CHALMERS

What are the effects of AEBS on collision avoidance?

Master's thesis in Engineering Mathematics and Computational Science and Applied Physics

LOUISE ROST
JOAKIM SÄLLBERG

Department of Mechanics and Maritime Sciences
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2019

MASTER'S THESIS IN ENGINEERING MATHEMATICS AND COMPUTATIONAL SCIENCE AND APPLIED PHYSICS

What are the effects of AEBS on collision avoidance?

LOUISE ROST
JOAKIM SÄLLBERG

Department of Mechanics and Maritime Sciences
Division of Vehicle Safety
CHALMERS UNIVERSITY OF TECHNOLOGY

Göteborg, Sweden 2019

What are the effects of AEBS on collision avoidance?
LOUISE ROST
JOAKIM SÄLLBERG

Master's thesis 2019:47
Department of Mechanics and Maritime Sciences
Division of Vehicle Safety
Chalmers University of Technology
SE-412 96 Göteborg
Sweden
Telephone: +46 (0)31-772 1000

Cover:
Collision Warning with Emergency Brake, i.e. the AEBS system in Volvo HGVs, alerts the driver when there is a risk of collision with a vehicle in front and activates the brakes if necessary. The blue light in the picture visualises the scope of the camera, while the radar is placed far down on the front of the HGV. The picture is obtianed from Volvo Trucks Images and Film Gallery (internal source).

Chalmers Reproservice
Göteborg, Sweden 2019

What are the effects of AEBS on collision avoidance?
Master's thesis in Engineering Mathematics and Computational Science and Applied Physics
LOUISE ROST
JOAKIM SÄLLBERG
Department of Mechanics and Maritime Sciences
Division of Vehicle Safety
Chalmers University of Technology

Abstract

Advanced Emergency Braking System (AEBS) is an active safety system for Heavy Goods Vehicles (HGVs) which aims to prevent rear-end collisions, i.e. when a vehicle drives into the rear of the vehicle in front. This report investigates the performance of AEBS in Volvo HGVs, and describes under what circumstances the system intervenes correctly and incorrectly respectively. Data from AEBS interventions by Volvo HGVs was analysed, and patterns of the incorrect interventions were identified. These patterns were translated into code, resulting in a program that automatically classifies the logged interventions as correct or incorrect. Different variables were investigated for the correct and incorrect interventions separately, for the purpose of finding factors that affect the performance of AEBS, i.e. under which circumstances the correct and the incorrect interventions occur. The majority of the incorrect interventions were found to be due to stationary targets, and only resulted in a short intervention with minor speed reduction. Therefore, it seems very unlikely that the incorrect interventions would cause collisions. The majority of the interventions were found to be true, and many of these interventions yielded a large speed reduction. Thus, AEBS interventions prevents many collisions.

Keywords: AEBS, collision avoidance, HGV, rear-end collision, Volvo Trucks

Acknowledgements

There are many persons who have helped us making the best of this thesis, and we are really grateful for their guidance and support during the project. We would like to thank our supervisor Ron Schindler for a great patience in answering all of our emails, and giving us feedback and constructive criticism on our report and presentation. We also want to thank our examiner András Bálint for always being available for questions, helping us with academic and administrative questions, encouraging us and giving us feedback on our work. Ron's and András' knowledge and guidance helped us find new perspectives and reach an academic height in our project.

We would also like to thank the people at Volvo GTT who have helped us with our thesis in different ways. A great thank you to Anna Theander, who has always been there for us, providing us with the material we needed, taking care of the administration and making sure that we are feeling well. Thanks to Andreas Ekfjorden for always being open for discussions and for spending time in supporting us, always with a smile on his face. We would like to give an extra thank you to Laurent Decoster who has given us the technical input needed for our thesis, guiding us with ideas and proposals for our work, and giving us clear answers and a great support. We wish you good luck with your baby! We would also like to thank Payam Chehrazi and Abhishek Karunagaran for the help with data handling and for being available for questions. Also, thanks to the other people at Volvo GTT for making us feel welcome and showing interest in our work.

Lastly, we want to thank our families and friends for your support and always being there for us, even though we have not had much time to spend with you during the spring. The encouragement to focus on our project and to believe in ourselves has been of great importance for us.

Contents

Abbreviations

ABS	Anti-lock Braking System
ACC	Adaptive Cruise Control
AEB	Autonomous Emergency Braking (passenger cars)
AEBS	Advanced Emergency Braking Systems (HGVs)
CW	Collision Warning
EB	Emergency Braking
ECU	Electronic Control Unit
FCC	False Classification Criteria
FCW	Forward Collision Warning
HGV	Heavy Goods Vehicle
TFC	True False Classification
TTC	Time to collision

Definitions

Acceleration pedal kickdown	When the driver pushes the acceleration pedal all way down, also passing a switch
False negative	A situation when AEBS should have been activated but was not
False positive	A situation when AEBS was inaccurately activated
Fusion	Data from two or more sensors are combined to avoid uncertainties in the data
Intervention	Activation of a safety system, here most often AEBS
Ghost target	When the radar or camera detects an object that does not exist due to a misinterpretation of the environment
HGV	Goods vehicle with a total weight of 3,5 metric tonnes or more [1]
Host vehicle	The HGV with the AEBS system implemented
Rear-end collision	Refers to when a vehicle crashes into rear of the vehicle travelling in front of it, in the same direction
Target	Object, e.g. a car, that is considered most critical by the system
True positive	A situation when AEBS was correctly activated

List of Figures

List of Tables

1 Introduction

Every year, around 1.4 million lives are lost due to traffic collisions [2]. World Health Organization classifies road injuries as the eighth leading cause of death in the world. The fatalities, injuries and property damage result in high societal costs, and measures are taken to reduce the number of collisions. This includes research about different accident types and their main causes, but also design of safe roads and the development of new safety systems for vehicles. One example of a relatively new safety system is Advanced Emergency Braking System (AEBS), which aims to reduce the number of collisions with a truck involved.

AEBS uses sensors to scan the road to detect potential collisions. If a potential collision is detected, the system automatically warns the driver, and if necessary it also applies the brakes of the vehicle equipped with the system. Volvo Trucks introduced AEBS in their Heavy Goods Vehicles (HGVs) in 2012, as a step in their strive towards zero serious collisions on roads, i.e. collisions leading to serious injury or death. Volvo's HGV are already equipped with a number of other safety systems such as Anti-Lock Braking System (ABS) and Electronic Stability Program (ESP), not to mention the seat belt, which all are saving numerous lives every year. Knowledge about how well a system performs, and in which situations it has effect, is relevant not only to the manufacturer but also customers and researchers. Since AEBS is a relatively new system, its real-world impact on collision avoidance is not yet known. Therefore, gaining detailed knowledge about the effects of the system is of interest.

This study is an investigation of the effects of AEBS on collision avoidance. The study is based on data analysis from real world AEBS interventions in HGVs with Volvo's AEBS system. The aim of the study is described in Section 1.1 and the limitations are described in Section 1.2.

1.1 Aim of the study

The aim of the project was to develop and implement a method to analyse and draw conclusions about under what circumstances correct interventions (true positives) and incorrect interventions (false positives) of AEBS occur, and to what extent AEBS decreases the number of rear-end collisions and reduces their severity. The project also aimed to point out connections between the false brake interventions and radar malfunction.

The project was divided into the following tasks:

- A literature review on the effects of the introduction of various accident avoidance systems on vehicles. Especially interesting is the question about if and how AEBS has been analysed before.
- Development of a method to perform an analysis of log data from equipped Volvo HGVs to make a quantitative analysis of the performance of Volvo's AEBS system. The method should be scalable to big data environments due to large data sets.
- Find quantitative results describing possible safety benefits from AEBS and find correlations of the true and false interventions.
- Understanding and describing environment limitations of the radar technology and propose a way to resolve these or suggest ways to reproduce scenarios for further development of the radar.

The following research questions were to be answered:

1. Has AEBS decreased the number of rear-end collisions?
2. Has AEBS decreased the severity of the rear-end collisions?
3. How many of the brake interventions by AEBS are true and false respectively? Is this affected by external conditions, e.g. time of day? The answer to the first part of this question cannot be published externally due to confidentiality.
4. What are the reasons of the false positives? Is there a correlation to external conditions such as time of day and geographic location?

5. Under what conditions do true positives occur? Is there a correlation to external conditions such as time of day and geographic location?

6. Can false radar targets derive from the surrounding objects or other factors, e.g. noise, other interfering signals or inappropriate system design of the radar?

1.2 Limitations of the study

The analysis that has been carried out is based on automatic braking systems in HGVs, and even more specifically Volvo HGVs. In the literature review both AEBS for HGVs and AEB for cars have been studied, but no data from AEB systems in cars has been analysed. The data sets used for the analysis contain many different variables. This analysis includes several of them, both separately and in combination with each other. However, there are many possible combinations of variables that can be investigated, but only some of them have been chosen for this analysis. When implementing the resulting programs, the aim was to make them as generic as possible and load in all relevant data, to allow for future analyses of other variables than the ones used in this thesis work.

1.3 Confidentiality

This thesis work was carried out at Volvo Group Trucks Technology and concerns the details about the performance of one of their products, which is a business secret. Therefore, several aspects of the thesis work are subjects to confidentiality and cannot be published in an official report. Thus, two versions of the report have been made – one official thesis report with general descriptions but without the details, and one company internal report with more details and scripts that have been created during the project. This is the official version of the report.

2 Background

In the EU, there were 23 900 fatal traffic collisions in 2016, resulting in 25 600 fatalities, according to the European Road Safety Observatory (ERSO) [3]. They also state that crashes in which at least one HGV was involved resulted in more than 4 000 fatalities, which is about 16% of all fatalities in traffic in the EU [4]. ERSO also states that of the fatally injured in the HGV-including collisions, 14% were HGV occupants, while 49% were car occupants, 16% pedestrians and 8% bicyclists. Thus, HGV occupant fatalities make up about 2% of all traffic collision fatalities. In other words, collisions involving HGVs are often serious or even fatal, but often other occupants are more seriously injured than those in the HGV. The large mass of the HGV is one of the main causes for this.

Of these fatal collisions in which an HGV was involved, 24% occurred on urban roads, 19% on the motorway and 56% on rural roads, while in 1% of the collisions the road type was unknown, according to the same report. Among the 570 HGV occupant fatalities, 12% were in urban areas while 88% were on rural roads or motorways [3], showing that HGV fatalities are more likely to occur on higher speed roads than on urban roads.

In 2015 a new commission regulation (No 347/2012) came into force in the EU, requiring HGVs with maximum mass exceeding 8 tonnes to have a collision warning and emergency brake system (AEBS) installed, with some exceptions [5]. In 2018, this requirement was extended to also include HGVs with a maximum mass below 8 tonnes. The tests performed on the new HGVs also have higher requirements in the second version. For example, the minimal speed reduction of the tested vehicle must be larger, and the speed of the target is higher than in the previous tests. Further information about the AEBS requirements and the regulation in general can be found in Commission Regulation (EU) No 347/2012 [5]. Similar systems already exist for cars but is then referred to as Autonomous Emergency Braking (AEB).

The main purpose of AEBS is to prevent and reduce the severity of rear-end collisions, i.e. collisions in which the front of a vehicle strikes the rear of another vehicle travelling in the same direction. Rear-end collisions are responsible for between a fourth and a third of all traffic collisions [6][7][8], making it the most common of all collision types. Rear-end collisions often result in damage to the vehicles involved and personal injuries, which in turn result in high societal costs. A rear-end collision resulting in fatality is less common, but still the rear-end collisions are to blame for around 7% of all fatal traffic collisions [6].

