MOVE_UK

Data Analysis Report Phase 2

November 2018

BOSCH TRL JAGUAR LAND ROVER THE FLOOW Direct Line Group ROYAL borough of GREENWICH

TSB Reference No.	102584
Project Title	MOVE_UK - Accelerating automated driving by connected validation and big data analysis
Deliverable No.	D7.4
Principal Author	Simon Morley (Bosch)
Contributing Authors	Kia Hafezi (Bosch), Tobias Mathony (Bosch), Michaela Semmler (Bosch), Jolyon Carroll (TRL), Richard Cuerden (TRL), Ciaran Ellis (TRL), Arun Kalaiyarasan (TRL), Saket Mohan (TRL), Matthias Seidl (TRL), Alex Livadeas (TRL), Ian Barlow (Jaguar Land Rover), Sam Chapman (The Floow), Mark Burke (The Floow), Dan Freedman (Direct Line Group), Ben Morris (Royal Borough of Greenwich), Chris Ndava (ETAS)
Date	13th November 2018
Issue/Revision No.	1.0

Data Analysis Report – MOVE_UK Phase 2 (Deliverable D7.4)

Table of Contents

Executive summary

Introduction

MOVE_UK is contributing to the progression towards automated driving. It enables this by connected systems validation and analysis of 'big data'. Specifically, MOVE_UK is trialling a new method of validating Automated Driving Systems (ADS) using a fleet of five Land Rover vehicles. These cars are due to complete in the region of 100,000 miles on public roads, predominantly in and around Greenwich, London. The project began in August 2016 and will conclude in July 2019.

Update on Phase 1 Use Cases

This report follows on from one published in January 2018, *Data Analysis Report - Phase 1 (D7.3)* , in which the design and initial data analysis results of four project 'use cases' were described: Autonomous Emergency Braking (AEB), Driver Harsh Braking (DHB), Traffic Sign Recognition (TSR) and augmented telematics (Telematics). This report covers Phase 2 of the project. However, updates on the data analysis results related to the Phase 1 use cases are contained within this report. In summary, a much more detailed analysis of this data was carried out with the support of sensor experts and a newly created team of data scientists.

In the case of camera based AEB, a number of highly relevant video sequences were recorded. A large number of DHB events were also recorded and used to identify a distribution of reasons for harsh braking manoeuvres and to investigate human braking behaviour using cluster analysis. Also, the number of traffic sign detections recorded by the TSR system rose from 30,000 in Phase 1 to over 85,000 towards the end of Phase 2. Considerable effort was spent identifying traffic sign clusters from the data collected in sFDE and using these clusters to conduct statistical analysis and an in-depth site investigation.

Phase 2 Methodology and Use Cases

A defining feature of MOVE_UK Phase 2 was the addition of front radar sensors to all trial vehicles, which extended the vehicles' sensor perception and allowed experimentation with methods for validating different sensor modalities. The integration of the radar sensors in the vehicles required both hardware and software adaptions. The project's data recording and analysis tools were also updated and extended to allow easy access to statistics and visualisations of the radar data collected.

The primary focus of this report is set on the four new use cases, which make use of the Phase 2 radar capabilities: Lead Vehicle Statistics (LVS), Radar-based Autonomous Emergency Braking (ARB), Cut-in Scenarios (CIN) and further telematics (Telematics 2).

LVS started from the requirement to understand better how drivers normally follow other vehicles. Conditions of particular interest are the gaps maintained between vehicles, including at what speeds and their duration, as well as what causes the ego driver to stop following the lead vehicle. The intention is that this information can help to understand how to make autonomous driving more comfortable. Emerging clusters in distance-speed distribution can be used to deduce driving behaviour regarding the gaps left between cars and what is common for most drivers. On the other hand, outliers in the point plot can inform about extreme cases that have to be anticipated.

The LVS use case serves as a proof of concept for using existing continuous 1 Hz collected CAN data for statistical evaluation rather unlike the CIN use case, where a trigger was designed and the collection of a new stream of dense CAN data was collected for a limited time at higher frequency (10 Hz).

The Cut-In scenario use case implemented in the MOVE_UK vehicles is used as an example to show the benefits of real-world event data. The use case is used to help understand driver behaviour (of both driver and surrounding drivers) in the particular situation of a Cut-In during a traffic jam on the motorway. A number of behaviours can be associated with a Cut-In situation, including but not limited to, drivers trying to cut across lanes in order to exit a motorway at a junction, changing lanes to avoid a long queue, or to get past a slow-moving vehicle. Understanding driver behaviour in these situations can benefit other Consortium projects by helping them define parameters for safe and human-like automated driving functions.

Due to time and performance considerations, the trigger conditions in the vehicles were kept simple. A first step in refinement is to consider which of the Cut-In sequences collected can be classified as real Cut-Ins and which represent false positives. At the moment, drawing any firm conclusion regarding the parameters for a safe Cut-In, or even to say which event is a true Cut-In or a cut-out, would be premature. Collection of such events and using big data analysis techniques to process, analyse and eliminate any false positives to ascertain true Cut-In events within the dataset is underway and results will be available in Phase 3 of the project.

As with the video AEB use case, the purpose of the Subcritical radar-based AEB use case is to develop, trial and demonstrate the capabilities required to perform silent connected validation for ADAS or ADS systems. In order to use synergies from Phase 1, the video AEB use case was modified appropriately to be applied to radar.

The parameters for the ARB function of the radar sensors are set to a more sensitive level than used in production in order to detect more situations and collect more data. Within the Subcritical ARB use case, 58 events were captured during the course of a nine-month period within Phase 2 of MOVE_UK.

With this ARB use case collecting subcritical ARB situations the following capabilities were successfully demonstrated: identifying real-world *false positive* situations; automated capturing of events which are activated by the system within the vehicle (on board); transmitting high volumes of data over the air; and the process of using real word data to re-simulate in order to optimise parameters that then can be re-flashed to the control unit in the field for the next iteration. In reaching our goal for these capabilities, methodologies have now been generated and an example infrastructure to validate actual and future ADS systems has been developed.

The analysis of the situations showed that the driver reaction time parameter for the trigger was set to a very subcritical level, resulting in many non-dynamic sequences with less value for the analysis. Therefore, a new subcritical parameter set closer to production setting was found through re-simulation of these sequences. In order to still gather a satisfactory amount of data, the parameter set is kept at a level of sub-criticality which will permit some sequences with low dynamics. A further, final adjustment of the parameter settings is planned after a sufficient amount of sequences is gathered. The final setting will be very close to production levels and only near critical situations will be captured. The collection of radar based subcritical ARB sequences will continue during Phase 3 of MOVE_UK.

Telematics data continues to be gathered alongside the core CAN Bus data from the test vehicles in Phase 2. This data gathering is geared to still support a number of key aims:

1. Comparison to existing risk understanding to help support development of improved risk estimation.
2. Comparison to existing event understanding to help support development of improved event data recorder (EDR) functionality that works alongside advanced vehicle technology.

3. To act as an independent data gathering to support statistical analysis for captured vs missed data.

In Phase 2, data gathering is checked to investigate geospatial bias, so as to ensure that meaningful analysis can be performed. This analysis helps to ensure geographically representative data, thereby enabling a meaningful statistical analysis of behaviour by locations. Study of the Greenwich test region (apart from the fleet base of operation which has disproportionately high test vehicle traffic) highlights a smooth and fairly even coverage of test vehicle traffic volume over the borough. The gathering region has been extended to enable capture of data in differing areas, thereby widening it to differing road types, out of the main urban testbed.

The ultimate aim of the telematics use case is to construct a method through which a comprehensive and synoptic description of risk can be computed. To achieve this, the MOVE_UK team is seeking to combine knowledge of telemetry, driver behaviour, the conduct of surrounding vehicles and geographical factors into a single parameterisation. However, to obtain such an understanding will require continued studies to obtain more data using a greater number of vehicles, ideally with the ability to distinguish between specific drivers. However, data collected during Phase 2 can be used to reach a good understanding of vehicles-in-front for incorporating into models of risk. Knowledge of the traffic context in which the vehicle is driving is clearly important, letting us observe not just how the driver is driving, but whether the observed behaviour is normal.

Next Phase of Project

During the final phase of the project, Phase 3, information will be combined from the additional corner Radar systems into the existing analysis; working towards a refined model of risk enabling a 360-degree understanding of proximity. The new Radars will provide positional information beyond that of the headway distance to leading vehicles and will also provide data from all objects observed by the sensors, not just those assigned to *target* status by the internal software, thus giving a more complete view of the entire surroundings of a vehicle in operation. The enriched data available in Phase 3 will grant us an enhanced view of vehicles during autonomous operation and will be particularly valuable for the study of the DHB, ARB and CIN use cases.

1 Introduction

MOVE_UK is a project contributing to the progression towards automated driving through 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, 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 will conclude in July 2019. Phase 2 of the project ran from December 2017 until September 2018.

This report follows on from one published in January 2018, *Data Analysis Report - Phase 1 (D7.3)*[1], in which the design and initial data analysis results of four project 'use cases' were described: Autonomous Emergency Braking (AEB), Driver Harsh Braking (DHB), Traffic Sign Recognition (TSR) and augmented Telematics. Updates on the data analysis results related to these Phase 1 use cases are contained in Section 3 of this report.

A defining feature of MOVE_UK Phase 2 was the addition of a front radar sensor to all trial vehicles, which extended the vehicles' sensor perception and allowed experimenting with validation methods for different sensor modalities. The project's data recording and analysis tools were updated and extended to allow easy access to statistics and visualisations of the radar data collected. These improvements are described in Section 2.

The primary focus of this report is on the four new use cases, which make use of the Phase 2 radar capabilities: Lead Vehicle Statistics (LVS), Radar-based Autonomous Emergency Braking (ARB), Cut-In Scenarios (CIN) and further telematics (Telematics 2). 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 2 use cases and related capabilities are described in more detail in Section 4.

Finally, Section 5 captures our interim conclusions from MOVE_UK Phase 2.

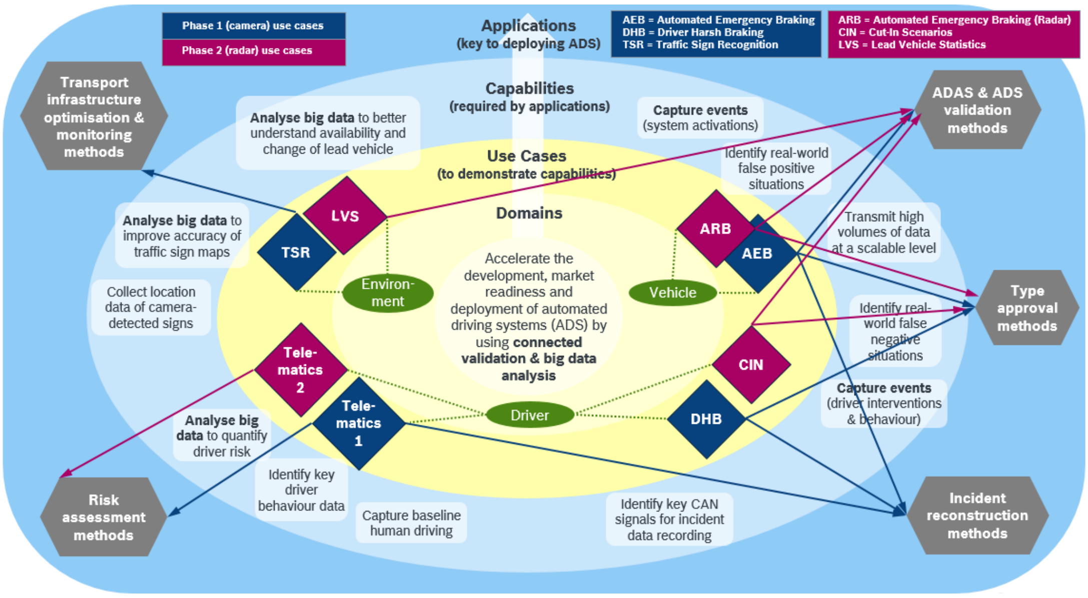

Figure 1 - MOVE_UK objectives and project strategy for Phases 1 and 2

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

2 Data Collection and analysis methodology for Phase 2

2.1 Additional Hardware & Software adaptions for Phase 2

For Phase 2, the Phase 1 setup of the vehicle fleet (five Land Rover Discovery Sport vehicles, model year 2016) was extended in each vehicle by the addition of one front RAdio Detection And Ranging (radar) sensor and the corresponding adaptions to hardware and software.

Radar sensors are able to transmit electromagnetic/radio waves. If these waves hit an object they are reflected back, creating an echo which the sensors are also able to receive. By analysing these echoes, radar sensors are able to precisely detect the range, angle and relative velocity of the encountered objects. Furthermore, radar sensors are largely unaffected by weather and lighting conditions. These advantages, however, come with the drawback that radar sensors have difficulty in determining the precise dimensions of objects.

Radar sensors are available with differing ranges:

- *Long-Range-Radar (LRR), which has a slim radio beam and a large range and is thus ideal for motorways,*
- *Mid-Range-Radar (MRR), which has a wider field of view with lower range and is therefore able to detect neighbouring objects at shorter distances.*

MRR was chosen for use in this project as it is more suitable for urban environments. Table 1 displays the technical data of the chosen radar sensor “MRR Plus” for forward-looking application.

Table 1 - Technical data of MRR Plus used for Phase 2 hardware setup

Technical data	
Frequency range	76...77 GHz
Detection range	0.36…160 m
Field of view (horizontal) **Main antenna** **Elevation antenna**	 ±6° (160 m) ±9° (100 m) ±10° (60 m) ±25° (36 m) ±42° (12 m)
Measuring accuracy **Distance** **Speed** **Angle**	 0.12 m 0.11 m/s ±0.3°
Object separation capability **Distance** **Speed** **Angle**	 0.72 m 0.66 m/s 7°
Cycle time	~60 ms
Modulation	Frequency modulation
Max. number of detected 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

The extension of the hardware setup with front radar sensors has enabled the implementation of Phase 2 use cases, including those of *Subcritical Radar Based Autonomous Emergency Braking* (ARB), *Cut-In* (CIN), and *Lead Vehicle Statistics* (LVS), which will be explained in further sections.

The integration of radar sensors in the vehicles required hardware and software adaptions. A hole was cut in the front bumper of each vehicle in order to enable the radar sensors to work without their fields of view being obscured. Mounting brackets were attached to the vehicle structure behind the front bumpers, into which the radar sensors were fitted – see Figure 2. Furthermore, the hardware setup was extended by the addition of a debug data cable and both public and private CAN connections to each radar sensor, the latter being combined with the private camera CAN. These cables were routed through the vehicle, from the front bumper to the rear luggage compartment.

Figure 2 - Hardware adaptions for integration of front radar

The software adaptions for the integration of the radar sensors involved a number of tasks. First, it was necessary to flash new software onto the sensor. Second, the correct settings had to be enabled on the sensor to ensure the required functionality for the Phase 2 use cases. Third, the sensor was calibrated to ensure its correct alignment. In addition, the Flea3 box configuration was adapted by adding the new CAN signals and the radar data was integrated as a new stream into the ADTF (Automotive Data and Time-Triggered Framework) measurement software, thus enabling the program to trigger radar sequences.

2.2 Field Data Collection

2.2.1 Changes to continuous CAN data collected for Phase 2

Experience gained during the first phase of the MOVE_UK project allowed for reconsideration of which CAN signals to collect in Phase 2. During Phase 1 of the project, some of the signals captured continuously at 1 Hz were found to be mostly empty or did not change and thus were excluded in Phase 2. Furthermore, due to the integration of radar in Phase 2, new signals were added to the dataset in order to support the relevant use cases for Phase 2.

Table 2 provides an overview of the signals which were collected in Phase 2. This signal list resulted from an analysis of the 262 signals recorded in Phase 1, of which 66 were removed for Phase 2, and 76 new signals were added. This resulted in a total of 272 signals being collected in Phase 2.

Table 2 - Overview of deleted and added signals for Phase 2

Signal list for Phase 2	
Signals recorded in Phase 1	262
Not recorded in Phase 1, but requested in Phase 2	4
Recorded in Phase 1, but not requested in Phase 2	-66
Phase 2 signals for radar	63
Additional signals from private CAN camera	6
Odometer signals	3
Total signals for Phase 2	272

The newly added radar signals for the Phase 2 use cases are mainly related to radar target objects. The built-in radar sensor can detect up to 32 objects; in the context of this project, the term target object refers to placeholders for 4 objects with specific positions in relation to the MOVE_UK vehicle (known as the 'ego' vehicle). Figure 3 illustrates a scenario in which the four target objects are assigned to four vehicles which are driving ahead of the ego, in its lane. The resulting signals include, amongst others, the distances to the target objects and their longitudinal relative velocity.

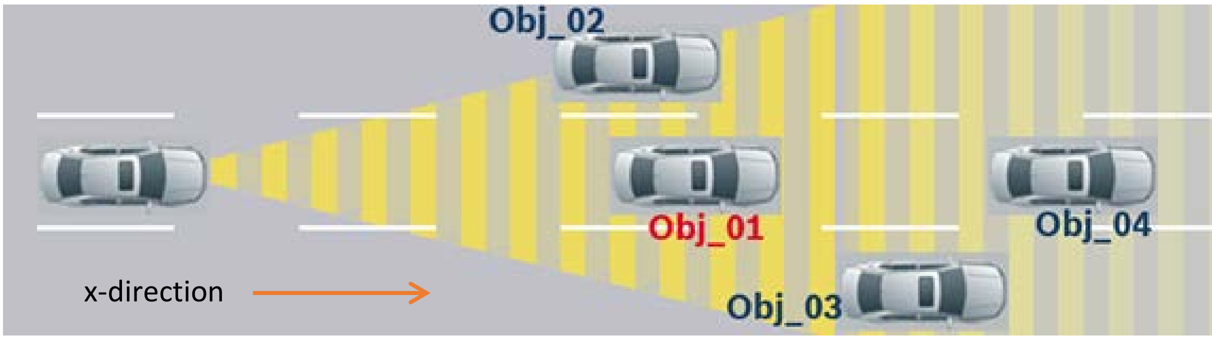

Figure 3 - Four target objects are assigned to four vehicles driving ahead of the ego vehicle

2.2.2 Event-based radar CAN data

For Cut-In events, which were of particular interest for further analysis in Phase 2, certain CAN data was captured at a higher frequency than the continuous CAN data (captured at 1 Hz).

For these events, the Consortium selected a set of, mainly radar based, CAN signals (approximately 30) to be recorded at high frequency (10 Hz) for a duration of 30 s; 15 s before and 15 s after the event.

2.2.3 Event-based video sequences with high-bandwidth video and radar data

For certain Phase 2 use cases, namely ARB, trigger conditions were activated to collect high-bandwidth video and radar data. These sequence files include the video images of the stereo video camera, radar data and high-resolution CAN data. When an event is triggered, the measurement system starts recording a 20 s sequence (15 s before and 5 s after the event trigger). Video, radar and CAN data is collected at the highest resolution available in order to allow re-simulation of the sequence with different algorithms or parameter sets of the ADAS.

This data is then converted to a readable format, stored on the MOVE_UK Cloud server and made accessible to the Consortium. The Consortium agreed to a video image resolution of 640x320 pixels and a cycle time of 66 ms for all 272 CAN signals pre-selected for continuous collection, (i.e. 15-times higher frequency than continuous CAN data). The capacity of the hard drive in the measurement system 'Movi-PC' is sufficient to store a large number of events per day, but the Wi-Fi connection and vehicle battery capacity limit uploads to a maximum of five sequences per vehicle per day. The upload of a single event sequence takes approximately 1 hour.

The use of the above described technique in MOVE_UK Phase 2, allowed for the sequence collection of potentially critical situations detected by the radar-based autonomous emergency braking function (subcritical ARB). This was in addition to the sequence collection of the MOVE_UK Phase 1 use cases AEB (subcritical video-based autonomous emergency braking function), and DHB (driver harsh braking).

2.2.4 Telematics data collected in Phase 2

Telematics data continues to be gathered alongside the core CAN Bus data from the test vehicles in Phase 2. This data gathering is geared still to support a number of key aims:

1. Comparison to existing risk understanding to help support development of improved risk estimation.
2. Comparison to existing event understanding to help support development of improved event data recorder (EDR) functionality that works alongside advanced vehicle technology.
3. To act as an independent data gathering to support statistical analysis for captured vs missed data.

In Phase 2, data gathering is checked to investigate geospatial bias, so as to ensure meaningful analysis can be performed. This analysis helps to ensure geographically representative data, thereby enabling a meaningful statistical analysis of behaviour by locations. Study of the Greenwich test region (apart from the fleet base of operation which has disproportionately high test vehicle traffic) highlights a smooth and fairly even coverage of test vehicle traffic volume over the borough (see Figure 4). In Figure 5 the gathering region has been extended to enable capture of data in differing areas, thereby widening it to differing road types, out of the main urban testbed.

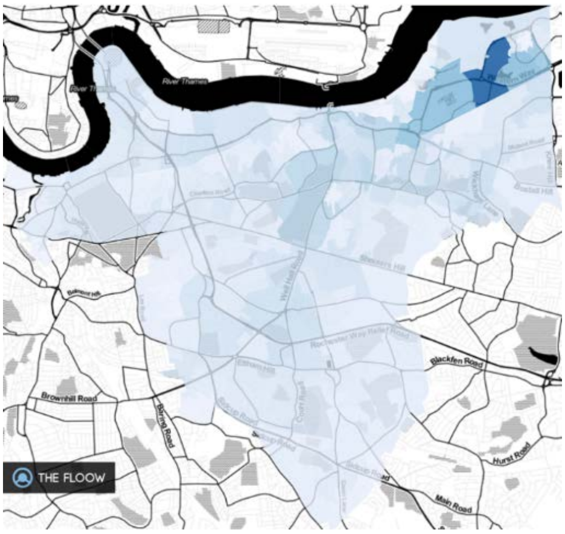

Figure 4 - The regions in which data is concentrated in the Greenwich area gathered from the passage of vehicles (areas shaded indicate by LSO region the amount of vehicle traffic in each region

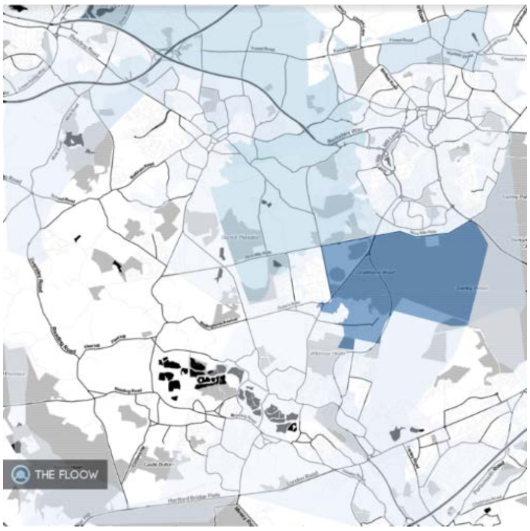

Figure 5 - The regions in which data is concentrated in the TRL test location (areas shaded indicate by LSO region the amount of vehicle traffic in each region)

2.3 Data storage and analysis tools

2.3.1 Improvements to data visualisation and analysis tools for Phase 2

Since publication of the Phase 1 report (D7.3), 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 radar-based use cases.

