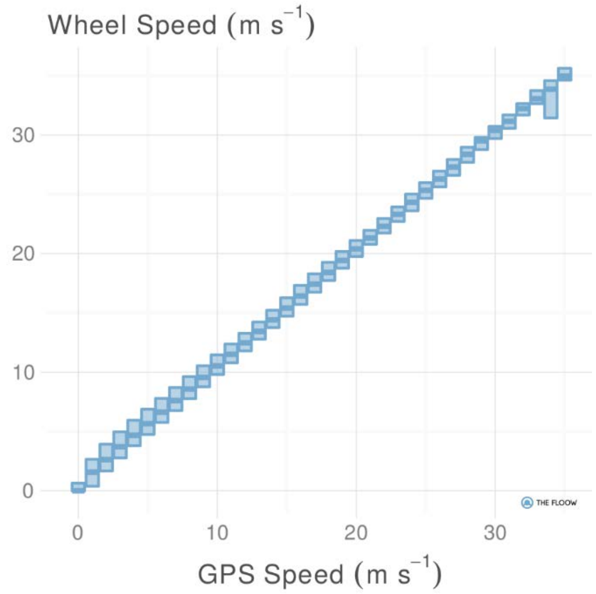

Figure 57 – Comparison of wheel speed and the speed measured by GPS

4.4.2 Future telematics

In the near- to medium-term future, most new vehicles will have the connectivity to enable the potential use of richer data. This connectivity could be provisioned by aftermarket devices or via a direct transmission from the vehicle. The exact make-up and operational parameters of these data will vary between manufacturers and by vehicle type. Such a large data lake will prove a boon to insurers because of its potential to power more insightful risk estimation. Inevitably, telematics providers will have to decide a trade-off between the amount of data they would like and the cost of obtaining such a dataset (a function of the cost to collect, store and process the data).

MOVE_UK has developed a methodology for evaluating the value of parameters for EDR and risk assessment purposes. This methodology is an essential step towards reducing the time required to develop new telematics solutions when accessing more data becomes financially viable. Ultimately, this approach will give telematics providers the ability to understand the extent to which a parameter, or set or parameters, provides meaningful insight into driver behaviour. Practically, this work will accelerate the creation and refinement of new metrics for grading the potential risk from vehicle operation.

The Floow chose several of the higher-value parameters as a starting-point from which to develop prototype behavioural metrics (hereafter, ‘scores’). These prototype scores fall into two broad categories; existing scores that are improved by using better data and new scores that are not inferable from typical telematics data.

At present, smoothness is inferred from the acceleration of the vehicle, which is calculated from GPS or accelerometer data. However, the MOVE_UK data contain many additional parameters that can be combined to produce a much more informed measure of smoothness. For example, we know the forward acceleration with greater accuracy because the wheel speeds are directly and independently measured, which also means that we can describe the exact motion of the vehicle through corners. In addition, we also have a direct measure of the driver’s interaction with the vehicle from the pressure applied to both the accelerator and brake pedals, as well as the dynamics of the steering wheel. Using this richer array of parameters enhanced scores can be created that offer more insightful and fine-grained views of risk-linked behaviours than is currently available in the telematics marketplace.

Entirely new scores were developed during the project. Most notably, The Floow used data from the radar and camera systems to produce a Tailgating score (see schematic in Figure 58), based on the distance a driver will maintain to the vehicle in front as a function of speed and road type. The Tailgating score was developed synoptically using the output of various investigations, such as determining the nominal distances to the vehicle-in-front as a function of speed, the consistency of the sensor systems, and the reliability of the radar's lane assignment software. Other examples of scores that reached the prototype stage include 'courtesy' (based on how the driver chooses to deploy the forward-facing lights), 'passenger distraction' (the extent to which a driver's behaviour is influenced by passengers) and 'dwell judgement' (a measure of the driver's ability to make decisions at junctions in light of oncoming road users).

