Conférence Européenne des Directeurs des Routes

Conference of European Directors of Roads

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

Stopped vehicle detection and reporting

Summary of research results

Deliverable D5.2

31 October 2022

CEDR Call 2019: Safe Smart Highways

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

D5.2 Stopped vehicle detection and reporting: Summary of research results

Submission date: 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 summarises results 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, 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 alerts and eCall voice.

eCall complements existing detection methods to create a much wider coverage of the road network. So, if MSDs are processed together with other 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.

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.

SHADAR work package 5 (detection improvement research) builds on the reviews of state-of-the-art (Huisken et al, 2021) and considers potential improvements in stopped vehicle detection. This report is a summary of the research results. A longer report on the research, with additional explanation, examples, and description of relevant external research, is provided in the companion report D5.1 (approximately 150 pages of detail).

1.2 Report structure

Each chapter covers a specific area of potential improvement in stopped vehicle detection, and can be read individually.

- How to harvest data from eCall (Chapter 0)
- Potential of improved radar detection (Chapter 3)
- Future/upcoming methods for stopped vehicle detection (Chapter 4)
- Data fusion (Chapter 5)
- Reporting alerts and performance (Chapter 6)

Chapter 7 provides a consolidated summary of ideas.

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 stopped vehicle detection (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.

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 should only increase. In the UK, 90-95% of the sales of top 20 selling cars now have eCall, and this should 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.

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.

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

2.2.1 Advantage of voice eCall

eCall as a voice channel has advantages over a 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.

With added MSD (minimum set of of data) it has highly reliable data that can be filtered to help NRA actions. MSD from 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
- Accurate location to a few meters, rather than as drivers often describe their location “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. In most newer vehicles, airbag deployment immobilizes the vehicle so the NRA may want to know about the vehicle to recover it, especially if stopped in a live lane.

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

Challenges for eCall voice

Having a public safety answering point (PSAP) aligned with the NRA traffic operations is key.

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. 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. This is a strategic decision for PSAP to tell police, other emergency services, NRA and other road authorities.

In using a voice plus metadata 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).

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.

2.3 eCall delivery path and response times

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.

Our 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.

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 as long as there is emergency network coverage.

2.4 eCall MSD for SVD alerts

Every eCall activation generates a MSD which must confirm with EN 15722 and which is sent along with the voice channel.

2.4.1 eCall MSD fields

SHADAR report D2.1 identified the fields; the following table adds further definition.

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.
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	(set on installation) Class of vehicle, from: • Passenger Vehicle • Buses and Coaches • Light Commercial Vehicle • Heavy Duty Vehicle • Motorcycles Currently only M1 (Passenger Vehicles) and N1 (Light Commercial Vehicles) are mandated to use eCall.
VIN	(set on installation) Vehicle Identification Number.

Field	Description and provenance of data
Vehicle Propulsion Storage Type	(set on installation) One or more from: Gasoline, Diesel, Compressed Natural Gas, Liquid Propane Gas, Electric Energy Storage, Hydrogen Storage, Other
Timestamp	Time of event. This is generated by the eCall unit.
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	(set on installation) The originating number of the eCall unit
Received at PSAP	The time of receipt at the PSAP.

2.4.2 eCall false alarms

The data in the MSD is mostly set on unit installation and hence reliable. Other fields are tightly defined (true/false or enumerated), which reduces the chances of error.

Automatic activations

Automatic activations are sent without human intervention when a suitable vehicle condition occurs, such as high deceleration or airbag deployment. 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.

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

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

- 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

Faulty eCall devices

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

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.

Vehicle location accuracy

eCall provides the last known location, direction of travel and optionally the previous two known locations. The locations may be the vehicle’s GPS position or a mobile phone location service.The MSD indicates whether there is 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 confidence is set high, study of GPS accuracy statistics suggests that the location accuracy would allow carriageway-specific (or at least neighbouring lane) location for both open skies and in GPS shadows.

However, 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.

Relevance to responder

Road networks in one country are often managed by multiple different authorities. As eCall activations can originate from any location, a responder who receives unfiltered eCall MSDs will need to disregard those outside their area of responsibility i.e. MSDs used in SVD should be geographically filtered for each responder’s areas of responsibility.

2.4.3 eCall rate analysis and false alarm rate estimation

The researchers sought access to eCall data for analysis but were not able to secure access. This limits the depth of analysis - we have used available 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. Access to data would enable deeper analysis.

Figure 1 shows voice calls from eCall activations in England from March 2019 to June 2021. 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.

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. when they are safe following an automatic alert or a false alarm. The unconnected calls average around 60% of the total.

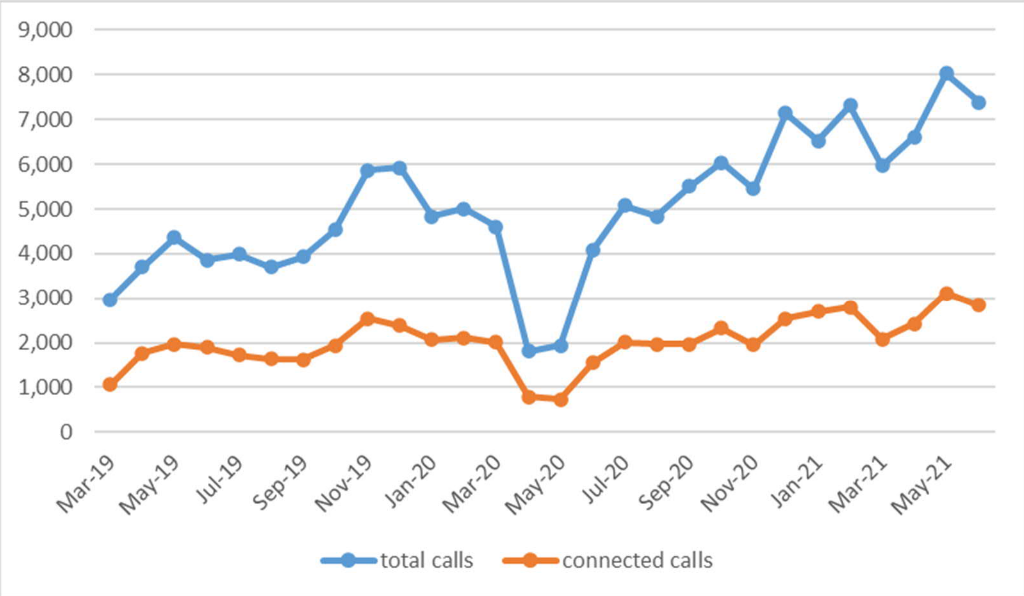

Figure 1: eCall totals and connected calls in England, Mar 19 - Jun 21

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 being automatic.

Assessment of Automatic eCall false alarm causes and rates

For analysis we make reasonable assumptions about false alarm rates:

Source	Description	False alarm rate
False alarm – faulty unit	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% (of automatic calls)
	Total Automatic eCall false alarm rate	5%

The overall false alarm rate contribution of automatic eCalls to total eCalls is the proportion of automatic calls (5%) * the propotion of false alarms (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**
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 calls.	40% (of manual calls)
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 contribution of manual eCalls to total eCalls is the proportion of manual calls (95%) * the propotion of false alarms (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 may still be worthy of investigation. 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 Call	**2 Calls**	**3 Calls**	**4 Calls**
Simple manual false alarm rate estimate	40%	16%	6%	2%

2.4.4 Methods for using eCall MSD for SVD alerts

From an understandiong of the characteristics of eCall MSD we now propose specific methods for its use in SVD. These methods fall into four areas and form a workflow process for MSDs (Figure 2).

- 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 alerts to the emergency responders

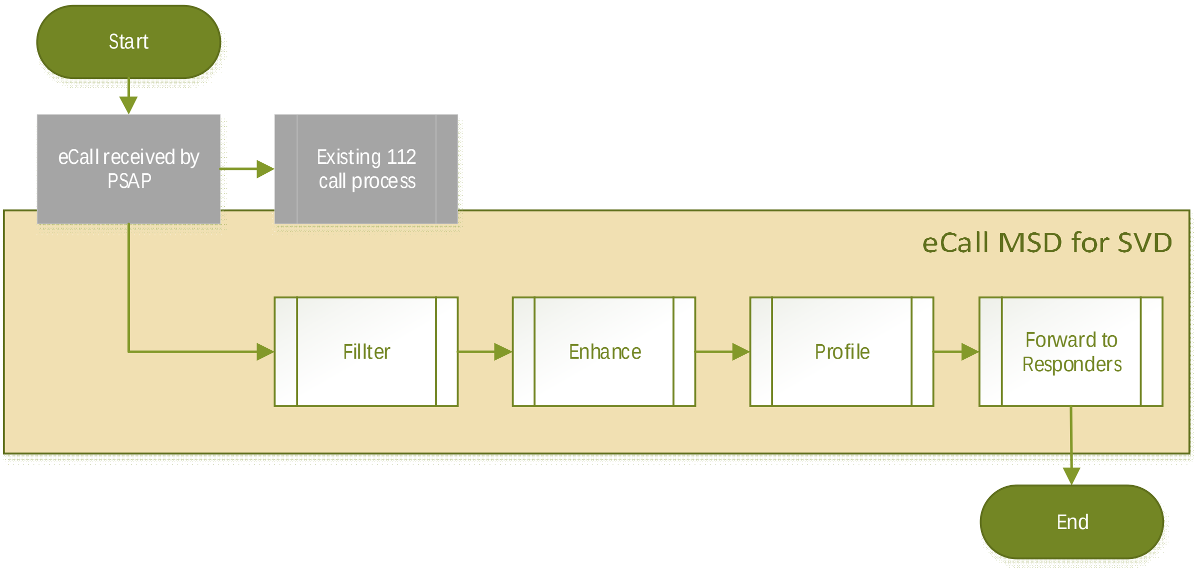

Figure 2: eCall MSD for SVD alerts

Filter

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

- VIN Blacklist

  Identify faulty units as a source of false alarms, by detecting unusual numbers of MSDs from a single vehicle and blocking 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 identify areas where the data can be enhanced to assist responders by making it operationally relevant.