2.3.2 Systematic Field Data Exploration (sFDE)

As described in the Phase 1 report (D7.3), 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) is 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 Phase 1 to display the use case data of Traffic Sign Recognition (TSR), Autonomous Emergency Braking (AEB) and Driver Harsh Braking (DHB). Figure 6 displays the Overview Dashboard of sFDE, which visualises the history of driven tours, total input files, input files for the current day and the latest vehicle positions. It also serves as a main page to navigate between the different views.

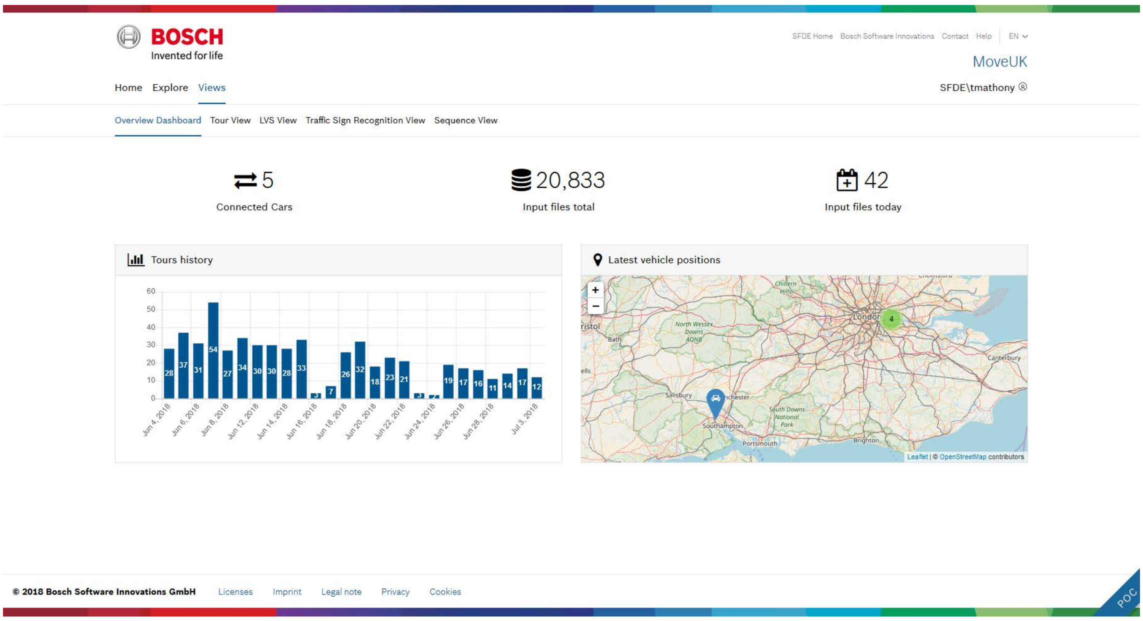

Figure 6 - Dashboard with overview for MOVE_UK

During the course of Phase 2, sFDE was extended with three new views to represent the data from the Phase 2 use cases of Lead Vehicle Statistics (LVS) and Cut-In scenarios (CIN). Furthermore, 'Sequence View' (described in section 4.3.3 of the D7.3 phase 1 report) has been updated to include radar data and associated views (see Figure 7). Besides these new views, the existing Tour View was enhanced with a mileage overview showing driven miles e.g. on different road types.

The Sequence View now visualises data for recorded events of the use cases AEB, DHB and ARB. It contains a view for the video sequence of the event, as well as a map and a visualisation of selected signals of the recorded CAN data. Furthermore, for the ARB use case, the Sequence View features a 'Bird Eye' View to display the objects detected by the radar.

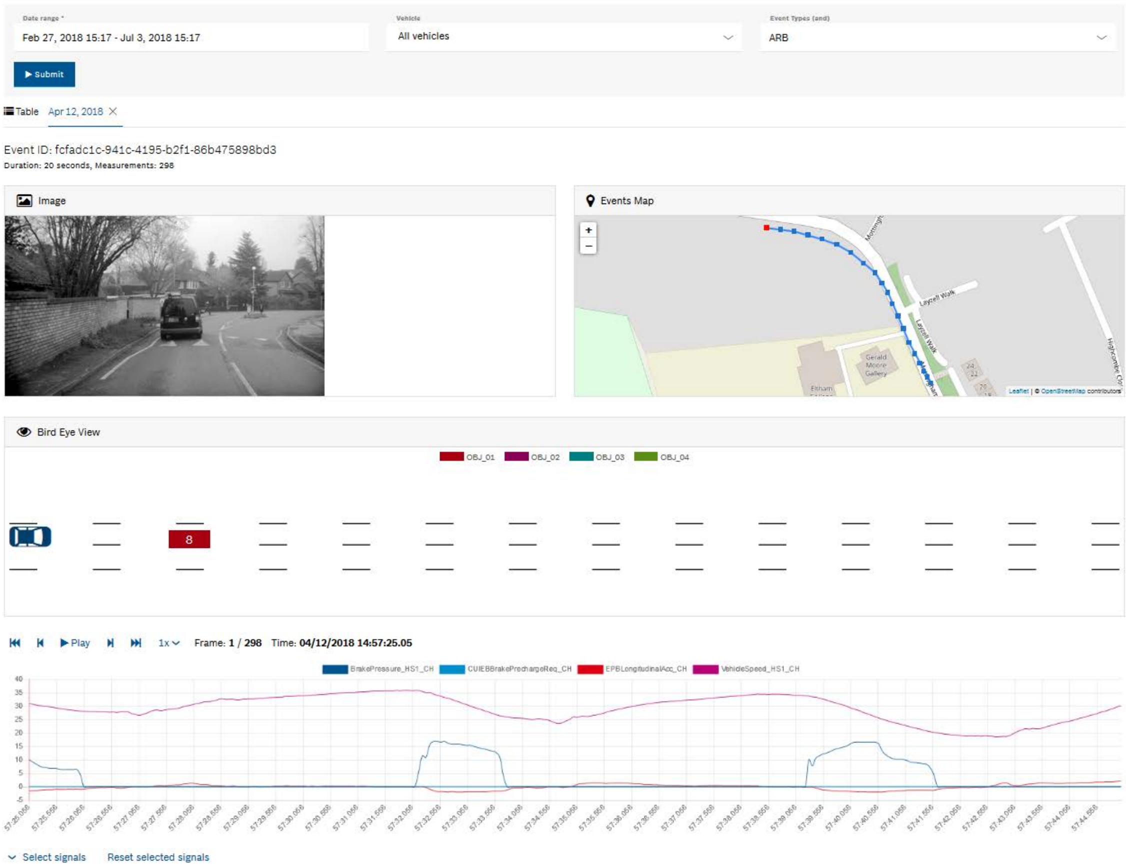

Figure 7 - Sequence View showing a sequence

The newly added LVS View features a detailed visualisation of statistics gathered within the LVS use case of Phase 2. Besides the distance versus speed distribution shown in Figure 8, the LVS View also contains visual representations of the snippet length and exit reasons (Figure 9). A snippet is a time period of a journey (tour) when the chosen pre-conditions are met (e.g. speed range, road type). Exit reasons comprise the reasons why a snippet ended.

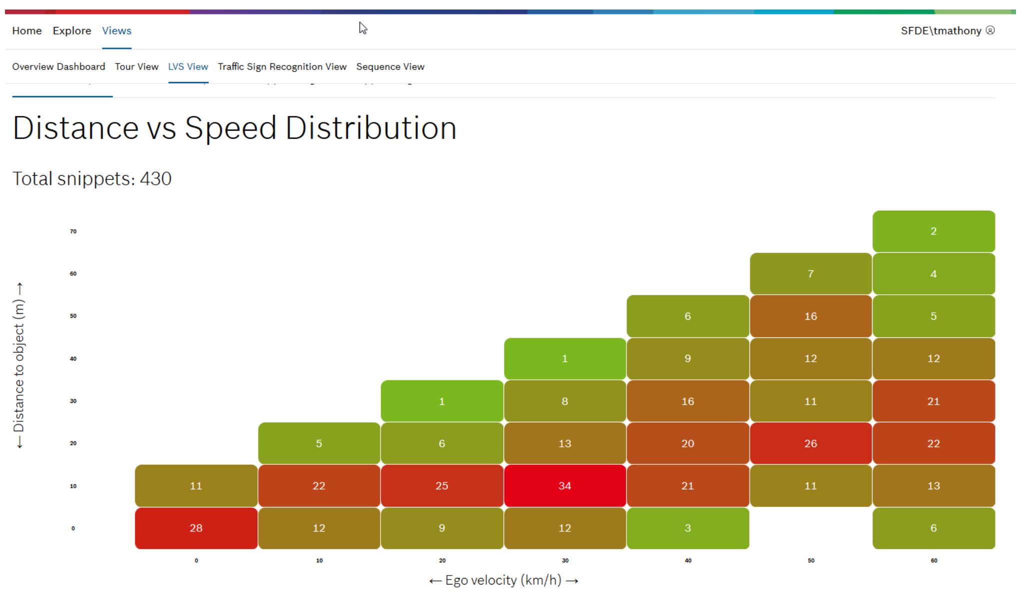

Figure 8 - Visual representation of the distance versus speed distribution in the LVS View

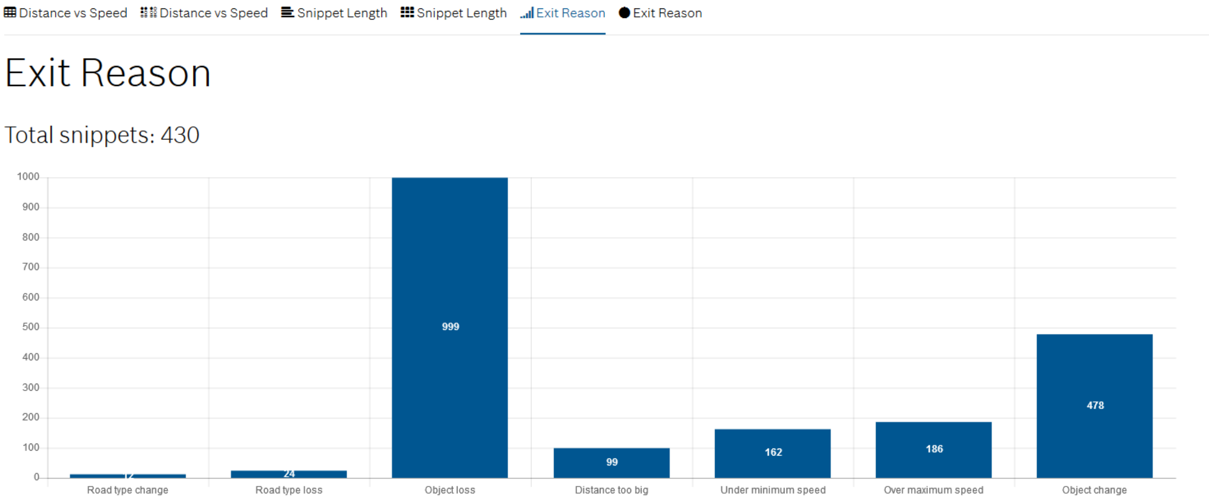

Figure 9 - Visual representation of the exit reasons in the LVS View

2.3.3 Enterprise Automotive Data Management (EADM)

EADM is a solution for management and analytics of automotive measurement data and metadata such as description of vehicles and test systems, developed by ETAS (a Bosch subsidiary company).

Within MOVE_UK, MATLAB has been used successfully by ETAS to interface with EADM and provide basic solutions for the TSR, AEB and DHB use cases during Phase 1 of the project. For example, MATLAB was used alongside EADM to determine the DHB use case trigger. This trigger is currently implemented on all MOVE_UK vehicles. In Phase 2, MATLAB was used to define the Cut-In use case trigger and, both MATLAB and Python were used to determine solutions for the LVS use case implementation. The Floow, ETAS and Bosch continue to use the EADM interface to carry out further and more advanced data analytics using R Language, MATLAB and Python, respectively.

No notable improvements have been made to EADM since the release of the phase 1 report (D7.3). However, work on the development of a new version of EADM is underway.

3 Update on Phase 1 Use cases

3.1 Overview

As previously described, use cases have been defined within the project which demonstrate different capabilities and relate to the various different applications (which are 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.

Four use cases have been developed and implemented previously, in MOVE_UK Phase 1:

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

For more detail on the purpose and design of these use cases, please refer to the *Data Analysis Report - Phase 1 (D7.3)*[2].

This section provides updates on the number of events, sequences and amount of continuous CAN data recorded during Phase 2. In addition, for AEB the nature of the two most-relevant sequences captured is discussed and the concept of sequence re-simulation introduced. For DHB, the analysis presented identifies the distribution of reasons for harsh braking manoeuvres and investigates human braking behaviour using cluster analysis. For TSR, an analysis of data collected around a number of hotspots is presented to assess the likelihood of traffic sign detection in relation to other continuous data variables collected. For Telematics 1, this report contains an update on the event data recorder (EDR) analysis and risk method.

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

3.2.1 AEB events in Phase 2

In Phase 2 of MOVE_UK, 10 events were captured within the Subcritical AEB use case during the course of a nine-month period. The following still frames (Figure 10 and Figure 11), which were captured during Phase 2 of the MOVE_UK 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 (i.e. are not added during post-processing).

<u>Example 1 (Subcritical AEB Sequence 32): Braking car in front due to pedestrian crossing the road</u>

A MOVE_UK vehicle is being driven behind another vehicle whilst approaching a pedestrian crossing, when suddenly the leading vehicle harshly brakes for a pedestrian who has turned abruptly left onto the pedestrian crossing (Figure 10). Due to the leading vehicle braking harshly, the driver of the MOVE_UK vehicle has to brake harshly also. During this sudden braking manoeuvre, when the time-to-collision fell, the stereo video camera detected and correctly classified the leading vehicle and activated AEB (in silent mode). The back of the vehicle in front was the object causing the trigger.

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

Figure 10 - Leading vehicle triggering Subcritical AEB Sequence 32 in March 2018. Still frame of the trigger extracted from the stereo video recording with boxes indicating camera-detected objects. Red arrow highlights the object causing the trigger.

<u>Example 2 (Subcritical AEB Sequence 33): surprisingly braking car in front</u>

A MOVE_UK vehicle is being driven behind a small truck whilst approaching a red traffic light. Just as both vehicles come to a halt, the traffic light turns green and the truck starts to move, as does the MOVE_UK vehicle. Immediately after starting up, the truck brakes unexpectedly and very harshly (Figure 11). Due to this, the driver of the MOVE_UK vehicle also brakes harshly. During this sudden braking manoeuvre, when the time-to-collision fell, the stereo video camera detected and correctly classified the truck and activated AEB (in silent mode). The back of the vehicle in front was the object causing the trigger. Even with the current production calibration, the AEB system would have triggered in this situation.

Figure 11 - Leading vehicle triggering Subcritical AEB Sequence 33 in April 2018. Still frame of the trigger extracted from the stereo video recording with boxes indicating camera-detected objects. Red arrow highlights the object causing the trigger.

3.2.2 Re-simulation of captured events

Another aspect covered in the data analysis of AEB sequences in MOVE_UK Phase 2 is the re-simulation of the situations captured using a different version of camera software (namely, production software). As the behaviour of camera software versions can vary, the re-simulation allows evaluation of camera software performance. The collected sequences are replayed in a simulation environment, including the camera with updated software. This produces new result files indicating whether the AEB was activated or not activated. The re-simulation of all collected sequences during the MOVE_UK Phase 2 trials showed that a series production AEB system would have triggered in a number of situations captured (see Example 2).

3.3 Use case Driver Harsh Braking (DHB)

3.3.1 DHB events captured in Phase 2

In Phase 2 of MOVE_UK, 55 DHB events were captured, which brings the total number of events up until the end of June 2018 to 89. In addition, three 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).

The monthly breakdown provided in Table 3 shows that significantly fewer events were recorded between November 2017 and March 2018, compared to the surrounding months. An inadvertent software misconfiguration in the Flea3 box was identified in March 2018 as the reason for this drop. The CAN monitoring cycle time for DHB trigger criteria (brake pressure and longitudinal acceleration) had been increased automatically to 500 ms (from its previous level of 100 ms) when changing a radar-related setting to this value, with the consequence that events where the trigger criteria lasted for less than 500 ms were likely to be missed. After the configuration was fixed on 9th April 2018, the number of events being recorded increased back to previously seen levels.

Table 3 - Number of DHB events captures 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	2	–	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	–	–
Total number	92			

For a considerable amount of time after the first month with fewer events it could not conclusively be decided whether the lower number of events recorded was due to a technical issue in capturing them, or due to natural variation in the number of events that occurred in reality. Considerable effort was then involved in identifying that an incorrect software configuration was causing the issue. This highlights the complexity of applying a system without a ground truth measurement against which to compare in type approval situations. Recording no events is not necessarily an indicator of there being no events, it could result from a failure in the recording technology. Thorough definition and oversight of the experimental setup and review of technical records will therefore be required to ensure consistency of approvals.

3.3.2 Reasons for drivers performing harsh braking manoeuvres

A review of the captured DHB videos was undertaken to determine the likely reasons for drivers braking harshly in each of the events. A classification was developed with the intention to capture the difference between events where the driver avoids an impending impact (for instance with a vehicle ahead which could be detected by the AEB camera) and a situation where the driver brakes to obey a give way instruction (e.g. come to a halt before a junction or traffic lights, mostly with vehicles approaching from directions that would not have been in view of the AEB camera). The system evolved further during the review of sequences and currently contains 16 categories, as shown in Table 4, which allowed classification of all events recorded so far. This breakdown of recorded events by reason for braking is visualised in Figure 12.

Table 4 - 92 DHB events split by reason for the harsh braking manoeuvre. Note that this breakdown includes the three AEB/ARB events, which also fulfilled DHB criteria

Reason for harsh braking manoeuvre	Number of events
Avoid imminent impact, with	
pedestrian	2
vehicle (incl. bicycle)	31
other	0
Give way, at	
junction (excl. roundabout)	5
pedestrian crossing	10
roundabout	4
slip road	1
traffic lights	6
other	0
Leave gap for oncoming traffic (narrow road section)	11
Low-speed manoeuvring (e.g. parking, three point turn)	5
Reduce speed, in preparation for	
upcoming turn	3
speed bump	7
other	0
Test sequence	2
Unknown	5

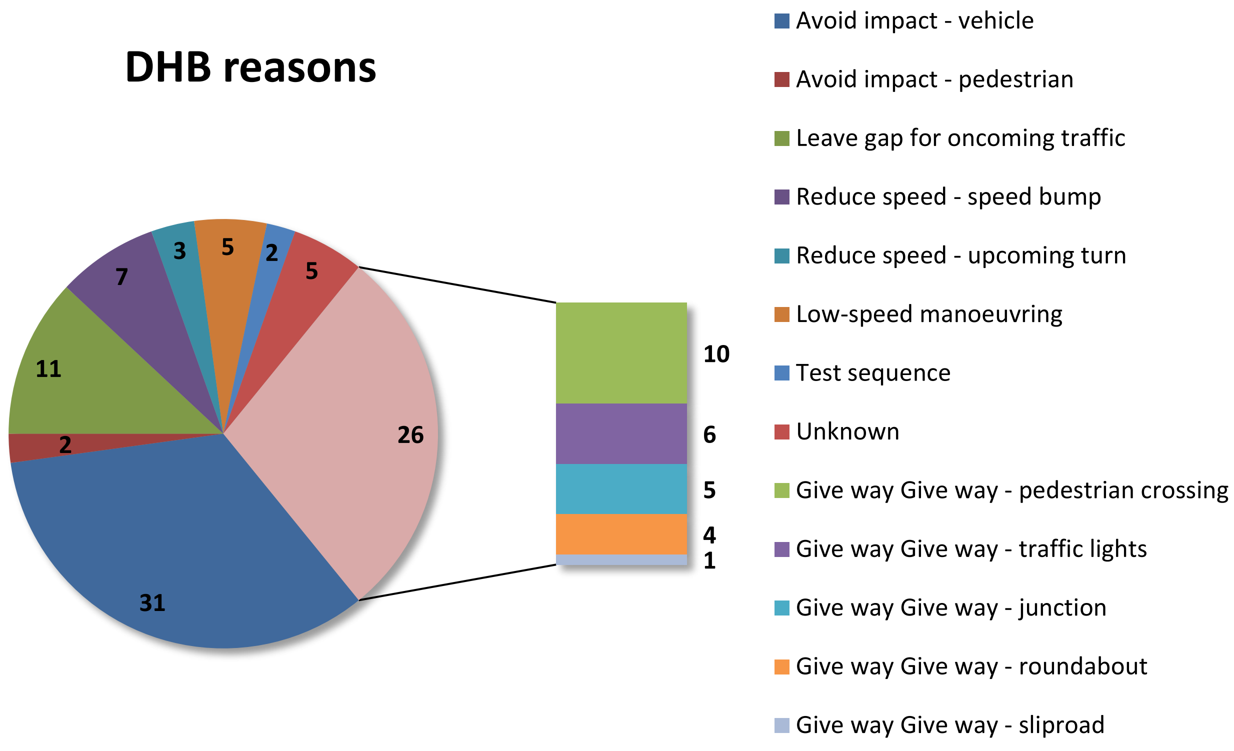

Figure 12 - Breakdown of 92 DHB events by reason of braking

It can be seen that the most frequent reason for a driver performing a harsh braking manoeuvre is to avoid an imminent impact with another vehicle. The review of sequences suggests that the majority of these events was probably not critical enough to warrant an AEB intervention, but had the driver not braked strongly an impact would have occurred. Table 5 shows that 23 of the 31 events had passenger cars as potential impact targets. Of the 31 events, 28 (i.e. 90%) would hypothetically have resulted in a frontal impact of the trial vehicle and therefore fall in the desired functional scope of an AEB system. Of these, 21 would have resulted in a front-to-rear shunt with another motor vehicle, which is the principal functional scope of most current vehicle-to-vehicle AEB systems.