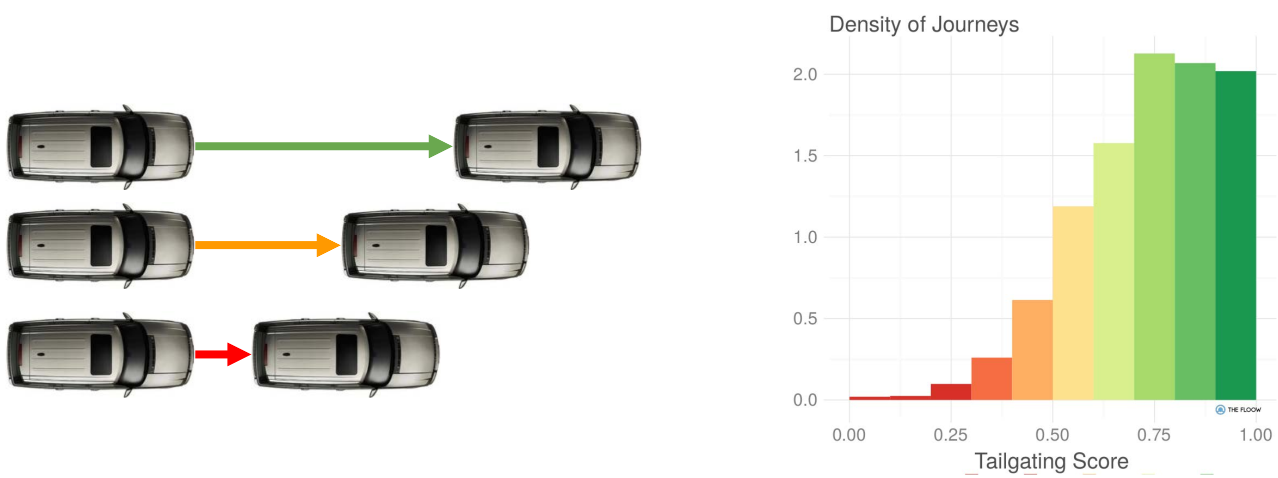

Figure 58 - Tailgating score based on the distance a driver will maintain to the vehicle in front as a function of speed and road type

Ultimately, this use case established the priorities and methodology for creating the next generation of scoring capabilities.

5 Conclusions

The use of event-based data collection, where only relevant data is recorded and stored has been shown to save a lot of data storage space and processing time compared with conventional data collection. This new approach enables shorter development cycles for automated driving systems, as only relevant data is processed. This reduces the time associated with moving, checking, reviewing and annotating the data prior to reprocessing the data with newer software, for instance in a Hardware in the Loop (HiL) simulation system. Therefore, event-based data collection was demonstrated to be a valuable method in reducing time and resources required for future Advanced Driver Assistance Systems (ADAS) and Automated Driving Systems (ADS) testing and development. It is the intention of JLR and Bosch to use this method to help develop its next generation of ADAS features.

The results of re-simulation of the captured Advanced Emergency Braking (AEB) events with production camera software were far better than expected and highlight the relevance of the data collected from the MOVE_UK trials. Critical AEB events, such as those collected in MOVE_UK, are very difficult to capture in the real world and are extremely useful for the development and validation of future ADS features. The main reason for their usefulness is that they often cover driving situations which the engineers/developers did not foresee and also situations which testing bodies (like Euro NCAP) do not specify currently. Using these events/sequences to test ADS software in a simulation environment, therefore, leads to more broadly effective and safer ADS features. It is the intention of Bosch to use the 23 critical AEB events captured to help test and develop its next generation of production AEB software.

Through the Traffic Sign Recognition use case, MOVE_UK developed an 'evidence-based methodology' that can determine the detection rates of speed signs and their change in detection likelihood over time. The methodology can be used to identify locations where the recognition system is performing poorly and to notify the location(s) to concerned parties for restorative interventions which would improve detection going forward (via a feedback loop system). This methodology could be adapted to work with other systems too, such as pothole detection and other road signs, road markings and infrastructure detection that support automated driving.

For Phase 3 of MOVE_UK, the Phase 1 and 2 vehicle modifications were extended in two of the five Land Rover vehicles with the addition of four corner radar sensors and the corresponding adaptions to hardware and software in each of the two vehicles. The corner radars were added to allow full 360-degree surround vehicle sensing. This enabled the following observations to be made.