- Add road names, direction, and reference points

  The MSD contains coordinates and optionally the previous two positions. These can be used in a map-matching process to identify the road number and direction, and nearest reference points such as a junction, markerpost or other feature.

- 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 (significant for recovery) and colour.

- Add event flags

  Information in the MSD may suggest a higher risk category, for example hazardous fuel. Higher risk cases can be detected by an algorithm and flagged.

- Add related events

  Note where other MSDs have already been received within a given radius and a given time, to allow responders to tie multiple alerts together.

- Handle retransmitted events

  A PSAP may request retransmission of an MSD. 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.

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

Forwarding the MSD SVD alert data directly to responders can significantly reduce response times. 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 above a certain priority.

TeCall

Chiltech and partners[1] have built a demonstration of this MSD processing capability, called TeCall, with a response time in the order of seconds.

Conclusion

These methods demonstrate that 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. The enhancement of the MSD with additional data provides an opportunity to improve the operational value for responders.

2.4.5 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

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

Very High) as our quantitative data is only from the UK and can only be indicative for other countries.

Area	eCall	Infrastructure
Road coverage	Very High 98% of UK roads (roads with sufficient mobile coverage.) Provides SVD where infrastructure 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 authorities there is no cost for the installation and operation. Relatively low cost for 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 of 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 plotted in Figure 3, with 0 = Very Low and 5 = Very High.

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

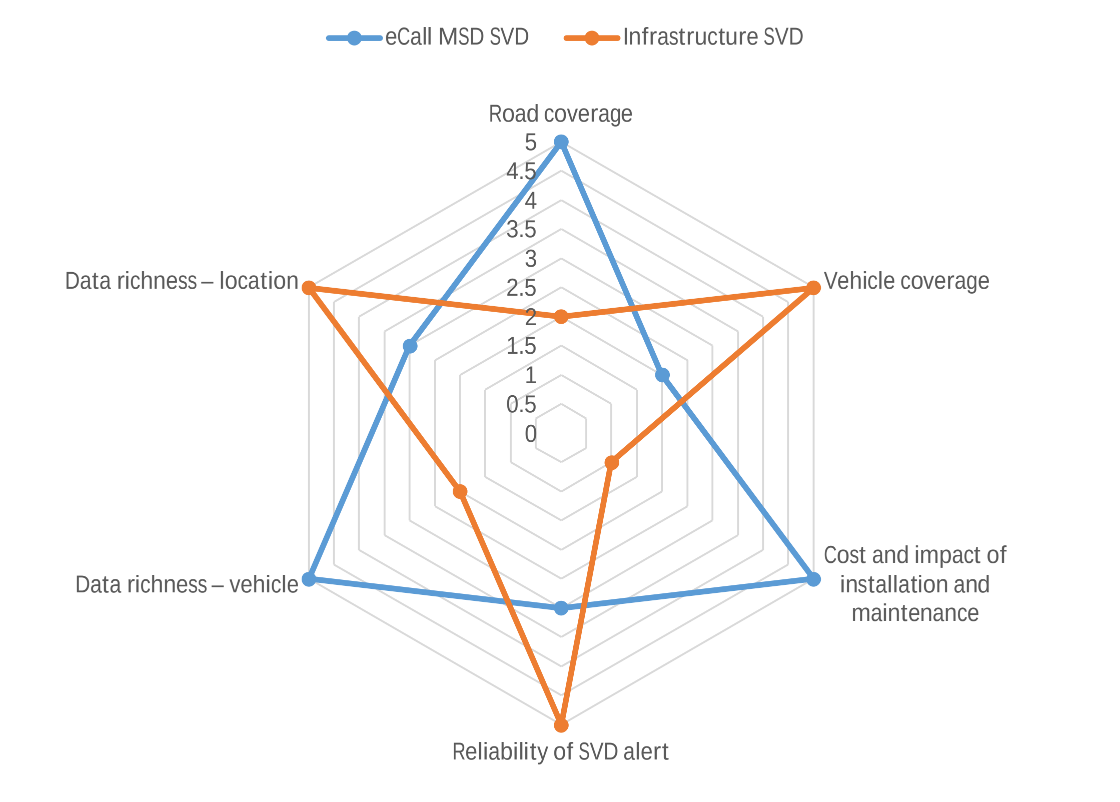

Figure 3 Relative values of eCall MSD SVD and Infrastructure SVD

2.4.6 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:

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

2.5 Conclusion - 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 relevance but we have demonstrated that these can be overcome.

eCall is not standing still; the next generation eCall standards already exist. Whilst the current eCall uses 2G and 3G and communication technology, the next generation will use 4G and beyond, and will add a capability to trigger additional sensors on the vehicle.

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.

It would be desirable to use larger datasets for such analysis, but acquiring live traffic data can be difficult due to sensitivity reasons, and acquiring test data has logistic challenges.

3.2 Lane discrimination information

The spacing between radar sensors influences capabilities and potential such as identification of the lane of the stopped vehicle. The lane could be useful in determining the hazard level for prioritisation, since a stop in a fast lane is more hazardous than a stop in a slower lane.

A rotating radar sensor can detect and track objects in both range and azimuth (angle) from the radar. Currently when a stopped vehicle is detected its longitudinal position along the road is indicated over a given 100m section length of the road, and the azimuth data is currently not utilised.

The precision of both range and azimuth is dependent on first the radar hardware, and then the software that processes the raw data. The rotating radar hardware can detect object at a range resolution of 17cm out to 500m in all directions. In contrast the azimuth precision is dependent on the beam width of the radar, which increases with distance from the radar. This means that the uncertainty of azimuth-based location 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 the centre of a temporarily closed motorway lane covered by a rotating radar system, making a series of stops, and GPS measurements were taken in the vehicle. 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.

Our first analysis compared the radar-derived points to GPS-derived points, but this did not show any relationship between the distance from the radar and the distance between the GPS and radar-derived points. Further study of the GPS data showed that it diverged from the lane which had been driven – it was not sufficiently accurate to support the planned analysis.

A second analysis used the differences between radar-derived positions and the centre line of the lane, as shown in Figure 4.

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. Gathering such data on the road is not logistically simple, and the data volume gathered in this project is not sufficient to establish statistical significance.

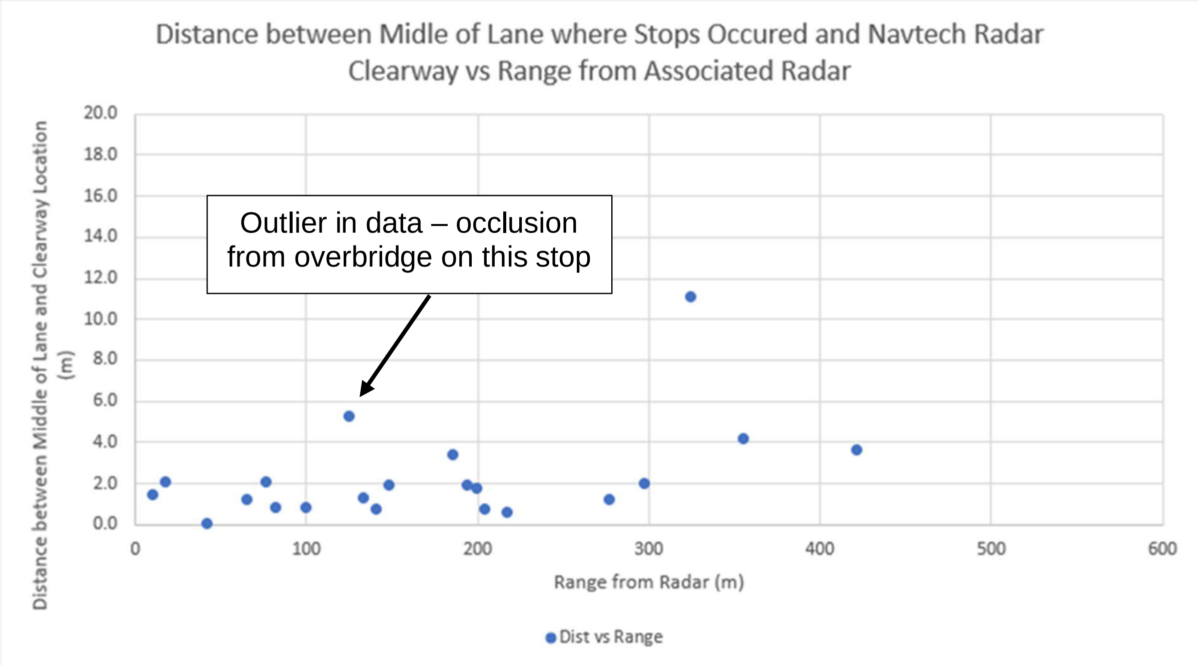

Figure 4: Distance between centre of the lane and rotating radar vs range from the radar

3.3 Quantifiable traffic parameters

The way in which a vehicle stops on a live lane is dependent on the 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. 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 of other motorists. This section explores whether these unusual traffic behaviours can be identified from radar data how additional data can help identify the risk of a stopped vehicle.

3.3.1 Traffic Speed

In the example explained above, a stopped vehicle may occur after a car slows down and causes 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 divides the road into 100m length sections, and traffic statistics such as speed and density can be calculated for each. Traffic speed data from this system is shown in the graph in Figure 5. This shows traffic speed from 16:00 to 17:00 gathered from one typical weekday afternoon. 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.

