Stopped vehicle Hazards – Avoidance, Detection, And Response (SHADAR)

Stopped vehicle detection and reporting

research results

Deliverable D5.1 Version 1.1

4 November 2022

CEDR Call 2019: Safe Smart Highways

Stopped vehicle Hazards – Avoidance, Detection, And Response (SHADAR)

D5.1 Stopped vehicle detection and reporting: research results

Due date of deliverable (at project start): January 2022

Actual submission date (v1.0): 14 July 2022

Submission date of this revision: 4 November 2022

Start date of project: 1/10/2020 End date of project: 16/12/2022

Author(s):

Ian Cornwell

Steven Boerma

Giovanni Huisken

Shayan Afshar

Danny Woolard

Alan Gentle

Andy Graham

Andy Rooke

Conférence Européenne des Directeurs des Routes
Conference of European Directors of Roads

Executive summary

The project “SHADAR” (Stopped vehicle Hazards – Avoidance, Detection, And Response) addresses the objective of “Preventing collisions with stopped vehicles in a live traffic lane”. Stopped vehicles on the highway network present a significant hazard with an impact on safety and the economy.

The SHADAR project aims to help reduce the risk of collisions with stopped vehicles on highway networks by improving the detection, reporting and management of these hazards. This is accomplished by establishing and sharing knowledge on current effective practices, and by researching potential improvements that can advance the current state of practice. This research proceeds in three inter-related strands – on detection and reporting technology, road user behaviour, and response from national road managers. The project identifies the state-of-the-art and then researches possible improvements.

This report is the output from SHADAR work package 5, which builds on the state-of-the-art reviews and considers potential improvements in stopped vehicle detection. The report considers potential improvements using radar, eCall and other connected vehicle sources, drones, human reporting, the fusion of multiple sources, and how the outputs can best be presented to traffic operators and technology managers.

The role of eCall in stopped vehicle detection

eCall is fitted to all newly type approved cars and light vans since 2018, with more types to come. It can be automatically activated, typically by air bag; or manually, by pressing an SOS button. When activated, a voice call is set up to a Public Safety Answering Point (PSAP), where an operator asks if assistance is required from the emergency services (police, fire or ambulance). This is similar to traditional 112 calls. However, eCall also contains a rapidly sent data packet, or Minimum Set of Data (MSD) that has vehicle ID, location and confidence in location, direction of travel, the number of occupants, fuel type and whether the alert was manual or automatic

eCall volumes are now increasing, with over 10,000 calls per month in the UK. But currently few NRAs are making the most of it. eCall is one of the first instances of data in traffic operations sourced directly from vehicles so if NRAs cannot make use of this data, it is not a strong foundation for more advanced use.

eCall could be a valuable tool to augment other detection technologies. So this report develops ‘best practice’ methods of deployment. It highlights opportunities to improve detection and advises how an NRA can exploit the MSD, without impacting emergency service responses and by fusing it with other data, create a new sensor for stopped vehicles.

To make the most of eCall, the MSD can be passed electronically through to responders to alert them to stopped vehicles. This will provide alerts in seconds rather than minutes, an order of magnitude faster than voice alone. The MSD does not replace the voice channels. Instead, it provides an “early warning”.

To assess the value of eCall data for stopped vehicle detection, we investigated reliability of data, false alarm rates, and accuracy. We found that automatic activations generated when a suitable vehicle condition occurs, such as high deceleration or airbag deployment, indicate a very strong likelihood of a real incident, and the MSD directly links the vehicle to the incident. The false alarm rate for these is very low as there is no human involvement. There are many scenarios where someone could manually press the button, such as to report a breakdown but also to demonstrate eCall in a car showroom. Hence manual activations have less confidence and need to be managed and filtered. And as eCall activations can originate from

any location – motorways, fields from 4X4s, country lanes, MSDs must first be geographically filtered for each responder's area.

We describe specific methods for use of eCall in stopped vehicle detection. These methods form a workflow process in an eCall process engine to filter, enhance, profile and forward processed eCall data. Enhancement of the MSD with additional data provides an opportunity to improve operational value, complementing infrastructure-based SVD alerts and eCall voice.

eCall complements existing SVD methods to create a much wider coverage of the road network. So, if MSDs are processed together with other SVD alerts, a much richer picture can emerge. For example, a manual MSD has a low confidence of an incident, but a manual MSD with a radar alert provides a high confidence and provides vehicle details and location details not available to radar.

Potential for improved radar detection

We explored several ways in which additional information could be provided by rotating radar systems to enhance the overall detection capability.

- The radar's azimuth resolution should support the determination of the lane of the stopped vehicle, not for the full operational range of the radar, but for approximately 150m of range. A small set of experimental results support this.
- Analysis of radar data for stopped vehicle events shows occurences of pre-stop and post-stop traffic speed reductions and queues, but so far not with sufficient volumes to clearly demonstrate correlation that could be used in stopped vehicle detection.
- Radar can provide a limited but potentially useful classification of vehicle type.
- Radar can detect and track pedestrians, and so could attach additional information about increased hazard level along with alerts.

Other connected vehicle sources

eCall is not the only source for stopped vehicle alerts from connected vehicles. Methods using vehicle sensors have the potential for fastest detection and the richest supporting information but have low levels of penetration in the vehicle fleet, although these are growing. Standardised cooperative ITS capabilities include stationary vehicle identification and warning, but still have low uptake beyond pilot projects. Several data providers offer traffic data commercially, and now the Data for Road Safety initiative aims to make safety-related traffic information available for all road users in Europe.

Recent changes in provision of location data reduce any potential of textual social media, but the traffic-specific application Waze has potential. Analysis of a large Waze dataset from the Netherlands revealed that reports of the incident are faster than the national registration in the Netherlands and they cover a much larger road network.

Aerial imagery

Images from unmanned aerial vehicles and satellites could both provide more accurate information on location and lane information and also on vehicle type but have practical disadvantages. Satellites do not make sufficiently frequent passes for useful real-time stopped vehicle detection coverage. Aerial vehicles are expensive and suffer from weather conditions, reducing the effective range. They may become feasible for targeted applications.

Data fusion

SHADAR report D2.1 showed that there are several different technologies for stopped vehicle detection, with varying performance on important metrics such as detection rate, false alarm rate, independence from environmental conditions, coverage of the road network, precision of location, timeliness and data content. Analysis suggests that every source is outperformed by another source on at least one metric. This suggests potential for data fusion to achieve better overall performance than can be provided by any one source.

Machine learning has become popular through success in many contexts. Machine learning of raw sensor data holds technical promise but may currently be working against the market in which technology providers aim to optimise their own detection offerings rather than provide raw data into a larger fusion system. A more practical route today for a roads authority is to fuse the outputs from these technology providers. This could also be done by machine learning, but SHADAR explored a simpler statistical method which may give equivalent value.

We show how the performance of a stopped vehicle detection fusion system can be determined, using probability theory. Such analysis of fusion schemes should help a road authority understand which fusion rules are appropriate, which data sources should be integrated, and what performance may be achieved. Better performance comes by fusing sources that behave independently. Sources with entirely different technical basis (such as eCall compared to radar) are likely to show high independence, while sources with some similarity (such as two methods detecting electromagnetic reflections) are unlikely to be entirely independent and may produce less improvement when fused.

Choosing a data fusion regime allows a choice between optimising the detection rate and optimising the false alarm rate. Data fusion can also provide a confidence level for every alert.

We applied the statistical data fusion techniques retrospectively to real stopped vehicle alert data from a highway in Europe, where two different technologies had been used in similar locations for three months. We did not have complete ground truth data, but we had a form of manual verification for all alerts. We had to make certain assumptions because our study was performed retrospectively rather than being built into the design and operation of the detection systems. The study showed that each source was missing true stopped vehicle events that were detected by the other source. Even without the expense of a full ground truth study with constant human vigilance, the analysis of two sources together provides knowledge about the performance of each source which was not otherwise apparent. Using these sources together in a data fusion system would have increased the detection rate and reduced the false alarm rate when compared with using a single source.

The fusion of multiple sources can be combined with operational user interface developments to help avoid additional workload for traffic management operators. This report shows user interfaces in which alerts can be grouped and the calculated confidence levels shown to help prioritise the workload. (The impact of such features on operational response is further explored in SHADAR report D6.1.) The routing of all sources through a data fusion system also enables reporting on the detection performance of each source, and this report shows examples of reports that could be used to support decisions on continued investment.

Human reporting behaviour

Interviews and experiments suggest that public behaviour for alerting police or road authorities is not uniform. Awareness of how to act when a stopped vehicle/incident is encountered seems low and diverse. Road authorities, governments or traffic safety organisations can contribute by making the appropriate information publicly available. (Road user behaviour is further explored in the next two SHADAR reports.)

Table of contents

1 Introduction

1.1 Purpose and scope

The project "SHADAR" (Stopped vehicle Hazards – Avoidance, Detection, And Response) addresses the objective of "Preventing collisions with stopped vehicles in a live traffic lane". Stopped vehicles on the highway network present a significant hazard with an impact on safety and the economy.

The SHADAR project aims to help reduce the risk of collisions with stopped vehicles on highway networks through improved detection, reporting and management of these hazards. This is accomplished by establishing and sharing knowledge on current effective practices, and by researching potential improvements that can advance the current state of practice. This research proceeds in three inter-related strands – on detection and reporting technology, road user behaviour, and response from national road managers. The project identifies the state-of-the-art and then researches possible improvements.

This report is the output from SHADAR work package 5, which builds on the state-of-the-art reviews (Huisken et al, 2021) and considers potential improvements in stopped vehicle detection.

The main goal is to identify options for improving the time to detect, the reliability of detection, the kinds of information that can be gathered and reported, and how that is reported to traffic operators and policy makers.

This D5.1 report has a counterpart summary report D5.2, which has been edited to less than half the length of the present document. The present document contains additional explanations, examples, and relevant quotations and summaries from other relevant research, for those who wish to take more time to consider one of its topics.

1.2 Report structure

The work package included 6 sub-tasks that each produced an internal technical note; each was reviewed by the project team and combined in the present report. The sub-tasks were:

- How to harvest data from eCall (Chapter 2)
- Potential of improved radar detection (Chapter 3)
- Future/upcoming methods for stopped vehicle detection (Chapter 4)
- Data fusion (Chapter 5)
- Human behaviour for detection and validation section (Chapter 6)
- Reporting alerts and performance (Chapter 7)

Each chapter covers a specific area of potential improvement in stopped vehicle detection, and can be read individually. Chapter 8 combines provides a consolidated summary.

1.3 Characterisation of research ideas

This report covers diverse ideas. As a supplement to the research findings, to help to illustrate and relate the ideas in a common and intuitive way, this report uses a technique known as a "How-Now-Wow" matrix. Ideas are placed on 2 axes: originality and (in)feasibility. Each axis can be divided into two parts, making 4 cells, as illustrated in Figure 1:

- Feasible, not particularly innovative (designated "Now")
- Feasible, high level of originality/innovation (designated "Wow!")
- Difficult/infeasible, high level of originality/innovation (designated "How?")

- Difficult/infeasible, not particularly innovative (of no further interest)

The categorisations are not based on clearly defined objective results, but rather are the educated opinions of the project team after considering the research results.

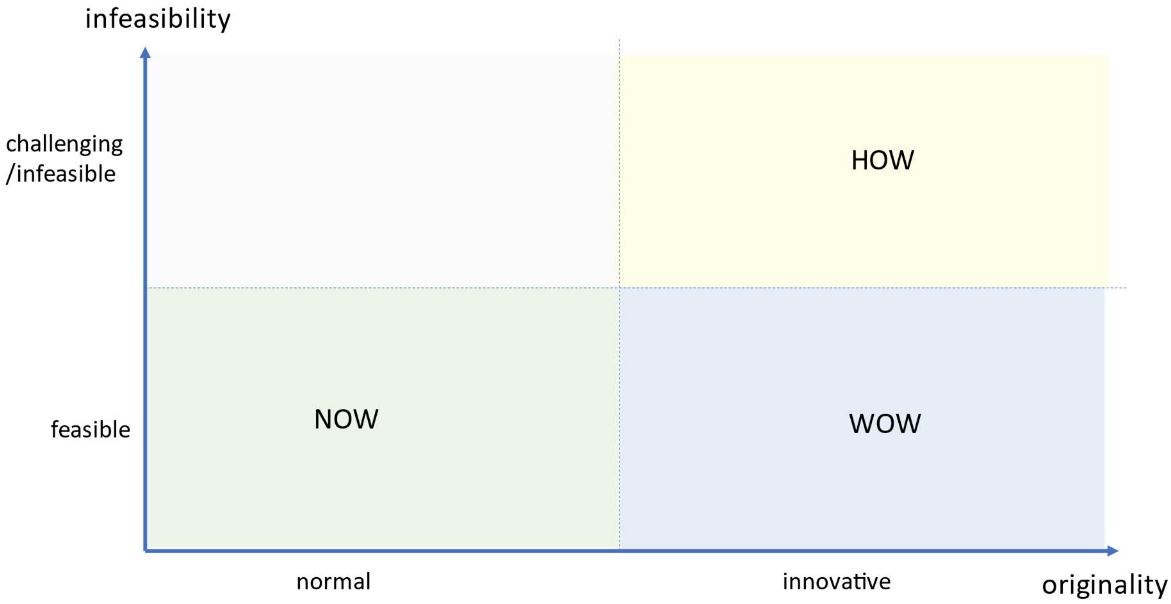

Figure 1 How-Now-Wow matrix used to characterise ideas

The How-Now-Wow matrix does not explicitly express the value in an idea (presence in NOW or WOW does not imply a recommendation to implement the idea), so we provide extra commentary in tables. Each research idea is also assigned a category representing the potential outcome: timeliness, reliability, accuracy, or information.

At the outset of this research, before the work described in the subsequent chapters had been conducted, the SHADAR project team held an initial brainstorm session to identify unmet needs for relevant stakeholders, then possible technical solutions were placed in a “How-Now-Wow” matrix. That early “How Now Wow” matrix is not itself a significant research result, but it is included in the Appendix. It helped to revalidate the scope of the work package and suggested significant potential in areas including eCall, connected vehicles, data fusion and increased integration of systems. Not every idea could be developed by the project. Each chapter concludes by placing the relevant ideas on a How-Now-Wow graph, and the concluding chapter provides a combined How-Now-Wow table to summarise research ideas in a single common way.

References in this report

This report follows the CEDR research report template in which all references appear at the end of the report. References follow the common academic referencing scheme known as Harvard, specifically the “Leeds Harvard” guidance from the University of Leeds. An exception is made for international and national official standards which are identified directly in the text by their number e.g. EN 50110-1. Inline hyperlinks to web pages are used as an alternative method where that enhances readability.

2 Harvesting data from eCall

2.1 The role of eCall in SVD

The deployment status of eCall in Europe varies across each member state. The implementation has been focused on providing 'crashed vehicle' detection and the voice channel for the emergency services. Research and knowledge of the deployed technology (Huisken et al, 2021) suggests that increasing use of the eCall data for detecting stopped vehicles could be a valuable tool to augment other SVD technologies.

This chapter develops 'best practice' methods of eCall deployment with a focus on use firstly of the voice element of the system operationally, both with and without the accompanying data, and then takes a detailed look at the data alone. It highlights valuable data such as vehicle identification number and last three GPS points. Finally, it identifies the opportunities to improve stopped vehicle detection both as a standalone method and in combination with other methods of stopped vehicle detection.

2.2 Introduction to eCall

eCall is a system fitted to all newly type-approved cars and light vans since 2018, with many more vehicle types to come. It is a system that:

- can be automatically activated, typically by air bag or other sensor activations; or
- manually, by pressing an SOS button

When activated:

- A voice call is set up to a Public Sector Answering Point (PSAP) where a trained operator asks if further emergency help is required. This is typically police, fire or ambulance but does not include NRA services such as breakdown or highway officers. This is very similar to traditional 112 calls from mobiles phones. However, in addition to the voice call the eCall contains a data packet.
- The data packet is rapidly sent containing a minimum set of digital data (eCall Minimum Set of Data (MSD)) that has:
    - o The vehicle ID as Vehicle Identification Number (VIN) number, which can be turned into its vehicle registration mark, vehicle model and type and colour using national databases such as the UK DVLA
    - o The location of the vehicle and confidence in that location (up to three recent GPS Traces can be sent)
    - o The number of occupants
    - o The fuel type
    - o Current direction the vehicle is facing
    - o Whether the alert was manual or automatic

Data and voice are sent together using an "in band modem" as a priority over any 2G/3G cellular network.

eCall volumes are now increasing, for example with over 9,000 calls per month in the UK. And as more and more vehicles become fitted, the volume of calls can only increase. In the UK, 90-95% of the sales of top 20 selling cars now have eCall, and this will increase as the Kia Sportage and VW Polo transition into new models with eCall.

eCall has developed from private services for premium vehicles to now include all new cars, but currently very few NRAs are:

- making the most of the voice content alone as an alternative to 112 calls and potentially emergency roadside telephones,
- using voice plus data, or
- extracting maximum value from the data as a digital feed separate from the voice, using its accuracy and timeliness, and rich content for NRA operations.

Section 2.3 describes how a NRA can make the most of eCall voice capability and adding MSD data to improve quality and richness for traffic operations.

Section 2.4 onwards describes how a NRA can exploit the MSD data alone, without impacting emergency service responses and by fusing it with other data create a new sensor for stopped vehicles.

eCall data is important as one of the first instances of data of vital use in traffic safety operations being sourced directly from vehicles. eCall is already in place. If NRAs cannot access and make use of this data then this does not set a strong precedent for more advanced connected vehicle data use. Finally, it is a free data set that must be made available to emergency responders by EU Directive, and is also compliant with GDPR.

2.3 Making the most of eCall voice, and voice plus MSD data

2.3.1 Advantage of voice eCall

eCall as a voice channel has many advantages over a hand-held or Bluetooth call from smartphones:

- It was designed for road safety by emergency practitioners, so is reliable and robust. For example, it provides location data in open GPS co-ordinates.
- Automatic activation will connect to the occupants even if they are injured or trapped, or don't have a mobile phone or their operator doesn't cover the area
- It is always a hands-free service for safety,
- The manual button can be used to report others' problems; for example, another vehicle on fire, or by the passenger if the driver is facing a medical emergency such as a heart attack.

With added MSD data it has highly reliable data that can be filtered to help NRA actions. MSD from the call, even just as meta data (data about the call) brings the additional benefits of:

- Time saving – as the details of the vehicle and its location do not have to be sought from the driver, who may not know where they are anyway (the so called "featureless highway")
- Accurate location to a few meters, rather than as drivers often describe their location as "on a motorway"
- Vehicle VIN and ID mean fewer transcription errors

Nevertheless, even just voice messaging alone is useful over a normal 112 call as it is a direct channel, especially when opened automatically after an airbag deployment. Note that in most newer vehicles, airbag deployment immobilizes the vehicle so whilst the occupants may not need for example medical care, the NRA may want to know about the vehicle to recover it, especially if stopped in a live lane.

2.3.2 Getting the eCall voice (or content of voice call) to the NRA

Challenges for eCall voice

Having a PSAP aligned with the NRA traffic operations team is key. There must be shared knowledge and understanding, perhaps via a knowledge Memorandum of Understanding, of how the link is made.

Where and who operates the PSAP/112 centre is a government decision, but how they pass on voice call and information, and to whom, would be key for the NRA to establish with the PSAP. What happens currently between the NRA and emergency services with 112 calls is a good model – how 112 derived voice information currently get into the NRA, if at all, is a good start.

If there is a clear model for how 112 calls (or the details of the call) are passed to an NRA, this is a good foundation for eCall. Areas to think about are:

- The NRA network definition - how will the 112/PSAP know the call is for the NRA (e.g., on a motorway) as opposed to in a road beside it?
- Calls for help not requiring emergency help may not reach the emergency services but may stop at the PSAP and go no further. Whilst “blue light” assistance may not be required, the NRA will certainly want to know especially if a vehicle is immobilized.
- Those 112 calls and eCall may go to a separate PSAP. There may be one national PSAP but many “level 2” PSAPs which may not tie up with where 112 calls are routed.
- Some NRAs support the emergency services in managing the scene and restoring normality of traffic flow. Early and accurate notification is paramount to an effective deployment.

Using a command-and-control system from the PSAP or 112 centre (if one exists) is an ideal way to transfer call details (rather than the voice) from PSAP and 112 centres to the NRA. In such a case, a key decision is whether the details are pushed by the PSAP, or the NRA has to request them from the PSAP.

The PSAP may not send the information to the NRA or other emergency service providers and may rely on the police (as is the case in the UK and most other European Member States, the notable exceptions are Czech Republic, Finland and France). This is a strategic decision for PSAP to tell police, other emergency services, NRA and other road authorities.

In using a voice plus meta data approach, transfer of data from the 112 command-and-control logs works best; transferring the actual voice call to the NRA is often not possible, as the communication architecture for eCall only envisioned that an eCall would only need to go as far as the level 2 PSAP (dispatch of rescue services). Email of log entries to the NRA is the usual default.

Recommendations for best practice eCall voice

Below are some recommendations to make the most of voice eCall and the MSD data:

- Link NRAs to PSAPs’ coverage, so eCall information is sent from the right PSAP to the right NRA centre. Mapping PSAP to NRA centre coverage is a first step here.
- Think about how quickly the critical information for a stopped vehicle can get to the NRA. For a stopped vehicle this may be by phone to the NRA.

To make the most of the combined eCall data and voice, the MSD should also be passed by PSAP. This is operationally useful as**:**

- Automatic vs manual activations need to be identified as good indicator of provenance and alert quality, and helps exclude false alarms

- It avoids delay and manual error in translating and transcription. The MAIT project showed delays of 4.5 minutes per transaction involving a keyboard.

Some issues for an NRA to consider are:

- How will the NRA take GPS data and map to an NRA road network? (It is not trivial, but NRAs may already do this for other operations.)
- Have you undertaken education on the public on when to use, and not to use, eCall, especially false calls?
- Influencing the script for the call taker at the PSAP so the operator can identify events that are of interest to the NRA. A stopped vehicle in a live lane should be notified to the NRA even if the caller is in a place of safety, and the event is not of interest to the police. For example, PSAP operators could ask:
    - Are you off the live traffic lanes?
    - Are you on a motorway?
    - Are you in danger?
    - Can you leave the vehicle safely without crossing any live lanes?
- Educating call takers in centres about the context of a call – for example broken down on fast roads vs broken down in a supermarket. They may be able to listen for the context of the voice call (e.g., background noise, horns….) as indicators of hazard, and the tone of the caller's voice.
- Educating users about the availability of eCall and how to use it, and not to use it. The Road Safety Authority (2022) in Ireland and National Highways (2022) in England have undertaken publicity about eCall.

But the above based on voice alone is not optimal. To make the most of eCall data, its speed and accuracy and to filter false alarms and triangulate with other sensors, a different approach is needed.

2.4 eCall data in SVD

eCall data can provide a ready and reliable source of stopped vehicle detection. In the most compelling scenario, a vehicle with an eCall unit that strikes an object with enough force to deploy the airbags will send an automatic call to the emergency services. A voice call is initiated as described above but of additional value is the data packet also sent, called the Minimum Set of Data (MSD). This MSD contains data about the location, direction of travel, unique vehicle identity and key vehicle details. If passed directly to responders, this data provides a very fast and reliable source of SVD. It also has distinct advantages over roadside infrastructure-based methods.

We take a detailed look at the potential role of eCall MSD in SVD. It introduces eCall MSD characteristics, identifies the sources of false alarms, and then draws on this to identify its value in the incident management lifecycle of stopped vehicles. The benefits to incident detection are laid out, followed by potential improvements to incident response and resolution to reduce false alarms, increase accuracy and improve relevance to responders. The benefit to downstream systems in the incident management lifecycle, such as active traffic management systems, are also identified.

Methods for assessment of eCall events to correctly identify stopped vehicles are presented, and the distinct advantages and disadvantages are discussed. This includes a brief discussion of the data protection aspects of eCall data. The use of eCall beyond SVD is briefly reviewed.

Finally, we look at improving eCall SVD through synthesis with other “big data” sources and completes with a brief vision of the future for eCall in stopped vehicle detection.

2.5 Introduction to eCall data

The SHADAR report D2.1 showed that on receipt by PSAP the voice channel will be answered by an operator, who will then pass the call and data to operators in different responders. We have identified that each country has different configurations of PSAPs and operators.

Regardless of configuration, each stage of this voice journey typically takes on average 7 minutes. With two operators in the chain, the average time is 14 minutes; with three, the average is 21 minutes. Therefore, the emergency responders and traffic management centres may not be alerted to a stopped vehicle event for some time if the voice channel alone is relied upon. For example, in England an eCall will be answered first by the PSAP 999 emergency operator, who will then transfer the call to the police, who will then notify National Highways. Only then can signs and signals be set to warn drivers.

The research identified that only two countries use the data packet to separately alert the emergency service responders. Using the data packet can reduce the time to alert responders to less than one minute. For England, this could alert the National Highways traffic control centres to our vehicle strike within 30 seconds of the event, rather than the average 21 minutes that exists with a reliance on voice only. This provides a compelling opportunity to improve not only stopped vehicle detection but the response time.

2.6 eCall delivery path and response times

The eCall makes the call using the 112 emergency network. This prioritises the call over non-emergency calls. The MSD is no larger than an SMS to ensure it can be delivered in areas of low network coverage. Even if the voice cannot connect, the MSD can be delivered.

This makes the MSD delivery reliable even in areas of poor network coverage. If there is emergency network coverage the MSD can be delivered.

SHADAR report D2.1 identified several eCall Architectures. Using a figure of 7 minutes average per leg of the journey, we can map the timeline for the models. Though each country may have different average response times, with some lower than 7 minutes, they will be in the order of minutes per leg.

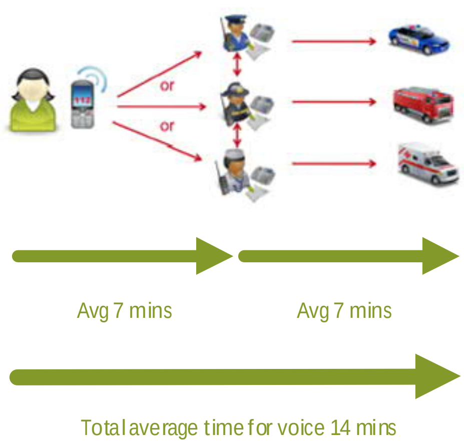

Figure 2 PSAP Model 1 – average time to notify responders by voice

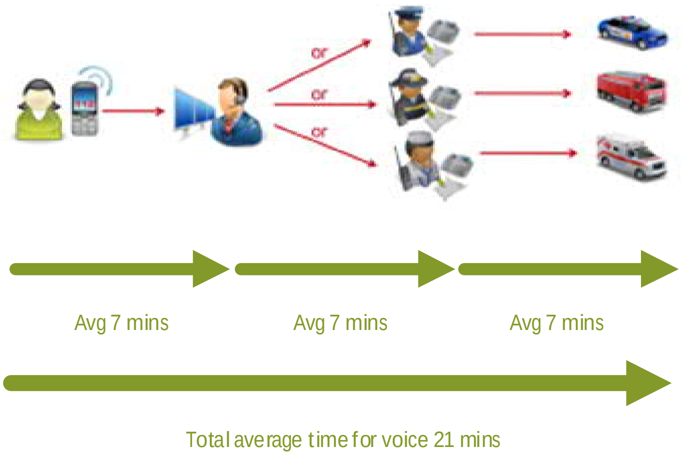

Figure 3 PSAP Model 2 – average time to notify responders by voice

For conciseness, the first two models are shown (similar analysis can be applied to the other models).

Very few countries use the MSD directly to alert responders. The data is passed verbally and therefore takes many minutes to pass through the chain. The eCall MSD can be passed electronically through to responders to alert them to stopped vehicles. This will provide SVD alerts in seconds rather than minutes, an order of magnitude faster than voice alone.

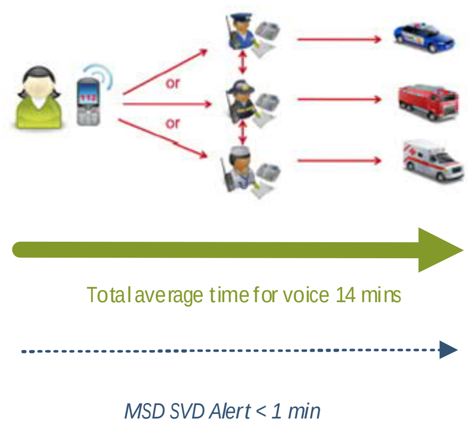

Figure 4 PSAP Model 1 – estimated time to notify responders by MSD SVD

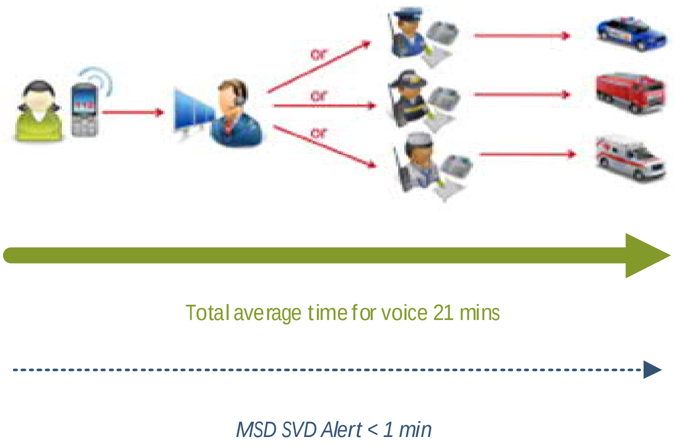

Figure 5 PSAP Model 2 – estimated time to notify responders by MSD SVD

2.7 Challenges for MSD SVD alerts

The MSD does not replace the voice channels. Instead, it provides an “early warning”. However, sending raw MSDs to responders directly raises some questions:

- How reliable is the MSD data for SVD?
- What are the false alarm rates and how do we detect and remove them? Note that the definition of “false alarm” is in the receiver’s viewpoint. For example, a broken-down vehicle might be regarded by blue-light services as a false alarm but if on a smart motorway may be of great operational importance to the NRA.
- Is the data in the MSD optimised for response or does it require enhancing?
- How do we assess the likelihood and impact of an incident and alert the responders accordingly?
- How do we only provide MSD SVD alerts that are relevant to a responder?

These questions are addressed in the following sections.

2.7.1 The eCall Minimum Set of Data (MSD)

Every eCall activation generates an MSD which is sent along with the voice channel from the eCall unit to a Public Safety Answering Point (PSAP). This data is called a Minimum Set of Data, or MSD and must conform with EN 15722.

This Minimum Set of Data (MSD) contains the following fields.

Field	Description and provenance of data
MSD ID	A sequence number commencing with 1, and incrementing with each requested retransmission. The PSAP can request a retransmission of an MSD. This field distinguishes each transmission; any MSD ID over 1 indicates a retransmission. The eCall unit generates this number automatically.

Field	Description and provenance of data
Automatic activation	True or False If the eCall detects an automatic activation event, typically an airbag deployment, this field is True. If the activation is instead caused by the user pressing the eCall button (typically the red SOS button) this field is False.
Test Call	True or False Under live conditions this will be set to True. For eCall unit tests this will be set to False, and typically this would be in controlled conditions where the MSD would not be sent to live PSAPs. However, the likeihood should never be discounted.
Position can be trusted	True or False "Low confidence" means less than a 95% confidence that the position is within a 150m radius This is provided by the eCall GPS device; either from the vehicle GPS or the device hosting the eCall software (dashcam, or smartphone), and is dervied from satellite visibility.
Vehicle Type	Class of vehicle, from: • Passenger Vehicle • Buses and Coaches • Light Commercial Vehicle • Heavy Duty Vehicle • Motorcycles Note that currently only M1 (Passenger Vehicles) and N1 (Light Commercial Vehicles) are mandated to use eCall. This value is set on installation.
VIN	Vehicle Identification Number This value is set in the eCall unit on installation.
Vehicle Propulsion Storage Type	One or more from: Gasoline, Diesel, Compressed Natural Gas, Liquid Propane Gas, Electric Energy Storage, Hydrogen Storage, Other This value is set in the eCall unit on installation.
Timestamp	Time of event This is generated by the eCall unit.

Field	Description and provenance of data
Vehicle Location	As Lat/Long coordinates This is provided by the eCall GPS device; either from the vehicle GPS or the device hosting the eCall software (dashcam, or smartphone). The eCall specification states that the horizontal error at 95% probability should be: - 15 metres in open skies - 40 metres in shadow
Vehicle Direction	Integer, in 2 degree steps (e.g. 179 = 358 degrees) This is provided by the eCall GPS device; either from the vehicle GPS or the device hosting the eCall software (dashcam, or smartphone)

The following fields are optional in the standard:

Optional Field	Description and provenance of data
No Passengers	For vehicle eCall units, indicative based on seat belt connections or similar
Recent Vehicle Location N1	As Lat/Long coordinates This is provided by the eCall GPS device; either from the vehicle GPS or the device hosting the eCall software (dashcam, or smartphone)
Recent Vehicle Location N2	As Lat/Long coordinates This is provided by the eCall GPS device; either from the vehicle GPS or the device hosting the eCall software (dashcam, or smartphone)

The PSAP also provides additional fields:

Field	Description and provenance of data
Sender	The originating number of the eCall unit This value is set in the eCall unit on installation.
Received at PSAP	The time of receipt at the PSAP This value is provided by the PSAP

A manufacturer can provide additional optional fields in the MSD, which allows for them but does not define them. As they are not part of the eCall standard and therefore cannot be relied on to be present we restrict ourselves to the MSD only.

2.8 Analysis of MSD for SVD alerts

2.8.1 eCall Data characteristics

To assess the value of eCall MSDs for SVD, we need to understand the reliability of the data, the false alarm rates, the accuracy, and the relevance of the data to emergency service responders.

Unreliable data reduces stakeholder confidence in the value of eCall MSDs correctly identifying an incident. Excessive false alarms reduce confidence in alerts. Inaccurate data slows down the incident response; for example, an inaccurate location results in wasted on-road deployments, and signs and signals not being set in the correct location.

Activations not relevant to a responder, for example outside the responder's area of responsibility, will increase the perceived false alarm rate for that responder.

2.8.2 Causes of false alarms

The data in the MSD is mostly set on unit installation and hence reliable. The Vehicle Identification Number, the Sender and Fuel Types are pre-set in the device for example. Other fields are tightly defined (true/false or enumerated), which reduces the chances of error.

However, the following areas require further elaboration as they can be significant causes of false, inaccurate, or irrelevant alarms:

a) Automatic and Manual activations

b) Vehicle location accuracy

c) Faulty eCall devices[1]

d) Relevance to responder

Automatic activations

Automatic activations are sent without human intervention and are generated when a suitable vehicle condition occurs, such as high deceleration or airbag deployment. These are identifiable by the Automatic Activation = True datum.

As such, they carry a very strong likelihood of an incident (but not 100%, given the chance of an eCall unit malfunction), and the MSD directly links the vehicle to the incident. The false alarm rate for automatic activations can be considered very low as there is no human involvement.

However, the relevance of an automatic eCall remains an issue; responders do not need to respond to all automatic eCalls, and drivers may not always want a response. UK call statistics indicate that a significant number of automatic calls results in no emergency dispatch, as the caller does not require emergency assistance but may well be in a situation the NRA would want to know about to intervene, e.g. a stopped vehicle following breakdown

Manual activations

It has always been the eCall proposition that a driver should manually press the SOS button to report an emergency, whether the call relates to themselves or another vehicle or person. These are the manual activations, identifiable by the Automatic Activation = False datum.

There is no consistency in the appearance and positioning of these buttons. Some are red with SOS on the button, some are not.

There are many scenarios where someone could press the eCall button, such as:

[1] Anecdotally, one faulty eCall unit in Ireland created an eCall activation every 30 seconds

- to report a breakdown of their vehicle
- to report an incident involving another vehicle seen while driving
- to report a non-vehicle emergency
- to find out what the SOS button does
- to demonstrate the SOS button in a car showroom

Vehicle location accuracy

All MSDs provide the last known location, direction of travel and optionally the previous two known locations. The time period between these points is not mandated in the eCall standard. By design, this prevents any calculations of vehicle speed.

The location is taken from the eCall unit location source. This may be the vehicle's GPS position or a mobile phone location service.

The MSD fields includes a Position Can Be Trusted flag, set to True or False. The definition of "Low confidence in position" shall mean that there is less than 95% confidence that the exact position is within a radius of ± 150 m of reported position.

Given a typical width of a 3-lane carriageway excluding hard shoulder is 11 metres (UK example), a low confidence does not provide sufficient accuracy to identify the location of a stopped vehicle. In dense areas of roads, it may not be able to clearly identify the road, let alone the carriageway.

If the location can be trusted, the accuracy of GPS locations then needs to be considered. GPS accuracy depends upon a number of factors; whether the vehicle is stationary or moving, whether it has a clear sky view, under GPS "shadows", for example in built up areas with reflections and tall buildings reducing satellite visibility, and ultimately, position accuracy.

According to a recent report by European GNSS Agency and Joint Research Centre (2019) which tested 15 manufacturers' eCall units, the average accuracy of GPS positions was as follows:

Condition	Average horizontal error at 95% probability	Threshold (as per EN15722 standard)
Static vehicle in open skies	1.47m – 1.84m	15m
Dynamic (moving) vehicle in open skies	2.95m	15m
Dynamic (moving) vehicle with GPS shadows (urban canyon conditions)	6.68m	40m

We can conclude that if the Position Can Be Trusted flag is True, the location accuracy is at a level that would allow carriageway-specific (or at least neighbouring lane) location for both open skies and in GPS shadows.

With further analysis of the previous two locations and the direction of travel, and taking into the account the local roadway geometry, we believe that location accuracy can be improved further. This is presented later in this document.

Comparison of eCall location accuracy against other methods

The provision of location by an eCall system can be compared against the traditional methods of incident location that are achieved by a mobile handset when using the single emergency number 999/112.

The three methods in use are:

- Caller defined location, which can be highly inaccurate as the caller may not know where they are, or makes use of locally known location descriptions. These local descriptions, for example “by the Brussels Cafe”, may not correlate to the standard address gazetteers employed by the emergency services.
- Cell referencing. This method uses the cells in the area to give a triangulation of the caller. This works reasonably well in an urban environment, however in the rural context the lack of cells make accuracy limited. In some cases, within 5km is the best accuracy achieved.
- Advanced Mobile Location (AML) now available across all mobile handset types giving sub-50m accuracy inside or outside. This is not universally employed by all emergency services, despite being mandated across Europe. In order to obtain the handset location through AML, the level 2 PSAP dispatcher has to make a request to the handset for the provision of the data. This is not automatic. The caller needs take no further action, but this second request to the handset is required.

Comparison of caller location is key to any emergency service response as can be seen from the above location capabilities. All have limitations, though AML is clearly a major advancement, and will improve as more emergency services gain more capability. The fact still remains that the operator has to ask the handset for the location, the user does not need to do anything, but nevertheless this adds another 20 seconds to the dispatch process.

By contrast eCall location referencing is provided automatically when the eCall is triggered. It is accurate, and the accuracy will only increase as more satellites become active.

Faulty eCall devices

eCall data from the English 999 emergency calls demonstrates that faulty eCall units can create a significant number of false alarms.

In June 2016, a single faulty eCall unit created approximately 2,500 false alarms compared to a monthly average of 900. In September 2017, another faulty device resulted in 1,659 calls compared to a monthly average of 1,600.

Faulty units, though uncommon, are a significant source of false alarms, and if not detected could overwhelm any automated SVD alerting system.

There is no uniform process defined in Europe to mitigate faulty eCall devices. It should be recognised that generally eCall devices are purchased in high volume by vehicle makers. So, there is a significant chance where a fault in one eCall unit is detected then it could be replicated across Europe or beyond in the same batch of equipment. It has been seen that motor manufacturers are now able to offer over the air fixes for some of these issues, which is an improvement, but does not resolve all of the issues.

It has yet to be decided it the eCall unit will form part of the annual fitness test as for most other safety devices.

Relevance to responder

Road networks in one country are often managed by multiple different authorities. For example, in England, National Highways is responsible for the strategic road network, but is not responsible for local roads.

As eCall activations can originate from any location – motorways, country lanes, off-road, driveways, car showrooms, scrapyards – a responder who receives unfiltered eCall MSDs will need to disregard those outside their area of responsibility. Though they may not be false alarms, for that responder they may be considered as such.

MSDs used in SVD must first be geographically filtered for each responder's area of responsibility to avoid diluting the value of the detection.

2.8.3 eCall rate analysis and false alarm rate estimation

Figure 6 shows voice calls from eCall activations in England from March 2019 to June 2021, with trends for both total calls made and connected calls. The reduction in traffic volumes due to the COVID-19 pandemic is clear in 2020. Though the data does not include other countries, it can provide an indication of connected alarm rates for eCall.

Connected v Unconnected Calls

A connected call is one where an operator passes the eCall voice channel to an emergency service. The operator may not connect a call to an emergency service if the driver does not request this, e.g., they are safe following an automatic alert, or a false alarm button press.

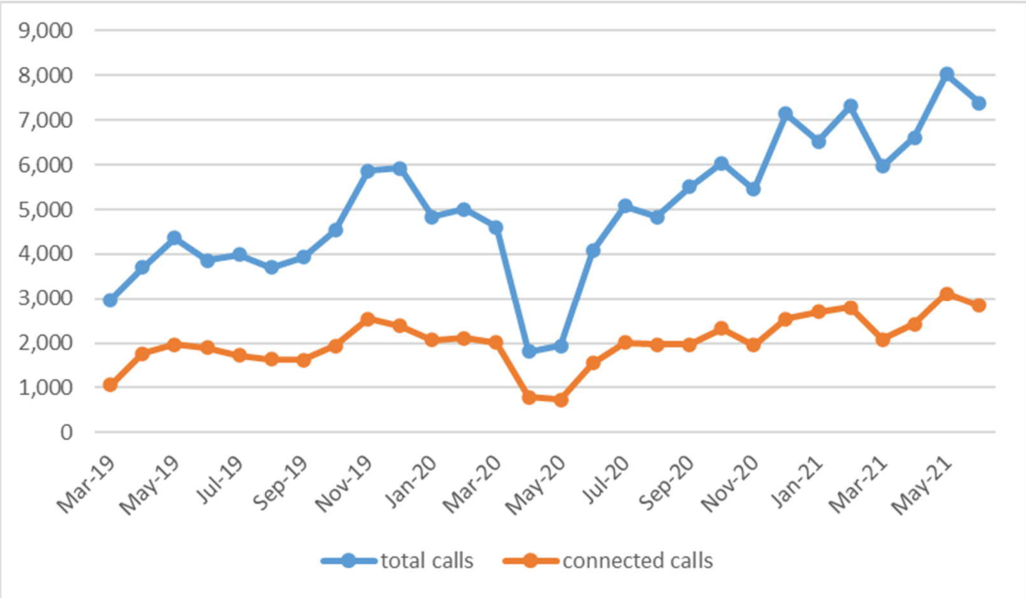

Figure 6: eCall totals and connected calls in England, Mar 19 - Jun 21

Figure 7 shows the unconnected calls as a percentage of the total, averaging around 60%. This is a useful metric when considering how the eCall MSD can support SVD detection.

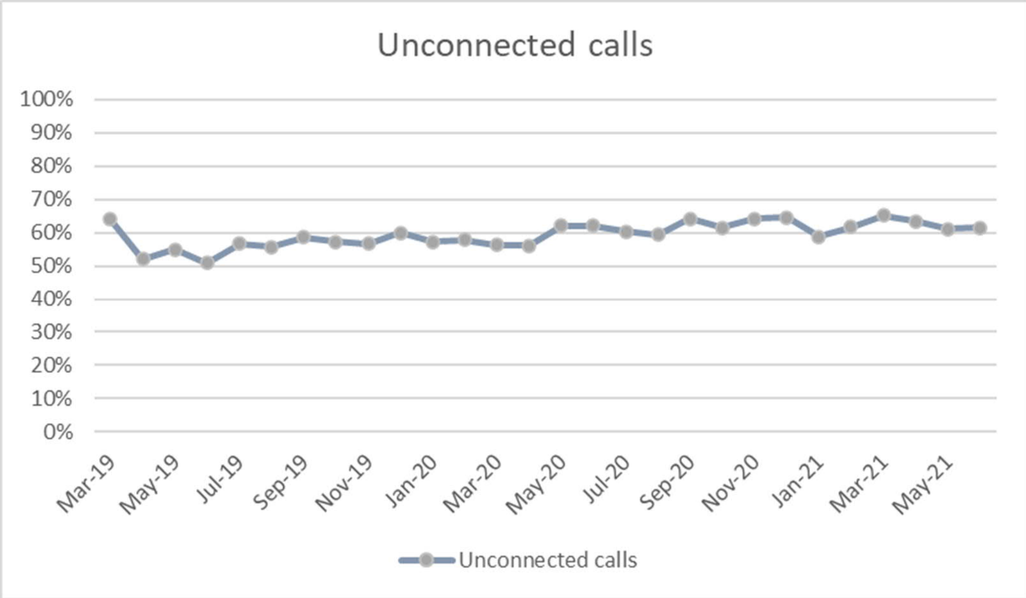

Figure 7 eCall unconnected proportion in England, Mar 19 - Jun 21

Manual v automatic calls

In the UK in June 2022, there were 10,947 calls, of which 528 were automatic calls. This provides us with a metric of 5% of all eCalls are automatic.

Assessment of Automatic eCall false alarm causes and rates

Of all eCalls 5% are automatic. Of these we can make some assumptions of the false alarm rates:

Source	**Description**	**False alarm rate (of automatic)**
False alarm – faulty unit	The eCall unit malfunctions and generates false activations. There have been two instances of faulty eCalls in the UK over the years mentioned above. As these are rare and can be screened out with eCall processing we discard these as a material source of false alarms.	0%
No emergency response required	The caller does not want emergency assistance for the incident, for example where an incident occurs close to the driver's home. However, although a driver may not want assistance, a responder may still wish to be informed of the call. For example if the airbag has deployed, the vehicle will be undriveable	5%
	Total Automatic eCall false alarm rate	5%

The overall false alarm rate for automatic calls is estimated at 5%.

The overall false alarm rate contribution of automatic eCalls to total eCalls, is:

- Percentage of all eCalls which are automatic (5%) x Automatic eCall false alarm rate (5%) = 0.25%

Assessment of Manual eCall false alarm causes and rates

Of all eCalls, 95% are manual. Of these we can make some assumptions of the false alarm rates, bounded by the metric that 60% of all eCalls are not connected to the emergency services. We have assigned estimates to each scenario. Without historical data available for detailed analysis, we applied our empirical knowledge, bounded by the available statistical data to estimate the rates for each scenario. These estimates are indicative only and should not be used authoritatively.

Source	**Description**	**False alarm rate (of manual)**
False alarm – human error	The caller manually activates the eCall not knowing its purpose. Anecdotally, this appears to be the greatest cause of false alarms; people not knowing what the eCall button is for. We estimate this at 40% of all calls.	40%
Silent call	The caller does not speak or cannot be understood. This could be due to a genuine incident with an unconscious driver, one who cannot be understood, or one unable to speak. Though this is unlikely to form a significant part of the unconnected call figures, it is an important use case that we will address later, as the MSD can help determine a genuine incident without relying on voice. We assert that silent calls should not be treated as false alarms.	0%
	Total Manual eCall false alarm rate	40%

The overall false alarm rate for manual calls is estimated at 40%.

The overall false alarm rate contribution of manual eCalls to total eCalls, is:

- Percentage of all eCalls which are manual (95%) x Manual eCall false alarm rate (40%) = 38%

In summary:

	Manual eCalls	Automatic eCalls
Call Volume	95%	5%
False alarm rate	40% of manual	5% of automatic
False alarm rate of all eCalls	38%	0.25%

Estimating multiple activation false alarm rates

With a 40% false alarm rate, a single manual activation is still worthy of investigation but requires confirmation through CCTV or on-road attendance.

Multiple manual activations in the same area and time will increase the likelihood that an incident has occurred. We would expect a major incident to result in many manual eCalls from observer vehicles, as well as automatic activations from the vehicles involved.

With some simple probability calculations, we can estimate the upper limit of the false alarm rate of multiple manual calls. The real false alarm rate for multiple calls will be much lower, when the probability of calls in the same location and time frame are factored in. For our purposes we can conclude that two manual eCalls in an area and time will have a low false alarm rate, with additional calls having exponentially lower rates.

	1 Manual Call	2 Manual Calls	3 Manual Calls	4 Manual Calls
Simple manual false alarm rate estimate	40%	16%	6%	2%

Table 1: Simple probabilities of false alarm for manual activations

Later we will consider how these values can be combined with other SVD alerts, as well as environmental and historical accident data, to reduce false alarm rates further through statistical analysis and "big data" fusion.

Relevance of eCall to an incident

A manual activation MSD may not be generated from a vehicle involved in the incident. For example, a driver may see an accident on the opposite carriageway but only decide to report it after a few minutes. In this example, the location of the activation and the vehicle details in the MSD will not be relevant to the original incident.

2.8.4 MSD data analysis

Detailed analysis of a large sample of eCall MSD data, such as a breakdown by automatic and manual, by location and road type, would provide further insights into the shape and trends of the data. It will help us to identify where SVD provides the greatest value.