The Time-To-Collision (TTC) for the following vehicle in a Cut-In scenario is generally between 1 and 10s, around the point of Cut-In, and rarely goes below the 1s mark. This shows us that the drivers of following vehicles are, for the most part, leaving a sufficient gap to the ego-vehicle in-front of them to avoid a collision should the ego-vehicle brake sharply when another vehicle cuts-in in front of it; this is a behaviour that autonomous driving vehicles of the future should emulate. However, there are still a significant number of occasions when the TTC for the following vehicle is below 2s, which shows that some drivers leave smaller time gaps than the recommended 2s, and therefore could become frustrated if automated vehicles are designed to maintain time gaps of 2s and above only. Further investigation into the gaps which automated vehicles should maintain is required to ensure that automated driving systems do not cause users to become frustrated. It is also important to note that if the time-gap between vehicles is less than 1s, an automated vehicle will have difficulty merging into a lane with other traffic. Again, further investigation of this kind of scenario is required.

The geo-fence based analysis approach was valuable in identifying sites for evaluation of traffic flow and driver behaviour. This approach has potential for being adapted for other use cases, as demonstrated in the determination of hotpots where TTC was observed to be low.

From the investigation of behaviours around specific locations, it is clear that vehicles fitted with surround sensors present a huge opportunity to learn about the decisions taken by human drivers. Data that can help the understanding of drivers' behaviours when interacting with other road users will be of value to those developing automated vehicles, to insurers for the purposes of identifying when customers participate in higher-risk manoeuvres, and to local authorities for providing an evidence-base to inform on improvements to infrastructure.

The high value of the forward radar data to both EDR (Event Data Recording) and risk assessment was demonstrated during Phase 2, and the analysis from the Phase 3 use cases strongly suggests that the surround data can add significant value to our understanding of both.

6 Glossary of terms

Abbreviations used	
ADAS	Advanced Driver Assistance System
ADS	Automated Driving System
ADTF	Automotive Data and Time-Triggered Framework
AEB	Autonomous Emergency Braking
ARB	Radar-based Autonomous Emergency Braking
AV	Autonomous Vehicle
CAN	Controller Area Network
CCAV	Centre for Connected and Autonomous Vehicles
CCU	Connectivity Control Unit
DHB	Driver Harsh Braking
CIN	Cut-in Scenarios
EADM	Enterprise Automotive Data Management
ECU	Electronic Control Unit
EDR	Event Data Recorder
ESC	Electronic Stability Control
GPS	Global Positioning System
HiL	Hardware in Loop
LDW	Lane Departure Warning
LRR	Long-Range Radar
LVS	Lead Vehicle Statistics
MRR	Mid-Range Radar
OEM	Original Equipment Manufacturer (vehicle manufacturer)
OSM	Open Street Map
RBG	Royal Borough of Greenwich
sFDE	Systematic Field Data Exploration
SUR-CIN	Surround sensing of Cut-In situations
SUR-CROSS	Surround sensing of Cross-traffic situations
TSR	Traffic Sign Recognition
UI	User Interface
VRU	Vulnerable Road User

Appendix 1. Additional information – Use case Driver Harsh Braking

A1.1. Classification of Driver Harsh Braking Events

TYPE 1 events were classified as 'avoid impact' and the object of the near-crash were specified as in Table 8.

NOTE: Only instances where there is a risk of imminent collision with a road user or an object were classified using this category.

Table 8 – Categorisation of Type 1 DHB events

AVOID IMPACT WITH MOTORISED VEHICLES	
Avoid impact – car	For passenger cars, SUVs, people carriers, family vans
Avoid impact – PTW	For powered two wheelers and scooters
Avoid impact – van	For light goods vehicles (up to 3.5 tonnes), including pick ups
Avoid impact – bus	For large size passenger vehicles, including minibuses, buses, coaches
Avoid impact - HGV	For heavy goods vehicles (over 3.5 tonnes), including bin lorries
AVOID IMPACT WITH VULNERABLE ROAD USERS	
Avoid impact – pedestrian	For pedestrians
Avoid impact – 2W (two-wheeler)	For cyclists and pedestrians using different types of scooters (e.g. push bikes, e-scooters), wheelchairs.
Avoid impact – other	Animals, drones and everything else

Type 2 events were classified according to the contextual factors requiring sudden braking. Contrary to Type 1 events, to be classified as a Type 2 event, a driver will have had to brake because of infrastructure or road layout (e.g. speed bumps, junctions, roundabouts, narrowing lane requiring driver to slow down for oncoming traffic).