Table 5 - Breakdown of the 31 DHB events classified as 'avoid impact with vehicle' by vehicle type of potential target and hypothetical impact configuration

Vehicle type about to be impacted	Likely impact configuration			
	Front-to-front	Front-to-rear	Front-to-side	Sideswipe
Bus	0	1	0	1
Car	0	18	4	1
Van	1	2	1	1
Bicycle	0	0	1	0

The subcritical AEB system activated in three of these 21 cases (two times Level 3 activation, i.e. brake activation; one Level 1 activation, i.e. brake pre-charge). From a preliminary review of the recordings no false negative AEB cases have been identified. The 18 cases where AEB did not activate were likely not critical enough. However, a more detailed investigation into whether there were any false negative cases at detection level (i.e. cases where a relevant vehicle target had not been identified at the time of harsh braking) will be carried out in the next phase of the project.

The second most frequent reason for harsh braking manoeuvres to be triggered is give-way situations where the driver has to stop to respect another road users right of way. The harsh braking manoeuvres most often occurred at pedestrian crossings and traffic lights. From reviewing the video recordings it can be speculated that these situations are often caused by late detection of the need to stop. Figure 13 shows an example of a visually obstructed pedestrian crossing in a right-hand bend. Two DHB sequences were recorded at this specific location, which might be indicative of hazardous placement of the crossing.

DHB event 25 – October 2017 | DHB event 40 – March 2018

Figure 13 - DHB sequences 25 and 40. Still frames extracted seconds before the harsh braking manoeuvre to demonstrate the obstructed driver view of the upcoming pedestrian crossing (arrow).

The third most prominent reason for DHB is to leave a gap for an oncoming vehicle through a narrow stretch of road – a situation that is often characterised by late detection due to poor oversight of the situation cause by parked vehicles. However, none of the situations observed appeared safety critical.

3.3.3 Event cluster analysis

3.3.3.1 Aims

The aim of this analysis was to investigate whether CAN signal data associated with harsh braking events could tell us anything about human braking behaviour. For example, are different styles of braking used in different driving situations? This information could lead to algorithms which could anticipate the type of braking expected from other cars in the same environment, or to predict behaviour in other environments.

3.3.3.2 Variables of interest

The event cluster analysis was carried out in early 2018 and was therefore performed using the 35 DHB events captured up until January 2018 (the two test sequences recorded in this period were removed for the analysis). From over 200 available CAN signals, 14 were included in the analysis. These variables included Event Variables which labelled the event and progress through it and Driving Variables (taken from the CAN signals) which allowed an evaluation of an individual's driving actions during the event (Table 6).

Table 6 - Initial list of CAN variables used for analysis with CAN signal names in brackets

Event Variables	Driving Variables
EventID	Speed (VehicleSpeed_HS1_CH)
Frame Counter	Steering Wheel Angle (SteeringWheelAngle_CH)
Time	Lateral Acceleration (LateralAcceleration_HS1_CH)
Trigger	Longitudinal Acceleration (EPBLongitudinalAcc_CH)
	Yaw Rate (YawRate_HS1_CH)
	Brake Pressure (BrakePressure_HS1_CH)
	Accelerator Pedal Position (PedalPos_CH)
	Distance to Object in Front (S6D_Obj_01_RadialDist)

The driving variables were chosen as the most likely to be of interest during a braking event. Preliminary analysis showed that these variables were often non-zero during a braking event and showed variation during the event.

Where multiple variables could be seen as representing the same phenomenon, the variable with the highest level of variation in the sample was used in the next step of analysis. For example, yaw rate, lateral acceleration, and steering angle are all expected to correlate with the turning of the vehicle. Of these variables, steering angle had the highest level of variation compared to yaw rate and lateral acceleration respectively and so only steering angle was used (standard deviation = 63.8 v 6.36 and 0.79 respectively). Similarly, accelerator pedal position and longitudinal acceleration are both a measure of forward acceleration. In this case, accelerator pedal position was found to be more variable (standard deviation = 18.6 v 1.80) and so this was included in the following analysis while longitudinal acceleration was excluded.

3.3.3.3 Data processing

Isolating the braking event from the data

Before analysing the data further, it was processed so only that from the period during which the driver was braking was included. This is important as for each DHB event, data is recorded and stored for 15 s before the trigger is reached and for 5 s after it. Without first isolating the data associated with the braking event itself, any noise from the surrounding data will interfere with the analysis.

The braking event was isolated by finding the point of maximum brake pressure, and working backwards in time until the brake pressure was zero (the point at which braking began), and forward in time until the brake pressure returns to zero (the point at which braking stops). In some cases the brake pressure did not return to zero before the end of the event recording (Figure 14, Events 5, 6 and 7).

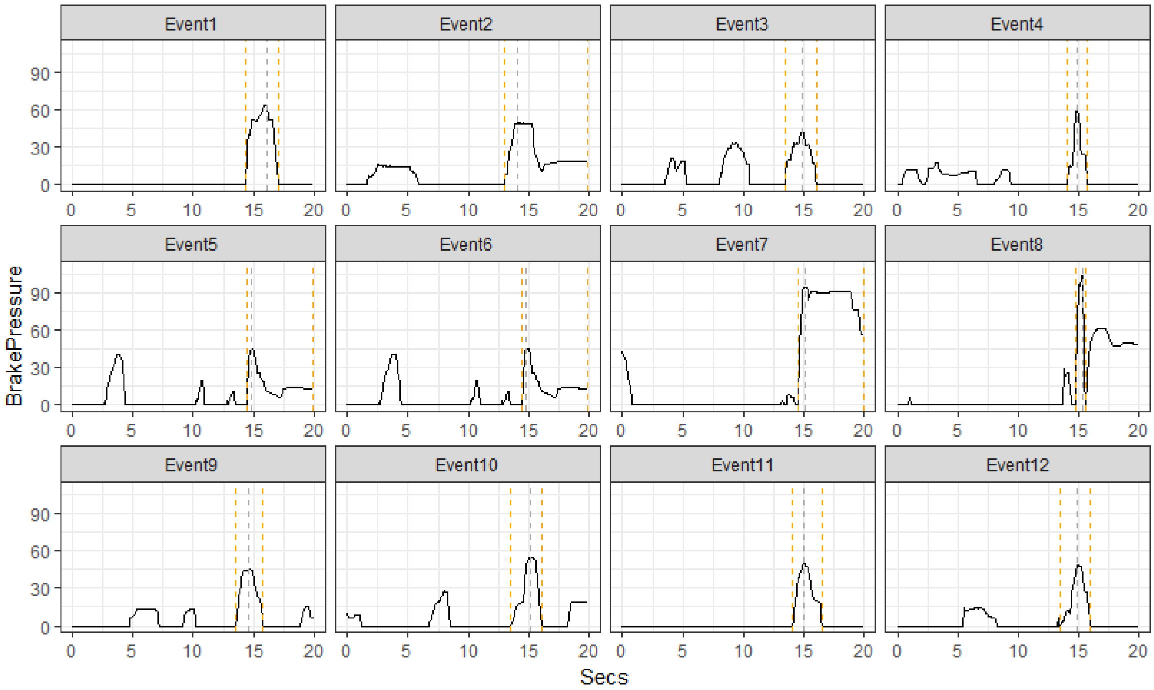

Figure 14 - Brake pressure against time for events 1 to 12 (subset of the 35 analysed events shown for illustration purposes). The beginning and end of the braking event is shown (dashed orange line), maximum brake pressure also shown (dashed grey line).

<u>Choosing the final variable set</u>

After the braking event had been isolated, a second series of variables was generated containing only data specific to the braking event (Table 7).

Table 7 - Variables generated from the isolated braking event

Variable	Group	Description
Maximum Pressure	Brake Pressure	Maximum recorded pressure during the braking event
Brake Time to Max	Brake Pressure	Time from start of braking to maximum pressure
Brake Total Time	Brake Pressure	Time from start of braking to end of braking (or end of recording, whichever is earlier)
Brake Pressure per Unit Time	Brake Pressure	$\frac{\text{Maximum Pressure}}{\text{Brake Time to Max}}$
Speed at Brake Start	Speed	Speed of car at start of braking
Change-over Period	Accelerator	Time between releasing the accelerator pedal and braking
Steer Angle at Brake Start	Steering	Steering wheel angle at start of braking
Distance to Object at Brake Start	Distance to Object	Distance to object in front of car at start of braking

As shown in Table 7, four of the braking event variables were related to magnitude or speed of braking (those in group 'Brake Pressure'). Some of these variables will contain the same information so inclusion of all of these variables in the analysis will not be necessary. Brake Total Time was excluded as for 24 out of 35 events these values were truncated (i.e. the data recording ended before the brake pressure reduced back to zero). In a larger sample this variable may be of interest and so it is still included in the list of variables to consider.

To ensure that the most parsimonious (i.e. statistically efficient) set of the remaining variables were included, a correlation plot was drawn to check for correlated variables and therefore those variables that share the same information (Figure 15).

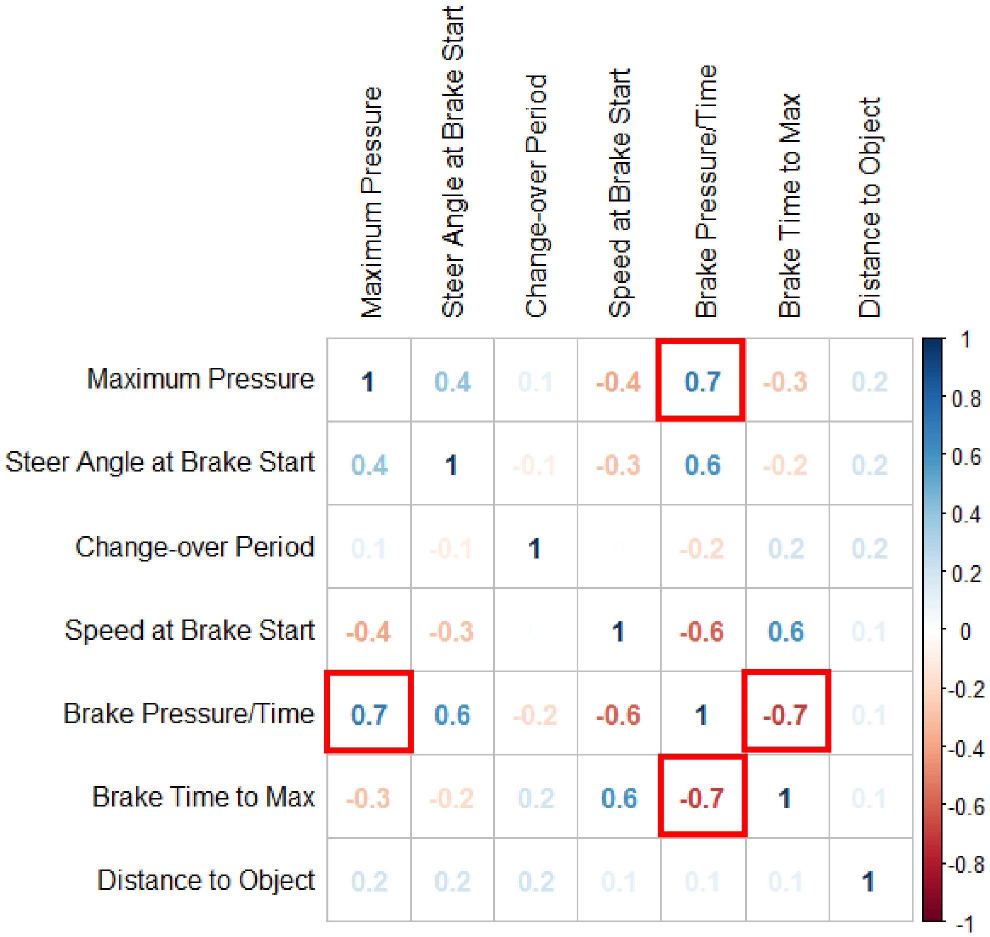

Figure 15 - Plot of correlations between variables

High (>0.7) correlations were found between both Maximum Pressure and Brake Time to Max, and Brake Pressure per Unit Time. Therefore it was only appropriate to use one of each of these pairs of variables in the final variable list. The variables were chosen after clusters were generated (Section 3.3.3.4).

3.3.3.4 Generating Clusters

Clusters (groups of similar braking events) were generated using K-means clustering in the programming language R. Elbow plots were used to determine the optimum number of clusters to generate by calculating the average distance between observations in a cluster and finding the minimum cluster number beyond which improvements in distances were considered marginal; in this case, 4. K-means clustering is a method for grouping data points based on similarity to other data points within a group; the algorithm works iteratively to find the centres (of each cluster) which minimise the distance of each data point from the cluster centre. This produces clusters of data which are appreciably different from other clusters, and similar within the clusters.

After four data clusters were generated the amount of variation in the clusters explained by each variable was determined via ANOVA (Analysis of Variance – a statistical technique to evaluate the causes of variation between groups). The results of the ANOVA were used to determine which of Brake Pressure per Unit Time, and Brake Time to Max, and which of Brake Pressure per Unit Time and Maximum Pressure to include in the analysis. Brake Pressure per Unit Time explained more variation between the clusters than either of the other variables and so the latter were excluded. This left a final variable set (in order of variation explained): Speed at Brake Start, Change-over Period, Brake Pressure per Unit Time, Distance to Object at Brake Start, and Steer Angle at Brake Start. Clusters were recalculated and the mean of each variable for each cluster was calculated and scaled. By scaling the values (the variables were normalised to an overall mean of 0 and a standard deviation of 1), the variables can be compared more easily and so the scale of the cluster deviation from the overall average can be seen (Figure 16). For example, cluster 2 contains braking events with a much higher change-over period (from accelerator to brake) than average.

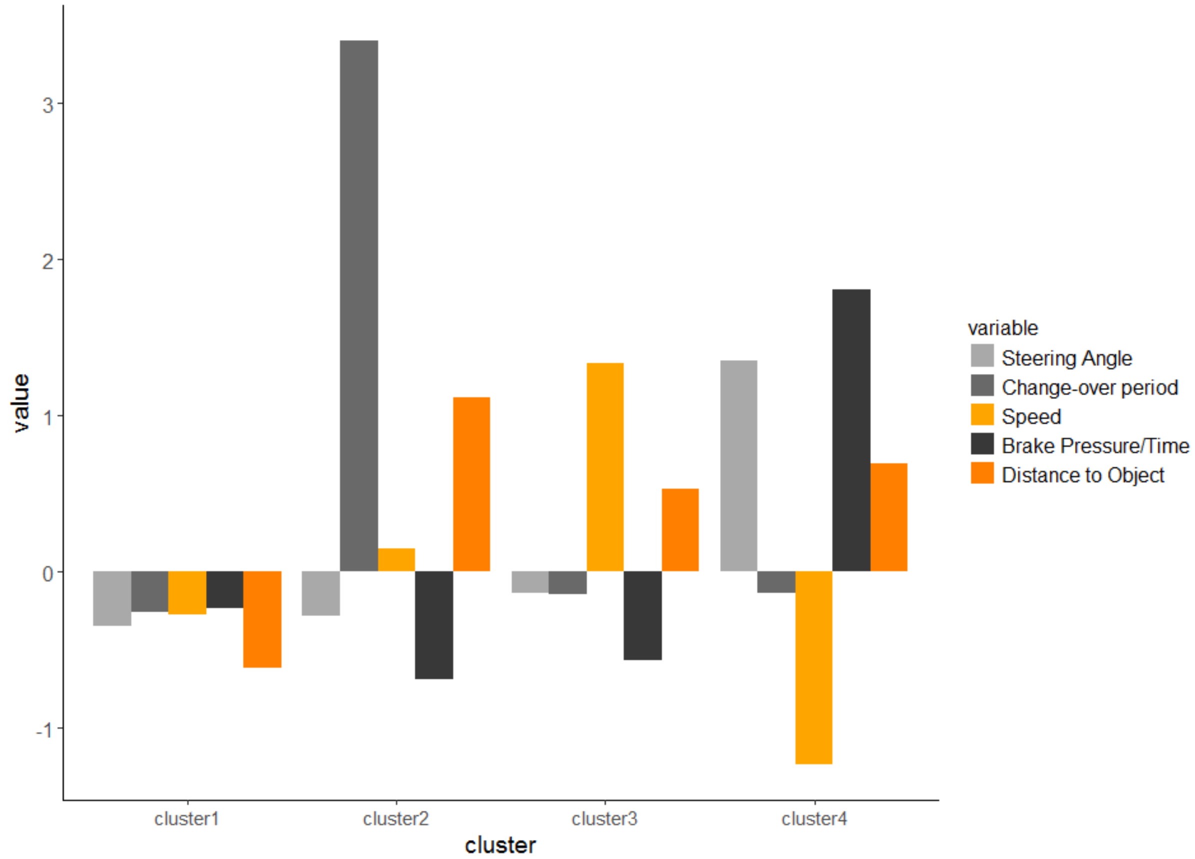

Figure 16 -Plot of mean values of braking variables for each cluster

3.3.3.5 Results: Mapping Braking Types to Clusters

While five variables were found to be important in differentiating the four clusters (Figure 17a), clusters can be distinguished when fewer dimensions are plotted. Figure 18 shows that the clusters can be broadly distinguished into three groups depending on the Speed at Brake Start and Brake Pressure per Unit Time during braking.

Cluster 4, includes events where Speed at Brake Start was very low and Brake Pressure per Unit Time was high; cluster one includes events where Speed at Brake Start was high and Brake Pressure per Unit Time was low, and clusters 2 and 3 inhabit intermediate positions. Cluster 2 only includes two events, but these can be distinguished from the others by the inclusion of Change Period, which is relatively high for events classified as cluster 2 (Figure 17b).

Graphing these in 3D is an option (Figure 18) but given the small number of events represented in cluster 2, the 2D representation including Speed at Brake Start and Brake Pressure per Unit Time is sufficient to describe the majority of variation between the clusters.

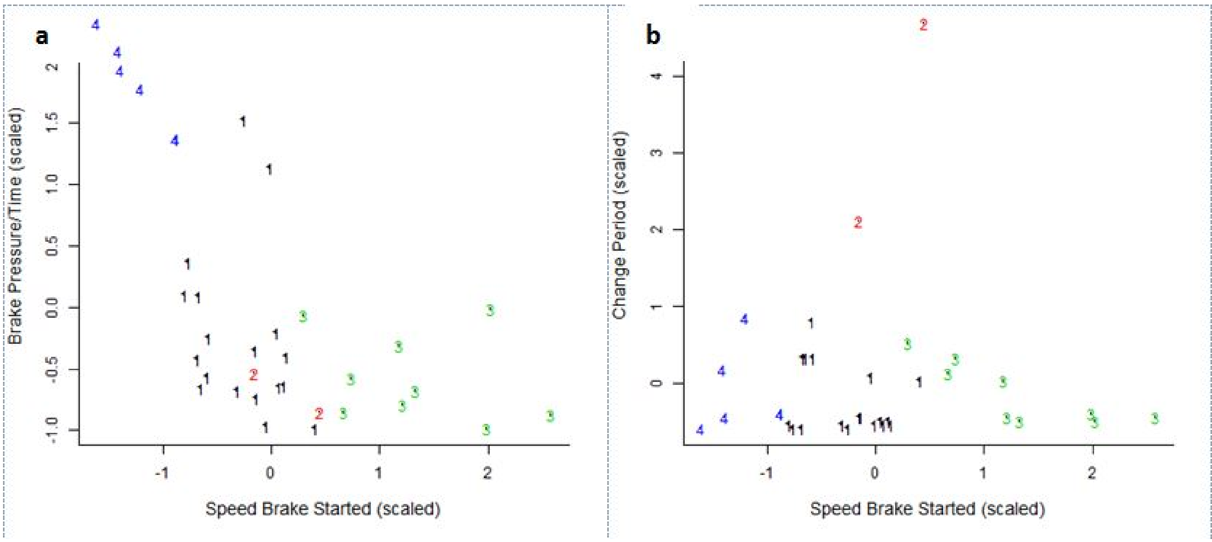

Figure 17 - Cluster ID (1–4) plotted against Speed at Brake Start and a) Brake Pressure per Unit Time, and b) Change Over Period

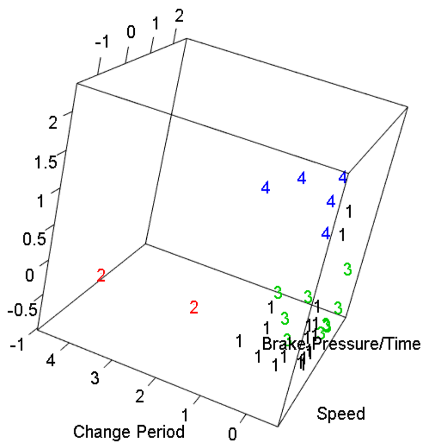

Figure 18 - Cluster ID (1–4) plotted against Speed at Brake Start, Brake Pressure per Unit Time and Change Over Period

When reasons for braking are mapped onto the graph, there is no clear pattern to suggest that specific types of braking occur when different situations are encountered (Figure 19). Braking for traffic lights is found in each cluster, though looks to be slightly more prominent in cluster 3 (high speed low application rate), whereas the only incidents of braking for slip roads occur at slow speeds and with relatively high application rate (cluster 4). Most braking to avoid a vehicle occurred at intermediate levels of Brake Pressure per Unit Time and Speed at Brake Start.