At Volvo Trucks, the Accident Research Team (ART) studies collisions to understand the cause and procedure. In *Volvo Trucks Safety Report 2017*, ART presents an overview of the distribution of the different types of serious collisions, i.e. collisions resulting in fatality or severe injury, where HGVs are involved. This is based on data from collisions in Europe. The authors concluded that 6.5-9.5% of serious collisions where a HGV is involved, are of the type where an HGV strikes the rear-end of a car or another HGV [9]. The same report also includes frontal collisions and collisions into the side of a car (either an oncoming car making a left turn across the HGV path or a car turning out from an intersection), and these add up to 17-29% of all serious collisions with an HGV involved. Another common collision type is when a car strikes the rear-end of an HGV, reaching 5-8%. Future AEBS/AEB systems might be able to avoid these type of collisions as well.

2.1 Functionality of AEBS

AEBS uses warnings and automatically initiates a braking to mitigate or completely avoid collisions. The system is designed to let the driver remain in charge of the vehicle as long as possible, and the emergency braking only intervenes when there is no other option left. Therefore, the system starts by warning the driver when there still is a chance to take action. The system alerts the driver if a collision is imminent by a visual and/or an audible warning. If the driver fails to respond to the warning, the system will automatically initiate the emergency braking. Section 2.1.1 and 2.1.2 describes the functionality of AEBS in Volvo HGVs and the limitations of the system.

2.1.1 AEBS in Volvo HGVs

Volvo's AEBS uses a radar and a camera to keep track of the objects ahead of the HGV [10]. A radar system can measure the distance and velocity to the tracked objects with high accuracy, but is not as accurate when it comes to determining angles or object type. A camera on the other hand, can determine angles and object type, but only give a rather rough estimation of the distance. If both the radar and the camera are tracking an object, and they agree on the values of the longitudinal and lateral position as well as velocity of the object, their information can be fused. Through the fusion, the advantages of radar and camera can be combined to acquire a clearer picture of the surroundings and thereby reduce the risk for misinterpretations. Thus, fusion makes the collected data more precise and trustworthy. Altogether, the sensor can keep track of several objects where some are radar only objects, some are camera only objects and some are fused objects. The functionality, advantages and limitations of the radar is further described in Section 2.3.

Beside keeping track of the surrounding objects, the sensor can also evaluate which object is the most critical from the AEBS system's point of view, and selects this object as *target*. Which object is considered the most critical depends on the time to collision (TTC), i.e. the estimated time it will take until the host collides with the target when assuming constant velocities or constant accelerations. The object with the lowest TTC, which is also in the path of the host vehicle, is selected as target. Only one object can be selected as target at a time. If a new object would become more critical, it will replace the previous object as the selected target. Knowing which of the tracked objects is the target is crucial when analysing logs to determine if the object detection was a true or false positive and if the brake intervention was correct or not.

The collected information from the sensor is evaluated in a so called Electronic Control Unit (ECU), which also receives information from other ECUs in the HGV, such as the velocity and yaw rate of the HGV. If the software of the control unit considers that a collision is imminent, a brake intervention is initiated. In Volvo HGVs, AEBS is a combination of a Forward Collision Warning (FCW) and an automatic emergency braking. How these are deployed depends on the traffic scenario. In most cases, the different parts of the system are activated gradually as described below and shown in Figure 2.1.

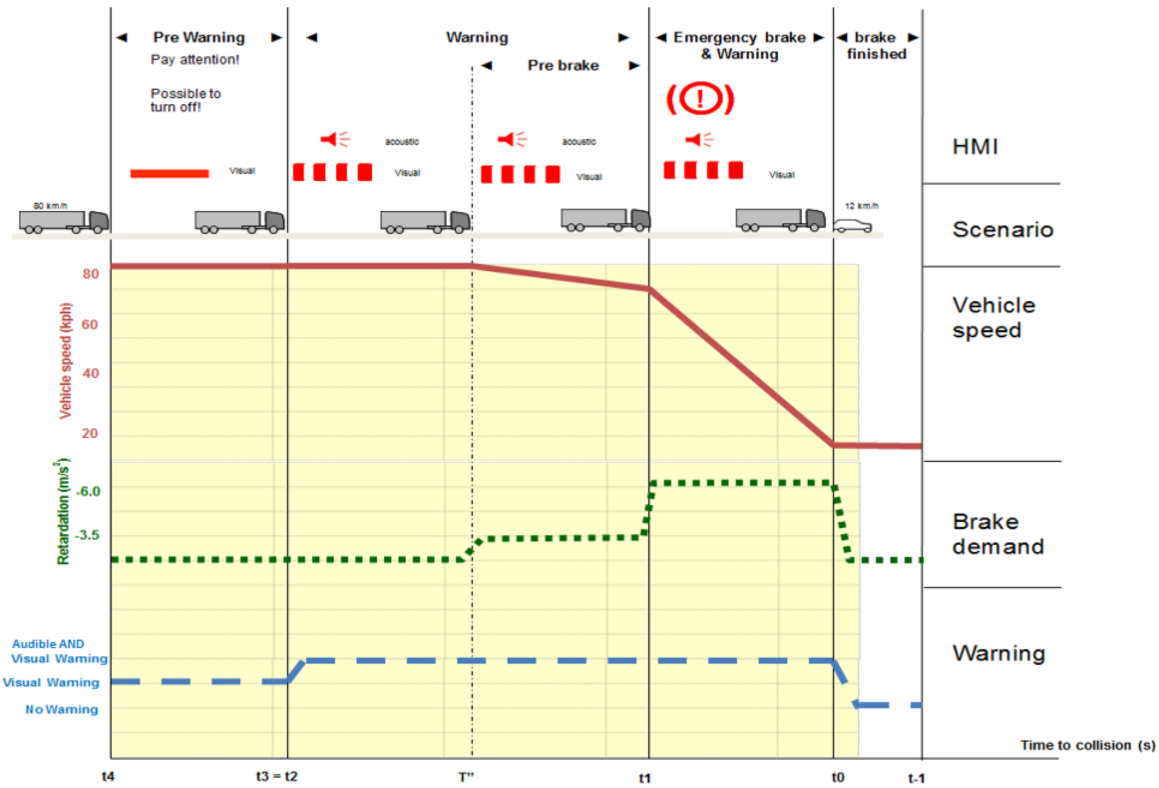

Figure 2.1: *The different steps of the AEBS intervention: pre-warning, warning, pre-brake and full brake [11].*

The first part of FCW is a visual signal to the driver of the HGV, in the form of a red LED light on the front windshield (see Figure 2.2). If the HGV keeps approaching the object ahead, a flashing light and an audible alarm signals the driver to brake. If there still is no sufficient reaction by the driver, the EB is activated. As a first step, the system initiates a light brake, called pre-brake, which both reduces the speed and triggers the braking system to be prepared for a potential full brake. Finally, if necessary, a full brake is initiated, which continues until the HGV has the same speed as the object ahead, or until a complete halt is reached. The EB is only initiated if a collision is inevitable without an immediate brake.

Figure 2.2: *Photo from the first part of an FCW intervention [12].*

Depending on how critical the scenario is, the procedure of the AEBS intervention can differ. For example, the system can immediately activate a full brake if a collision suddenly becomes inevitable, e.g. due to a cut in manoeuvre.

2.1.2 Limitations of the system

Volvo's AEBS system brakes for moving targets as well as stationary targets if the camera and radar data are fused, but not yet for Vulnerable Road Users (VRUs). The system has been introduced in three different stages, released in 2012, 2015 and 2018 respectively. The system intervenes independent of the host speed, and can completely avoid stationary targets at host speeds up to 50 km/h, 80 km/h and 90 km/h for the three versions respectively. The system is activated when the host speed is above 15 km/h for the two first versions, and 5 km/h for the last version. Moreover, for the system to be activated it is required that the HGV has a functional ABS in the tractor and also the trailer, if there is one.

There are several ways in which the system can be deactivated. First of all, there is a button that the driver can simply push to turn it off. If the sensors are misaligned or covered in some way, e.g. by snow, rain or dirt, and thus have limited functionality, the system is automatically turned off and the driver will be informed about this. The driver can also perform a so called acceleration pedal kickdown, which is when the driver gives full gas and then also an extra push to the acceleration pedal. This action deactivates the system for 15 seconds. However, if there has been a kickdown for more than one minute, the system is reactivated. Moreover, the system is deactivated if there have been three interventions from AEBS. Reactivation is performed when the HGV visits a service center. To conclude, there are a number of ways in which the system can be deactivated and thus it is possible that there are HGVs driving around without AEBS activated.

Though the advantages with AEBS are many, there are also risks with the automatic braking. For example, it could make the driver lose control of the vehicle, distract the drivers nearby by the sudden brake or possibly even cause other collisions. Overall, the advantages of AEBS still weigh much higher than these risks. The system has a few weaknesses which can cause false detections, so called ghost targets, as well as missed targets.

2.2 Previous studies on AEBS and AEB

There are some studies on the effects of AEBS, and even more about AEB, but as these are performed in different geographical regions and account for different factors, their conclusions vary and can not trivially be compared. Moreover, the studies on AEB might not be comparable with the result from this analysis due to the differences of cars and HGVs, e.g. regarding the physical attributes, the number and types of collisions, and the functionality of AEB and AEBS. Yet, some of the studies contain interesting findings relevant for the analysis performed in this thesis, and also ideas for the analysis. Some of these studies are discussed below.

2.2.1 Analyses of the effectiveness of AEBS

In a study from 2016, Grove et al. analysed 150 trucks equipped with AEBS from two external suppliers (i.e. not developed by the truck makers) [13]. The trucks drove their normal routes between 2013 and 2015, resulting in 2.5 million miles (ca 4 million kilometers) of logged driving. The two systems had generated 30 and 234 EB interventions respectively. Out of these, 8 and 1 brake interventions were false, corresponding to 27% and 0.4% of the total number of interventions respectively. The same study also included an analysis of the duration of the brake interventions, the maximal deceleration and the total velocity reduction. The analysis was made for true and false positives individually. The results showed clearly that the false positives on average have a shorter duration, lower maximal deceleration and a lower total velocity reduction. It is worth mentioning that this might only be valid for the systems used in the study and not necessarily applicable to other emergency braking systems.

One way to determine the effect of AEBS/AEB is to compare vehicles that have the system installed with models that do not, e.g. in number of rear-end collisions reported to the police or to insurance companies. In a study conducted by Isaksson-Hellman and Lindman, insurance claims from rear-end collisions with one car model only (Volvo V70 model year 2010-2015) in Sweden were studied to investigate the performance of emergency braking (AEB) in combination with collision warning, brake support and adaptive cruise control (CW + ACC). With only low-speed AEB (limited to 30 or 50 km/h depending on generation of the system), the number of rear-end collisions was reduced by 25%, while in combination with CW + ACC the reduction was 28% [14]. In another study by the same authors, with the same conditions except they studied the Volvo XC60 with the AEB not limited to low speed, they noted a 37% overall reduction of rear-end collisions [15]. The highest reduction was then among medium and high severity collisions (defined as an impact speed of 5-15 km/h and above 15 km/h respectively).

Fildes et al. summarised the results from several studies of AEB and concluded that it would reduce the number of rear-end collisions by 38% [16]. In addition they noted that there was no significant difference between high speed or low speed (above or below 60 km/h), and they recommended a widespread fitment throughout the vehicle fleet.

2.2.2 Velocity reduction and driver experience of AEBS

Fecher et al. performed a real life study in which they analysed and compared the velocity reduction in critical situations when having AEBS or AEB installed (i.e. both cars and trucks were used) and also when not having any of the two systems [8]. The test drivers' experience of the brake interventions were also analysed. The brakings were of two intensities, partial or full brake. Since the energy in a collision increases quadratically with the velocity, it is essential to reduce the velocity as much as possible if a collision is imminent. Thus, the velocity reduction was measured in the different cases, i.e. car or truck and partial or full brake. These numbers were compared to the results from test cases where no automatic braking system was activated. The speed reductions for the cars were 8-51 km/h without any braking system, 42-63 km/h with partial brake and 44-65 km/h with full brake. The corresponding speed reductions for the trucks were 15-32 km/h, 34-36 km/h and 36-40 km/h respectively. Thus, it was clearly beneficial to have an automatic braking system activated. However, the difference between partial and full brake was not noticeable for cars. In trucks, the difference was of greater importance. This motivates the importance of an automatic braking system, especially in trucks. Also, since the severity of a collision is higher when a truck is involved, it is essential to reduce the velocity as much as possible.