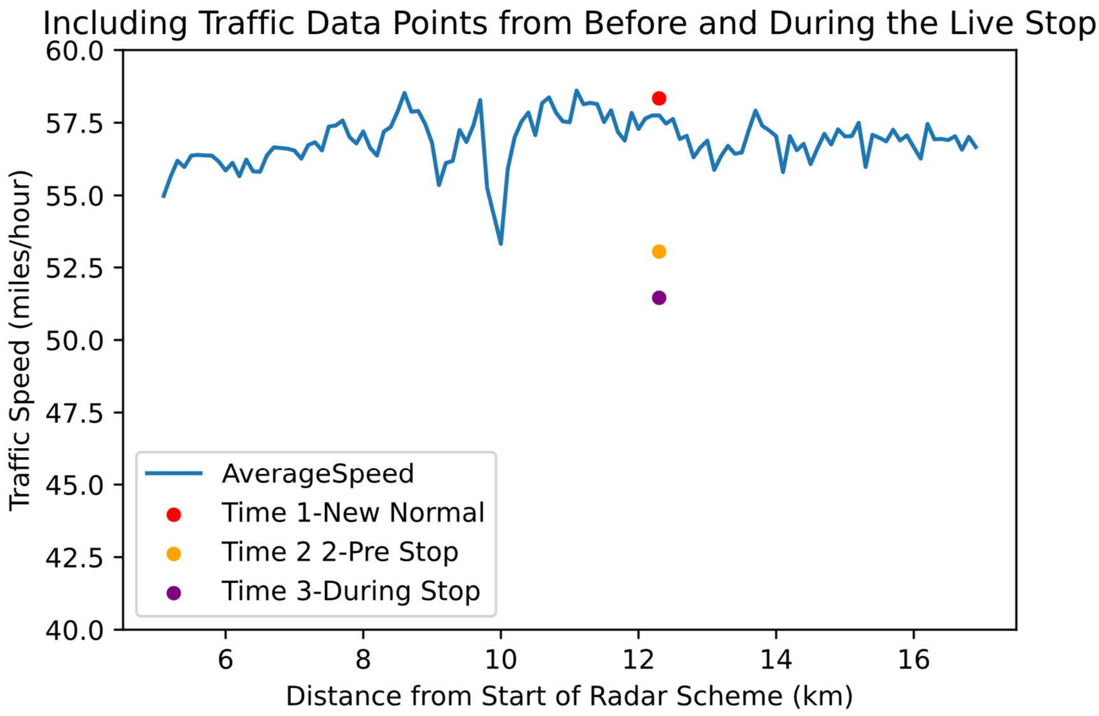

Figure 5: Depicting what "normal" traffic looks like, from 16:00-17:00 on a particular scheme that the Navtech Radar Clearway system is deployed in. Coloured points show traffic data at 12.3km into the scheme, before and during a particular live stop

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 (fifteen minutes up to one minute before the stop), Time 2 – pre-stop (the minute before the stop), Time 3 - during the stop (duration of three minutes).

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.

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 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 a stop on a narrower road.

In another example in a 2-lane tunnel scheme, two roadwork vehicles were seen to stop in the live lane. 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 as shown by the blue line in Figure 6. The speed remains relatively constant at around 80 km/h. The orange line shows the average speed during the period that there was a stopped vehicle, and the green line shows the speeds for the 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 line (during stop) at 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. In the green line, a smaller drop can be seen in the adjacent section. This may signify the slowing down of traffic as the vehicles approach a stop. Four minutes was the minimum duration retrievable from the system; if this time could be reduced, it is possible that more significant speed reductions would be observed.

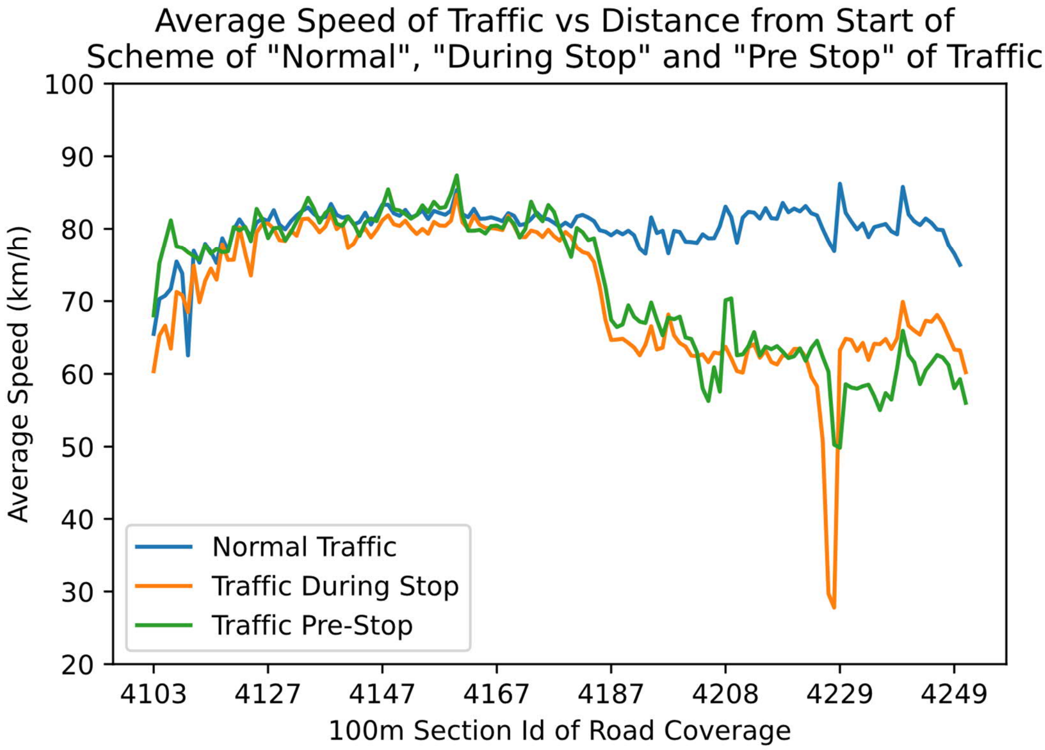

Figure 6: Traffic normal, during stop and pre-stop

A further example was acquired on another scheme on a single-lane carriageway, where one would expect a much greater impact on the traffic behaviour. Traffic speeds are shown in Figure 7. The average speed of the traffic after the occurrence of a stopped vehicle drops almost to 0 miles/hour (blue). Traffic speed recovered gradually. Data was not available prior to the stopped vehicle event in this example.

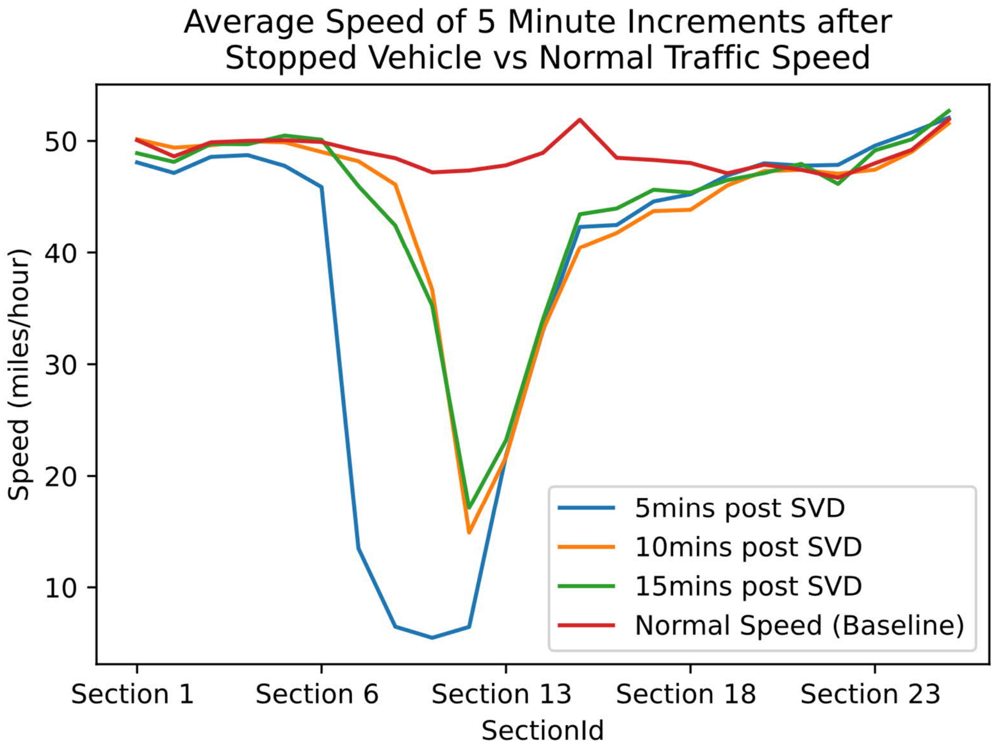

Figure 7 - Average speed of traffic in a 1 lane scheme

These examples illustrate how speed data might be useful in supporting the detection of an event of a stopped vehicle. On many-lane roads the changes are not likely to be considered significant enough to cause an alert, but they might be used as supporting evidence for an alert in a data fusion system. On single-lane roads the behaviour change is more significant and may trigger alerts.

A related research question is whether detection of individual slowing vehicles can be useful in stopped vehicle detection. 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* detects 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 our data, approximately 10% of stopped vehicle alarms had an associated slow vehicle alarm. This result is quite low, and not considered significant enough to formulate a direct correlation between these two alerts. The slow vehicle alarm configuration was then altered to detect a vehicle travelling between 2 and 25kph for just 2 seconds within each 100m section. This new configuration was run on the system for a period of approximately 3 weeks. With that configuration, approximately 50% of stopped vehicles had a preceding slow vehicle alarm. This suggests that slow vehicle detection might be used to support the confidence of detection of a stopped vehicle. It does not seem pactical for such slow vehicle alerts to be raised to the operator independently from a stopped vehicle alarm, as a large number of slow vehicle alarms was generated.

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 8 below.