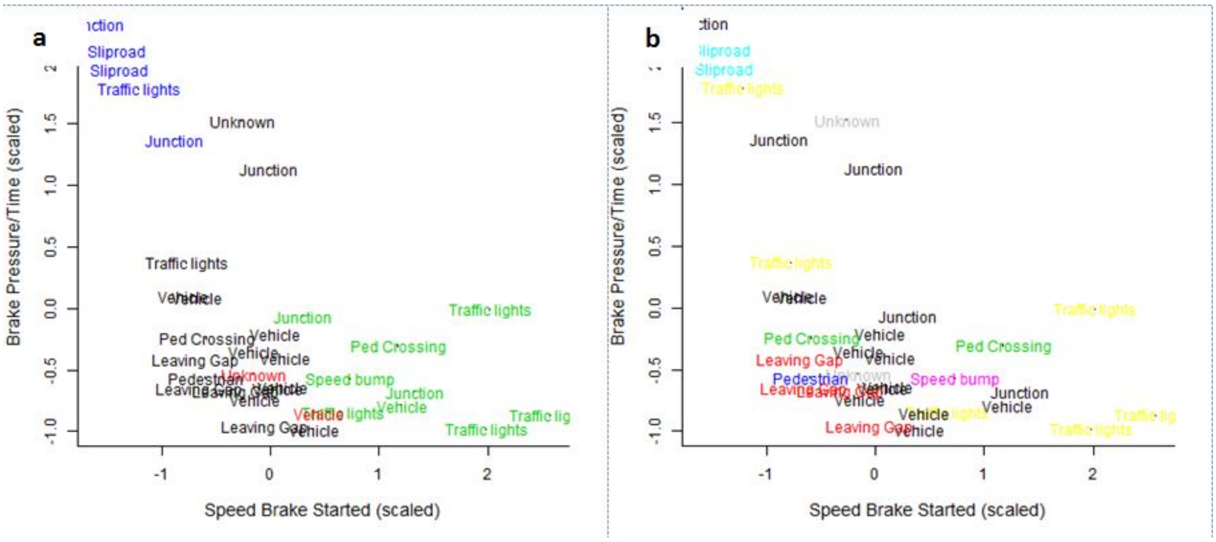

Figure 19 - Reason for braking plotted against Speed at Brake Start and Brake Pressure per Unit Time, with colours defining (a) Cluster ID, (b) reason for braking

Braking events were classified into whether the event was one where AEB would be expected to come into play under more critical circumstances. AEB-relevant events were those where the driver braked to avoid imminent impact with another vehicle and which would have resulted in a front-to-rear impact. These classifications when mapped onto the graph, show that AEB-relevant events occurred in the moderate position on the graph, i.e. intermediate speed at braking and lower than average brake pressure over time (Figure 20).

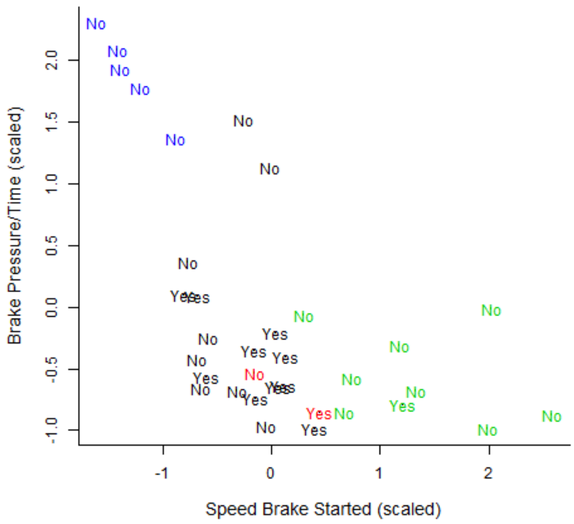

Figure 20 - Braking events classified as AEB-relevant, plotted against Speed at Brake Start and Brake Pressure per Unit Time, with colours defining Cluster ID

3.3.4 Conclusions

The analysis of reasons for harsh braking manoeuvres found that 28 of the 92 events recorded would hypothetically have resulted in a frontal impact and therefore fall within the desired scope for an AEB system. Further analysis of the geometry (distances, angles and overlaps) and speeds characterising these situations could be used to devise future test setups for AEB systems based on real-world data. The distribution of harsh braking events could also be used to inform priorities for future driver assistance systems. The high occurrence of give-way situations could be seen as an indication that drivers might benefit, for instance, from some form of red light, or right of way, assistance. The high prevalence of harsh braking before pedestrian crossings observed in the sample could suggest unused potential for optimising the infrastructure layout to increase conspicuity of the crossings.

The event cluster analysis suggests a way of partitioning braking events into types of behaviour before and during braking. In the sample used, it was difficult to conclude that different styles of braking are used in different driving situations, but some patterns did emerge. Braking on slip-roads (in the small sample) was characterised by low speeds and high brake application rate, and braking events associated with AEB style braking were associated with a small range of two key variables (Brake Pressure/Time and Speed at Brake Start). A larger sample of hundreds of such events could elicit stronger patterns allowing us to predict the braking situation based on behaviour characteristics. This could eventually lead to algorithms allowing an autonomous vehicle to anticipate the braking behaviour of others in a given environment.

3.4 Use case Traffic Sign Recognition (TSR)

In phase 1 of the TSR use case, the MOVE_UK team looked at TSR relevant data collected from the vehicles. A method was developed within the sFDE user interface, under the Traffic Sign Recognition view, to display the spread of all the traffic signs detected by the vehicle and their associated GPS coordinates on an open street map (OSM). At the same time, a clustering methodology was developed to calculate actual and missed detections. Phase 2 of the TSR use case carries forward the work completed in Phase 1, whilst at the same time collecting more traffic signs. The work focuses on identifying traffic sign clusters from the data collected in sFDE for conducting statistical analysis and an in-depth investigation.

Statistical analysis allowed us to understand the effects of external and vehicle related factors on the detection probability of a traffic sign, and the probability of detecting the correct speed. Evaluated factors included:

- Environmental condition such as rain
- Lateral acceleration of the vehicle
- Speed of the vehicle,
- Fog light on/off
- Head lights on/off
- Driving on different road class type such as motorway, artery roads and city roads etc.
- Wipers and other associated factors

The difference between detection probabilities between clusters for location-specific reasons was also examined. The outcome of this analysis will support the Royal Borough of Greenwich in maintaining an up-to-date database of traffic signs and identifying issues relating to the current positions of traffic signs, thereby potentially assisting with improving road-side infrastructure.

3.4.1 Identification of clusters for Statistical analysis

The following terminology is used in section 3.4 of this report:

Location – is the position of a place based on a fixed point on earth. The most common way to identify the location is using GPS coordinates such as latitude and longitude.

Cluster – a collection of traffic signs in a given location. Each cluster is characterised by a specific location number and a traffic sign number; for example, cluster 7_0 represents location 7 and 0 is the first of the total number of traffic signs contained within that cluster.

This analysis identified factors affecting detection of a traffic sign and its classification i.e. true or false detections. In order to do so, a relative sample of data points (i.e. speed limits and associated CAN signals) were collected and processed using ML libraries such as DBSCAN and K-means. The speed limit signs were located across Greenwich and detected by the MOVE_UK fleet over a period of 6 months. This amounted to 9,587 journeys in total and resulted in generating 11 clusters.

The distribution of these clusters contained a combination of different road classes, junction types and road geometry which were manually selected using the map view in sFDE. The selection was made on the basis of the clustering methodology page in sFDE around a particular junction, roundabout or road type. This implementation in sFDE helped in selecting clusters which had significantly higher detection rates than the ones which were not selected and vice versa. For each cluster, detection and missed detection events were recorded, along with a selection of CAN signals thought likely to be relevant to TSR.

R, a well-known statistical tool was used to conduct analysis for this purpose.

3.4.1.1 Purpose of the analysis

The purpose of the analysis was to see which factors affect the likelihood of detections (the system detecting a road sign when one is present) and the likelihood of true detections (the system detecting the correct speed on a detected speed limit sign).

3.4.1.2 Classifying Detections/Missed Detections

The use of sFDE Traffic Sign Recognition view to download relevant data and the use of big data analysis techniques helped to classify TSR data into two broad categories, namely detections vs missed detections, and true vs false detections. Only true detections were included in the final detection/missed detection data set. Further information on detections and missed detections can be obtained from section 4.4.2 of the report published for Phase 1 of the MOVE_UK project.

3.4.1.3 Classifying True/False Detections

Each time a MOVE_UK vehicle passed through a cluster, an event was generated and characterised as including either a detection event or a missed detection. For journeys where a sign was detected, these detection events were characterized as “True” or “False” detections.

True detections were classified as those which matched the modal (most frequent) speed sign reading for that cluster, and all other speeds were counted as false detections. For example, if 10 journeys through a cluster detected a 20 mph speed sign, and eight journeys detected at 30 mph speed sign, the correct speed was deemed to be 20 mph and all other detections would be classified as false detections.

This methodology may be sufficient when each cluster only includes one road sign. However, some clusters may include more than one speed sign, making it possible that two could be legitimately detected and read within the same cluster. In this case the number of false detections will be overestimated. The analysis conducted so far takes into consideration some limitation in the methodology itself but it does not take into account the direction of travel. For example, if 2 vehicles were traveling in opposite directions and correctly detected true speed limit signs on each side of the road, they would still be detected and read within the same cluster, thus leading to a low count of true detections.

Reducing cluster sizes and taking into account the direction of travel would go some way to resolving this but would increase model complexity. There is no analytical way to avoid these false negatives with the current method of deriving clusters. By conducting site visits (i.e. physical checks), it was possible to see whether multiple road sign detections within one cluster were present in the data.

3.4.1.4 Data Processing

The data was cleaned before analysis began by removal of unreliable results. For example, GPS uncertainty error of above 100 metres, undeterminable direction of movement, and some other anomalies) were removed. The initial dataset size had included 9,587 journeys and 222 clusters and cleaning reduced it to 6,594 journeys and 65 clusters. Before evaluation, any observations where the speed was less than 6 kph were also excluded, as the traffic sign recognition system used in MOVE_UK does not work below this speed.

Two new datasets were generated from the same cleaned dataset:

1. Detection/missed detection data set (6,122 journeys including 5,415 detections and 707 missed detections from 64 clusters).
2. True Detection/False Detection data set (5,887 journeys including 5,415 true detections and 472 false detections).

3.4.1.5 Factors influencing Detection/Missed Detection Probability

The detection/missed detection dataset was used to build a statistical model to determine which factors influenced the probability of missed detections, and by how much. Binomial regression models allow data with binary categories to be analysed, with both discrete and continuous variables used as inputs. Initially, 12 independent variables were fitted, and non-significant variables were then excluded one by one until the highest quality model was found (as measured by Akaike's Information Criterion).

Average detection rates were high at 96%, however there were factors which decreased the likelihood of a detection. The most significant of these was speed. A 10 kph drop in speed (relative to the average) reduced the probability of detection by 12%. Distance from the object in front also had an influence, and for every 10 meter reduction, the probability of detection reduced by 2%. These results indicate that cars in traffic jams are less likely to detect speeds signs than those on clear roads. The practical implications of this may be small given that cars in slow-moving traffic are unlikely to be exceeding the speed limit.

A number of environmental variables were tested including fog status (i.e. fog present or not present), lighting status and wiper status. Of these, only the wiper status had an influence, and that was relatively small and not statistically significant (therefore it cannot be ruled out that this result was due to chance alone). In the data analysed, a wiper status of 'on' (indicating rain) was associated with a 2% drop in probability of detection.

Road type was also influential, with types 4 and 5 associated with a 4% increased probability of speed sign detection above that of journeys made on road type 3. Road type 4 and 5 are roads that provide for a high volume of traffic movement at moderate speeds between neighbourhoods. Road type 3 (also referred to as secondary highways) are roads which connect city roads to major highways and provide for moderate volumes of traffic movement.

The effects of the various factors are illustrated in Figure 21.

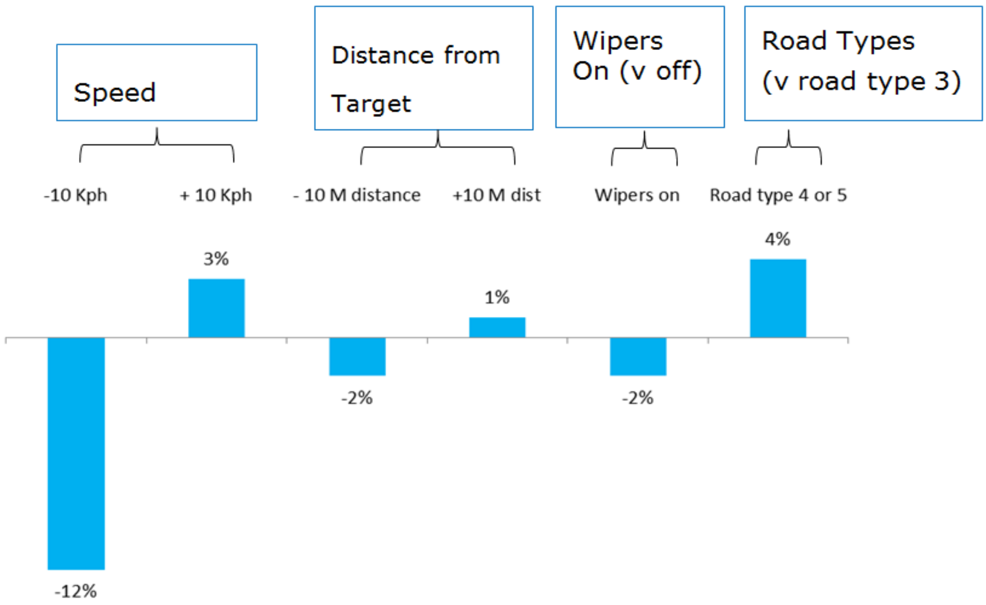

Figure 21 - Influence of various factors on probability of speed sign detection

As well as the effects already covered, cluster identity also had a significant influence on probability of detection. For the average cluster the probability of detection was 88% but there was a large amount of variation between clusters and 11 out of 64 clusters had a probability of detection of below 80%.

The differences between clusters are unaccounted for; it is possible that road configuration played a part, or damaged or defaced signs, or other factors. To find out more detail about the possible reasons for these differences, a ground-truth exercise was undertaken and 38 clusters were visited where road signs and unusual features were noted.

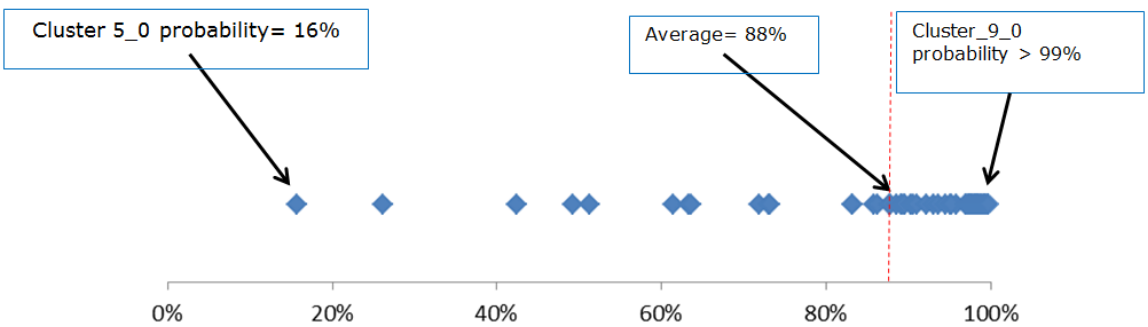

Figure 22 - Predicted detection rates for each cluster

3.4.1.6 *Factors influencing True/False Detection Probability*

A similar modeling approach was taken to see which factors influenced the likelihood of true or false detection. In this case the most influential factor was road type, with road types 4 and 5 giving improvements in the probability of true detection of 9% and 7% when compared to type 3 roads.

Positive lateral acceleration (turning or bearing left) was associated with a lower probability of true detection than a lateral acceleration of zero or negative lateral acceleration. This does not have an obvious interpretation, however, it may indicate when a car is turning into a side road and therefore possibly seeing a new road sign.

Speed also had a small effect on the probability of a true detection, with a 10 kph decrease associated with a 2% reduction in the probability of a correct detection.

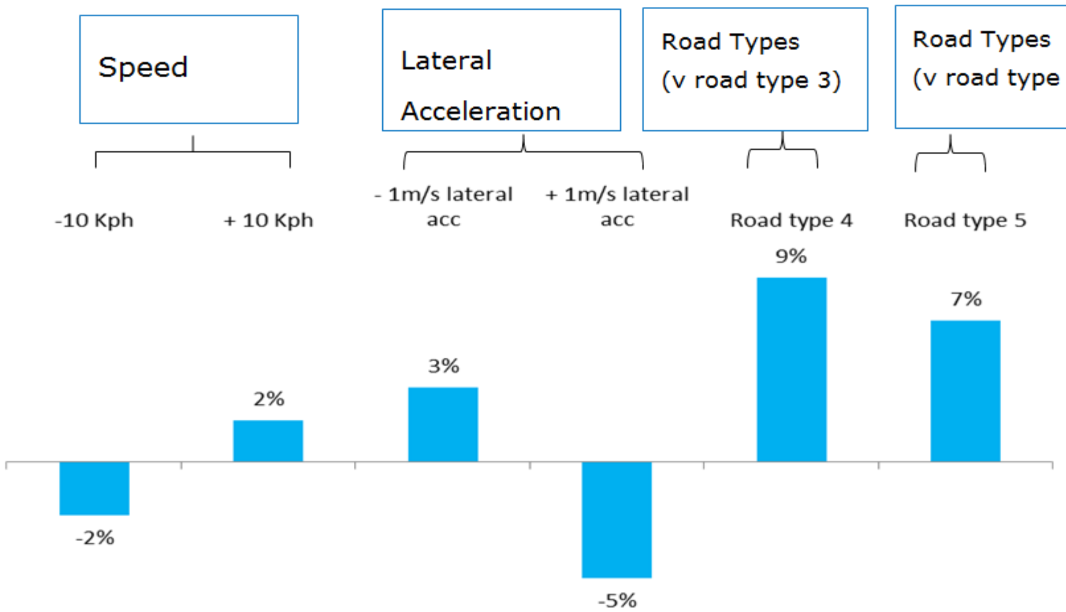

Figure 23 - Influence of factors on probability of true detection

The average cluster had a probability of true detection of 76%, however there was a wide range of variation between clusters with 24 out of 64 clusters having a rate of correct detection of below 80%. As above, a ground-truth exercise was carried out to find out why some clusters were much more likely to lead to accurate speed sign detection than others.

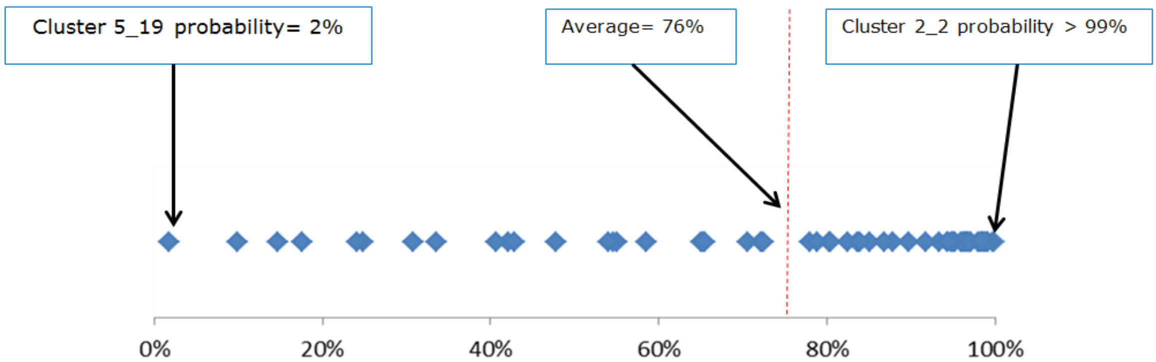

Figure 24 - Predicted true detection probabilities for each cluster

3.4.1.7 Information from ground truth exercise

A number of clusters were visited within Greenwich as shown in Figure 25. Different colours were used to represent different clusters on the map; for example, green symbols are used for cluster 1 and white symbols for cluster 7. A circle in a cluster represents a positive (and true) detection probability and a star in a cluster represents a negative or a low detection confidence probability. When a cluster was located, photos were taken from a number of different directions (2 directions on a street with no junction, and up to 4 directions at crossroads). A number of sources of potential ambiguity were found:

- Multiple signs visible within one cluster
- Signs obscured by street furniture
- Poorly angled signs on side streets
- Signs obscured by vegetation

Figure 25 - Locations of clusters across Greenwich

In some cases, multiple speeds signs were situated close together within a cluster, so there was some ambiguity as to what the 'correct' speed should be, something which is not helped by the limitations of the current cluster-generating algorithm which does not take into account the vehicle direction of travel. An example of this ambiguity was seen in cluster 7_0, where two signs were positioned in close proximity to each other (Figure 26). For the vast majority of journeys (97%) the 30 mph sign was recognised by the TSR system, but in 3% of journeys this area was recognised as a 20 mph zone. In this case the correct speed limit was 30 mph in the direction of travel.

Figure 26 - Cluster 7_0: ambiguous sign placement

Similarly, multiple instances were observed of speed signs signifying the speed limit for a side road, but which were angled well towards the main road and so could be detected by the vehicle from there (Figure 27 & Figure 28).

Figure 27- Cluster 5_0: 30 mph sign on separate junction also visible within cluster

Figure 28 - Cluster 5_0: 20 mph signs are visible from the main street

While the interpretation of these signs is straightforward for a human driver, it is more difficult for an AD system which can easily detect them from the main road and incorrectly determine them to be relevant. The magnitude of this effect was variable and depended on the exact road configuration: in one case this situation led to a low probability of detection, and in another the probability of detection was very high (almost all cars which passed it detected the sign as relevant to them - Figure 29). Other instances of poor detection could have been due to partial obscuration of signs by street furniture; for example a 30 mph sign was partially obscured by a traffic light (Figure 30).

Figure 29 - Cluster 8_0: nearly all cars passing this cluster picked up 20 mph as the relevant speed

Figure 30 - Cluster 5_0: 30 mph sign obscured by traffic light

In contrast, there were some clusters where the probability of detecting a road sign was very high, and the levels of speed ambiguity were very low. For example, cluster 2_2, which all cars identified correctly as 50 mph (Figure 31). This cluster included one road sign, which was on a straight stretch of road, with no street furniture obscuring it.

Finally, multiple instances were also observed of signs that were obscured by vegetation such as the ones in clusters 8_0 and 5_6 (Figure 32). These obscured signs would be extremely difficult to spot and read even for a human driver; therefore for a vehicle to do the same would also be difficult. Locations such as these lead to clusters with lower identity as the traffic signs detected by vehicles passing through them come with an extremely low probability.

Figure 31 - Cluster 2_2: 50 mph sign on a straight road was detected in all journeys passing the cluster

Figure 32 - Clusters 5_6 and 8_0: showing growing vegetation in front of the traffic signs

3.4.1.8 Conclusions from the Ground-Truth Exercise

The site visit uncovered issues with both speed sign placement and configuration, and issues with the methodology itself.