NOTE: Judgement of whether events should be classified as Type 1 or 2 were made based on the distance from other vehicles, VRUs or objects as per notes in red below. Understandably, in this case 'the relevant object of potential impact column' should be left blank (e.g. n/a).

NOTE: Distinction between give way to vehicular vs. pedestrian traffic were made based on the presence of pedestrians waiting to cross. If a driver brakes to give way at a junction with traffic lights and marked pedestrian crossing, but no pedestrians are waiting, then classify as 'give way – junction with traffic lights'; if there are pedestrians waiting to cross then classify as 'give way - crossing with traffic lights'.

Table 9 – Categorisation of Type 2 DHB events

GIVE WAY TO VEHICULAR TRAFFIC	
Give way – junction with traffic lights	When a driver is reducing speed/stopping to give way to other vehicles at junctions (including roundabouts) with activated traffic lights for vehicular traffic (red or amber) (Classify as Avoid impact-vehicle if there is another vehicle in front that is braking/stopping in front of the traffic light)
Give way – junction without traffic lights	When a driver is reducing speed/stopping to give way to other vehicles at junctions (including roundabouts) without traffic lights for vehicular traffic (Classify as Avoid impact-vehicle if there is another vehicle in front that is waiting for a gap before joining traffic at a junction/roundabout; if already engaged in turning and a vehicle appears – then classify as Avoid impact- vehicle)
GIVE WAY TO PEDESTRIAN TRAFFIC	
Give way – crossing with traffic lights	When a driver is reducing speed/stopping at activated traffic lights for vehicular traffic (red or amber) to give way to pedestrian/s crossing (If there is a vehicle in front then classify as Avoid impact - vehicle; classify as Avoid impact – pedestrian if a pedestrian has already started crossing – foot on road)
Give way – crossing without traffic lights	When a driver is reducing speed/stopping to give way to pedestrian/s crossing at uncontrolled crossing (without traffic lights but with some form of road markings) (If there is a vehicle in front then classify as Avoid impact - vehicle; classify as Avoid impact – pedestrian if a pedestrian has already started crossing – foot on road)
Give way – other	Any other scenario where a driver brakes to give way to another vehicle or VRU, that doesn't fall under the above categories
LEAVE GAP	
Leave gap – oncoming traffic	When a driver needs to slow down/stop to leave gap for oncoming traffic due to narrow lanes (e.g. extended sidewalks, presence of bollards, bike lanes, middle islands, parked cars on the side etc.) (If there is a risk of imminent impact – then classify as Avoid impact- vehicle)
Leave gap – overtaking vehicle	If driver brakes to leave sufficient gap for a vehicle from oncoming direction to complete overtaking
Leave gap – other	When a driver needs to brake to leave sufficient gap for another road user under any other situation (e.g. temporary lane closure, road works, fencing)
REDUCE SPEED	
Reduce speed - turn	Before turning (not at a junction)
Reduce speed – vertical deflection	When a driver brakes due to vertical deflection calming measures (e.g. speed bumps, humps, cushions, tables etc.)
Reduce speed – other	Looking for a parking space, wrong turn, 3-point turn, low speed manoeuvring, narrow lanes (parked cars on the side), and all other brake manoeuvres that don't fit into one of the other categories.

A1.2. Time to Collision

A1.2.1. Defining the key time period

The first important question that was considered was how to define the key time period for each 20 second sequence included in the final sample. In general, the relevant braking event only lasted for a small portion of the sequence, so there were going to be periods which were not relevant to the analysis. In addition, as has already been accounted for, the trigger for each event type always occurred at 15 seconds into the sequence, so the relevant phase of braking should occur around that point. For each sequence, the key time period was defined as follows:

1. In the 20 second sequence, find the frame where the maximum level of brake pressure is applied. This is time point 2;
2. From this frame, work backwards through the sequence until the brake pressure is equal to zero (i.e. the start of the braking phase which includes the point of maximum brake pressure). This is time point 1; and
3. The key time period is defined to be the part of the sequence that lies between time point 1 and time point 2.