Figure 8: 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 8, 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.

Figure 9 shows the locations of heads and tails of queues (yellow) and isolated stopped vehicles (red dots) on a Navtech Radar Clearway scheme. 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.

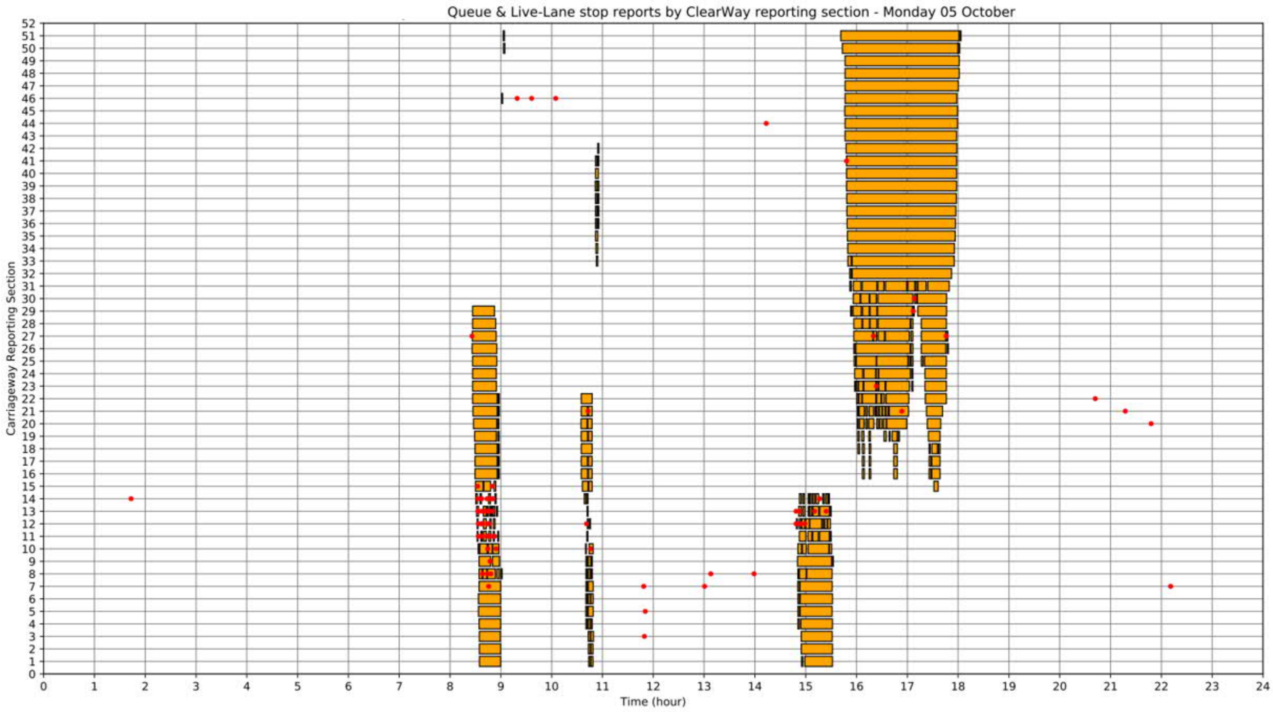

Figure 9: Navtech Clearway Scheme - queue length (yellow) and isolated stopped vehicles (red dots) by section vs hour of the day

3.4 Pedestrian information

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.

Pedestrian detection is a core function of the Navtech Radar Clearway system. This uses the 360° rotating radar to detect and track a pedestrian in the radar's coverage area.

This passenger location information can be considered crucial in certain cases of stopped vehicles that pose greater danger for the passengers and road users. There can be cases where visibility is low due to extreme weather or smoke, but radar can still detect the pedestrians.

Pedestrian detection is illustrated through data gathered by another existing Navtech Radar scheme 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. Figure 10 is a camera image capturing this scenario.

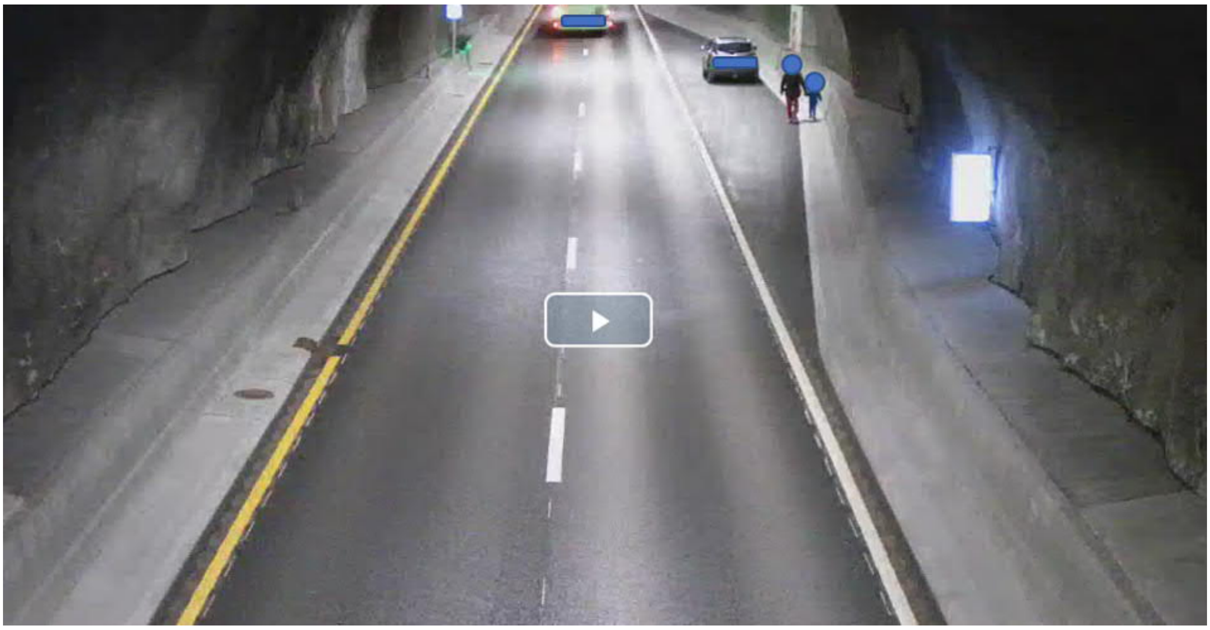

Figure 10:- Stopped vehicle in tunnel, with motorist and child walking to a specific location

The Navtech Radar Clearway system was able to detect and track this as shown in Figure 11, with the pedestrian tracks of both individual pedestrians (yellow and green) extracted and then plotted on an X vs Y coordinate plot, and the stopped vehicle track represented in blue.

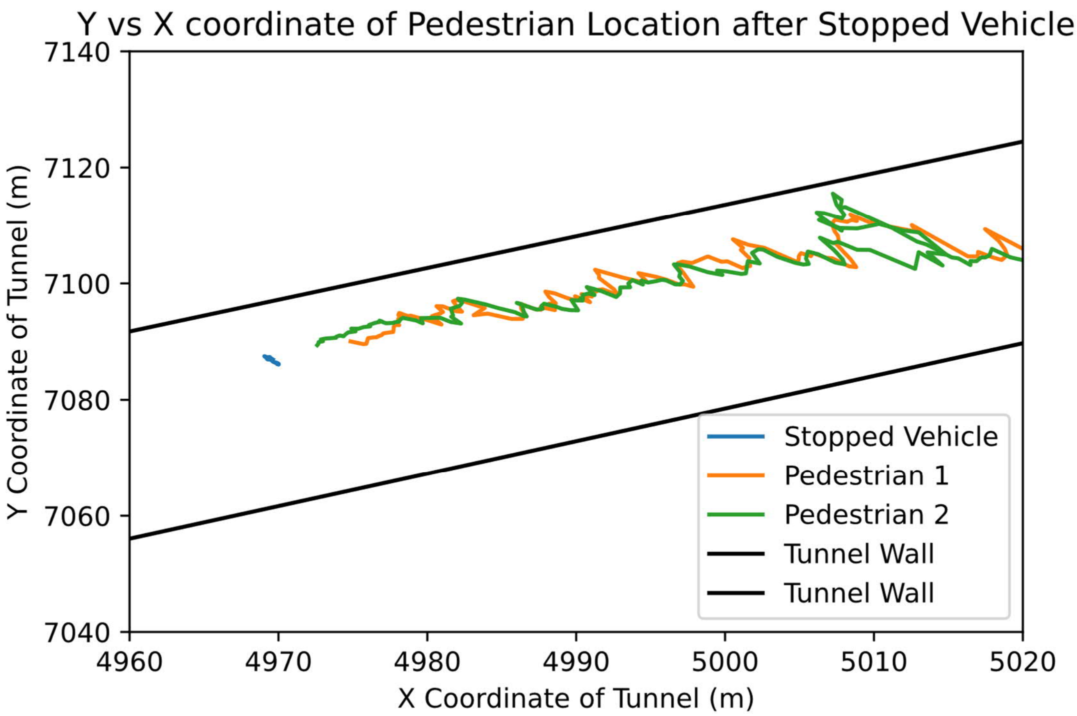

Figure 11: Pedestrian and stopped vehicle tracks of tunnel pedestrians

The graph shows that the vehicle had stopped in a section of the tunnel, with both pedestrians disembarking, and walking to another part of the tunnel. The radar capability 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 emergency personnel could be informed to better understand their recovery plan. The information may also influence the priority with which the alert is raised to operational personnel, and could support more precise messaging to the other motorists about the incident.

3.5 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 a 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 potential impact risk level can 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, the impact risk level can be further increased.

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, which includes "obstacles… on the road".

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 based on obtaining sensor data from the vehicles which is enriched and aggregated to obtain information usable by the road authorities (Ismail, 2020). Data that can be sent from vehicles includes manually and automatically triggered breakdown calls and eCalls, numeric sensor values from which a stop might be inferred, and external object detection which might include another vehicle stopped.