Examples of the insights we seek are:

- Proportion of manual to automatic activations
- Geographic distribution of activations
- Analysis of activations by road network:
    - Managed motorways
    - Non-managed motorway
    - Major non-motorways
    - Local roads
    - Off network locations
- Analysis of location accuracy, including how closely locations map to roads
- Number of units with repeated calls, indicating the presence of faulty units
- Clustering of calls by location indicating potential sources of false alarms such as vehicle showrooms and scrapyards
- Clustering of calls by location and time, to assess characteristics of multiple eCall events

A SHADAR project partner has sought access to data for England through National Highways but has not been able to secure access. This prevents any further insights and we have had to rely on summary statistics to estimate values. However, we have worked with partners to generate test MSD data which backs up some of our qualitative assessment of the data such as location accuracy. Although we do not believe it is critical to this analysis, we recommend that MSD sample data is obtained to identify further patterns and trends.

2.8.5 Methods for using eCall MSD for SVD alerts

With our assessment of the characteristics of eCall MSD we can now propose specific methods for its use in SVD. These address the challenges raised in section 2.7.

These methods fall into four areas and form a workflow process for MSDs (Figure 8). These processes would be implemented in an eCall process engine which we name "TeCall".

- Filter – to detect and reduce the instances of false alarms
- Enhance – to improve the operational value of MSD data
- Profile – to prioritise activations based on the likelihood and severity of an incident
- Forward to Responders – to provide the relevant MSD SVD alerts to the emergency service providers

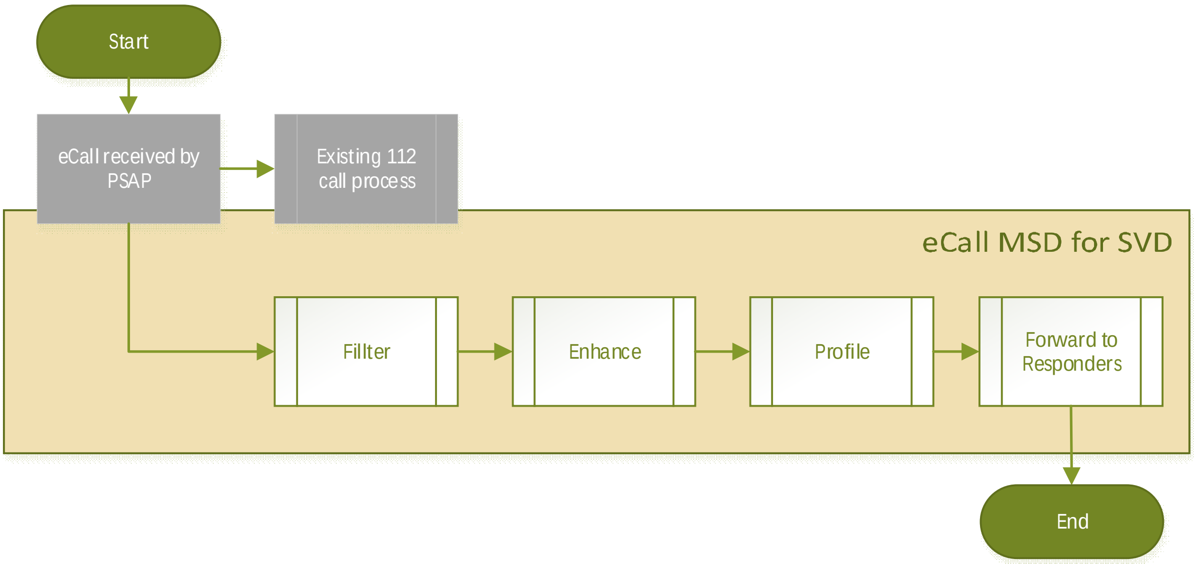

Figure 8: eCall MSD for SVD alerts

Filter

The filter process prevents clear false alarms from being forwarded on to responders.

- VIN Blacklist

  Identify faulty units as a particular source of false alarms, where a single unit can create thousands of alerts. This filter detects unusual numbers of MSDs from a single vehicle and blocks further activations for a period.

Enhance

The MSD is an intentionally small dataset, designed to be delivered over low bandwidth networks. It is not optimised for responder operations. For example, the MSD contains the VIN but does not include the make or model, which is important in identifying the incident on CCTV and coordinating recovery. We have identified areas where the data can be enhanced to assist responders by reducing the time taken to "decode" the MSD data to make it operationally relevant.

- Add roadway names and direction

  The MSD contains the latitude and longitude position and optionally the previous two positions. With a reference network model these data can identify the closest matching roadway and include this with the MSD.

Where coordinates match two or more roadways, resolved by using the vehicle direction and the previous two positions to determine which roadway the vehicle is travelling along.

- Add reference points

  In addition to the roadway name, the nearest reference point such as a junction, feature or marker post can be determined using a GIS database lookup. This provides a more granular location than the roadway name alone.

- Add vehicle details

  The MSD VIN can be decoded using online vehicle data service providers to provide descriptive vehicle details including make, model, vehicle registration, number of doors, automatic or manual transmission (important for recovery arrangements) and colour (important for visual acquisition of the vehicle with CCTV)

- Add event flags

  Some information in the MSD may identify specific risk factors for the event. These can be algorithmically identified from the MSD and included as one or more event flags. Example of event flags are hazardous fuel, potential stopped vehicle strikes, and multiple vehicle incidents.

  In particular, the analysis of the last three position could indicate potential crossovers, where the first two positions are matched to one carriageway and the final position is on the opposite carriageway. The opportunity to explore potential for location analysis in this way needs to be pursued through access to live eCall data.

  The authors wish to highlight the importance of further research into eCall data with real datasets.

- Add related events

  Where MSDs have already been received within a given radius and a given time, the MSD can be enhanced with references to the earlier MSD alerts. This would allow responders to tie multiple alerts together. A single manual MSD activation has an indicative false alarm rate of 40%; two manual MSD activations in the same place has a false alarm rate of 16%.

- Add multiple events

  If a PSAP requests a retransmission of an MSD, the MSD SVD can add a reference to the earlier MSD SVD alert to help the responder to tie the retransmissions together. These MSDs need to be detected and managed differently to avoid creating multiple SVD alerts.

Profile

With limited resources, responders need to prioritise SVD alerts over other demands. Given the enhanced data available we can prioritise SVD alerts using two risk-based criteria – likelihood of an incident, and impact of the incident - to generate a risk-based priority.

For example, eCall SVD alerts could be prioritised as in the following table.

Type of activation	Likelihood	Impact	Priority
Single manual activation	Low	Low	Low
Potential stopped vehicle strike: Single automatic activation, with three locations in same place	High	High	High

The location accuracy and relevance may also determine the priority. Any MSD that does not have the Position Can Be Trusted Datum = True will not have a reliable location and may be assigned a lower priority. Manual activation MSDs may carry the location of an observer rather than the incident, whereas an automatic MSD will carry the incident location. Each can be prioritised separately.

There is also an opportunity to fuse the SVD data with other environmental and situational data to refine the profiling - this topic is explored in Chapter 5.

Forward to responders

It has been identified that response times can be improved by an order of magnitude over voice calls, through forwarding the MSD SVD alert data directly to responders. However, each responder may have specific areas of responsibility, and SVD alerts need to be filtered by the geographic location. With the road identified in the “Enhance” stage, the responder can be identified by a lookup which may have been prepared using provided data or by GIS “geofencing” techniques.

In addition, as we are prioritising SVD alerts using risk factors, a responder may only want to receive SVD alerts over a certain priority. We should recall that MSD SVD alerts are provided in addition to the voice call. A responder may only want to receive data for the high priority SVD alerts and rely on the voice calls only for the low priority alerts.

TeCall

Chiltech and partners[2] have built a demonstration of this MSD processing capability, called TeCall. This successfully shows that MSDs can be processed as described above, with a response time in the order of seconds.

Application to eCall Use Cases

We identified four challenges with eCall for SVD. The processing above demonstrates that they can be treated, and the issues mitigated to great extent. The challenges and mitigations are summarised in the following table.

[2] Centras Associates, White Willow Consulting and ShadowFocus Consultancy

Challenge	Processing
Automatic and Manual activations	Automatic and manual MSDs can be managed and prioritised in different ways
Vehicle location accuracy	Location accuracy and relevance to the incident can be prioritised based on the MSD data and the type of activation. In addition, the three locations in the MSD can be matched to a network model to identify the carriageway and direction of travel.
Faulty eCall devices	MSD filtering can automatically identify and block units that produce large numbers of faulty eCalls.
Relevance to responder	Each responder can receive only the eCall MSD SVD alerts that are in their area of responsibility and filtered by those above a minimum priority.

In addition, MSDs provide an opportunity to enhance with vehicle and location data to make it operationally relevant, reducing the response times.

Conclusion

These methods demonstrate that causes of false alarms identified in 2.8.2 Causes of false alarms can be addressed through MSD processing. The speed of response can be increased by an order of magnitude over voice eCall. In addition, the enhancement of the MSD with additional data provides an opportunity to improve the operational value for responders.

2.8.6 Comparison of MSD SVD with infrastructure SVD and voice eCall

It is useful to compare eCall MSD SVD alerts to alerts generated by on-road infrastructure-based solutions such as radar and CCTV.

Relative values have been assigned (Very Low to Very High) as our quantitative data is only from the UK (we do not have access to data for other countries, which would require anonymisation before it could be shared to us) and can only be indicative for other countries.

Area	eCall	Infrastructure
Road coverage	Very High 98% of UK roads eCall SVD alerts are generated on any road with minimal mobile coverage. According to the RAC (2022), 2% of roads have no mobile coverage. eCall provides an SVD system where infrastructure SVD is either too expensive or cannot be installed.	Low Only on roads with infrastructure installed. 18% of England's smart motorways have SVD and smart motorways account for 7% of the strategic road network
Vehicle coverage	Low Only vehicles fitted with eCall units. In the UK 16% of vehicles are estimated to be fitted with eCall.	Very High SVD detects all vehicles.
Ease of installation and maintenance	Very High For Road Transport Authorities there is no cost for the installation and operation There would be a relatively low cost for the operation of eCall processing for SVD.	Very Low Requires installation of equipment, and ongoing maintenance costs for power and communications, including closure of road for routine maintenance and repairs.
Reliability of SVD alert	Medium Automatic eCall carries a high level of confidence, of location and incident Manual eCalls carry a low level of confidence or location and incident	Very High Alerts carry a high confidence of location and incident
Data richness – vehicle	Very High The MSD, enhanced with vehicle details, describes the vehicle in good detail	Low Alerts can sometimes identify the vehicle registration plate but may not include any further details of the vehicle or passengers
Data richness – location	Medium The MSD, enhanced with roadway details, describes the vehicle location	High Infrastructure at known locations

These values are shown on the following chart, with 0 = Very Low and 5 = Very High.

Figure 9 Relative values of eCall MSD SVD and Infrastructure SVD

In addition, eCall SVD has specific advantages over the voice-only eCall that is in place for almost all road authorities:

Area	eCall MSD SVD	eCall voice
Speed of response	Very High Data-based SVD alerts are an order of magnitude faster; an SVD alert can be processed and forwarded in less than a minute	Very Low Requires voice handoffs between initial PSAP call taker and further emergency responders, typically 14 minutes for two parties
Dependence on human factors	None The MSD SVD is not dependent on human communication, and can detect incidents where drivers cannot talk through incapacity or language.	Medium Relies on voice call and ability to communicate

These comparisons show that eCall SVD alerts are complementary to infrastructure-based SVD alerts and eCall voice calls. The addition of eCall SVD complements existing SVD methods to create a much wider coverage of the road network.

2.8.7 Data protection for eCall MSD

Chiltech et al has separately undertaken work to assess the Data Protection implications of eCall and review existing EU reports and legislation.

The key conclusions from the documents consulted are:

1. MSDs contain only tenuous linkages to Personally Identifiable Information (PII).
2. MSDs may only be used for the resolution of emergency situations.
3. Enhancing the MSD with data such as vehicle make and model to improve the emergency response is permitted.
4. MSDs can be recorded for statistical purposes. However, the VIN should be removed from historical data to reduce any residual risk.
5. Level 2 emergency responder such as National Road Authorities, have a right to the MSD data

2.8.8 Use of historical eCall data

Aside from the operational benefits of eCall, the eCall data provides a rich historical dataset that can help:

- Inform infrastructure investment decisions by identifying areas with high incident rates
- Inform incident analysis by vehicle type, road type, time of day and date
- Identify growth of eCall volumes as vehicles are replaced or upgraded with eCall capability

2.9 The future of eCall and SVD Detection

eCall provides a ready and growing source of stopped vehicle detection. There are challenges with false alarms, accuracy and relevancy but we have demonstrated that these can be overcome.

In addition, eCall as a technology is not standing still; the next generation eCall standards are already in existence.

Whilst the current eCall uses 2G and 3G and communication technology, the next generation will use 4G and beyond.

The next generation of eCall will retain the same MSD that is currently received at the PSAP but there will be a capability to trigger additional sensors on the vehicle. This could include video or any other form of digital sensor. The method of data transmission will change from circuit-switched to packet-switched, with the limitation being the data capability of the network, and of course the receiving point. The change to eCall will also coincide with changes at all PSAP across Europe, from Public Switched Telephone Network (PSTN) to Session Initiation Protocol (SIP), which in the simplest terms transfers all calls to the internet. This will include 999/112. With the location and destination of the call defined by a header with an IP address, this provides infinite possibilities for the exchange of data.

2.9.1 The opportunity of eCall MSD SVD

There is an immediate opportunity to use the MSD element of the eCall to improve response times by an order of magnitude, improve operational response, and integrate with operational systems.

By fusing the MSD with environmental data and other SVD sources such as radar and CCTV, there is a much greater opportunity to create integrated SVD response systems. Such systems could receive SVD alerts from multiple sources, collate them, assess them with environmental data, and provide higher quality SVD alerts than a single source can do.

2.9.2 Breakdown Notification (bCall)

Breakdowns are a common reason for a vehicle stopping on a live lane. In the UK some motoring organisations have recognised that the information generated by eCall meets needs for information required when a vehicle breaks down.

The recovery services need to know at the very highest level:

- What is it (VIN)
- Where is it (GNSS Location)
- Fuel type
- The number of people in the vehicle

Some motoring organisations have now revised their call handling systems to provide fields to be completed by the call handler, when the member calls in to report a breakdown, so that the above information can be entered into the command-and-control system for the dispatch of a recovery vehicle. The same is also true where motoring organisations provide an App to the member, the required information is then provided automatically.

Some motoring organisations in the UK and beyond also have direct contact with the NRA to make sure that they are informed of potential obstructions on the network. This could be automated, if the NRA so chose.

2.9.3 A vision for a new Emergency Data Service

Today we have telephone-based emergency call handling services; we call 112 or 999 and speak to people. But where does emergency data like eCall go? It currently uses the call-handling channel. Is this the best way of using this data, and where will future emergency data go from the next generation of vehicle-based sensors?

The authors can envisage a potential future of a data-based “emergency data service” that gathers, assesses, fuses, prioritises and forwards emergency data alerts such as SVD to NRAs and other responders. This is the digitisation of emergency response, and it could realise the full opportunity of connected vehicles and in-vehicle data for emergency response.

2.10 Characterisation of research ideas

Table 2 and Figure 10 place the research ideas discussed in this chapter into a How-Now-Wow matrix as described in the introduction.

	Subject	Category	Description
NOW	eCall Voice	Timely	Today's baseline in some countries
	Link PSAPs to NRAs	Timely	Reduces delays with SVD going to wrong control centre
	Educate road users in eCall	Reliability	Reduce false alarms
		Timely	Drivers more likely to use it when needed
	Optimise call handler processes, scripts and training	Accuracy	Get the right information from the right source
		Timely	Reduces delays in getting the right information
		Information	Get the right information at the right time
WOW	**Subject**	**Category**	**Description**
	Automatic eCall data processing	Accuracy	Highly accurate
		Reliability	Very high SVD indication
		Timely	Very fast (<1 min)
		Information	Provides details of vehicle and location
	Manual eCall data processing	Accuracy	Highly accurate
		Reliability	Medium SVD indication
		Timely	Very fast (<1 min)
		Information	Provides details of vehicle and location
	Automatic and Manual eCall data fusion	Accuracy	Highly accurate
		Reliability	Very high SVD indication, with greater coverage than just auto eCall
		Timely	Very fast (<1 min)
		Information	Provides details of multiple vehicles and locations
HOW	**Subject**	**Category**	**Description**
	eCall and bCall data fusion	Timely	Faster responses due to greater SVD coverage
		Information	Much richer data for SVD

Table 2: Now-Wow-How categorization harvesting data from eCall

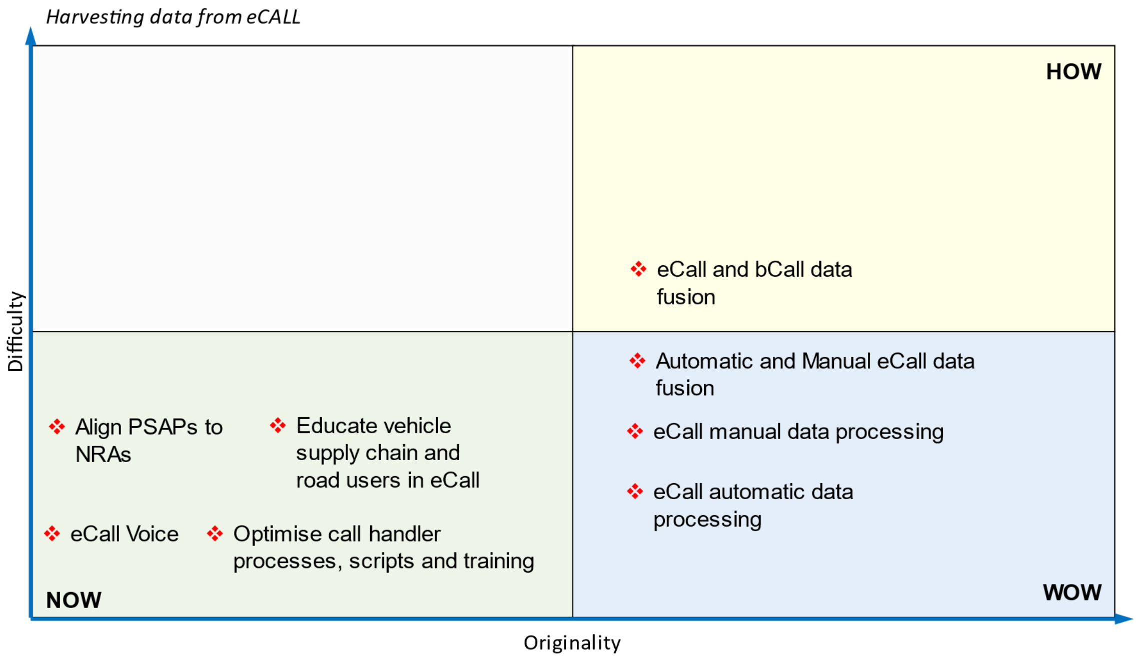

Figure 10; Now-Wow-How matrix harvesting data from eCall

3 Improvement of radar detection

3.1 Introduction

This chapter explores the potential for improvements in rotating radar for stopped vehicle detection. The facts quoted and the experiments performed used Navtech Radar equipment, but the principles could in theory apply to any rotating radar equipment.

Rotating radar has been deployed in at least the following countries, for the contexts identified:

- Norway – highways, tunnels
- UK – highways, tunnels
- Sweden – tunnels, highways, bridges
- Switzerland – highways, tunnels
- Finland – highways
- Italy - highways
- Thailand – highways
- Australia – highways
- New Zealand – highways
- Argentina – highways
- Canada – highways (animal detection)
- Austria – highways
- Netherlands - tunnels
- Slovenia – tunnels
- China – tunnels
- Egypt – tunnels

3.2 Lane discrimination information

The rotating radar unit is currently deployed in many projects worldwide for stopped vehicle detection. Each of these projects include different specifications and requirements, which dictate the spacing between sensors. The distance between these units influences capabilities such as confidence levels to identify the exact position that the stopped vehicle has occurred in – potentially enabling the identification of the lane.

Lane information can be important to communicate. Currently when a stopped vehicle is detected, its longitudinal position along the road is indicated over a given 100m section length of the road. Azimuth data, which could be used to determne the lane the vehicle is positioned in, is currently not utilised. This information could be extremely useful in identifying the hazard and impact level of each individual stopped vehicle in a live lane. For example, a stopped vehicle in the slowest lane is safer than one that is in a faster lane. The reporting of the lane data could influence the severity of the alert, and the nature of the response. To achieve this level of detection, further consideration on sensor parameters and spacing is required.

A rotating radar sensor can detect and track objects in both range and relative azimuth from the radar, as depicted in Figure 11. The location accuracy of detection can therefore be split into these two key parameters.

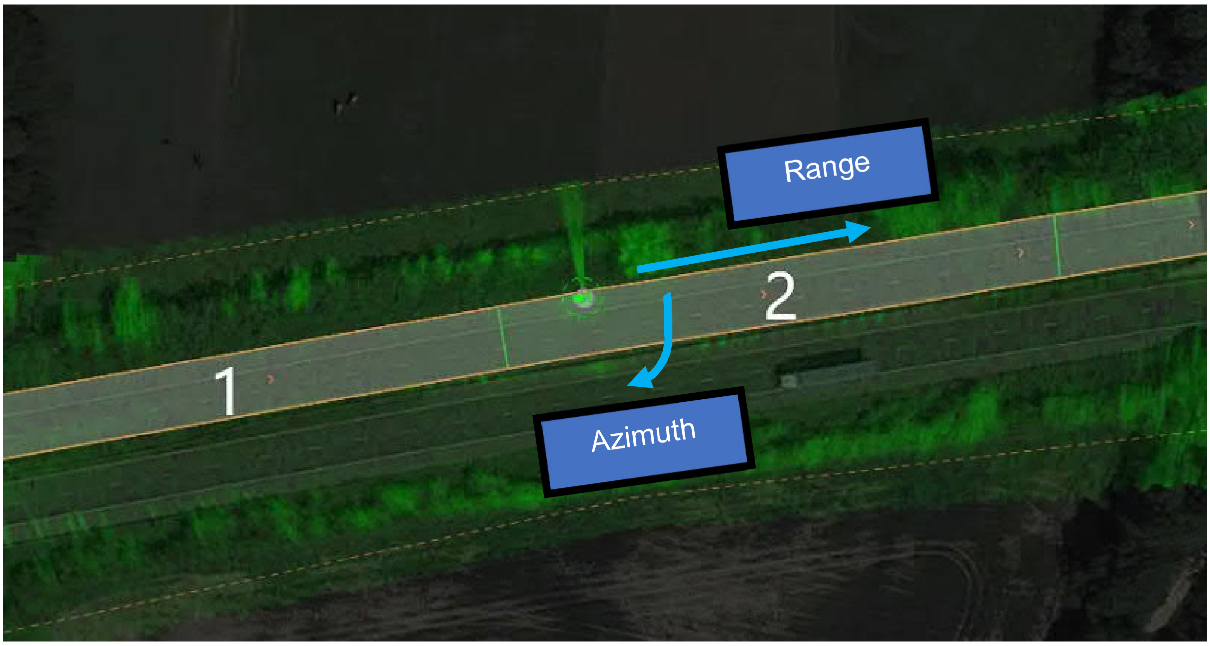

Figure 11: Rotating radar data

For the rotating radar, the locational accuracy of both range and azimuth is dependent on first the radar hardware, and then the software processes that take the radar raw data and convert it into vectorised tracks. These tracks contain information about the stopped vehicle, including centroid and location. The rotating radar hardware can detect object at a range resolution of only 17cm, out to 500m in all directions. This provides great accuracy of stopped vehicle locations in range. However, the azimuth positioning is dependent on the beam width of the radar sensor. The beam width of the radar source is a fixed value, but this causes a virtual beam width increase with distance from the radar. This means an uncertainty of azimuth-based location that increases with distance from the radar source. Theoretically, the rotating radar sensor should be able to determine the correct lane within a range of approximately 150m. At greater range, the uncertainty is more than width of a typical lane. In our knowledge, this theory had not previously been tested.

Experiments were performed in which a vehicle was driven up a temporarily closed motorway lane covered by a rotating radar system, making a series of stops, and GPS measurements were taken in the vehicle.

Figure 12 shows the graphical user interface used for testing. A total of 14 stops were completed on this two radar segment of road. For each of the stops, the rotating radar system detected and tracked the vehicle, and produced a corresponding latitude and longitude location.

It was expected that the error within the GPS system should be independent of the distance from the radar for each stopped vehicle, whereas the radar system would have an error dependent on range from the radar. To test this theory, the distance between both radar and GPS latitude and longitude was calculated at each stop, and then plotted with respect to distance from the radar. It was expected that as the distance increases from the radar (x axis) that the distance between the points would also increase, showing a positive correlation, since only one of the data sets has a dependent relationship with range from the radar.

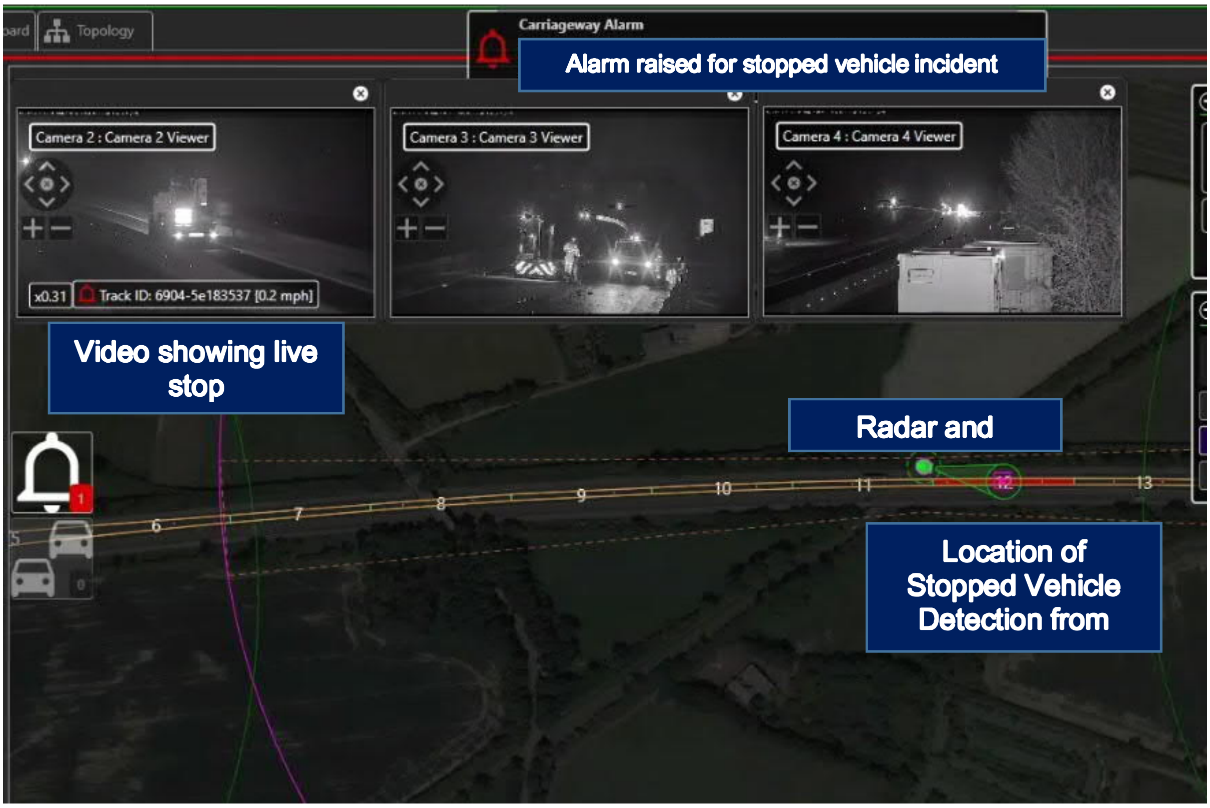

Figure 12: Navtech Radar ClearWay User Interface showing radar and live stop acquired through testing

Figure 13 shows the plotted graph of the distance (m) between the GPS data and the rotating radar location, versus the range from the radar that it was detected. As stated above, it is expected that as range from the radar increases, the distance between GPS and rotating radar detection also increase. This relationship however, was not observed in the graph. The points of the distance between these two location sources did not positively increase as the distance from the radar increased. Neither did any relationship form past 150m from the radar.

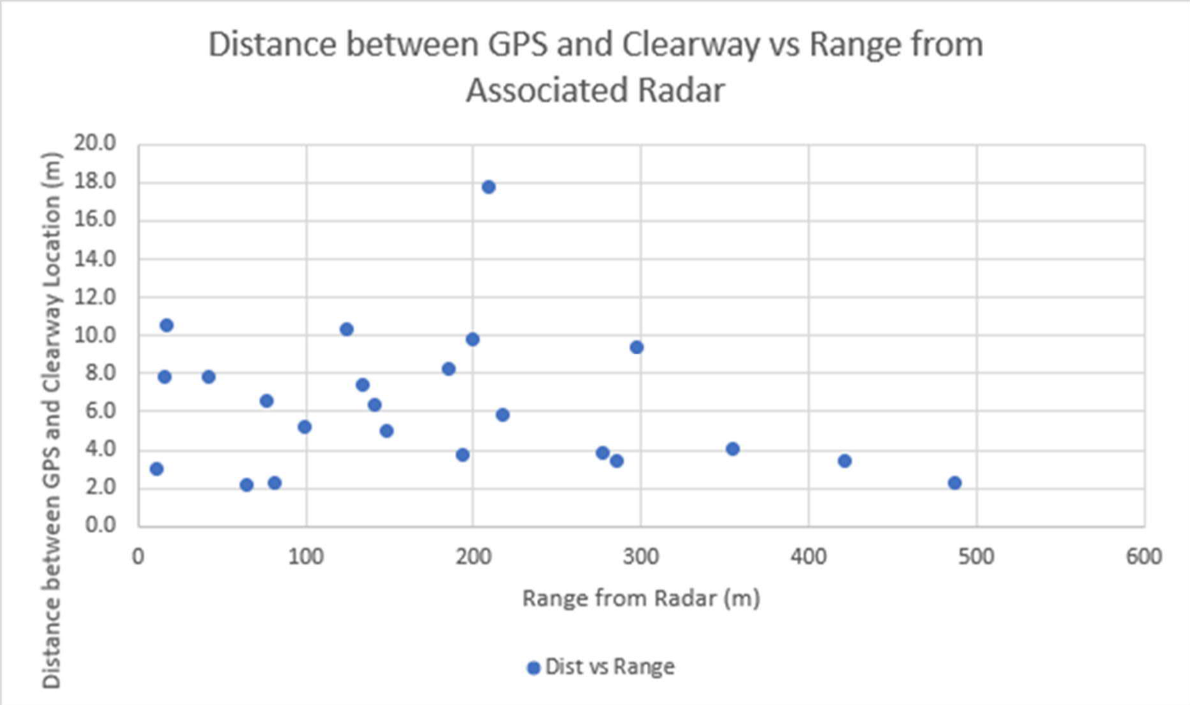

Figure 13: Graph of distance between GPS and rotating radar vs range from the associated radar

The p-value of this data set was also calculated. This is a hypothesis test to assess how strong the evidence from the dataset is to arrive at the necessary conclusion. With the dataset above, the p-value was 0.9. This is extremely high, and shows that the data does not support the expected conclusion.

An alternative visualisation of this data is by plotting on a map, as in *Figure 14* which shows two examples of these stops, with the black representing the GPS and red representing the radar detection. The camera symbol represents the location of the radar. Inspection of maps showed that the GPS did not always position the vehicle in the lane in which we knew it had been physically located. Further investigation confirmed that the GPS source stated a +/- tolerance of 5m in positional accuracy. The GPS should not be used as the ground truth for this analysis as it was not sufficiently accurate.

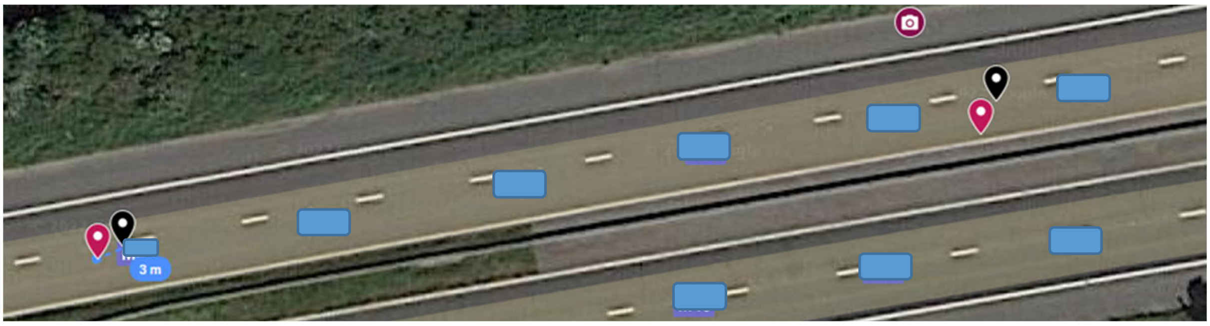

Figure 14: Position of Radar ClearWay (Red) vs GPS (black) stops (blue boxes used to hide distinguishing information)

It was known that vehicle attempted to drive up the centre of the lane, so a separate comparison was made between radar stops and the centre line of the lane.The positions for the latter were manually obtained from map software. Calculating the differences between these points produced the graph in Figure 15.

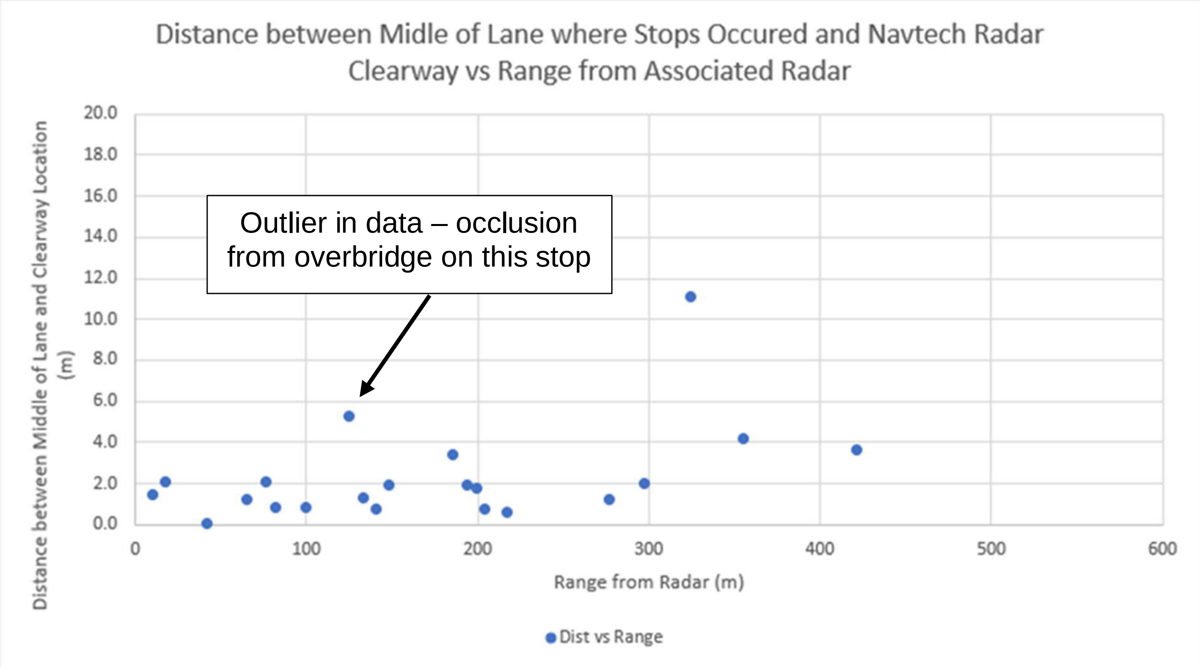

Figure 15: Graph of distance between middle of the lane and rotating radar vs range from the associated radar

Successful lane identification is assumed to occur where the distance is approximately 2m or below. With that assumption, lane identification was successful for all points within 150m of a radar with one exception due to known occlusion from an overbridge. Up to 250m from the radar, lane identification is successful in the majority of cases. Beyond that range, the variance is much higher.

This result shows a clearer indication of the relationship that was proposed in the hypothesis. However, running this through another correlation test, shows the result and correlation still being statistically insignificant. A larger sample may confirm the result.

3.3 Quantifiable traffic parameters

The method in which a vehicle stops on a live lane is dependent on the exact cause of the issue it is experiencing. For example, a vehicle with a punctured tyre may on average display different behavioural characteristics to someone with an engine failure, or involvement in an accident. The radar improvement in this section is to understand how additional data can help identify the risk of a stopped vehicle.

Vehicles stopping in live lanes do not necessarily stop abruptly. For example, a driver may put their hazard lights on, start to slow down, change lanes into a safer one, and then eventually come to a slow stop, all the while causing changes in traffic conditions, which can include queuing or slowing down of other motorists. If this data can also be captured to identify and quantify these unusual traffic behaviours, this can feed into decision making. Over time, this learned behavioural data can be assessed to see what occurs before, during and after a stopped vehicle event, with the potential of incident prediction.

3.3.1 Traffic Speed

In terms of unusual traffic behaviour, a specific metric that can be identified is vehicle speed. In the example explained above, a stopped vehicle may be caused after a particular car slows down and creates more cautious driving by other vehicles.

Data gathered from a live project from Navtech Radar's ClearWay system on a 5-lane highway shows an event where a stop occurred in a live lane. The ClearWay system operates by creating a virtual carriageway that covers the road or real carriageway within the radar's detection area. These virtual carriageways are separated into 100m length sections, where traffic statistics such as speed and density, can be calculated for each.

Using this concept, traffic speed data was collected and is shown in the graph in Figure 16. This shows the traffic speed from 16:00 to 17:00 gathered from one typical weekday afternoon, as a function of coverage distance that the Navtech Radar ClearWay system has on this scheme. Apart from a drop approximately 10km into the scheme, the average traffic speed during this time generally sits between 56 and 58 miles per hour.

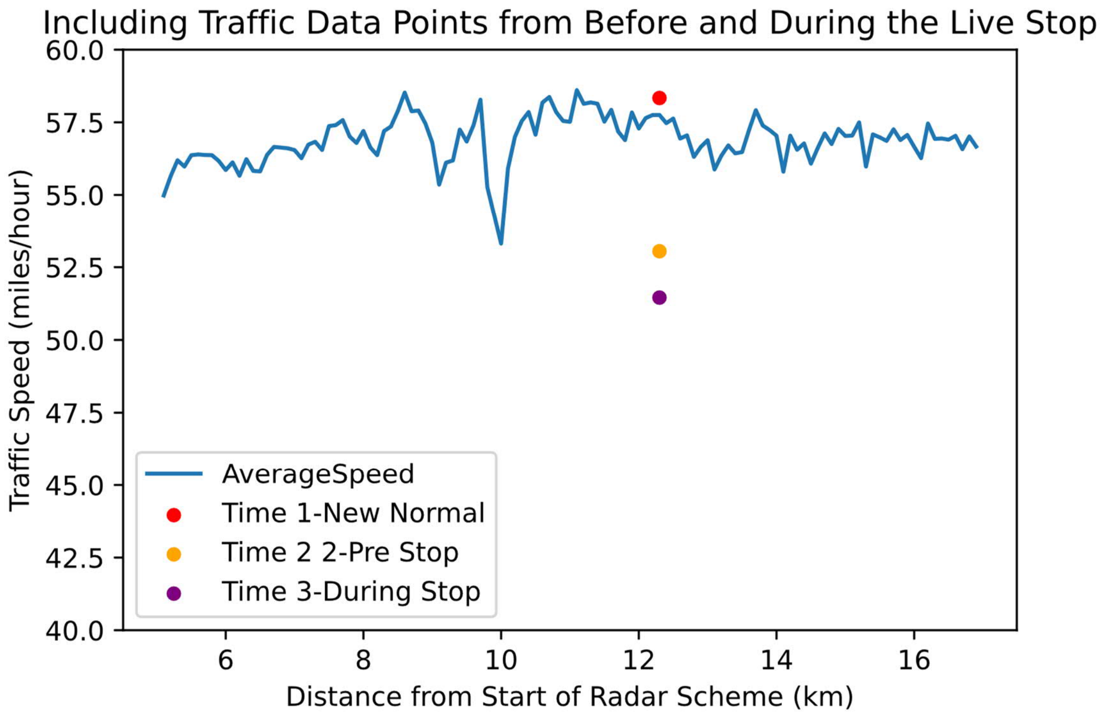

Figure 16: Depicting what "normal" traffic looks like, from 16:00-17:00 on a particular scheme that the Navtech Radar Clearway system is deployed in. Coloured points show traffic data at 12.3km into the scheme, before and during a particular live stop

In the instance where a stopped vehicle occurs on this motorway, one can identify any changes in this traffic behaviour away from a baseline. One example of this can be to look at the traffic speed directly prior to a stopped vehicle occurring.

Using the Navtech Radar Clearway system, a real stopped vehicle was identified on this scheme, during the same time period (16:00 to 17:00), on the same day, but in a different week. This stop was seen to occur 12.3km into the scheme.

The traffic data for this was gathered and aggregated solely in this specific section of the scheme at 3 separate intervals in time. Time 1 – Normal Traffic, Time 2 – Pre Stop, Time 3 - During the Stop. The time intervals of these points are shown in the table below, with the live stop occurring at 16:47 to 16:50.

Time Interval	Label
16:30-16:46	Normal
16:46-16:47	Pre Stop
16:47-16:50	During Stop

These points are plotted on the same graph as above. Point 1 (red) is from a time before the incident had occurred and affected the traffic. That point is within approximately 2% of the normal traffic threshold. Point 2 is the minute before the stop, and Point 3 is the 3 minutes when the vehicle was stopped.

It can be seen that at the pre-stop point (yellow), the traffic speed has decreased approximately 9%. The purple point, which is during the stop, shows a decrease of average traffic speeds in that particular section of the scheme of approximately 11%. This can be considered a change in traffic behaviour that can be identified prior to the stopped vehicle event occurring. However, a drop in traffic speed of only 9% can be seen as too insignificant for use as a direct alert on its own.

It is worth noting that this particular stop was in lane 1 of a 5 lane motorway, with a part of the vehicle hanging in the verge, so this stop was not as obstructive as another stop might have been on a a smaller motorway. The effect can be considered more significant on schemes with a lower number of lanes, as the impact of the stopped vehicle will be greater on the traffic behind it. It is desirable to use larger datasets for such analysis, but acquiring these can be difficult due to sensitivity reasons.

The following example illustrates a similar concept using data from another Navtech Radar Clearway scheme with a lower number of lanes.

In a 2-lane tunnel scheme that the Navtech Radar system is deployed in, two roadwork vehicles were seen to stop in the live lane. The image of this can be seen in Figure 17.

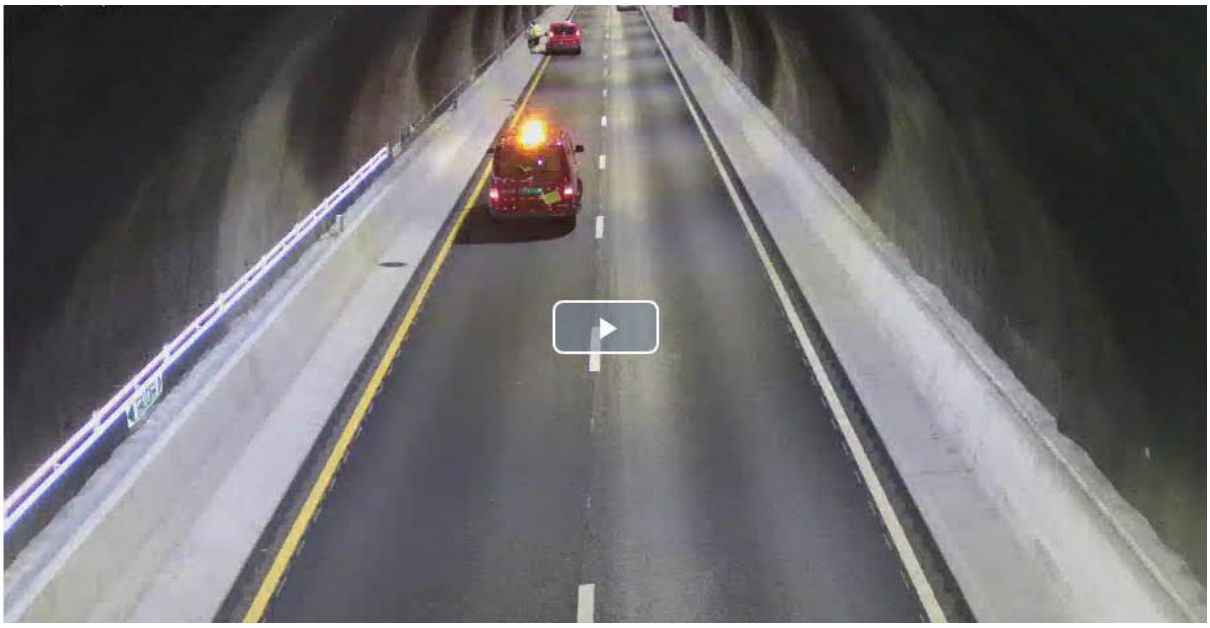

Figure 17: Vehicles stopped in live lane of a tunnel

The traffic conditions were gathered from the Navtech Clearway system during this time and also before the stop. These stops occurred at 13:36 until 13:42. Data was gathered for each of the 12 days prior to this event, during the hours from 13:00 to 14:00, and aggregated over that time to quantify an appropriate baseline of normal traffic in the tunnel. This is shown in Figure 18 below where each coloured line in this graph shows the average speed of a different day from the times of 13:00 to 14:00, with respect to the Section ID in the Clearway System. Each of these Section IDs relate to a separate 60-100m of the tunnel, starting from 4101 and finishing at 4249.

It can be seen that generally the average speed over this time remains at approximately 80 km/h. There are however some dips and anomalies within some of the datasets. For example, section Id 4194 and 4211 have some uncharacteristic drops across some of these data sets. Section 4194 and 4211 shows these dips consistently across all the data sets. These may be due to an anomaly in the reporting radar data and setup configuration, and have been ignored in the following analysis.

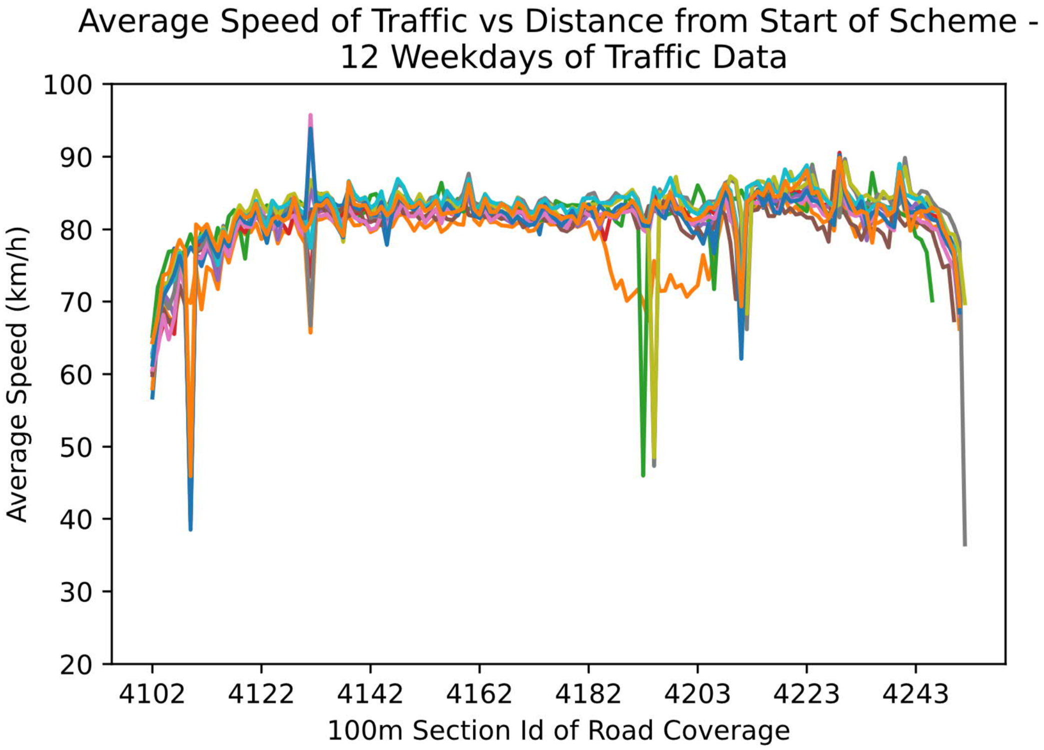

Figure 18: Traffic data showing average speed vs Section ID in Navtech Radar ClearWay Tunnel Project

Figure 19 below then shows the normalisation and averaging of all of the plots with sections 4194 and 4211 omitted. This uses a baseline of 12 weekdays for its normalisation data. From Figure 19 it can be seen that the baseline of “normal traffic” in the tunnel stays relatively constant during the times of 13:00 to 14:00, at 80 km/h. Small fluctuations do exist within this profile, however it is expected that with the aggregation of more data over time, this will eventually flatten this out.