The detection and correct identification of speed limit signs was a particular challenge in residential areas, at junctions and cross-roads, where side roads and main roads had different speed limits. From observations of traffic sign installations throughout Greenwich it seems as though TSR detection may be improved by angling the speeds signs so that they are perpendicular to the road for which they are relevant, rather than, as in some cases, angled to be visible from the main road. That said, it is recognised that signs which are visible from the main road but relevant to a side road are likely to help human drivers who are about to turn into that side road. It may be necessary, therefore, to consider an alternative signage system which allows some warning to human drivers but is less ambiguous to machine reading, whilst still allowing human drivers to monitor and adapt to the speed limit.

Issues where signs were partially obscured would be more straightforward to deal with; road signs should be clear of any obscuring features such as traffic lights or trees.

In some cases clusters included a number of signs and it was impossible to ascertain which sign was relevant to each journey. Ideally, only one sign would be represented by each cluster. To achieve this the GPS data would have to be accurate within a very small range (5 to 10 meters), and the direction of travel would also need to be known so that journeys made in the opposite direction would not be included in the analysis due to a sign that was irrelevant to them.

3.4.2 Automation potential

There is potential for automating the methodology for calculating the probability of detection and true detection so that low detection probability and low true detection signs could be flagged. However, a number of barriers would need to be overcome first:

1) The cluster-generating algorithm would have to be improved so that only one sign is identified per cluster. This could include any (or all) of the following features:
    a. Improvements in GPS accuracy or the exclusion of results where a particular GPS accuracy was not met.
    b. Smaller cluster sizes.
    c. Knowledge of vehicle direction of travel when clusters are generated (so that only vehicles travelling in one direction are included within a given cluster).
2) Data cleaning steps would need to be automated (including removing data where GPS uncertainty was high).
3) The statistical methodology would need to be integrated with the cluster-finding algorithm and also with sFDE (currently the data collection, processing and statistical analysis are carried out in three separate steps).

If these barriers were overcome then it would be possible to output detection probability values for each road sign that a MOVE_UK vehicle passes. What would be more difficult is identifying the reasons for failure of detection in a systematic way. For this use case 38 cluster locations were visited to determine reasons for the detection probabilities. With a sufficiently large dataset (requiring many days of on-site surveying) patterns may start to be seen which may lead to more generalised predictions; for example, the prediction that a sign with a poor probability of detection at a minor road junction may have been poorly angled, but it would be impossible to know with any degree of certainty without a site visit. The output of the model could, however, generate a list of candidate signs with low detection probabilities for councils to follow-up on.

3.4.3 Next Steps

Clearly, the above analysis has identified areas of improvement such as: taking a relatively large sample size of data; removing unwanted data; refining and improving the coefficients of the clustering algorithms; including bearing information (i.e. direction of travel) in the analysis to obtain a better list of candidate signs. The next step would be to carry out a re-run of the analysis using the above-mentioned parameters to generate a model which may be integrated into the TSR maps in sFDE and also be used externally by Local Authorities.

3.4.3.1 Infrastructure update

Information gathered and processed above clearly shows a clear path for upgrading and planning current and future roadside furniture. This exercise has been extremely helpful in identifying locations which may need attention to clear obscured signs etc. Clearly this information can be extremely useful for any agencies which plan, install and maintain roadside infrastructure. The automation potential adds a new opportunity to speed up a process which has, until now, been done manually, and a method may be developed which separates spurious from relevant data, which could be used by RBG and others to update their roadside infrastructure.

3.5 Use case Telematics – Phase 1

3.5.1 Summary of Phase 1 work

In Phase 1 of the project extensive examination of signals from the MOVE_UK vehicles was made in order to understand the value of data in giving improved ability to understand risk and incidents. This analysis developed an evaluation approach and framework enabling the value of data to be judged. Phase 2 work develops this position, bringing new understandings for risk and incident evaluation based on new data and further analysis of it.

3.5.1.1 EDR method and Phase 1 position

In order to assess data in Phase 1, analysis was undertaken across all available data fields, aiming to understand the value of each in the context of future potential incidents. This work evaluated data from its values and distributions both alone, and in conjunction with other fields. For example, a brake pre-charge value was found to be involatile over longitudinal analysis in all but a few single-second instances; however, despite this involatile behaviour, the data proved to be of high value for EDR triggering, even when not occurring in all extreme situations. This, when found, was judged to be highly helpful.

3.5.1.2 Risk method and Phase 1 position

In order to assess a new data value in terms of understanding risk in Phase 1, analysis was undertaken across all available data fields aiming to understand the potential value of each data field for risk prediction. As this requires outcome data of crash incidents, a process was developed whereby the value of data fields was judged for its potential use as a risk discriminator. This process had to largely discard static and unchanging values as it was essential to use volatile data to judge and understand behavioural discriminators. For example, the pressure of the brake pedal and some of the CAN Bus accelerometer channels proved to be of high value to risk understanding, above and beyond base telemetry data.

3.5.1.3 Extension of work from Phase 1

Within Phase 2, the work from Phase 1 was further extended with additional signals and analysis to improve risk and EDR understanding for the Insurance sector. During Phase 1 it was shown that positional information regarding other road users had much higher value towards both EDR and risk. Therefore, assessment of the new parameters focussed on those that relate to the monitoring of surrounding objects; data from the newly-mounted forward-facing radar systems and additional data from the stereo video camera system.

3.5.2 Extension of Phase 1 EDR analysis

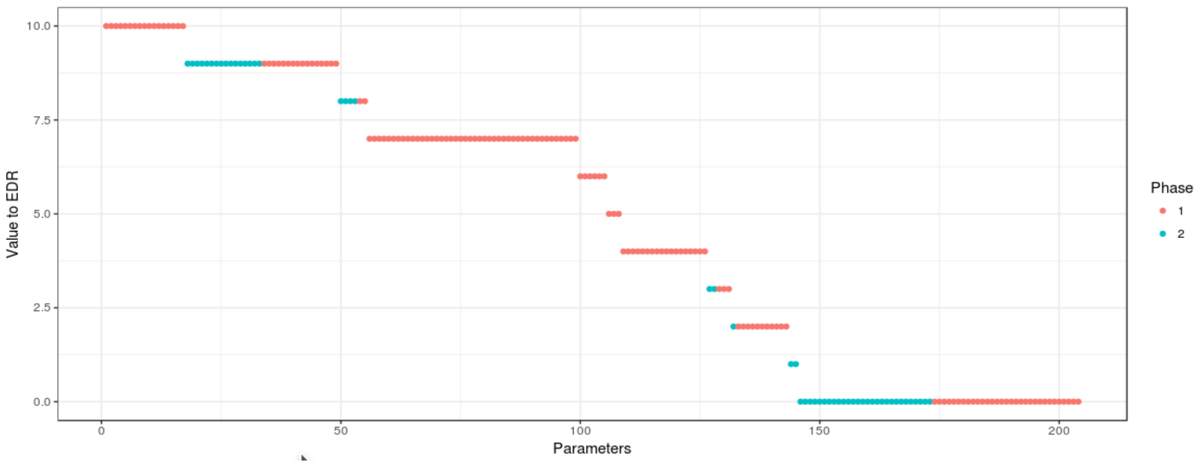

Figure 33 - Ordered ranking of EDR value for each parameter. Phase 2 parameters highlighted in turquoise

Figure 33 shows the ordered ranking of all parameters in terms of value to EDR. While many of the new parameters were found to be highly volatile, and therefore of value to EDR, other parameters were found to be static.

Many of the highest-ranked parameters related to the additional information from radar. This continued radar monitoring frames the *context* of a journey; the difference between moving at speed on an open road compared to a densely populated highway.

It is recognised that Figure 33 implies that over 50 parameters are of zero value to EDR, however, it should be noted that several of these parameters are those that were not fully implemented on the MOVE_UK vehicles, such as the oncoming object Radar parameters. It is reasonable to expect that, once implemented, many of these parameters would actually be of much higher value to EDR purposes.

3.5.3 Extension of Phase 1 Risk analysis

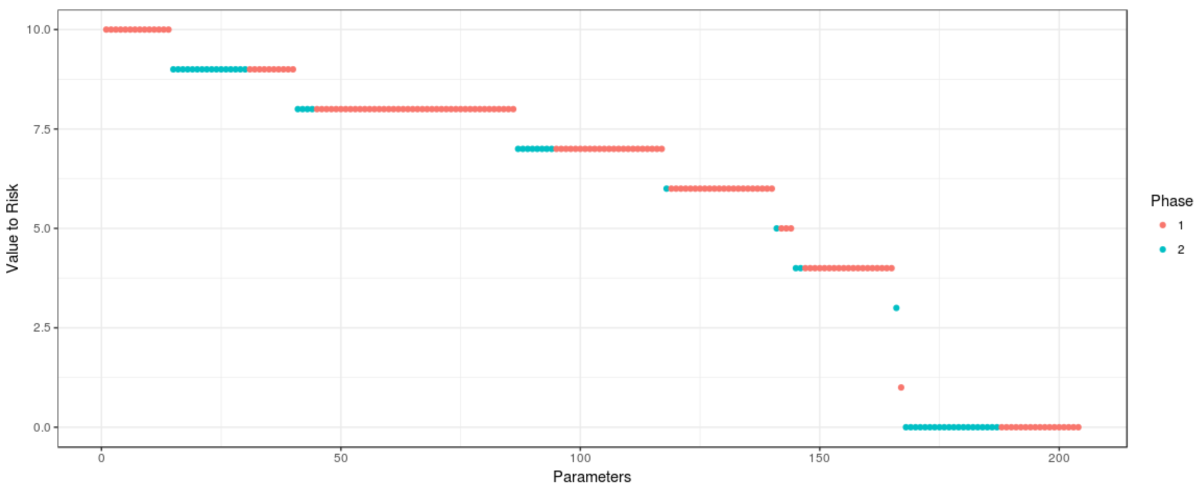

Figure 34 - Ordered ranking of Risk value for each parameter. Phase 2 parameters highlighted in turquoise.

The lack of variation observed in many fields means that they have negligible value for the assessment of long-term risk (Figure 34). In particular, those fields associated with the ACC and AEB system display a strong lack of the variability required to enable a risk segmentation from the infrequency of alteration. As with EDR, the highest value fields are those relating to positional information of other road users supplied by the radar and camera systems. While these are of higher value, the independence of the two systems occasionally leads to two different, competing perceptions of the leading vehicles. These differences make the challenge of interpretation more difficult, though ultimately create a richer and more complex picture of the surrounding environment at a given time. This is useful, in that both systems provide added value in their own right, and also collectively provide additional insight into the environment surrounding the vehicle; each have their strengths and weaknesses, but are complementary.

3.5.4 Conclusions of extension of work

The forward-facing radar allows for a precise mapping of the positions and speeds of vehicles moving in the same direction as the ego vehicle, when within its field of vision.

Radar and camera data provide an extremely valuable contextual understanding of a journey at any given time. However, some aspects of the implementation clearly impact the trust that can be placed in each data stream. For example, our analysis of object identification made by the radar firmware identified that occasionally the in-lane, leading vehicles are incorrectly identified. This occurs at instances when the signal disappears from either optical or radar systems due to occlusions, road curvature or inclinations (e.g. speed bumps). The techniques and applications developed here for identifying and evaluating such issues with vehicle detection have wide applicability to the specification and testing of future systems.

4 Phase 2 Use cases and capabilities

4.1 Overview

In the introduction (Section 1) the updated version of the objectives diagram was presented, with potential areas of application for the methods developed in MOVE_UK. With the limited number of vehicles, the project does not intend to perform any of these applications, but rather to develop, trial and demonstrate certain capabilities that will be needed for these applications.

This project does not intend to use the developed method of connected silent validation to validate the production level ADAS features fitted to the trial vehicles (this has already been done; those features are already in production). Instead, MOVE_UK intends to demonstrate that silent connected validation is feasible and can be performed at a large scale for the development and validation of future ADS.

Examples of capabilities to demonstrate in this context are:

- Identification of relevant events using on-board vehicle sensors.
- Transmittance of high-bandwidth data at a scalable level.
- Connectivity of a vehicle fleet in order to swiftly collect the required data.
- Re-simulation of recorded events with modified parameters.

As already described, the technical project work has been organised in 'use cases', which relate to the different applications and demonstrate different capabilities. A use case consists of a description of the events considered to be relevant (which is 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.

In MOVE_UK Phase 2, four use cases have been developed and implemented:

- Subcritical Radar AEB Activation (ARB)
- Lead Vehicle Statistic (LVS)
- Cut-In scenarios (CIN)
- Telematics – Phase 2

Figure 1, the objectives diagram presented in the Introduction, provides an overview of how each use case relates to the applications described in the phase 1 Report (D7.3). The use cases are covered in more detail in the subsequent individual sections. The following provides guidance on terminology throughout these sections:

When validating the activation pattern of a vehicle system, the potential errors can be thought of as falling into one of two categories:

- False positive activation: The system activates where it should not. For example, the ARB system would react to a 'ghost' object and apply the brakes.
- False negative activation: The system does not activate where it should. For example, the ARB system would not detect a vehicle in front and not apply the brakes in a critical situation.

Accordingly, correct system behaviour can be classified as:

- True positive activation: The system activates where it should. For example, the ARB system intervenes to prevent a front-to-rear collision with a leading vehicle.
- True negative activation: The system does not activate where it should not. Most of normal driving: No critical situation is present and the ARB system does not activate.

The Phase 2 use cases were started in parallel, but implemented at different times, which is why the amount of data recorded and the stage of analysis varies between them. Additional use cases will be developed in Phase 3.

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

For the purposes of the MOVE_UK project, radar based AEB functionality is referred to as ARB in order to distinguish it from video based AEB.

4.2.1 Purpose

As with the video AEB use case, the purpose of the Subcritical radar-based AEB use case is to develop, trial and demonstrate the capabilities required to perform silent connected validation for ADAS or ADS systems. In order to use synergies from Phase 1, the video AEB use case was modified appropriately to be applied to radar.

The following paragraph is analogous to the Video AEB use case described in D7.3 report of Phase 1, as it applies identically to the radar-based AEB use case. In order to achieve a comprehensive and complete understanding of the radar-based AEB use case, this paragraph is presented once more:

With traditional validation methods, the matter of *false positive* ARB activations particularly requires high test drive mileages. This is because the situations causing false activations occur infrequently and have a wide range of potential causes, which do not follow a systematic pattern. To demonstrate silent connected validation of the ARB system in respect of false positive activations, this use case aims to collect sequences of all real-world situations encountered during the trials where ARB would activate. The activation parameters of the ARB algorithm in MOVE_UK are calibrated to react more sensitively and earlier than they would in production vehicles, in order to capture a wide set of potentially relevant situations (false positives as well as true positives). Note that the ARB system is operating in silent mode, i.e., it does not activate the brakes in the vehicle but only creates the relevant CAN signals. The recorded set of collected data is then used to simulate system reaction with production version calibration to validate whether any of the situations would have caused an ARB braking event, where they should not have.

This use case is intended to demonstrate the following capabilities of the validation approach developed in MOVE_UK:

- Trigger sequence recording using a radar-generated CAN signal.
- Selectively capture a set of real-world system activations which are of high relevance for the Consortium.
- Transmit high volumes of data at a scalable level (high-bandwidth data, including high-resolution video images and CAN data).
- Initial classification of the recorded sequences into false positive and true positive situations by analysing the driver braking behaviour observed.
- Re-simulation of relevant situations based on the data recorded.

4.2.2 Design

The parameters for the ARB function of the radar sensors are set to a more sensitive level than in production in order to detect more situations and collect more data. The production setting of these ARB parameters would lead to no or to very few triggered situations during the project lifetime. This is because the production setting is designed to eliminate *false positive* activations (i.e. where the system incorrectly applies full braking). The calculation for the decision to brake or not to brake is realised inside the radar system. No additional computer or other devices are involved. In comparison with a camera system, the radar sensor applies different levels of AEB depending on the criticality and dynamics of the situation. In these different levels, different deceleration values are requested from the brake system in order to ensure appropriate reactions to the individual situations. This is made possible due to the exact measurement of distances and relative velocities by the radar sensors.

For the triggering criteria in this use case, the thresholds of some of the internal parameters of the ARB function (e.g. driver reaction time, lane and object probability of target object) were slightly changed. With these parameters, a criticality value is calculated and when the criticality value rises above a certain threshold, ARB is activated.

Identically to the AEB use case, when the criticality value reaches the defined area, a trigger signal is sent from the radar system to the measuring system in the vehicle. A video and radar sequence and corresponding CAN data is collected, stored and transferred to the cloud. These event-based video and radar sequences are used by the Consortium for further analysis.

4.2.3 Synthesis and analysis

Within the Subcritical ARB use case, 58 events were captured during the course of a nine-month period within Phase 2 of MOVE_UK.

The data analysis has so far covered two aspects:

1. Identification of the objects causing an ARB trigger
2. Re-simulation of the situations captured to classify them as false positive or true positive using a production ARB calibration

The following still frames (Figure 35 - Figure 36), which were captured during Phase 2 of the MOVE_UK trials, show different examples of the camera view at the moment of a Subcritical ARB trigger.

The radar sensor uses the current speed, trajectory and relative distance to the target objects to calculate if a collision would be unavoidable if no braking manoeuvre were to be initiated. In this eventuality, deceleration is requested. The level of the requested deceleration depends on the relative distance to the target object and the dynamics of the situation. Note that all three of the following example situations were triggered because of the changed driver reaction time parameter. The red arrow highlights the object which caused the trigger. Note that the colour red is not meant to convey any meaning in this context.

Example 1 (Subcritical ARB sequence 48): lead vehicle slowing down for lane switch

A MOVE_UK vehicle is driving behind a vehicle whilst approaching a roundabout, when the lead vehicle slows down to switch to the left lane (Figure 35). Due to the lead vehicle braking, the driver of the MOVE_UK vehicle has to brake also. The radar sensor calculates that, with the pre-configured driver reaction time, a crash would be unavoidable, thus a deceleration is requested (in silent mode). Due to the large distance to the target object and the low dynamics, the requested level of deceleration is low.

Figure 35 - Leading vehicle triggering Subcritical ARB Sequence 48 in May 2018. Still frame of the trigger extracted from the stereo video recording. Red arrow highlights the object causing the trigger.

Example 2 (Subcritical ARB sequence 27): lead vehicle braking unexpectedly

A MOVE_UK vehicle is stationary behind a vehicle in a traffic jam (Figure 36). When the lead vehicle starts up, the driver of the MOVE_UK vehicle also does so. Almost immediately after starting up, the lead vehicle brakes unexpectedly, which leads to a braking manoeuvre from the driver of the MOVE_UK vehicle. The radar sensor calculates that with the pre-configured driver reaction time that a crash would be unavoidable, thus a deceleration is requested (in silent mode). Due to the short distance to the target object and the high dynamics, the requested level of deceleration is high.

Figure 36 - Leading vehicle triggering Subcritical ARB Sequence 27 in March 2018. Still frame of the trigger extracted from the stereo video recording. Red arrow highlights the object causing the trigger.

Example 3 (Subcritical ARB sequence 32): motorbike in front braking due to speed bump

A MOVE_UK vehicle is being driven behind a motorbike whilst approaching a speed bump (Figure 37). When arriving at the speed bump, the motorbike in front brakes, which leads to a braking manoeuvre from the driver of the MOVE_UK vehicle. The radar sensor calculates that with the pre-configured driver reaction time that a crash would be unavoidable, thus a deceleration is requested (in silent mode). Due to the large distance to the target object and the low dynamics, the requested level of deceleration is low. The target object which triggered was a motorbike, which is noteworthy as it is a rather small object.

Figure 37 - Leading vehicle triggering Subcritical ARB Sequence 32 in April 2018. Still frame of the trigger extracted from the stereo video recording. Red arrow highlights the object causing the trigger.

4.2.4 Discussion

The following paragraph is analogous to the Video AEB use case described in the D7.3 report of Phase 1, as it applies identically to the radar-based AEB use case. For reference, this paragraph is presented once more here.

With this ARB use case collecting subcritical ARB situations the following capabilities were successfully demonstrated: identifying real-world *false positive* situations; automated capturing of events which are activated by the system within the vehicle (on board); transmitting high volumes of data over the air; and the process of using real word data to re-simulate in order to optimize parameters that then can be re-flashed to the control unit in the field for the next iteration. In reaching our goal for these capabilities, methodologies have now been generated and an example infrastructure to validate actual and future ADS systems.

The radar sensors as a decision unit in the vehicle successfully identified subcritical situations. The trigger signal for these events was generated over 58 times. The analysis of the situations showed that the driver reaction time parameter for the trigger was set to a very subcritical level, resulting in many non-dynamic sequences with less value for the analysis. Therefore, a new subcritical parameter set closer to production setting was found through re-simulation of these sequences. In order to still gather a satisfactory amount of data, the parameter set is kept at a level of subcriticality which will permit some sequences with low dynamics. A further, final adjustment of the parameter settings is planned after a sufficient amount of sequences is gathered. The final setting will be very close to production settings and only near critical situations will be captured.

Moreover, it was noticeable how many events were triggered due to speed bumps. Most of these triggers had low dynamics and high distance to the target objects leading to sequences with less value for data analysis. For a real-world scenario of autonomous driving cars, AD systems which take better account of speed bumps, or removal of speed bumps from roads altogether, should be considered.

The measurement system in the vehicle captured and stored the relevant situations automatically, and the Consortium are satisfied with the way it performed. During the trials the time of transferring the data from the vehicles, via Wifi hotspot, to the cloud was again substantially improved.

The transmission unit including the cloud infrastructure sent and received the high bandwidth and high volume data and processed it for further analysis. As with video AEB, the result is a situation database where developers of ADAS systems can work with and improve algorithms and finally create safe and better quality products.

The collection of radar based subcritical ARB sequences will continue during Phase 3 of MOVE_UK.

4.3 Use case Cut-In (CIN) Scenarios

4.3.1 Purpose

A Cut-In, which is characterised by a vehicle moving in front of the ego vehicle from an adjacent lane is a situation which a highly automated vehicle should react to in a way which a human driver would expect and would feel comfortable with.