A proof of concept was held in 2019-2020 (van Rij, 2020) which included the provision of STRI data by several companies. The providers retain ownership of the information which is provided under conditions to the participating NRAs. Evaluation has focussed on the Netherlands – there the data included over 100,000 alerts of obstructions and over 80,000 alerts of vehicles in difficulty.

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:

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

For alerts of broken-down vehicles, the time-saving amount was up to 7.5 minutes compared with the available data from reported incidents in the public sector dataset.

Further study of the penetration rate, fleet size, and growth rate of these connected vehicles should give more insight into the potential of this data source.

4.1.2 SRTI data feed service providers

Irrespective of the current and future extent of data provision through DfRS, some companies currently offer commercial safety-related data services.

TomTom

The traffic incidents API from TomTom includes the categories "Accident" and "Broken Down Vehicle". The data includes the length and delay of a possible traffic jam, estimated intensity level, and the likelihood of a report being true. TomTom collects this data from sources covering 79 countries including over 600 million devices: GPS devices, mobile phone signals and sensors. The costs (2021) are 5 cents for 1000 tile-based API requests or 50 cents for

1000 non-tile-based API requests.

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 floating car data.

HERE

The HERE incidents API, similarly to the TomTom incidents API, provides a list of reported incidents within a given area, including categories “Accident” and “DisabledVehicle”. Data includes the possible length of a possible traffic jam, delay, severity level, whether the report has been verified, and whether a response vehicle is on the way. HERE collects this data from multiple sources, including connected car probes, roadway sensors and live operations centres. Annual cost is based on population. Indicative estimates were 50 k€ per year for NL and double for UK per year (without any discount).

4.1.3 C-ITS safety related services

C-ITS services are standards-based exchanges of data between vehicles, the roadside infrastructure, control and service centres, and other road users. There is significant activity in research and pilots, with the first national implementations in progress. C-ITS services include stationary vehicle warnings. Relevant C-ITS messages, both to and from vehicles, are identified in Table 1.

	Message		Description
CAM	Cooperative Awareness Message	Day1	Maintain awareness of cooperating objects including ther positions. Can indicate a stopped vehicle.
DENM	Decentralized Environmental Notification Message	Day1	Alert of road hazard or abnormal traffic conditions, which may include a stopped vehicle.
CPM	Collective Perception Message	Day2	Informs of locally detected objects to improve situational awareness. A CPM indicating a detected stopped vehicle may lead to a DENM alert.
MCM	Maneuver Coordination Message	-	Coordinate maneuvers between stations (early stage of development) (V2X). Could help other vehicles safely navigate past a stopped vehicle hazard.
IVIM	In-Vehicle Information message	Day1	Conveys information in the form of signs or pictograms. Could be used to warn of a stopped vehicle ahead.
MCDM	Multimedia Content Dissemination Message	-	Sharing of multimedia content between ITS stations to improve environmental perception. Could convey an additional warning of a stopped vehicle ahead.

Table 1: Messages relevant to stopped vehicle hazards

C-Roads implementation status

Slow or stationary vehicles use cases are implemented on around 20 pilot sides in 13 countries through Europe. The status of the pilot sites in 2019 was reported by Kernstock (2019). Results/evaluation of the pilot sites was reported by Gruber (2021). We are not aware of published consideration of effectiveness for stopped vehicle detection in terms of coverage, detection rate or false alarm rates.

C2C-CC

The CAR 2 CAR Communication Consortium (C2C-CC) guides C-ITS developments, and is working on the definition and development of Day 2 services and technologies which promise to enable the sharing of information about objects detected in the environment.

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. The CAM, DENM, MCM messages were extended.

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 (Wijbenga, 2021), (Wijbenga, 2020). Simulations used mixed traffic scenarios (automated, cooperative, and conventional vehicles), initially with ideal communication then 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.

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 Coverage of roads and vehicles

Notifications from connected vehicles have the potential to cover all roads where there are suitable communications signals, but not all vehicles are yet covered. eCall coverage has been considered 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 in general has been predicted to have a compound annual growth rate of 16.9% (Mordor Intelligence, 2020), with 237 million vehicles globally in 2021.

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, 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

Past research found that traffic incidents like SVD could be classified and georeferenced to create a practical filtered feed for human interpretation (Grant-Muller, 2015). This depended on the precise geo-tagging of tweets, which was available for 1% to 2% of tweets at the time of the research.

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. 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 2.

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 2: 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 quantitative and qualitative difficulties severely limit the 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.

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 3. The full list of the NDW situation record types is available in the documentation.

SituationRecord	SituationRecordType
VehicleObstruction	brokenDownVehicle brokenDownHeavyLorry
Accident	accidentInvolvingHeavyLorries overturnedHeavyLorry accident

Table 3: 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 4 based on their relevance to SVD. Only these subtypes were considered in further analyses.

Type	Subtype
Accident	Accident_Minor Accident_Major
Hazard	Hazard_On_Road Hazard_On_Road_Car_Stopped Hazard_On_Shoulder_Car_Stopped

Table 4: Selected Waze alert types

Waze also provides a breakdown assistance feature consisting of three types of reports:

“Fellow Wazers”, “Emergency call”, and “breakdown help” (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, but these two subtypes are not available in Europe 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).

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%) were matched to at least one Waze alert. Vice versa, 31% of the Waze alerts on the IM network were 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 5 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, 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 6 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 7). 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 5: Comparison of the relative amount of NDW situation types between all situations and those not matched to a Waze alert

Alert type	Waze alerts on IM network: Total	Waze alerts on IM network: 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 6: Comparison of the relative amount of Waze alert types between all alerts on the IM network and those not matched to a NDW situation

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

Table 7: Comparison of the relative amount of Waze alert types between all alerts and those outside the IM network

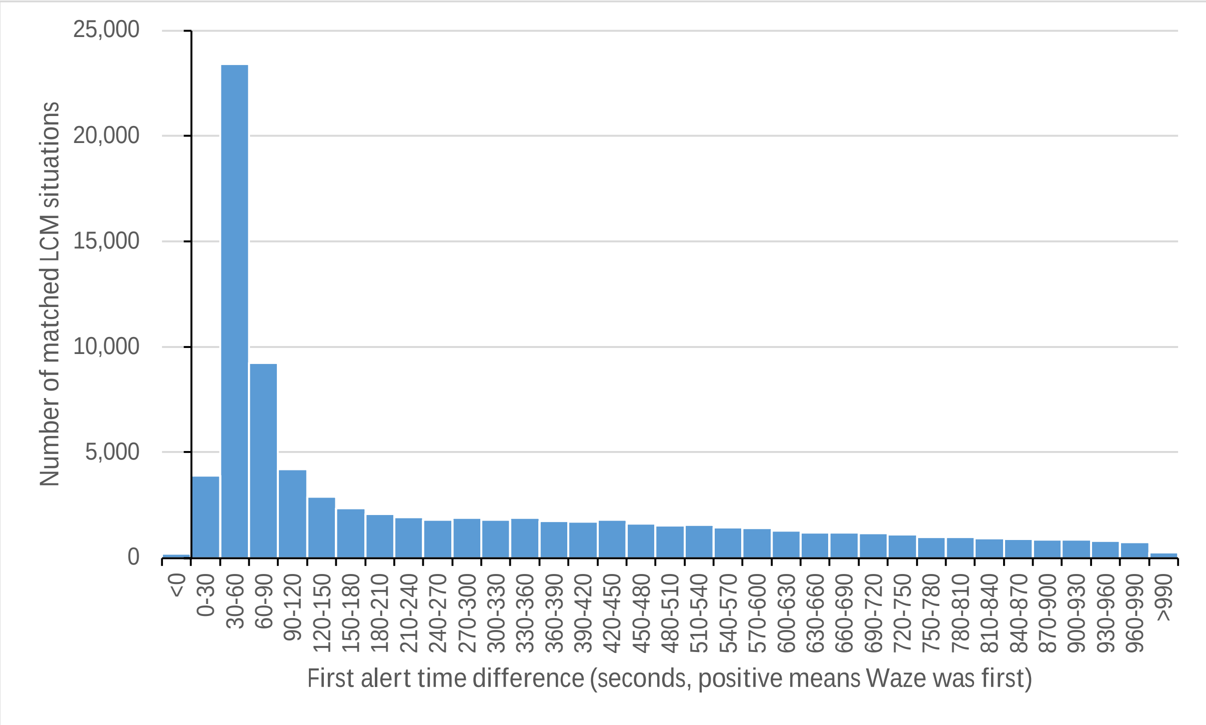

Figure 12: Distribution of the first alert time difference between NDW and Waze

Figure 12 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 TMCs until the end of 2020. This combination is a strong indication that TMCs use Waze alerts in their operations and that both data sources are not independent.