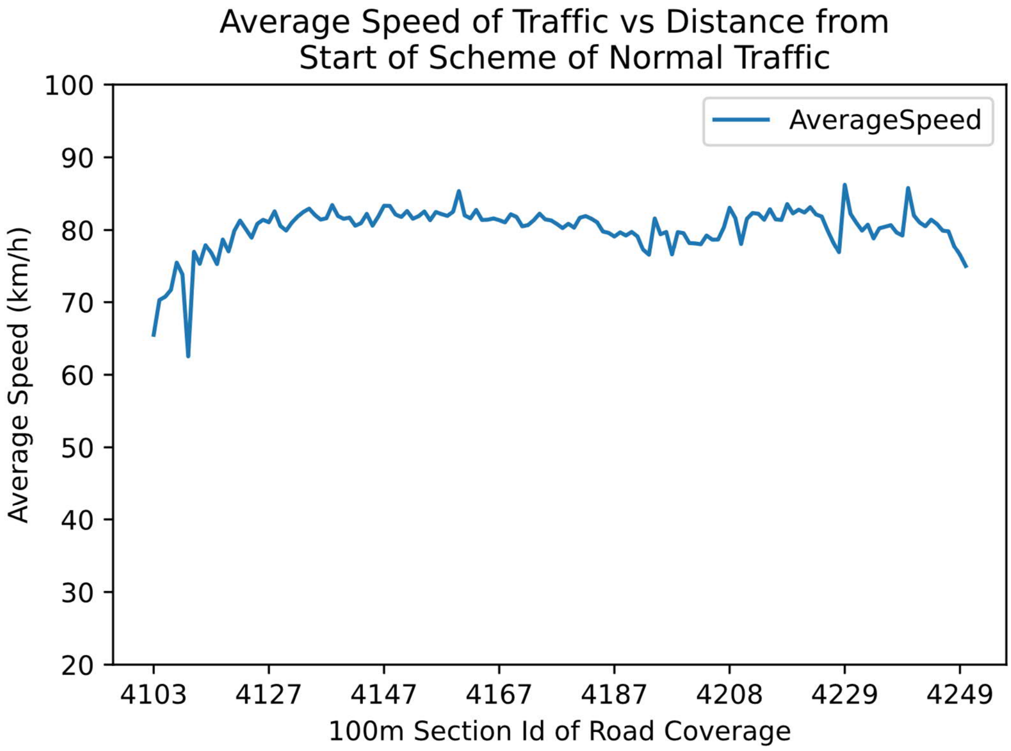

Figure 19: Quantifying "normal" traffic in Navtech Clearway Tunnel Project

Figure 20 below then illustrates the impact the stopped vehicles pictured above had on the traffic speed. The orange graph measures the average speed vs Section ID during the period that there was a stopped vehicle, and the green graph shows this 4 minutes before the stop, with the stopped vehicle occurring in Section ID 4227. It should be noted that the operators managing this tunnel changed the speed limit of the tunnel in these sections from 80 to 60km/h during this time, as a response plan.

In the orange graph, it can be seen that in Section ID 4227, there is a sharp decrease in the average speed, falling significantly below the 60 km/h speed limit to 30 km/h - a 50% decrease. Likewise in the green graph, a smaller drop can be seen in the adjacent section. This can signify the slowing down of traffic as the vehicles approach a stop.

In Figure 20 the pre-stop data is from the full 4 minutes prior to the stop occurring, as this was the minimum duration the system could export. If this time was reduced, there could possibly be more significant speed reductions.

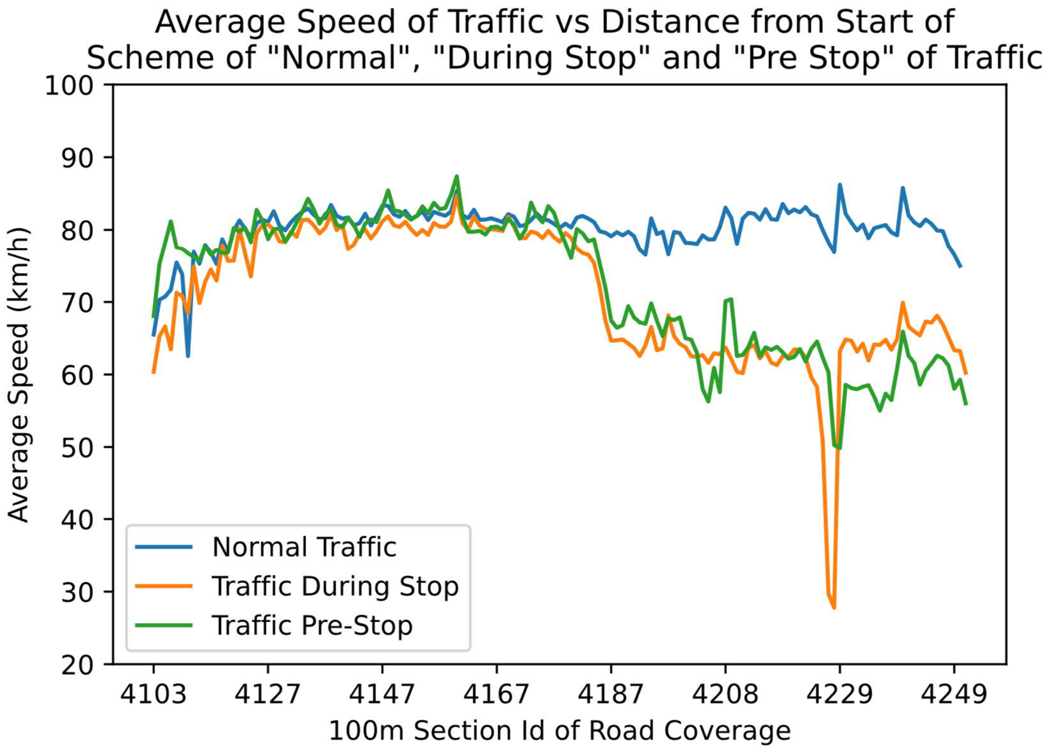

Figure 20: Traffic normal, during stop and pre-stop

A further example was acquired on another scheme that the Navtech Radar ClearWay system is deployed on, where this time it consists of a 1 lane system. A stopped vehicle in the live running lane should have a much greater impact on the traffic behaviour, as there is little to no room to manoeuvre around the vehicle, since there are no extra running lanes. In this case, a stopped vehicle will most likely result in queues and lower traffic speeds from the head of the queue. The use of speed or queue data is far more significant in this instance, and can be more reliably used for additional information of an event occurring on the live lane.

The significance of this result can be seen in Figure 21. Compared to the other graphs shown above, the average speed of the traffic after the occurrence of a stopped vehicle drops almost to 0 miles/hour (blue). After more time, it was observed that traffic was attempting to pass the vehicle, albeit much slower, since there was no extra running lane. Data was not able to be gathered prior to the stopped vehicle event in this instance.

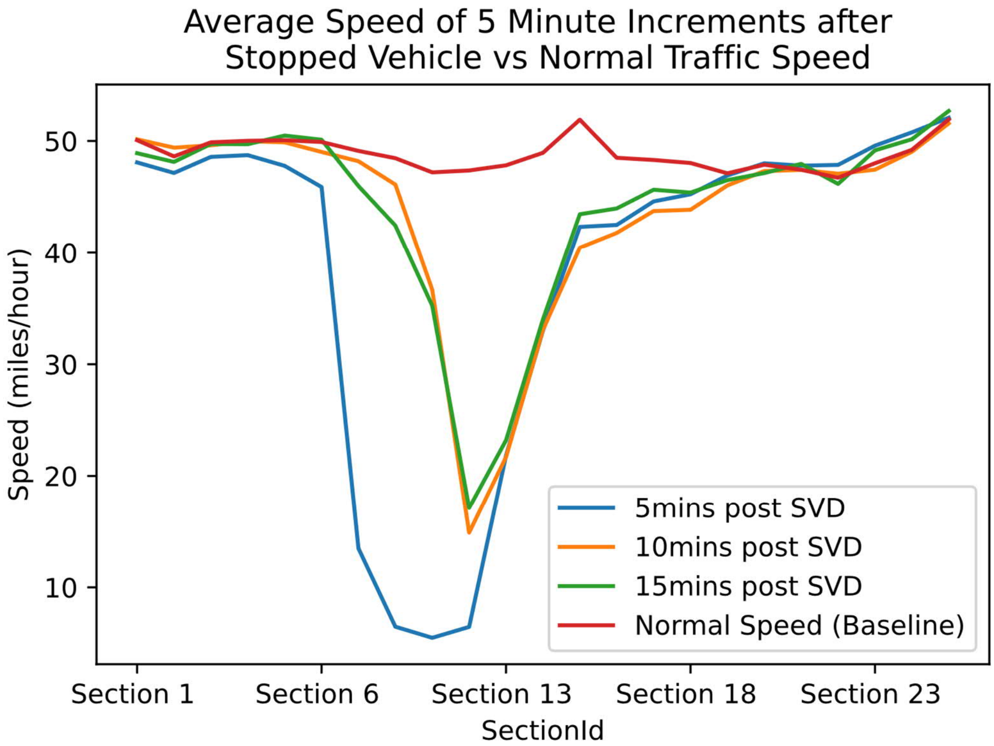

Figure 21 - Average speed of traffic 5, 10 and 15 minutes after a stopped vehicle, in a 1 lane scheme. Compared to normal traffic conditions (red) seen on the same scheme at the same time, but a different day

These plots presented above can be used to further support the detection of an event of a stopped vehicle. In particular, if there exists a misdetection for a stopped vehicle, the use of traffic speed can provide support for this. With traffic speed falling significantly below what is deemed as expected, this can be used to raise a supplementary alert for this incident, or provide a notice of an unusual traffic behaviour. In schemes that include smaller numbers of lanes, this behaviour is amplified, and the traffic speed differences can become a more valuable asset in the detection of a stopped vehicle.

This concept of looking at traffic behaviour before a stop can become more significant in the case where an individual vehicle track is analysed prior to the stop, as opposed to the average speed of the whole traffic prior to the stop. If an individual slow vehicle is identified, can that then lead to a stopped vehicle? If so, how often is that the case for it to become meaningful data?

In a 2-lane, 30km Navtech Radar ClearWay system, a *Slow Vehicle Alarm* was set in conjunction with the Stopped Vehicle Alert. The *Slow Vehicle Alarm* works to alarm the situation where a vehicle drives between 10 and 25 kph, after travelling an initial distance of 30m, plus an additional sighting time of 4 seconds in each 100m virtual section. Over a period of time, both alarm logs for stopped vehicle and slow vehicle were exported, with the aim of identifying if prior to a stopped vehicle, a slow vehicle alarm was triggered in the same Section Id, within a timeframe of 1 minute.

It was seen that approximately 10% of stopped vehicle alarms, had an independent slow vehicle alarm associated with it. This result is quite low, and not considered significant enough to formulate a direct correlation between these two alerts. However, the slow vehicle alarm can be seen to be quite strict for this case, as it is designed to alarm on the case that there is a slow moving vehicle over a larger distance. With this knowledge, the slow vehicle alarm was altered and made much more sensitive. The settings were that the vehicle had to travel between 2 and 25kph and a requirement to only break this rule for 2 seconds within each

100m section. This new rule was run on the system for a period of approximately 3 weeks.

Both stopped vehicle and slow vehicle alarms were then exported after this 3 week period. It was then seen that approximately 50% of stopped vehicles had a preceding slow vehicle alarm attributed to it, which is quite a significant increase from the initial results. It does not capture all the stopped vehicles, but it might still be used to aid the confidence of detection of a stopped vehicle, if a slow vehicle is detected 20 seconds prior, and in the same section. This result cannot however be used in isolation, as there was a large number of slow vehicle alarms generated over this short time span, which if viewed in isolation (ie not in conjunction with the stopped vehicle rule) can overwhelm any party viewing this information.

3.3.2 Queues

On another scheme that the Navtech Radar Clearway system protects, a live stop occurred during busy traffic in the middle lane of a motorway. This can be seen in Figure 22 below.

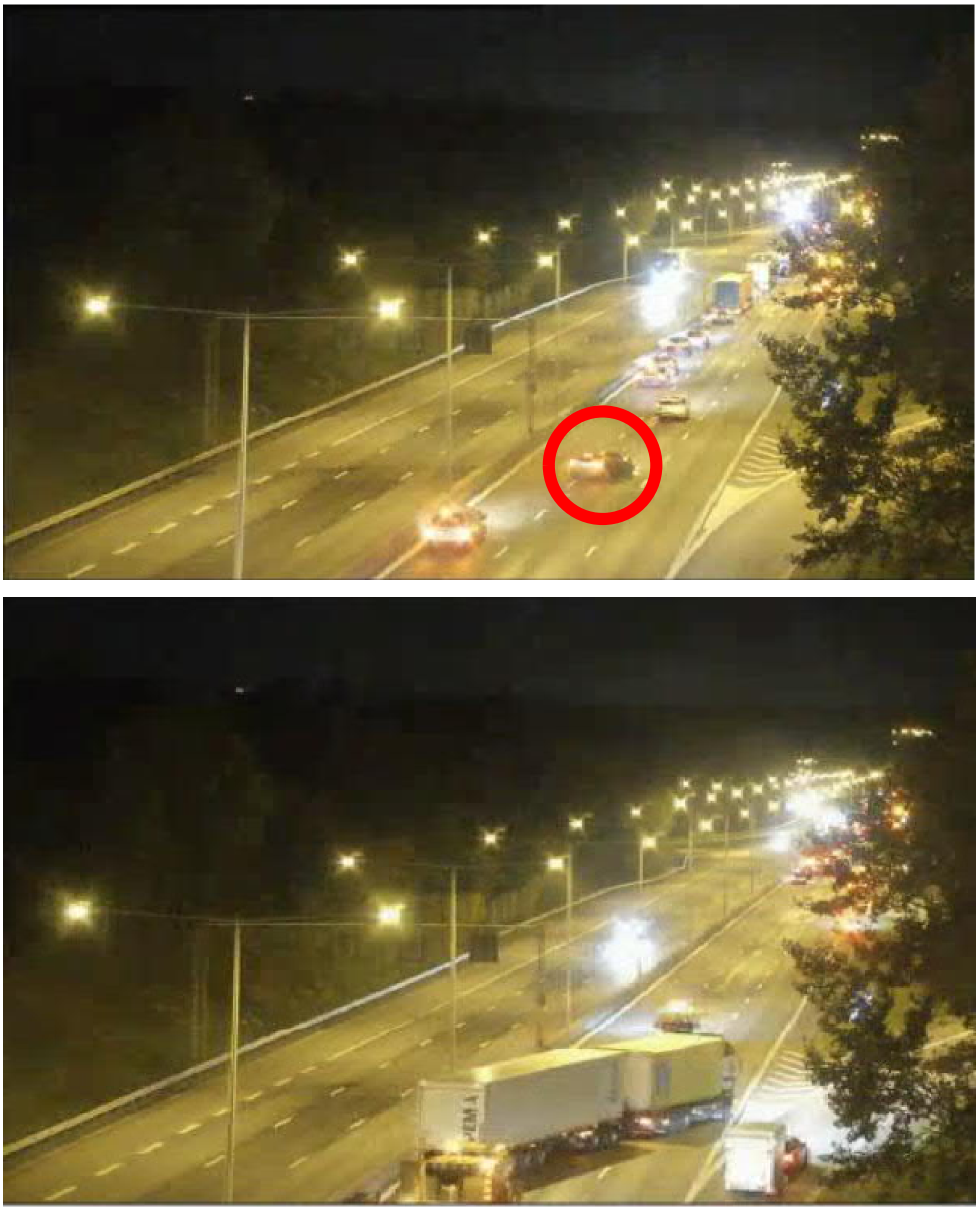

Figure 22: TOP - Live stop in the middle lane on a motorway (pictured within red circle) that is covered by the Navtech Radar Clearway system. BOTTOM – Live stop causing build up of traffic behind it, with large truck requiring to manoeuvre to get past

After this, it can be seen in the bottom picture of Figure 22, a build-up of traffic ensues as the large truck attempts to navigate around the vehicle. The process of detection of queues can also aid in stopped vehicle detection, after the incident has occurred, in order to build more confidence in the alert. If both a stopped vehicle and queuing traffic behind is detected, the confidence in the alarm can be much higher.

A representation of this method is shown in Figure 23 and Figure 24. Figure 23 shows the mean speed by 100m section vs the hour of a particular day of another scheme that the Navtech Radar Clearway system is positioned in. The colour scale on the right shows the different colours representing different speeds in m/s – the darker colours representing slower speeds. The increasing section numbers indicates travel closer to the city centre, where Section 52 being the closest to the city centre. The black points in the plot show areas where data was not able to be collected for that particular time.

It can be seen that during the expected “rush hour” times, that the average speed tends to decrease. This is especially the case of traffic moving into the city from 7.30am, and traffic coming from the city from approximately 3pm on this particular day.

Figure 24 shows the locations of heads and tails of queues (yellow) and isolated stopped vehicles (red dots) on the same day and locations as Figure 23. The queues in yellow show a similar pattern to the Mean Speed plot in Figure 23. This is due to the fact that the queue setting is a function of a speed threshold, and a traffic density per section. This kind of graph could help identify whether stops trigger queues and decreases in speed, but this particular dataset appears inconclusive. A queue forms after a stop at section 41 at 16:00, but the queue appears to propagate from downstream sections, so this may be coincidence.

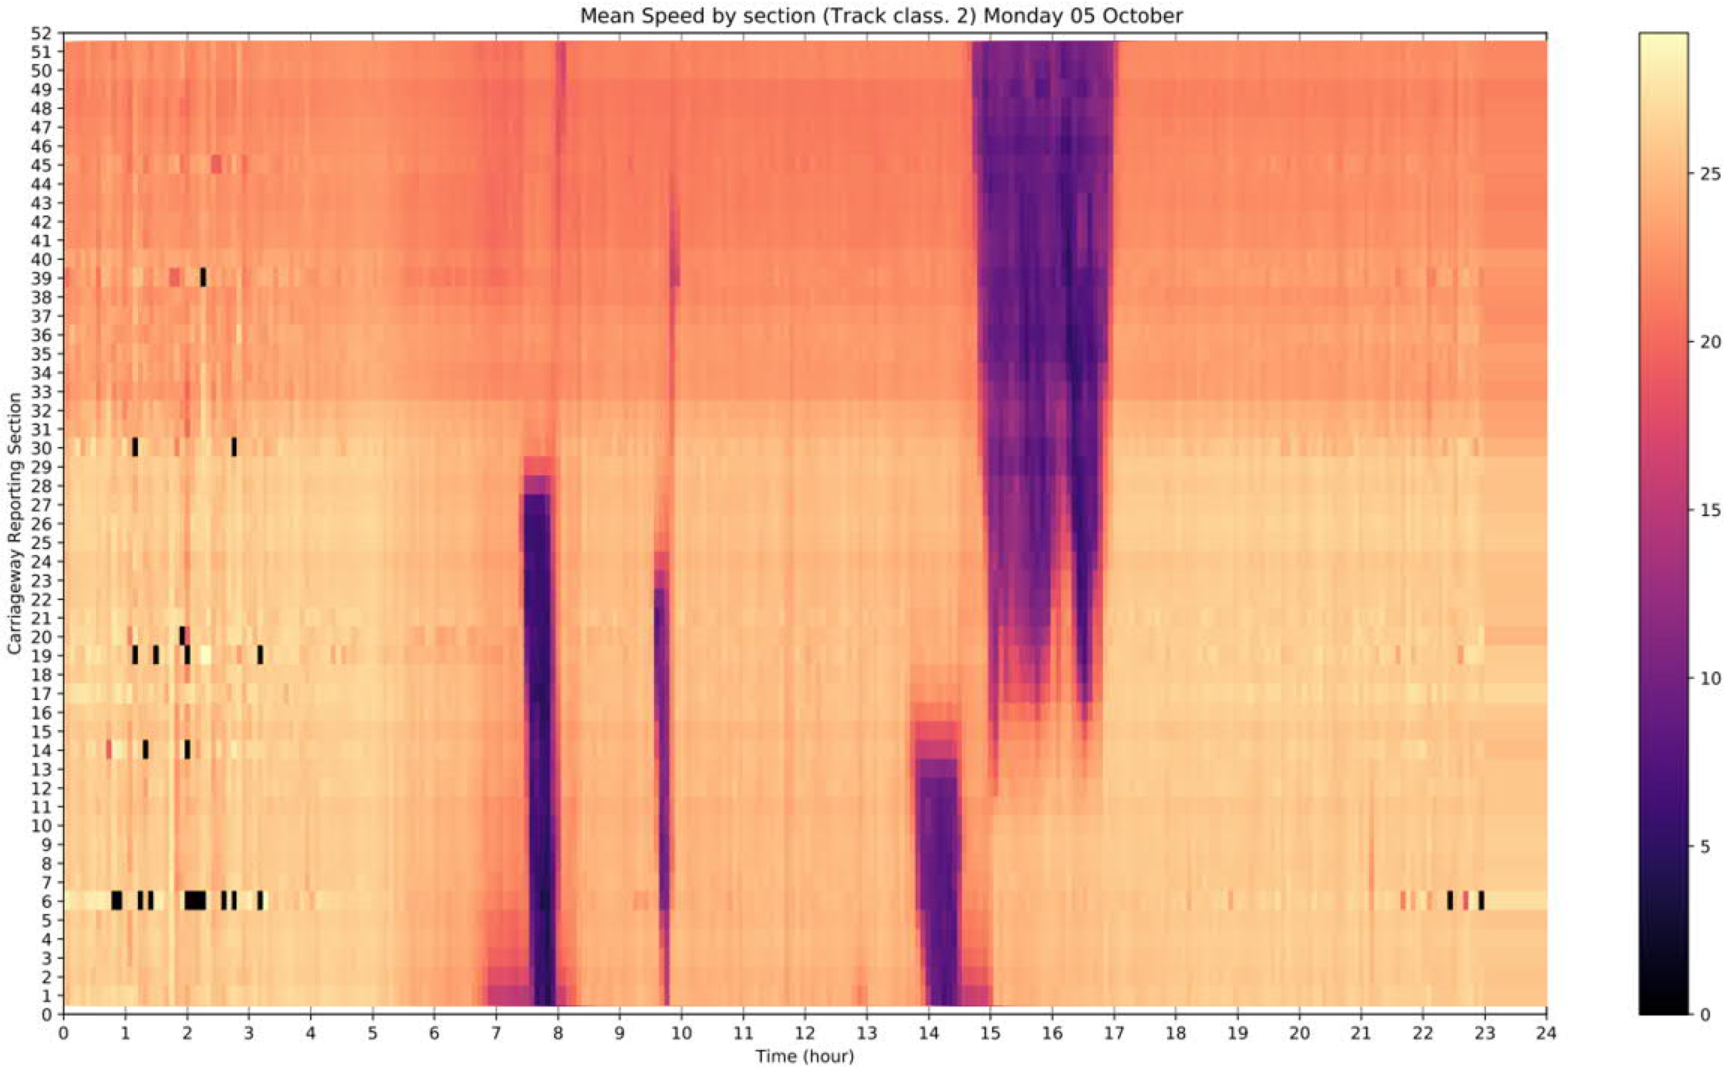

Figure 23: Navtech Clearway Scheme - mean speed by section vs hour of the day

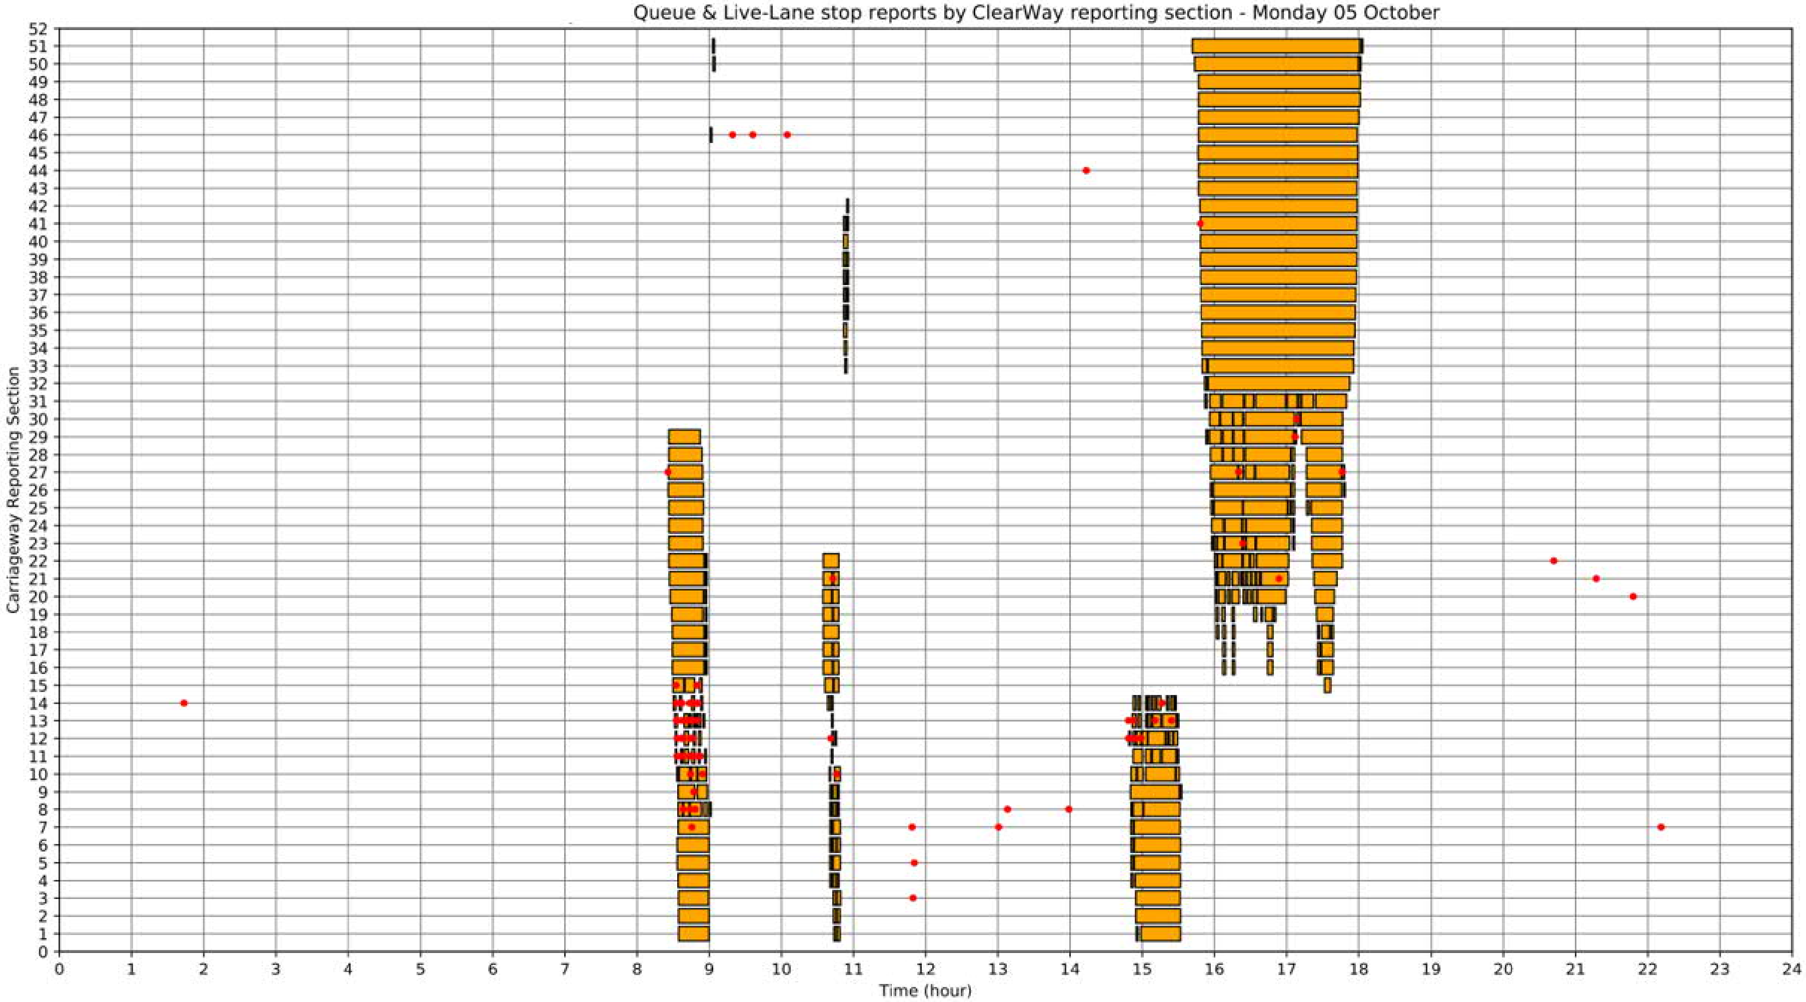

Figure 24: Navtech Clearway Scheme - queue length (yellow) and isolated stopped vehicles (red dots) by section vs hour of the day

3.3.3 Limitations and future work

Knowledge of queues and speed reductions may aid stopped vehicle detection, but limitations and complexities include:

- Base case not consistent enough
    - If the base case that quantifies normal traffic is not consistent, then there will be added difficulty in quantifying what unusual traffic behaviour will look like
- How many “unusual” traffic behaviours lead to a stopped vehicle. Is there a positive correlation for this?
- Added complexity
    - How useful is this information in complex environments?
- Noisy data sets
    - Deriving valuable information and filtering out noise in the data sets

To address the listed limitations above, the next steps for this research topic would be to find more examples of stopped vehicles that have displayed this type of behaviour, to find useful correlations. It may also be useful to derive more quantifiable parameters and to assess the correlations of them in association with real stopped events, or even to explore unsupervised learning methods. Collection and labelling of the individual data sets is a time-consuming activity, complicated by the sensitivities around certain data sets.

3.4 Stopped vehicle classification

A rotating radar system can provide a classification of vehicle type. This information can be valuable. For example, knowing that the stopped vehicle is a fuel tanker may result in a different operator reaction.

The classification of the stopped vehicle can be conducted through an analysis of several inputs collected during the tracking of the object. These parameters include signal strength, size and weight. This however does not include information regarding the engine type, or whether it is carrying any dangerous substances. The overall system however has potential to take information from other sources, to fuse in order to provide comprehensive messaging regarding the type of stopped vehicle.

3.5 Pedestrian information

Further potentially useful information is about passengers of the stopped vehicle. In many cases of a stopped vehicle in a live lane, the passengers of the vehicle will vacate the vehicle, and locate to a safer setting. In some cases, this may not be possible. Therefore, it is important to understand the location of the passengers in each of these events, to understand the direct risk to life and the scale of response required. This will allow emergency personnel and operators to better plan their response.

The Navtech Radar Clearway system can detect pedestrians and other objects within the radar detection area, from a size of 50x50x50cm upwards. Detection of debris, animals and pedestrians has been proven, though distinction between several classifications in the same installation is challenging. Pedestrian locations can be communicated to safety personnel, who can then aid recovery and communication. Therefore, passengers exiting a stopped vehicle can be detected, and their whereabouts communicated.

This passenger location information can be considered crucial in certain cases of stopped vehicles that pose greater danger for the passengers and road users. These are cases when visibility is low on the road, such as extreme rain and weather, where there needs to be a reliance on technology.

A specific example of this is seen in Figure 25 below. This is an image from a motorway in Sweden, an existing Navtech Radar Clearway system scheme. Here, there is a broken down stopped vehicle in the middle of a live lane, which is on fire. This is an extreme case of a broken-down vehicle/stop on a live lane, and a critical safety case that requires urgent action. The location of the passengers of that vehicle here is extremely important, in order to ensure that they have vacated the vehicle before it ignited. As well as that, the thick smoke emanating from the vehicle is of a particular safety concern, as it can significantly reduce visibility, and air quality around the vehicle. This case becomes even more extreme when considering this in a tunnel environment. Therefore, the passenger detection is of great importance in this case, in order to ensure that the emergency crews can quickly and accurately locate them.

Figure 25: Stopped vehicle in live lane in Sweden Motorway, detected by the Navtech Radar Clearway system

The passenger information can therefore be extremely important as part of the message to communicate the general risk of the stopped vehicle, which can then alter the prioritisation of the alert that comes through to an operator.

The example above can also be represented through data gathered by another existing Navtech Radar scheme, which is inside a tunnel. In this example, a motorist had stopped their car in the tunnel, exited their vehicle with a child, and walked inside the tunnel. An image of this can be seen below.

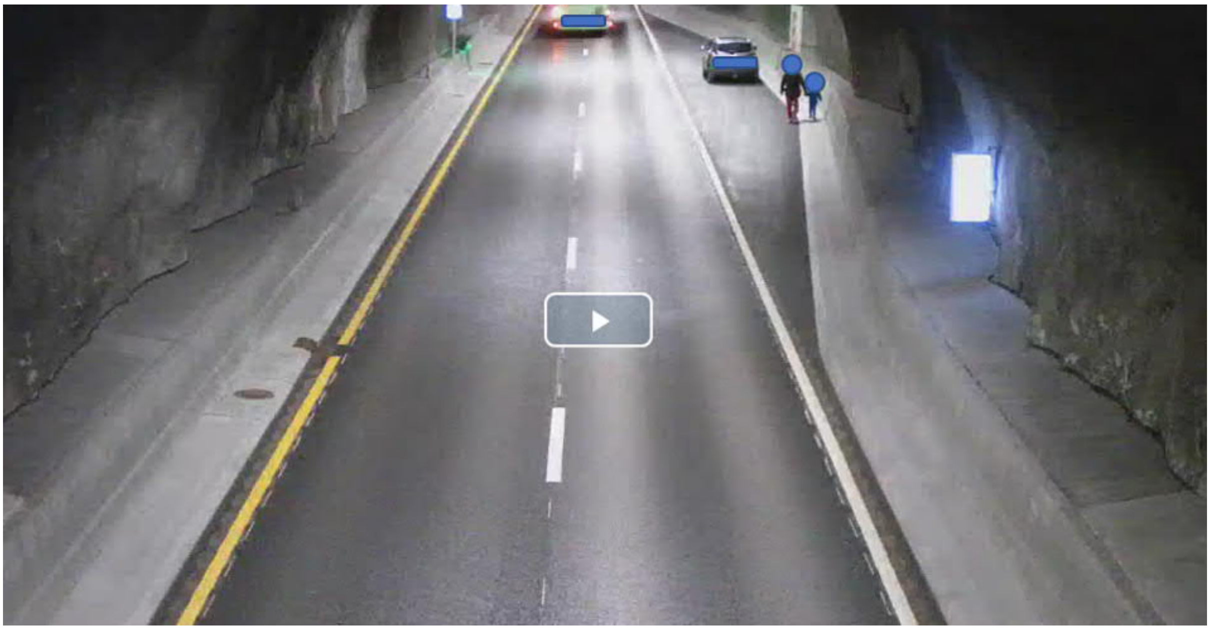

Figure 26:- Stopped vehicle in tunnel, with motorist and child walking to a specific location

The Navtech Radar Clearway system was able to detect and track this, with the pedestrian tracks of both individual pedestrians (yellow and green) extracted and then plotted on an X vs Y coordinate plot, with the stopped vehicle track represented in blue. The X, Y coordinates are arbitrary, expressed in metres and are dependent on the reference points set within the system itself. This can be seen below in Figure 27.

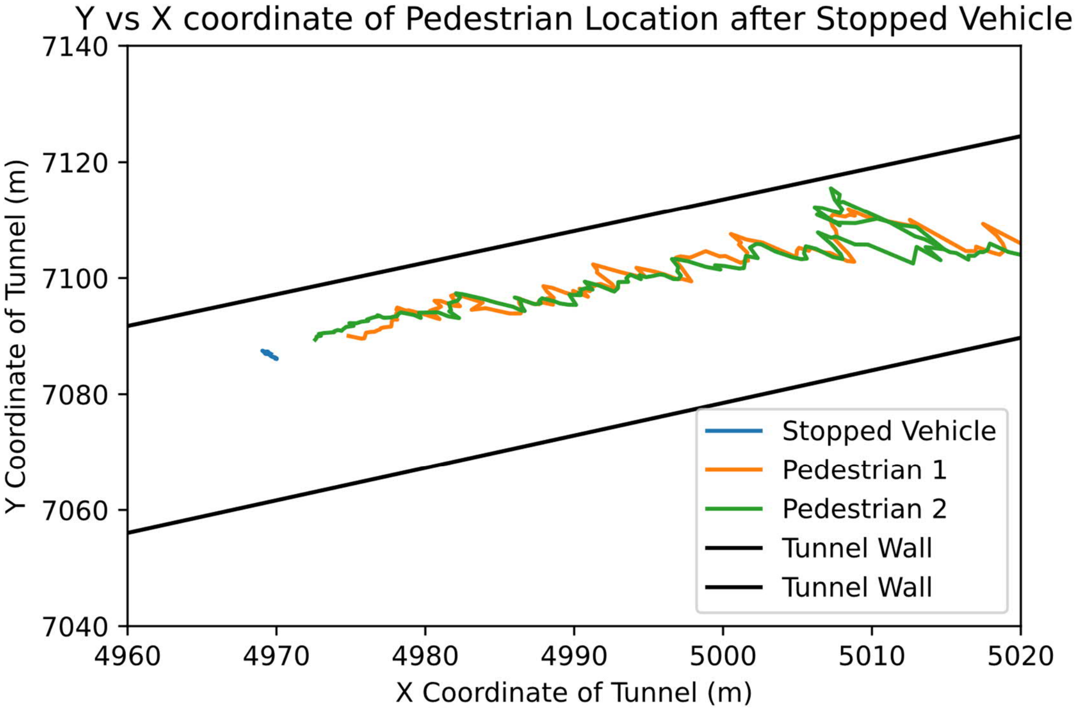

Figure 27: Pedestrian and stopped vehicle tracks of tunnel pedestrians

It can be seen in the graph that the vehicle had stopped in a section of the tunnel, with both pedestrians disembarking, and walking to a specific location of the tunnel. With the Navtech Radar system capable of tracking both stopped vehicle and pedestrians, it can be used to allow emergency personnel and reporters to better plan their response. For example, if the pedestrians in the example above enter a predefined “higher risk” zone of the carriageway, such as a busy and dangerous corner of a carriageway, a tricky bend in the tunnel, an area of low visibility for other motorists, or a section of road without a crash barrier, then it could provide emergency personnel with this information to better understand their recovery plan. This might come in the form of a potential risk level for an alert that could help operators apply priority, and support more appropriate/accurate messaging to the other motorists about the incident.

3.6 Conclusion

The additional data that might be derived from a rotating radar system can be used to move from a simple binary stopped vehicle alerting source towards a more comprehensive dataset to feed into a fusion system as discussed in the following chapter.

The additional data can allow consideration of likelihood and impact risk to better inform operators and emergency personnel and support prioritisation. The probability and impact risk can dynamically increase or decrease from the rotating radar sensor over time based on the events that happen during initial detection, and after. Based on the location and the number of lanes, speeds and queues may increase or decrease the probability of the event. The impact risk level will be calculated from assessing the stopped vehicle’s location and the type of stopped vehicle. If a pedestrian is then detected in the same area and a specified time after the stop, then that can be used to inform this impact level. The location of the pedestrian can then be tracked, and if they enter any pre-defined hazardous areas, this value can be further increased. A flow chart showing this vision as a process is given in Figure 28.

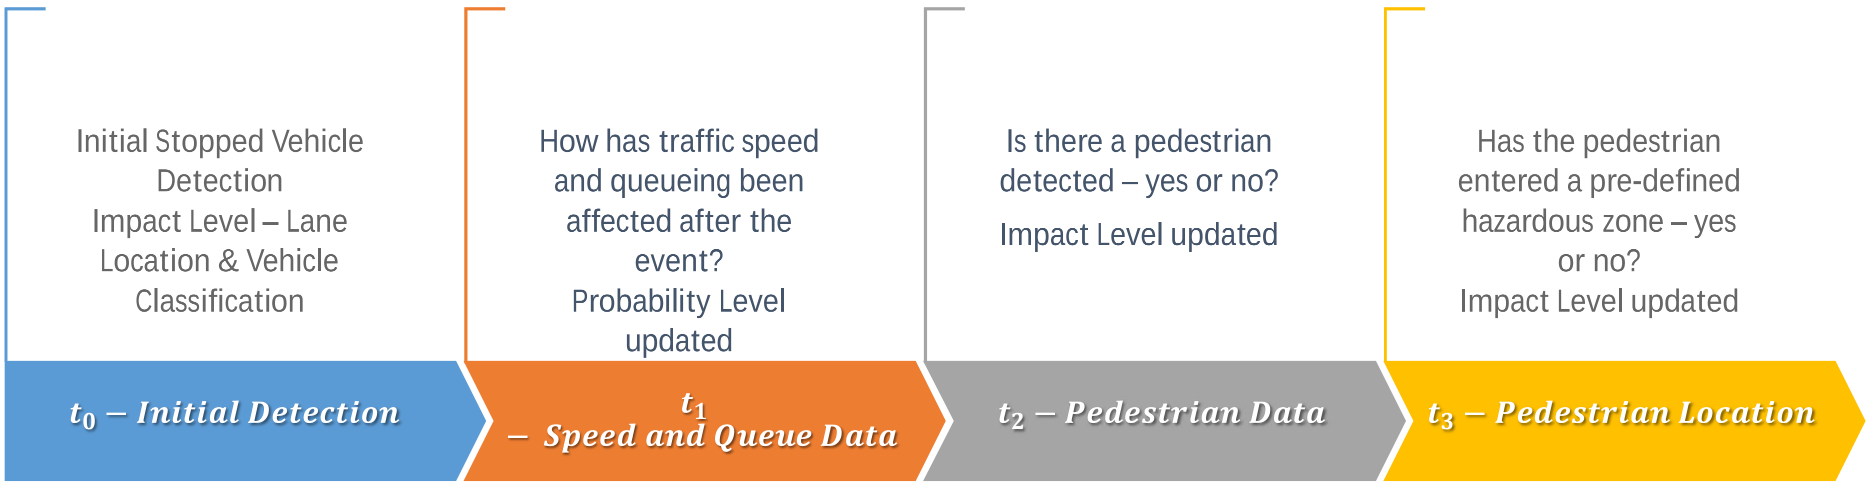

Figure 28: TImeline of Navtech Radar ClearWay detection, with accompanying datasets informing probabilities and impacts

Figure 29 and Table 3 place the research ideas discussed in this chapter into a How-Now-Wow matrix as described in the introduction.

	Subject	Category	Description
NOW	Lane definition	Accuracy	More accurate location information across the width of a carriageway/road
NOW		Information	Retrieving vehicle lane positions helps the TMC respond to an incident more effectively, and can better plan emergency responses and recovery
NOW	Vehicle classification	Information	Retrieving vehicle classification helps the TMC to respond to an incident more effectively, and can better plan emergency responses and recovery
NOW	Pedestrian information	Information	Retrieving pedestrian information after the SVD event helps the TMC respond to an incident more effectively, and can better plan emergency responses and recovery
WOW	**Subject**	**Category**	**Description**
WOW	Impact levels from additional radar info	Information	More information about the impact of the alert, and thus a higher priority can be assigned. This information could influence operator response.
HOW	**Subject**	**Category**	**Description**
HOW	Correlation of traffic parameters	Information	Using data to describe traffic behaviour before, during and after an SVD event can help provide useful insight into the impact of traffic parameters on an alert, which may ultimately feed into AI models for improved/predictive detection
HOW	Confidence and probability levels from additional radar info	Reliability Accuracy Information	Confidence and probability levels in alert could increase detection rate and reduce false alarm rates. More accurate reporting of the incident.

Table 3: How-Now-Wow categorization for radar improvement.

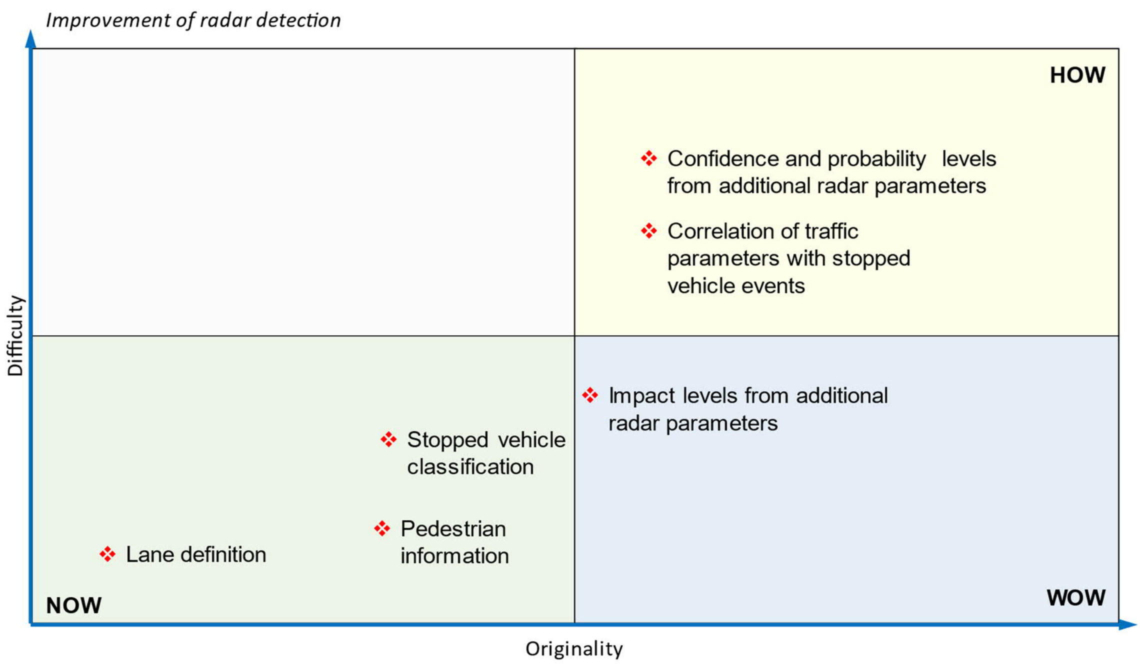

Figure 29: How-Now-Wow matrix improvement of radar detection

4 Future and upcoming methods

This chapter complements SHADAR report D2.1, which described existing deployments of stopped vehicle detection on European roads. The chapter examines methods that are not yet currently deployed, or which are the subject of current initiatives that suggest they will grow in scale and practicality.

4.1 Availability of Safety-Related Traffic Information

Bound by the European regulation 886/2013 of 15 May 2013, public road operators, service providers and broadcasters dedicated to traffic information are obligated to share safety-related traffic information. In the regulation, the type of information and the shape is prescribed. The traffic information should consist of at least one of the following categories:

1) temporary slippery road;
2) animal, people, obstacles, debris on the road;
3) unprotected accident area;
4) short-term road works;
5) reduced visibility;
6) wrong-way driver;
7) unmanaged blockage of a road;
8) exceptional weather conditions.

4.1.1 Data For Road Safety

Road authorities and other data providers are now participating in the Data For Road Safety (DfRS) initiative (www.dataforroadsafety.eu) which aims to make SRTI available for all road users in Europe. The initiative is the successor of the Data Task Force.

The SRTI ecosystem is based on obtaining sensor data from the vehicles (called "Level 2" or "L2" data) which is enriched and aggregated to obtain information usable by the road authorities (called "Level 3" or "L3") (Ismail, 2020). In Figure 30 the SRTI ecosystem is shown. The eight categories mentioned earlier from the delegated regulation are adopted.

- In-Vehicle User Interface Element triggered by customer - Example: wiper, manual breakdown call)
- In-Vehicle User Interface Element triggered by vehicle regularly - Example: ABS Lamp, Stability Program Lamp)
- In-Vehicle User Interface Element triggered by vehicle rarely - Example: automatic e-call, automatic breakdown-call
- Simple Sensor Reading, minimally processed - Example: temperature, friction value representing a known physical value
- Locally simple combined sensor data - Example: sending ABS only if brake force <x
- Locally complex fused sensor data - Example: rain density by locally fusing wiper frequency and rain sensor data with speed and windshield angle
- Complex object detection - Example: object detection by camera

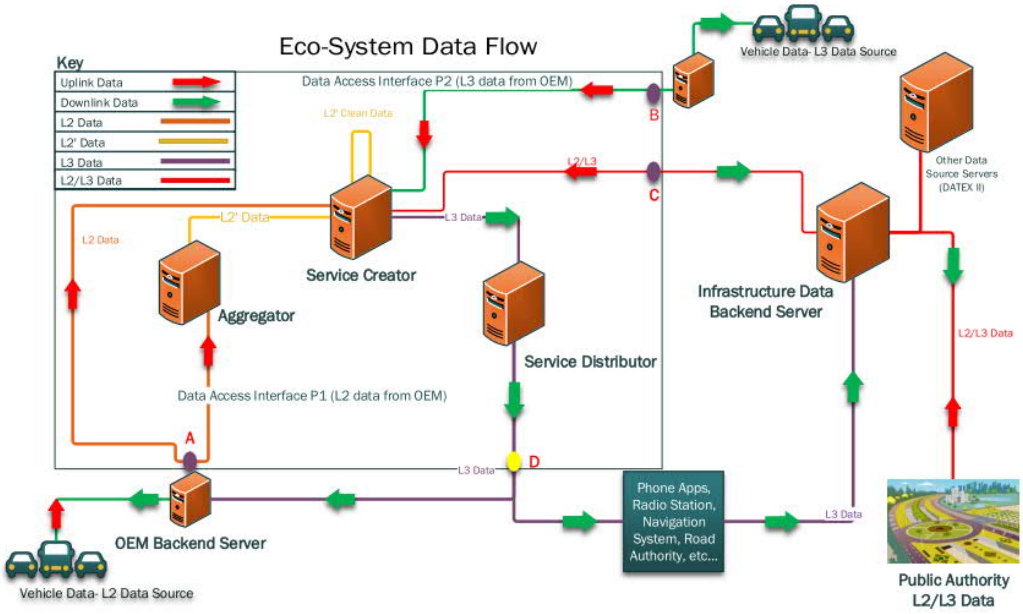

Figure 30: DfRS SRTI Ecosystem (Ismail, November 2020)

A proof of concept was held from June 2019 until October 2020 (van Rij, 2020) to test the exchange of STRI messages between private and public parties. The vehicle makers retain ownership of the information, but it is provided under certain conditions to the participating NRAs. Level 2 data was supplied by BMW, Daimler, Ford and Level 3 data was supplied by Nira, TomTom, and Volvo. The NDW in the Netherlands provided the central testbed.