After the tests the drivers were questioned about their experience of the system. The majority of the test drivers considered the system to be very effective, both for cars and trucks, but the truck drivers had an even higher percentage of satisfied drivers.

The test also included some faulty activations of the systems. The questionnaires showed great annoyance of these faulty activations among the drivers. It is thus very important that AEBS works correctly. The system should brake when a collision is imminent to avoid or mitigate the collision, but faulty activations can annoy the driver and make him or her turn off the system. Thus, it is desired to maximise the number of true positives, while minimising the false positives, which is not a trivial task.

2.2.3 Combination with other safety systems

In his article "Autonomous Emergency Braking AEB (city, inter-urban)", Saadé mentions that vehicles with both AEB and FCW have a 20% higher risk of getting struck than cars without the two systems [17]. However, the risk of striking is reduced by 50%, which thus makes it convenient to have both systems installed.

Traffic situations are very complex, and there are a lot of aspects that affect the outcome of a traffic incident. Similarly, the performance of AEBS is affected by many different circumstances, such as the curvature of the road, the speed of the host vehicle as well as the target and the surrounding environment [17]. The efficiency of the system also depends on which other safety systems the car or HGV has installed, e.g. collision warning and cruise control.

Jammes et al. performed an experiment to investigate if the reaction time is affected by the use of cruise control (CC) in case of the need of an emergency brake [18]. The experiment included eleven test drivers who were told to brake at the end of each driving session. In total 20 braking tests were performed on each individual, of which 10 were driven with CC and 10 driven manually. The results from the tests with CC were compared to the ones without for each individual. The authors' concluded that the reaction time was significantly longer with the use of CC, and the time from that the driver touched the brake pedal until reaching the peak brake force was also longer. Both these factors contribute to a lengthened braking distance. Thus, the use of CC implies a need of automatic braking to prevent the increased risk of collisions that otherwise comes with the activation of CC.

2.3 Automotive radar

As previously mentioned, the AEBS system in Volvo HGVs uses a combination of radar and camera. This combination has several benefits. The radar system can measure distances with a high precision but with the disadvantage of a low angular resolution. A camera has much better angular resolution but cannot determine distances with high precision. By combining the radar and camera, the benefits of both sensor types can be used for better positioning of surrounding objects. However, some of the disadvantages of the sensors are difficult to avoid completely and they might fail in detecting objects or detect so called "ghost targets", i.e. misinterpret the surrounding and cause a false emergency braking. Even though measures are taken to reduce erroneous or missed detections, the AEBS system will not perform flawlessly always.

By looking into the principles of an automotive radar, potential weaknesses are derived. These weaknesses are discussed further and compared with the findings from the data analysis. Finally, a few test cases are proposed to test how well the radar performs in these scenarios, which hopefully can give ideas for possible improvements of the radar.

This section discusses the working principles of a radar and the physics behind it. Then radar design is described, and the radar equation is derived. The main information source for this section (except the derivation of the radar equation) is *Automotive RADAR* by H. Winner [19] if not indicated otherwise.

2.3.1 Principles of a radar

The main principle of a radar is to measure distances by transmitting electromagnetic radiation that is reflected by surrounding objects, and measure the time for the radiation to return. Since the radiation traverses through air, its velocity is approximately the same as the speed of light in vacuum. Thus the distance to the object can easily be derived by dividing the time for the radiation to travel with the light speed and divide by two, since it travels back and forth.

The radar uses radio waves, for which metal materials act as nearly perfect reflectors, but some other materials might also cause reflections. The shape of an object also plays a key role as it determines in which angle the radiation is reflected. Convex shapes have at least one side perpendicular to the propagation, and will therefore reflect back some of the radiation while flat and concave surfaces will be more dependent on how they are angled (see Figure 2.3). E.g. corners with right angles will reflect back all radiation towards the source. Sometimes, the radiation can be reflected several times, so called secondary reflections (see Figure 2.4), which can result in faulty distance and angular measurements. Primary (normal) reflections and secondary reflections might interfere with each other and sometimes the phase is shifted enough for them to cancel each other out.

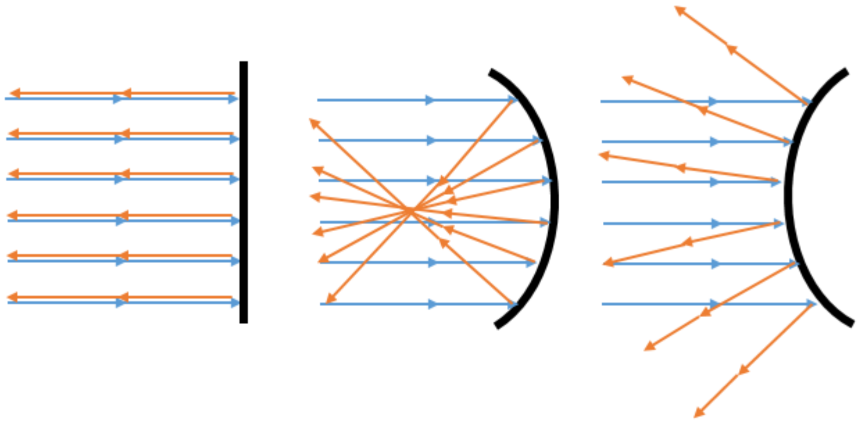

Figure 2.3: *How radiations is reflected by different shapes. The blue lines shows transmitted radiation from the radar and the orange lines show how it is reflected. A flat surface (left) reflects with an angle equal to the incident angle, making the shape visible for a radar pointing straight at it. A convex shape (middle) focuses and then disperses the radiation, causing the fraction of radiation back at the radar to vary greatly. A convex shape only disperses the light making it visible to the radar from most directions, but not at very long distances.*

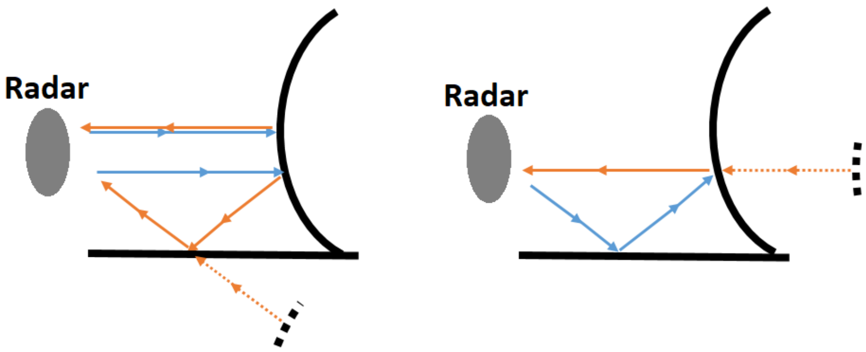

Figure 2.4: *Primary and secondary reflections, and how a radar interprets them. The blue lines shows transmitted radiation from the radar and the orange lines show how it is reflected. The solid black line shows surface of reflecting objects. In a primary reflection (upper left), the radiation traverses the shortest distance between the reflecting object and the radar. Since the radar is unable to determine if the radiation has been reflected more than once, i.e. a secondary reflection (lower left and lower right), its interpretation of the environment becomes incorrect, either in terms of distance or/and angle. The dotted orange and grey lines shows the radar's interpretation of radiation and reflecting objects respectively.*

Besides reflection, other physical phenomena such as refraction and attenuation can occur. Attenuation is a damping of the signal and does not cause faulty detections but can limit the range of the radar. The attenuation is increased in inclement weather such as rain, snow and hail. Refraction means that the propagation direction of the radiation changes when the radiation passes through the interface between two materials. Refraction of the radiation can for example occur when radiation passes through the interface between water and air, and if the radar is covered by an uneven layer of water it can result in a so called lensing effect, where the radiation is focused or dispersed. Thus, refraction can lead to reduced angular precision and distorted measurements.

In addition to ranging, a radar can estimate the size of an object by measuring how much radiation is reflected by an object (and by that also determine the object type). Besides, the Doppler effect can be used to determine the velocity of the object from the change in frequency.

2.3.2 Design of a radar

In order for a radar to achieve accurate measurements of the position and velocity of numerous objects simultaneously it must be properly designed. This includes beamforming, which is explained below, but also signal processing, the casing around the radar and the placement on the vehicle.

Beamforming means that the radiation is directed instead of spreading isotropically (i.e. equal power density in all directions). This is usually carried out with an antenna array, whose electric fields can interfere constructively or destructively in different directions depending on distance and phase. In the far-field, i.e. at a large distance from the antenna compared to the wavelength, the radiation beams form so called lobes. It is possible to steer a large proportion of the radiation in a certain direction and achieve a so called main lobe, but there will also be minor side lobes. The power density (energy) achieved in a certain direction is often divided with power density that would be achieved if the emitted power was isotropically distributed in all directions, and this is called the antenna gain. The maximal gain is found in the direction of the main lobe, and this is called directivity. A high directivity and small side lobes is highly desired in an automotive radar antenna.

If the transmitting antenna array also is used as receiver, the same radiation pattern will be valid for the reflected signals, resulting in high gain for signals received from the direction of the main and side lobes.

2.3.3 The radar equation

To derive the ranging capacity of a radar one can start by considering a spherically radiating antenna. The transmitted power P_t is then distributed over the surface of a sphere. At a distance r the power per unit area is then $P_t/(4\pi r^2)$. When the radiation reaches an object, a proportion of the radiated power is reflected. This proportion depends on the material, geometry and orientation of the object. Since the radiation is reflected in different directions in a complex manner, it is common practice to simplify by replacing the object with a uniformly reflecting metal sphere that would yield the same power reflected back to the radar antenna. The power is then uniformly distributed, and at the antenna the power density has decreased by a factor of $4\pi r^2$.

Using this simplification, the power density reaching such a sphere is dependant on the cross section area of the sphere, often referred to as the radar cross section (RCS) and denoted as σ, which might differ a lot from the size of the actual object. For example, a human has a radar cross section of around 1 m^2 and typical cars around 100 m^2. A perfectly reflecting flat metal surface of 1 m^2 perpendicular to the direction of propagation can have a radar cross section of nearly 10^6 m^2. However, if the plate is placed 60 m from an automotive radar and oriented just 1 degree from perpendicular, there will be practically no power density reflected back to the radar.

When propagating, the radiation is also attenuated exponentially, which is described with an attenuation constant k that depends on the signal frequency and propagation medium. If the attenuation is expressed in dB/km, a distance r (measured in meters) will result in a power decrease according to $10^{-kr/1000}$. For signals travelling to an object and back, the attenuation will cause a power decrease by a factor of $10^{-2kr/1000}$. Finally, the dimension A and efficiency K_A of the antenna will determine how much power the radar can pick up. Altogether, the relation between the received power P_r and the transmitted power can now be written as

$$P_r = \frac{P_t \sigma 10^{-2kr/1000} A K_A}{(4\pi)^2 r^4}, \tag{2.1}$$

An actual radar often uses beamforming to increase the power gain G_t in a certain direction. The reflected signals will also experience gain according to [20]

$$AK_A = \frac{G_r \lambda^2}{4\pi}. \tag{2.2}$$

If the same antenna is used for transmitting and receiving, one can replace $G_r G_t$ with G^2. Finally the radar equation is derived.

$$P_r = \frac{P_t 10^{-2kr/1000} \sigma \lambda^2 G^2}{(4\pi)^3 r^4} \tag{2.3}$$

With a certain output power P_t, the radar equation yields information about the received power depending on radar cross section and distance. To detect an object, the received signals must be higher than the noise in the receiver. This is often measured in signal-to-noise ratio (SNR), and a typical automotive radar has an SNR requirement of 6-10 dB, i.e. around 4-10 times higher.