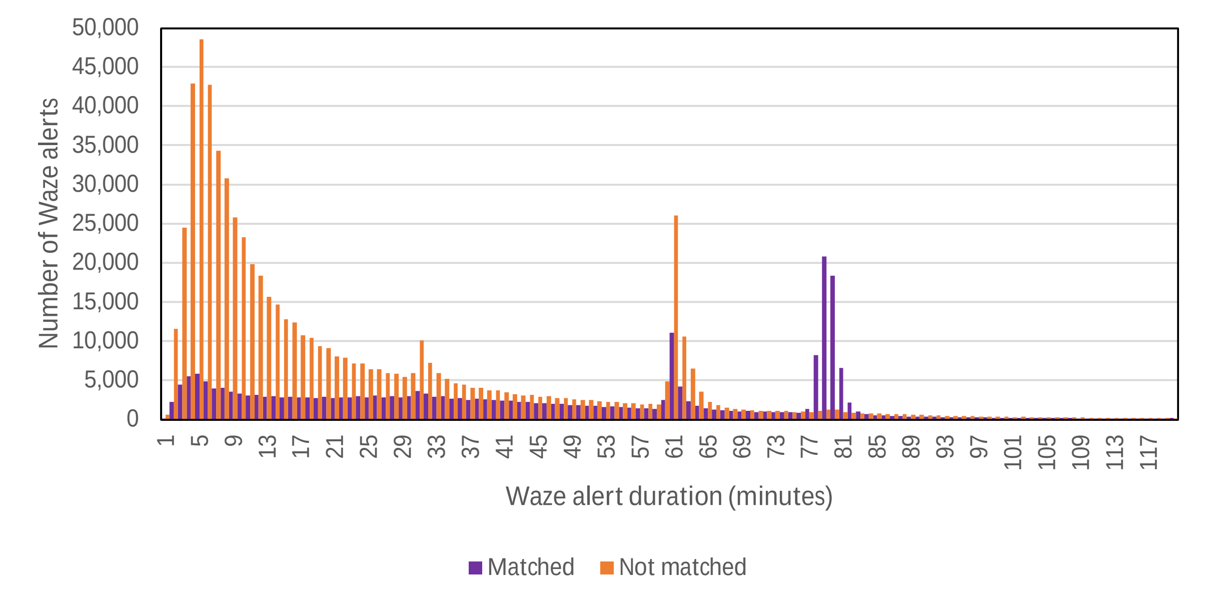

Figure 13: Distribution of the Waze alert duration

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 13). 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 13 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 14 and Figure 15. Figure 14 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 15 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.

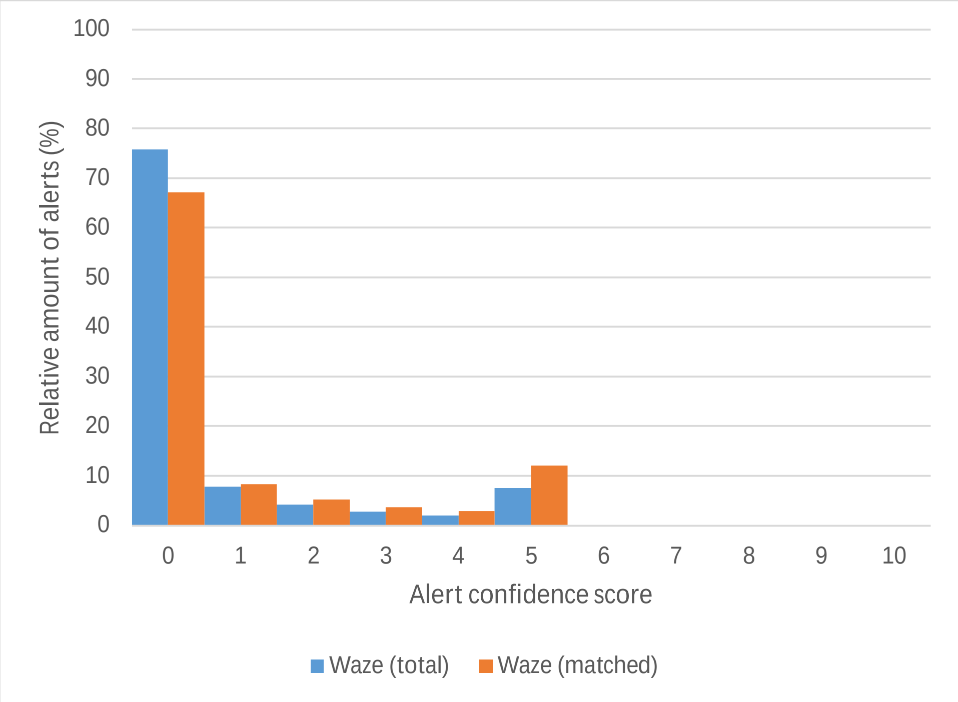

Figure 14: Relative amount of Waze alerts per confidence score

Figure 15: Average duration of Waze alerts per confidence score

These analyses 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)

Aerial imagery has the potential to cover large areas with a high spatial resolution and without installations by the road authority. Challenges include cost and weather dependence. Each kind of source of aerial imagery brings its own advantages and disadvantages which are further explored below. Most of the research focuses on image processing algorithms to detect and track vehicles, rather than hardware and deployment issues.

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”.

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. The figures show that 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.

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.

Weather, especially clouds, also significantly hampers 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

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.

Research on vehicle detection methods suited for UAVs (Unmanned Arial Vehicle) shows promising results. One research (Xu et al, 2017) reaches up to 98.43% correctness and 96.40% completeness rates. Other recent research (Barmpounakis, Sauvin, & Geroliminis, 2020) shows 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. 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

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 are different. Research has typically focused on finding and testing methods for moving vehicle detection, and studies have mainly been limited in scale, so it is not yet possible to identify which methods would be best suited for stopped vehicle detection.

Conclusion

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 promise 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, this could make them feasible for targeted situations and locations.

4.4 iWKS development Netherlands

In the Netherlands, Rijkswaterstaat is working on the next generation roadside unit called iWKS (Rijkswaterstaat, 2021). 5700 new systems will be installed on Dutch highways. iWKS will maintain existing functions but will also enable extension with new applications or data sources. Such roadside devices have often been difficult to change, but iWKS is designed to enable change quickly with automated integration and deployment processes. Traffic management applications are stored in a container and are automatically distributed on the designated roadside units. In this way new software related to stopped vehicle hazards could be deployed to the roadside, for example interfacing to new detection or C-ITS communications.

4.5 Conclusions

New methods have potential to improve detection in timeliness, reliability, accuracy, and information content. The potential of connected vehicle sources is growing rapidly as the vehicle coverage is growing. The variety of different detection 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 16 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.

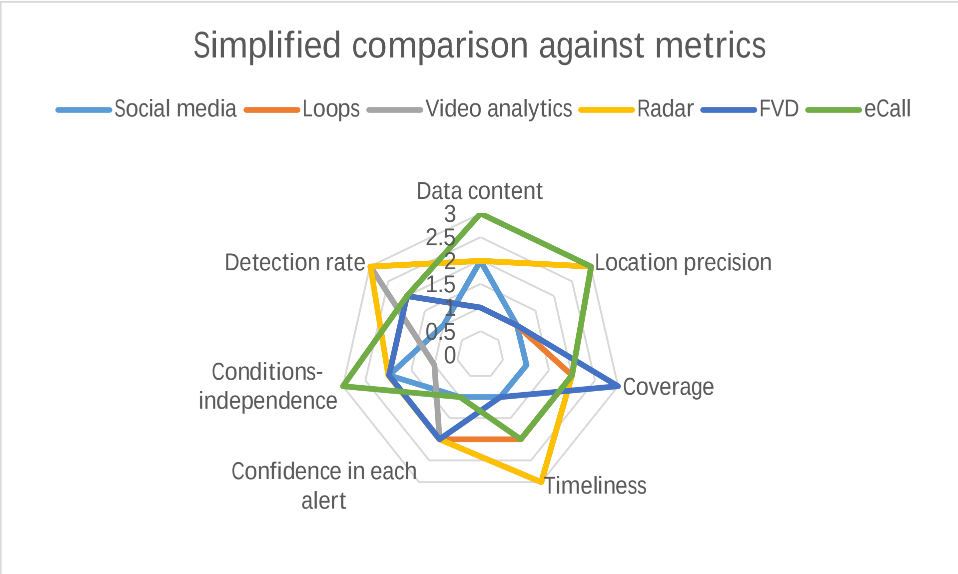

Figure 16 Simplified comparison of SVD technologies

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 17.

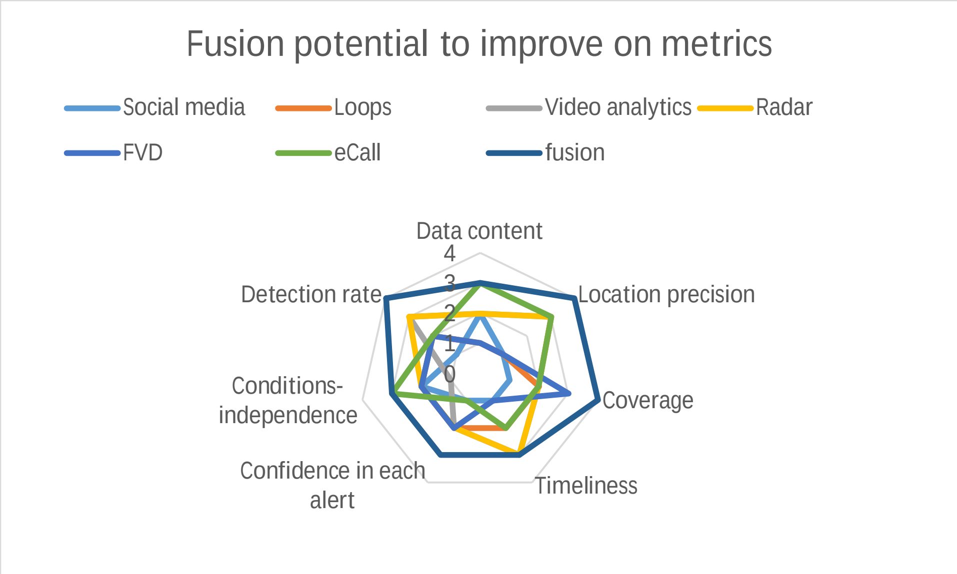

Figure 17 Potental of fusion to improve on individual SVD technologies

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) who noted an increasing focus on fusion through deep learning. A CEDR review by Cornwell et al (2016) noted two contrasting kinds of fusion for incident detection: (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.

(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. This approach has achieved success in other domains. 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. The full version of SHADAR D5.1 notes some considerations for stopped vehicle detection by machine learning.

(ii) **Fusing outputs of multiple detection sources.** This approach has less 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.