The Cut-In scenario use case implemented in the MOVE_UK vehicles is used as an example to show the benefits of real-world event data. The use case is used to help understand driver behaviour (of both driver and surrounding drivers) in the particular situation of a Cut-In during a traffic jam on the motorway. This situation is chosen for Cut-In evaluation due to its well-defined characteristics such as low speed driving in one-way traffic (i.e. no oncoming traffic) and on a dedicated road type which helps narrow down data selection for big data analysis. A number of behaviours can be associated with a Cut-In situation, including but not limited to, drivers trying to cut across lanes in order to exit a motorway at a junction, changing lanes to avoid a long queue, or to get past a slow-moving vehicle. Understanding driver behaviour in these situations can benefit other Consortium projects by helping them define parameters for safe and human-like automated driving functions. The detection of Cut-Ins can be used in the context of development of autonomous vehicles to form a basis for decision making, including eventual precautionary safety measures. Of particular interest in this regard are real-world parameters describing the behaviour of both the ego driver and the surrounding drivers, such as relative speed and distance between the vehicles during the manoeuvre, or the amount of braking and steering by the ego vehicle driver.

The detection of the Cut-In traffic scenarios is based on sensor data from the radar, which is fitted in the front of the MOVE_UK vehicle. The real-world parameters to be analysed are available on CAN and gathered around the time event of the detected Cut-In.

4.3.2 Design of the data collection

Before describing the setup of the Cut-In use case, the terminology used should be clarified:

Target Objects: The term "Target Objects" refer to objects (vehicles) detected by the front radar sensor and located in specific positions relative to the ego vehicle. In this project, four target objects OBJ_01 – OBJ_04 were recorded and analysed, all of which were travelling in the direction of the ego vehicle (x-direction). Their relative positions are shown in a so called 'bird's eye view' (as seen previously in Figure 3). Target Object OBJ_01 is the vehicle directly in front of the ego vehicle (the 'lead' vehicle) and is the object of most interest. Target Object OBJ_02 is the nearest vehicle in front of the ego vehicle in the left lane. Target Object OBJ_03 is the nearest vehicle in front of the ego vehicle in the right lane. Target Object OBJ_04 is the vehicle in front of the lead vehicle OBJ_01.

Vehicle ID: In a lower level of computation, the radar sensor detects 32 different objects in its vicinity and assigns each object a number or 'Vehicle ID' from 1 to 32. Each Target Object is given the Vehicle ID of the object that fulfils the criteria for that Target Object at a particular point in time. If, as the vehicles move around, a new object fulfils the criteria for a Target Object, then that Target Object is assigned the Vehicle ID of the new object. The corresponding CAN signal for the Vehicle ID on the CAN bus is called 'OBJ_0X_nr', where 'X' is the Target Object number and 'nr' is the Vehicle ID associated with that Target Object at that point in time.

During a Cut-In manoeuvre, the presence of a new lead vehicle (i.e. the vehicle that has cut in) can be observed through a change in the vehicle ID of OBJ_01. If, for example, the vehicle to the left of the ego vehicle (OBJ_02) performs a lane change and cuts in in front of the ego vehicle, then that vehicle becomes OBJ_01 and the Vehicle ID in the signal OBJ_01_nr changes to the Vehicle ID of the cut-in vehicle. At the same time, the old lead vehicle becomes OBJ_04 and the Vehicle ID in the signal OBJ_04_nr changes to the Vehicle ID of the old lead vehicle (see Figure 38 below). Note, a new OBJ_01 does not necessarily have to be classified as OBJ_02 or OBJ_03 prior to moving in front of the ego vehicle in order for the manoeuvre to be considered as a Cut-In. The new OBJ_01 could equally be one of the other 32 objects detected but not previously classified as a Target Object or another object previously undetected by the radar.

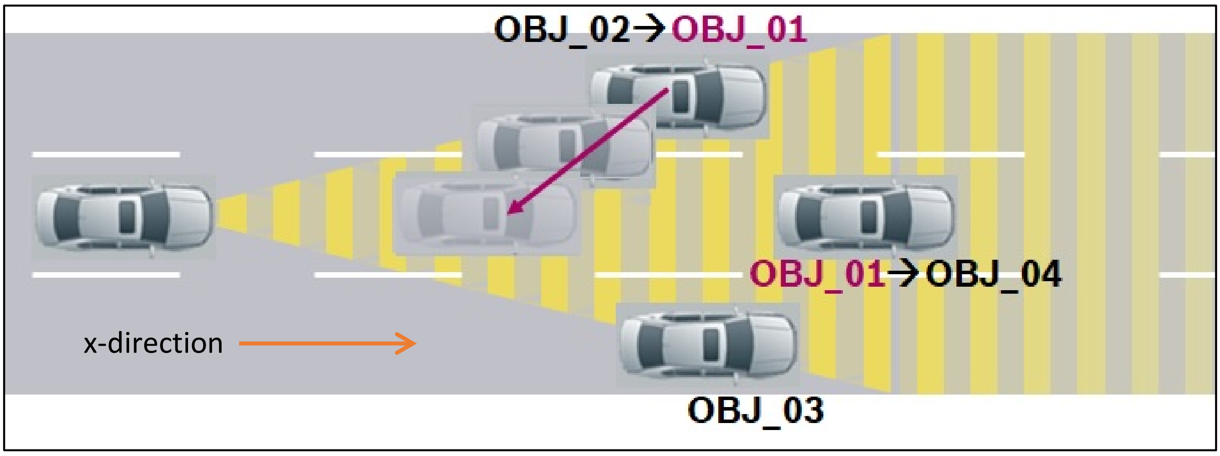

Figure 38 – Example of a Cut-In manoeuvre

The task of identifying and defining a suitable trigger for the CIN use case has been a step-by-step approach. The process started with the definition of an initial trigger which was then refined on the basis of the knowledge gained through the collected data. The steps are described below, in the order of their implementation.

4.3.3 Identification of signals recorded at 10 Hz

In order to enable the MOVE_UK analysts to define a meaningful Cut-In trigger, data from relevant situations had to be captured as a first step. It became apparent that due to the speed at which objects change position, the analysis of the 1 Hz continuous CAN data collected could not result in the identification of such a trigger, since too much information is lost between sample points. Therefore, the MOVE_UK consortium carefully selected a list of CAN signal to be recorded at 10 Hz (i.e. 100 ms between samples). These signals were chosen for their relevance to the Cut-In use case and, for example, included "Distance to OBJ_01", the vehicle IDs, and relative velocities for all 4 Target Objects.

4.3.4 Implementation of the preliminary *Traffic Jam Trigger*

Due to performance considerations, the dense 10 Hz data could not be recorded in a continuous manner. Therefore, it was decided to only trigger the recording of data in traffic jam situations which are of particular interest for this use case. This 'Traffic Jam Trigger' was implemented on the Flea 3 box as a preliminary step to gather dense data for analysis and was defined according to the following trigger conditions: speed of the ego vehicle below 65 kph, OBJ_01 detected, restriction to certain major road types, and a maximum distance of 70 m to OBJ_01.

4.3.5 Implementation of the Time Gap Formula on the CAN Gateway

The first important step in refining the Traffic Jam Trigger towards an actual Cut-In trigger was to replace the trigger condition "maximum distance to OBJ_01" with "time gap to OBJ_01". This time gap was calculated from the ego vehicle speed and the distance to OBJ_01 and was preferred because, unlike distance, a time gap takes account of the velocity of the vehicles involved and is the critical factor in collision avoidance.

The Flea 3 box has certain limitations calculating new signal values as well as comparing one signal value with another. Therefore, calculation of the time gap as a new signal 'TimeGap' was implemented on one of the hardware CAN gateways in the vehicle. The Flea 3 box was then able to monitor this new signal and a trigger condition "maximum TimeGap to OBJ_01 of 4s" was chosen to replace the previous trigger condition "maximum distance to OBJ_01 of 70 m".

The implementation of the time gap signal on the CAN gateway mainly consists of the simple division of "distance to OBJ_01" by "ego vehicle speed". However, to prevent division by zero, the algorithm checks the distance to OBJ_01 in the case of ego vehicle speed below 7 kph. If the distance to OBJ_01 is too great, a time gap of over 4s is set by the gateway in order to prevent a trigger by the Flea.

4.3.6 Implementation of the ID Change Signal on the CAN Gateway

The second important step towards a Cut-In trigger was to refine the actual point in time when a Cut-In is detected. As previously described in 4.3.2, a change in the CAN signal OBJ_01_nr (i.e. the Vehicle ID associated with OBJ_01) can be used to identify a time point when a different vehicle is detected in front of the ego vehicle, which is a good indication that a Cut-In manoeuvre may have taken place.

Since the Flea is also limited in relation to time-based calculations, another additional signal denoting the change in Vehicle ID of OBJ_01 was implemented on the CAN gateway and called "ID_Change_Trigger". The algorithm used to generate this signal was enhanced by introducing a time delay function to filter out "spikes" in the OBJ_01_nr signal of less than 2s. Also, a comparison of the previous and new distances to OBJ_01 was added to avoid a trigger in case of the vehicle in front cutting out of the lane instead of in.

When a valid Vehicle ID change occurs, the signal ID_Change_Trigger holds the value 1 for 200 ms. The condition of "ID_Change_Trigger = 1" was, therefore, added to the trigger conditions in the Flea 3 box.

4.3.7 Summary of final Cut-In Trigger conditions implemented in the Flea 3 box

Following the steps mentioned above in sections 4.3.3 to 4.3.6, a final Cut-In trigger was achieved based on the following trigger conditions:

- Ego vehicle speed below 65 kph
- OBJ_01 detected
- Road Type denoting major arterial roads
- Time Gap signal value below 4 s
- ID Change Trigger signal value of 1

The duration of the triggered Cut-In measurement was chosen to be 15 s prior to and 15 s post the trigger point. However, in order to compensate for the 2s delay of the spike filtering functionality, a duration of 17 s prior to, and 13 s post, the change of the ID Change Trigger signal was implemented in the Flea 3 box.

Following the successful testing of the Cut-In trigger on one vehicle, an updated software configuration was deployed to all vehicles in the MOVE_UK fleet.

4.3.8 Visualisation of Cut-In events in sFDE

For the analysis and visualisation in sFDE of the Cut-In data captured from the MOVE_UK fleet following deployment of the Cut-In trigger, the use case technical group decided to give the end user two possibilities:

1. To view individual Cut-In events in full detail ("CIN Sequence View"), and
2. To receive statistical evaluations of certain interesting parameters (such as distance and relative velocity to OBJ_01 around the point of Cut-In,) based on all available Cut-In data ("CIN Statistics View").

4.3.8.1 CIN Statistics View

To visualise the behaviour of the ego driver and surrounding drivers (vehicles cutting in) the chosen statistics can be separated into two categories; “Ego Driver Reaction Statistics” and “CIN Vehicle Statistics”. For the Ego Driver Reaction Statistics, the behaviour of the ego vehicle driver in reacting to a Cut-In is analysed. Statistics are generated for the maximum brake pressure and the maximum absolute steering rate exercised by the ego driver during a certain evaluation period (currently starting 7s before the Cut-In trigger and ending 7s after the trigger) . See Figure 39.

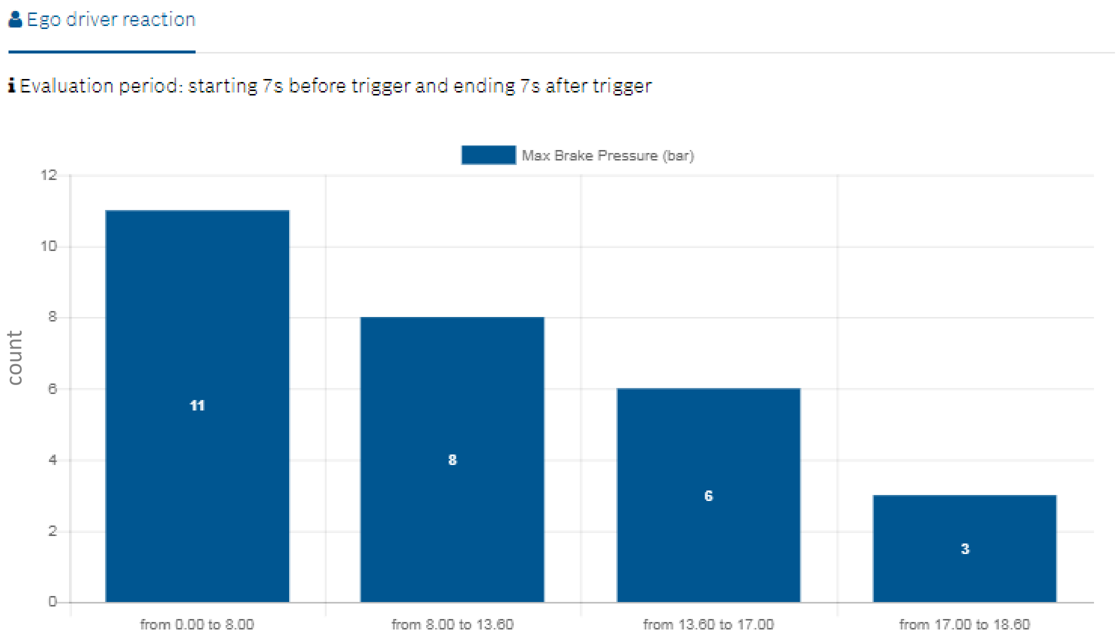

Figure 39 - Ego driver reaction statistics

For the CIN Vehicle Statistics, the behaviour of the driver of the vehicle executing the Cut-In is analysed. The minimum distance from the ego vehicle to the Cut-In vehicle (OBJ_01), both in metres and seconds, and maximum relative negative velocity between the ego vehicle and Cut-In vehicle (OBJ_01) are evaluated within a defined evaluation period after the Cut-In trigger. All point pairs of “TimeGap to OBJ_01 (s)” vs “Relative velocity of OBJ_01 (kph)” are plotted together as a scatter plot (Figure 40).

The main difference between these two categories is the evaluation period which, in the case of Ego Driver Reaction Statistics, begins at 7s prior to and ends 7s after the Cut-In trigger, while for CIN Vehicle Statistics it starts with the trigger time itself and ends 10s afterwards. The reason for this distinction is the fact that information about the Cut-In vehicle is only available from the point when it actually “becomes” OBJ_01, while the ego vehicle driver potentially has reacted already before the Cut-In manoeuvre concludes.

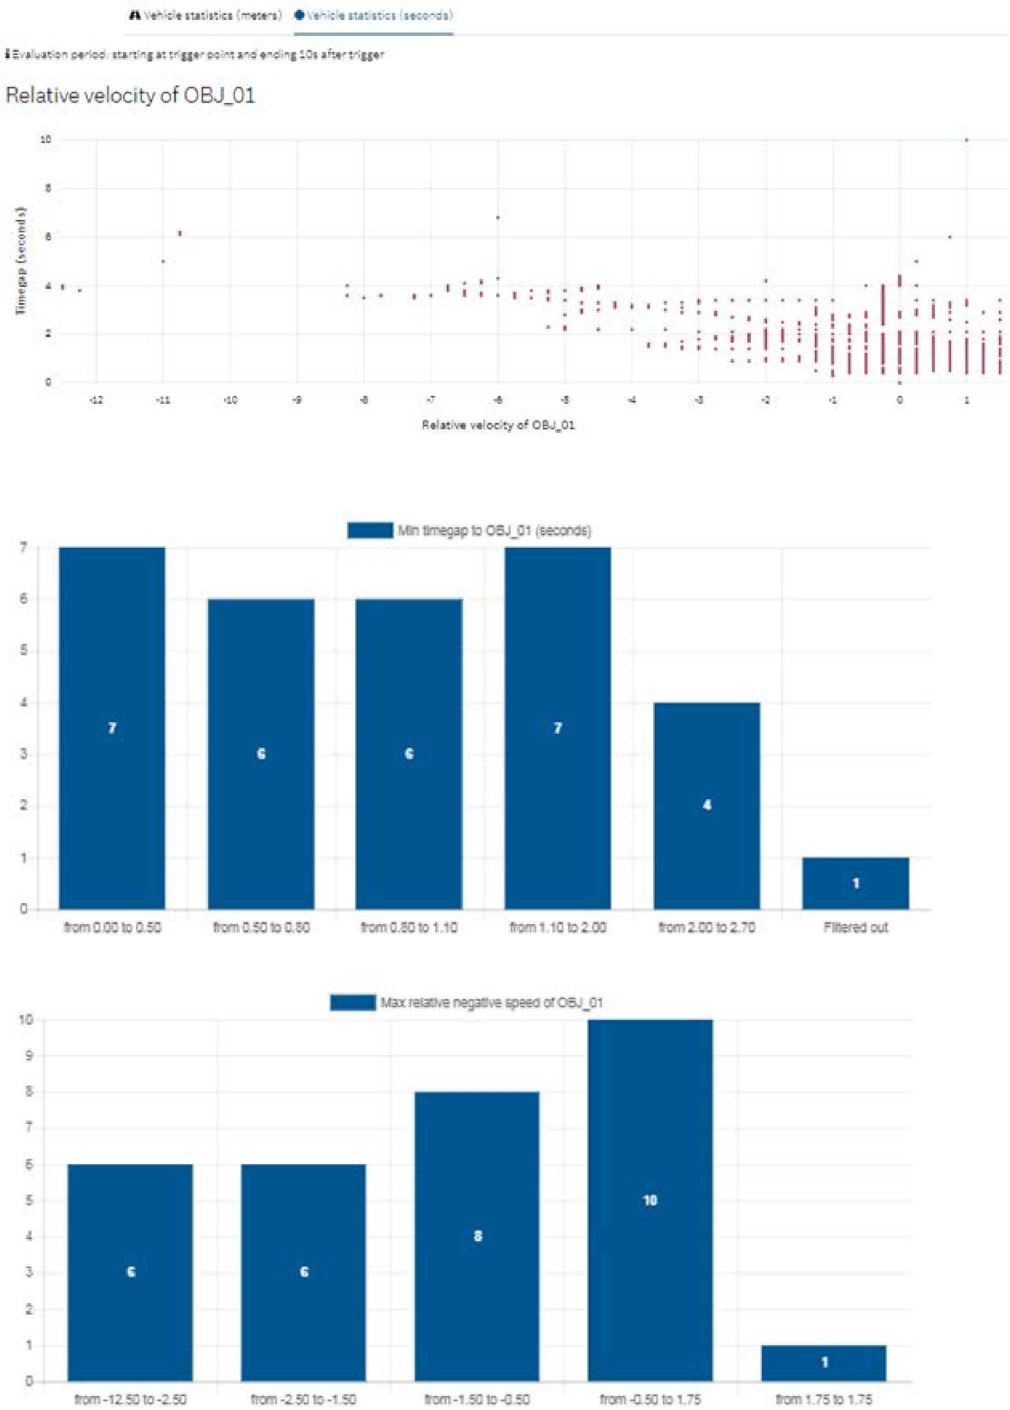

Figure 40 - CIN vehicle statistics

4.3.8.2 CIN Sequence View

The CIN Sequence View page displays the individual CIN events and associated CAN signals for each data point during the triggered measurement. The visualisation consists of the part of the map where the triggered sequence was captured, a radar "bird's eye" view of the relative positions of OBJ_01 – OBJ_04 and a list of pre-selected CAN signals (Figure 41).

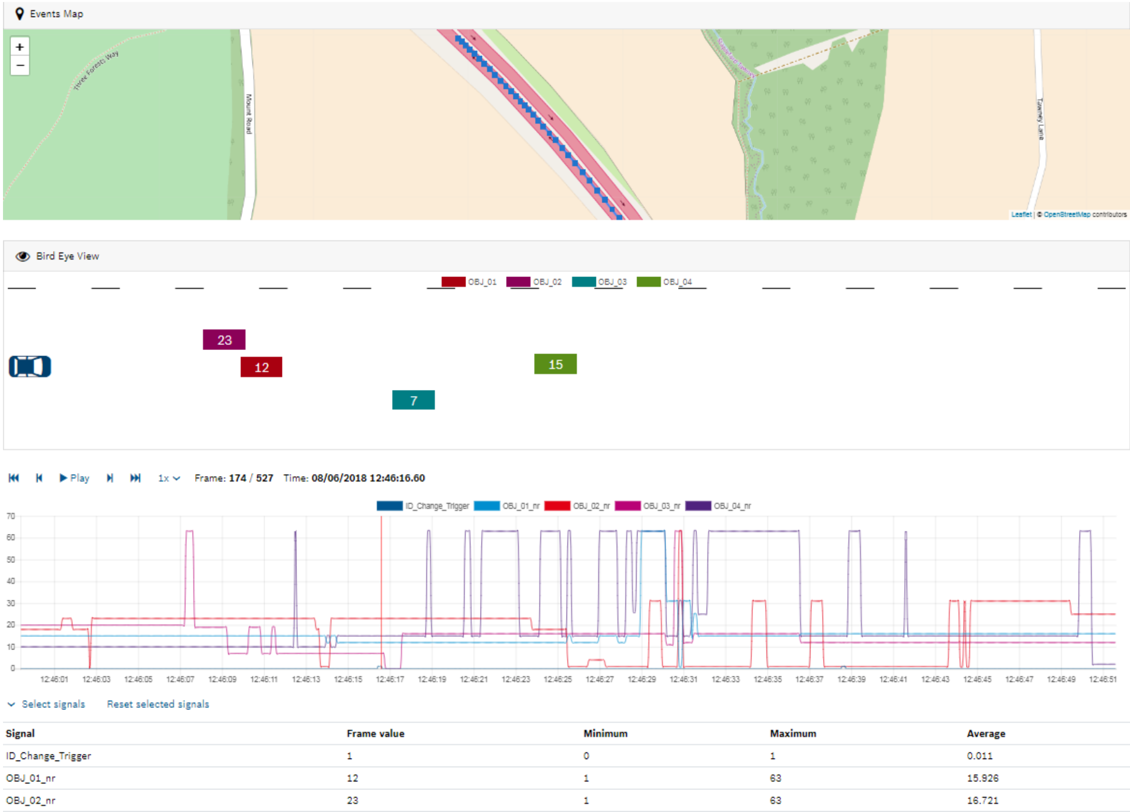

Figure 41 - CIN sequence view

4.3.9 Synthesis and analysis

4.3.9.1 Analysis of single Cut-In measurements

For an in-depth analysis of an individual Cut-In manoeuvre, the CIN Sequence View in sFDE can be used. In the following, an analysis of two possible Cut-In candidates is described.

The first CIN sequence was recorded during a drive on the M25 motorway. A snapshot of this sequence can be seen in Figure 42 below.

Figure 42 - A snapshot of Cut-In candidate 1 taken from sFDE's birds eye view

Figure 43 shows the progression of the vehicle IDs of OBJ_01 to OBJ_04 over time for the entire duration of the Cut-In sequence. It also shows the ID_Change_Trigger signal which changes from 0 to 1 after the Vehicle ID for OBJ_01 (OBJ_01_nr) changes; note, this signal is delayed by 2 seconds compared with the other signals shown.