Figure 59 shows the level of brake pressure over the 20 second duration for one of the DHB sequences, with time points 1 and 2 matching the descriptions as above indicated.

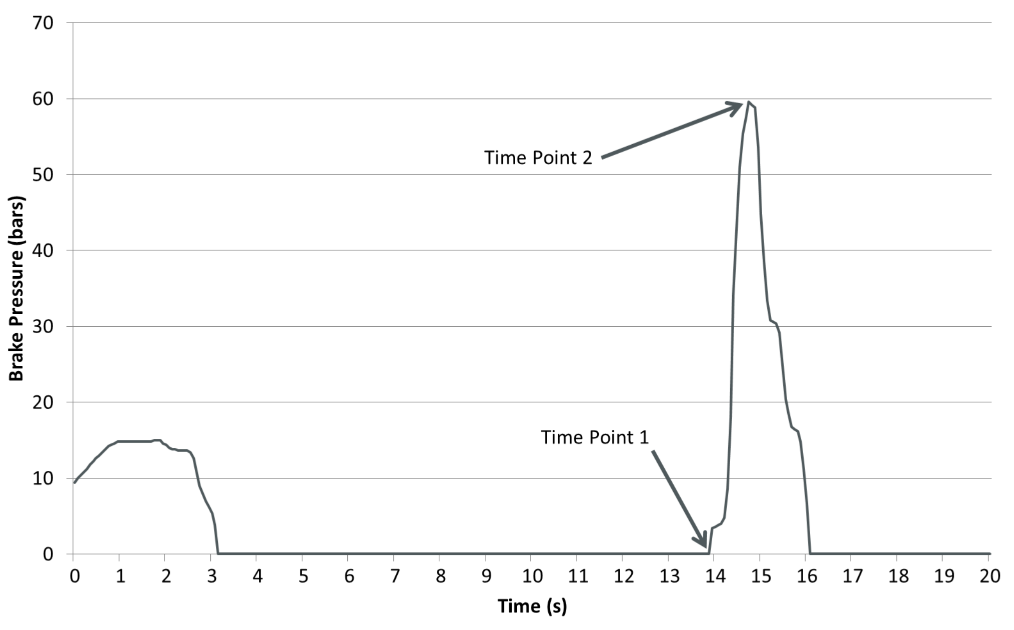

Figure 59 – Level of brake pressure over 20 second DHB sequence with time points 1 and 2 indicated

In addition to the three parts of the definition described above, there was a need to account for sequences where a vehicle was not detected as OBJ_01 at time point 1, but then became the most relevant object (OBJ_01) at some point during the braking phase, before the maximum level of brake pressure was reached at time point 2. Table 10 presents the number of sequences of each event type for which this was the case. This included the majority of DHB sequences and some of the ARB sequences. Furthermore, Figure 60 shows that when a different object was detected, this typically happened for most of the key time period.

Table 10 – Sequences where OBJ_01 changes during key time period

Event Type	Number of sequences in final sample	Number of sequences where OBJ_01 changes during key time period
DHB	27	22
AEB	16	0
ARB	92	16