The evaluation report focuses mainly on the Netherlands because all parties provided data for this purpose. Some parties provided data for the whole of Europe while others only provided data for the Netherlands. Dataflows gradually came available during 2020. In the evaluation, only data from June and July 2020 are analysed due to privacy and commercial reasons.

The notifications per party are clustered into the 8 types of SRTI messages. The notifications "ABS active" (found both in braking and accelerating) "vehicle in difficulty" (unclear) and "emergency vehicle" (also unclear) were included separately because it was not clear to which of the 8 types these belong. During the trial mainly data about unprotected accident areas and animals/people/obstacles/debris on the road (broken-down vehicles) were reported. Figure 31 shows the number of notifications per party.

Number of messages per type and party (The Netherlands)

	Level 2			Level 3		
	Party A	Party B	Party F	Party C	Party D	Party E
Adverse weather condition	32.790	662.896				
Animal, people, obstacles, debris on the road		92.105	6.048		48.305	
Reduced Visibility	1					
Slippery road				16.541		2.051
Unprotected accident area		664			100	
ABS active		41.262				
Emergency vehicle					2	
Vehicle in difficulty						81.530

Figure 31: Number of notifications per party and type within the Netherlands (van Rij, 2020)

Reports on slippery roads and exceptional weather conditions could not immediately be passed through as an SRTI message due to the large amount and the necessity of post-processing. Vehicle data does not yet contribute to the following SRTI types:

- Short-term road works;
- Unmanaged blockage of a road;
- Wrong-way driver.

In most notifications, the L2 parties also provide vehicle traces, with the last series of positions before a notification/incident. The different data providers each supplied different numbers of positions in vehicle traces. Van Rij concludes that the quality of vehicle traces was usable from two of the three parties.

Average number of points in vehicle trace

	Party		
	Party 1	Party 2	Party 3
Adverse weather condition		3,00	20,98
Animal, people, obstacles, debris on the road	24,48		21,00
Reduced Visibility		3,00	
Slippery road		3,00	
Unprotected accident area			3,81
ABS active			20,98

Figure 32: Average number of points in vehicle trace (van Rij, September 2020)

During the trial the timeliness was examined between the registration of the incident by the vehicle and the available message at the national access point (NDW):

- 52% within 5 seconds
- 85% within 1 minute
- 96% within 5 minutes

Figure 33 shows the latency per party. Party VI obtains data through the NDW, improves these to L3 data and returns them to NDW. The original time is retained by Party VI.

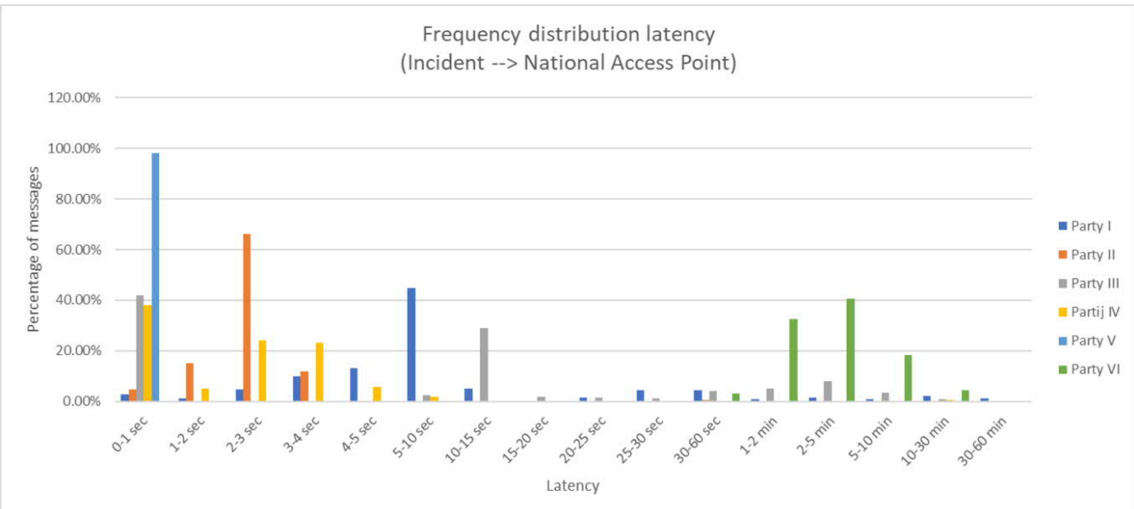

Figure 33: Latency frequency distribution (van Rij, 2020)

The report does not mention any analysis of latency, however mentioned earlier in the report, temporary slippery road notifications are not immediately forwarded, mainly due to a large number of notifications that require post-processing. There might be a relationship between latency and type of message, which is however not clear (yet).

Looking at the results regarding broken down vehicles and accidents, only one party delivered reports of accidents during the trial. In the evaluation, the timeliness was compared with existing sources. For a strict selection of broken-down vehicles, the time-saving amount was up to 7.5 minutes compared with the available data from reported incidents in the NDW dataset.

One of the outcomes of the trial is that it is unclear which sensors trigger the generation of L2 messages. Generated messages differ between car manufacturers and are not always in line with the categories established in European regulation. For stopped vehicle detection more insight into message generation could help support filtering for relevant messages. Further study of the penetration rate, fleet size, and growth rate of these connected vehicles should give more insight into the usability of this data source.

4.1.2 C-ITS safety related messages

C-ITS services are standards-based exchanges of data between vehicles, the roadside and urban infrastructure, control and service centres and other road users (ISO/CEN, June 2020).

In the C-ITS domain many organisations are working on investigating, testing, simulating and performing pilot implementations as well as on nationwide implementation of day 1 and day 1.5 services. Short desk research examined the current state of services and the expected usability for stopped vehicle detection within the next five years.

From C-Roads (2022) the following safety-related services are in focus:

- Wrong-way driving
- Emergency brake light
- Other hazardous notification
- Obstacle on road
- Emergency vehicle approaching
- Slow or stationary vehicle(s)
- (unprotected) accident area

In Europe and beyond many research and pilot projects have been or are working on the C-ITS topic, and the first national implementations are in progress. In the deliverable D5.1 of the H2020 program TransAID (Rondinone, 2018) an overview of C-ITS messages is provided and analysed for use in the TransAID project. The following messages, see Table 4, are relevant for stopped vehicle hazards.

Acronym	Message	Depl.	Description
CAM	Cooperative Awareness Message	Day1	Maintain awareness of each other and support cooperative performance using the road network (V2X) for Day info about Speed and slow-changing vehicle data
DENM	Decentralized Environmental Notification Message	Day1	A facility layer message with road hazard or abnormal traffic conditions (event and position)
CPM	Collective Perception Message	Day2	CP messages about Locally detected objects to improve situational awareness
MCM	Maneuver Coordination Message	-	Coordinate maneuvers between stations (early stage of development) (V2X)
IVIM	In-Vehicle Information message	Day1	Information about infrastructure-base traffic services (e.g. dynamic road and traffic signs)

MCDM	Multimedia Content Dissemination Message	-	Sharing of multimedia content between ITS stations to improve environmental perception (not yet considered by C-roads and C2C-CC for deployment)

Table 4: Messages relevant to stopped vehicle hazards, source: (Rondinone, 2018)

CAM

The Cooperative Awareness Message (CAM) is a message that creates and maintains awareness of ITS stations between themselves and support cooperative performance. A CAM provides information about presence, position, dynamics and basic attributes of the originating station. Based on the position and dynamics an ITS station can calculate for instance, a collision risk. The H2020 MAVEN project designed an extension for enabling platooning. In the I2V extension, also the number of occupants was included, which is relevant safety related information. CAM informs more about the general parameters of an ITS station and is not specifically an SRTI message type.

DENM

A DENM message contains information related to a road hazard or abnormal traffic with type of event and position. The message is used to alert other road users about the occurrence of an unexpected event. An optional data container of DENM is the Fields of impact container which can include the position of occupants. DENM is typical the message type to be used for SVD messaging for an ITS station.

CPM

For SVD the CPM message set is of interest due to the fact that CPM standardizes communication via an abstract representation of detected objects instead of raw non-standardized sensor data. In addition, the CPM messages can contain an abstract description derived from a single sensor as well as the result of a (local) sensor fusion algorithm. This provides a flexible implementation. The Collective Perception is designed to use both Vehicle and Road Side Unit (RSU) detections. For the RSU a reference point is used with east and north alignment. To allow mapping origin stations will always transmit information about their coordinate system. Also, every origin station will communicate their detection capabilities in terms of the sensors Fields of View by which a cross-check can be done if by the receiving station if the originating station has sensors covering its direction.

CPM is typically a message type that will/can be part of an SVD detection. When detecting a hazardous situation, the information CPM communication can lead to a DENM message.

MCM

The MCM (Maneuver Coordination Message) message set is currently in the definition phase (ETSI TC ITS). The MCM can be used to coordinate manoeuvres between ITS stations in the future when connected and automated vehicles encounter a stopped vehicle and can get assistance to pass by the incident location. MCM is a supportive message which can be used to let connected vehicles cope with a stopped vehicle hazard.

IVIM

The In-Vehicle Information Message (IVIM) is an infrastructure-to-vehicle communication to convey infrastructure-based traffic messages focussing on road safety and traffic efficiency. For the first deployment C-Roads, the C2C-CC agreed on adopting the IVI profiling. The EN ISO 14823 graphic data dictionary standard is adopted which contains standardized codes for existing signs and pictograms for traffic and traffic information. IVIM can bring the warning of a stopped vehicle into the connected vehicle.

MCDM

ETSI TC ITS is defining the Multimedia Content Dissemination Message (MCDM). The target is to share multimedia content between ITS stations like pictures and videos regarding obstacles on the road, traffic conditions to improve the environmental perception but also regarding local services and products. MCDM can convey a warning for a connected vehicle regarding a stopped vehicle notification.

4.1.3 C-ITS implementation status

This section describes the status of C-ITS piloting and implementation that may be relevant to stopped vehicle detection.

C-Roads

Within the C-Roads platform, a large number of use cases are supported. HLN – Stationary vehicle (Working Group 2, October 2020) is the one that covers the signage of Stopped Vehicle Detection. Slow or stationary vehicles use cases are implemented on around 20 pilot sides in 13 countries through Europe. A slow-moving or stationary vehicle signals its presence to other vehicles and also, in general, ITS stations that could be roadside units that can convey the information to the TMC.

Pilot site	SRTI related info
Austria	Deployment of hazardous Location Notification is set up in 2021-2022 in which the TMC of Salzburg can notify connected vehicles of dangerous situations
Belgium Flanders	Main objective is to connect road users with the TMC to create ad direct interaction. Combination of Here Location Cloud and TMC allowing 1000 test drivers to receive SRTI messages. Pilot location is on the core network in Flanders. 650 users were reached.
Belgium Wallonia	Pilot is set up in Liege with communication by RSU and on the highway network with service provider Coyote. End of 2020, the sites entered the last phase before deployment. Other hazardous notification and Slow or stationary vehicle(s) are among the use cases
Czech	Other hazardous notifications and Slow or stationary vehicle(s) are among the use cases. Intensive testing was done on the communication part and users throughout the Czechia.
French	No SVD relevant use case was analysed
Germany (Hessia, Lower Saxony)	In Hessen and Lower Saxony no SVD related use case will be deployed other than Emergency Vehicle Approaching Service Deployment (EVA)
Greece	Greece is a newcomer. SVD related use cases are -Stationary Vehicle (HLN-SV) - Obstacle on the Road (HLN-OR). The development of test sites is ongoing starting now on the procurement and development
Hungary	Mainly focussing on road safety at work zones. The existing M1 pilot will be upgraded, and a lot of use cases/sub use case will be tested (not clear at the moment what….)
Denmark	Participation through NordicWay
Finland	Evaluation of C-ITS services is completed looking at technical feasibility, ecosystem, business models, socio-economic impacts. Slow or stationary vehicle(s) & traffic ahead warning was tested on two pilot sites and cooperative collision risk warning on one site.
Italy	Italy’s use cases are not SVD-related
Ireland	Working on test sites to enable and evaluate a range of day 1 and day 1.5 services
The Netherlands	No SVD related use cases
Norway	See Finland

Portugal	A number of sub-activities covers the development of a SPA mobile app for slow or stationary vehicle and other hazardous location notification. The organisation and rollout of systems is ongoing
Sweden	In the Swedish national pilot Hazardous Location notification emergency brake light, Emergency vehicle approaching, and other hazards will be tested in NordicWay pilot 3.
Slovenia	Mainly the rollout of the C-ITS landscape
Spain	No evidence found related to results of SVD detection
UK	Focus on cross-border harmonisation (InterCor)

Table 5: Short description of status of C-Roads Pilot sides 2020 related to SVD

The status of the pilot sites was reported in the beginning of 2019, (Kernstock, 2019). Results/evaluation of the pilot sites became available November 2021 (Gruber, 2021). A short summary is put into Table 5. Whether any of the pilot sites implemented a use case with reporting towards a TMC of some sort is not elaborated in the documentation. Analysis of the detection rate or false positives in the use cases are not mentioned in the annual report. The focus is on the roll-out of C-ITS pilots and on technical and functional operation.

C2C-CC

The CAR 2 CAR Communication Consortium (C2C-CC) guides C-ITS developments with the aims of supporting accident-free traffic (vision zero) at the earliest possible date. It further aims to support the highest safety level at improved traffic efficiency anywhere, anytime at the lowest cost to the end-user and the environment.

Its definition and development phase of Day 2 services and technologies is ongoing. Its ”Roadmap Day 2 and beyond” (2021) sets out the next steps in the deployment. *‘In the Day 2 phase, vehicles and RSUs will take advantage of being equipped with environment-sensing technologies to share information about detected objects. This capability will enable receiving vehicles to be aware of obstacles they would not otherwise detect with their own sensors (e.g. pedestrians or cyclists hidden behind a corner in intersection areas). In this way, enhanced safety applications compared to the Day 1 can be introduced.’*

In the Day 2 development and the rollout phase especially attention to the extended DENMs and the further development of CPMs and CPS allows sensed driving and advanced warning use cases.

C2C-CC, alongside DFRS, produced in January 2021 a new version of a document for safety-related EC high-level categories with the means to provide insight and understanding of the current state of SRTI event selections throughout Europe and how these events can be expressed in the mainstream standards in use (C2Car CC, January 2021).

TransAID

The H2020 project TransAID investigated and developed traffic management procedures to support connected Automated Vehicles (CAVs) to enable the coexistence of automated, connected and conventional vehicles. Protocols, communications and message sets were developed.

Use cases were developed, simulated and tested in the real world. During implementation and testing traffic measures, TransAID also identified and created the needed messages sets and protocols to implement V2X communication. The CAM, DENM, MCM and MAPEM were extended. The use of CPM messages for enabling Collective Perception Service (CPS) was investigated. TransAID proposed policies to further optimize the CPM regarding content and transmission triggering to achieve the necessary level of redundancy and minimize the impact on the implementation, stability and scalability. Work on analysing types of sensors and their properties and performance and fusing capabilities of vehicles and roadside equipment was used to assess the way CPM is deployed and CPS is enabled (Correa, 2020).

The use of CPM for Collective Perception and DENM messages to alert the driver about upcoming hazardous situations was tested in several services. The use case dealt with potentially hazardous situations by providing guidance to connected and connected automated vehicles to be able to deal with the occurring situation or to support a timely transition of control and/or guidance to a safe spot. Figure 34 gives an overview of the services and use cases of the TransAID study (Wijbenga, 2021), (Wijbenga, 2020):

- Use case 1.1 - Provide path around road works via bus lane
- Use case 1.3 - Queue spillback at exit ramp
- Use case 2.1 - Prevent ToC/MRM by providing speed, headway, and/or lane advice
- Use case 2.3 - Intersection handling due to incident
- Use case 3.1 - Prevent ToC/MRM by traffic separation
- Use case 4.2 - Manage MRM by guidance to safe spot (& lane change assistant)
- Use case 4.1+5.1 - Distributed safe spots along an urban corridor
- Use case 5.1 - Schedule ToCs before no-AD zone

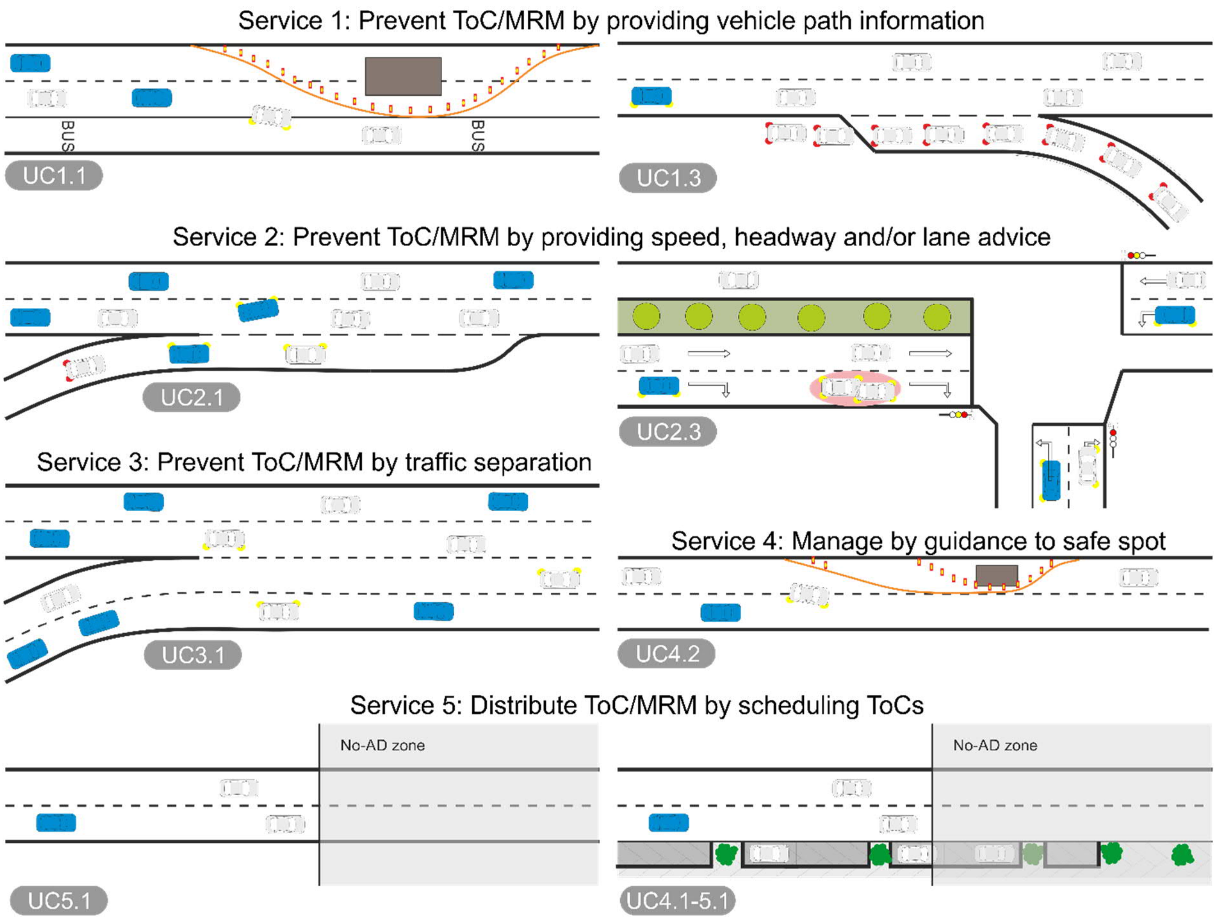

Figure 34: TransAID service and investigated use case overview (Wijbenga, 2021)

All the use cases covered potentially hazardous situations and tried to deal with this in different simulated mixed traffic scenarios (automated, cooperative, and conventional vehicles). This was simulated with ideal communication and in a second iteration also with realistic communication (with errors, distance problems between V2X etc.) to see the impact on the traffic flow coping with non-standard situations for the connected vehicles (Lücken &

Schwamborn, 2021). When automated vehicles cannot cope with the situation a take-over request is made to the driver. If a transition of control is not issued in time, the vehicle will perform a minimum risk manoeuvre coming to a full stop. Roadside units can assist the manoeuvre to guide a vehicle to a safe spot. If not, the full stop will be performed within the traffic flow.

The concept of using sensor information (CPM messages) to obtain a common ground for the CPS and the use of DENM are of interest for SVD. The further implementation and rollout of these C-ITS services should be followed, and there may be a case for increasing the priority of stationary vehicle hazard-related applications within C-ITS implementation projects.

C-ITS and Data for Road Safety

DfRS and C-ITS are, at the moment different eco-systems, albeit with some overlap. DfRS addresses the business/cooperation model which business partners can join. The C-ITS eco-system has a different architectural paradigm. Neither has yet achieved full pan-European vehicle coverage.

4.1.4 SRTI data feed service providers

This section elaborates on the safety-related data services available from some commercial service providers.

TomTom

The traffic incidents API from TomTom provides a list of reported incidents within a given area. TomTom gives the ability to track the following road incidents in this API:

- 0: Unknown
- 1: Accident
- 2: Fog
- 3: Dangerous Conditions
- 4: Rain
- 5: Ice
- 6: Jam
- 7: Lane Closed
- 8: Road Closed
- 9: Road Works
- 10: Wind
- 11: Flooding
- 14: Broken Down Vehicle

The TomTom incidents API does allow filtering on category, which means specifically for SVD, it is possible to filter on categories 1, 14 and possibly 0. Additionally, the TomTom incidents API logs the length and delay of a possible traffic jam, as well as estimated intensity (1-3 for minor to major, 0 for unknown, 4 for undefined), and logs the likelihood of a report being true, within the ACI field. The full content of the TomTom incidents API response is shown in Table 6.

TomTom collects this data from a wealth of traffic information, covering 79 countries and collecting from over 600 million devices, including GPS devices, mobile phone signals and sensors. The costs for retrieving incidents are shown in Figure 35.

Traffic

	Endpoint /Service	PRICE PER 1000 REQUESTS
Traffic API		
	Traffic incidents non-tile	$0.50
	Traffic Incidents tiles	$0.05
	Traffic Flow non-tile	$0.50
	Traffic Flow tiles	$0.05

Figure 35: Costs per request for traffic information, TomTom (2021)

BeMobile

The BeMobile incident API provides a list of both reported incidents and automatically detected incidents. The reported incidents are collected through their Flitsmeister app, which has nearly 2.3 million monthly users across the Netherlands. BeMobile additionally detects incidents based on FCD (Floating Car Data). Their data is supplied with fields such as location coordinates, heading and category coding in AlertC or DATEX II format.

BeMobile provides data on accidents, traffic jams, open bridges, incidents, roadworks, road closures, etc. for countries like Slovenia, Norway, Denmark, Sweden, Turkey, Luxembourg, France, Germany, Belgium, Netherlands, Poland, Finland.

HERE

The HERE incidents API, similarly to the TomTom incidents API, provides a list of reported incidents within a given area. HERE gives the ability to track the following incidents in this API:

- Accident
- Congestion
- DisabledVehicle
- RoadHazard
- Construction
- PlannedEvent
- MassTransit
- OtherNews
- Weather
- Misc
- RoadClosure *
- LaneRestriction

The HERE incident API also allows filtering on category, similar to TomTom. The HERE incidents API logs the possible length of a possible traffic jam, as well as the delay, and severity (0 for critical to 3 for low impact). The API identifies whether the report has been verified or not. The HERE incidents API provides a few extra fields for extra possibilities, such as the political boundary, an edge_ID for the NAVTEQ map and whether or not a response vehicle is on the way. The full content of the HERE incidents API response is shown in Table 6.

HERE collects this data from multiple sources, including connected car probes, roadway sensors and live operations centres.

Annual cost is based on population. As an indication this will be 50 k€ per year for NL and the double for UK per year (without any discount).

Source	Response Content	Remarks

TomTom incidents API	{
	"incidents": [
	"properties",
	"geometry",
	"type"
	],
	"properties": [
	"roadNumbers",
	"events",
	"delay",
	"iconCategory",
	"from",
	"endTime",
	"id",
	"length",
	"aci",
	"startTime",
	"magnitudeOfDelay",
	"to"
	],
	"events": [
	"description",
	"code"
	],
	"geometry": [
	"coordinates",
	"type"
	],
	"aci": [
	"lastReportTime",
	"numberOfReports",
	"probabilityOfOccurrence"
	]
	}
		The TomTom incidents API should also return TMC (Traffic Message Channel) information according to specification but was not visible in our testing of the API.

Here	{ "Traffic_Item": ["Product", "End_Time", "Traffic_Item_Status_Short_Desc", "Abbreviation", "Mid", "Comments", "Traffic_Item_Id", "Location", "Traffic_Item_Detail", "Criticality", "Start_Time", "Verified", "Traffic_Item_Description", "Entry_Time", "Traffic_Item_Type_Desc", "Rds-Tmc_Locations", "Original_Traffic_Item_Id"	Some of the fields in the Here incidents API have a deeper description than has been given here but have been omitted for brevity.

```
    ],
    "Location": [
      "Length",
      "Defined",
      "Navtech",
      "Political_Boundary",
      "Intersection",
      "Geoloc"
    ],
    "Traffic_Item_Detail": [
      "Road_Closed",
      "Alertc",
      "Incident",
      "Event"
    ],
    "Event": [
      "Scheduled_Construction_Event",
      "Event_Item_Cancelled"
    ],
    "Alertc": [
      "Event_Code",
      "Urgency",
      "Update_Class",
      "Description",
      "Quantifiers",
      "Alertc_Q_Binary",
      "Extent",
      "Alertc_Direction",
      "Traffic_Code",
      "Alertc_Duration",
      "Alertc_Q",
      "Supp_Traffic_Code",
      "Duration",
      "Phrase_Code"
    ],
    "Traffic_Item_Description": [
      "Type",
      "Value"
    ],
    "Rds-Tmc": [
      "To",
      "Length",
      "Direction",
      "Alertc",
      "Primary_Offset",
      "Origin"
    ],
    "Incident": [
      "Accident_Incident",
      "Disabled_Vehicle_Incident",
      "Congestion_Incident",
      "Response_Vehicles",
      "Miscellaneous_Incident"
    ],
    "Diagnostic": [
      "Info",
      "Sfile",
      "Pdd"
    ]
}
```


Table 6: Traffic Incident API information

The DfRS OEM SRTI-data feed from BMW is available via HERE under Creative Commons License (CCL). The expectation is that all EU public sector entities related to traffic information will be the users of the data. BMW is offering a subset of data category 1 - Road safety (out of 4 data category types) which is defined in the ACEA (European Automobile Manufacturers Association) position paper on access to the vehicle and vehicle generated data (19.09.2016). The focus is on the social benefits. Anonymized data is exchanged between contributing parties (including public authorities) to enable a significant improvement in traffic safety. The data available under CCL is event data as shown in Table 7. Data is available for the European continent.

Event
ABS or DSC signal lights on/off
Breakdown-Call or Emergency-Call
Rear fog light on/off; Wiper frequency is increasing
Rear fog light on/off; Wiper frequency is increasing; ABS or DSC signal lights on/off;

Table 7: OEM SRTI CC data

4.1.5 Coverage of roads and vehicles

SHADAR report D2.1 showed that stopped vehicle detection deployment differs across European countries. Where traffic is dense on motorways, detection, traffic management and cameras are often available. When an all lane running or dynamic hard shoulder scheme is in place, additional safety measures are in place, including additional detection. In more rural areas there is little or no detection available. Where there is no detection infrastructure, most of today's detection relies on human observation such as a 112 call, police, or traffic officers.

Notifications from connected vehicles have the highest potential coverage. eCall has been covered in Chapter 2. For the coverage of the SRTI-connected vehicles within the DfRS ecosystem, an educated guess based on conversation with HERE is 1% of the fleet and growing. Growth of connected vehicles has been predicted to have a compound annual growth rate of 16.9% (Mordor Intelligence, 2020).

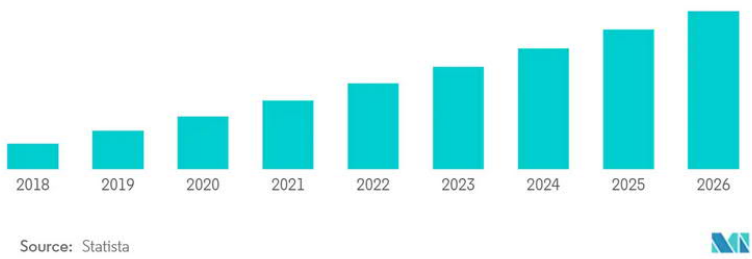

Figure 36: Relative growth of connected cars (Mordor Intelligence), 2021 = 237 million globally.

For the C-ITS services, from 2019 first vehicles are equipped with V2X safety functions. Availability is low and safety, security and data protection issues have to be resolved before enabling of sharing DENM and CPM type of data for SVD. Security and data protection are subjects of current work in the Netherlands by Rijkswaterstaat, RijksDienst Wegverkeer (Netherlands vehicle authority) and others.

Overlap between these three ecosystems is high, as new vehicles will increasingly be equipped to support all three, but the data could differ as it may come from different sources in the vehicle and different data processing before presentation as relevant information for stopped vehicle detection.

4.2 Social media/apps

Twitter

As mentioned in SHADAR report D2.1, the CEDR UNIETD project investigated the use of Twitter data and found that traffic incidents like SVD could be classified and georeferenced to create a practical filtered feed for human interpretation (Grant-Muller, 2015). The research mentions that only 40 out of 8 million selected tweets in the analysed dataset mentioned one of the road numbers in the study area where several were in another country. This highlights that the potential use of tweets for SVD therefore depended on the precise geo-tagging of tweets, which was available for 1% to 2% of tweets at the time of the research. The textual content of Twitter messages has the potential for use for traffic-related information, but there are costs involved with using the global Twitter firehose or other commercial products.

SHADAR performed an updated investigation of current potential of Twitter messages for stopped vehicle detection, in a heuristic way to point out the possible benefits and drawbacks from a data fusion/cross-referencing perspective

To assess the feasibility of Twitter data for SVD, five million tweets were programmatically collected free of charge from approximately one week of Dutch twitter data in October 2020. This resulted in a subset of 130.000 tweets containing a spatial location, which was further reduced by excluding non-human Twitter accounts and TMC Twitter accounts. A dictionary-based text analysis method classified these tweets based on three criteria:

- presence of a Dutch road number (motorway or trunk road)
- a relation to a physical road or part of a physical road
- a relation to a vehicle

Several thousand of tweets scored on one of these criteria and were therefore deemed relevant. Only a single tweet scored on all three criteria and actually described a stopped vehicle situation. The tweet text contains a reference to a road number and driver location sign number, but the geotagged coordinates from the tweet lie more than 13 kilometres from this location. It is highly plausible that the geotagged coordinates originated from the bounding box of the administrative boundary of the region the tweet was sent from. Based on the road number and driver location sign number, the tweet could be manually combined with a corresponding Waze alert, which is shown in Table 8.

The situation that is described in the tweet could not be found in data from TMC's. This is surprising, because the Road operator (Rijkswaterstaat) is expected to dispatch a road officer to secure the potentially dangerous situation on this kind of road. The reason for this situation not being in the TMC data can be that (1) it was not detected upstream or (2) it was detected upstream but not classified as being relevant.

Source	Date and time	Content	Coordinates
Twitter	24-10-20 16:59	"A broken down lorry has been stationary on the hard shoulder for 6 hours on the N280 at driver location sign 14.2 Right Lights on, no sign of the drivers #dangerous"	More than 13 km from the driver location sign

Waze	Start: 2020-10-24 10:05:14 End: 2020-10-24 20:30:47	HAZARD_ON_SHOULDER_CAR_STOPPED	Within 10 m of the driver location sign on the correct carriageway

Table 8: Example of a tweet with a corresponding Waze alert

The example in table 6 illustrates the potential (limited) value of tweets for SVD. While the tweet was posted six hours later than the Waze alert, thereby limiting its use for early detection, the tweet content does add valuable context to the situation. The text may be from a local road user that passed the situation several times. It is conceivable that other cases exist where tweet content provide unique content missed by Waze.

Overall, a small fraction of the collected Dutch tweets proved to be useful for SVD (0.001%, a single relevant and useful tweet per week) compared to the traffic situation messages related to SVD from TMC's (more than 2000 per week on average). This result is in line with the work of Grant-Muller et al. (2015). The trade-off between effort and potential use of tweets for SVD may be different per country or region depending on their existing data sources.

A key factor in this consideration is that the accuracy of the geotag coordinates that are attached to Twitter messages is often inadequate for the purpose of location referencing. This is most likely influenced by the decision of Twitter to remove precise geotagging in 2019 (Guerrero-Ibanez et al 2020). Twitter users can currently add the precise coordinates of their location to a tweet only by taking a photo using the Twitter app or by sharing a geotagged message from a third-party app. The practical implication is that the majority of geotags do not represent the precise physical location of the Twitter user. Instead, the attached coordinates originate from a general perimeter of a place, such as a city or a region. In summary, the aforementioned quantitative and qualitative difficulties severely limit the added value of Twitter data for stopped vehicle detection.

Traffic messages and Waze datasets NL

The national access point in the Netherlands (NDW) publishes five sources of traffic situation messages on opendata.ndw.nu. NDW also provides an ITS access point with multiple sources containing two safety-related sources, SRTI-data for the Highway network and enriched data combined by BeMobile. An overview is given in Table 9.

Description	Sources	DATEX II SituationRecords
Traffic messages gebeurtenisinfo.xml.	TMC's Mostly Highway network Validation by VCNL	*Operatoraction* • SpeedManagement • RoadOrCarriagewayOrLaneManagement • GeneralInstructionOrMessageToRoadUsers • GeneralNetworkManagement • ReroutingManagement • RoadSideAssistance • MaintenanceWorks *TrafficElement* • AbnormalTraffic • Accident • WeatherRelatedRoadConditions • NonWeatherRelatedRoadConditions • PoorEnvironmentConditions • AnimalPresenceObstruction • GeneralObstruction • InfrastructureDamageObstruction • VehicleObstruction • EnvironmentalObstruction • AuthorityOperation

		• PublicEvent • DisturbanceActivity • EquipmentOrSystemFault *NonRoadEventinformation* • TransitInformation • RoadSideServiceDisruption • CarParks
Safety related messages srti.xml.gz.	Combined service for SRTI related data with direct influence on traffic safety (EU-regulation) Coverage Roads Rijkswaterstaat Quality thru protocol (IM-process) Weather from KNMI	• SpeedManagement • RoadOrCarriagewayOrLaneManagement • GeneralInstructionOrMessageToRoadUsers • GeneralNetworkManagement • ReroutingManagement • RoadSideAssistance • AbnormalTraffic • Accident • WeatherRelatedRoadConditions • NonWeatherRelatedRoadConditions • PoorEnvironmentConditions • AnimalPresenceObstruction • GeneralObstruction • InfrastructureDamageObstruction • VehicleObstruction • EnvironmentalObstruction • AuthorityOperation • EquipmentOrSystemFault
Breakdown and accidents incidents.xml.gz.	Data essential to incident management Source insurers and salvage companies Coverage IM-network?	• VehicleObstruction • Accident
NAP-ITS Enriched data Only under license	BeMobile Multi source (NDW, own community, feedback information services, FCD etc.) Semi manual integration Coverage Main road network Ring roads around major cities Provincial road network Urban road network	• Temporarily slippery road surface • Animals, people, obstacles or debris on the road • Unsecured Accident Location • Short-term road works • Decreased visibility • Ghost Rider • Unsecured roadblock • Exceptional weather conditions
Road works and events		Not relevant for SVD
Bridge opening		Not relevant for SVD

Table 9: NDW situation messages

Traffic situation messages from 3 different NDW sources from 2020 were collected, filtered, and aggregated to represent 124,253 unique traffic situations relating to stopped vehicles. The spatial extent of this data is limited to the Dutch incident management network. Only the NDW alerts with a subtype relevant to SVD have been assessed, as shown in Table 10. The full list of the NDW situation record types is available in the documentation.

SituationRecord	SituationRecordType
VehicleObstruction	brokenDownVehicle brokenDownHeavyLorry
Accident	accidentInvolvingHeavyLorries overturnedHeavyLorry accident

Table 10: Selected NDW alert types

Within the tooling of the IM-viewer of MAPtm, a historical dataset of user generated Waze alerts is available. Approximately 5 million Waze alerts on Dutch roads were selected from 2020. A subset of Waze alert types is given in Table 12 based on their relevance to SVD. Only these subtypes were considered in further analyses. The full list of the Waze alert types is available in the documentation.

Report	Type	Category	Comments
Hazard	On road	Stopped vehicle Road work Object Bad road surface Dead animal Defective traffic light	Comment picture
	Hard Shoulder	Defective vehicle Animals Damaged board	Comment picture
	Weather	n/a	n/a
Accident	Small Large Other side	-	Comment picture
breakdown assistance	Fellow Wazers Emergency call breakdown help		

Table 11: Waze alerts

As shown in Table 11, WAZE also provides a breakdown assistance feature consisting of three types of reports (Waze, 2021). All three subtypes relate to a broken-down vehicle and are therefore relevant for SVD. As the name suggests, the “Fellow Wazers” alert subtype notifies other Waze users that help is required. The emergency call and breakdown help subtypes share information with local third parties, such as local emergency services or roadside assistance service providers. These two subtypes are only available in the United States, Canada, and Brazil at time of writing. All breakdown assistance reports are active for a maximum of 30 minutes. A key difference is that this alert type is generated by Waze users about themselves (primary source), as opposed to Waze users reporting on other traffic situations (secondary source).

Type	Subtype
Accident	Accident_Minor Accident_Major
Hazard	Hazard_On_Road Hazard_On_Road_Car_Stopped Hazard_On_Shoulder_Car_Stopped

Table 12: Selected Waze alert types

In order to assess the added value of Waze for SVD, a comparison has been made between NDW situation messages and Waze alerts by matching items from both datasets in space and time. The Waze dataset has been manually limited to the IM network to allow for a meaningful comparison, resulting in 94% of all Waze alerts. A total amount of 115,045 NDW situations (93%) was matched to at least one Waze alert. Vice versa, 31% of the Waze alerts on the IM network was matched to at least one NDW alert. This means that several hundred thousand of Waze alerts can be of added value. On the one hand, matched Waze alerts are useful for their fusion potential or their potential to give faster information or additional content within the IM network. On the other hand, unmatched Waze alerts may lead to the detection of new situations outside the IM network.

Table 13 shows that the majority of situation records is made up of the broken down vehicles situation types, with less than 20% of all aggregated NDW situations describing some sort of accident. This accident figure increases to 25% for the 8053 unmatched situations. The unmatched NDW situations may be a result of the pre-selected Waze alert types (see Table 11), but a manual search for such examples gives no indication for this pattern. The limitations of the matching method that has been applied can also be a potential source of false negatives in terms of matches between NDW situations and Waze alerts. A sensitivity analysis of the spatio-temporal variables in the matching implementation can create a better understanding of the interaction with the matched percentage.

Similar statistics are given in Table 14 for Waze alerts. In total, the bulk of the Waze alerts are of the HAZARD_ON_ROAD_CAR_STOPPED and HAZARD_ON_SHOULDER_CAR_STOPPED subtypes. 6% of the alerts describe some kind of accident. Outside of the IM network, the most common Waze alert type is Accident (36%), followed by generic hazard subtypes (see Table 15). Only 6% of alerts were classified as one of the stopped cars subtypes, which may be due to the absence of a hard shoulder on the type of roads outside of the IM network.

	NDW situations	
Alert type	Total	Not matched to Waze alert
Accident (general)	18%	24%
Accident involving a lorry	1%	1%
Broken down vehicle (general)	76%	70%
Broken down lorry	5%	4%
Total	100%	100%

Table 13: Comparison of the relative amount of NDW situation types between all situations and those not matched to a Waze alert

	Waze alerts on IM network	
Alert type	**Total**	**Not matched to NDW situation**
Accident	5%	2%
Hazard (general)	2%	2%
Car stopped on road	36%	38%
Car stopped on shoulder	58%	58%
Total	100%	100%

Table 14: Comparison of the relative amount of Waze alert types between all alerts on the IM network and those not matched to a NDW situation

	Waze alerts	
Alert type	**Total**	**Outside IM network**
Accident	6%	36%
Hazard (general)	1%	30%
Car stopped on road	1%	29%
Car stopped on shoulder	36%	1%

Table 15: Comparison of the relative amount of Waze alert types between all alerts and those outside the IM network

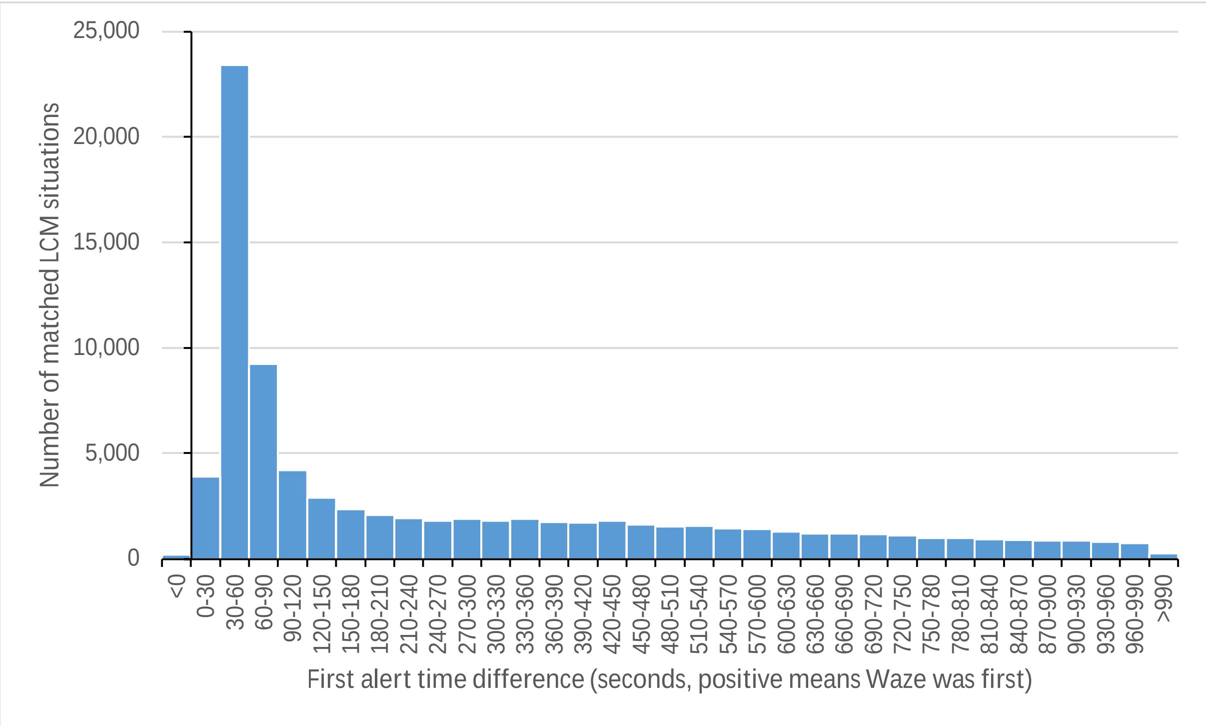

Figure 37: Distribution of the first alert time difference between NDW and Waze

Figure 37 depicts the distribution of the first alert time difference between NDW and Waze for matched NDW situations per 30 seconds. NDW published a situation message earlier than Waze for less than 200 matched NDW situations. On average, Waze alerts related to stopped vehicles are published 1.5 minutes before NDW publishes theirs. In addition, specific IM tooling using Waze data was in operation at Dutch TMC's until the end of 2020. This combination is a strong indication that TMC's use Waze alerts in their operations and that both data sources are not independent.

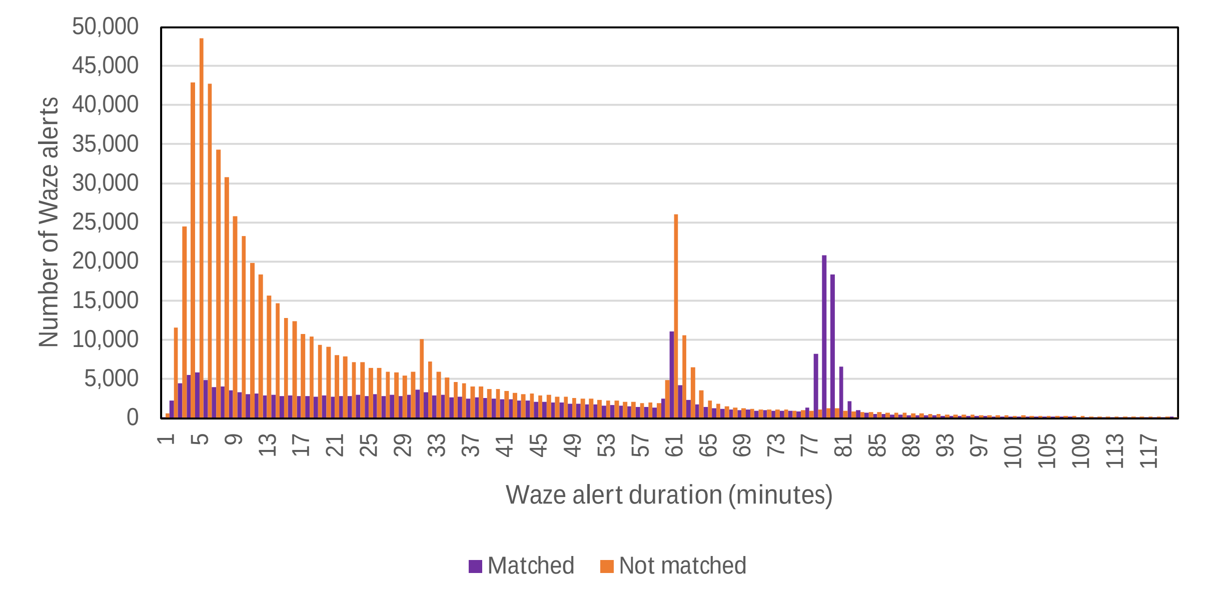

Figure 38: Distribution of the Waze alert duration

The Waze alerts on the IM network that were not matched to an NDW alert were predominantly incidents with a relatively short duration compared to Waze alerts that were matched to an NDW alert (see Figure 38). The majority of the unmatched alerts had a duration of 15 minutes or lower, whereas the timespan of matched Waze alerts regularly ranges up to 80 minutes. The average duration of all unmatched Waze alerts on the IM network amounts to 50 minutes, as opposed to the 65-minute average duration of matched Waze alerts. This pattern may be caused by short-term and low impact incidents that are reported in Waze, but are not caught by the detection methods that lead to an NDW situation message.

Figure 38 also shows surprising peaks at 30, 60, and 80 minutes. These regularities hint at a predetermined and automatic cut-off period for Waze alerts that are not ended manually.

Each alert in the Waze dataset also includes a confidence score, which is a a discrete number between 0 and 10 and serves as an accuracy measure based on user feedback. A higher score indicates more positive feedback from Waze users and is an indication that the alert corresponds with the real traffic situation on the road. The potential use of this confidence score is explored in Figure 39 and Figure 40. Figure 39 shows the relative amount of Waze alerts per confidence score. The general pattern is that Waze alerts that were matched to an NDW situation received more positive feedback from users compared to all Waze alerts. The average confidence score for all Waze alerts is 0.69, whereas the average score for matched Waze alerts is 1.00. Figure 40 combines the average duration of Waze alerts with the assigned confidence score. This does not show a clear pattern or correlation. The historical Waze data as presented in this section also includes a reliability score based on the experience level of the reporter, which is a potentially useful addition to data fusion in a later stage.