From looking at Figure 43, it can be seen that a vehicle with Vehicle ID 31 is the lead vehicle until t=15:52:03 when a vehicle with ID 18 (previously classified as OBJ_02) seems to cut in from the left side and replaces Vehicle ID 31 as OBJ_01. At the same time, Vehicle ID 31 seems to move to the right hand lane to become OBJ_03 and a Vehicle ID is no longer assigned to OBJ_02. From this, it can be concluded with a fair amount of certainty that this is actually an example of the ego vehicle cutting out instead of another vehicle cutting in.

Figure 43 - A view from SFDE showing the change in CAN signals over time for Cut-In candidate 1

The second CIN sequence analysed ("Cut-In candidate 2") represents a true Cut-In. This sequence is illustrated in Figure 44 and Figure 45. As can be seen, the radar sensor detects an OBJ_02 with a Vehicle ID of 15 which later (at t = 45:44:00) cuts-in in front of the ego vehicle to become the new OBJ_01, replacing the previous OBJ_01 Vehicle ID of 10. At the same time, the vehicle with ID 10 becomes OBJ_04 (see Figure 44). From this, it can be concluded, with a high degree of certainty, that this sequence constitutes a true Cut-In event.

Figure 44 - SFDE bird's eye view of Cut-In candidate 2 before and after the cut in has taken place

Figure 45 - A view from SFDE showing the change in CAN signals over time for Cut-In candidate 2

4.3.9.2 Analysis of Cut-In statistics

The Cut-In sequences collected are also processed to provide the statistics described earlier in section 4.3.8.1. The following figures below represent statistics of the ego driver reaction and the vehicle related statistics from a sample of the 24 Cut-In events triggered so far in Phase 2 of the project.

As seen in Figure 46, there have been three events where the driver applied the brakes at a maximum brake pressure between 17 - 18.60 bar. Such high levels of braking indicate an aggressive cut in manoeuvre, possibly into a small gap between the ego vehicle and the lead vehicle. Lower levels of braking were seen more frequently with 10 events where the maximum brake pressure was between 0 - 8 bar.

In terms of the ego driver steering rate, there have been 10 events where the maximum absolute steering rate was between 4 – 12 degrees/sec. This represents a gentle steering manoeuvre, compared to 2 events where the maximum steering rate was between 92 - 124 degrees/sec which indicates a strong ego driver reaction to a cut-in event.

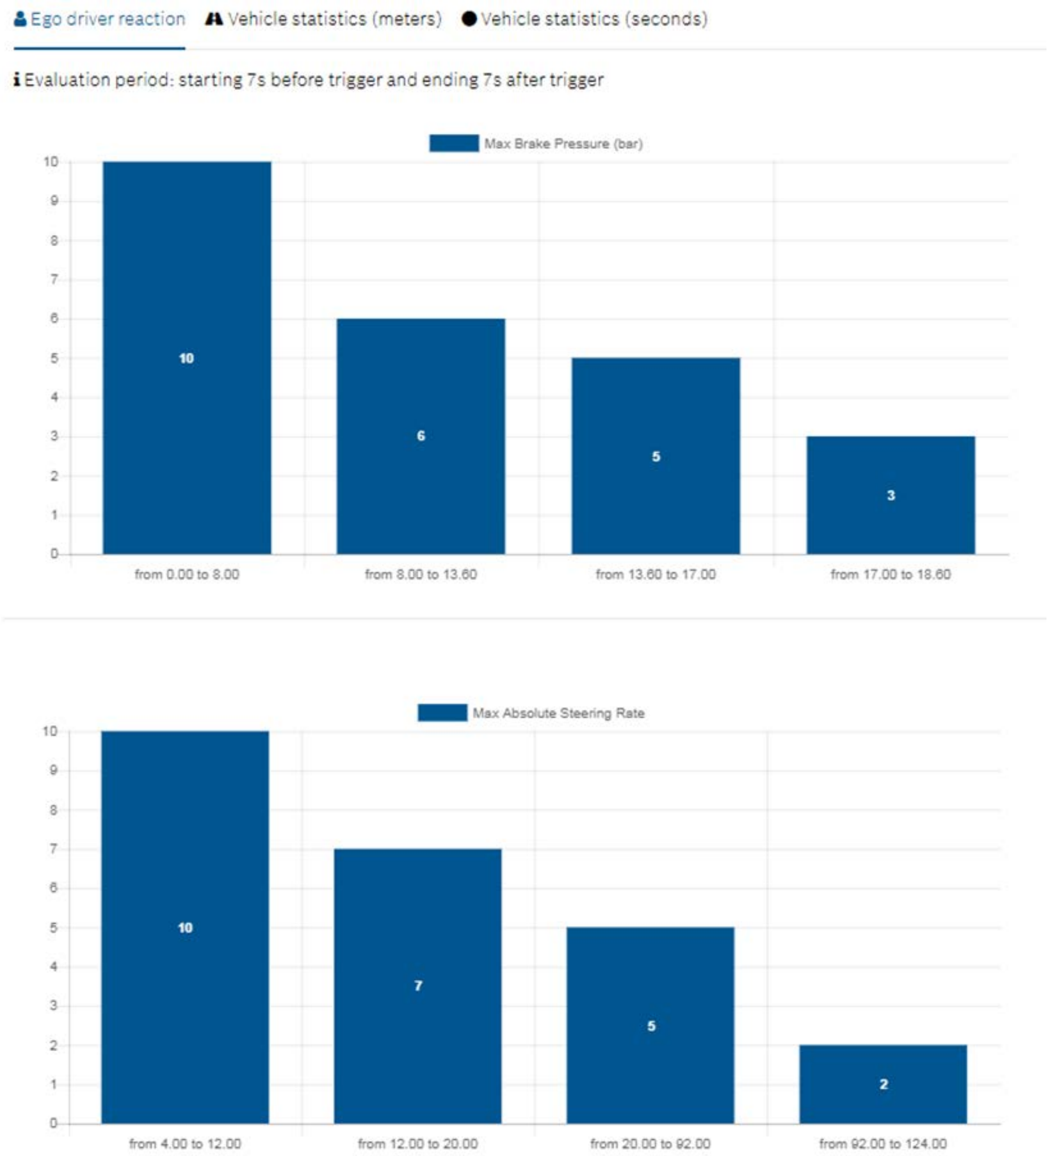

Figure 46 - Ego driver reaction statistics

Figure 47 below shows a scatter plot of TimeGap in seconds on the y-axis vs relative velocity of OBJ_01 on the x-axis. It can be observed that most of the activity cluster forms in the region of -3 to 3 on the x- axis and between 0-2 s on y-axis. This represents normal driving situations where surrounding drivers try to cut-in, in front of the ego vehicle. A few instances can be observed where the relative speed is beyond -10 which is quite high, possibly representing a sudden or unexpected manoeuvre of another vehicle pulling in front of the ego vehicle. These kinds of manoeuvres may force the ego vehicle driver to either apply maximum brake pressure or to change lane by steering at a high rate to allow the other driver enough space.

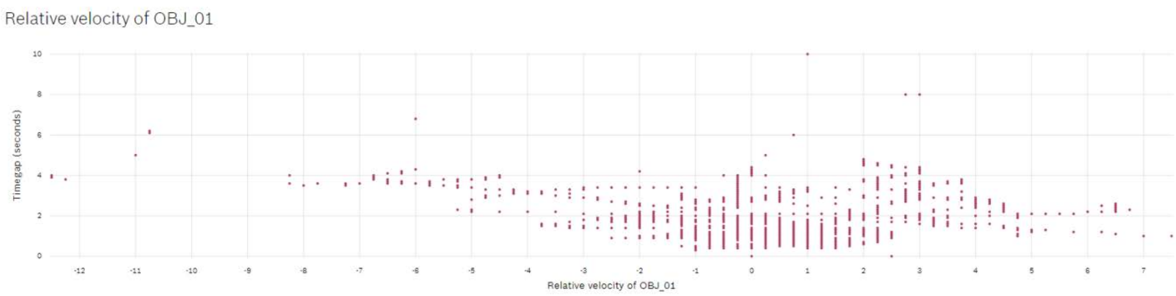

Figure 47 - Scatter plot of TimeGap (in seconds) vs relative velocity of OBJ_01

4.3.9.3 Discussion

Due to time and performance considerations, the trigger conditions in the vehicles (i.e. in the Flea 3 box and the CAN gateway) were kept simple. A first step in refinement is to consider which of the Cut-In sequences collected can be classified as real Cut-Ins and which represent false positives. The knowledge gained by such analysis could be used in the future to formulate rules to filter out these false positives. These rules could then be included in the sFDE back-end to enable more accurate statistical analysis.

Some of the possible situations where the similarity in dynamics of the vehicles involved in Cut-Ins can lead to false positives are as follows:

- Jumping back and forth of the signal OBJ_01_nr between two vehicle IDs due to probability based lane estimation. This behaviour is mainly observed in situations where the vehicle is driving at greater distances from others, or when the road includes curves.
- As described above, the ego vehicle cutting out can show similar behaviour in the vehicle IDs to a Cut-In. These situations can in the future potentially be filtered out using the steering wheel angle, noting that Cut-Ins could develop a critical steering response from the ego driver which would need to be captured.

Since the MOVE_UK fleets have only recently been deployed with the Cut-In trigger, the above statistics arise from a relatively small sample of Cut-In events. Hence, drawing any firm conclusion towards the parameters for a safe Cut-In, or even to say which event is a true Cut-In or a cut-out, would be premature. Collection of such events and using big data analysis techniques to process, analyse and eliminate any false positives to ascertain true Cut-In events within the dataset is underway and results will be available in Phase 3 of the project.

4.4 Use case Lead Vehicle Statistics (LVS)

4.4.1 Purpose

LVS started from the requirement to understand better how drivers normally follow other vehicles. Conditions of particular interest are the gaps maintained between vehicles, including at what speeds and their duration, as well as what causes the ego driver to stop following the lead vehicle. This can help to understand how to make autonomous driving more comfortable.

This information is primarily of importance for motorway driving and will, over time, also provide data for ongoing historical comparison (without complicating lower speed motorway driving with urban type driving). Ideally this will give a better understanding of behaviours going into, during, and coming out of traffic jams on motorways. These situations will be of assistance to other Consortium partner projects to help define parameters for safe and 'human-like' level 3 and 4 features.

4.4.2 Design

Unlike the CIN use case, where a trigger was designed and a new stream of dense CAN data was collected for a limited time (30 seconds) at higher frequency (10 Hz), the LVS use case serves as a proof of concept for using existing continuous 1 Hz CAN data for statistical evaluation.

The collected data is filtered in sFDE 'LVS View' using both fixed pre-conditions and configurable conditions. In this use case the following fixed pre-conditions were chosen:

- OBJ_01 available/detected, i.e. the radar sensor assigns a vehicle in front as target object 1,
- OBJ_01_nr constant, i.e. no Vehicle ID change, thus excluding Cut-Ins by other vehicles and cut-outs by the ego vehicle
- TimeGap to OBJ_01 below 4 s, i.e. the vehicle in front is not too far away.

The configurable conditions include the selections of the date range, the vehicle from which the data is collected (i.e. all, or a specific vehicle), the road type and the speed range. Figure 48 displays the different configurable conditions for LVS in sFDE. For road type, selection is possible between road type class 1 (motorways) and road type class 2 (main arterial roads, potentially in urban context), or both. Speed range includes the following options:

- 0 - 65 kph (traffic jams)
- 0 - 37.5 kph (low velocities)
- 37.5 - 65 kph (traffic jams excluding low velocities)
- 65 - 120 kph (unrestricted motorway travelling)

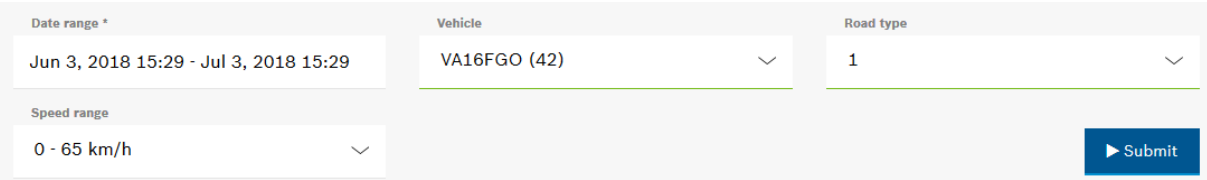

Figure 48 - LVS conditions selection in sFDE 'LVS View'

By selecting the conditions, only the 'snippets' of data (from all tours/journeys) for which the chosen conditions are fulfilled will be considered for statistical evaluation (i.e. not the whole continuous data log), as illustrated in Figure 49.

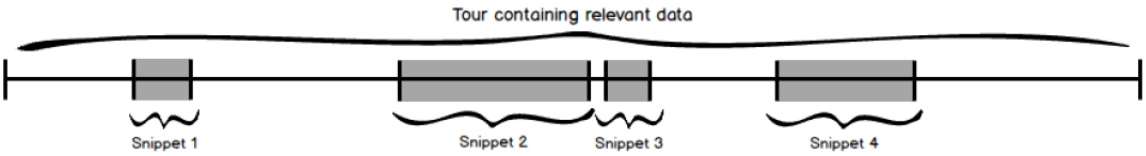

Figure 49 - Snippets of data resulting from querying with specific conditions

The following paragraphs represents all statistic types and the display options available for each.

In the context of drivers following other vehicles, it is of particular interest to gain meaningful statistics from the distances which the drivers keep relative to the speed at which they travel. For this purpose, a dedicated view in sFDE is provided. It consists of a heat map of the distance-speed pairs of all calculated snippets based on the selected conditions. Here, the emerging distribution highlights the behaviour of the majority versus outliers by clustering the data, as depicted in Figure 50.

Offline, the MOVE_UK data analytics team also looked into the possibility of a scatter plot, which is shown in Figure 51. This view plots the distance-speed pairs individually as points, resulting in more resolution with exact values while losing the cluster characteristics. Due to performance reasons, it was decided to not include this display method into sFDE, but keep it as an offline analysis tool.

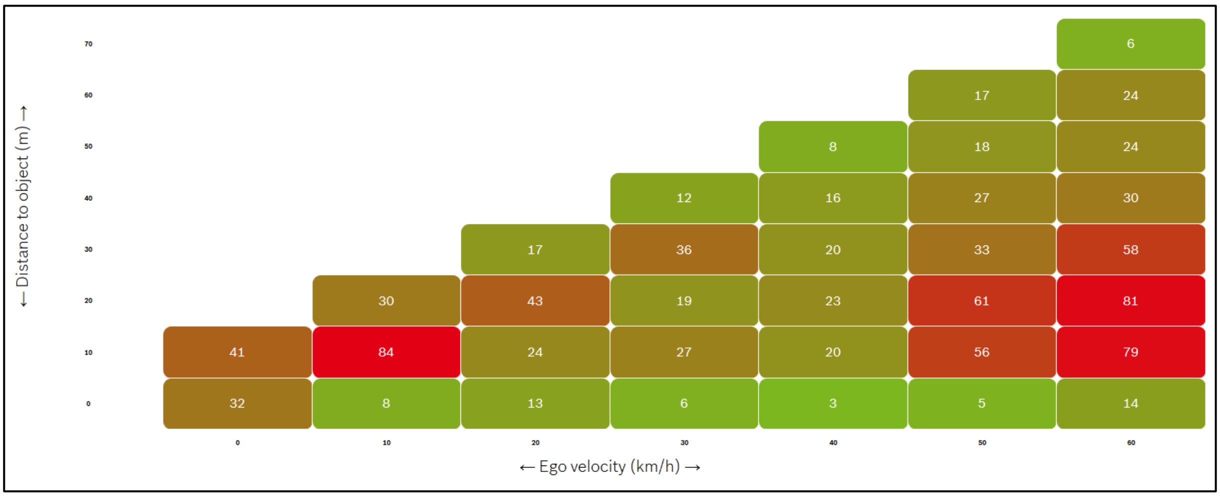

Figure 50 - Heat map: distance vs. speed in sFDE 'LVS View'

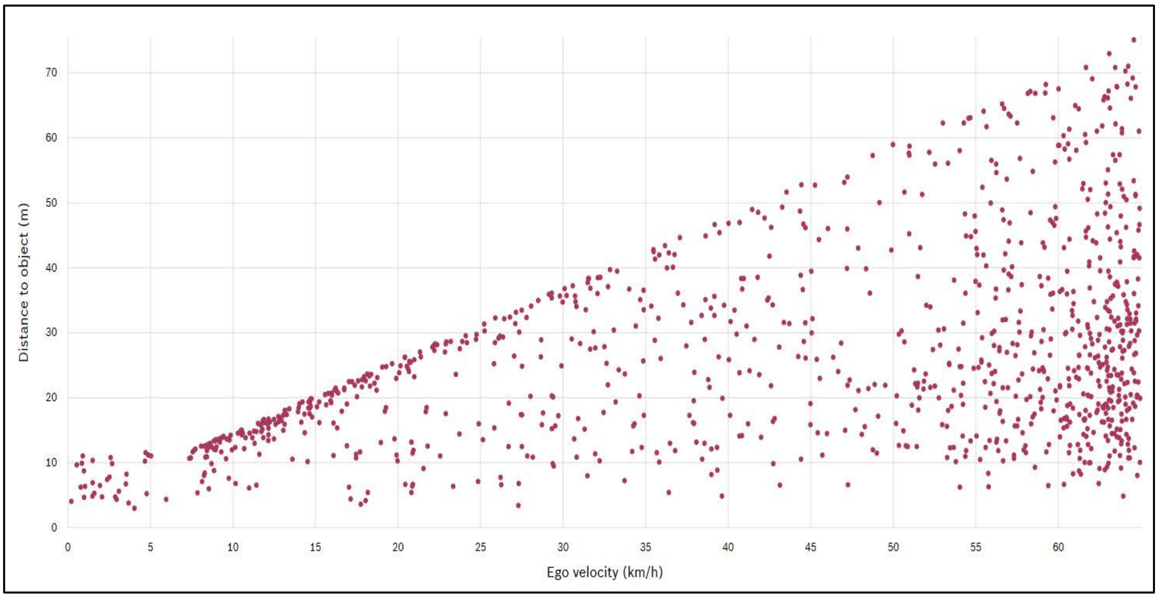

Figure 51 - Scatter plot: distance vs. speed created offline

In order to gain an insight into the length of time for which drivers follow other vehicles on large roads, the snippet length distribution was the second statistic to be examined. Analogous to the distance vs. speed distribution, two views are provided in sFDE to visualise this statistical data. One of these presentations consists of bar charts which count how many snippets of each length (in seconds) are available based on the chosen conditions, as depicted in Figure 52. The second view displays the same information in a view where quantity is represented by tile size, as visible in Figure 53.

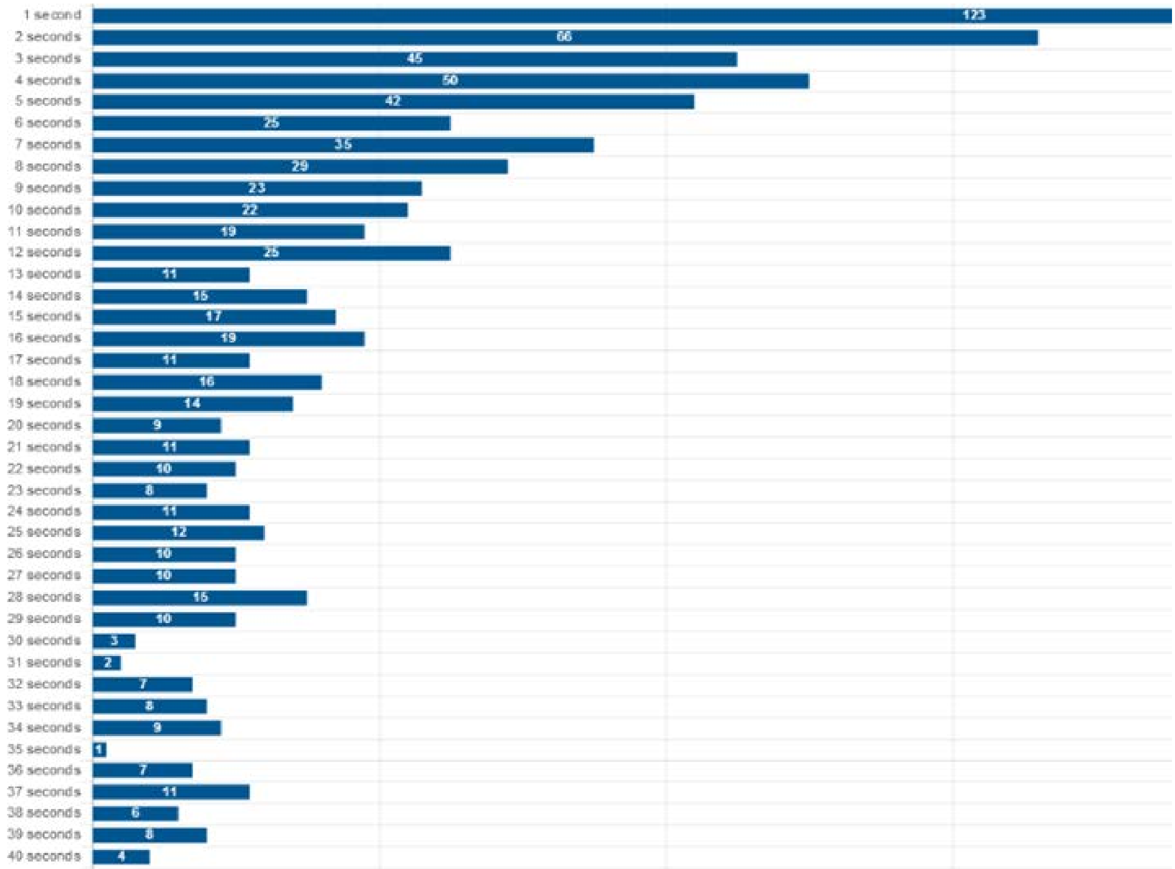

Figure 52 - Bar chart: snippet length in sFDE 'LVS View'