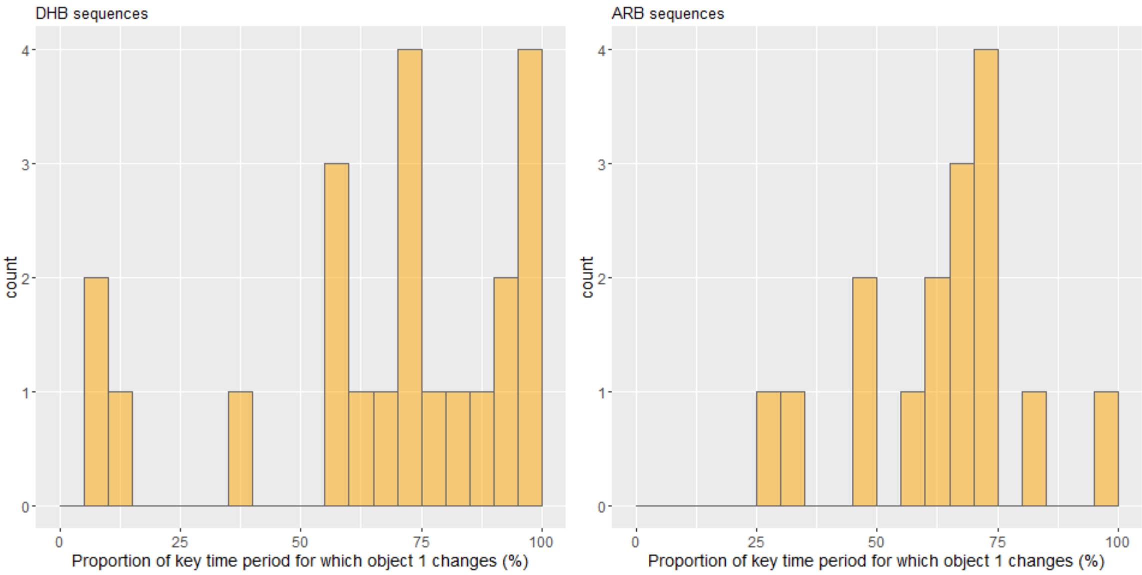

Figure 60 – Proportion of the key time period for which a different OBJ_01 is detected, for DHB (left panel) and ARB (right panel) sequences where this is the case

For DHB sequences, this meant that for 22 out of the 27 events, the wrong object was considered to be the target for at least part of the key time period. For the majority of these events, this was case for more than 50% of the key time period. There was less of a problem for ARB events – only 16 of the 92 events were impacted, and on average a smaller percentage of the key time period was affected.

In order to ensure that TTC was not calculated based on misleading information, the following approach was undertaken for these 22 DHB and 16 ARB sequences. For each sequence where the problem was encountered, the minimum TTC was calculated by only considering the parts of the key time period where the object identified as OBJ_01 by the CAN system was the lead vehicle, or more precisely, the object identified as OBJ_01 at time point 2. The other parts of the key time period were ignored.

This methodology relied on the assumption that at the most critical point of the sequence (time point 2), the lead vehicle was the object considered to be the most relevant. This was true for the vast majority of cases, however the higher percentage bars on the left-hand panel of Figure 60 demonstrate that for some DHB events, this may not necessarily have been the case.

A1.2.2. Calculating TTC

The second question that was considered was how to calculate TTC from the CAN and radar signals fitted to the MOVE_UK vehicles. The two pieces of information which were required to do this were the distance between the ego vehicle and the lead vehicle, and the rate at which this distance was changing. These were obtainable either from the CAN signals, S6D_OBJ_01_RadialDist and S6D_OBJ_01_RadialVel, or from the radar signals, OBJ_01_dx and OBJ_01_vx. However, data from the radar signals was generally more prone to dropping in and out during the sequence, making it less reliable. At the point of maximum brake pressure, around 20% of sequences did not give a value for either OBJ_01_dx or OBJ_01_vx. Therefore, it was decided to focus on data from the CAN signals for this analysis.

The radial distance measures the longitudinal distance between the ego vehicle and OBJ_01. In the sequences considered here, this is always the lead vehicle. As Figure 61 shows, the distance was measured from the rear axle of the ego vehicle. This was compensated for by subtracting 3.586m, the approximate distance from the rear axle to the front of the ego vehicle.