5.3 Study using probability to characterize fused detection

This section defines 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. The full version of SHADAR D5.1 provides an expanded description with worked examples.

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.

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 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 5.3.5. Alignment in location is possible by adopting a coarse resolution for the fusion system – see 5.4 for further discussion. The analysis defers consideration of the influence of weather and traffic state until section 5.6.

5.3.1 Fusion regime A – alert if any source alerts

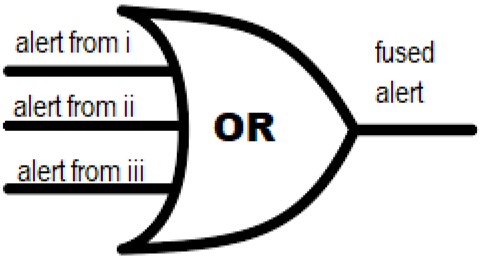

Figure 18 OR fusion regime

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) …

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.

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).

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:

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

Fusion regime A will tend to improve the detection rate at the cost of an increased false alarm rate. 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.

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

Time to detect

TTD(fused) = min(TTDi, TTDii, …)

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

5.3.2 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.

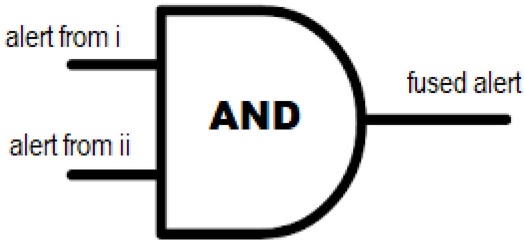

Figure 19 AND fusion regime

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.

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.

5.3.3 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.

5.3.4 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.

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

To calculate the fused false alarm rate we would need to know additionally how many times false alarms occur simultaneously.

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

5.3.5 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.

P(event given one source detects while another does not) = the proportion of times an alert unique to the detecting source represents a real event, out of all alerts unique to that 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.

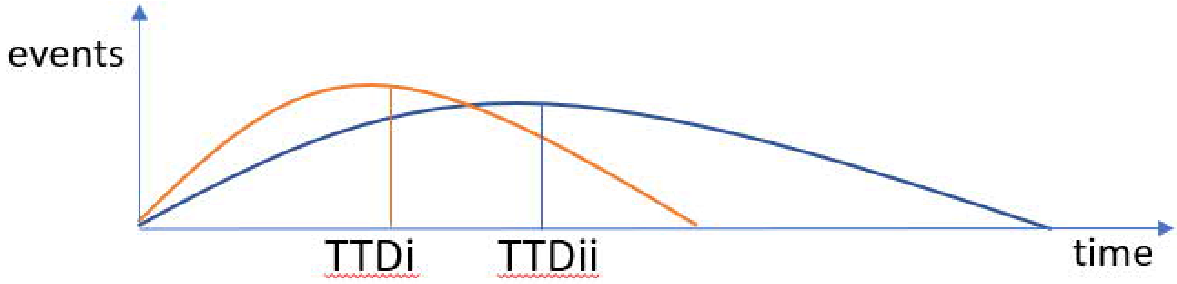

Figure 20 detection time profiles

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.

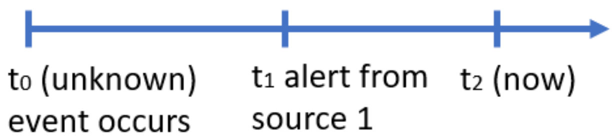

Figure 21 occurrences and times relevant to a stopped vehicle event and its alerts

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.

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

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

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.

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

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 5.6.

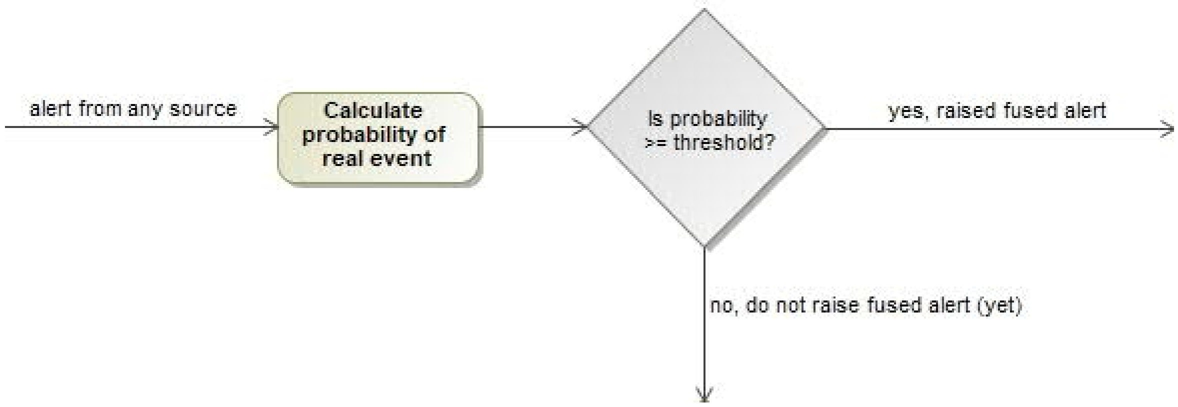

Figure 22 Selective data fusion regime

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 probability of an event may 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 5.6.

For each alert, we calculate P(event | this permutation of sources alerting and not alerting) = true alerts from this permutation / all alerts from this permutation.

To calculate a full DR(fused) and FAR(fused) would require data on 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.

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) – sum for all sources(P(alert unique to source))

Compared to having no fusion, this fusion regime maintains false alarms at an operationally tolerable level, at the cost of reduced detection rate.

The confidence probabilities can be pre-calculated for permutations of sources; dynamic calculation and decision-making at runtime is considered in section 5.6.

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.

A worked example illustrating statistical data fusion using a combination of infrastructure-based and connected vehicle detection sources is included in the full version of SHADAR D5.1.

5.3.7 Summary of fusion regimes

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.

5.4 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 6).

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.

5.5 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 a period of 3 months 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.[2]

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 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

---

2 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.

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. 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.

5.6 Fusing and decision-making dynamically at runtime

Section 5.3 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.

5.7 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 5.6. The mathematics should produce equivalent results, so it is more a question of the ease of thinking about and managing the data.

5.8 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 6).

6 Reporting alerts and performance

6.1 Operational user interfaces for alerting

This section considers how operational user interfaces might support the kinds of features explored in previous chapters, such as multi-sensor data fusion.

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.

User interface mock-ups were constructed to explore function and feasibility, and key features are summarised here.

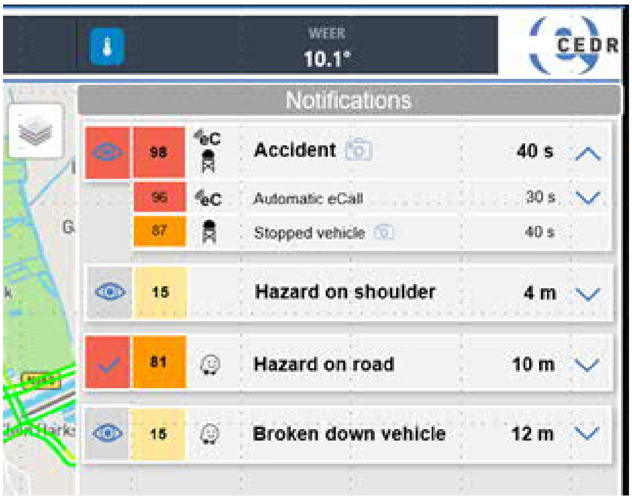

Figure 23 Notification pane

Figure 23 shows a notification pane which could appear in a traffic management system to convey alerts from a data fusion system.

- Alerts from multiple sources that are judged to relate to the same stopped vehicle event are grouped together. Groups could be expanded or collapsed by the user as required.
- Confidence is indicated through both colour and a percentage value (which could be derived using the techniques of chapter 5.
- The overall fused confidence is displayed, and the confidence attached to the individual sources is also displayed.
- The time since the alert is displayed – both for the overall fused alert and for the individual source alerts.

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.

Contextual information available to the fusion system could be integrated in the operational display – for example weather measurements could be shown either in side-panes or as colour effects superimposed on the geographic map. Relevant CCTV feeds could be shown in response to alert selection – this could save time in validation.

Additional data from new detection sources could be shown with alerts – for example vehicle and propulsion types obtained by enhanced eCall alert processing. Vehicle traces available from eCall data could be highlighted on the geographic map when an eCall alert is selected.

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 SHADAR report D6.1.

6.2 Performance statistics reporting

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 a simple report 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. Such data can also be used to inform incident analysis e.g. identifying road sections with high incident rates.

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. The statistics would be presented for a selected range of time and location.

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 24: 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, and 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).