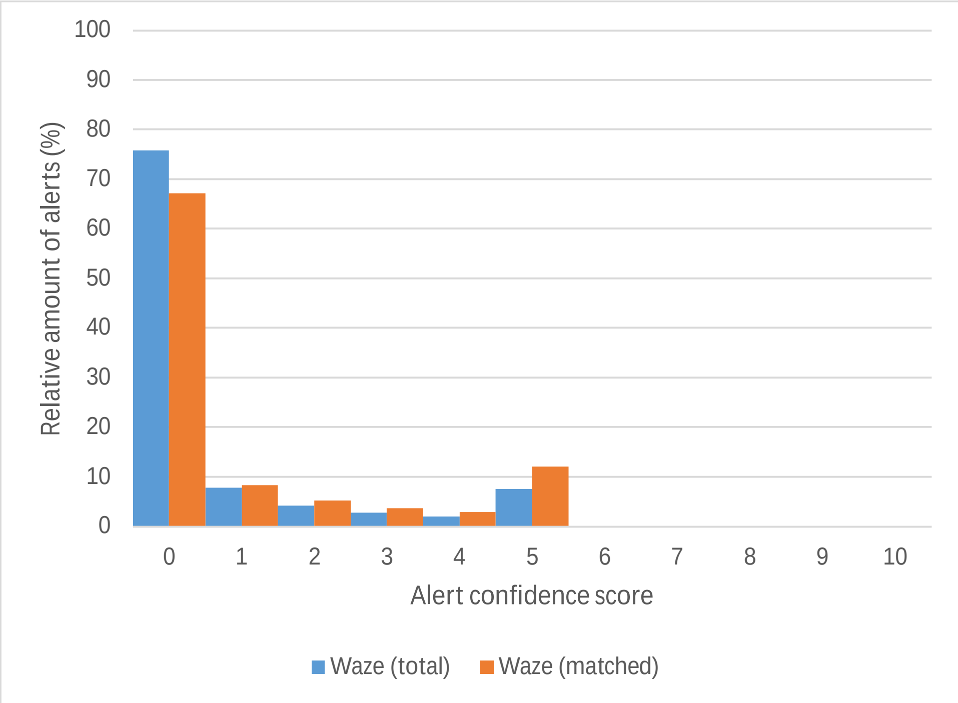

Figure 39: Relative amount of Waze alerts per confidence score

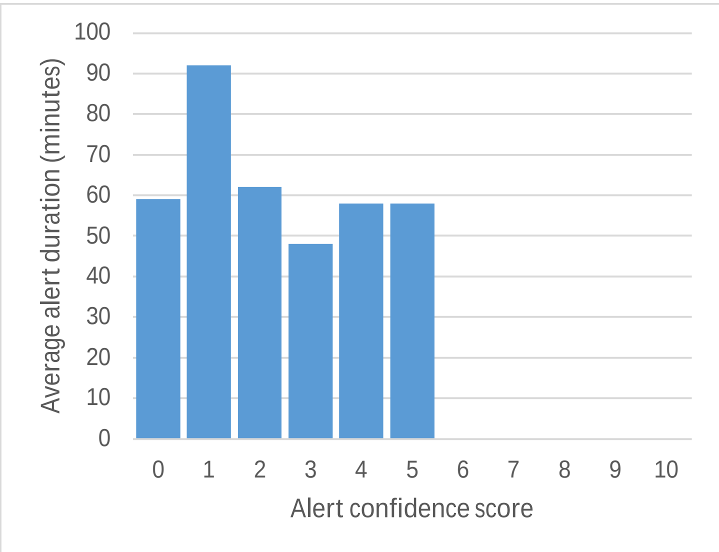

Figure 40: Average duration of Waze alerts per confidence score

Various analyses of the Waze dataset are presented in this section, including alert duration, type, spatial distribution, and alert confidence score. These demonstrate that user-generated Waze alerts are a valuable data source for SVD. This is because the dataset covers the vast majority of the main road network (and more), the alerts are published relatively quickly, and the alerts have high spatial accuracy. The most notable downside that comes with the Waze dataset is data noise. A key example is vehicles making a brief stop on the hard shoulder, which are typically not registered in the NDW dataset. However, this negative effect can be mitigated by filtering the Waze dataset by alert type, location, duration, and a confidence score to match the desired use-case. Another potential avenue of research is combining Waze alerts with data from other service providers, such as TomTom, Be-Mobile, or HERE. However, this approach requires issues such as cross pollination among data sources to be addressed. Ideally, ground truth data is established to verify the quality of service provider alerts. NDW data comes closest to this level of confidence in a practical application in the Netherlands

4.3 Aerial imagery (drones, satellites, other aircraft)

One of the areas of opportunity in the near future for (stopped) vehicle detection, is aerial imagery. This area of SVD comes with a few advantages and disadvantages that differ from roadway sensor technologies. A few of these are listed:

Advantages:

- Large area coverage
- No installation needed
- Detection quality is less dependent on hardware
- High spatial resolution

Disadvantages:

- Moving imagery
- Reliant on image processing algorithms
- Weather dependent (impact varies)
- Costly methods

Between different sources of aerial imagery, there are a few advantages and disadvantages as well, such as satellites being much easier to use on large scale and during the full day period, while drones can provide much more accurate imagery on a small scale and focus their detection purposes on key areas. On the other hand, satellite imagery may only be available on a daily basis or real time but very costly.

When researching these methods of SVD, it becomes clear that most of the research describes different image processing algorithms to detect and track vehicles, rather than describing different hardware components. Satellites are not easily upgraded to higher resolution and framerate hardware, with a reliance on open access satellite imagery, and drones already have access to high-quality imagery. One major advantage can be seen in this, which is that the possibilities of these methods of vehicle detection can be greatly explored by exploring image processing methods alone. A recent article (Outay, Mengash, & Adnan, 2020) discusses specifically the recent advances and challenges around applications UAVs and mentions that "*The main challenges lie in the process of information extraction from the videos as well as deploying a system that is fool-proof so that drones can carry out their function*".

Satellites

Recent research (Stuparu, Ciobanu, & Dobre, 2020) shows that using satellite images to detect vehicles can obtain a "very good detection accuracy and a very low detection time". A model using the RetinaNet architecture reached a detection time as low as 300ms. Other research shows similar worthwhile results such as a 94.7% accuracy using a background subtraction method for moving vehicles (Ahmadi & Ghorbanian, 2019), or promising results using R-CNN (Regional Convolutional Neural Networks) on the CPEC route (Ibrar et al., 2018). Even an old study from 2002 shows vehicles can be counted on 1-m resolution satellite imagery (Sharma, 2002). Another research shows how high-resolution video from satellites can be processed to provide information about traffic congestion (Khalil et al 2017). As mentioned earlier, weather conditions can affect or completely block the ability to capture usable imagery, as mentioned by most of these researchers.

While image processing for satellite imagery is advancing rapidly and reaching the capabilities necessary for SVD, the availability of real-time imagery is equally important. In this area, advances are slowly being made in this area, as the commercial space sector is growing. One of these projects is SOAR, giving access to imagery from the SkyMap50 satellite, at a rate of US$12/km² for newly tasked (on-demand) imagery, with a resolution of 50cm, although with a delivery time of up to 48 hours. In a higher price category (estimated up to US$100/km²), we can find MAXAR. MAXAR provides a direct access program that allows for short term task scheduling, and with satellites that visit multiple times per day, allow for only a few hours waiting time. They are expected to launch Worldview Legion in 2022, which will increase their rate to up to 15 visits per day. Similar high-frequency imagery can be delivered by Planet Skysat, which falls into a similar price category as Maxar.

These results tell us that currently, satellites are not yet suitable for real-time SVD. In the current best-case scenario, 1 visit each 1.5 hours, with a delay of 15-30 minutes to receive imagery could be possible. While these images could be used to detect stopped vehicles, it would not be fast and frequent enough to practically apply to the use case of SVD and possible data fusion.

However, with the recent commercialization of satellite imagery, and more satellites being launched into orbit, it is possible that satellites become a viable source of SVD within the next decade, especially when multiple sources are combined. However, suitable satellite imagery (1-m or better), is still very costly, and might not become affordable for high-frequency usage on a large scale.

Unfortunately, weather, especially clouds also significantly hamper the availability of satellite imagery, and while there are processing methods to remove clouds, these processes still require multiple images from moments where they can bypass the clouds and are unable to be applied for SVD. This problem is expected to continue until other satellite sensor styles also become commercially viable in high frequency.

UAVs

Research into detection methods suited for UAVs (Unmanned Arial Vehicle) show similar promising results. One research (Xu et al, 2017) reaches up to 98.43% correctness and 96.40% completeness rates, through the method of Faster R-CNN to detect moving cars, showing that the method used in the CPEC research provides an even stronger result when used for drones. Older research from 2010 uses a joint probabilistic relation graph approach to detect and track large numbers of moving vehicles at once. It reaches an accuracy of 85% during tracking of multiple objects (moving vehicles) at once, for low frame rate and spatial resolution aerial videos. (Xiao et al, 2010). Other research attempts to combine methods to create a hybrid method (Xu et al, 2016) or explores an enhanced detection method resistant to the movement of the UAV, resulting in a detection rate of 90% and false alarm rate of 1% of moving vehicles (Najiya & Archana, 2018), but also mentions DNN (hybrid deep convolutional neural networks) as a possible direction for future research, which is also discussed in earlier research (Chen et al, 2014). Other recent research (Barmpounakis, Sauvin, & Geroliminis, 2020) manages to reach high accuracy on lane detection and lane-changing identification and goes as far as to implement a prediction algorithm to predict lane-change manoeuvres.

In contrast to satellite imagery, UAV imagery can already be accessed, and drones can be deployed on demand. To make a brief distinction, we mention both drones and UAVs in this section. While UAVs are also drones, drones are not always UAVs. In the common interpretation used here, UAV refers to high altitude unmanned drones that often fly without being directly controlled, while “drones” typically refers to low altitude drones that are controlled by hand. UAVs still are very costly however, and drones are even more affected by weather and suffer from low flight times. Drones do have the option to be deployed in swarms and more selectively monitor high risk or high traffic areas, and are in a more affordable price category, with a price range of 2.000 € to 30.000 € for professional drones.

Object detection methods

As mentioned in (Outay, Mengash, & Adnan, 2020): *"as the drone technology becomes improved and flights are more stable, now focus is more on advancing software capabilities to improve the methods used in image processing and scene reconstruction"*, it is necessary to understand the advances, advantages and disadvantages of different image processing methods and object detection (and tracking) methods. As the article also mentions, however, that *"Almost all of the studies/efforts are applied on a limited scale for testing and validation"*, it is not yet possible to say which methods would be best suited towards the application of SVD. Additionally, the focus of the research lies mostly in finding and testing methods most suited for moving vehicle detection, however, capabilities to track and detect stopped vehicles do seem possible.

Since aerial imagery brings an additional challenge to the processing of sensor data, compared to roadside sensors, namely the moving sensor, whereas roadside sensors are stationary, the object detection methods for aerial imagery are not easily compatible with the methods for roadside technology, although advances have been made to close the gap.

The image processing methods that have been used in vehicle detection based on aerial images, seem to fall into either the algorithmic category (PCA, BBT, gradient-based, Viola-Jones, HOG + SVM) or the neural network category. (CNN, R-CNN, Faster R-CNN, RetinaNet). Some of these methods, such as the PCA and BBT methods rely on a clear background image and detect the objects based on differences. In the table is a short description of each method.

Method	Description	Remarks
Algorithmic methods		
PCA	The principal component method attempts to realign a set of information bands, such that the separation between the vehicles and pavements is maximized in at least one of the resulting principal component bands, where the information bands are different image properties to exploit.	Accuracy dependant on variability in background gray-scale values.
BBT	The Bayesian Background Transformation (BBT) method is used to classify pixels in an image as either stationary or dynamic. Stationary pixels are the ones that show little or no change in gray value when compared to a historic background estimate, while the dynamic pixels are the pixels that show a significant change	Performance is resistant to noise.
Gradient-based	The gradient-based method works primarily on the gray level differences along the vehicle boundary.	Needs high contrast images
V-J	The V-J scheme is based on multiple cascaded Haar-like classifiers. The basic concept is to use a conjunctive set of weak classifiers to form a strong classifier. In practice, Haar-like features are computed as the sum of differences of the pixel intensities between different rectangular regions at a specific location in a detection window.	
HOG + SVM	The histogram of oriented gradients (HOG) is a feature descriptor used in computer vision and image processing for the purpose of object detection. The extraction of these HOG features occurs in 5 subsequent steps.	Invariant to geometric and photometric transformations
Neural Network methods		
(Fast(er)) R-CNN	Regional CNN is a method that uses selective search to extract regions of interest from an image, and applies CNN to each of these regions to detect and classify objects. R-CNN is a slow method, which Fast and Faster R-CNN improve upon by changing how the regions of interest are detected.	Faster R-CNN can achieve near real-time detection.

RetinaNet	RetinaNet is a one-stage object detection model, using 2 other such models: FPN (Feature Pyramid Networks) and Focal Loss. RetinaNet does not inherently rely on Neural Networks, but CNN is used in the FPN layer of the RetinaNet.	Proven to work well with dense and small-scale objects. Incorporation in ArcGis
HDNN	A deep neural network is a neural network that has multiple layers between the input and output layer. The hybrid DNN divides the maps of the last convolutional layer and the maxpooling layer of DNN into multiple blocks of variable receptive field sizes or max-pooling field sizes, to enable the HDNN to extract variable-scale features.	Expected to perform better in difficult and complex tasks.

Table 16: Image processing methods for vehicle detection – definitions are from the respective papers.

Conclusion

In the end, it should be noted that each form of aerial imagery has downsides. Satellites still have low availability and high costs for the amount of coverage that can be obtained, although they do show promises for future improvements, with commercial companies launching more suitable satellites, increasing coverage to become closer to real-time. High flying UAVs are extremely costly (possibly more costly than satellites) for on demand deployment, or are unlikely to reach real-time coverage for open access. Drones, on the other hand, have low deployment times and ranges.

All of the above sources additionally suffer from weather conditions, reducing the range of applications. Alternate sensor technologies could potentially reduce the limitations by weather conditions, although literature on this topic is not easily found. Moreover, due to the already limited availability of satellite and high-flying UAV imagery, this area of research is not likely to be available in the near future.

Drones are at least relatively low cost and an option for flexible and quick deployment, including possible usage of various sensors, to mitigate weather issues. While they might not be a suitable option for large-scale SVD detection, this could make them feasible for targeted situations and locations.

4.4 iWKS development Netherlands

In the Netherlands, MTM (Motorway Traffic Management) is used as a safety measurement on the Dutch highways, stretching over almost 3000 km. The basis of the system was defined in the 1970's. Rijkswaterstaat is working on replacing the installed base over the coming years with the next generation roadside unit called iWKS (Rijkswaterstaat, 2021). 5700 new systems will be installed. First of all, iWKS will maintain the functions of MTM but will also enable extension with new applications or data sources. The maintained functions are:

- AID (Automated Incident Detection), signaling (above the highway, upstream) will react on slow driving traffic showing an automated speed limit. When traffic restores to normal the measurement will be automatically cancelled.
- Safeguard traffic incidents by crossing of a lane / multiple lanes, combined with a speed limit and attention signage.
- Support road works and other hindrance-related situations by putting measurements on the signage (speed limits, pre-warning, expulsion arrows etc.)

The iWKS system is a newly developed platform that contains traffic management software used for tasks at the roadside. The life cycle of the traffic management software can be short with many small iterations and remotely distributed. Automated processes are used (CI/CD, continuous integration, continuous delivery) to quickly deploy or change functionalities in the iWKS. The iWKS architecture is based on IoT (Internet of Things). Practical implementation of processors, sensors, actuators and communication hardware to collect and act on data that is acquired from the iWKS environment. The IoT iWKS architecture is comparable with the environment of a smartphone. A smartphone has sensors, actuators and numerous up-to-date-apps that can use these sensors and actuators running on a software platform that can be installed on different types of hardware.

Traffic management applications are stored in a container and are automatically distributed on the designated roadside units. As software deployment is seamless new methods of detection or applications can be added.

As an example, a proof of concept was done with floating car data on a stretch of road on the A9 to replace loop detection for the AID functionality. Other use cases will follow in the future.

For SVD the usages of iWKS can be found in the area of:

- Extension in the detection methods to detect stopped vehicles (e.g. rotating radar) with location and lane-specific information
- Local hazard warning communication
- ITS G5 C-ITS two-way message services.

4.5 Conclusions

The potential benefits of topics explored in this chapter can be characterised as one or more of the following:

- More timely information
- More reliable information
- More accurate information
- Additional relevant incident information

Timely information

Data obtained from eCall and SRTI-Datafeed should shorten the detection time compared to manual detection methods. At the moment the coverage of vehicles is low but will grow.

C-ITS solutions exist but coverage is really low. If C-ITS fulfils its promise then it should be a major future source for detecting stopped vehicles. Not only receiving SVD but contributing to the cooperative awareness with roadside equipment will contribute to the value of the C-ITS eco-system.

For the short-term at least, data coming from service providers, using detection methods through analysis of data and/or reporting by users, is most valuable. Analysis of the Waze dataset found that reports of the incident are quicker than the national registration in the Netherlands and cover a much larger road network.

More reliable information

The use of multiple data streams to detect and verify SVD helps to obtain more reliable detection of real incidents. Waze provides indications of trustworthiness by classifying the reliability of the reporting Wazer and the numbers of thumbs given by fellow Wazers for a reported incident.

However, research on Twitter information shows that retrieving geospatial information in combination with traffic-related tweets gives a low success rate, not justifying further research or investment.

More accurate information

The use of UAVs and satellites can provide a source that could provide more accurate information on location, lane, and vehicle type. At the moment only drones could be feasible for SVD but have a small range of sight. UAVs and satellites could be, in the future, of help as a reactive measure to create a view of the incident site and identify the type of SVD on the location in order to estimate what resources should be deployed.

Through the set of coordinates in the SRTI-message it is possible to determine the lane a vehicle is in. From the evaluation, it is not clear if this data is accurate enough.

It is possible to obtain more accurate data from C-ITS message sets, like speed, direction, collision risk, emergency electronic brake warnings.

More relevant incident information

From C-ITS messages (and also eCall) additional attributes may be present or derived, such as the number of occupants, dangerous goods, type of human problem, type of propulsion. These could help determine what response measures are needed at the stopped vehicle site.

The SRTI eco-system could also be extended with this kind of information but is outside the dataset prescription of the EU legislation regarding DfRS.

As mentioned in the previous paragraph UAV/satellite/drones could be an additional source in the future as a reactive measure to obtain additional information regarding a potential SVD.

Originality and feasibility

Table 17 and Figure 41 present the research ideas in a Now-Wow-How matrix the position in the matrix gives an educated guess of feasibility and originality.

	Subject	Category	Description
NOW	Waze / commercial traffic information	Timely	Extension of accuracy by alerting and fusion in TMC
NOW		Reliable	Categorisation and trustworthiness Waze feed as extra source
NOW	Harvest Data for Road Safety	Timely	High road coverage, low vehicle coverage, but growing.
NOW	C-ITS safety-related	Timely, reliable, information	Potential high quality data direct from vehicle sensors. However, coverage still at R&D levels.
WOW	**Subject**	**Category**	**Description**
WOW	Platform for easy function extension roadside (iWKS)	General	Seamless function enabler on roadside and central applications using sensors and actuators. Despite appearing in the WOW category, by itself this idea does not deliver value for SVD, it must be combined with other advances.
HOW	**Subject**	**Category**	**Description**
HOW	C-ITS extension	Timely	Extend the use of CPM/DENM/ to warn based on detection SV and other CVs
HOW	UAV/satellites	Accuracy	High accurate info of SVD and location but a future use. Drones can be deployed to view the incident site

Table 17: Now-Wow-How categorization upcoming methods.

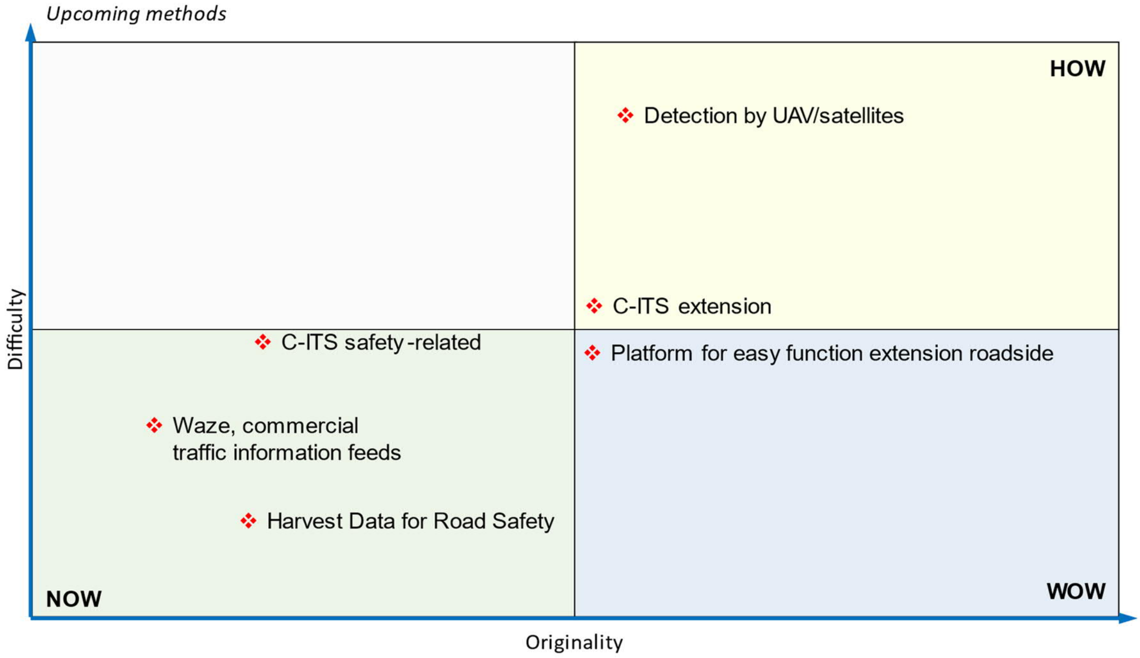

Figure 41; Now-Wow-How matrix upcoming methods

The variety of sources explored in this chapter suggests potential of data fusion for cross-checking and increasing confidence, a topic explored in the next chapter.

5 Data fusion

5.1 Purpose of data fusion for stopped vehicle detection

This chapter explores the research hypothesis that data fusion can improve stopped vehicle detection, in at least one of the following: detection rate, false alarm rate, coverage of locations, or richness of alert data content.

SHADAR D2.1 presented a simplified comparison of various kinds of stopped vehicle detection sources using a variety of metrics. Figure 42 plots a subset of that data on a single diagram, where low performance has been mapped to radius 1, medium performance to radius 2, and high performance to radius 3, so the outer region represents a better performance. It is notable that every source outperforms another on some metric, and every source is exceeded by another source on some metric.

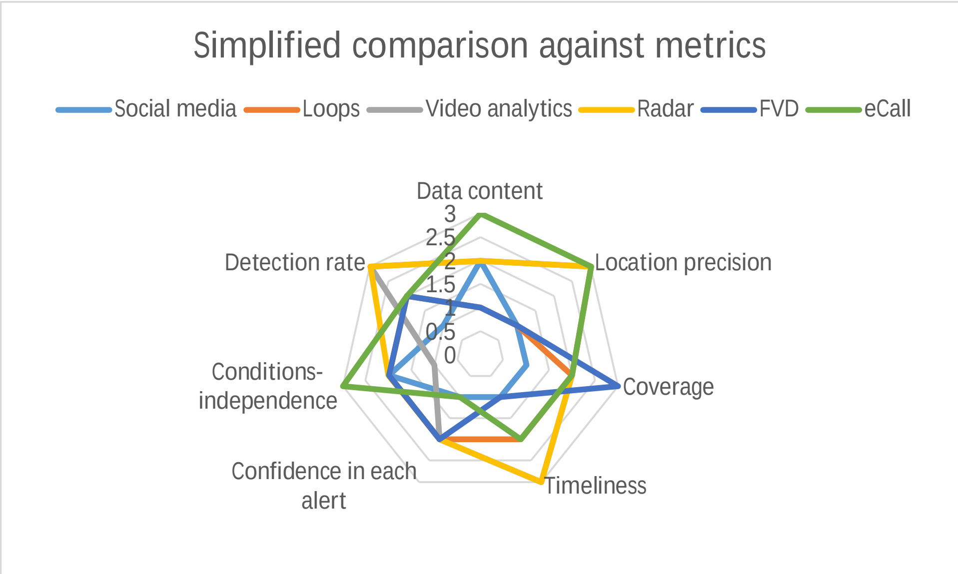

Figure 42 Simplified comparison of SVD technologies

This suggests that fusion of data from multiple sources could achieve a better overall performance than from any of the individual sources, as illustrated by adding a fusion polygon to the chart in Figure 43.

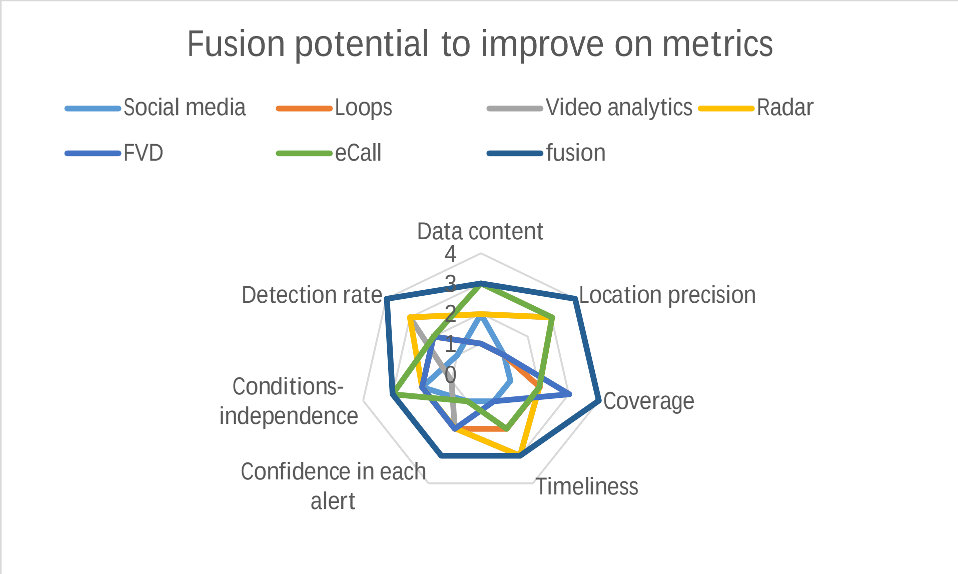

Figure 43 Potental of fusion to improve on individual SVD technologies

An intuitive example is coverage of locations. Methods using fixed infrastructure as sensors can potentially detect all events in their locus of coverage, which is only on road sections where the infrastructure is installed, while methods using connected vehicles can potentially detect on all roads, but only where there are suitably equipped vehicles. By combining the available data from both types of sources, a greater coverage of stopped vehicle events can be achieved.

The coverage of connected vehicle sources will change over time: overall as a set they will grow, but individual types or brands of connected vehicle sources may fade as particular technologies or businesses are overtaken by others. Fusion of multiple different connected sources could mitigate this variability.

On a metric such as data content, a fusion of data from different sources evidently has the potential to preserve the best from each source.

For metrics such as detection rate and false alarm rate, this chapter will explore the hypothesis that fusion can improve these even beyond the best rates achieved by individual sources.

5.2 Data fusion methods

SHADAR report D2.1 pointed out recent fusion works El Faouzi and Klein (2016), Klein (2019), and most recently Cvetek et al (2021).

In a CEDR review Cornwell et al (2016) noted an increasing focus on traffic theory to fuse individual vehicle and detector measurements, with rule-based techniques becoming less frequent. In 2021, Cvetek et al noted an increasing focus on fusion through deep learning, which they expected due to the recent popularity of deep learning across many domains. The CEDR review also noted that fusion for incident detection used two contrasting kinds of fusion: (i) fusing the raw data to improve a single detection process (ii) allowing multiple detection algorithms to reach a conclusion and then fusing their outputs. The subsequent fusion work identified in the Cvetek review appeared to focus on traffic state estimation and did not directly inform the state-of-the-art for incident detection, where the most effective techniques are not yet known.

The choice of fusion method depends on the data sources available. Studies have typically compared techniques given known data sources, or compared methods independently from specific data sources. We have not seen research that supports both the choice of fusion methods and the individual data sources in an integrated study, but NRAs to some degree have that choice – there is a large palette of potential data sources that could be deployed or procured or harnessed, NRAs will not deploy all of them, so in a data fusion regime it is the sources that have the biggest positive effect on the overall fused detection that should be pursued first (although in some cases the data sources are not dedicated to stopped vehicle detection and so there may be independent benefits for their deployment or procurement).

Refining the observation of the two contrasting fusion methods:

(i) **Fusing raw sensor data into a single detection process.** Stopped vehicle detection can be considered a classification problem for which machine learning can be applied. In other domains the success experienced using machine learning suggests it to be the most obvious method. Features from every raw source could be input to a supervised machine learning method. Machine learning projects typically acquire large amounts of manually classified training and testing data, and then test the relative success of different kinds of classifiers. The SHADAR project did not have access to substantial sources of raw data, but this could be tackled by future research. However, each detection technology supplier is already optimizing their own detection, very often using machine learning, so bypassing this valuable detection learning by accessing raw data would mean some duplicated or wasted effort, albeit a multi-source classifier would learn in different ways. The established learning in individual sources is also a reason that some suppliers may be reluctant to provide raw data which they may see as a lesser service. **So, although experience suggests that this method promises technical success, it may not be the most practical solution.**

(ii) **Fusing outputs of multiple detection sources.** This has the least architectural impact and from today's starting points should be cheaper than successfully fusing raw data.

Approaches for this kind fusion include manipulation of probability, a heuristic scheme based on confidence (Dempster-Schafer is the best researched, or bespoke heuristic schemes can be invented), or fuzzy logic. The mathematical foundations for all approaches to manage uncertainty have been questioned, but the approach with the most widely accepted (least challenged) foundation is probability theory including Bayes theory. Section 5.4 presents a further study of how this could be employed in stopped vehicle detection fusion.

5.3 Considerations for stopped vehicle detection by machine learning

Practical reasons identified above led to SHADAR not further researching machine learning, but this section briefly notes thoughts for potential future research.

Machine learning requires significant training data where ground truth is known along with outputs from all the detection sources to be fused. The ground truth data set should include many examples of true positives. The most obvious way to acquire the ground truth data is to use video cameras and have humans record the times when genuine stopped vehicles occur. For best accuracy, the video monitoring would be throughout a study period. A shortcut would be to perform human verification only at times that at least one source raises an alert, but if the data fusion classifier is taking features from a data source at regular intervals, not just when that source raises an alert, then the ground truth data set would lack the class of real stopped vehicle events not detected by any individual source.

The kinds of feature extraction used for each source may vary significantly. For a source like eCall, and most of the other connected vehicle sources, there is only the alert and its associated data – at other times there is no data. For fixed infrastructure sensors there is the possibility to extract data periodically. While the entire data set (such as every image in a video feed) could be used as an input to a classifier, machine learning typically extracts features from the original data. The suppliers and experts for each technology would be best placed to advise on the kinds of features which are most significant for stopped vehicle detection.

The radar improvement studies in Chapter 4 suggest there may be potential in including features beyond the most obvious ones, for example to detect stops from vehicle behaviour immediately preceding a stop. Also using pre-detection parameters, research might also explore whether there is any possibility to distinguish the characteristics of a likely breakdown from a likely temporary stop. If that could be done with sufficient chance of accuracy, an alert from a suspected temporary stop might be withheld from operators (or raised with lower priority in the user interface) for a period to avoid overloading operators, then if the stop persisted it could be alerted with full priority.

If the practical problems of acquiring raw data could not be overcome, machine learning could also be applied to the outputs of individual classifiers, but from our intuition based on experience of classification problems this seems unlikely to significantly outperform a probabilistic approach (explored in subsequent sections) to be worth the significant investment.

5.4 Study using probability to characterize fused detection

This section analyzes how probability can characterize a data fusion system and therefore support decisions about what sources to use and how to use them in a data fusion system.

Much of the following analysis is about **a priori characterization of data fusion systems**, rather than about calculations that happen dynamically at runtime to determine dynamically how to treat an alert from a data source. Its purpose is to allow understanding of what is being achieved when fusing multiple sources in different ways, to show what will be the a priori characteristics of the fused system and how that compares to the original a priori characteristics of the individual sources.

Using detection performance on past examples to predict detection performance on future examples assumes that the past examples are representative of future examples (i.e. or in other words the past examples allow estimation of a model that can be used for out-of-sample forecasting). This is not a trivial assumption, but even if the past examples are not perfectly representative, and the forecasts not perfectly accurate, the technique can be useful.

This analysis uses probability theory. No known current practical detection source is 100% accurate 100% of the time - every alert from every detection source can be considered to represent a probability of a real stopped vehicle event.

Expecting Bayesian probability to be a useful solution for data fusion problems, we undertook a literature search for this. While we did find published work on Bayesian probability for data fusion, we found it applied to slightly different data fusion problems to ours, for example how to calculate detection probabilities given sources which change their mind about a classification over time. The following is our own application of probability to the problem of stopped vehicle detection data fusion. It uses conditional probabilities but does not require Bayes' theorem. Although it is simpler than most research on Bayesian data fusion, and uses nothing more than elementary probability theory, its application to the analysis of the properties of event detection fusion systems may be novel.

We reuse definitions from D2.1:

- "true positive" – a stopped vehicle event is detected, producing an alert
- "false positive" – a stopped vehicle alert is reported but there is no stopped vehicle
- "false negative" – a real stopped vehicle event is not detected i.e. there is no alert

A "true negative" would be when there is no stopped vehicle and no alert.

$$Detection\ Rate\ (DR) = \frac{True\ Positives}{True\ Positives + False\ Negatives}$$

$$False\ Alarm\ Rate\ (FAR) = \frac{False\ Positives}{True\ Positives + False\ Positives}$$

The Time to Detect (TTD) is the total interval between the event occurring and the reporting of an alert.

Given DR, FAR and TTD of individual sources, it is possible to calculate the DR, FAR and TTD of fusion schemes based on these sources (given assumptions), and therefore understand the value of fusing the sources. With further data from sources detecting for the same times and places, these assumptions can be refined and the calculated fused DR and FAR will become more accurate.

A naïve assumption used at some points of the analysis is that the chance of detection of a stopped vehicle by one source is independent from the chance of its detection by another source. This may be reasonable if fusing a source like radar with a fundamentally different source like eCall, but if fusing say video and radar sources, there could be correlation between the chances, for example both may be highly likely to detect the same event in good conditions but both might be less likely to detect the same event in extreme atmospheric conditions (albeit that radar might be less affected by such conditions than video). This assumption would therefore be valid for some combinations of detection sources but not all combinations.

This naïve assumption can be removed by studying data from the sources so that *conditional* probabilities are known e.g. the DR/FAR of one source in cases when another source raises an alert. If certain defined contextual conditions seem to have a significant impact on DR/FAR, those conditional probabilities can also be calculated and used e.g. the DR/FAR of a source when optical visibility is known to be poor due to heavy precipitation or fog.

Simple analysis of an event-detecting data fusion system can consider a crisp binary output – either an alert is reported to an operator, or it is not. In practice there is no need to restrict a data fusion system to have a binary output. Instead the concept of priority can be used if it is supported in operational user interfaces for traffic managers. We assume that traffic management operators have limited resource and may not be able to react instantly to every indication that there may be an incident, especially if we allow lower probability indications of an incident. Operational user interface features for prioritisation can allow high probability alerts to be raised with prominence, perhaps with audible alarms, while making lower probability indicators accessible should the operator have enough time to explore them at times of quieter workload.

D2.1 observed that there can be a trade-off between detection rate, false alarm rate and time-to-detect, and that detection sources can be calibrated to achieve a desired balance. The same is true for fusion schemes – the rules of fusion can be chosen to optimise detection, or false-alarm rate, or to balance between these.

The following sections contain examples using three illustrative stopped vehicle alert sources identified as sources i, ii, and iii.

Source	DR	FAR	TTD
I	85%	30%	10s
Ii	80%	20%	20s
Iii	12%	20%	30s

The following analysis assumes that it is possible to identify two alerts from different sources as representing the same stopped vehicle event, even if they initially occur in slightly different locations at slightly different times. Most sources report not only when an event starts but when it clears, so no matter the precise time that they raise an alert, there may be periods when both sources are asserting the existence of the same event. Time is considered further in section 5.4.6. Alignment in location is possible by adopting a coarse resolution for the fusion system – see 5.5 for further discussion. The analysis defers consideration of the influence of weather and traffic state until section 5.7.

5.4.1 Fusion regime A – alert if any source alerts

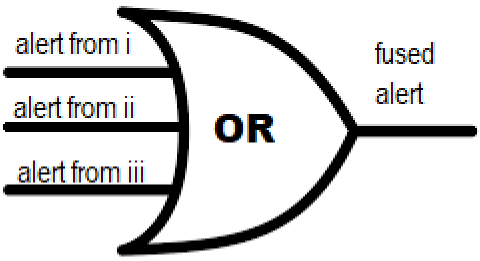

Figure 44 OR fusion regime

Detection rate

Detection rate = 1 - p(all sources fail to detect)

With the naïve assumption that each probability is independent (which we will explore removing later in this analysis):

P(all sources fail to detect) = p(source 1 fails to detect) * p (source 2 fails to detect) …

P(source fails to detect) = 1-DR for that source

So P(all sources fail to detect) = (1-DRi) * (1-DRii) …

With our example:

P(all sources fail to detect) = (1-0.85) * (1-0.8) * (1-0.12) = 0.0264

So DR(fused) = 97.36%

We have increased from the max 85% rate of one source to over 97%.

We can also calculate the incremental effect of adding sources. For example, just by fusing sources one and two we reach 15% * 20% = 3% missed detections, DR 97%, so the incremental effect of also fusing source iii is only adding a further 0.36% (still with the assumption of totally independent detection probabilities).

An alternative expression of the formula for 2 sources is DR(fused) = DRi + DRii – (DRi * DRii).

An alternative expression with 3 sources is DR(fused) = DRi + DRii + DRiii - DRi * DRii - DRi * DRiii - DRii * DRiii + DRi * DRii * DRiii.

In our example, using only the first 2 sources, DR(fused) = 0.85+0.8-0.85*0.8 = 0.97.

False alarm rate

In fusion regime A, the fused false alarm rate FAR(fused) is the proportion of false alarms out of all alarms raised by at least 1 source.

The combined FAR cannot be calculated from the false alarm rates alone, because sources do not report alerts with equal frequency, so each source FAR does not have an even weighting in the overall likelihood of occurrence of a false alarm – the frequency with which a source reports alerts must be considered.

Alerts consist of true detections plus false alarms, so the true detections are the proportion (100%-FAR) of total number of alerts, so true detections = (1-FAR) * alerts, so the number of alerts = true detections/(1-FAR). In our example, source i would raise ~121 alerts for every 100 real events, while source ii would happen to raise 100 alerts.

Of these alerts, for a single source, the number of false alarms is number of alerts minus true detections (e.g. 121-85=36 for source i).

alerts = events * DR / (1-FAR)

false alarms = events * (DR/(1-FAR) – DR)

To get total numbers of fused alerts across sources, we cannot simply add the numbers of detections from individual sources together because many of them will overlap in time, and adding would count those twice.

If we continue the naive assumption that non-detection is independent across sources, then the number of true alerts is the sum of individual true alerts minus those detected by both sources. The proportion detected by both sources is simply DRi * DRii.

Total alerts (fused) = events * (DRi / (1-FARi) + DRii / (1-FARii) - DRi * DRii).

The total number of false alarms is the sum of individual source false alarms minus false alarms that occur from sources simultaneously. However, individual DR and FAR numbers say nothing about the chance of simultaneous false alarms happening – that would need further data about the distribution of alarms and false alarms in time.

For the purposes of this simple initial analysis, if we make a further naïve assumption (which we will explore removing later in this analysis) that false alarms never occur simultaneously then the total number of false alarms is simply the sum of individual false alarms.

i.e.total false alarms (fused) = events * (DRi/(1-FARi) – DRi + DRii/(1-FARii) – DRii)

The false alarm rate being the number of false alarms over the total number of alerts, and with the number of events in top and bottom cancelling out, then with the stated naive assumptions, the false alarm rate in this fusion regime is:

$$FAR(fused) = \frac{\frac{DRi}{1-FARi} - DRi + \frac{DRii}{1-FARii} - DRii}{\frac{DRi}{1-FARi} + \frac{DRii}{(1-FARii)} - DRi\ .\ DRii}$$

In our example with two sources and say 100 real events:

FAR(fused) = (0.85/0.7 – 0.85 + 0.8/0.8 – 0.8) / (0.85/0.7 + 0.8/0.8 – 0.85*0.8) = ~37%

- Source i detects 85 events and raises 36 false alarms
- Source ii detects 80 events and raises 20 false alarms
- 68 of the events are raised by both sources
- 17 are unique to source i while 12 are unique to source ii, making 97 true alerts
- With assumption of no simultaneous false alarms, there are 56 false alarms

- There are 97+56 = 153 alerts in total.
- 56/153 =~37% (fractional differences arise because fractional false alarms have been rounded above).

So, by fusion regime A, if the naïve assumptions are true then we have achieved a much improved DR of 97%, but at the cost of an increased false alarm rate of 36.8%.

As the number of sources increases, in this fusion regime the false alarm rate continues to increase, while the detection rate increases only relatively modestly as illustrated in the section above.

Recall FAR(fused) is the proportion of false alarms out of all alarms raised by at least 1 source.

P(unique true alarm from source iii) = P(true alarm from source iii) – P(true alarm detected by both iii and i) – P(true alarm detected by both iii and ii) + P(true alarm detected by all three sources), the last term removing double counting from the previous two terms.

In this example source iii would raise 12 alerts, but the proportion unique to source iii would = 0.12 – (0.12 * 0.85) – (0.12 * 0.8) + (0.85*0.8*0.12) = 0.36%, i.e. on average with 100 events source iii would add no further unique detections. It would however also raise 3 false alarms.

The false alarm rate therefore rises to about 59/156, about 38%.

This modest increase in rates may not seem surprising for a source with a low detection rate. If source iii had a detection rate of say 70% then the same arithmetic would produce an overall fused detection rate of 99% but with an increased fused false alarm rate of 43%.

Time to detect

TTD(fused) = min(TTDi, TTDii, …)

This fusion regime alerts as soon as any individual sources alert.

With the example sources, the TTD could be as low as 10s, but only if source i does raise an alert. The mean TTD would be (detections by source i * 10s + other detections by source ii * 20s + detections unique to source iii * 30s)/total detections = 11.3s.

5.4.2 Fusion regime B – alert if both sources alert

While this could be considered a special case of the more general “alert if all sources alert”, that approach would become increasingly unlikely to be practical with 3 or more sources, so we consider explicitly the fusion of two sources in this regime.

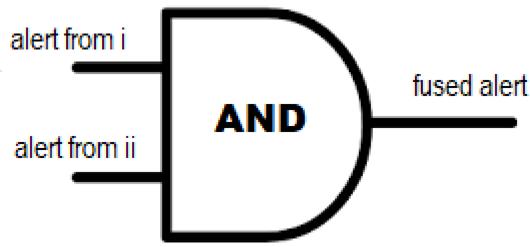

Figure 45 AND fusion regime

In contrast to the “OR” fusion regime which passes all indications as fused alerts, this “AND” fusion regime implies that information may be suppressed. While that model is used for characterisation of the scheme in this analysis, in practice the information need not be discarded – the same logic may instead be used to decide between higher or lower priority treatment of an alert in a user interface for example.

Detection rate

DR(fused) = DRi * DRii

False alarm rate

A false alarm in this regime would be where both sources reported a false alarm simultaneously.

However, as noted above, individual DR and FAR numbers say nothing about the chance of simultaneous false alarms happening – that would need further data about the distribution of alarms and false alarms in time.

The regime A analysis illustrated the assumption that false alarms never occur simultaneously – in that case the false alarm rate in regime B would be zero.

FAR(fused) = 0

If stopped vehicle events at a given location are relatively rare in time, then that assumption is not totally unreasonable. At the opposite end of the spectrum if stopped vehicle alerts at a given location are always present then:

FAR(fused) = FARi * FARii.

Continuing the example, fusing sources i and ii by fusion regime B would reduce the detection rate to 68%, but with the benefit of a false alarm rate somewhere between zero and 6%, probably much nearer zero.

Time to detect

TTD(fused) = max(TTDi, TTDii).

This fusion regime waits for both sources to alert, so has a longer TTD than the faster of the single methods.

5.4.3 Other fusion regimes

With more than 2 sources, other fusion regimes are possible. Indeed, extending regimes A and B becomes increasingly impractical as further sources are added: “alert if any one source alerts” suffers from increasing false alarm rates, while extending regime B “alert if both sources alert” to “alert if all sources alert” suffers from declining detection rates.

Once again, in the following analysis the implication is that fusion logic can suppress an alert, but in practice the decision could be about prioritisation of an alert in the operational user interface.

One simple alternative is “**alert if more than one source alerts**”.

A alternative approach that attempts to be more precise is “alert if the chance of this alert being a false alarm is below a threshold”, or phrased positively “**alert if confidence is above a threshold**”. Say the FAR of sources i and ii were individually both much higher than assumed so far, say 70%, and this put too high a load on operational staff, so a 50% threshold on false alarm probability was set i.e. a fused alert should only be raised (with high priority) if the fusion process derives an alert with less than 50% probability of being a false alarm, or phrased positively, the alert should be raised if confidence in it is above 50%.

When an alert is received from one individual source but not another, the chance of this alert being a false alarm is not simply the false alarm rate of the alerting source, because in fusion we have extra information from the non-detection by the second source. Mathematical treatment of this scheme is given below after re-examination of the assumptions.

Mixed heuristic approaches are also possible without use of probabilistic analysis, for example always alert on one source that is trusted most, and conditionally alert on other sources depending on corroboration by each other. These approaches can be seen to approximate "alert if confidence is above a threshold", but without a statistical understanding of the properties of the fused scheme. Incorporation of contextual factors into the decision to alert is discussed in 5.7.

5.4.4 Vehicle-based sources with potential for multiple reports

Fixed sensors should raise at most one alert for one event, but where connected vehicles can report, reports may be come from more than one vehicle for the same event. This could be where vehicles collide and each raise an automatic eCall alert, or it could be through people in passing vehicles raising a Waze alert, or sensors in passing vehicles detecting a stopped vehicle and informing a vehicle information service that is consumed by the road operator. It would be possible to treat each individual report as an alert to input to data fusion, but it may be easier to manage the fusion system if each source pre-aggregates its individual reports and uses these to update its reported confidence in an event. For example, a single manual eCall activation occurs may result in the eCall input to data fusion reporting an alert with 60% confidence, while two automatic eCall activations at the same time and location may result in the eCall input to data fusion reporting an alert with 100% confidence.

5.4.5 Improving on the naïve assumptions

Above we used a naïve assumption that the probability of sources i and ii detecting an event are independent, so the probability of both occurring at the same time was simply DRi * DRii. For sources with entirely different basis this may be reasonable e.g. the fact that an eCall alert has been raised may not obviously increase or decrease the chance a roadside radar device has of detecting the same vehicle, compared to the normal chance of the radar detecting a stopped vehicle event that has happened. However, not all sources will be independent, and at the extreme two different sources using similar technologies might spot exactly the same cases and miss exactly the same cases which do not suit their technology. In that extreme DR(fused) = DRi = DRii i.e. fusion has no benefit.

If there is alert data available from the sources from the same locations and times, further data analysis can be performed to improve on the assumptions of independence and lack of knowledge of simultaneous appearance of false alarms. In the data we can see the number of times that both sources detected the same event, and we can see the number of times (if any) that both sources simultaneously reported a false alarm.

Say we have historic data of alerts raised (or not) for a period with 100 real stopped vehicle events. Say as before source i detects 85 events and raises 36 false alarms, source ii detects 80 events and raises 20 false alarms. But say this time that 75 (rather than 68 in the earlier example) events are detected by both sources. Although this is not a big sample size, the numbers suggest that the DRi and DRii probabilities are not independent.

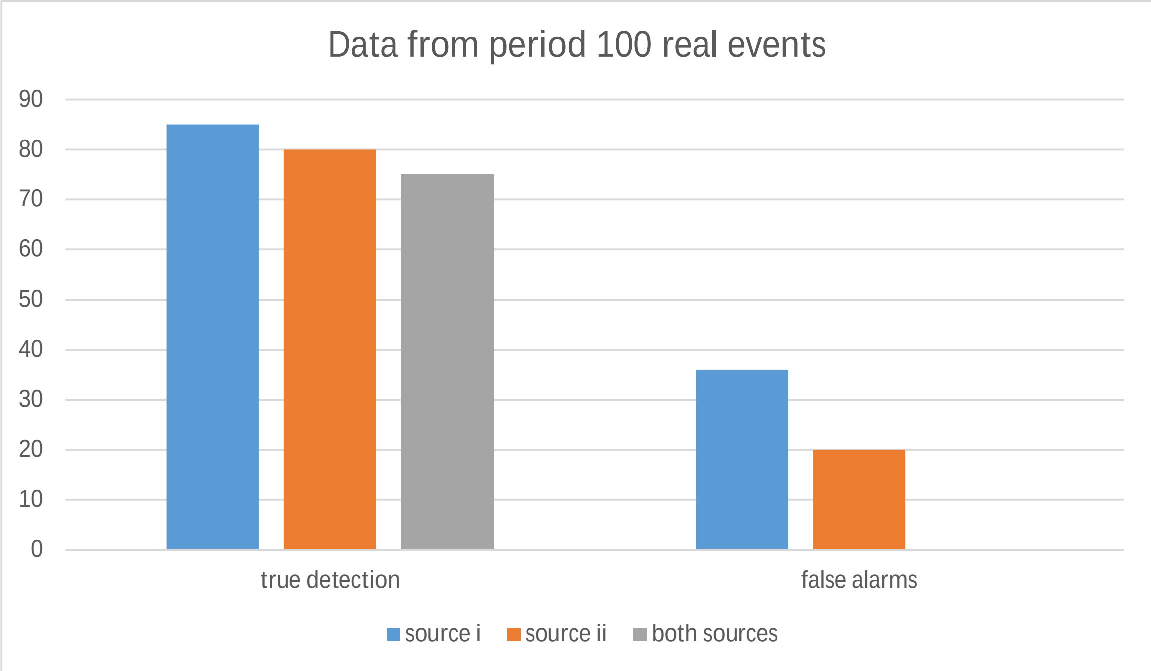

Figure 46 Characteristics of two illustrative alert sources

Although DR(fused) and FAR(fused) for each fusion regime could be noted directly from the empirical data, it may be useful to show the formulae at work.

In fusion regime A, DR(fused) has the form P(A or B) = P(A) + P(B) – P(A and B). With independent events P(A and B) = P(A) * P(B), i.e. DR(fused) = DRi + DRii – DRi * DRii. With dependent events, P(A and B) =P(A) * P(B given A), where "B given A" is usually written as "B | A".

If we define DRii|i as probability of detection from source ii given detection from source i, then:

DR(fused) = DRi + DRii – DRi * DRii|i

In fusion regime B, DR(fused) = DRi * DRii|i

With the numbers above DR(fused) can be observed to be 75% in regime B, most simply by directly noting that 75 events are detected by both sources.

DR(fused) in regime A is 0.85+0.8-0.75 = 90%.

The above numbers still do not allow refinement of fused false alarm rate – for that we would need to know additionally how many times false alarms occur simultaneously. Say in this example that there were 12 cases where false alarms from both sources occurred simultaneously. (So the other 24 raised by source i were unique to that source, while the other 8 raised by source ii were unique to that source).