Modern radars also utilise other techniques to increase the capacity. On example is pulse compression, which focuses the energy of a signal to a shorter period of time, and thus makes it exceed the required SNR.

2.3.4 Choice of frequency and interference from other sources

Currently there are four frequency bands allowed for automotive radars, one at 24-24.25 GHz, one at 76-77 GHz, one at 77-81 GHz and an ultrawide band at 21.65-26.65 GHz. At these frequencies the attenuation is around 1 dB/km, which is low compared to for example a frequency of 60 GHz for which the attenuation is up to 15 dB/km. Therefore, the attenuation does not cause any major problems for objects at distances relevant for AEBS.

Since most automotive radars use the same frequency bands, there could potentially be interference from signals deriving from different vehicles, which could disturb the systems. However, since the radiation travels with lightspeed, it only takes around 1 μs for light to travel 200 m forth and back. Therefore, most signals from other sources can be suppressed by only receiving signals in a short interval after the transmission. Since measurements are taken often, it would not be critical if one measurement is affected by interference. In addition, the measurements are not made with the exact same time spacing since this would allow for other sources to repeatedly interfere. Instead, the measurements are made at random intervals.

3 Method

With data from a large number of brake interventions, conclusions can be drawn about the effects of AEBS. However, there is a risk that the AEBS system misinterprets the surrounding environment of the HGV and intervenes when it should not. To draw correct conclusions about in which scenarios the AEBS avoids or mitigates a collision and what causes false interventions, the data was divided into true and false interventions before any further analysis was made. To manually analyse logs and classify them as true or false positives is possible but very time consuming and thus not suitable for large data sets. Therefore, a script was developed to automatically classify logs as true or false. This procedure is henceforth referred to as the true false classification (TFC).

To determine if an intervention is true or false is a complicated task, and the data has a limited number of variables. Therefore it was assumed that the TFC script would require continuous testing and improving and thus a method was formed that quickly yields feedback on how well the script performs and the information needed to easily find where improvements are needed. More information about the data sets, the data extraction and the development of the TFC can be found in Sections 3.1 and 3.2.

The output from the TFC, i.e. the classification information of the data logs, was later used when importing the log data into MATLAB for a deeper analysis of the true and false interventions respectively. This analysis consisted of a data preparation, and calculations and graphical visualisations of different variables, and comparisons of the different data sets. The deeper analysis is described further in Section 3.3. For some of the logs classified as false, the environments where the interventions occurred were analysed to understand what causes ghost targets for the radar, which is described in Section 3.4.

3.1 The data

The data used for the analysis was provided by Volvo Trucks, Each log contains data from an AEBS intervention, either only FCW or both FCW and a brake intervention. The data was of two different logging types, which resulted in the data sets being structured in different ways, and stored either in Excel or MAT-files. Thus, different methods were required to extract the data and make it well-shaped for the analysis. In this section, the provided data sets are introduced and the data extraction is explained.

3.1.1 Data logging

Two types of data logging have been used to create the data sets used in this study. For simplicity, the logging types are referred to as α and β. Depending on which type of data logging is used, the data is stored in different ways. If logging type α is used, the log from an AEBS or FCW intervention will contain a number of variables, including a snapshot of variables in the moment of the intervention, and also a log sequence covering three seconds before and after the start of the intervention, with a time resolution of 0.2 seconds. The snapshot contains the GPS position of the host HGV and the status of different systems that affect the AEBS intervention, e.g. if a trailer is connected and if the ABS is activated in the tractor and trailer. The log sequence contains the velocity and acceleration of the host HGV, longitudinal and lateral distance and velocity of the target, the status of FCW and AEBS etc.

The extraction of the stored data can be done when the HGV is at a service center, and the data can further be transferred to Volvo. It is also possible to perform a wireless transfer of data when the HGV is not at a service stop. There is however a risk that the data retrieval or transfer fails. This can cause data from some interventions to be missing, incomplete, corrupt or duplicated, which needs to be considered before the data is used for analysis.

The logging type β is also known as MLog, which is a logging equipment that logs signals sent over the CAN network. This means only data that is sent between ECUs can be logged while the internal variables in each ECU are unavailable. The data is organised in a different way compared to the data logged by α, since signals are sent at different times and in many cases might need interpolation or extrapolation to organise the data in an appropriate format. The MLog data contains more information about an intervention than the α

data. Only a handful of Volvo HGVs used in field tests are equipped with the MLog system. The driving was carried out by customers driving their normal routes, i.e. naturalistic driving. In that sense, the field test data is a good estimate of real world traffic.

3.1.2 Data sets

The provided data consists of four different data sets, denoted by A, B, C and D. The two smallest, A and B, use a logging type denoted as α and are stored in Excel files. The difference between these two data sets is the geographic location, and thus also the infrastructure. The infrastructure in the geographic location of data set A is known to be more difficult for the system, and the data set is thus not representative for how well the system performs in general. These difficulties were considered to be useful in the analysis of the radar. Data set B is considered to be more representative for the performance of the system in general, but is small in size. The logs in data set C are collected with MLog, here referred to as β, and are stored in MAT-files. These HGVs are equipped with a newer AEBS system and a different log equipment than those in data set A and B. Data set D is very large, and is logged with the logging type referred to as α. These HGVs are equipped with either the older or the newer AEBS system from Volvo, referred to as system versions 1 and 2. This data contains the same variables as data sets A and B, but does not look the same initially, requiring an extra step in the data extraction (described in Section 3.1.3). Since data set D in general contains more recent logs than the other data sets, it uses a newer system software version, which probably will yield more up-to-date numbers of the system performance. See Table 3.1 for an overview of the four data sets.

Table 3.1: *Overview of the four data sets used in the analysis.*

	Data set A	**Data set B**	**Data set C**	**Data set D**
Logging type	α	α	β	α
System version	1	1	2	1 & 2
Size	Small	Small	Medium	Very large
File format	Excel	Excel	MAT	Excel (encoded)

3.1.3 Data extraction

The Excel files contain information about the host and the selected target with a lower fixed sampling rate than the MAT-files. The different variables in a log are all sampled in the same points in time. All logs contain the same amount of variables, and thus the same amount of data.

The MAT-files contain information about the host and multiple objects that are being tracked, sampled with a higher variable sampling rate and covering a longer time span than the Excel files. Thus, the MAT-files contain a larger amount of data. Each variable is logged with its own timesteps, which can vary in length, i.e. there is no fixed timestep. Instead there is a corresponding time vector for each variable.

The scripts for the data extractions from the Excel and MAT-files, as well as the TFC, were written in Python. In Python, a dictionary is an unordered data structure which contains keys and belonging values. A dictionary template was used to define a structure of the data that was to be extracted. The names of the logged variables were set as keys, with their belonging values left blank to be filled during the extraction. The variable values in data sets A and B are initially encoded in a hexadecimal string, and in data set D in a base64 string. To extract the data in data set D, the base64 string is first converted to a hexadecimal string. Then, the data extraction procedure can follow the same algorithm for data sets A, B and D. The hexadecimal string is decoded to decimal values of the variables, which is inserted into the dictionary until the dictionary is filled with values for each variable in each timestep. This dictionary was to be used in the TFC. Pseudo-code for the data extraction from the Excel files is found in Algorithm 1.

initialisation of *dictionary* to store variable values
for each *variable* defined in *dictionary template* **do**
  read and convert value of current *variable* in each timestep
  save values of *variable* in *dictionary*
end
return *dictionary*

Algorithm 1: *Pseudo-code for data extraction from Excel files.*

The overall picture of the data extraction from the MAT-files looks similar to the data extraction from the Excel files, but is actually more complicated. Since the MAT-files contain more variables and more timesteps, the data needs to be converted such that it can be used in the same manner as the Excel data. First a dictionary is initialised, containing almost the same variables as the dictionary used for the Excel files, with the exception that some variables are not included here since they are not needed in the TFC.

The next step is to extract the variable values from the MAT-files. The variables related to the host vehicle have a corresponding variable in the MAT-files, and can thus be trivially extracted from the MAT-file. The remaining variables, i.e. the ones related to the target, are more complicated to extract due to the different structure of the data and the need to distinguish the target from all other objects before the extraction. There is a variable defining which of the objects, defined by unique object IDs, is the target in each timestep. Moreover, the sensor has a certain number of "tracks", and in each of these, one object can be tracked. The sensor tracks are given one number each, i.e. a sensor track ID. For each sensor track, there is a variable defining which object (defined by the object ID) it is tracking in each timestep. Thus, with these variables, the sensor track logging the current target can be found. When the sensor track logging the current target has been found, variable values, such as velocity and distance, can be extracted. Moreover, for some of these variables there are two alternatives for the same sensor track, depending on if the tracked object is stationary or moving. Both the target and which sensor track is used to track the target can change during the duration of the log, making the extraction even more complicated. To conclude, for each variable related to the target, it is necessary to find which object is the current target, which sensor track is used for the current target, and for some variables if the target is stationary or moving. One and the same variable in the MAT-files can have one variable for each sensor track, and sometimes even twice as many due to the stationary/moving classification.

Figure 3.1 shows an example of what the data extraction procedure from a MAT-file can look like. In the first four timesteps, the object with ID 10 is the target, while in the three following timesteps, object with ID 7 is the target. We find that object 10 is logged in sensor track 1, while object 7 is logged in sensor track 4 in each corresponding timestep. Once the sensor track is defined, the variable values can be extracted. Here the speed values of the target in each timestep is extracted. Thus, to find the speed of object 10 in the first four timesteps, we look in the array defining the speed of sensor track 1, and similarly we look at array "speed 4" to find the speed of object 7 in the three following timesteps. In this way the "target speed" can be defined, regardless if the target is switched or not. Only the variable values of the current target is of interest since the target is the object that will possibly trigger a brake intervention in the HGV.

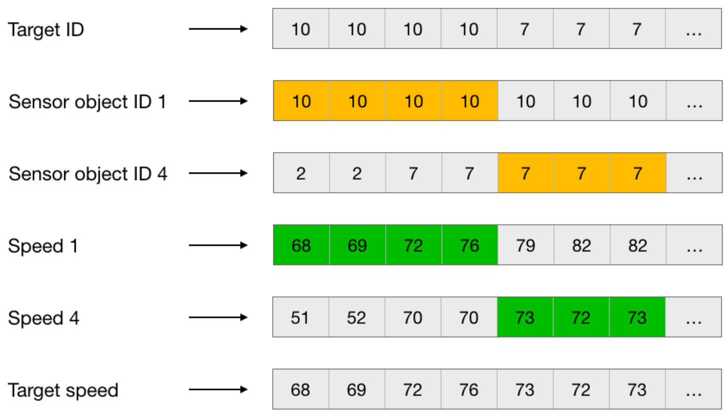

Figure 3.1: *Simplified example of extraction of the speed of the target from MAT-files.*

The example shown in Figure 3.1 is simplified since the length of the timestep is not fixed in the log. The inaccuracy of the timesteps of the different variables makes the extraction even more complicated. To solve this problem, an interpolation is performed on each extracted variable value. A time array with the same timesteps as the Excel files is used for the interpolation. When the interpolation has been performed on the extracted variable values, they can be inserted into the dictionary. Finally, the dictionary will have the same format as the dictionary resulting from the data extraction from the Excel files, and thus the TFC can be used similarly for all data sets.

initialisation of *dictionary* to store variable values
definition of *time array* to use for interpolation
for each *variable* defined in *dictionary* **do**
  read value of current *variable* in every timestep
  interpolate values of current *variable* according to *time array*
  save values of *variable* in *dictionary*
end
return *dictionary*

Algorithm 2: *Pseudo-code for data extraction from MAT-files.*

3.2 True false classification

When the data has been extracted from the logs, they are to be classified as true or false positives before the further analysis is done. A program that performs this classification was implemented, and named True False Classification (TFC). Most of the data logs contain a brake intervention, but some only contain an intervention from FCW. Since the analysis is only to be performed on the brake interventions, the logs are first checked if they contain a brake intervention or if they are so called ”FCW only”. Each log that contains a brake intervention is then classified as true, false or possibly false (i.e. a grey zone containing interventions that are difficult to classify). This section contains a description of how the TFC was implemented.