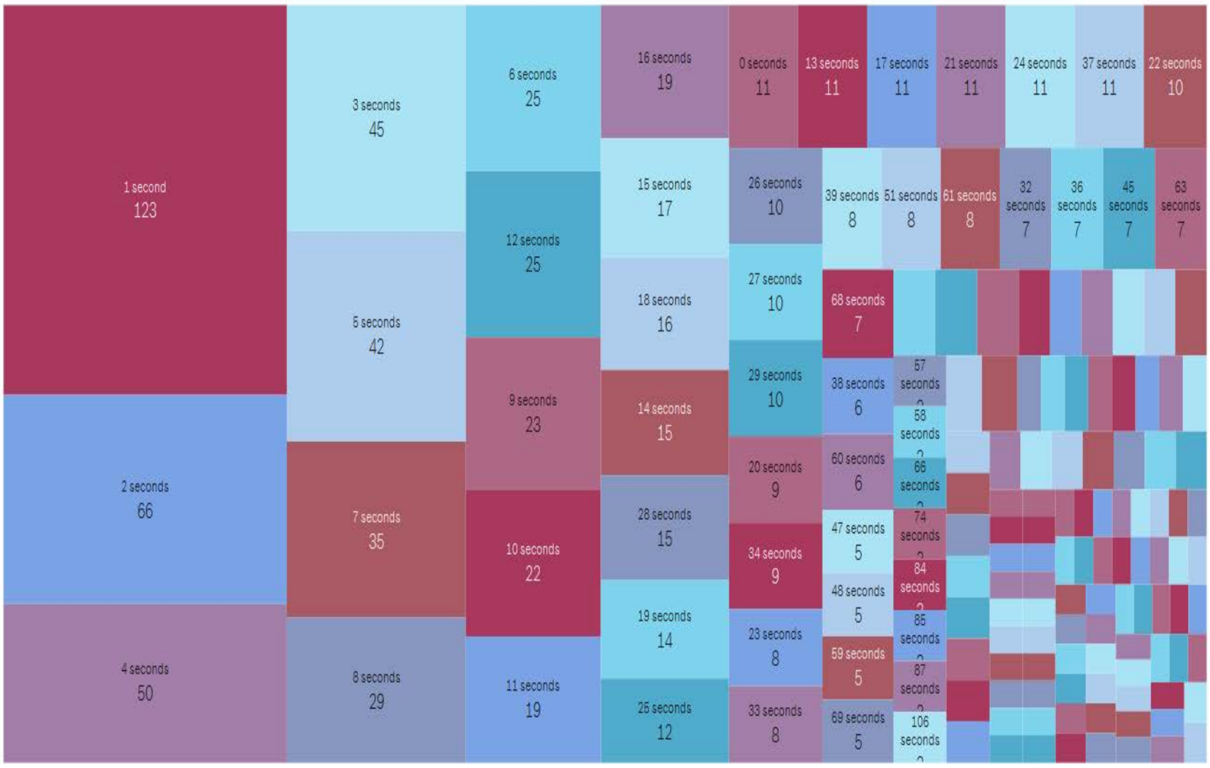

Figure 53 - Tile chart: snippet length in sFDE 'LVS View'

Last but not least, the LVS use case considers the exit reason i.e. which condition was violated and caused a snippet to end. This allows conclusions to be made about driver behaviour, as well as which pre-conditions might not be well-defined. For the LVS use case, the chosen possible exit reasons are the following:

- road type change, i.e. the ego vehicle changes from driving on road type 1 to 2 or vice versa
- road type loss, i.e. the ego vehicle changes from driving on the selected road type to a road type other than 1 or 2 (i.e. to smaller roads)
- object change, i.e. OBJ_01_nr changes to another valid vehicle ID
- object loss, i.e. no target object OBJ_01 is detected
- distance too big, i.e. TimeGap to OBJ_01 is bigger than 4 s
- over maximum speed, i.e. the selected speed range is exceeded
- under minimum speed, i.e. the ego vehicle speed is below the selected speed range, which is not applicable for 0 - 65 kph or 0 - 37.5 kph

To visualise this data, two views are again provided: a bar chart (Figure 54) and a display of percentages (Figure 55).

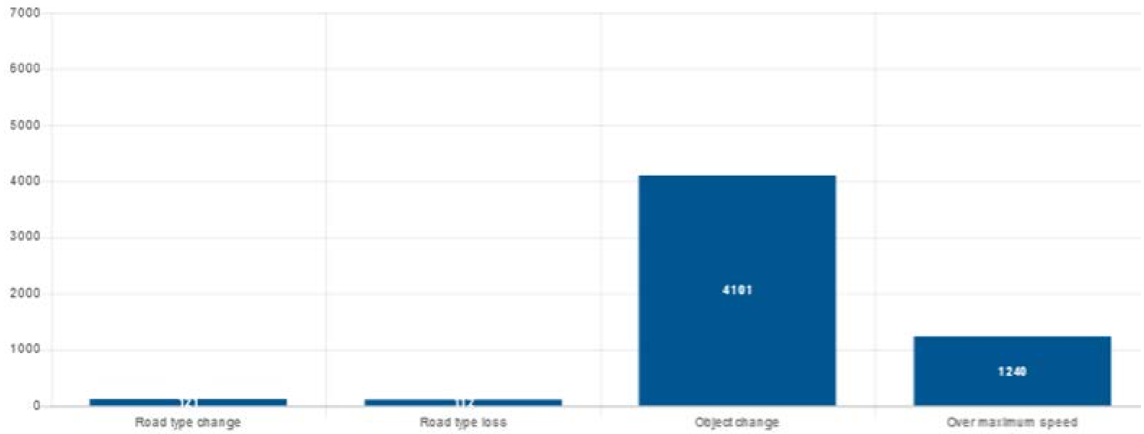

Figure 54 - Bar chart: Exit reason in sFDE 'LVS View'

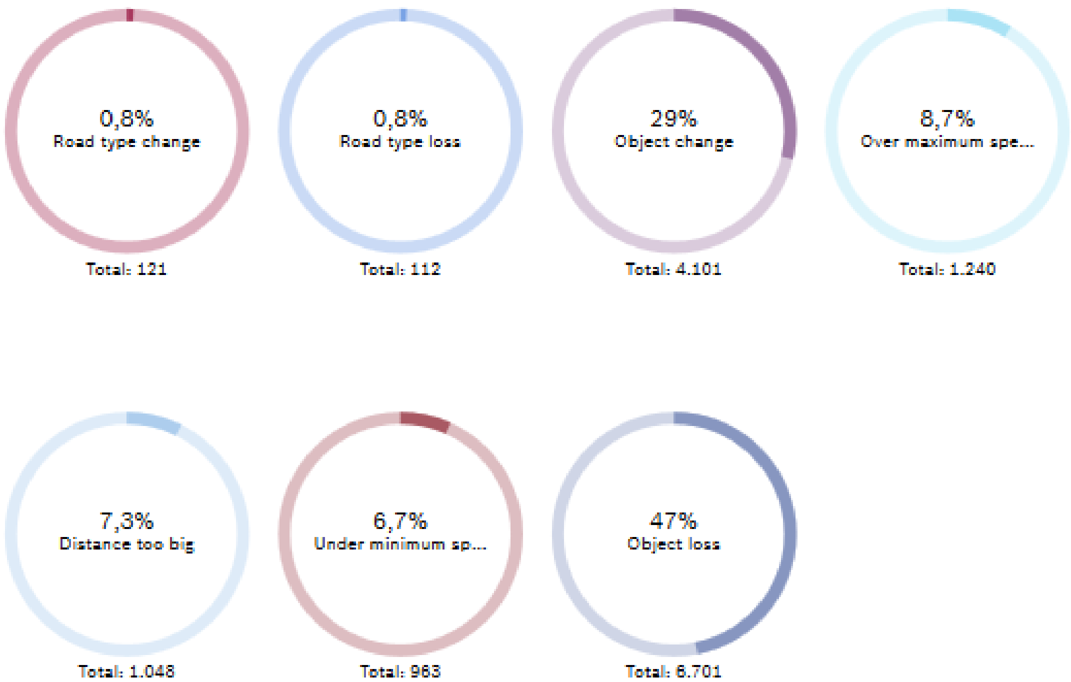

Figure 55 - Percentage view: exit reason in sFDE 'LVS View'

4.4.3 Synthesis and analysis

The LVS View in sFDE is still under development. However, a preliminary analysis is already possible for both improvement of implementation and to start to draw conclusions. The pre-condition of "TimeGap to OBJ_01 below 4 s" is illustrated especially well in a scatter plot of distance vs. ego velocity as a crisp cut-off line. In the heat map, a cluster emerges at higher velocities at the upper end of the range (55-65 kph) with distances between 10-35 m.

Regarding the snippet length distribution, a lot of very short snippets are visible. In the tile chart it can be concluded that approximately a third of the data lasts less than 8 s.

Analysis of the exit reasons shows that the majority of snippets ended due to object change or loss.

4.4.4 Discussion

The LVS use case has shown that a statistical analysis on the basis of the previously captured 1 Hz data is possible by just defining search criteria. Preliminary conclusions can also able to drawn as follows.

Emerging clusters in distance-speed distribution can be used to deduce driving behaviour regarding the size of gap left between cars by most drivers. On the other hand, outliers in the point plot can inform about extreme cases that have to be anticipated.

It is also evident that some refinement of the existing pre-conditions would be beneficial in order to obtain improved statistics from which even more relevant conclusions can be drawn; this is something that would need to be done following the end of this project. For example, it was noticeable that the amount of really short snippets does not tally with one's own experience while driving. This could be caused both by the formulation of the pre-conditions and the data quality. All pre-conditions have hard limits, e.g. if the speed of the ego vehicle is in a margin around the velocity conditions, each slight crossing causes a new snippet. For data quality, especially at larger distances, the radar sensor can briefly drop the object as target due to the probability-based calculation, also causing a new snippet. The assumption that this characteristic might erroneously cause some of the large number of short snippets is also supported by the fact that the most common exit reason is object change or loss. However, without further analysis it is not possible to conclude that the driver behaviour is actually such that lane changing happens all the time. In both cases, more computationally expensive methods such as hysteresis might prove necessary, which exceeds the scope of this use case as a proof of concept.

4.5 Use case Telematics – Phase 2

4.5.1 Purpose

Current commercial understanding of driver risk is derived entirely from vehicle telemetry. An objective of MOVE_UK is to reach an enhanced understanding of behaviour from richer and more comprehensive mobility data. The perpetual aim of risk estimation research is to identify new clusters of behaviour and factors that can aid prediction for understanding the likelihood and outcomes of incidents or real risk scenarios. In other words, existing telematics are used to seek new ways to distinguish journeys that appear very similar, but which actually present divergent levels of risk when using additional data.

From the early camera data in Phase 1 of MOVE_UK a strong need was identified to gain increased visibility around the vehicle. Detailed information about the behaviour and positioning of surrounding vehicles from camera and Radar systems has a revolutionary potential for EDR and risk compared to observation of telemetry alone. In considering the full range of potential camera data, line detections for example proved to be a less reliable feature owing to intermittent obscuration by parked vehicles, other obstructions, and variation in the lighting and weather conditions. A generalised approach to single event reconstruction or analysis cannot easily make use of data with such volatility. Accordingly, 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.

4.5.2 Design

Initial experiments were designed to encompass data from all journeys. All data obtained during the majority of Phase 2 (December 1st 2017 to June 29th 2018) was collated and assessed in terms of the available parameters. Other vehicles identified in the radar data are expressed in terms of an assigned

Vehicle ID, tracking its movements relative to each MOVE_UK vehicle. This Vehicle ID enables the tracking of movements between lanes, which are then tidied in terms of the total observation of that given vehicle as it moves across lanes.

4.5.3 Synthesis and analysis

For the first time, an evaluation is made of what constitutes the nominal relative position for surrounding vehicles, and further aggregate this by other parameters of interest. For example, Figure 56 shows the contrasting Radar positions for vehicles in positions 1-4 (i.e. OBJ_01 to OBJ_04) between when the MOVE_UK vehicle is moving at varying speeds.

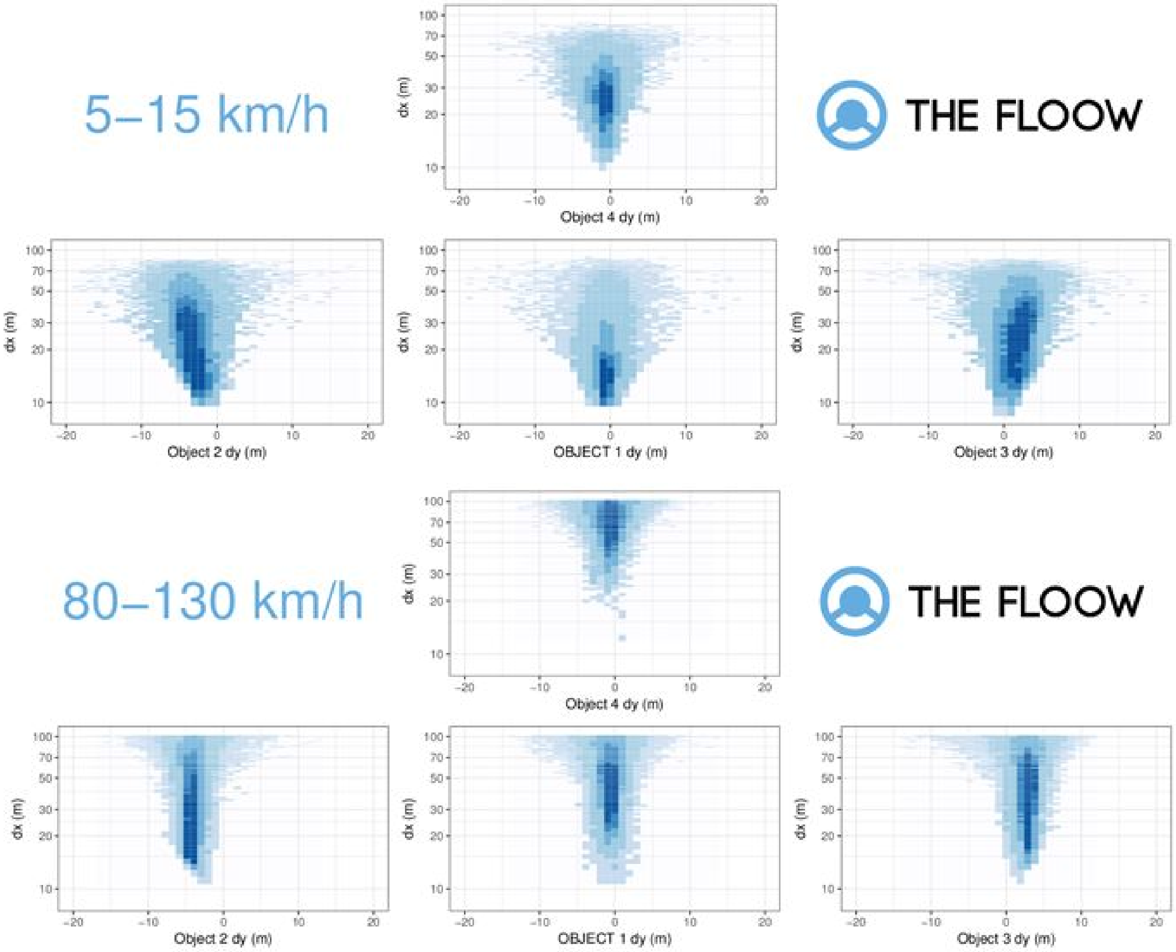

Figure 56 - Probability density contours showing MOVE_UK data marginalised over Radar object position (distance in front dx, distance to sides dy) for different vehicle speeds. Increasing depth of shade indicates progression from 99% to 50% confidence regions.

Object 3 demonstrates that the spatial distribution of leading vehicles is far narrower when the MOVE_UK vehicle is travelling at high speed. A more orderly, processional behaviour of vehicles is to be expected at greater velocities, when lateral movements are less viable for the safe driver. The true value of this positional information is that extreme behaviour can be directly identified from knowledge of that which constitutes normal behaviour. For example, 99% of the time when the MOVE_UK vehicle is travelling faster than 80 kph the relative distance to the lead vehicle is greater than 10 m, so lead vehicles consistently observed at shorter distances could be highly indicative of risky driving practices. This simple understanding of position opens up various further avenues for exploration, such as re-framing the distance to the lead vehicle in terms of a time-to-impact parameter, which was particularly useful for supporting the CIN use case. For example, we can calculate the time-to-impact for leading vehicles over the first instant that they are observed by the Radar, which with large volumes of data can be used to form a probabilistic picture of the risk posed by leading vehicles.

An appraisal was made of the ability of Radar and video camera systems to track vehicles for extended periods and the consistency between the two sensors. The capability of the camera system to identify and track many more types of surrounding features and objects means that there is often information available from the camera when there are no detections in the Radar data. Indeed, it was found that there are optical detections for 85% of the time when the Radar assigns no target objects. However, optical detections of cars or trucks exist for only a small percentage of the time for which there are no Radar detections. While this could indicate deficiency on the part of the Radar system, the wider field-of-view and more elevated position of the camera will mean that a greater portion of the surrounding area is visible to the optical system.

An improved comparison of sensor data from the observed properties of the neighbouring vehicles can be achieved and the question asked to what extent the two systems agree about the movement of traffic in front. The spatial positions of simultaneous detections are compared in order to find a mutual set of objects. The result is that the lead vehicle, as designated by the camera, is seen as the lead vehicle by the Radar system 85% of the time. The radar assigns the optical lead vehicle to a neighbouring lane 15% of the time, and detects an object at closer quarters for fewer than 0.5% of joint detections.

Assessing changing risk of using forward-facing sensors also involves understanding the variation of visibility as a function of local geography. Initial steps have been made into defining a score to identify roads that are particularly problematic for vehicle sensors when the forward-facing field-of-view is limited. An example of this analysis can be seen in Figure 57.

Figure 57 - Risk arising from use of forward-facing sensors with restricted field-of-view, mapped onto Greenwich and surrounding areas. Roads found to be of higher risk appear more red.

4.5.4 Discussion

The ultimate aim of the telematics use case is to construct a method through which a comprehensive and synoptic description of risk can be computed. To achieve this, the aim is to combine knowledge of telemetry, driver behaviour, the conduct of surrounding vehicles and geographical factors into a single parameterisation. However, to obtain such an understanding will require continued studies to obtain more data using a greater number of vehicles, ideally with the ability to distinguish between specific drivers.

Data collected during Phase 2 can be used to reach a good understanding of vehicles-in-front for incorporating into models of risk. Knowledge of the traffic context in which the vehicle is driving is clearly important, letting us observe not just how the driver is driving, but whether the observed behaviour is normal.

During Phase 3, information will be combined from the additional corner Radar systems into the existing analysis; working towards a refined model of risk enabling a 360-degree understanding of proximity. The new Radars will provide positional information beyond that of the headway distance to leading vehicles and will also provide data from all objects observed by the sensors, not just those assigned to *target* status by the internal software, thus giving a more complete view of the entire surroundings of a vehicle in operation. The enriched data available in Phase 3 will grant us an enhanced view of vehicles during autonomous operation and will be particularly valuable for the study of the DHN, ARB and CIN use cases.

5 Conclusions

As MOVE_UK heads towards the end of the second of three project phases, it can be reported that significant further steps on the way to achieving the project objectives have been completed in Phase 2.

A forward facing radar sensor was added to each of the trial vehicles to enhance their ability to detect objects in front of them and allow experimentation with new use cases which rely on radar based data. Corresponding changes to the system setup were successfully made to support the collection of the additional radar data and several new views were added to the sFDE web interface to allow visualisation and analysis of data relating to the new use cases developed in Phase 2; namely ARB, CIN, and LVS.

Throughout Phase 2, data corresponding to the use cases developed in Phase 1 continued to be collected. This data included an updated selection of continuously recorded CAN signals which was produced from experience gained in Phase 1. A new team of data scientists was created and, with the support of sensor experts, a much more detailed analysis of the Phase 1 use case data was carried out. In the case of camera based AEB, a number of highly relevant video sequences were recorded, which when re-simulated using production camera software proved to be true positive activations, meaning that these events would have activated the production AEB system. A large number of DHB events were also recorded and used to identify a distribution of reasons for harsh barking manoeuvres and to investigate human braking behaviour using cluster analysis. The number of traffic sign detections recorded by the TSR system rose from 30,000 in Phase 1 to over 85,000 towards the end of Phase 2, and considerable effort was spent identifying traffic sign clusters from the data collected in sFDE and using these clusters to conduct statistical analysis and an in-depth site investigation. The outcome of this investigation will be followed up in Phase 3 of the project.

The three new use cases added in Phase 2 have demonstrated further the capabilities identified at the beginning of the project and have also provided a better understanding of the behaviour of surrounding vehicles which could be fundamental for the development of safe and ‘human-like’ level 3 and 4 automated driving features. The ARB use case, which builds heavily on the work started in Phase 1 with AEB, has proved that subcritical situations can be successfully identified using radar and captured using the connected validation method in order to identify real-world false positive situations. ARB has resulted in the capture and analysis of a large number of relevant events. The CIN use case, which has involved the most work in Phase 2, has led to the development of a useful trigger for capturing ‘cut-in’ events. Initial analysis of the limited number of events captured so far has begun and development of big data analysis techniques to process, analyse and eliminate any false positives is underway. Finally, the LVS use case has demonstrated that useful statistics about vehicle behaviour can be developed from relatively low frequency (1 Hz) continuous CAN data; specifically, the analysis of the data allowed for a better understanding of the gaps maintained between the ego and lead vehicle, particularly in traffic jam situations. Further capture and analysis of ARB and CIN events, as well as LVS data, will continue in Phase 3.

Regarding risk assessment and incident reconstruction methods, the additional telematics and continuous CAN data collected in Phase 2 (which included improved camera object data) was used to extend the EDR and risk analysis completed in Phase 1 of the project. As in Phase 1, the highest value signals were found to be those relating to positional information of other road users, the additional radar data allowing for more precise mapping of the positions and speeds of vehicles moving in the same direction as the ego vehicle, within the field of vision of the radar. The new radar data collected in Phase 2 was also used to gain a better understanding of what constitutes normal and extreme behaviour of vehicles travelling in front of the ego vehicle, something which can be incorporated into risk models. Phase 3 work will focus on combining additional sensor information into the existing analysis in order to refine the risk model even further.

Phase 2 of the project ends with a sound basis in place for the addition of corner radars to two of the MOVE_UK vehicles in Phase 3. This will provide a 360° understanding of surrounding objects which is vital for the successful development and validation of more complex automated driving features and for meeting MOVE_UK's ultimate goal of accelerating the development, market readiness and deployment of ADS using connected validation and big data analysis.

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
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
TSR	Traffic Sign Recognition
UI	User Interface