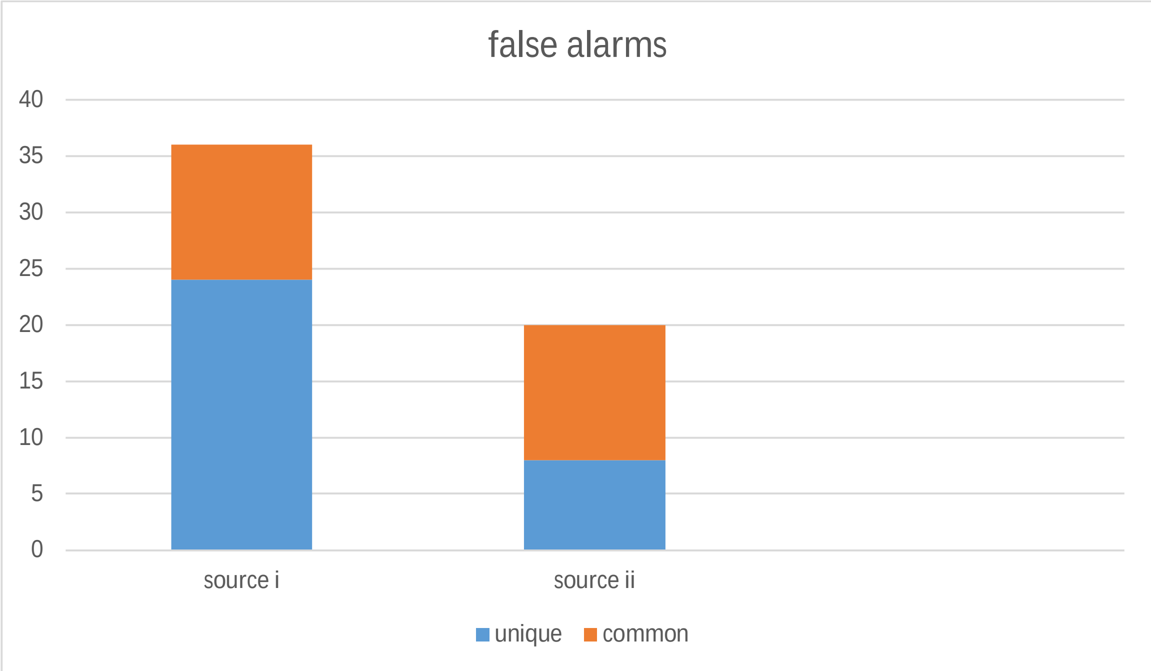

Figure 47 False alarms from two illustrative alert sources

The one additional datum of 12 common false alarms allows the total alert and false alarm populations to be visualised.

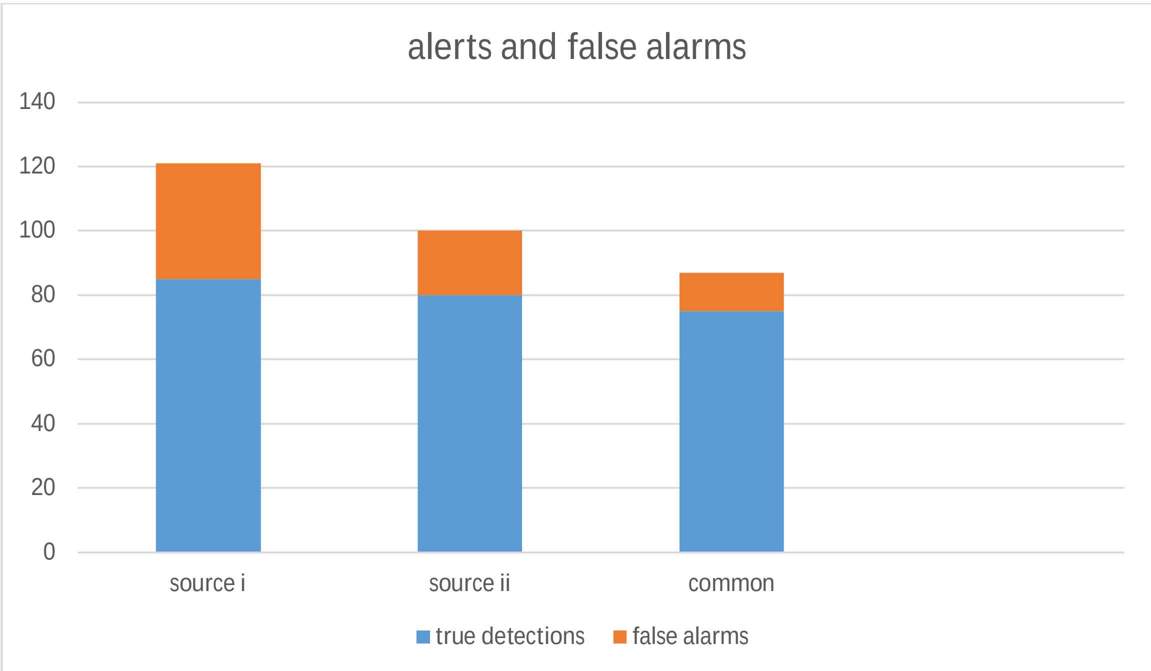

Figure 48 Alerts and false alarms from two illustrative alert sources

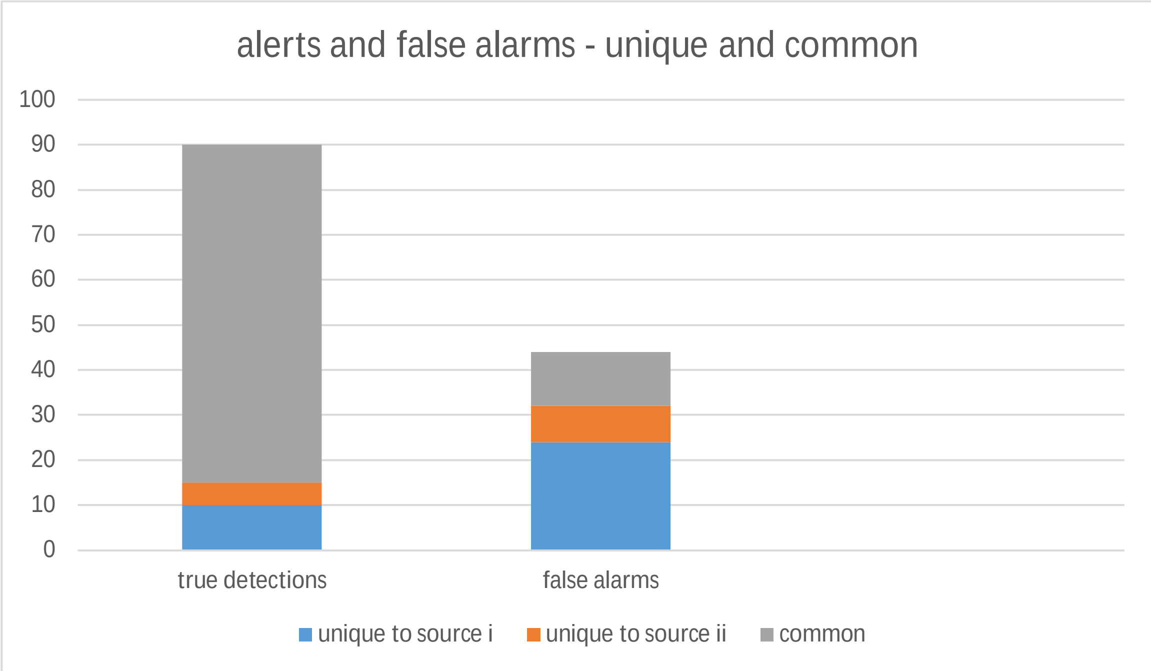

Figure 49 Unique and common alerts from two illustrative alert sources

FAR(fused) = number of fused false alarms / number of fused alerts.

In fusion regime B only the grey common true detections and false alarms lead to fused alerts. 12 common false alarms are raised, out of 12+75 fused alerts, so FAR(fused) = 12/87 = 14%

In fusion regime A, all true detections and false alarms are raised. The fused false alarms are the common ones plus the unique ones = 12+24+8 = 44, so the false alarm rate is 44/(44+90) = 33%.

These numbers allow consideration of the benefit of fusion for the choice of sources and fusion method.

With regime A (OR) which prioritises detection over false alarms, fusion has increased DR from individual maximum 85% to 90%, at the cost of a slightly increased false alarm rate from individual maximum 30% to 33%. These increases are more modest than the ones achieved with completely independent sources.

With regime B (AND) which prioritises avoiding false alarms, fusion has reduced false alarm rate to 14% from individual maximum 30%, at the cost of a decreased detection rate of 75% compared to individual maximum 85%. These reductions are not as significant as the ones achieved with completely independent sources.

The scale of the effect of fusion on the DR and FAR clearly depends on the number of common detections and common false alarms. This is illustrated in the following pair of graphs which expand on the example data used in this section by showing on the horizontal axis a range of possibilities for the number of common alerts and false alarms from the two sources.

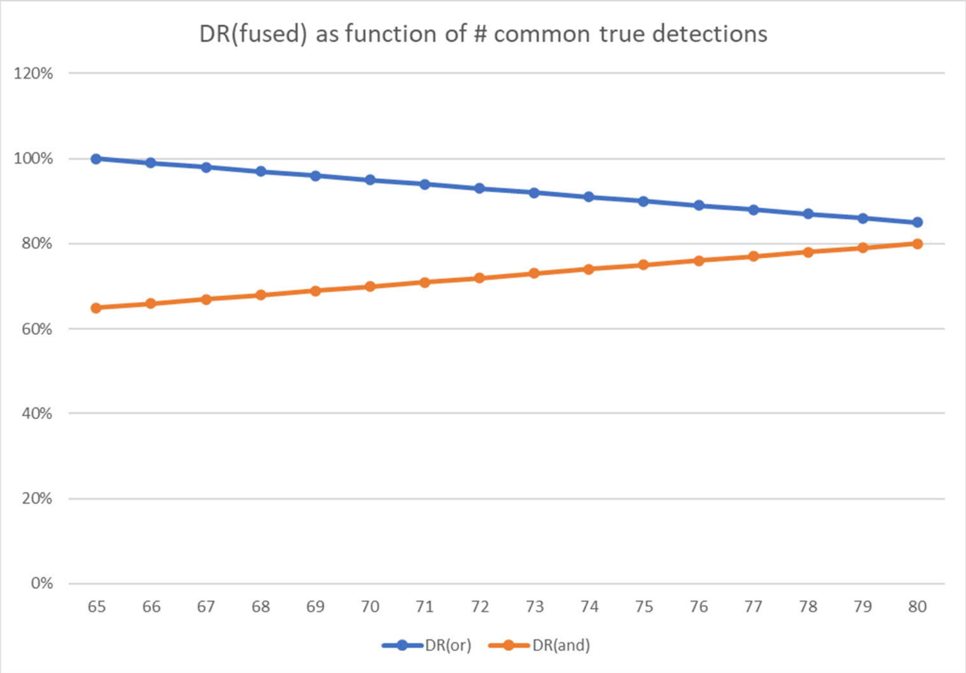

Figure 50 Detection rate trends in fusion regimes

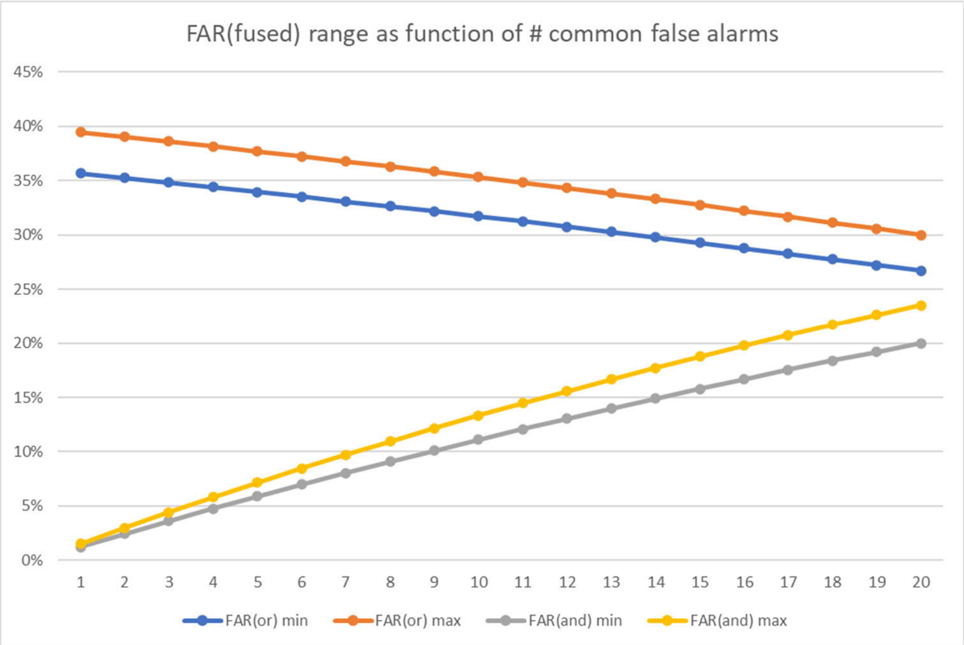

Figure 51 False alarm rate trends in fusion regimes

This analysis illustrates the intuitive fact that fusion would bring most benefit when the sources are most independent.

5.4.6 Significance of alerts reported by some sources and not others

If we know one source has alerted but that another source covering the same location has not, how does this alter the likelihood of the one alert present representing a real event? That depends on the range of times to detect for the missing source. To simplify the analysis initially, take the situation where we are confident that if a source was going to alert, it would have done so by now.

Consider P(event given i not ii) = the proportion of times an alert unique to source i represents a real event, out of all alerts unique to source i.

With the example of sources acting independently: there are 85-68=17 unique true detections, and say the full 36 false alarms are unique, so P(event given i not ii) = 17/53 = 32%.

Conversely P(event given ii not i) = 80-68 unique true detections / (80-68+ 20 unique false alarms) = 37.5%

Source ii detects slightly less often but is less prone to false alarms, so when it *does* detect, we have slightly more confidence in it than in source i. If we only had that single source in the first place, confidence in the alert would be 80%, but the fact that we also have source i and it has not detected an event reduces the confidence to 37.5% (if source i had agreed, confidence would have increased to 97% as noted above).

In the example of sources that demonstrate dependency:

P(event given i not ii) = 85-75 true detections / (85-75 + 36-12 unique false alarms) = 10/34 = 29%

P(event given ii not i) = 80-75 true detections / (80-75 + 20-12 unique false alarms) = 5/13 = 38%

A roads authority might consider using these numbers in different ways:

- In a priori analysis of the sources and potential fusion schemes, to better understand the choice of fusion regime and sources and how they will operate, to supplement the DR(fused) and FAR(fused) derived above. However, the extra insight does not seem very significant.
- To display the appropriate number, or a derived qualitative level, as a confidence indicator for the operator with each alert, and perhaps influencing prioritisation of the alert in the operator's user interface or workflow system. This concept is further consider in Chapter 7.
- In a sophisticated fusion scheme with more than 2 sources using the regime "alert if confidence is above a threshold" – the subject of the next subsection.

This analysis used the simplifying assumption that we have reached a time where we are confident that we have passed the time to detect of each source. However, the time to detect from a source may not be a constant fixed interval – there is likely to be a distribution of detection times, so that at any instant after a real event that is going to be detected by a source there is a (yet another) probability that the alert would be raised by now.

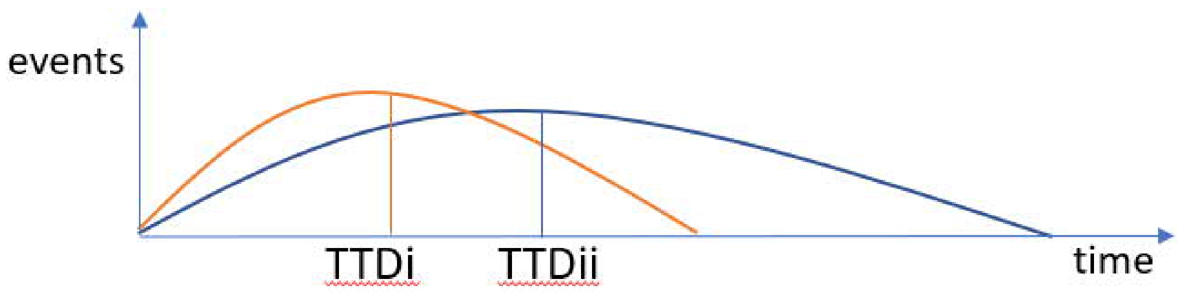

Figure 52 detection time profiles

If data is available on that population of detection times from each source then that data corresponds to probability distributions, but these cannot be applied directly into the calculation of probability of an event from alerts because for each new event we do not know the origin time at which event occurs. Considering just two sources, if we wish to use their time distributions, we would have to consider both the time distributions together to calculate, given that one source has alerted at time t_1, the probability that the other source, if it was going to alert, would have done so by the current time t_2.

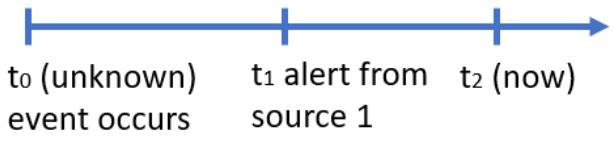

Figure 53 occurrences and times relevant to a stopped vehicle event and its alerts

Given that the data on past populations of detection times may not be available (it would require ground truth knowledge of the actual event occurrence times), this additional accuracy is not further explored here. A simpler and less accurate scheme is to not put *any* weight on a non-reporting source until some time selected from the relationship in average times to detect from the sources, and then assume as above that a non-reporting source is not going to report. Using the difference in mean time to detect of the sources as an offset from the first alert time would be rather harsh because half of all alerts by the slower source would not be expected by that time anyway. Another option is to use the mean time to detect of the non-reporting source as an offset from the first alert. Any such choice is a heuristic and not perfectly accurate.

A further complication is that not all sources provide clearance of alerts. When an alert arrives, it can be treated by the fusion system as described in this study, but at any time after that it is not known whether the alert condition persists. A heuristic solution could be used to reflect this. The crudest solution would be to assume the alert condition remains for a configured fixed time interval. Liu and Xiao (2019) describe the “Credibility Decay Model”, and a potential improvement, in which confidence in the reports from sensors decreases with time.

5.4.7 Fusion regime C – alert if confidence is above a threshold

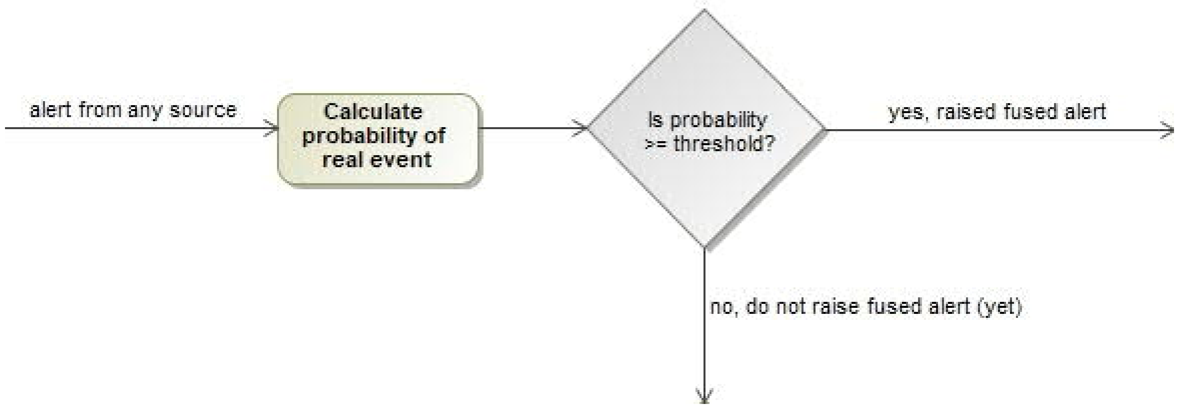

Figure 54 Selective data fusion regime

After receiving any alert, using knowledge of whether each source has alerted or not, the probability of the alert representing a real event can be calculated, and used to decide whether (or with what priority) to raise an alert to an operator.

The characteristics of this fusion regime are illustrated through a further example.

In this example say we have source i DR 80% FAR 70%, source ii DR 75% FAR 60%, source iii DR 20% FAR 20%, and say we have no data yet on their co-located performance, so we have to use naïve assumptions of independence. In this example say a threshold of 50% confidence (so 50% chance of false alarm) has been set. Assume for simplicity that all sources report at the same time (an assumption we remove later in the analysis).

The probability of an event may also depend on current conditions, including daylight, weather and traffic state. Adding these further dimensions makes the analysis more complex, and to do so properly requires further data that may not be available. Consideration of these factors is deferred until section 5.7.

On each alert, we calculate P(event | this permutation of sources alerting and not alerting).

Spelling this out for this three-source example:

alerts raised	P(event) if all sources equally fast	P(event) as a proportion
i only	P(event | i only)	number of times an alert raised by only i represents a real event / all alerts raised by i
ii only	P(event | ii only)	number of times an alert raised by only ii represents a real event / all alerts raised by ii
iii only	P(event | iii only)	number of times an alert raised by only iii represents a real event / all alerts raised by iii
i and ii only	P(event | i and ii)	number of times an alert raised by i and ii only represents a real event / all alerts raised by i and ii only
i and iii only	P(event | i and iii)	number of times an alert raised by i and iii only represents a real event / all alerts raised by i and iii only
ii and iii only	P(event | ii and iii)	number of times an alert raised by ii and iii only represents a real event / all alerts raised by ii and iii only
i, ii and iii	P(event | i, ii and iii)	number of times an alert raised by i, ii and iii represents a real event / all alerts raised by i, ii and iii

P(event) is also equal to 1 – P(false alarm). While one system of equations could express the whole problem, for simplicity the following sections illustrate the behaviour of the fusion regime with the example sources by using a simple way to calculate the probability in each kind of case.

Report by only one source

When an alert is raised by only source i or only by source ii – we know individually their false alarm rate is above the permitted threshold, and while corroboration by other sources would increase confidence, the lack of reporting by other sources will only decrease the confidence, so without showing the calculation here it can already be seen that a fused alert should not be raised.

When an alert is raised by only source iii – considered individually this source has a 20% false alarm rate, which is better than the threshold, but here we also know that sources i and ii have not detected an event, which reduces confidence in this alert.

Confidence is P(event given iii but not i or ii) = the proportion of times an alert unique to source iii represents a real event, out of all alerts unique to source iii.

In this example, for 100 real events source iii raises 20 true positive alerts (DR 20%) and 5 false alarms (FAR 20%, 5 is 20% of 20+5) . Without data on past false alarm combinations, one possible naïve assumption is that all 5 false alarms are unique (which may be likely if for example this is a connected vehicle source being fused with electromagnetic detection sources). Above it was stated that P(*unique* true alarm from source iii) = P(true alarm from source iii) – P(true alarm detected by both iii and i) – P(true alarm detected by both iii and ii) + P(true alarm detected by all three sources), in this case = 0.2 – 0.2 * 0.8 – 0.2 * 0.75 + 0.8 * 0.75 * 0.2 = 1%. Confidence = 1/(1+5) ~ 17%. The absence of alerts from sources i and ii has reduced confidence in a source iii alert from 80% to 17%, and it would not pass the threshold test – no fused alert would be raised.

In this fusion regime in this example, an alert will require corroboration from at least one other source before it is raised.

Report by all sources

Confidence in an alert from all three sources is 1 minus the false alarm rate of a fusion requiring all sources to alert. As shown above, accuracy of fused FAR requires further data on the co-occurrence of false alarms. At one extreme where events are rare and false alarms never occur simultaneously, fused confidence is 1.0. At an opposite extreme where alerts are always present, the false alarm rate would be FARi * FARii * FARiii so the confidence in this example would be 1-0.7*0.6*0.2 = 92%, still well above the threshold for raising a fused alert.

Report by two sources

Confidence in the case of an alert raised by only two from three sources is illustrated in the case that sources i and iii raise an alert when source ii does not.

If source ii was not present, an alert raised by sources i and iii has somewhere between zero and 0.6*0.2 = 12% chance of being a false alarm (which could be narrowed if further historic data was available to analyse), at least 88% confidence. If we assume that false alarms never occur simultaneously, then an alert raised by sources i and iii cannot be a false alarm – the non-detection by source ii does not affect confidence. But if we had data indicating a chance of co-occurrence of false alarms, then an event is not certain, and confidence is reduced by the non-detection by source ii.

Say we had historic data from 100 events showing that there were 2 false alarms common to sources i and iii without a simultaneous false alarm from source ii.

Confidence is P(event given i and iii but not ii) = the proportion of times an alert raised by only sources i and iii represents a real event, out of all alerts raised by only sources i and iii.

= (true alerts raised by i and iii – true alerts raised by i, ii and iii) / (numerator + false alarms raised by i and iii)

With detection independence this equals (100*0.8*0.2 – 100*0.8*0.75*0.2) / (numerator + 2) = 67%.

The corroboration of sources i and iii give enough confidence for a fused alert to be raised, despite non-detection by source ii.

Properties of regime C

To calculate a full DR(fused) and FAR(fused) for this example would require the full set of data for false alarms in common to permutations of the three sources.

Fusion regime C may for many sets of detection source characteristics turn out to behave equivalently to “alert if more than one source alerts”, but this depends on detection source characteristics and will not always be true. For example if sources i and ii often produced false alarms at the same time, such that 70 of the false alarms they would raise in this scenario were common (including 1 also raised by source iii), confidence for an alert raised by only sources i and ii would be (48/48+69)=41%, and there would no fused alert raised given the required confidence of 50%.

In scenarios where this regime behaves as “alert if more than one source alerts”, the detection rate (with assumption of independence) can be calculated as:

DR(fused) = P(any alert) – P(alert unique to source i) – P(alert unique to source ii) – P(alert unique to source iii) where P(any alert) is equivalent to DR(fused) in fusion regime A.

The probability of unique alerts is DR for the source minus alerts this source detects in common with other sources, for example for source i in the case of independence = DRi – DRi*DRii – DRi * DRiii + DRi * DRii * DRiii. In this scenario DR(fused) = 0.96 – 0.16 – 0.12 – 0.01 = 67%

Compared to having no fusion, this fusion regime maintains false alarms at an operationally tolerable level, at the cost of reduced detection rate in comparison to the individual sources i and ii.

Compared to fusion regime B with only sources i and ii (which with the DRs in this scenario would give DR(fused) = 60%), the inclusion of source iii and introduction of this more flexible fusion regime has increased the fused detection rate while keeping a tolerable false alarm rate. (In contrast if source iii was added by extending the fusion regime B approach of requiring all sources to alert, the fused detection rate would be only 0.8*0.75*0.2 = 12%).

The confidence probabilities can be pre-calculated for permutations of sources; dynamic calculation and decision-making at runtime is considered in section 5.7.

The above assumed that sources report at the same time, so knowledge of non-detection could be used, but if reporting speeds are significantly different then a road operator may not want to wait for a slower source before considering an alert from a faster source. In that case, when an alert from a faster source arrives, confidence that this is a real event is simply 1-FAR, as it would be if this fast source was the only source.

Time to detect

There is no simple single expression for time to detect in this fusion regime. A different expression can be stated for each permutation of alerting sources.

For example say that with 3 sources any 2 sources together produced sufficient confidence to raise a fused alert. Then in the case that only sources i and ii alert, TTD(i and ii) = max(TTDi, TTDii). In the case that all 3 sources alert, we only need 2, so TTD(i,ii, and iii) = min(max(TTDi, TTDii), max(TTDi,TTDiii), max(TTDii,TTDiii)).

Tool support

It would be feasible to create tools to supporting the kind of analysis illustrated in this section, with implementations of all the necessary equations, to help a national roads authority to select and characterise data sources and data fusion methods, given basic properties of the individual sources and their common performance if known.

5.4.8 Another fusion system example

SHADAR (in D2.1 and 5.3) has already quoted some DR and FAR numbers from real sources and national specifications, so this section presents an illustrative example of how a NRA might consider fusion if it had such sources available. While we wished to avoid making the example too large, it seems important to recognise some of the connected vehicle services sources which are available for exploitation. In this example the NRA does not intentionally cover locations with multiple sources, but geographic overlap of sources arises through the natural coverage of connected vehicle sources and the ability to reuse existing CCTV camera feeds for video-analytics-based detection.

The following sources are considered for the example, based on real archetypes.

1. The National Highways specification for radar-based SVD on open highways requires at least an 80% detection rate and at most a 15% false alarm rate. In past trials one source outperformed this requirement on both metrics, while other sources were not progressed and we might assume did not meet these requirements. In this example we will assume a

radar source meeting the National Highways specification but performing no better than that. Call this source i, which covers a portion of the road network, with DRi = 80% FARi = 15%, TTDi = 20 seconds.

2. The only complete video-based dataset we have seen reported shows 99.5% with zero false alarms. In other data sets the number of missed detections (false negatives) is not known because only the alerts have been verified by humans. Another dataset shows FAR 7.3% (including camera limitations, with unknown DR). Yet despite these impressive figures, and reuse of existing camera assets, this source has not become operational, so rather than assuming this performance will remain generally achievable in this example we use the figures from the National Highways generic specification for stopped vehicle detection solutions (TR 2643) which requires an 85% detection rate and a 15% false alarm rate. Call this source ii, which covers a larger portion of the road network than source i, with partial overlap, and DRii = 85%, FARii = 15%, TTDii = 20 seconds.
3. We have heard anecdotally of fixed detector installations with much higher false alarm rates than 15% (this might occur if calibration is not yet satisfactory, but the source is considered sufficiently useful to include in a fusion regime) so in this example a source iii with DRiii = 85% FARiii = 70%, TFiii = 20 seconds is available on a different part of the network from sources i, but partly overlapping with the coverage of the camera-based source ii.
4. eCall – 5.3 estimates FAR 40%; DR is not estimated but 16% of UK vehicles are equipped giving an upper bound at present which is increasing; one would expect currently only a small percentage of stops to involve eCalls even where equipped, but this would improve if driver education was to take place. In this example we say DRiv = 5% and FARiv = 40%. A key strength of eCall is that it can go where there is no detection infrastructure – it covers the whole network. In this example, a service is exploiting the eCall data rather than relying on manual handling of the voice call, and TTDiv = 60s. As 5.3 notes, serious incidents are more likely to produce multiple calls so DR would increase and FAR would decrease, but here we consider the more common stop cases without secondary incidents other than queues.
5. The NRA consumes a service from a connected vehicle service provider with stationary traffic warnings. Because it relies on detecting secondary effects, it has a lower detection rate than infrastructure, and a longer time to detect, but a relatively low false alarm rate. 5.1 showed that safety-related events detected by a vehicle typically reach a consuming NRA within a minute, and often much faster, but for stopped vehicles, unless the stopped vehicle happens to be the one reporting then those initial detections will be offset in time from the original stop event. In this example, considering a plausible proportion of live lane stops which cause clearly measurable secondary effects we take DRv = 50% FARv = 5% TTDv = 180 seconds. While it may not cover all the minor roads in a country, due to insufficient traffic, it may cover all the highways of interest to the NRA.
6. SHADAR D2.1 quoted the Netherlands specification for detection in tunnels, which required at least a 99.9% detection rate and at most a 1% false alarm rate. A tunnel represents a more constrained environment in which it may be easier to achieve such performance than on an open highway. It seems clear enough without further analysis that in locations where such an accurate source is available, alerts should be raised automatically without requiring fusion. Fusion of further lower probability sources could have additional value, but it is more marginal in this context. A tunnel may have dedicated systems rather than using a generic road SVD fusion system. In this example we therefore consider open highway locations where such an accurate source is not available. Source 6 is not further used in this example.

To keep the example simpler, we have not included a Waze source, but it is not unrealistic that a road operator may also fuse Waze events, as in the case of NDW in the Netherlands. 5.1 does not give a figure that represents DR, but the reported 93% match with other known events from any source suggests that detection rate is high. FAR is not known, but one might expect this to be lower than with eCall because accidental reporting is less likely. TTD is not known – it requires actions by road users, often with some time delay after the event. A road operator may narrow the number of connected vehicle sources to the one or two they consider are best, but since these data sources have wider applications beyond stopped vehicle detection they may be selected for different reasons, and then judged to be worthwhile including in data fusion. For further consideration of ongoing management and selection of data sources see section 7.6.

Summarising the alert source characteristics:

Source	Type	DR	FAR	TTD
i	radar	80%	15%	20s
ii	video analytics	85%	15%	20s
iii	radar	85%	70%	20s
iv	eCall	5%	40%	60s
v	private service	50%	5%	180s

There are therefore various possible permutations of sources:

- Tunnels, where the dedicated system is used – these are not further considered.
- On minor roads where there is insufficient traffic and no detection infrastructure, there is only eCall, but say in this example that these are not of interest to the highway authority, and these are not further considered.
- On roads of interest with no infrastructure, there is coverage from the connected vehicle sources (iv and v).
- On roads with some infrastructure coverage, there are the following permutations (each also having coverage from the connected vehicle sources):
    - source i
    - source ii
    - source iii
    - sources i and ii
    - sources ii and iii

The figure shows this graphically, using a completely arbitrary network (the geographic and topological relations are not significant). The connected vehicle sources (iv and v) are assumed to be available everywhere on this network, even where there is no infrastructure.

Conférence Européenne des Directeurs des Routes
Conference of European Directors of Roads

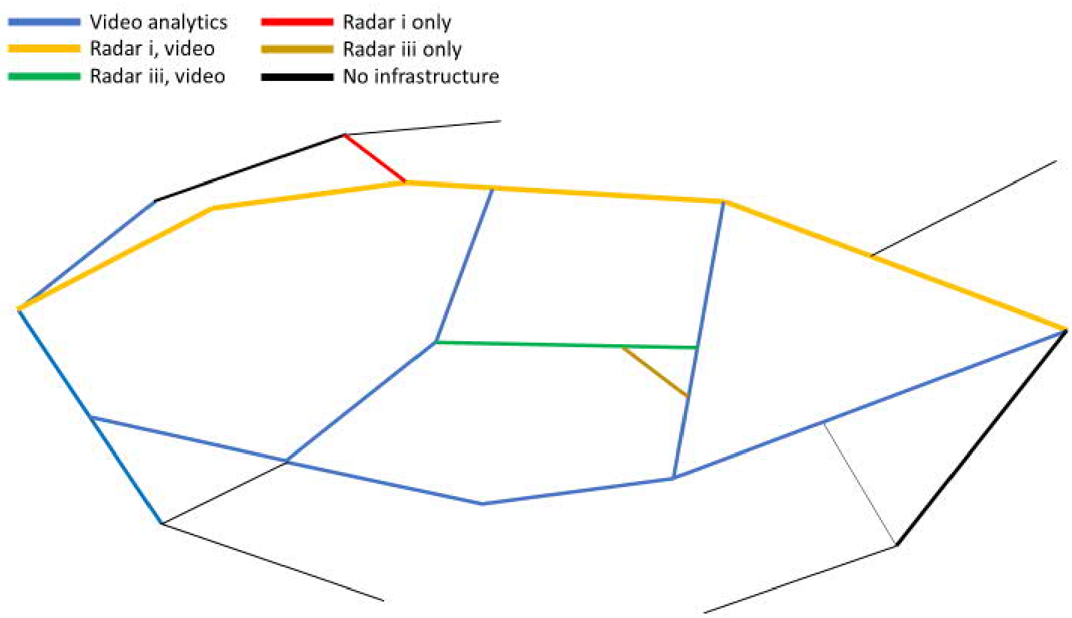

Figure 55: Illustrative network fragment with different kinds of coverage

In this fictitious but plausible illustrative example, through a combination of study data and estimates by considering the properties of each source[3], the road operator has estimated the following inter-relations between sources, expressed as rates per real event:

Sources	common true alerts	common false alarms
iv and v	2%	0
i and iv	4%	0
i and v	42%	1%
i, iv and v	2%	0
ii and iv	4%	0
ii and v	44%	1%
ii, iv and v	2%	0
iii and iv	4%	0
iii and v	44%	2%
iii, iv and v	2%	0
i and ii	75%	5%
i, ii and iv	4%	0
i, ii and v	4%	0
i, ii, iv and v	2%	0
ii and iii	77%	6%
ii, iii and iv	4%	0
ii, iii and v	41%	0
ii, iii, iv and v	2%	0

[3] For example, until there is co-located data to derive these numbers empirically, a temporary estimate could be made by considering the probability if the sources were independent (P(A)*P(B)) and adjusting that by considering the likelihood of common detection given the kinds of technology involved.

The following sections consider the properties of the fusion system in each part of the network, using the formulae described in earlier sections. It is assumed that no source reports confidence, only crisp alerts, giving less scope for more sophisticated fusion approaches.

No infrastructure – connected vehicle sources only

Fusion regime A (alert if any source alerts) would give: DR 53% FAR 10%. The TTD is as low as possible with these sources – 60s where eCall reports otherwise 180s. Although the increased detection by fusing eCall data is modest, it also gives improved detection times in the cases where it alerts.

Fusion regime B (alert if both sources alert) would give: DR 2.5% which seems not useful, despite FAR 0%.

Fusion regime C does not add anything to A or B when there are only two sources.

Since the eCall detection time is significantly faster, the lack of a source v alert when receiving an eCall alert would not affect confidence in the eCall alert, but the other way round: when receiving a source v alert, we may allow the lack of a corresponding eCall alert to affect confidence. In this case though, the effect is negligible due to the relatively low detection rate of eCall and would not be expected to have any practical consequence.

Source i plus connected vehicle sources

This part of the network has a radar source performing to a reasonable specification, plus connected vehicle sources.

Fusion regime A (alert if any source alerts) would give: DR 89% FAR 18%.

Other fusion regimes may be considered not to add significant benefit with these three sources: the false alarm rates are low except for eCall, yet to wait on corroboration for eCall would remove any benefit of that source.

Source ii plus connected vehicle sources

This permutation behaves very similarly to the one above. The simple fusion regime A would give: DR 92% FAR 17%.

Source iii plus connected vehicle sources

Although this uses the same mathematical formulae as the two cases above, the high false alarm rates of input sources mean fusion regimes might be considered. Fusion regime A would give: DR 92% FAR 68%.

The high false alarm rate creates an unwanted operational burden. Requiring corroboration – any two sources alerting produces a fused alert – would reduce the false alarm rate but would not be making best use of source v which only has a 5% false alarm rate. Regime C (alert if confidence above threshold) could provide a different balance between false alarms with detection rate. The threshold can be chosen to effectively be to allow connected vehicle alerts from source v directly, but require corroboration of sources iii and iv by one another in order to add further alerts to those raised by source v. In other words the fusion rules would be “v or (iii and iv)”. That scheme would slow the detection, since an alert from source iii would never be raised until an alert arrived from a slower source.

When an alert arrives from source iii, it has (1-FAR) = 30% confidence, which is not affected by non-detection by other sources since source iii is the fastest method. Alternatively, if an alert arrives from source iv and there has not been an alert from source iii, the normal 60% confidence reduces to 29%. With these statistics and speeds, there is no threshold that allows single source iv alerts directly, yet rejects single source iii events, and this lack is justifiable because a single alert from source iv, given non-detection from source iii, has been shown to be no more likely to represent a real event than a single alert from source iii.

Sources i and ii plus connected vehicles

This part of the network has 4 sources. The mathematical principles are the same but for simple regime A we expand using P(A or B or C or D) = all singles – all pairs + all triples – quadruple. This gives DR 95% and FAR 23%.

When alerts arrive from the fastest sources (i and ii), false alarm rate is calculated as if these are the only the sources. If both sources alert, which is the most common occurrence, the confidence is 94%. If only source i alerts, confidence is affected by the non-detection by source ii, and is only 35%. If this is not satisfactory for operational efficiency, the fusion regime could require corroboration by a further source iv or v, which would restore confidence in the alert being a real event, but the fusion system would have to wait until the slower sources reported.

Sources ii and iii plus connected vehicles

The method is similar to the case above, but this time the basic fusion method A gives DR 97% FAR 68%.

If both of the fastest sources (ii and iii) alert, then despite the high false alarm rate of source iii the overall confidence in a fused alert is high. However if only source iii alerts, the confidence is only 4%! Spending operational resources on an alert with 4% confidence is not likely to be considered worthwhile, and so this unreliable source should be used only in a more selective fusion regime. Again, if the alert was corroborated by the slower connected vehicle sources, then its confidence significantly increases and may be worth raising for operational attention.

If the connected vehicle sources were themselves highly reliable then the addition of source iii in a selective fusion regime would have no benefit.

Summary of the fusion example

In summary across all permutations: the simplest fusion method (A), raising an alert if any source alerts, seems likely to give the best results when the sources are relatively reliable, but when false alarms become a problem then the selective method (C) improves on the false alarm rate while maintaining the alerting performance of any more reliable detection sources. Fusion regime C would tend to be beneficial where there is more than one unreliable source, or where an unreliable source provides extra information that could help response.

While it should be useful for a road authority to consider what permutations of sources exist, and what fusion methods could give the best balance of results, this may not entail the configuration of different fusion settings for different parts of the network, because fusion method C can act like method A if the threshold is set appropriately i.e. it may be possible to use fusion method C on an entire network with a carefully chosen confidence threshold that would not cause suppression of alerts from reliable sources but only unreliable ones without corroboration.

If the simple statistical analysis presented in this section appears to a road operator to be too complicated, an even simpler heuristic approach could be adopted instead – such as configuring all sources as requiring corroboration or not, based on their perceived reliability. However, a heuristic approach may not provide an understanding of the optimum balance of detection rate and false alarm rate.

5.5 Location identity – a practical challenge for data fusion

When there are two true alerts at similar locations from different sources (or even from the same source using multiple individual traveller reports), they may represent the same stopped vehicle event or two different ones.

A practical approach is for the alerting fusion system not to care how many events there are, but rather whether there are any events in a given section (e.g. 100m section). This would be consistent with operational policy used by England's National Highways where their dedicated SVD systems report only in terms of whether there is an alert in a 100m section, then operational staff manage the details in that section.

If there were actually two separate events in same road section, fusing incorrectly into one alert in most cases would not be significantly detrimental, because follow-up with cameras or other operational actions would discover this. Considering the risk that two stopped vehicle locations judged to be in the same section are on either side of one camera, the section boundaries should be chosen to minimise this risk, and/or the operational staff could be required to check 2 cameras. If the risk still occurs, or there is no camera coverage, a mitigation is that further operational action may still discover the second event.

The opposite risk – that there is one real event but location accuracy from the detection sources has them some distance apart – that is one reason that a coarse resolution such as 100m may be used within the fusion system. If the fusion regime being used would supress alerts from either single source (for example if they have high false alarm rates) then it becomes more important to mitigate this risk, and to choose a relatively higher size of fused alerting section, allowing the sources to corroborate each other. When the fusion system sees close but separated alerts, it could note uncertainty of whether they are the same event, and this could even be expressed in an operational user interface (although that is the kind of feature that may not add enough operational benefit to justify the complexity and use of screen space – see Chapter 7).

While operationally the requirement for longitudinal accuracy seems not too demanding compared to the capabilities of technology, the need to identify the correct carriageway is very important. In tracking-based methods like radar and video, carriageway identification is inherent in how these technologies work. In connected vehicle methods it should not be taken for granted – naïve location matching from a single GPS position could identify the wrong carriageway. The connected vehicle technologies may use further data such as previous positions to provide a more reliable carriageway identification, but if they simply give a single set of coordinates then any fusion regime that might suppress an uncorroborated alert should consider the possibility that the event is on the adjacent opposite carriageway.

Using fusion to improve location accuracy or precision

While the basic requirement for SVD is to identify the correct carriageway and a coarse-grained longitudinal section, it also could be of benefit if the lane can be determined accurately, which may enable automated signal settings for lane management for example. In general fusion of agreeing detection sources increases confidence in their individual outputs. We therefore briefly considered whether fusion of two sources that individually have insufficient confidence in lane determination could produce enough confidence in lane determination to trigger automated signal settings. For example there might not be enough confidence from radar or eCall alone, but if they agreed on a lane, the confidence might be considered sufficient.

Where there is infrastructure for lane management, there is likely to be good camera coverage, so this kind of location fusion may only be worthwhile where the alerts are quicker than the time taken for human operational response to verify the lane and set signals.

To determine the value in such a regime, the method would require ground truth lane data, for example from human observation, to derive the probability of correct lane assignment for each individual source and the increased probability when the sources agree on the lane. This method adds to the effort of creating a fusion system and in many cases may not be of sufficient benefit for consideration.

5.6 Study of fusion of two real data sets

To explore and illustrate the statistical data fusion techniques described above, we obtained and examined two real historic data sets which overlapped in time and in location.

For a limited trial period in 2020-2021, two sensor-based stopped vehicle detection systems were employed on the same highway in Europe. Each system used a different detection technology. The exact location and nature of the sources has been anonymized; the purpose of this study is to explore the potential of data fusion, not to identify the performance of specific detection sources (which were private trials in this case).

The SHADAR project obtained records from the period:

- Source A: Data recording each true positive and false positive from source A in the 3-month period, which had been verified by manual checking of camera footage.
- Source B: Data from a traffic management system showing alerts raised by source B and related operational actions.

Because these data sources were not designed to support the kind of data analysis undertaken by SHADAR, they were (not surprisingly) not ideal for that purpose. Correlating the sources is not straightforward and requires assumptions. Data fusion would be simplified if reporting and logging were designed with a requirement to support data fusion.

This study did not have access to a complete record of ground truth. We have assumed that the manual checking of alerts from source A is correct, but no 24x7 check was performed so we do not know what stopped vehicle alerts may have been missed.

Nevertheless, analysis can still derive information about each source, and about data fusion, which is not apparent from each individual source alone.

Apparent characteristics of source A – before consideration of source B

The data covered the period from 12th October 2020 to 13 January 2021 and used 14 devices: 4 in a tunnel and 10 on open highways. Our study used only the data for the 10 devices on open highways.

There were 640 true positive alerts and 30 false alarms in total, but these include not just alerts about stopped vehicles but also on some kinds of congestion. In this study we take only the alerts designated as detected stopped vehicles. There were 587 of these: 564 true positives and 23 false alarms. The false alarm rate was therefore 4%.

The detection rate for this source cannot be determined because there is no ground truth – we do not know how many events were missed. Nothing at all can be inferred about detection rate without looking at other data sources.

Apparent characteristics of source B – before consideration of source A

Data from the same period was obtained from a traffic management system that processes stopped vehicle alerts from a detection system. These data sets are large, because a wider set of locations are covered, multiple events may be logged for one potential alert, and there are more alerts per location than reported in source A. While there is no ground truth, every alert raised by the system would normally be investigated by an operator, and the result logged should determine whether a stopped vehicle event was confirmed or whether the alert is considered a false alarm for operational purposes.

Overall (not limiting to locations also covered by source A) there were:

- 6447 alerts confirmed by operators to correspond to a stopped vehicle.

- 31244 alerts where the operator said no stopped vehicle was found. These could be false alarms, or they could be very transient events in which a vehicle stopped then moved away again before the operator could examine the location and the system had not cleared the alert in time to prevent the operator from looking. Further exploration of this very high number led to a report that the traffic management systems had not been correctly handling reported clearances in some cases, an error not fixed until after our sample period. Spot-checking comparison of the two log sources seemed to confirm the reported behaviour. These cases should not be considered as false alarms and should be removed from consideration. We have not confirmed the number of cases, but an upper limit is 8394, so a lower limit for the number remaining is 25351.
- 2501 alerts which the operators considered an "invalid event" (rather than "no event"). An example cited to help explain this category was debris on the road.

The operational experience equates to a false alarm rate of 33745/(33745+6447) = 84%. However, due to the apparent problem with clearances of transient stops, this cannot be taken as the performance of source B at the time. Excluding the alerts raised by the traffic management systems in error (using the upper limit, since we have not counted those cases) the recoded operational experience of false alarm rate was 25351/(25351+6447) = 80%. Whether the remaining cases classified by an operator as "no event" should be considered false alarms is arguable. Study of the data shows it is quite common to have an alert raised and then cleared within 1-2 minutes, but where the operator reaction is even faster, often within a minute of the alert. Perhaps the operator reaction is faster than the service level for clearance by the detection system - then legitimate but transient stopped vehicle events will lead to the "no event" classification. These are not confirmed false alarms like those confirmed by human study of camera footage. If these are excluded from the data set, and we retain the assumption that the "invalid event" classification does represent a false alarm, the false alarm rate from source B was 2501(2501+6447) = 28%. The following analysis takes "invalid event" entries to represent false alarms but excludes the "no event" classification for which neither positive detection nor false alarm can be confirmed.

Correlation of the two sources

Due to the data sources not being designed to support our purpose, the correlation was not straightforward in either time or location.

Source A had incomplete time data – each record used a 12-hour clock without saying whether it was a.m. or p.m. We took an optimistic view when looking for matches – if there was a match in either a.m. or p.m. then we assumed that was the time intended.

Entries within 5 minutes of each other, and at a matching location across the two sources, were assumed to be describe the same stopped vehicle event.[4]

The source A data that we obtained did not identify the location of the stopped vehicle, only the location of the device (using the nearest 100m marker post).

The source B data did not identify the device but the location of the 100m marker post nearest to the stopped vehicle.

Source A used a uni-directional sensor but we did not have data on the direction. Again, we took an optimistic approach: we looked up to the reported range of a source A sensor in both

[4] Further study of a smaller sample of times has shown that matched alerts from the two sources were often over 3 minutes apart, so we may have obtained more matches if we had widened the time matching criteria.

directions from the device location and the direction in which we found most matches with source B was assumed to be the direction of the source A device (there was always one direction with significantly more matches, so this assumption seems reasonable).

In a real operational deployment of data fusion, the assumptions made in this section should be improved upon by further study of the source system details, but they are practical for the purposes of this analysis whose purpose is not to illustrate the characteristics of individual sources or technologies but to illustrate the additional information and understanding that can be gained through data fusion.