To start with, data sets A and B, containing true and false brake interventions as well as FCW interventions, were provided for an initial data analysis. Some of these logs had manually been classified as true or false positives by an experienced Volvo employee. This manual work is very time consuming, and thus it would be unreasonable to perform this kind of analysis on larger data sets. Therefore, a program to perform an automatic classification was to be implemented, i.e. the TFC.

The manually classified subset of data sets A and B were manually analysed to find patterns of false positives, and how these could be distinguished from the true positives. To do this, an understanding of the actual traffic scenario, i.e. how the host and target are moving, was needed. Thus, a MATLAB script that extracts data from the log file and visualises the trajectories of both the host and the target was created. Examples of visualisations of a true and a false intervention are shown in Figure 3.2a and 3.2b respectively. In the true intervention, the same object is selected as target throughout the whole log. The host approaches the target and gets too close, resulting in an intervention from AEBS. In the false intervention the target appears right in front of the host and is only tracked for a short period of time, which one can see by the small number of logged positions of the target. If it would have been an actual vehicle that was tracked at these positions, it would most likely not appear and disappear that sudden.

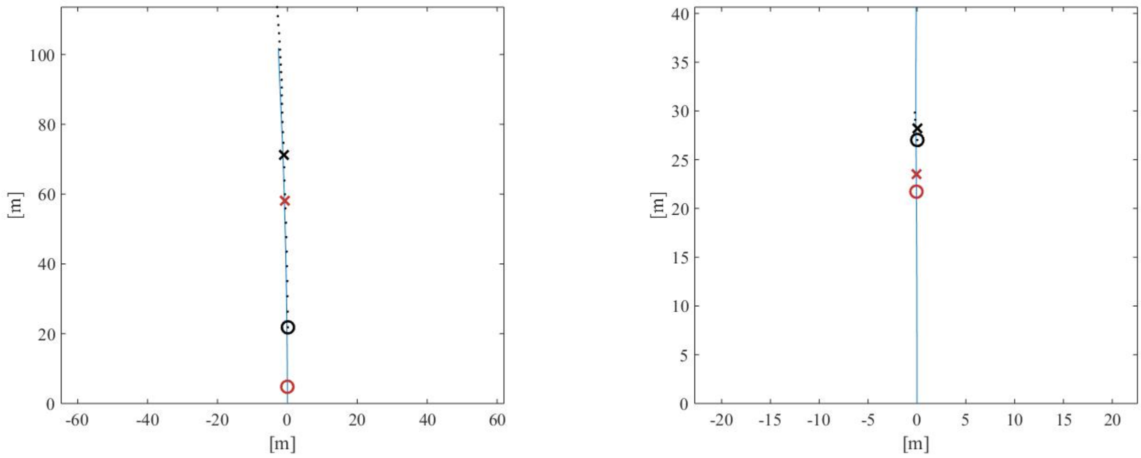

(a) *True intervention* (b) *False intervention*

Figure 3.2: *Visualisations of the trajectories of the host (blue line) and the target (black dots). The red and black circles are the positions of the host and the target respectively, when the target is first detected, whereas the crosses correspond to the start of the brake intervention.*

Using the trajectory script and the variable values in the logs, each log that had been manually classified as a false positive by the Volvo employee was analysed and the findings were noted in a structured manner. While performing this analysis, all ideas of possible ways to distinguish between false and true positives were noted in conjunction with the notes about the other findings. A few true positive logs were also analysed to better understand the differences between true and false positives.

The findings from the manual analysis of the logs were reviewed and reformulated into criteria for false positives that can be implemented in code. These so called false classification criteria (see detailed explanations in Section 4.1) were based on assumptions and estimations of the traffic scenario and the driver's response. Thresholds for different variables were set based on the observations of the true and false positives. Due to the complexity of traffic scenarios, there are no threshold values that always yield a perfect classification. However, after having iteratively tested the results from the implemented false classification criteria, the TFC script was found to work in a satisfactory way, i.e. the logs were classified similarly to the manual classification performed by the Volvo employee. The output from the TFC was formatted to be appropriate for both testing and further analyses (see Table 3.2).

The false classification criteria and the thresholds in the TFC were at this stage based on a part of data sets A and B, and were not strictly defined. Thus, the script was now run on whole data sets A and B, and the result was analysed. The thresholds were tuned based on the analysis of the result. However, there is a trade-off between not detecting a false positive and misclassifying a true positive as a false one. In most cases it was thus not possible to strictly define thresholds to define the logs as true or false interventions, since the false classification criteria with certain thresholds could be fulfilled by both true and false positives. To avoid misclassification of the interventions, a grey zone called "possibly false" was added to the script, defining an interval which could contain both true and false positives. With this grey zone, the risk to misclassify the logs could be decreased. Only the true and false positives were used for the further analysis.

It is of great importance that the TFC performs well since the true and false interventions are to be analysed separately in the main analysis and the results are to be compared. Therefore, a lot of testing and manual analysis was required to validate the performance of the script, that the false classification criteria were valid and that the thresholds were well-defined.

To facilitate the main analysis, the output was formatted such that it can easily filter logs based on classification, the different false classification criteria, as well as variable values such as the host velocity, target velocity, relative velocity, duration of the intervention etc. The output is formatted as an Excel file with the structure shown in Table 3.2.

Table 3.2: *Illustration of the output from the true false classification.*

Log	Classification	False criteria	Possibly false criteria	Host velocity [km/h]	Target velocity [km/h]	Duration of intervention [s]	...
File path 1	FALSE	3	1, 4	25.4	20.4	0.4	
File path 2	TRUE			17.3	2.5	0.2	
File path 3	FCW only						
File path 4	TRUE			75.2	40.2	1.6	
File path 5	Possibly FALSE		2	19.2	0.2	0.2	
File path 6	FALSE	3		20.0	4.5	0.8	
...							

To conclude, in the TFC each log that contains a brake intervention is investigated and categorised. There are three categories; false, possibly false and true. First the criteria for the log to be considered as a possibly false intervention are checked. If the log does not fulfil any of the criteria, the logged intervention is considered to be true. If the log does fulfil one or several of the criteria, the log is checked once more, this time with stricter values of the thresholds, to find if it can be considered as definitely false.

The file paths of the logs and their corresponding classification (true, false or possibly false) are written to an Excel document. If the log is false or possibly false, the fulfilled criteria is also defined in the Excel document. This Excel document can then be used in the further analysis.

3.3 The difference between true and false interventions

The output file from the TFC is used to divide the logs into one group of true interventions and one of false, to further analyse the circumstances of the interventions in the two groups respectively. The data from the logs was read and then pre-processed before any further analysis was made. The pre-processing included a check for duplicates and erroneous logs. Duplicates are found by comparing if different logs have occurred at the same time and place, using the time and GPS coordinates of the intervention. The check for erroneous logs consisted of checking that the data is structured in the correct way for the scripts to automatically read in the data. If the data is structured correctly, it is also checked for faulty values (NaN, erroneous decoding when switching format etc.).

After the pre-processing, the actual analysis started, where different variables are investigated and compared between the groups of true and false interventions, and also for the different data sets. To find under what circumstances the true and false interventions occur, different variables are to be investigated for the two groups of logs respectively. Also combinations of variables could be of interest. The distribution of the values of the chosen variables are plotted as histograms or pie charts. This is done for true and false interventions separately to find if the result differs. If there is a noticeable difference in the investigated variable, further investigations can be of interest. By inspecting the histograms and pie charts, conclusions about true and false interventions can be drawn, and similarities and dissimilarities can be detected and further investigated. The environment where true interventions have occurred were also analysed using Google's Street View to better understand which environments the AEBS system has the biggest impact.

3.4 Radar analysis

To find potential weaknesses of the radar, a literature study was made to understand the technology, which is summarized in Section 2.3. From the knowledge gathered during this study, objects and environments that potentially can cause problems are derived, and then compared to environments where false interventions actually have occurred.

If something in the environment has caused ghost targets, it is possible that it has occurred more than once. Therefore, all the logs that were classified as false or possibly false were gathered to compare the GPS coordinates of the false interventions. If the longitudinal and latitudinal distances between two interventions were less than certain threshold values, the logs were considered to occur at the same location. For the latitude coordinate, the threshold was set to 0.001 degree, which corresponds to around 110 meters. However, for the longitude coordinate it is different, since the coordinates correspond to different distances depending on the distance from the equator. Therefore, the threshold had to be scaled with a factor of $1/\cos(latitude)$, to make it equal to the latitude threshold.

These locations were later investigated further using Google's Street View, to find objects or patterns in the environment that are likely to cause the ghost targets. For this purpose, the coordinates were written to a KML-file, which is a file format created by Google to save GPS positions for an easy visualisation in Google Earth. In addition, several locations where a true intervention had occurred were investigated to compare with the locations where multiple false interventions had occurred.

4 Results

This section describes the result of the manual analysis of true and false interventions, in the form of a TFC program with its false classification criteria and information about its efficiency. Distributions of different variables related to the intervention scenarios are then presented as histograms, pie charts and explanations of discoveries about what characterises true and false interventions in Sections 4.2-4.12. Due to the small sizes of data sets A and B, the histograms tend to be irregular, especially the histograms corresponding to the false interventions in data set B. In contrast, data set D contains a very large number of interventions and hence its related histograms are more smooth.

An analysis of false interventions possibly caused by ghost targets is presented, with the aim to describe possible weaknesses of an automotive radar. Some examples of real road environments that are likely to have caused false interventions are shown. Finally, a short description of the road environment of true interventions is presented, to find at what road types the system makes the biggest difference.

4.1 True false classification

In the manual analysis of interventions, four false classification criteria (FCC) were found to catch the false interventions. These criteria will henceforth be referred to as false classification criteria 1, 2, 3 and 4 respectively, abbreviated FCC1, FCC2, FCC3 and FCC4. A description of each criterion is found in Table 4.1.

Table 4.1: *Descriptions of the false classification criteria.*

False classification criteria	Description of criteria
FCC1	**Short time as target:** If an object is selected as target for a short period of time, it is likely to be a faulty detection. The short time as target is a typical behaviour of ghost targets, which is the kind of false positives this criterion mainly detects in the TFC. Also, inaccurate measuring could result in the object being target for only a short time. For example, if an object is faulty selected as target due to erroneous measuring and a brake intervention is initiated, but the measuring is soon updated and the brake intervention is terminated, it is likely that the object was only target for a short period of time.
FCC2	**Acceleration pedal kickdown:** An acceleration pedal kickdown is an action that the driver is well aware of that he/she is performing. Thus, if a brake intervention is terminated due to an acceleration pedal kickdown by the driver, it is most likely that the driver intends to interrupt the brake intervention. This criterion is thus based on driver behaviour and the concrete cause is not known.
FCC3	**Acceleration no brake, excluding cut out scenarios:** The third false classification criterion is partly based on driver behaviour. In true positive interventions, the driver usually manually brakes, either before the brake intervention is initiated or when he/she notices the reason for the brake intervention. Thus, if the driver accelerates and does not brake noticeably during the intervention, it is likely to be a false positive. However, there are a couple of exceptions. One case is the so called "cut out scenario", i.e. when the target vehicle switches lane out of the lane where the host is driving or completely turns off the road. If an intervention is initiated in such a scenario, the intervention is considered to be a true positive despite the acceleration of the host vehicle. In these scenarios, the host driver can predict that the target vehicle will turn off, and accelerates to overtake the vehicle. From a system point of view, the intervention is still correct, as the host comes too close to the target.
FCC4	**Target misclassified:** When the radar and camera are not fused, the data is not as accurate and reliable, and therefore, the system should not brake for stationary targets. If there is no fusion, the longitudinal speed of the target is low, and the driver does not brake noticeably, it is likely to be a stationary target that the sensors interpret as a moving target. Logs fulfilling these criteria are considered to be false positives.