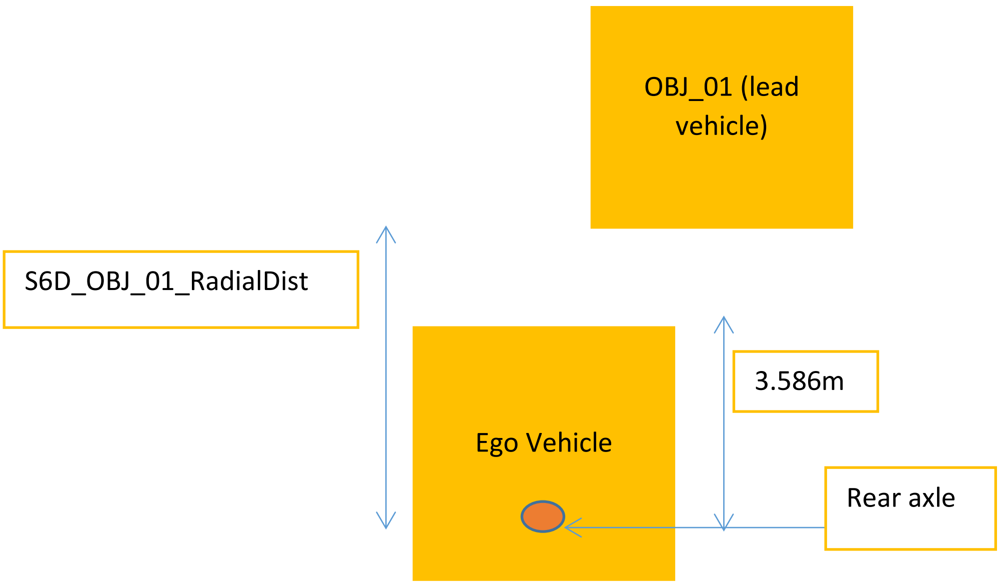

Figure 61 – How radial distance between two vehicles was measured

The radial velocity measures the rate of change in the radial distance. Therefore, these two signals provided the following simple calculation for TTC at any given frame in the sequence:

Time to Collision = Radial Distance / Radial Velocity

This calculation was performed at each of time points 1 and 2, as defined earlier. In addition, it was performed at every frame during the key time period in order to obtain the minimum TTC over the period of interest, and the rate of change in TTC over the same time.

Once this methodology was implemented, plausibility checks were performed on a subset of relevant sequences to ensure that the values of radial distance and radial velocity, and subsequently of TTC, corresponded to what was seen in the videos of the sequences. It should be noted that there was limited information available on the level of accuracy of the CAN signals, and so when applying this approach more generally, care should always be taken to ensure that any outlying or erroneous values are identified and removed.

A1.3. DHB cluster analysis – generating clusters

Clusters (groups of similar braking events) were generated using K-means clustering in the programming language R. The main improvement that was made to the methodology since Phase 2 was to use a technique known as bootstrap clustering to more robustly determine the optimal number of clusters, K, that the data should be grouped into. This method assesses the stability of the clusters that are formed for a range of possible values of K. The greatest value of K that results in sufficiently stable clusters was selected.

Initially, K-means clustering was performed using data from the six variables on all 213 DHB events from across the course of the project. Unlike in the previous analysis, events from Phase 1 were now included as they contained data for each remaining variable. Of this set of events, nine were not included in the analysis due to being classified as false positives, and a further two events were removed because data for the change-over period was missing. This left 202 DHB events in the analysis.

Bootstrap clustering identified K=2 as being the most appropriate number of clusters to fit to the data. This was then implemented using the K-means algorithm. Figure 62 shows the mean value of each variable within the two clusters that were generated.