Restricting the source B data set to the locations potentially covered by source A, the following numbers of alerts were present:

	Source A	Source B
alerts in common locus	587	1930
human-verified	564	1355
false alarm/unverified	23	575

The figures for source B exclude 4991 “no event” cases which include both alerts raised due to the traffic management system error and potential transient unverified alerts where the operator response was faster than the detection system clearance.

The numbers of matches in time and location, using the criteria identified above, are as follows:

	Source A	Matched in Source B	%Matched in Source B
true alerts	564	276	49%
false alarms	23	3	13%

The reasons for lack of matches may include over-simplification of the location matching. The locations of devices in 8 out of 10 cases were assessed to be very similar across sources, but in one case the location of the nearest source B device was not confirmed, and in another case the devices were some distance apart with an intervening bridge. Restricting the data set to 8 pairs of co-located devices, the rate of matching increases only slightly:

	Source A	Matched in Source B	%Matched in Source B
true alerts	540	268	50%
false alarms	20	2	10%

The level of matching was not even across device pairs. The rates of matching per device pair were: 44%, 31%, 31%, 100%, 72%, 45%, 23%, and 73%, and for the 2 less similarly located pairs 30% and 36%. If assumptions were to be investigated and verified, remaining low levels of matching may suggest where to focus attention on the performance and calibration of specific sensors.

Inferring detection rates

The figures allow the inferences about the detection rates of each source, which were not possible when considering each dataset in isolation, although this requires assumptions.

- If one was to assume that between the sources, all stopped vehicle events are found, then the detection rates of each source can be calculated. That assumption may not be valid, so a detection rate calculated is a maximum (if other assumptions are true).
- Using an event detected by one source as part of the calculation of detection rate for another source that missed the event requires a common definition of what should

constitute a stopped vehicle event, and it assumes that both sources are required to detect all such events anywhere in the locus of overlapping detection coverage.

Using only the 8 similarly located device pairs:

	Source A	Source B	Matched	Total
Validated events	540	1216	268	1488

These data correspond to detection rates as follows:

	Source A	Source B
Inferred detection rates, given assumptions	540/1488 = 36%	1216/1488 = 82%
False alarm rates, given assumptions	4%	31%

Applying data fusion in real-time

This section explores the characteristics that would have been achieved if the two sources had fed a data fusion system generating alerts to the operators. All 10 open-highway devices from source A and their locus are considered in this section.

With an “OR” regime (regime A described above), *with the same significant assumptions used above to express rates for each source*, the fused detection rates and false alarm rates would be as follows:

	(max) DR	FAR
Source B alone	82%	30%
Source B fused with source A (OR regime)	100%	27%

The fusion in this scenario appears entirely beneficial because it detects more events and even although there would be a higher absolute number of false alarms, the false alarm rate FAR (which is relative to the total number of alerts) decreases because source B had a much lower false alarm rate.

With an “AND” regime (regime B described above), *with the same significant assumptions used above to express rates for each source*, the fused detection rates and false alarm rates would be as follows:

	(max) DR	FAR
Source B alone	82%	30%
Source B fused with source A (AND regime)	17%	1%

The severe “AND” regime almost eliminates false alarms, but at the expense of a much lower detection rate.

With only two data sources, fusion regime C (alert when confidence is above a threshold) is not very useful, but the calculations on confidence may be useful to consider. An alert from source A would initially have a confidence of 96% (1 – FAR) before considering source A. Assuming a simplified approach to detection time, and knowing that source B reports quickly, say a time had elapsed at which we expected source B would have alerted if it was going to alert: then if source B has not alerted, the confidence in the source A alert is reduced to a level that could be precalculated individually for each device, or globally for the scheme or system. The table below shows that the confidence using scheme-wide statistics would be 49%, using the equation for regime C. Deriving the equivalent figures for source B is problematic in this study due to the large portion of the alerts from the ‘no event’ class for which we have incomplete knowledge. Simply excluding this class of alerts, the confidence would be 70% initially before considering source A. Lack of a source A alert, after a sufficient period of time, would reduce confidence to 56%. (Alternatively, including the ‘no event’ alerts and taking confidence to mean the probability that a real event will be observed by the operator, the confidence from source B would initially be 20% from data in this study, dropping to 16% from absence of source A.) Confirmation by source A would increase confidence to 99%. This knowledge could be used to configure the operator’s systems, for example to influence the priority with which alerts are displayed.

Confidence:	Initially	Absence of other source confirmed	Both sources alert
Source A first	96%	49%	99%
Source B *(known subset)* first	70%	56%	99%
Source B *(all in study)* first	20%	16%	99%

Time-to-detect

Traffic management systems producing logs with timestamps are typically synchronized using Network Time Protocol, but it was not confirmed that the specific data sources in this study used synchronized clocks.

Study of samples of source B alert data shows that the alerts often come in series, perhaps due to standing traffic forming and moving over time, and that characteristic coupled to the lack of ground truth data makes time-to-detect more difficult to analyze. Examining a smaller set of source A alerts from one device: there are corresponding source B alerts within 4 minutes in all cases; source B raised alerts before source A (10 seconds to 3 minutes earlier) in two thirds of the examples, but source A was earlier in the other third of examples studied. This suggests that the use of both sources together could give benefit through earlier detection, but the different levels of sensitivity and the lack of ground truth make it difficult to confirm or quantify.

Summary of benefits

Using two co-located sources brings significant knowledge about those sources that was not apparent from using each source alone, without the expense of a full ground truth study requiring constant human vigilance of the entire set of locations.

False alarms rates from each source were already apparent after human investigation of alerts, but the detection rates were unknown. Was each source finding all there is to find? Study of the data together shows they were not, and shows possible detection rates from each source given certain assumptions.

Use of these sources together in a data fusion system would have increased the detection rate and reduced the false alarm rate when compared with a single-source operational regime.

5.7 Fusing and decision-making dynamically at runtime

Section 5.4 above considered the a priori characteristics of fused detection systems using probability. Although it would be possible to perform the probability calculations dynamically as alerts arrived, it would also be entirely possible to perform all of the probability calculations in advance, for all permutations, and then define fixed rules for how to raise alerts. The number of permutations of sources is not high. For example, when fusing two data sources there are only three permutations to consider when alerts are received: {only source i detects, only source ii detects, both sources detect}; even with 4 sources there are only 15 permutations, whose significance in terms of probability could be pre-calculated and considered for the definition of operational rules. The probabilistic analysis is done a priori to inform the rules, rather than being used calculated dynamically to decide the impact at runtime.

Calculating probability at runtime becomes useful if the effect of significant contextual factors is known. Experience suggests that the most significant contextual factors are **weather** and **traffic state**. Daylight versus darkness may also be a factor in performance of camera-based sources. Inclusion of these factors requires:

- the conditions to be measurable (or reliably predictable)
- knowledge about the impact of such conditions on the performance of the detection sources.

When particular conditions are present, the significance of an alert or a non-alert changes. For example, say we video is our primary alert source with excellent performance in good visibility, but its performance is known to badly degrade in thick fog, and say we have a secondary source from connected vehicles which has lesser performance but is known to be relatively unaffected by fog. If thick fog is present, and we have an alert from a connected vehicle but not from video, we do not want the non-detection from video to exert undue influence on the fused result, whereas with good visibility the non-detection would reduce the likeliness judged from the connected vehicle detection. If fog was the only environmental variable, and there were only two sources, then again it seems more practical to predetermine rules than to calculate any influence dynamically, but if the road operator has identified a higher number of significant contextual factors and has many sources, it may be more practical to calculate the probabilities as the alerts arrive and use the resulting confidence to determine whether to raise the fused alert.

How would this kind of detailed contextual knowledge emerge? When detection sources miss real events or raise false alarms, it might be assumed that their technology suppliers would be keen to investigate some cases to allow improvement in their products so they can meet KPIs. Through this or other mechanisms, data on contextual factors influencing detection performance can be built. For example, the set of false alarms in a video-based data set from an operational trial were categorised by the supplier as follows:

Reason	Occurrences	Implications for alert processing
Lights and reflective road markings	9	Further explanation would be needed to understand whether this correlated with any measurable condition.
Sun glare and lens flare	10	May correlate with measurable weather data.
Objects on the road	5	May correlate with other traffic events known to the wider integrated traffic management system.
Software error	4	Only for supplier consideration
Water on lens	1	May correlate with measurable weather data.

If there is enough ground truth data, a reason that occurs in significant numbers, and a correlated data source that can be integrated into the traffic management system, then the conditional probabilities of non-detection or false alarm when the condition is present or absent can be calculated from known ground truth data. Then when an alert arises from any source, the presence or absence of the correlating condition can be used to select the appropriate probabilities to determine the correct approach to this alert in the fusion regime.

As an additional feature, if a source was to provide numeric or qualitative confidence with an alert, rather than the binary choice between an alert or silence, this could be fused at runtime to produce an overall confidence or probability of alert. Although it seems reasonably likely that some detection technology providers use such figures internally, none has so far to our knowledge offered to provide this as an output.

5.8 Correlation with more general event sources

The analysis has so far focussed on sources that directly detect stopped vehicles, but as D2.1 recognised there are also sources which report secondary effects such as queues, which in some cases may be due to a stopped vehicle.

The influence of these sources on stopped vehicle data fusion may depend on operational policy. If every report of a queue leads to an operator immediately verifying the situation with a camera, that would identify any stopped vehicle without need for any data fusion. Alternatively, without that operational policy to immediately check alerts raised by a queue detector, there could be merit in considering the queue detector as an additional stopped vehicle source.

If treating a general event source exactly like an additional stopped vehicle source, an event (such as queue) for a reason other than a stopped vehicle would be considered to behave in the fusion system like a false alarm! This is not saying that the queue warning is a false alarm – it is probably a real queue – just that within the fusion system it is treated in the same way as false alarms from other stopped vehicle detection sources to achieve the correct influence on the probability of stopped vehicle events.

Since stopped vehicles are not the most typical cause of queues, it would seem an unsuitable choice to employ fusion regime A (any queue would result in a fused stopped vehicle alert, with a high false alarm rate), or regime B (a stopped vehicle alert would never be alerted until a queue formed), but regime C could be useful – a reported queue would increase confidence in an alert reported by another detection source, which would be useful if the other detection source was not sufficiently reliable and no fused alert had yet been raised. However, by the time a queue has formed, a large part of the safety hazard has already occurred, in other words waiting for secondary effect detection is much later than is ideal.

Further study could consider the relative merits of this approach compared to the approach of treating traffic state (including queues) as context affecting the stopped vehicle event probability as described in section 5.7. The mathematics should produce equivalent results, so it is more a question of the ease of thinking about and managing the data.

5.9 Conclusion from this data fusion study

This study has illustrated methods that allow NRAs to understand the performance they would achieve from fusing candidate data sources.

Study of characteristics can influence the choice of data sources to invest in – the sources that are the most independent bring the most benefit when fused together.

The study has shown how confidence a fused alert can be calculated from various factors. This information can be used to determine how an alert is presented to an operator (further explored in Chapter 7).

Table 18 and Figure 62 summarise the findings using the “Now-Wow-How” categorisation of the apparent feasibility and originality of the ideas explored in this research.

	Subject	Category	Description
NOW	Harness multiple detection sources	Accuracy	Some sources provide more accurate location
		Timely	Get alerts as fast as the fastest source in each case – but at cost of increased operator workload
		Information	Connected vehicle sources, potentially enhanced by lookups, can provide extra information e.g. vehicle type.
	Fusing alerts	Reliability	In addition to benefits of using multiple sources, fusion can increase detection rate and decrease false alarm rate, potentially without significant impact on operator workload.
	Multi-source ground truth study	Information, Accuracy, Reliability	A ground-truth study provides knowledge of relative and absolute performance of detection sources, allows more accurate assessment of confidence for future alerts, and therefore more reliable decisions about how to prioritise.
WOW	**Subject**	**Category**	**Description**
	Probabilistic fusion influencing priority	Timely	Alerting user interface gives higher priority to alerts with higher confidence and can help save operator time.
HOW	**Subject**	**Category**	**Description**
	Machine learning on raw data	Accuracy, Reliability	Machine learning fusion of raw sensor data seems likely to provide good accuracy and reliability, but seems less easily feasible than fusing the current outputs from technology providers.

Table 18: How-Now-Wow categorization fusion

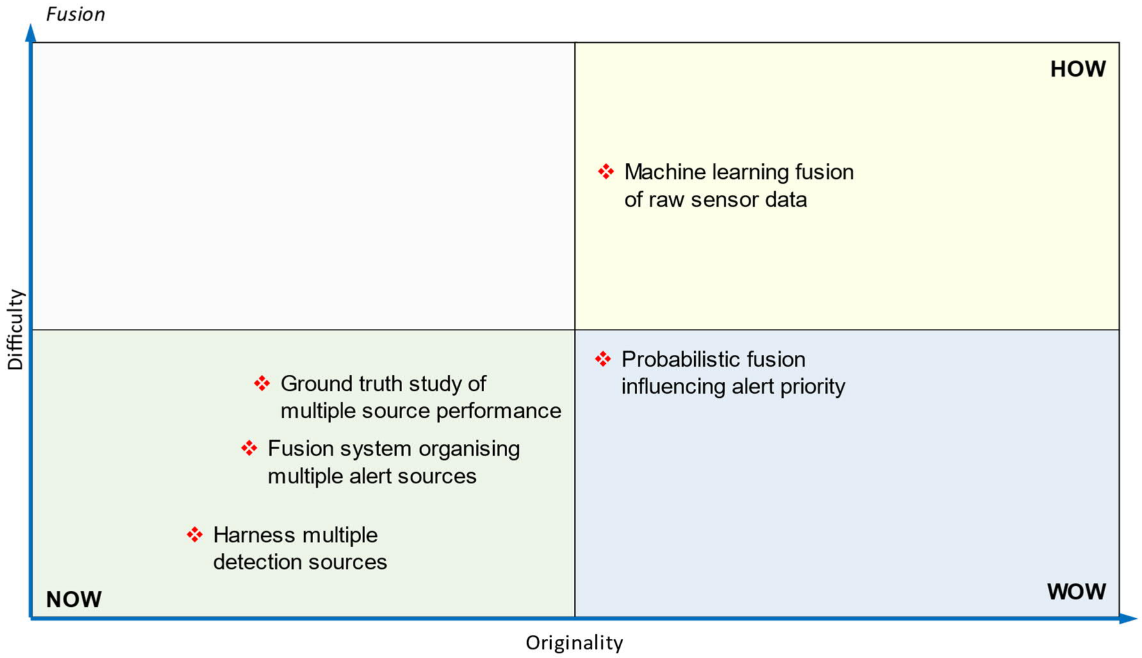

Figure 56 How-Now-Wow matrix fusion

6 Human reporting behaviour

Various public service websites have been created to guide drivers, such as smartdriving.co.uk, passmefast.co.uk, idrivesafely.com. The advice focuses on drivers that will stop to provide aid or assistance of some kind, and where the person contacts the authorities through the emergency number or SOS phone alongside the road, if available. General guidelines for drivers when approaching and passing an incident/accident are:

- Approach with caution
- Warn others
- Keep yourself safe and don't aggravate the situation
- Only stop for help if you can do this safely
- Prevent further danger
- Get help (SOS phone, 911,999,112)

The detection of stopped vehicles through drivers/bystanders/passers-by is still one of the most valuable contributions to detect stopped vehicles, especially in areas where no roadside detection is available. Still, drivers are barely trained in how to react to an incident/accident situation. Availability and content of information on the topic differs per country.

6.1 Means of reporting

For the road user/driver and his/her passengers means for reporting an incident are:

- (hands-free) phone call
- Navigation apps of service providers with report abilities (Waze, TomTom, BeMobile…..)
- Manual eCALL
- Manually triggered C-ITS or SRTI message

Phone call

There is a distinction between stopping at the location of the incident or being a passerby. Reporting as a passerby (preferably a passenger, not a driver), the precise location can be more difficult to convey.

Phone calls may be trusted until otherwise proven. A policy of passers-by reporting incidents through emergency numbers would create a high workload for the emergency centres where the information of one or two calls would be sufficient for assessing the report.

Navigation apps

When using an app like TomTom or Waze the driver can create an incident notification. When doing so the chosen type of accident combined with GPS coordinates are sent to the service provider. This data can be shared with the TMC (Waze data is publicly available). Unlike with a phone call, there is no opportunity to collect additional information about the incident through dialogue. On the other hand, the reliability of the reported incident is enhanced when multiple reports are created in the same area and time interval or other passers-by confirm a previous report, and this fusion can be automated rather than requiring additional validation workload.

Manual eCall

The use of eCall is only recommended when a driver is stopped in a live traffic lane and requires emergency services (National Highways, 2022).

Manual C-ITS/SRTI trigger

While this channel is theoretically possible, it seems not commonly provided or used.

The technology-based reporting methods that involve data transfer provide more precise location, at least as coordinates, but as has been observed in a previous chapter one pair of coordinates on their own are not always sufficient to correctly match the location to the network.

6.2 Human behaviour experiments

Within the SHADAR workpackage “Road user behaviour”, the reaction and behaviour of drivers in situations with stopped vehicles on the road was explored using virtual reality simulations. At the time of writing, the first batch of simulation experiments was complete. Regarding gathering traffic information, 68% of the participants use multiple sources during the journey of which 46% use navigation and 46% use a smartphone. When encountering an incident ahead (and drivers were informed about it) most participants adapt their behaviour and 27% even consider leaving the motorway if possible. Additional information is appreciated and 58% would like to receive instructions. Behaviour for alerting police or road authorities by the participants was not uniform. Some participants would alert police or road authorities, some would not, others would first assess the situation or stop to make a call. Most of the participants did not know which number to call or which organisation to alert. Mostly participants would alert the police by calling them. Furthermore, the test persons would only call if there was a passenger in the car who could make the call or if they would have a hands-free device available to them. Reporting the incident via an app was also considered by some test persons. It was also assumed that the incident was already reported (smart highway, cameras) or that nowadays everybody has a mobile phone and therefore the driver would have already called for help himself.

One of the most eye-catching phenomena is that the driver's awareness of how to act when a stopped vehicle/incident is encountered seems low and diverse. This also seems not or only marginally addressed during driver training for a driver’s license. It has been suggested that advice for safely encountering and reporting incidents should be incorporated in driving lessons, including the use of the eCall facility and in the near future the C-ITS solutions. NRAs, governments or traffic safety organisations can contribute by making the appropriate information publicly available and launch awareness campaigns on using eCall and incident reporting.

Human reporting and its influence on efficient road operator response is also considered in the forthcoming SHADAR report D6.1, and road user behaviour when encountering a stopped vehicle is the main focus of the forthcoming SHADAR report D4.1.

6.3 Findings

Table 19 and Figure 57 show the corresponding How-Now-Wow matrix.

	Subject	Category	Description
NOW	Driver awareness campaign	Reliability	Driver awareness of reporting including eCall (campaign, part of driving lessons)

Table 19: How-Now-Wow categorization - human behaviour

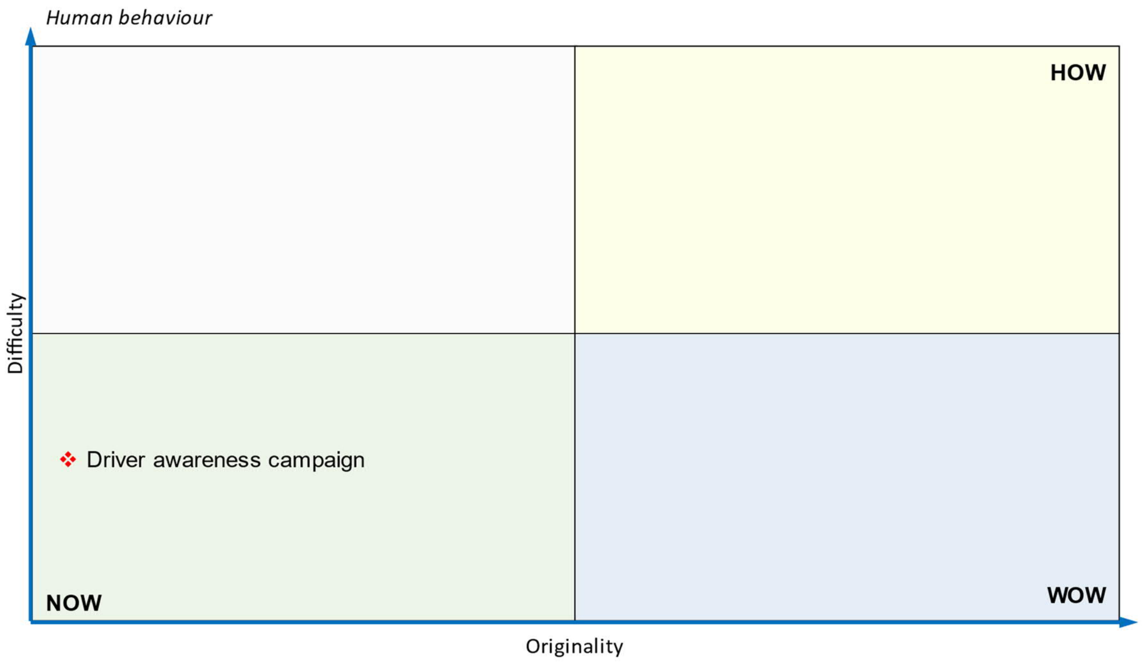

Figure 57; How-Now-Wow matrix human behaviour

7 Reporting alerts and performance

7.1 Introduction

This chapter considers how operational user interfaces might support the kinds of features explored in previous chapters, such as multi-sensor data fusion. It also considers how dashboards could provide insights into the performance of technologies to allow managers to understand what they have been getting for their investments in detection sources and in data fusion.

7.2 Storyboards

To provide a context for the user interface mock-ups illustrated in the next section, a set of storyboards have been elaborated. For practical reasons these are provided as a separate file. There are six different scenarios presented. Each uses a different combination of stopped vehicle event type and detection technology coverage. In each case it is hypothesised that a data fusion system receives alerts from individual detection technologies, including eCall and Waze, and manages the display of fused alerts to operators.

As well as providing context for the user interface mock-ups in the next section, the storyboards provide discussion material for the concurrent SHADAR workpackage on operational response improvement.

A brief summary of the six scenarios follows:

1. Rural road, no roadside tech, accident occurs, multiple low-level reports leading to high confidence report
    a. Manual eCall activation
    b. Waze activation (type hazard)
    c. Waze activation (type accident minor)
    d. Reports are fused to indicate a higher likelihood event
2. Major road breakdown, radar and automatic eCall, better event details
    a. Radar activates, tells us the location and lane
    b. Automatic eCall (containing vehicle details) in same location
    c. Operator knows the location (from radar) and the vehicle details (from eCall), operators know which vehicle to look for, could initiate recovery more quickly
    d. eCall tells us the vehicle is electric, allowing the operator to inform responders
3. Major road multiple collision, no radar in area
    a. Traffic detectors (e.g. MIDAS) activate traffic slowdown (low likelihood, low impact event)
    b. Manual eCall received (now medium likelihood, low impact event)
    c. Multiple Waze events - hazard on the road, accident major (now high likelihood, high impact event)
4. Critical incident, multiple collisions
    a. Radar alert
    b. 3 automatic eCall activations (3 details of vehicles)
    c. Multiple manual eCall activations
    d. Multiple Waze activations
    e. Additional alerts (which could be many) could be suppressed in the area to avoid distracting operators
5. Breakdown at night, rural area, bad weather

   a. Waze report, accident (normally low likelihood, medium impact but raised to medium likelihood and high impact due to weather and night time)
6. Waze alert in radar area, no radar activation
   a. Traffic detectors (e.g. MIDAS) activate traffic slowdown (low likelihood, low impact event)
   b. Waze alert – accident (medium likelihood)
   c. Radar does not activate (which suggests more likely a false alarm)
   d. Investigation confirms an event

7.3 User interface representation

This section explains the individual features used in the full mock-ups of the next section.

It is not our intention to prescribe specific user interface design, as each country will have its own existing traffic management systems with its own icons and conventions, but rather to explore the kinds of new features that might be added to existing traffic management user interfaces to support data fusion and increased data integration.

Notifications

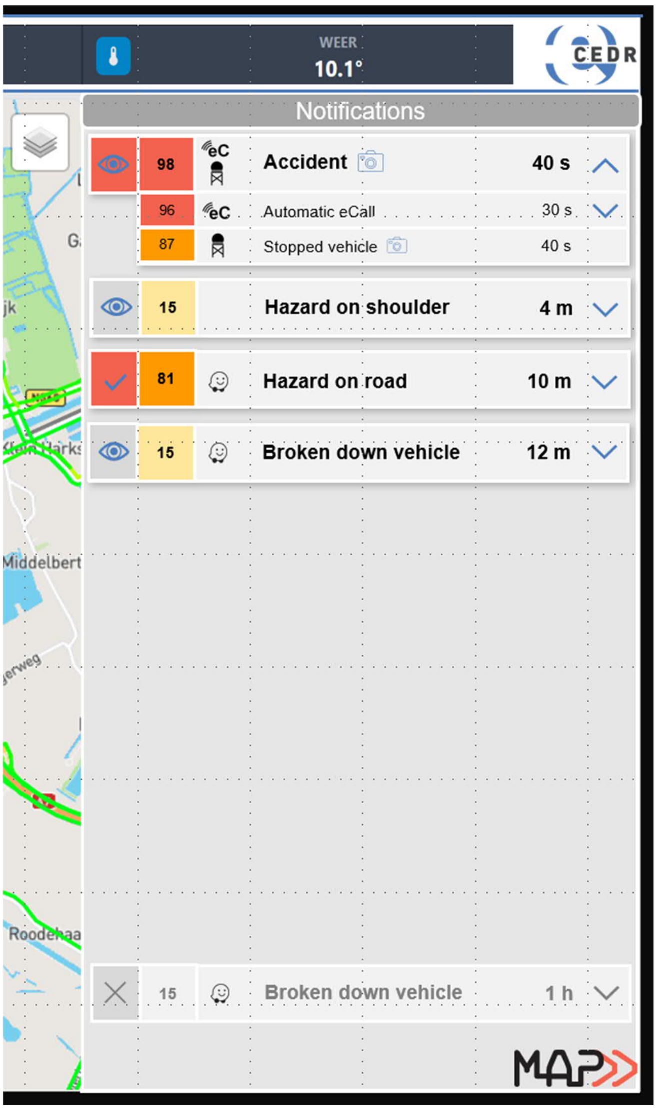

Figure 58 Notification pane

Figure 58 shows a notification pane with icons explained in Table 20. This is an example setup and can be adapted to the own incident/event statuses of a TMC.

Icon	Meaning	Description
Status field 1		
	Not addressed yet	Operator has not interacted with the notification
	Under surveillance of TMC	Operator is aware of the notification and is taking action
	Being handled	Response process is underway
	Finished	Response process is finished
	Dismissed	Operator assessed the notification as irrelevant and is not taking action
Colour	Severity	low probable definite major high
	Changing status by dropdown on the icon	
Status field 2		
98	Reliability percentage	
Colour	Reliability class	NaN, 0-50%, 50-75%, 75-90%, 90-100%
	Calculated based on fusion	NaN= no values/empty
Sources		
	Waze	
eC	eCall	
	Radar	
	Emergency service	
	Phone call	
MTM	MIDAS/Motorway Traffic Management	
	C-ITS message	DENM, CAM
Available verification		
	Visual information available	
	Emergency service on scene	
	Emergency service en route	
Additional information		
101.1r	Hectometre location and roadside	Location details recognisable for operators with lane discrimination up to lane-specific information

Table 20: Icons notification bar

Iconographic

Icon	Meaning
	Weather conditions at higher altitudes
	Weather conditions at lower altitudes
	Sunny weather conditions
	Cloudy weather conditions
	Showers
	Thunderstorms
	Fog
	Snow
	Freezing rain
	Heavy winds
	Incident Management alerts
	Waze notifications
	Weather
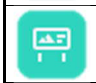	Variable message signs
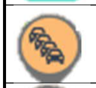	Traffic jam
	Closed lane
	Accident
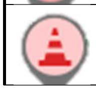	Vehicle on roadway or shoulder

Table 21: Icons header and map

7.4 User interface mock ups

The following mocked screen shots illustrate screens that could occur during management of the scenarios elaborated in the storyboards, or similar.

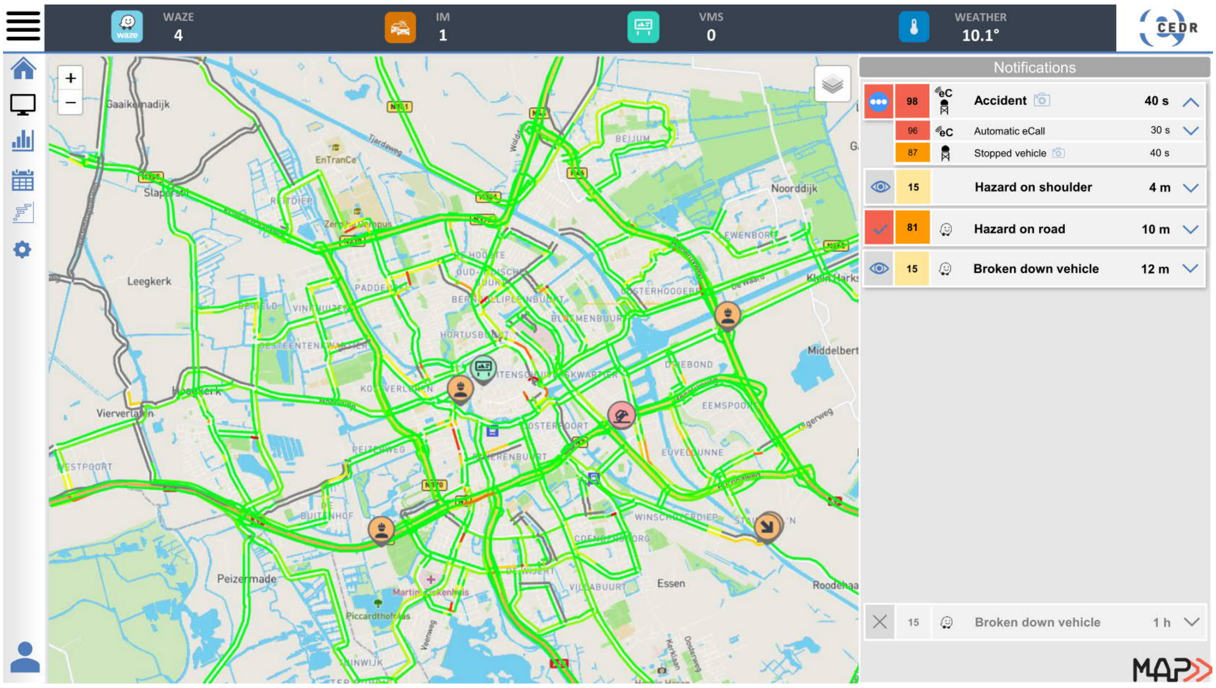

Figure 59: Mockup, major accident eCALL, Radar

It is assumed that a fusion system has received the alerts from multiple sources, and where they cover the same location it has fused them into a single composite notification. Information about each original source alert is still present, but it is grouped, and the group could be shrunk or expanded as required. Confidence in each alert is presented, here by colour and also by a numerical value which could be a probability calculated as described in Chapter 4.

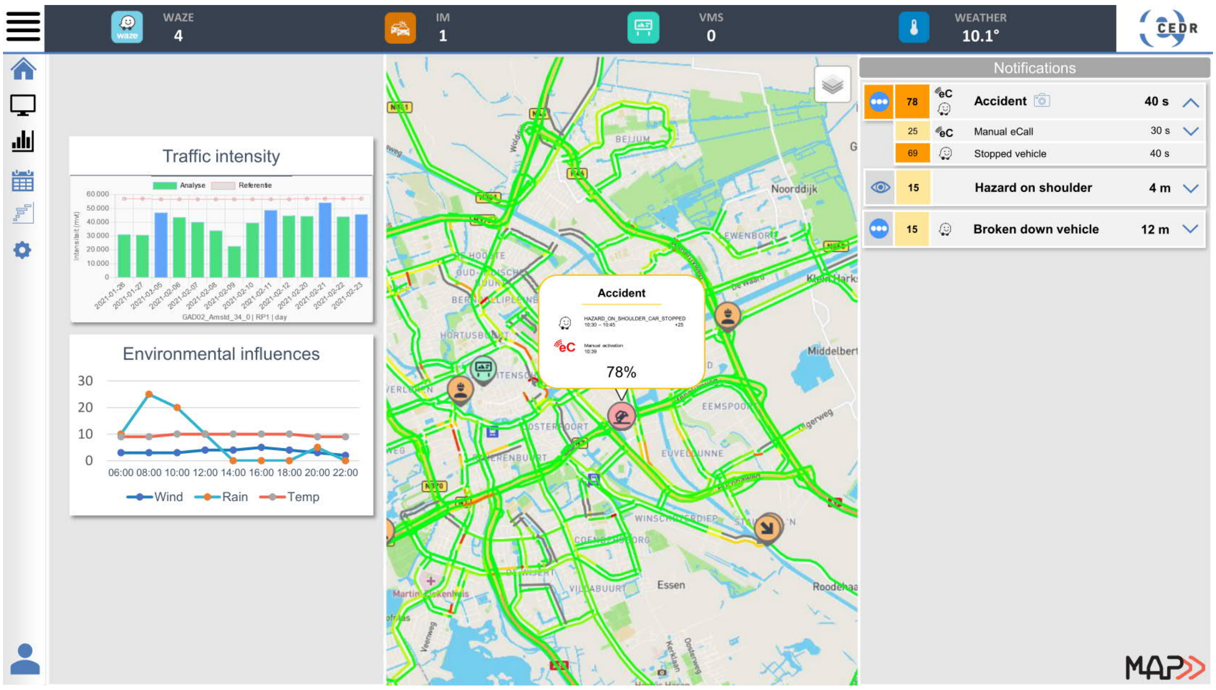

Figure 60: Mockup, side menu with traffic intensity and Environmental influences

In addition to the alerts *Figure 60* shows contextual information including traffic and weather trends.

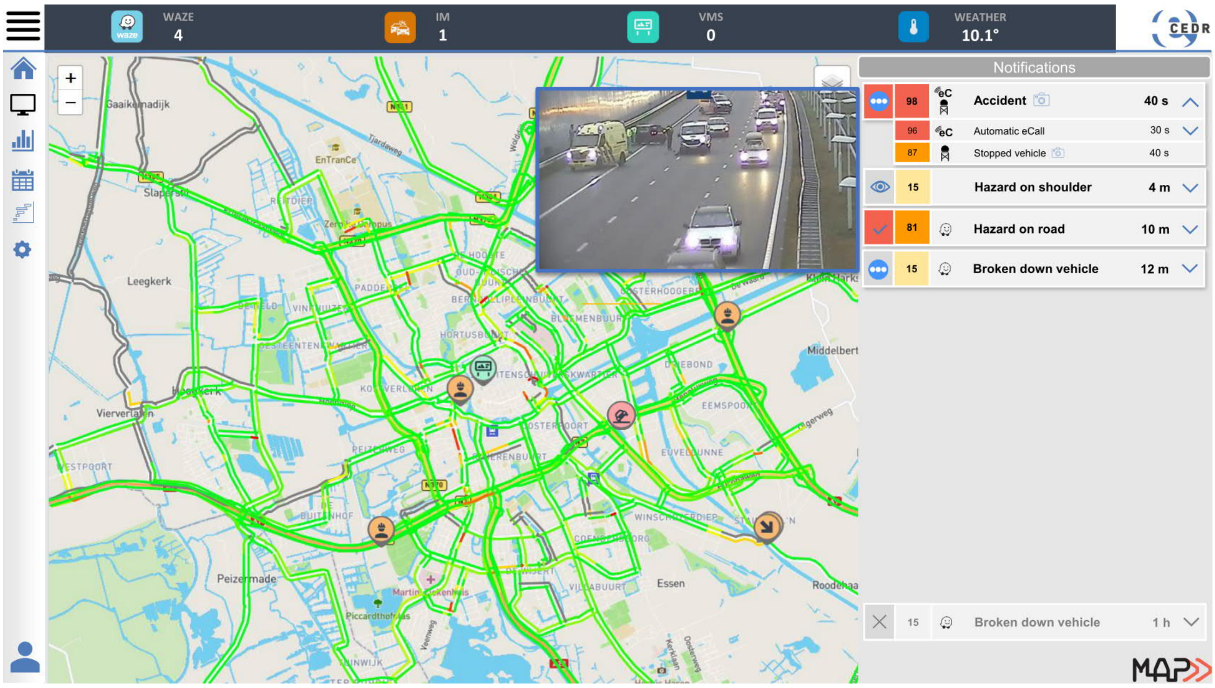

Figure 61: Mockup with camera footage

Efficient integration of the relevant CCTV camera feed promises to save time in validation. *Figure 61* illustrates the appropriate camera feed being shown in response to selection of the corresponding alert.

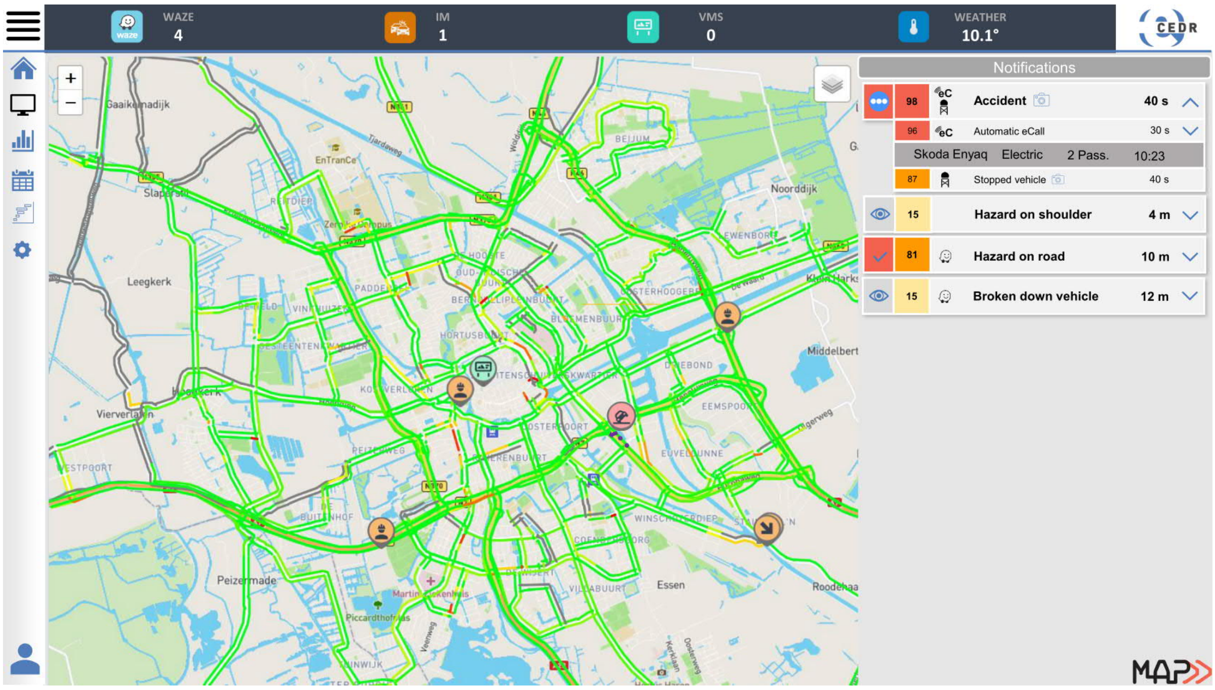

Figure 62: Mockup drill down with eCall data available including trace in geo area

eCall data includes vehicle identification which can support a look-up to obtain vehicle type information, which has been shown in *Figure 62*. eCall also includes the last 3 positions of the vehicle, which could be used to achieve more reliable network matching; here they are also plotted directly (dark dots).

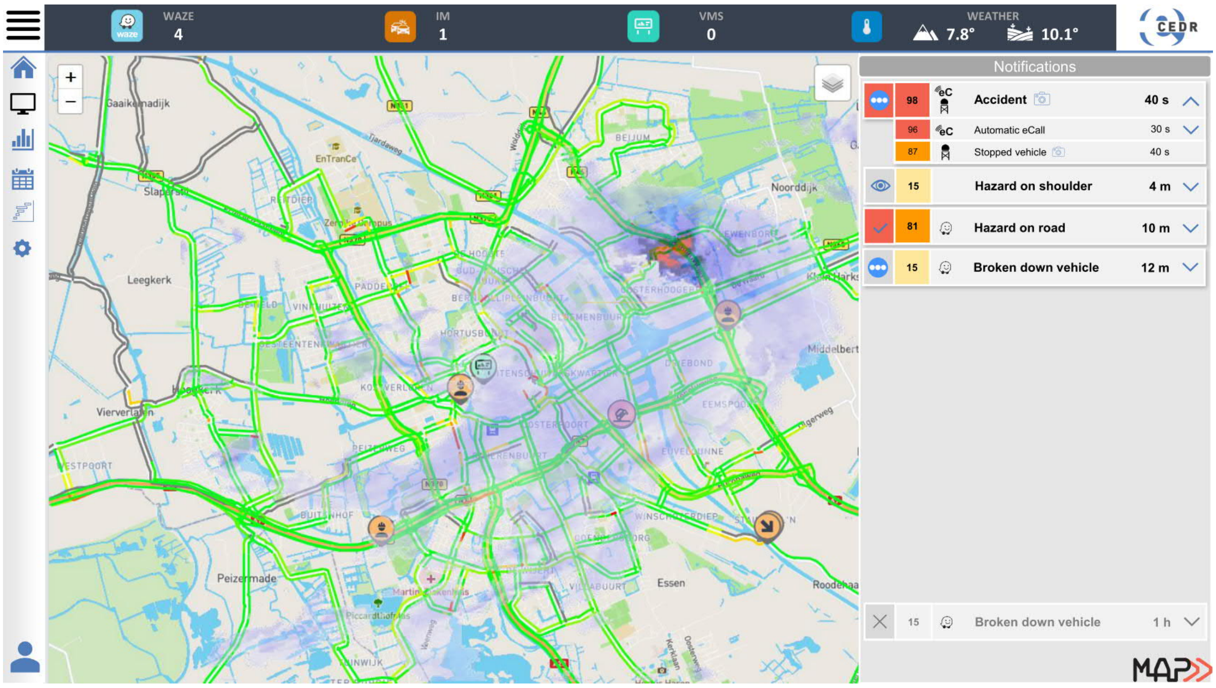

Figure 63: Mockup, weather info in geographical representation

One possibility for presentation of relevant weather information is shown in Figure 63 where rain radar is superimposed on the network map.

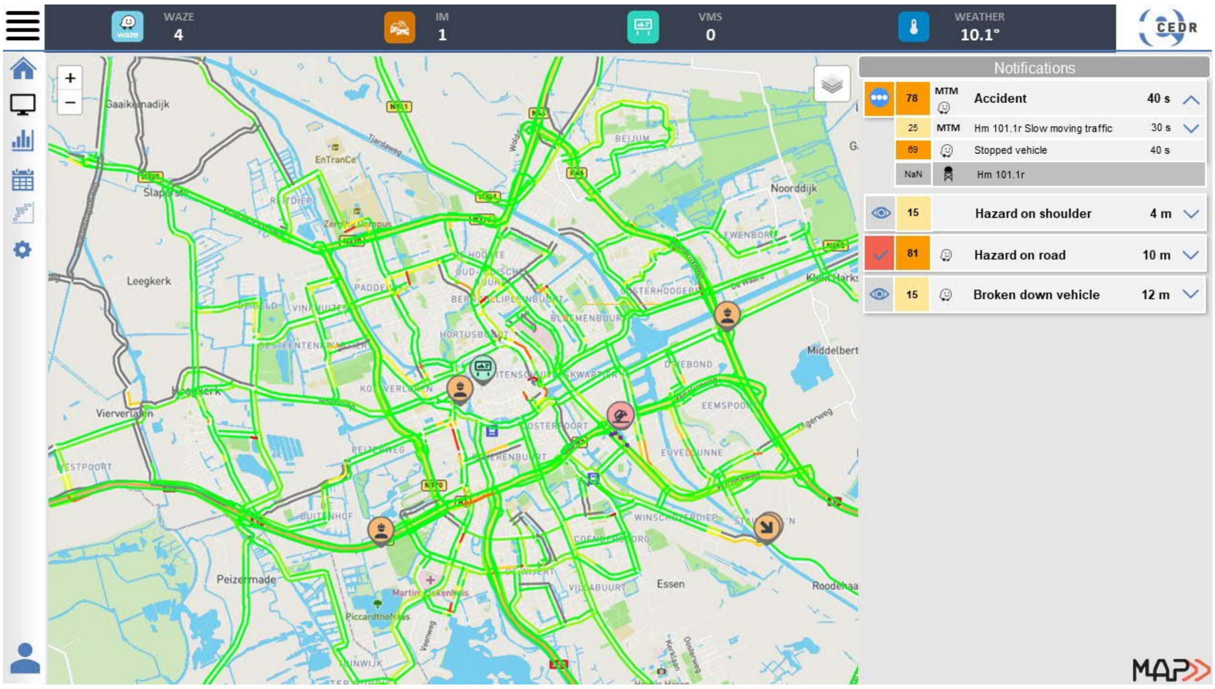

Figure 64: Mockup, registration of accident with no conformation from radar source

When detection infrastructure covers a location but has not detected an alert raised by another source, that is useful information that might be shown together with the alert, because it increases the possibility that the alert is a false alarm.

The illustrated features convey information that is potentially useful, but they might also be found difficult to understand or use and may be judged not to support the most efficient workflow for operators. The topic of whether such features could support efficient operational response is further explored in the forthcoming SHADAR report D6.1.

7.5 Use of existing traffic management interfaces

Even without additional enhancement of user interfaces, existing traffic management systems may already include useful concepts that can be adapted to convey the kinds of information discussed in this report, such as the priority of an alert. For example in the UK the CHARM active traffic management system supports an interface to receive alerts and present them to operators with an indication of priority, as in Figure 65.

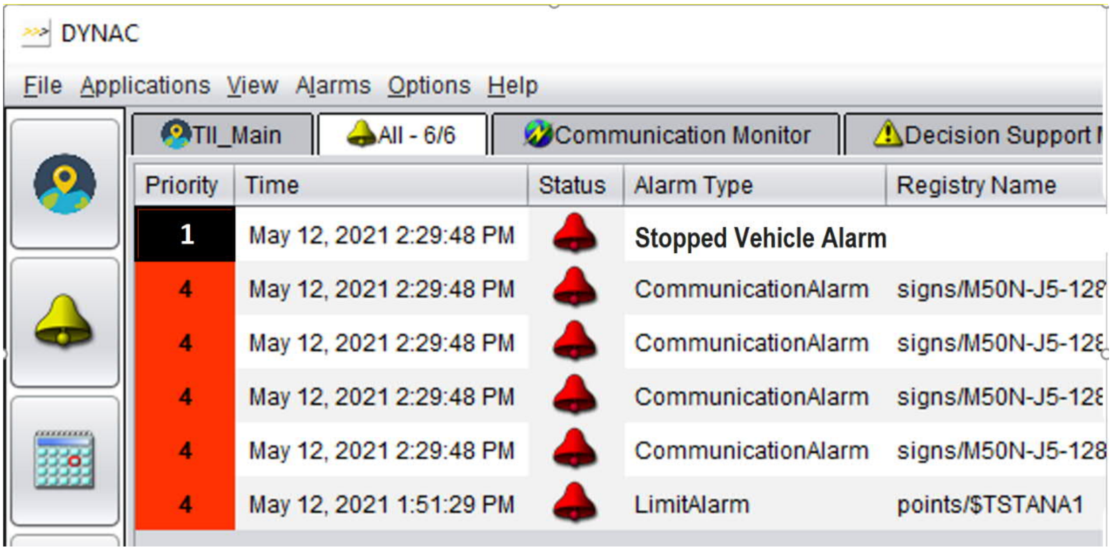

Figure 65 Mock-up of an SVD alert presented in the CHARM alarm banner[5]

Driver information systems may also benefit from receiving SVD alerts. For example, there are systems emerging that can present a virtual message sign in vehicle satellite navigation systems. SVD alerts could generate these virtual messages and pass them to in-vehicle information devices.

7.6 Statistics dashboard

In a regime with several kinds of stopped vehicle detection sources and data fusion, a technology manager may want to see how each source is performing, using simple reports that could be reviewed periodically or on demand. This may be especially useful for connected vehicle sources whose impact may grow or shrink over time as technologies and/or brands grow or shrink in popularity.

A technology manager may want to jump straight to statistics, but if their remit is for a wide area then they may wish to start with a map view in which a region or stretch could be selected.