The four false classification criteria were implemented into the TFC. The result from running the TFC on data sets A, B, C and D is confidential information to Volvo, and can thus not be included in this report. The fraction of false interventions fulfilling each FCC, with respect to the total number of false interventions in each data set, is visualised in Figure 4.1. Note that an intervention can fulfil more than one FCC, and thus the fractions of false interventions fulfilling each FCC will sum up to more than 100% for the different data sets.

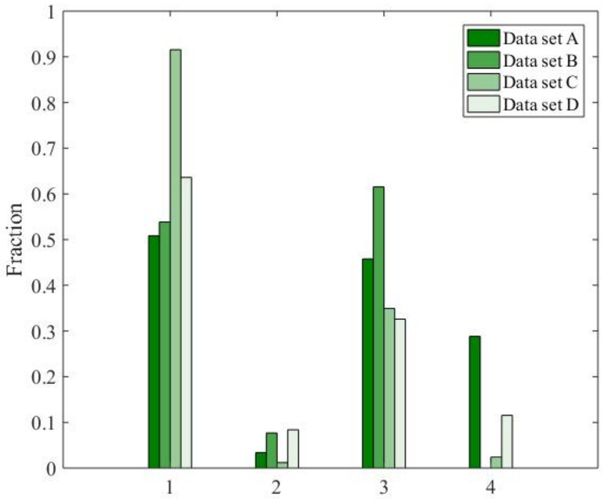

Figure 4.1: *Distribution of false interventions fulfilling each false classification criteria*

The thresholds in the FCC were set such that the categories true and false only should contain true and false logs respectively, and the third category, possibly false, is supposed to catch all the remaining logs. With this classification framework, the categories can be used to draw conclusions about the performance of the system. The actual fraction of false interventions should reasonably be at least equal to the fraction of logs classified as false by the TFC, and as most be equal to the fraction of false and possibly false logs together (based on the classification from the TFC). The actual fraction of true interventions can thus be found analogously. The lowest fraction of false logs was found in data set D, which to the most part contains recent interventions, and thus a newer system software than the other data sets in general. The highest fraction is found in data set A. This is likely due to the infrastructure in the geographic area of data set A, which is believed to differ from the infrastructure in the other data sets. The actual numbers are left out due to confidentiality.

Short time as target (FCC1) is the most frequent reason for a false positive (see Figure 4.1). FCC1 is often combined with FCC3, i.e. acceleration no brake. This combination is an indication that the target is possibly a ghost target. Acceleration pedal kickdown (FCC2) is the least common of the four FCC. However, since the driver is well aware of if (s)he performs an acceleration pedal kickdown, an intervention which is interrupted by a kickdown is very likely to be a false intervention. Thus, the criterion is justified. The existence of misclassified targets (FCC4) is also unusual in data sets B, C and D. Nevertheless, in data set A this criterion is fulfilled by a notable fraction of the interventions. Once again, the infrastructure is believed to be the cause for this difference.

Table 4.2: *Table showing the average run time of the TFC for each data set, based on 10 runs.*

	A	**B**	**C**	**D**
Run time (s)	75.8	29.0	245.5	65.9

In Table 4.2, the average run time of the TFC for the four data sets is shown. The data extraction is identical for data sets A and B, and in both cases one Excel file is read for each log. The file paths of the logs in data set C are also retrieved from an Excel file, and computations are required to extract the data (see Section 3.1.3), which lengthens the run time. The extraction from data set D only requires reading from an Excel file twice for the whole data set. The reading of Excel files is mainly what lengthens the run time of the TFC, and explains why data set D has a relatively low run time even though it contains many more logs.

4.2 Time to collision

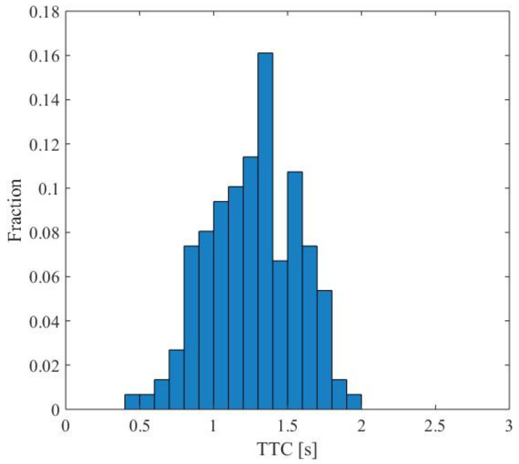

(a) *True interventions in data set A*

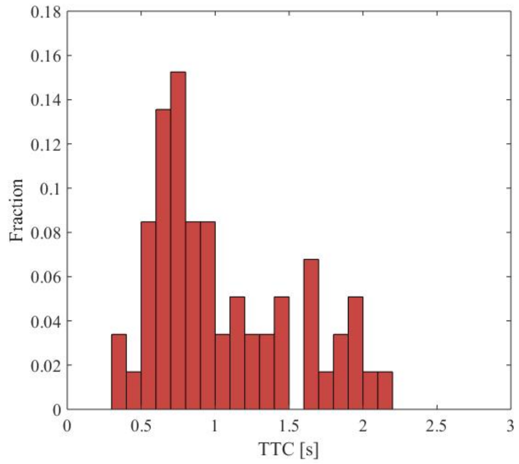

(b) *False interventions in data set A*

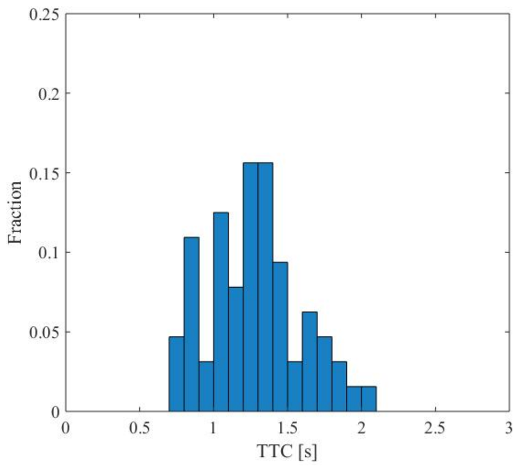

(c) *True interventions in data set B*

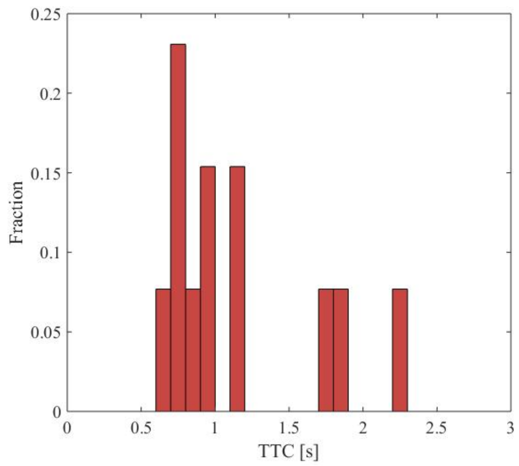

(d) *False interventions in data set B*

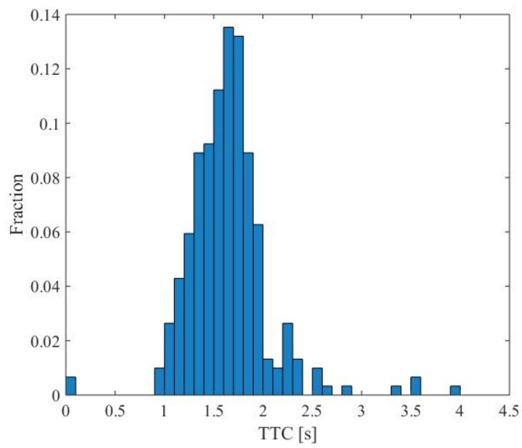

(e) *True interventions in data set C*

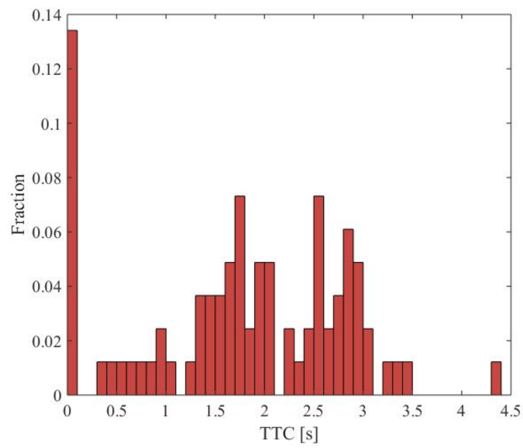

(f) *False interventions in data set C*

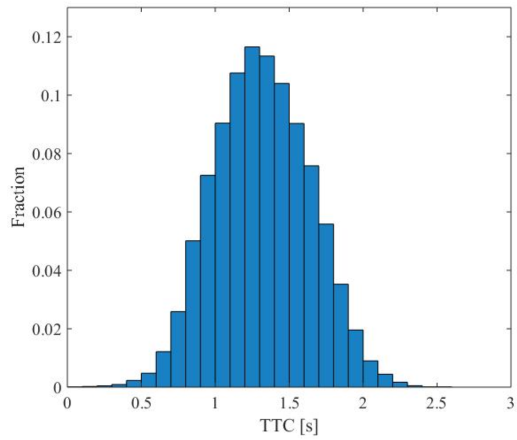

(g) *True interventions in data set D*

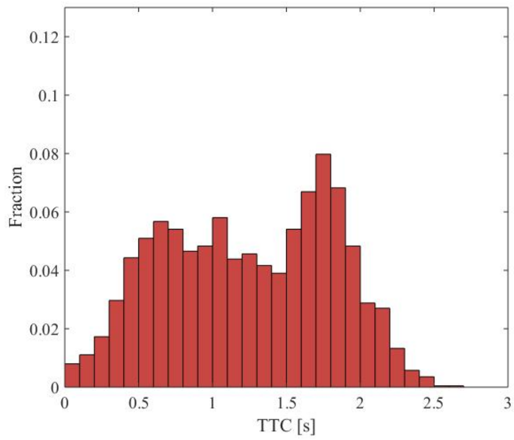

(h) *False interventions in data set D*

Figure 4.2: *Histograms showing distribution of time to collision at start of intervention.*

In Figure 4.2, the time to collision (TTC) at the start of the intervention is visualised for the true and false interventions respectively, in the four different data sets. The TTC is a calculation of how long it will take until the host collides with the target if no braking occurs. How this time is calculated can differ, mainly depending on if the velocity of the host and target are assumed to remain constant or not. In Volvo's system, the calculation is complicated and accounts for multiple different factors. Therefore, a simplified calculation of the TTC has been used. In case of steady-state, i.e. when the target has no acceleration, the TTC is computed as

$$TTC_1 = \frac{d}{v_h - v_t} \ ,$$

where d is the longitudinal distance to the target, v_h is the longitudinal speed of the host and v_t is the longitudinal velocity of the target. With the acceleration of the target, denoted a_t, taken into account, the TTC is instead computed as

$$TTC_2 = \frac{v_h - v_t}{2 \cdot a_t} + \sqrt{\frac{(v_h - v_t)^2 - 4 \cdot a_t \cdot d}{4 \cdot a_t^2}} \ .$$

In Figure 4.2, the minimum of TTC_1 and TTC_2 is used for each log. This means that if the target decelerates, the deceleration is accounted for, while if the target has constant speed or accelerates, the speed is assumed to remain constant.

Since the emergency brake should only be initiated when a collision is becoming inevitable, the TTC should correspond to exactly the required time of braking to avoid a collision. However, depending on the scenario, the system might start with a short pre-brake, which results in a slightly higher TTC at the start of the intervention.