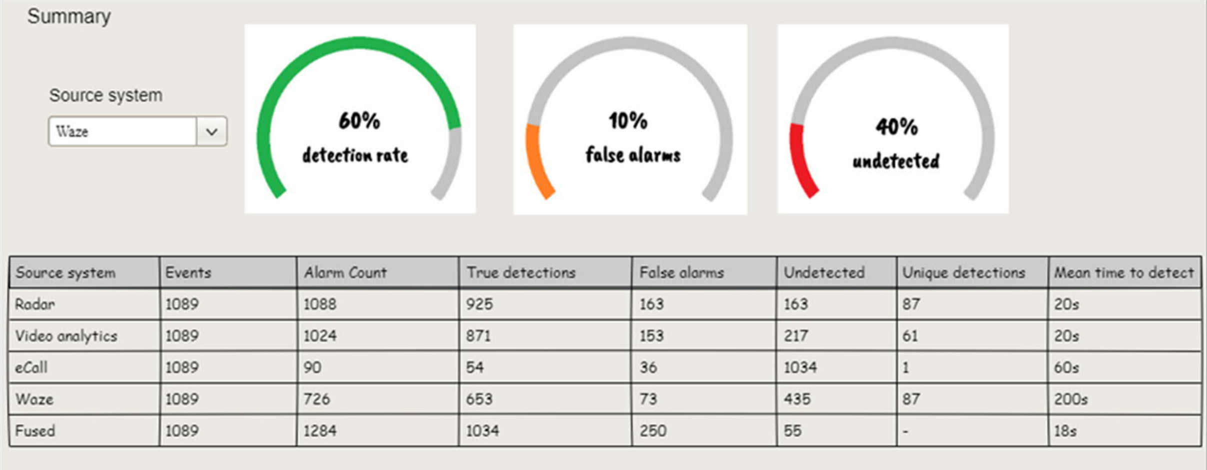

Source system	Events	Alarm Count	True detections	False alarms	Undetected	Unique detections	Mean time to detect
Radar	1089	1088	925	163	163	87	20s
Video analytics	1089	1024	871	153	217	61	20s
eCall	1089	90	54	36	1034	1	60s
Waze	1089	726	653	73	435	87	200s
Fused	1089	1284	1034	250	55	-	18s

Figure 25: Additional statistics calculable when ground truth data is available

6.2.1 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 Conclusion

This report covers diverse ideas. To help to illustrate and relate the ideas in a common and intuitive way, the ideas are presented here using 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 26:

- 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.

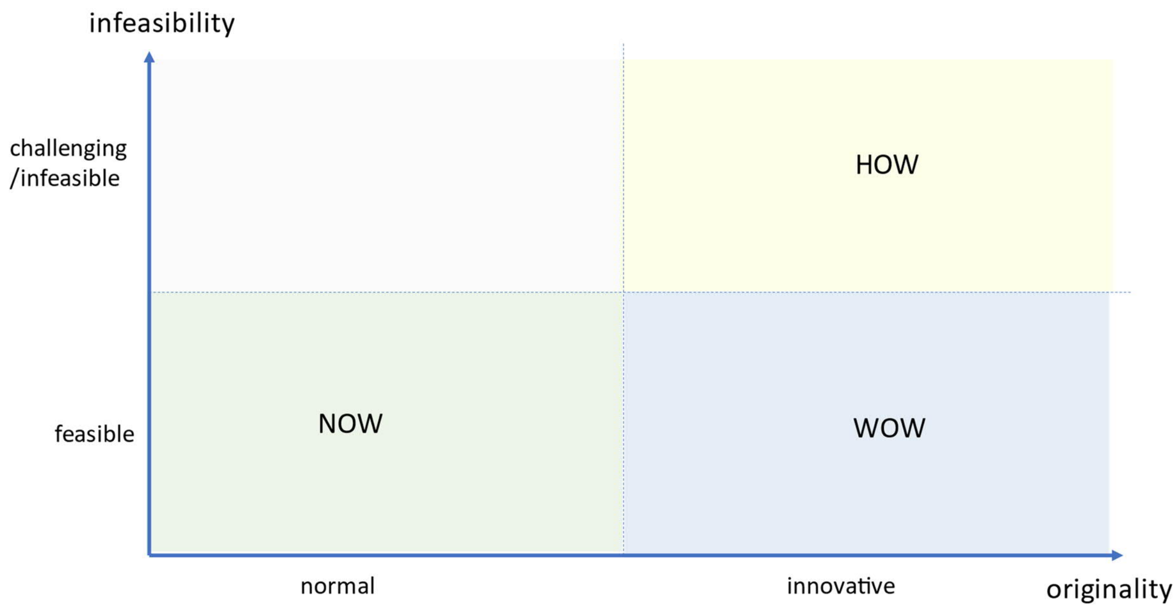

Figure 26 How-Now-Wow matrix used to characterise ideas

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.

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) 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	Timely	Reduces delays with SVD going to wrong control centre

NRAs		
(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 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.
(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.
Driver awareness campaign	Reliability	Driver awareness of reporting including eCall (campaign, part of driving lessons)
(6) Alert source combination	Timely	Avoids separate operator investigation of separate related alerts

	(6) Integrated weather presentation	Information	Weather notifications per environment
	(6) 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.
	(6) 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 27 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.

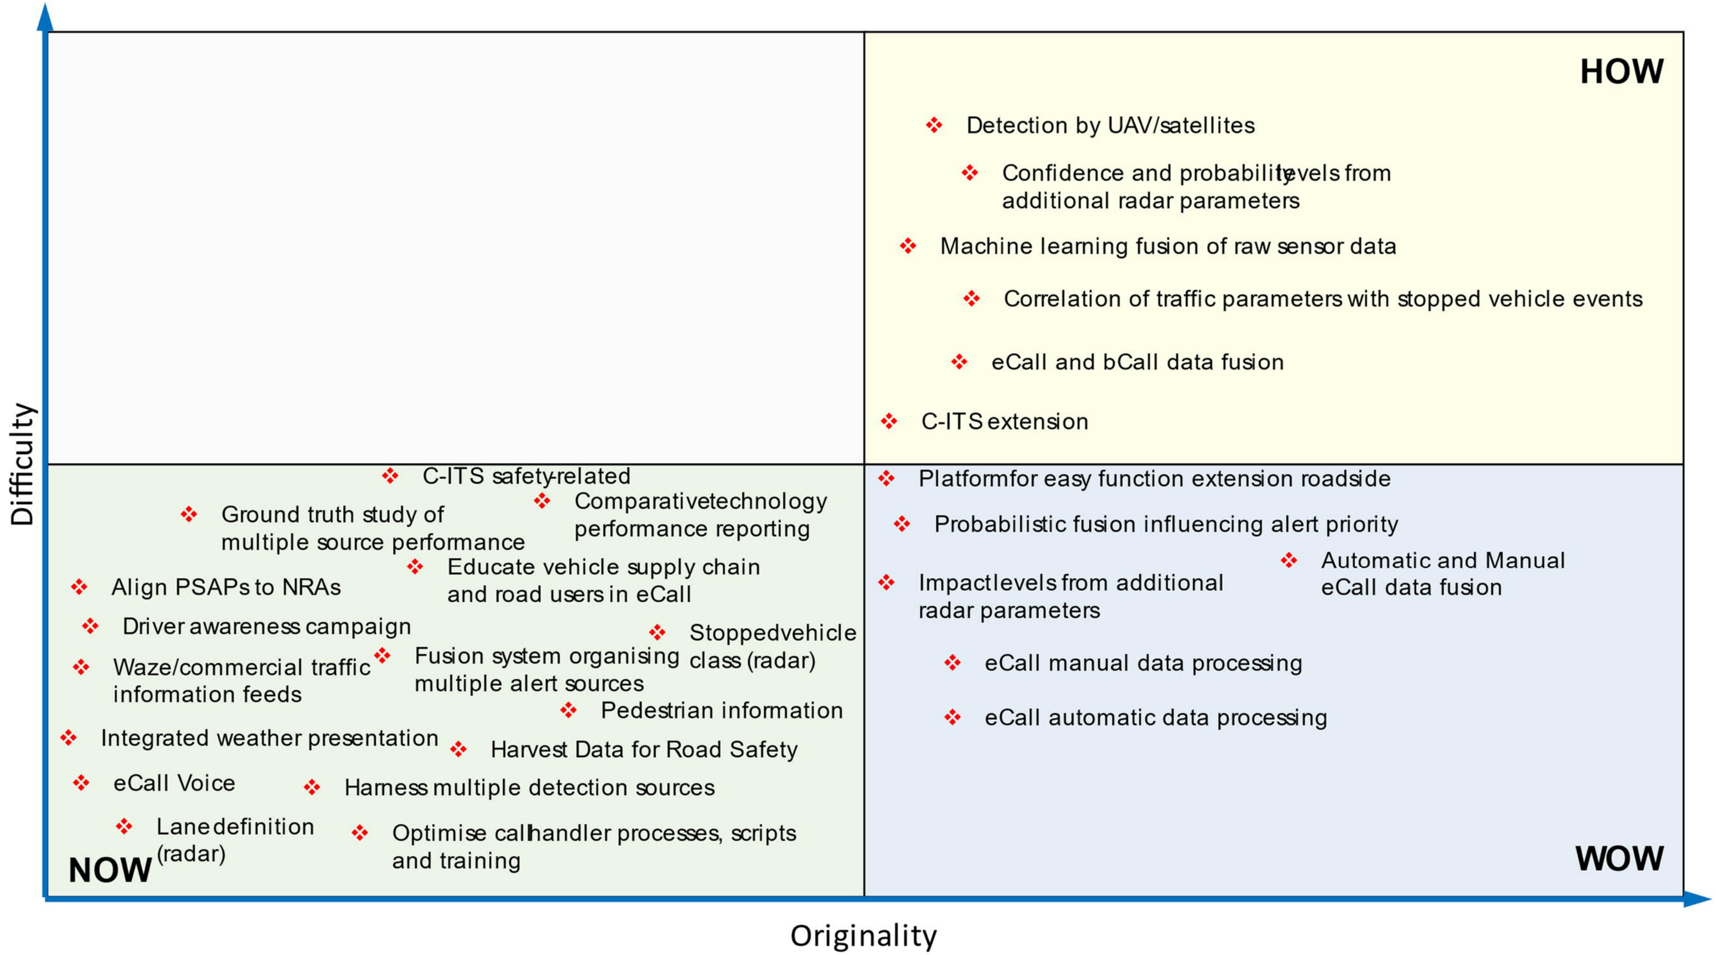

Figure 27 Combined How Now Wow matrix (positions adjusted for legibility)

8 References

Barmpounakis, E., Sauvin, G. M. and Geroliminis, N. (2020). Lane Detection and Lane-Changing. *Transportation research record*, 1-15.

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.

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.

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.

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

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

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.

National Highways (2022), ”eCall SOS”, https://nationalhighways.co.uk/road-safety/ecall/, 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.

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/

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.

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*.