[5] With permission of Kapsch TrafficCom

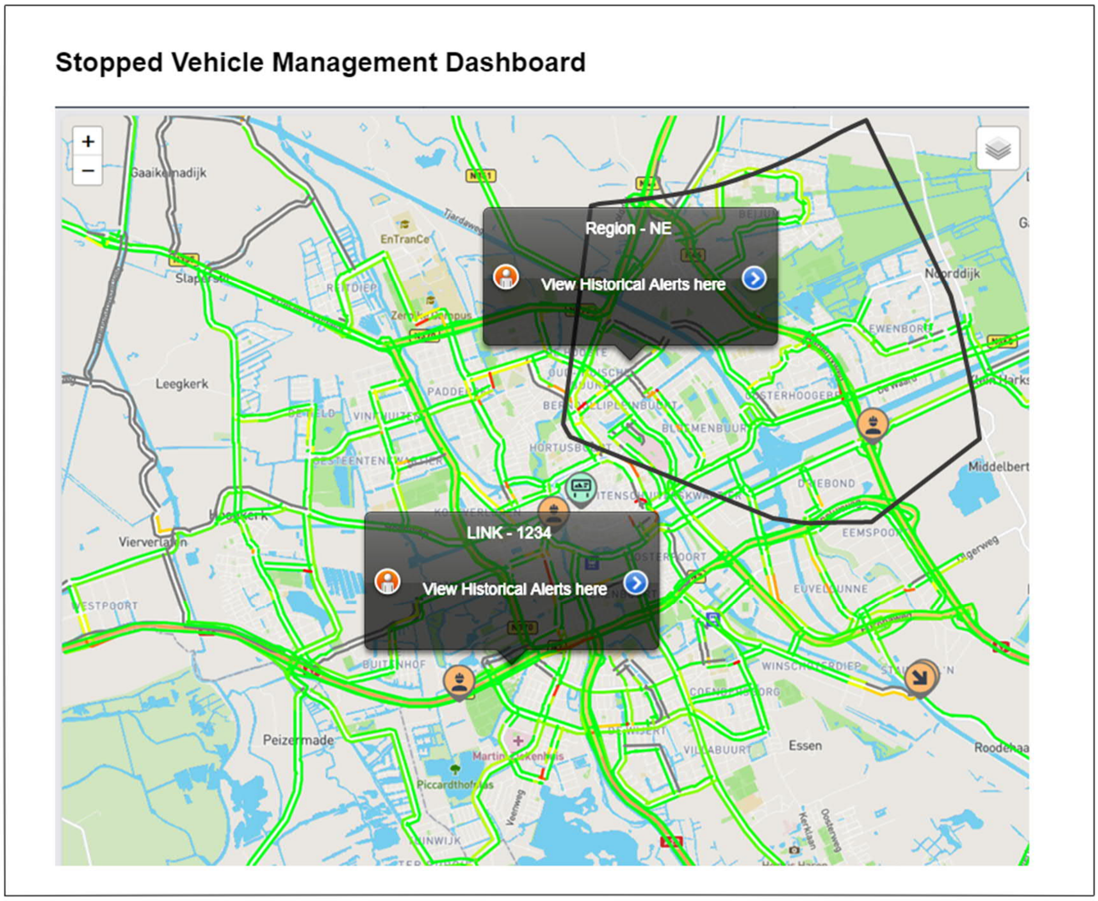

Figure 66: Stopped vehicle management dashboard – entry point by location

Once a region of interest is selected, performance statistics can be identified.

The statistics that can be derived depend on whether ground truth data is available.

Normally there will be no source of complete ground truth data because that needs special effort to collect. Complete ground truth data would definitively confirm whether a vehicle was stopped at any location and time. Human verification of alerts does not constitute complete ground truth because it is not known whether any stopped vehicle events went totally undetected.

Without ground truth, the detection rate cannot be computed definitively.

However, useful statistics can still be computed. Assuming each alert is investigated by a human operator and its status as a confirmed stop or apparent false alarm recorded in the incident management system, this data could be used to present comparative performance of different sources as shown in the following mockup.

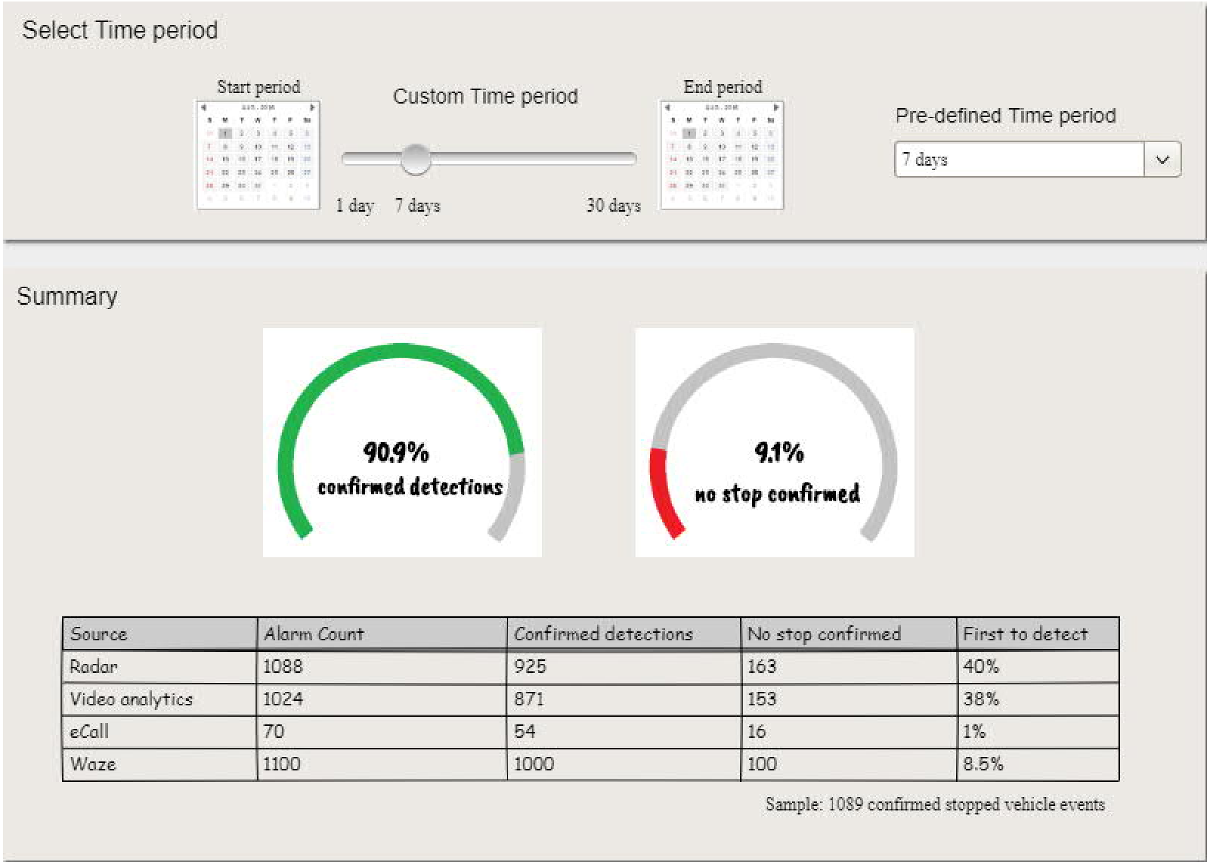

Source	Alarm Count	Confirmed detections	No stop confirmed	First to detect
Radar	1088	925	163	40%
Video analytics	1024	871	153	38%
eCall	70	54	16	1%
Waze	1100	1000	100	8.5%

Sample: 1089 confirmed stopped vehicle events

Figure 67: Compared preformance of sources

Using assumptions about location and time, it is possible to classify alerts from different sources as representing the same stopped vehicle event, and therefore derive statistics on the timeliness of each source, for example identifying which source was the first to detect each event, as shown above.

If ground truth data is available (which might be true in a limited study period for example), the detection rate and false alarm rate statistics can be derived. Performance could be reported in a dashboard screen such as the one below:

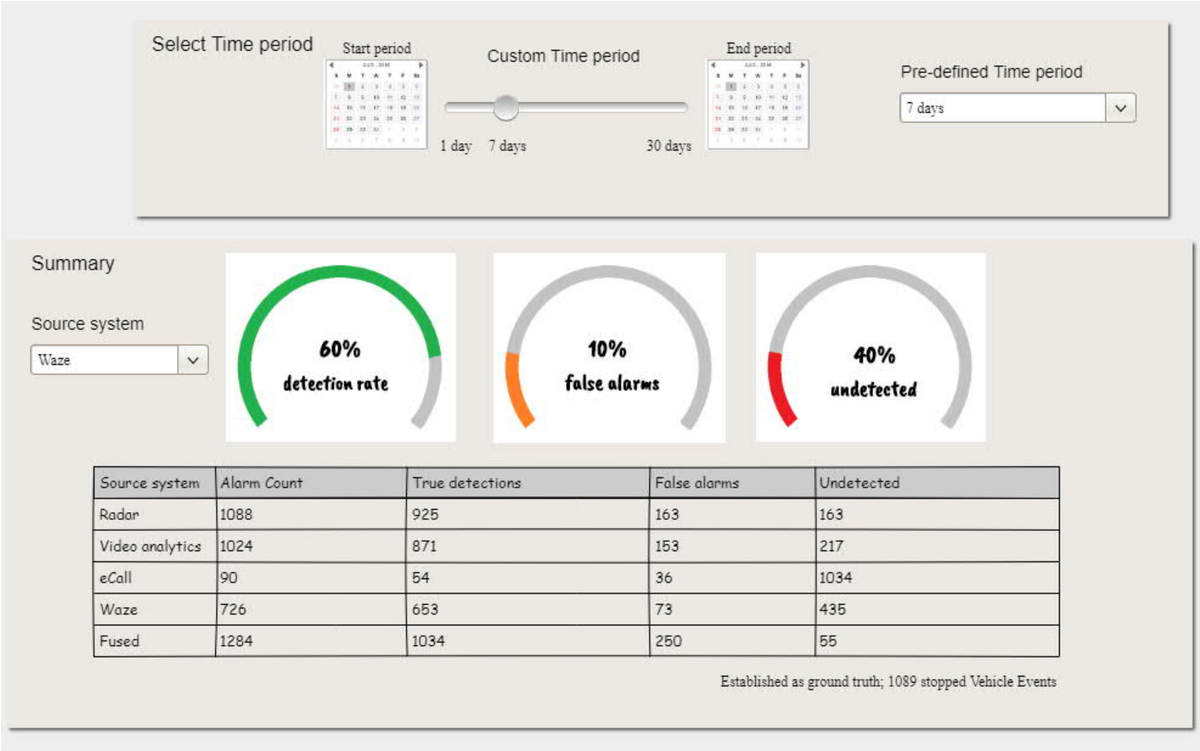

Source system	Alarm Count	True detections	False alarms	Undetected
Radar	1088	925	163	163
Video analytics	1024	871	153	217
eCall	90	54	36	1034
Waze	726	653	73	435
Fused	1284	1034	250	55

Figure 68: Additional statistics calculable when ground truth data is available

With ground truth it would also be possible to calculate and display mean-time-to-detect, and the number of stopped vehicle events detected only by one source (“unique detections” in the following screen).

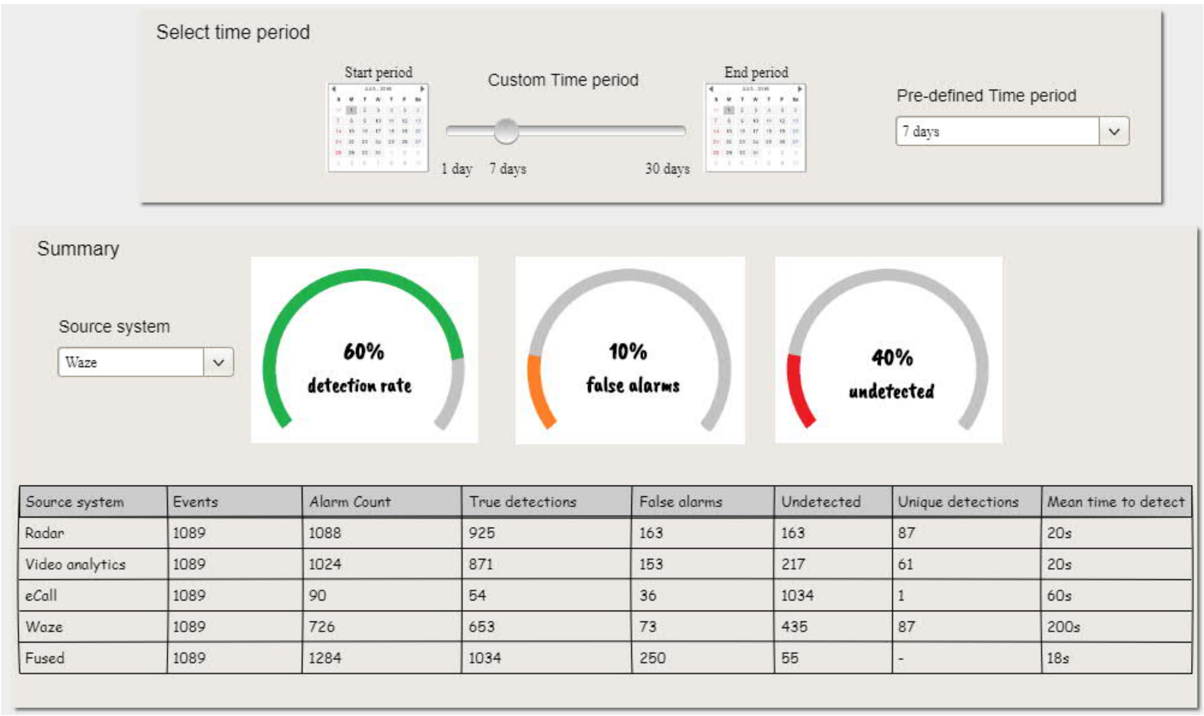

Source system	Events	Alarm Count	True detections	False alarms	Undetected	Unique detections	Mean time to detect
Radar	1089	1088	925	163	163	87	20s
Video analytics	1089	1024	871	153	217	61	20s
eCall	1089	90	54	36	1034	1	60s
Waze	1089	726	653	73	435	87	200s
Fused	1089	1284	1034	250	55	-	18s

Figure 69: Further additional statistics calculable when ground truth data is available

Although a metric for operational performance rather than technology performance, a manager might also want to know operational response times (e.g. time between the event or the alert and a confirmed operational response such as sign-setting), in comparison to a target figure, as shown:

Figure 70: Mockup operational response times

Data could be viewed by region:

Figure 71: Mockup hotspots per region

The performance could be presented in an alternative way, such as pie chart, for a selected source which might the fusion system overall (as in the following figure) or an individual source.

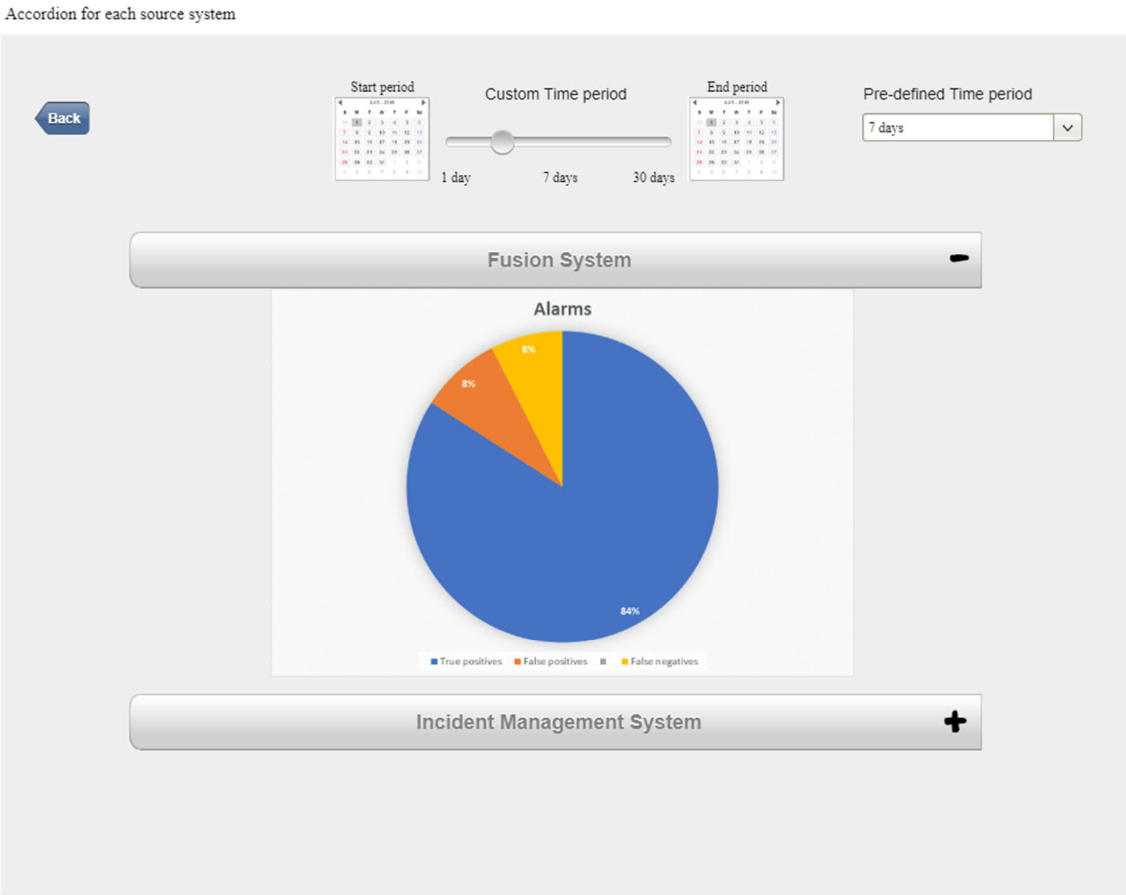

Figure 72: Mockup alternative representation

7.6.1 Reactions of road authorities

Representatives of national road authorities gave the following opinions in interviews:

Purposes and usefulness

- A performance management dashboard was generally considered helpful for gaining insights to support road management and response.
- The comparison of the performance of different technologies in the same terms was generally considered useful.
- Two potential purposes are optimisation of existing technology and informing new or continued investment decisions.
- Some NRAs would use these annually or 6-monthly.
- Some NRAs would want to see trends over time (e.g. sets of monthly changes) to support NRAs in being able to understand current and future trends and help support road management strategy.
- Changes in the performance of data sources might signal a need for improvement, not only in detection but perhaps also in verification processes.
- Performance data informs confidence in the data sources.
- There is a significant distinction between sources from infrastructure of the NRA and third party external sources - the former can be optimised by the NRA.
- The statistics allow the purchaser to give concrete feedback or requests to improve to the technology providers.

Choice of metrics

- Seeing non-detections and possible false alarms by specific technologies is interesting and could be used for improvement.
- Limits to what can be presented without ground truth data are important to understand.
- Ground truth data (and the richer statistics that it supports) is valuable – especially when a technology is first introduced.
- The statistics available when ground-truth data is available are more useful than those available without.
- Of statistics computable without ground-truth data, the number of times that a source is first-to-detect and number of detections unique to a source seem particularly useful.
- Seeing incident response performance time statistics could be useful for performance improvement (this is already done by some NRAs).
- Seeing statistics for specific locations is considered likely to be useful for multiple purposes – resource planning, identifying new or growing hot spots, identifying gaps or problem locations requiring optimisation, calibration, or troubleshooting.

These findings inform requirements for any such reporting developments by NRAs.

7.7 Findings

Table 22 and Figure 73 show the corresponding How-Now-Wow matrix.

	Subject	Category	Description
NOW	Alert source combination	Timely	Avoids separate operator investigation of separate related alerts 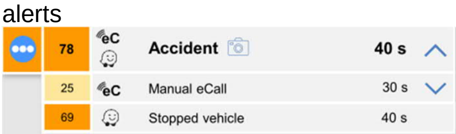
NOW	Integrated weather presentation	Information	Weather notifications per environment
NOW	Comparative technology performance reporting	Information	Potential bonus of a data fusion system – enables comparative reporting of technology to give new insight for investment decision-making.
WOW	**Subject**	**Category**	**Description**
WOW	Confidence/priority levels indicated	Timely/Information	Operators may prioritise high-confidence alerts.

Table 22: Now-Wow-How reporting alerts and performance

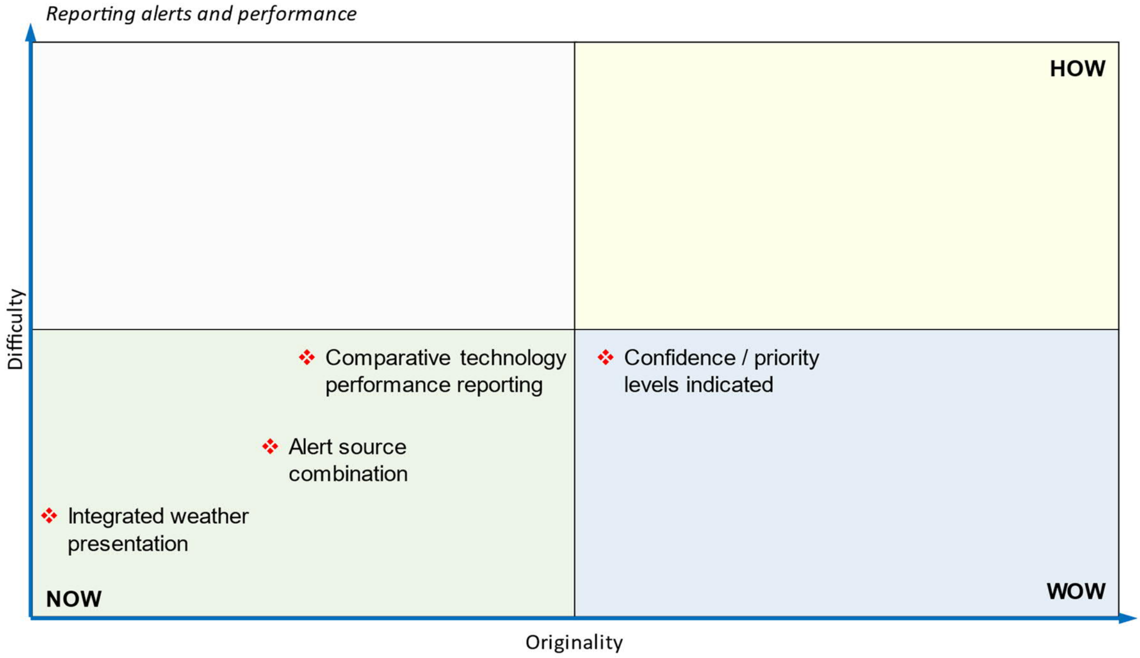

Figure 73 Now-Wow-How matrix reporting alerts and performance

8 Conclusion

This research has explored a number of ideas which each represent a further opportunity for investment by a NRA. In this concluding section they are categorized by the class of benefit and placed in a How-Now-Wow table as explained in the introduction.

The (number) in parentheses in the subject column corresponds with the following subjects and more details can be found in the sections with the same number: (2) eCall, (3) Radar, (4) Upcoming methods, (5) Fusion, (6) Human behaviour, (7) Reporting

Ideas in the “NOW” table have been assessed as *relatively* feasible to implement. Their originality is relatively low.

NOW	Subject	Category	Description
	(2) eCall Voice	Timely	Today’s baseline in some countries
	(2) Link PSAPs to NRAs	Timely	Reduces delays with SVD going to wrong control centre
	(2) Educate road users in eCall	Reliability	Reduce false alarms
		Timely	Drivers more likely to use it when needed
	(2) Optimise call handler processes, scripts and training	Accuracy	Get the right information from the right source
		Timely	Reduces delays in getting the right information
		Information	Get the right information at the right time
	(3) Lane definition	Accuracy	More accurate location information across the width of a carriageway/road
		Information	Retrieving vehicle lane positions helps the TMC respond to an incident more effectively, and can better plan emergency responses and recovery
	(3) Vehicle classification	Information	Retrieving vehicle classification helps the TMC to respond to an incident more effectively, and can better plan emergency responses and recovery
	(3) Pedestrian information	Information	Retrieving pedestrian information after the SVD event helps the TMC respond to an incident more effectively, and can better plan emergency responses and recovery
	(4) Waze / commercial traffic information	Timely	Extension of accuracy by alerting and fusion in TMC
		Reliable	Categorisation and trustworthiness Waze feed as extra source
	(4) Harvest Data for Road Safety	Timely	High road coverage, low vehicle coverage, but growing.
	(4) C-ITS safety-related	Timely, reliable, information	Potential high quality data direct from vehicle sensors. However, coverage still at R&D levels.
	(5) Harness	Accuracy	Some sources provide more accurate location

	multiple detection sources	Timely	Get alerts as fast as the fastest source in each case – but at cost of increased operator workload
		Information	Connected vehicle sources, potentially enhanced by lookups, can provide extra information e.g. vehicle type.
	(5) Fusing alerts	Reliability	In addition to benefits of using multiple sources, fusion can increase detection rate and decrease false alarm rate, potentially without significant impact on operator workload.
	(5) Multi-source ground truth study	Information, Accuracy, Reliability	A ground-truth study provides knowledge of relative and absolute performance of detection sources, allows more accurate assessment of confidence for future alerts, and therefore more reliable decisions about how to prioritise.
	(6) Driver awareness campaign	Reliability	Driver awareness of reporting including eCall (campaign, part of driving lessons)
	(7) Alert source combination	Timely	Avoids separate operator investigation of separate related alerts 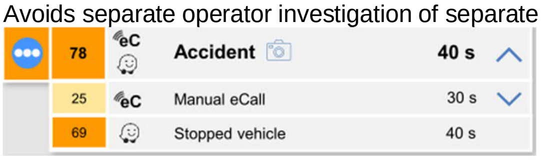
	(7) Integrated weather presentation	Information	Weather notifications per environment
	(7) Comparative technology performance reporting	Information	Potential bonus of a data fusion system – enables comparative reporting of technology to give new insight for investment decision-making.

Ideas in the WOW table have also been assessed as *relatively* feasible to implement and are more original or innovative.

WOW	Subject	Category	Description
	(2) Automatic eCall data processing	Accuracy	Highly accurate
		Reliability	Very high SVD indication
		Timely	Very fast (<1 min)
		Information	Provides details of vehicle and location
	(2) Manual eCall data processing	Accuracy	Highly accurate
		Reliability	Medium SVD indication
		Timely	Very fast (<1 min)
		Information	Provides details of vehicle and location
	(2) Automatic and Manual eCall data fusion	Accuracy	Highly accurate
		Reliability	Very high SVD indication, with greater coverage than just auto eCall
		Timely	Very fast (<1 min)
		Information	Provides details of multiple vehicles and locations

	(3) Impact levels from additional radar info	Information	More information about the impact of the alert, and thus a higher priority can be assigned. This information could influence operator response.
	(4) Platform for easy function extension roadside (iWKS)	General	Seamless function enabler on roadside and central applications using sensors and actuators. Despite appearing in the WOW category, by itself this idea does not deliver value for SVD, it must be combined with other advances.
	(5) Probabilistic fusion influencing priority	Timely	Alerting user interface gives higher priority to alerts with higher confidence and can help save operator time.
	(7) Confidence/priority levels indicated	Timely / Information	Operators may prioritise high-confidence alerts.

Ideas in the HOW category are considered more difficult to implement, but they have higher originality, which may indicate that they are worth further research.

HOW	Subject	Category	Description
	(2) eCall and bCall data fusion	Timely	Faster responses due to greater SVD coverage
		Information	Much richer data for SVD
	(3) Correlation of traffic parameters	Information	Using data to describe traffic behaviour before, during and after an SVD event can help provide useful insight into the impact of traffic parameters on an alert, which can ultimately feed into AI models for improved/predictive detection
	(3) Confidence and probability levels from additional radar info	Reliability Accuracy Information	Confidence and probability levels in alert could increase detection rate and reduce false alarm rates. More accurate reporting of the incident.
	(4) C-ITS extension	Timely	Extend the use of CPM/DENM/ to warn based on detection SV and other CVs
	(4) UAV/satellites	Accuracy	High accurate info of SVD and location but a future use. Drones can be deployed to view the incident site
	(5) Machine learning on raw data	Accuracy, Reliability	Machine learning fusion of raw sensor data seems likely to provide good accuracy and reliability, but seems less easily feasible than fusing the current outputs from technology providers.

Figure 74 combines all ideas in a single matrix (except for two pairs unified where the same concept is addressed under data fusion and reporting), but with the positions within each box changed slightly for legibility.

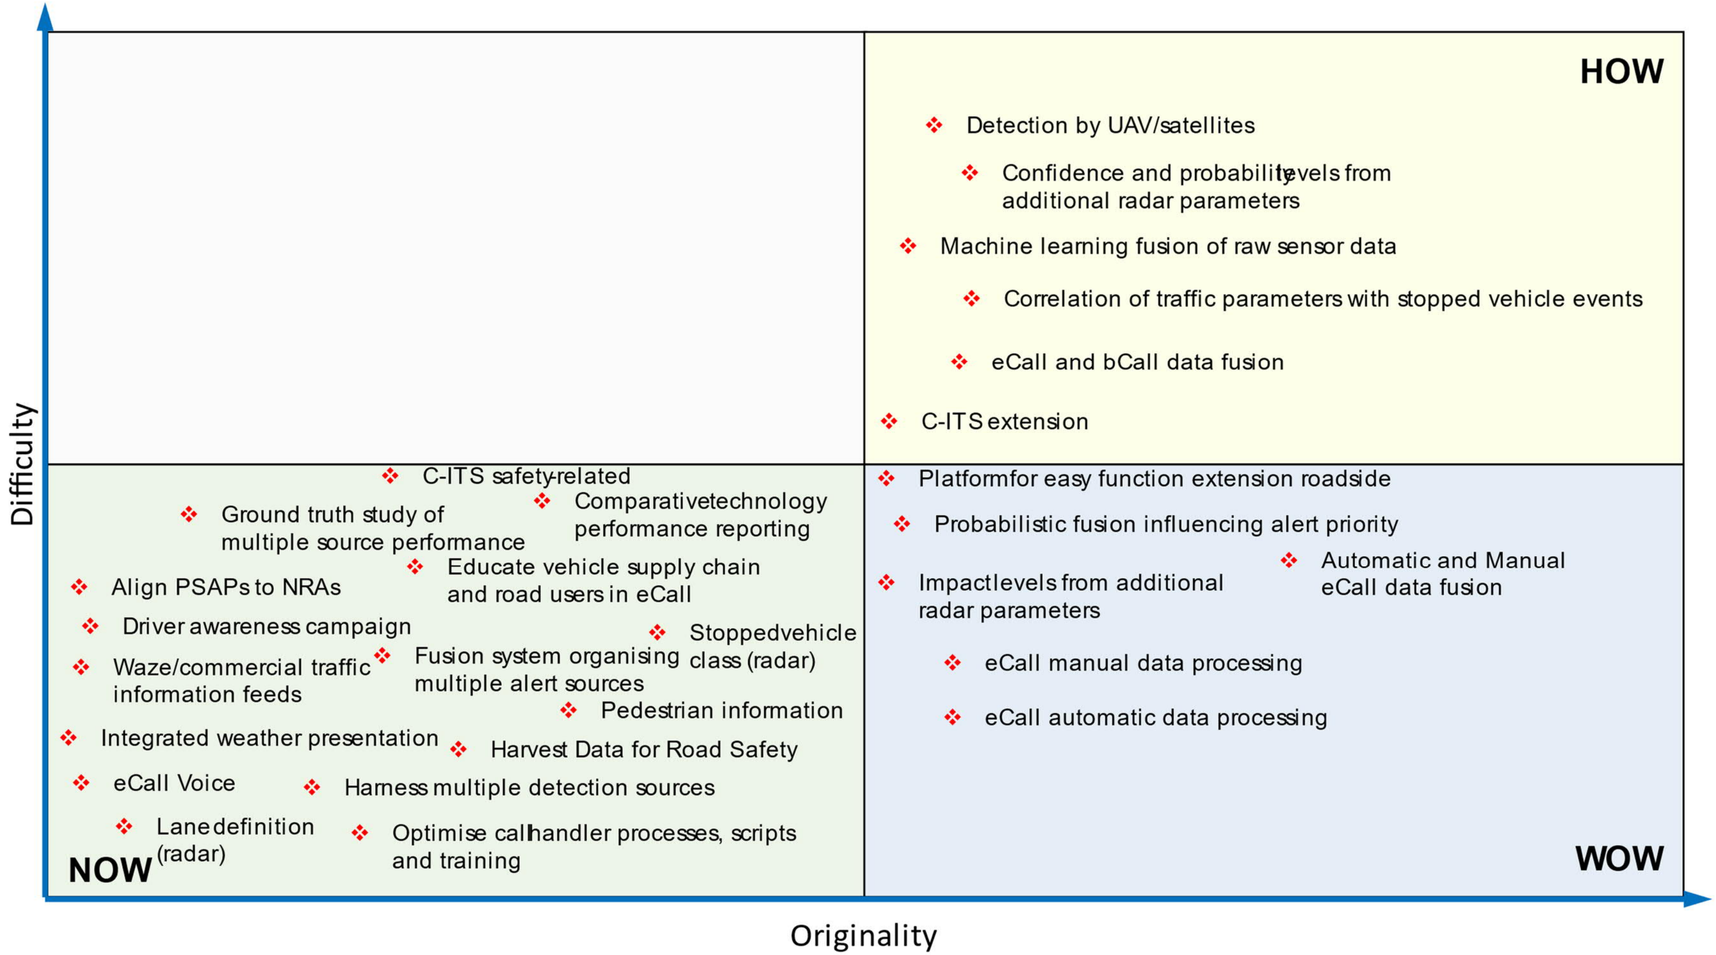

Figure 74 Combined How Now Wow matrix (positions adjusted for legibility)

9 References

Ahmadi, S. A. and Ghorbanian, A. (2019). Moving vehicle detection, tracking and traffic parameter estimation from a satellite video: a perspective on a smarter cit. *International Journal of Remote Sensing*, 8379-8394.

Barmpounakis, E., Sauvin, G. M. and Geroliminis, N. (2020). Lane Detection and Lane-Changing. *Transportation research record*, 1-15.

Be-Mobile. (2021). Be-Mobile. Be-Mobile - EN. Retrieved December 3, 2021, from https://be-mobile.com/developers/products/traffic#traffic-events

Car 2 Car Communication Consortium. (2021, January). Safety related message sets – Selection of DATEX II Situations, DENM and TPEG2-TEC Causes and TMC Events for EC high level Categories (ITSTF20001 v1.5).

CAR 2 CAR Communication Consortium. (2021, July). Guidance for day 2 and beyond roadmap (No. 2072, version 1.2). https://www.car-2-car.org/fileadmin/documents/General_Documents/C2CCC_WP_2072_RoadmapDay2AndBeyond_V1.2.pdf

Chen, X., Xiang, S., Liu, C.-L., & Pan, C.-H. (2014). Vehicle Detection in Satellite Images by Hybrid. *IEEE GEOSCIENCE AND REMOTE SENSING LETTERS*.

Cornwell, I., Wright, D., Elias D., Nadler, B., Nadler, F. (2016), *Fusion of third party and existing traffic data*, CEDR Call 2013 project "UNIETD" report, may be available from CEDR on request.

Correa, A. (2020). D5.2 V2X-based cooperative sensing and driving in Transition Areas (v2.0). TransAID. https://usercontent.one/wp/www.transaid.eu/wp-content/uploads/2017/Deliverables/WP5/TransAID_D5.2_V2X-based-cooperative-sensing-and-driving-in-Transition-Areas.pdf

C-Roads (2022), Implemented Services: C-Roads, https://www.c-roads.eu/pilots/implemented-services.html, retrieved July 2022.

C-Roads Working Group 2, Technical Aspects. (2020). Common C-ITS Service and Use Case Definitions (2.0.0 WG 3).

Cvetek, D., Muštra M., Jelušić, N., Tisljaric, L. (2021) "A Survey of Methods and Technologies for Congestion Estimation Based on Multisource Data Fusion", *Applied Sciences*, 11, 2306.

El Faouzi, N., Klein, L. (2016). "Data fusion for ITS: Techniques and Research Needs", *Transportation Research Procedia*, Volume 15, December 2016, Elsevier.

European Commission. (2013, May 15). EUR-Lex - 32013R0886 - EN - EUR-Lex. Eur-Lex.Europa.Eu. https://eur-lex.europa.eu/eli/reg_del/2013/886/oj

European GNSS Agency and Joint Research Centre (2019) GSA/JRC ECALL CONFORMANCE TESTING CAMPAIGN ECALL DEVICES ASSESSMENT REPORT, https://www.euspa.europa.eu/sites/default/files/content/gsa-jrc-ecall-conformance-testing-campaign.pdf

Grant-Muller, S., Cross, P., & Clarke, M. (2015). *UNIETD: Understanding New and Improving Existing Traffic Data: Assessing the relevance of Social Media Content* (Deliverable 3.1). CEDR; Conference of European Directors of Roads.

Gruber, D. A. (2021, June). *Annual pilot overview report 2020* (V1.0). C-Roads. https://www.c-roads.eu/fileadmin/user_upload/media/Dokumente/M16_Annual_pilot_overview_report_2020_v1.0.pdf

Guerrero-Ibanez, J., Zeadally, S., Contreras-Castillo, J., Hu, Y., & Wang, R. Q. (2020). Understanding the removal of precise geotagging in tweets. *Nature Human Behaviour*, 4(12), 1219-1221.

Huisken, G., Rodrigues, N., Vreeswijk, J., Boerma, S., Afshar, S., Baucutt, S., Cornwell, I., Zacharopoulos, G., Woolard, D., Rooke, A., Graham, A., Kaufmann, C. (2021) "Stopped vehicle detection and reporting – current methods", CEDR Call 2019 project report, may be available from CEDR on request.

Ibrar, M., Mi, J., Karim, S., Laghari, A. A., Shaikh, S. M., & Kumar, V. (2018). Improvement of Large-Vehicle Detection and Monitoring. 3D res 9.

Ismail, M. (2020). INCREASING ROAD SAFETY BY SHARING ROAD SAFETY RELATED DATA IN PUBLIC AND PRIVATE COOPERATION (Annex 1). Data Task Force. https://www.dataforroadsafety.eu/images/Documenten/Annex_1_-_Data_For_Road_Safety__Technical_Documentation_Version_1.01.pdf

ISO/CEN. (2020). Cooperative intelligent transport systems (C-ITS) Guidelines on the usage of standards. https://www.itsstandards.eu/app/uploads/sites/14/2020/10/C-ITS-Brochure-2020-FINAL.pdf

ISO 21217:2014, Intelligent transport systems — Communications access for land mobiles (CALM) — Architecture

Kernstock, W. (2019, April). Annual pilot overview report 2018 (V1.0). C-Roads. https://www.c-roads.eu/fileadmin/user_upload/media/Dokumente/Annual_pilot_overview_report_2018_v1.0.pdf

Khalil, M., Li, Jianping, Sharif, A., & Khan, J. (2017). Traffic congestion detection by use of satellites view. 2017 *14th International Computer Conference on Wavelet Active Media Technology and Information Processing*. Chengdu: IEEE.

Klein, L. (2019), *Sensor and Data Fusion for Intelligent Transport Systems*, SPIE Digital Library

Liu, G., & Xiao, F. (2019) Time Series Data Fusion Based on Evidence Theory and OWA Operator, *Sensors*, 19(5):1171

Lücken, L., & Schwamborn, M. (2021). D6.2 Assessment of Traffic Management Procedures in Transition Areas (V2.0). TransAID. https://usercontent.one/wp/www.transaid.eu/wp-content/uploads/2017/Deliverables/WP6/TransAID_D6.2_Assessment-of-Traffic-Management-Procedures-in-Transition-Areas.pdf

Mordor Intelligence (2020), Europe Connected Cars Market Size, Trends, Growth (2022 - 27), https://www.mordorintelligence.com/industry-reports/europe-connected-cars-market.

Najiya, K. V., & Archana, M. (2018). UAV Video Processing for Traffic Surveillence with Enhanced Vehicle Detection. *2018 Second International Conference on Inventive Communication and Computational Technologies (ICICCT)*. Coimbatore: IEEE.

National Highways (2022), "eCall SOS", https://nationalhighways.co.uk/road-safety/ecall/, retrieved July 2022.

Outay, F., Mengash, H. A., & Adnan, M. (2020). Applications of unmanned aerial vehicle (UAV) in road safety, traffic and highway infrastructure management: Recent advances and challenges. Transportation Research Part A: Policy and Practice, 116 - 129.

RAC (2022), https://www.racfoundation.org/motoring-faqs/mobility#a35, retrieved July 2022.

Rijkswaterstaat (2021), Intelligent Wegkantsysteem langs de weg, https://www.rijkswaterstaat.nl/nieuws/archief/2021/01/het-intelligent-wegkantsysteem-meer-dan-een-computerkastje-langs-de-weg, retrieved July 2022.

Road Safety Authority (2022), eCall on vehicles, https://www.rsa.ie/road-safety/road-users/car-drivers/resources/e-call-on-vehicles, retrieved July 2022.

Rojas, J. J. (2019). Using Traffic Data with Maps and Routes. Opgehaald van developer.tomtom: https://developer.tomtom.com/blog/build-different/using-traffic-data-maps-and-routes

Rondinone, M., & Correa, A. (2018). Definition of V2X message sets (D5.1, V2.0). TransAID.https://usercontent.one/wp/www.transaid.eu/wp-content/uploads/2017/Deliverables/WP5/TransAID_D5.1_V2X-message-sets.pdf

Sharma, G. (2002). Vehicle Detection and Classification in 1-M Resolution Imagery. Ohio State University: OhioLINK Electronic Theses and Dissertations Center. Opgehaald van http://rave.ohiolink.edu/etdc/view?acc_num=osu1420208314

Stuparu, D.-G., Ciobanu, R.-I., & Dobre, C. (2020). Vehicle Detection in Overhead Satellite Images Using a One-Stage Object Detection Model. Opgehaald van https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7696426/

TomTom. (2021). Pricing. TomTom Developer Portal. Retrieved December 3, 2021, from https://developer.tomtom.com/store/maps-api

van Rij, M. (2020). Evaluation of the PoC 'Data for Road Safety' (3/70822/ SWNL0266365). Data For Road Safety. https://www.dataforroadsafety.eu/images/Documenten/Microsoft_Word_-_PoC_DTF_-_monitor__evaluation_report_Sweco.pdf

Waze (2021). Changing the Way We Get — and Give — Roadside Help. Retrieved from https://medium.com/waze/sending-out-an-sos-with-roadside-help-4bc1804b6950

Wijbenga, A. (2020). D8.2 Meta-analysis of the results (v1.0). TransAID. https://usercontent.one/wp/www.transaid.eu/wp-content/uploads/2017/Deliverables/WP8/TransAID_D8.2_Results-of-the-meta-analysis.pdf

Wijbenga, A. (2021). D8.3 Guideline and Roadmap (D8.3). TransAID.

Xiao, J., Cheng, H., Sawhney, H., & Han, F. (2010). Vehicle detection and tracking in wide field-of-view aerial video. 2010 *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. San Francisco: IEEE.

Xu, Y., Yu, G., Wang, Y., Wu, X., & Ma, Y. (2016). A Hybrid Vehicle Detection Method Based on Viola-Jones and HOG + SVM from UAV Images. Sensors.

Xu, Y., Yu, G., Wang, Y., Wu, X., & Ma, Y. (2017). Car Detection from Low-Altitude UAV Imagery with the Faster R-CNN. *Journal of Advanced Transportation*.

Appendix: Original How-Now-Wow

At the outset of this research, before the work described in the body of this report had been conducted, the SHADAR project team held an initial brainstorm session to identify unmet needs for relevant stakeholders, then possible technical solutions were placed in a “How-Now-Wow” matrix (described in Chapter 1). That early “How Now Wow” matrix is not itself a significant research result, but it is included in this Appendix.

The primary purpose of the session was to inspire the project team and share thoughts in preparation for the research described in this report. It helped to revalidate the scope of the research and suggested significant potential in areas including eCall, connected vehicles, data fusion and increased integration of systems.

The matrix is shown in Figure 75 and in an easier-to-read tabular form in Table 23. These represent the results produced in the session, and therefore represent subjective opinions of participants which may not always have been fully moderated by the other participants. Not every idea could be developed by the project. The ideas are not further defined, but the topics considered most valuable and feasible for the SHADAR project to research can be found within the main body of this report.

<table>
<tr><td></td><td>HOW</td></tr>
<tr><td></td><td>Detection
Transition period CAV, CV
I2V guidance CAV along SVD location
Seamless detection of vulnerable periods
Prediction of risky traffic situations
eCall like providing of detection
C-ITS application hazard warning based on eCall like data
Location references standardised
V2I standardised communication on SVD (TransAID)
Characterisation to get the right type of response
C-ITS hazard warning between CAVs
Machine learning traffic behaviour
Reporting
Use of risk profiles of stopped vehicle occurrence
Feedback time to incident emergency vehicles
Standardised accurate location referencing
Space tailored strategy for national needs (different situations and systems)
Online real-time traffic model
Inform about vehicle in live lane (not only closed lane info)</td></tr>
<tr><td>NOW</td><td>WOW</td></tr>
<tr><td>Detection
Aggregation of eCall, auto and SOS
Trace multiple detections of one SVD
Incident type identification
Automatic focussing cameras in SVD location</td><td>Detection
QKZ model to characterise pref. requirement
Historical data analyses
Use confidence level of alert
Stopped vehicle location broadcast</td></tr>
</table>

• Technology handling a moved away stopped vehicle Reporting • Road weather condition information • Alarm dashboard provides comprehensive overview of alarms and false positives • Informing driver help underway • Interface descriptions • User management dashboard at network level • Easy to use dashboard • Web-based info for en-route incident planning	• Lane info from eCall and detection • Data fusion info from SVD, eCall, cameras • Data fusion to prioritise and characterise SVD also about weather etc. • Vehicle details for arranged recovery • Personnel and injury details for welfare Reporting • Published tech roadmap for integrations • Transition period CAV/CV's and non-connected vehicles • Integrating signage with SVD

Table 23: Categorised in-session content from initial How-Now-Wow brainstorm session

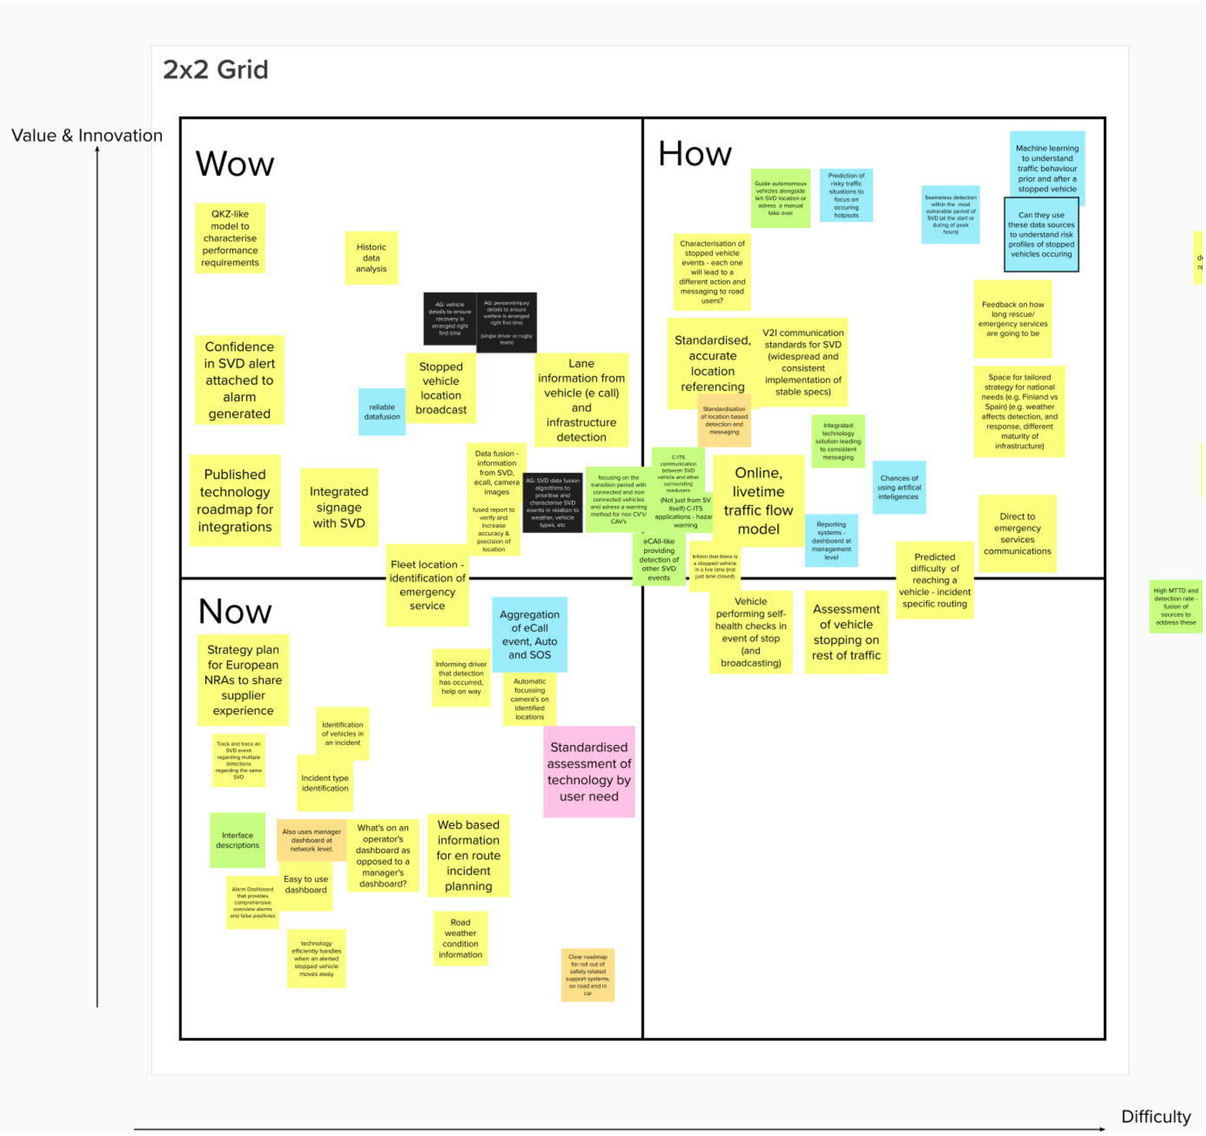

Figure 75 Original How Now Wow session content