4.3 Duration of intervention

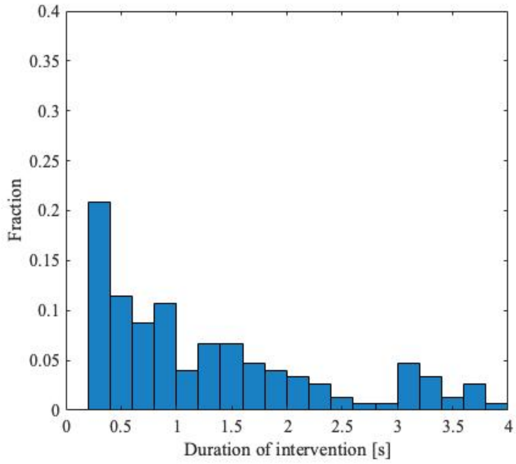

(a) *True interventions in data set A*

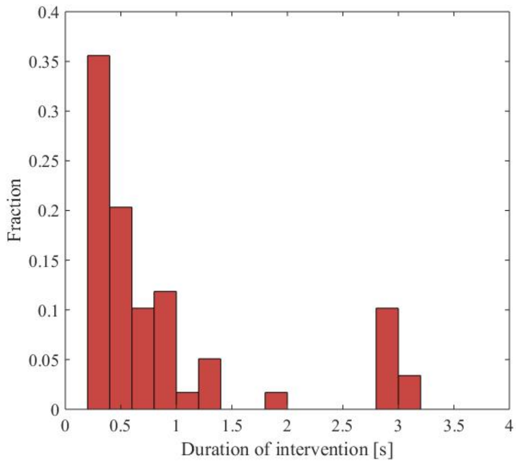

(b) *False interventions in data set A*

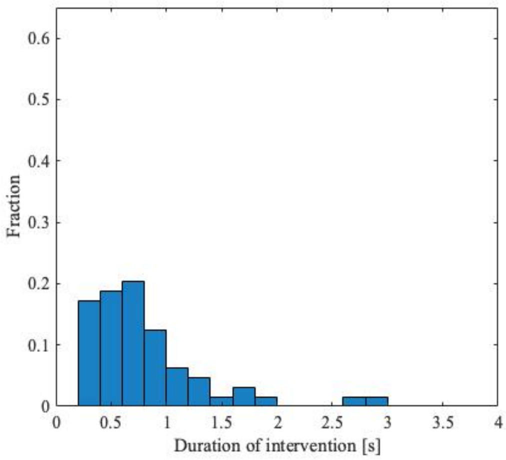

(c) *True interventions in data set B*

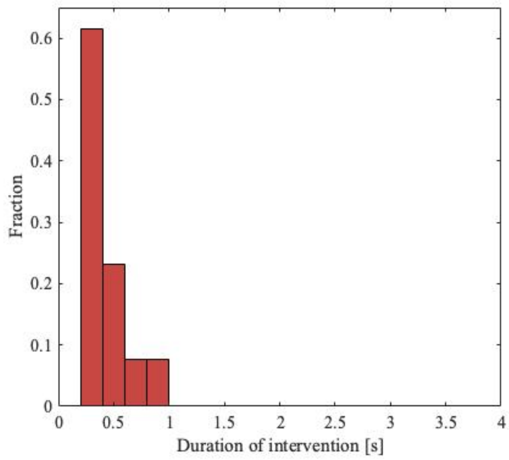

(d) *False interventions in data set B*

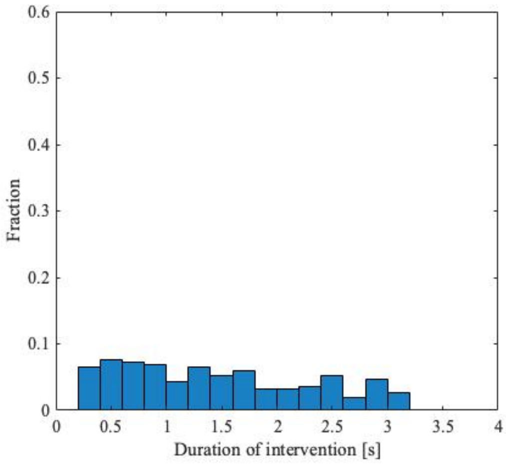

(e) *True interventions in data set C*

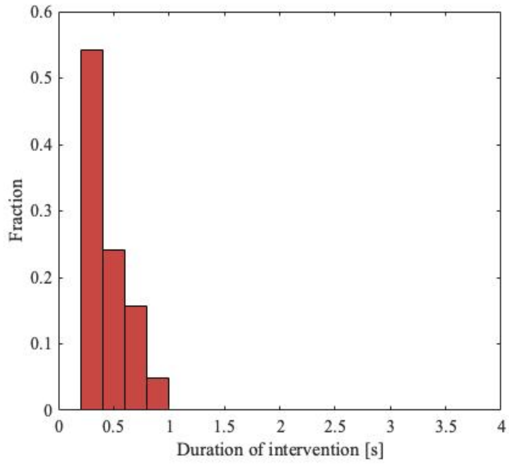

(f) *False interventions in data set C*

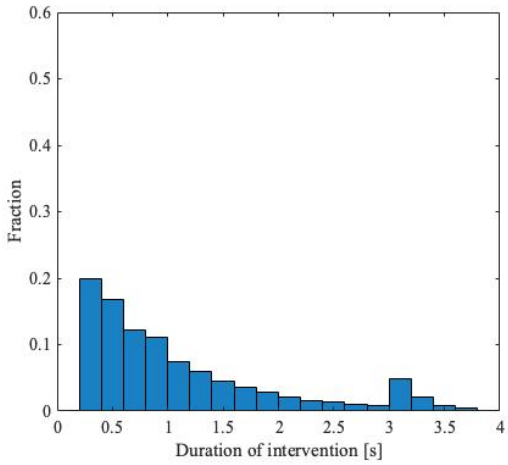

(g) *True interventions in data set D*

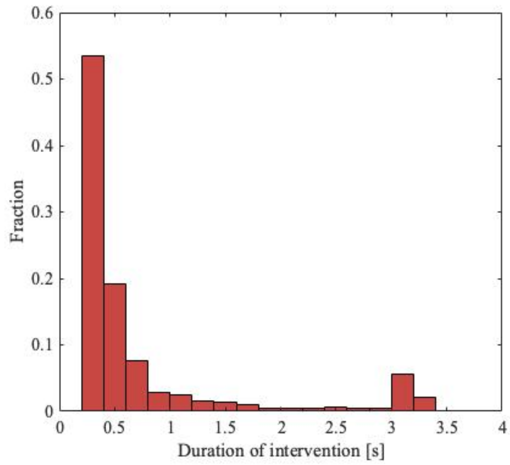

(h) *False interventions in data set D*

Figure 4.3: *Histograms showing distribution of the duration of the interventions.*

The duration of the interventions are shown in Figure 4.3. The duration was defined as the time when the AEBS system was in a pre-brake or full-brake state, and was calculated as the number of timesteps multiplied by the length of a timestep, i.e. 0.2 s. In the four data sets, the false interventions tend to be short, mostly below 1 second, while the duration of the true interventions seem to be above 1 s in around 50% of the interventions.

A peak can be seen at a duration of around 3 seconds, which most likely depends on the fact that most logs are 6 seconds long and have the start of the intervention in the middle of the log, i.e. there data only contains information about the three first seconds after the start of the intervention. This means that most logs with an intervention that lasts longer than 3 seconds still will appear in the bar at 3 seconds. However, in some of the logs, the start of the intervention is not exactly in the middle of the log, which might result in a few logs with long duration end up in the bars at 2.8 or 3.2 s. However, below 2.5 seconds, the histograms should be reflect the reality.