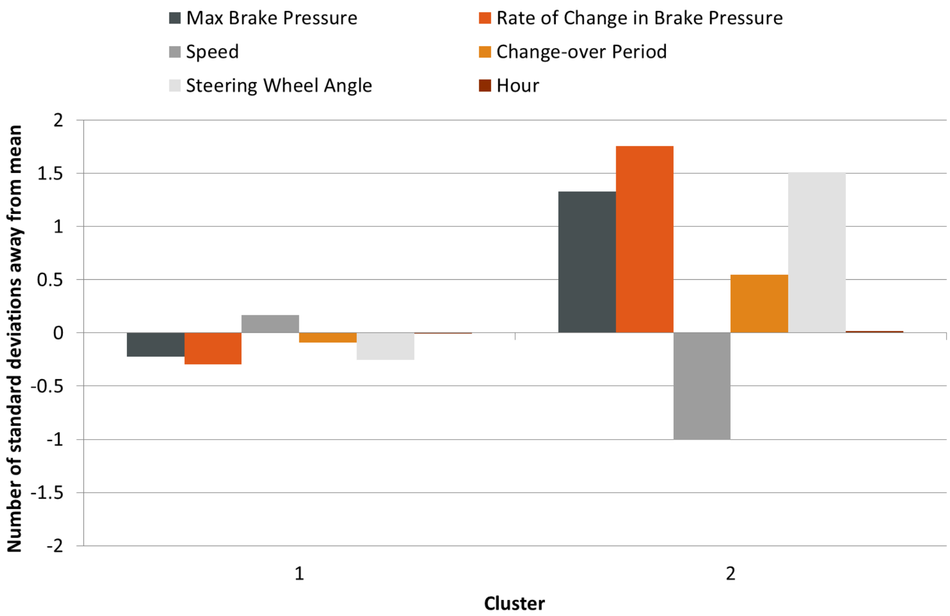

Figure 62 – Plot of mean values of each variable in the two initial clusters

The results show that the time of day, represented by *Hour*, was on average very similar between the two clusters, and so did not contribute very much to the results. Therefore, it was decided to remove this variable from the analysis and repeat the K-means clustering on the data from the remaining five variables. Again, K=2 was chosen as the most appropriate number of clusters, using the bootstrap method. The K-means algorithm was re-run to produce two clusters whose mean values are presented in Figure 63.

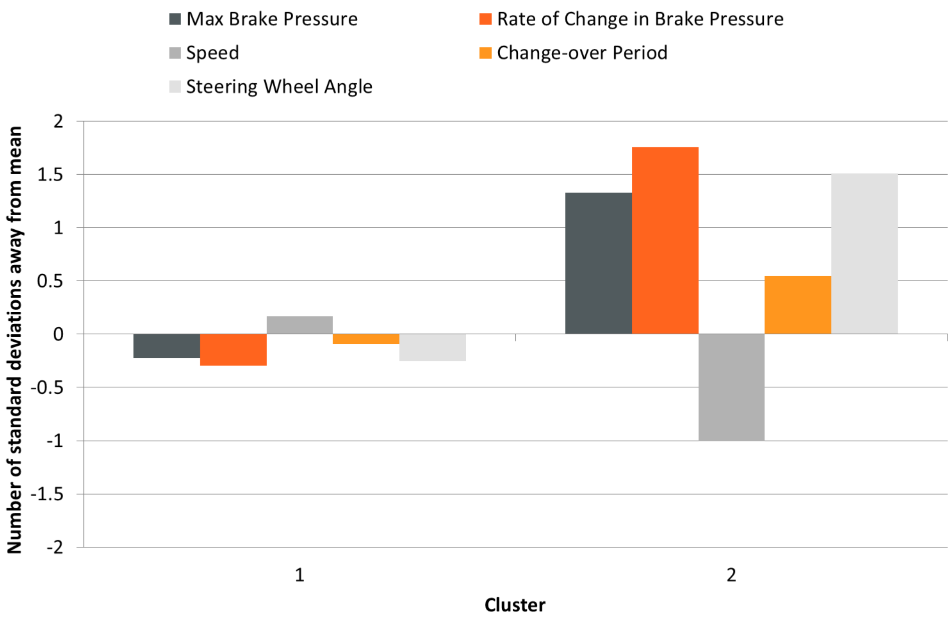

Figure 63 – Plot of mean values of each variable in the two final clusters

The results show that there were differences between events in the two clusters with respect to each variable. Compared to events in Cluster 1, those in Cluster 2 had, on average:

1. A higher maximum brake pressure;
2. A higher rate of change in brake pressure;
3. A lower speed at the start of the braking event;
4. A longer change-over period; and
5. A higher steering wheel angle at the start of the braking event.

Out of the 202 DHB events, 173 (86%) were contained in Cluster 1, and the remaining 29 (14%) were contained in Cluster 2. This suggests that with respect to the key variables, the majority of the DHB events had reasonably similar driver behavior. However, there were some events, generally with a lower speed and a greater level of steering taking place, where the driver had to brake considerably more harshly than normal.