4.4 Speed reduction

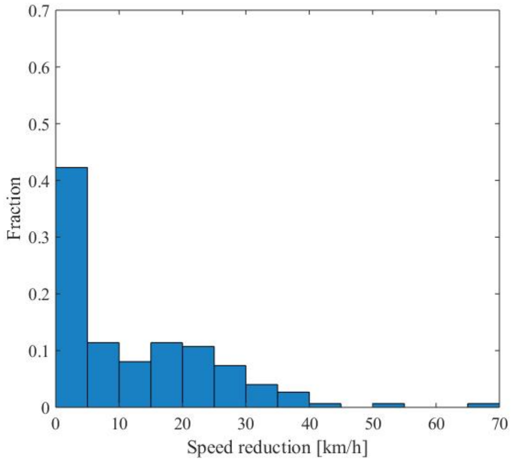

(a) *True interventions in data set A*

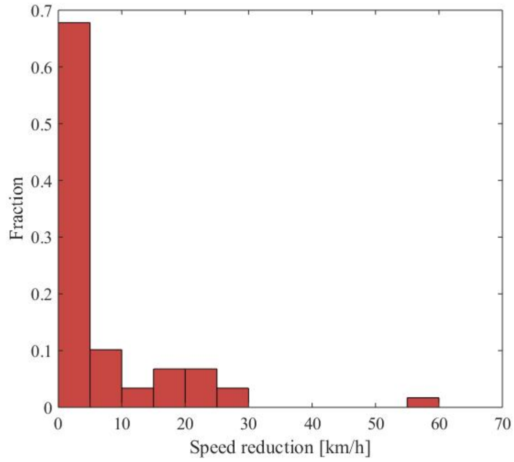

(b) *False interventions in data set A*

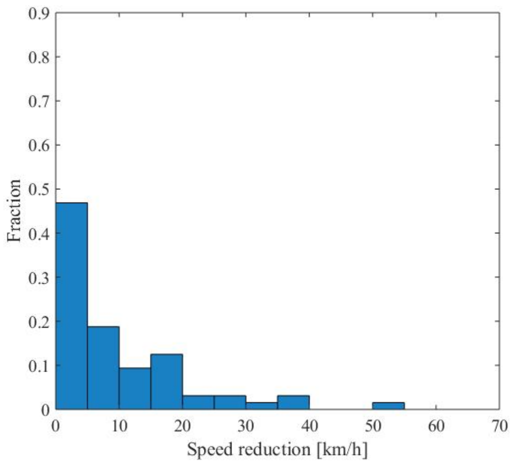

(c) *True interventions in data set B*

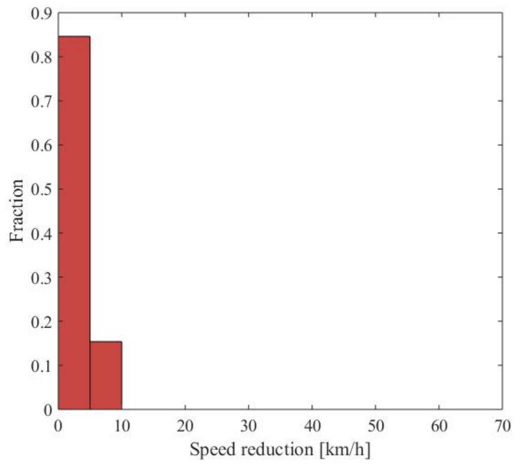

(d) *False interventions in data set B*

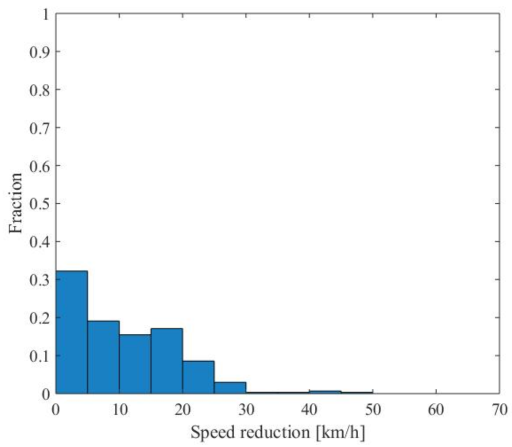

(e) *True interventions in data set C*

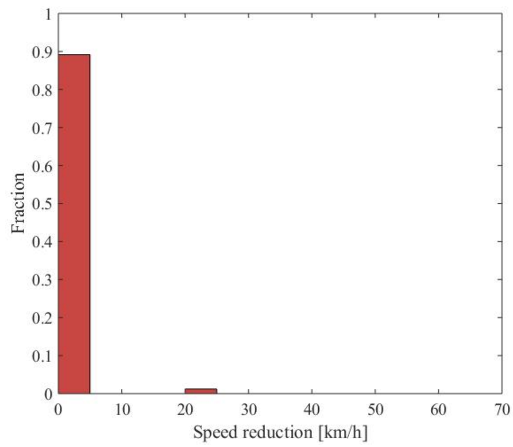

(f) *False interventions in data set C*

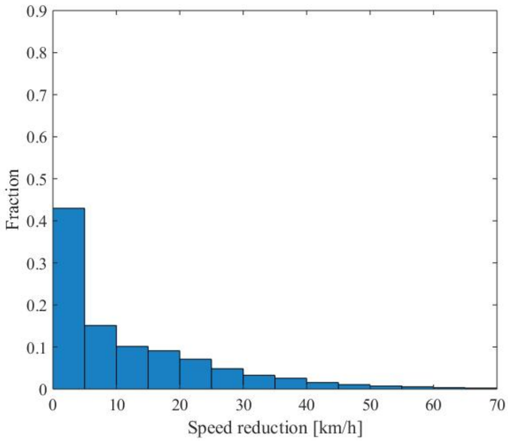

(g) *True interventions in data set D*

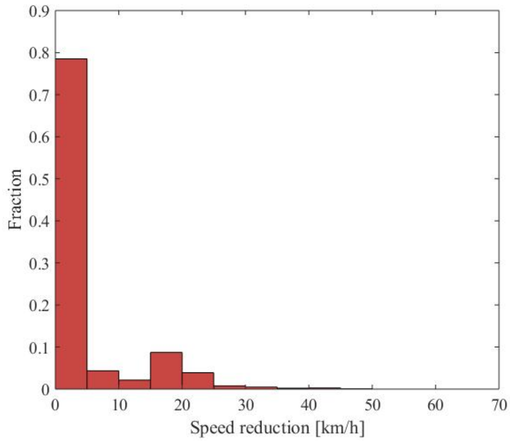

(h) *False interventions in data set D*

Figure 4.4: *Histograms showing distribution of the speed reduction of the host from the start until the end of the intervention.*

The speed reduction shown in Figure 4.4 was measured as the difference in host speed between the start and the end of the intervention. Among the true interventions, around 30-50% have a speed reduction below 5 km/h, whereas the proportion among the false is 70-90%, i.e. almost twice as high. This applies to all data sets, and data set D, which is the largest, seems to be somewhere in the middle of both intervals. Thus, for the true interventions, the fraction of logs with a speed reduction higher than 5 km/h is much larger than for the false. Besides, among the true interventions, it is not unusual with speed reductions of 30 km/h, in contrast to the false interventions where this seems extremely rare.

4.5 Life length of target

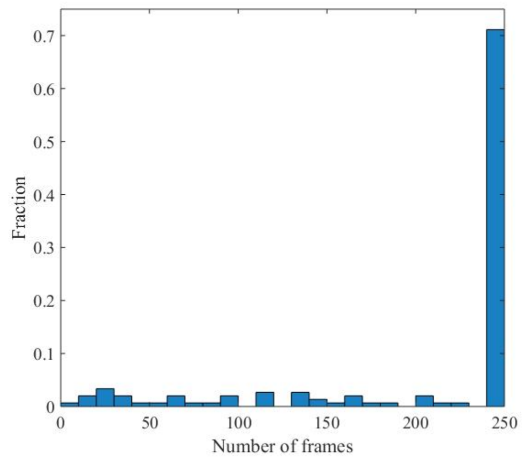

(a) *True interventions in data set A*

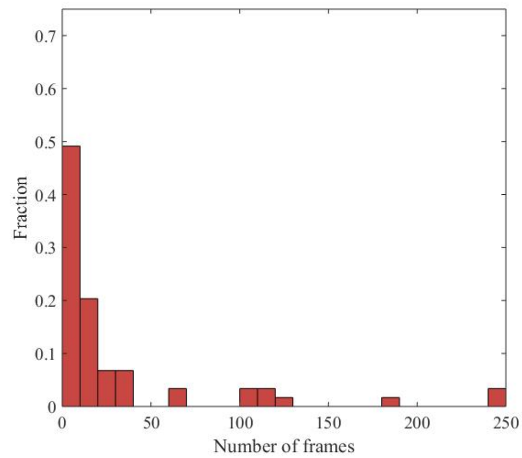

(b) *False interventions in data set A*

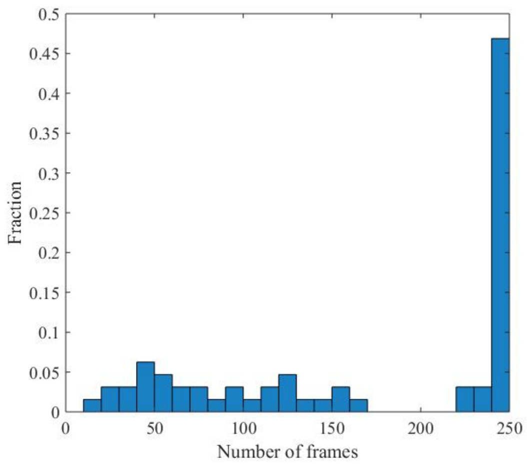

(c) *True interventions in data set B*

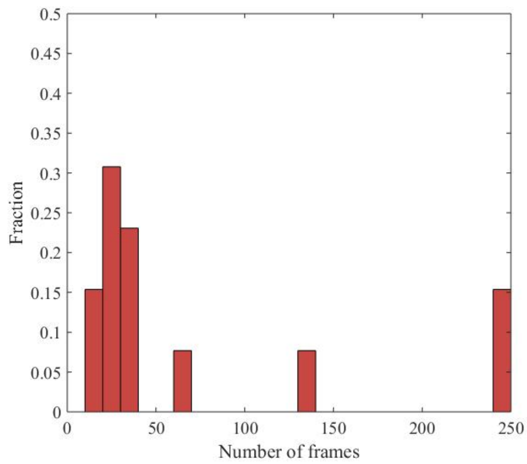

(d) *False interventions in data set B*

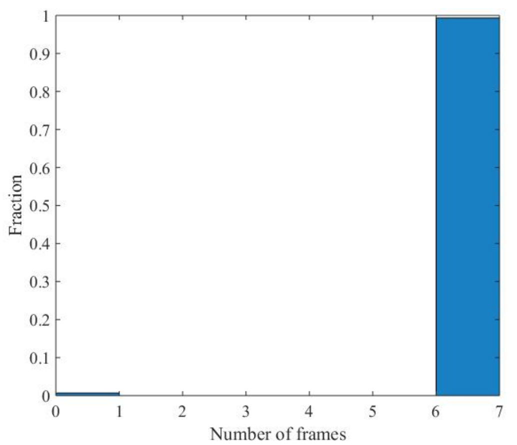

(e) *True interventions in data set C*

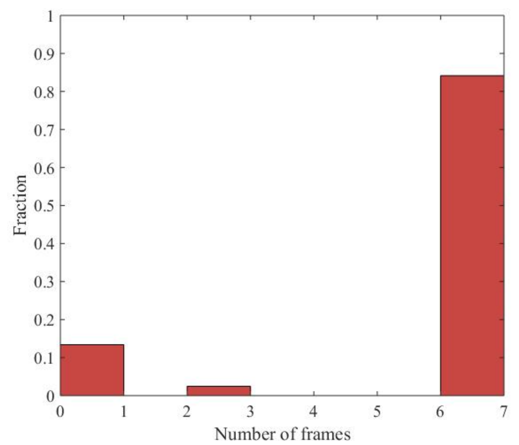

(f) *False interventions in data set C*

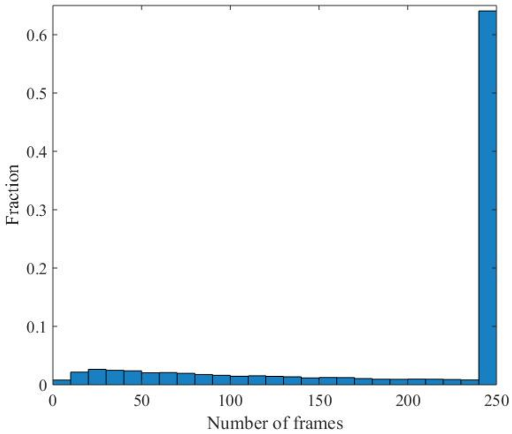

(g) *True interventions in data set D*

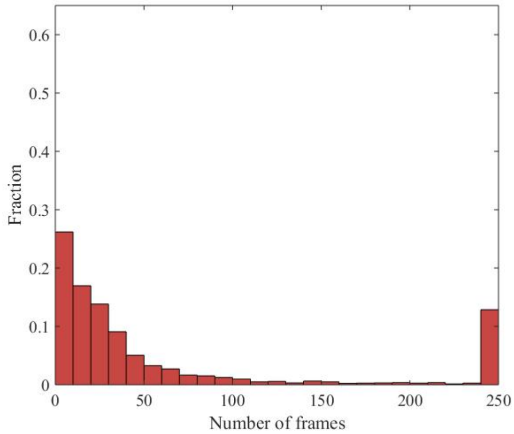

(h) *False interventions in data set D*

Figure 4.5: *Histograms showing distribution of life length of target at start of intervention. In data sets A, B and D, one frame corresponds to one fifteenth of a second and the maximal value is 250 frames, whereas in data set C, one iteration corresponds to 30 ms.*

The life length of the target is defined as the time that it has been tracked by the system at the start of the intervention, including any time it was tracked before being selected as target. For data sets A, B and D the life length increases every fifteenth of a second until it reaches a maximum value of 250. For data set C, it increases every 30 ms until it reaches a maximum value of 7, i.e. 0.21 s. In the majority of the true interventions, the life length of the target reaches its maximum value in all data sets, i.e. the target has been tracked for a longer time. Among the the false interventions, there is also a tendency of targets with maximal life length, this is small in comparison with the fraction of newly detected targets, i.e. tracked for less than 3 s.

4.6 Fusion of the sensors

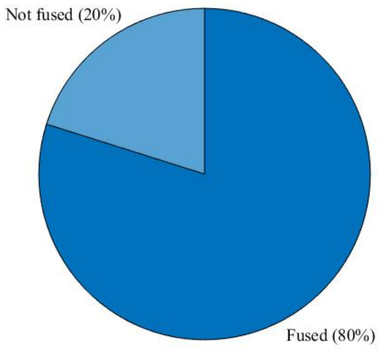

(a) *True interventions in data set A*

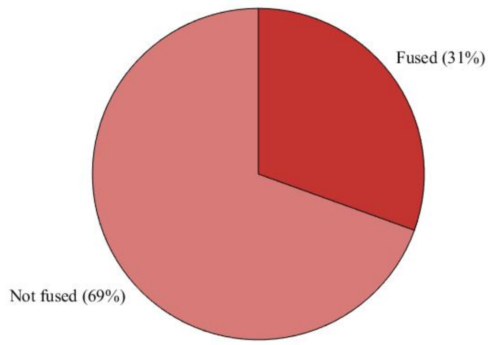

(b) *False interventions in data set A*

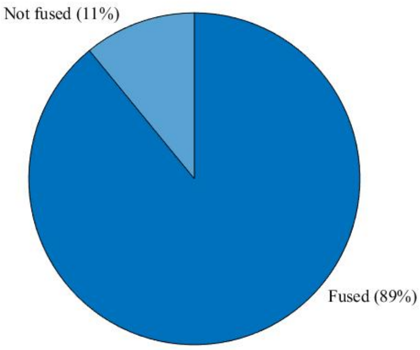

(c) *True interventions in data set B*

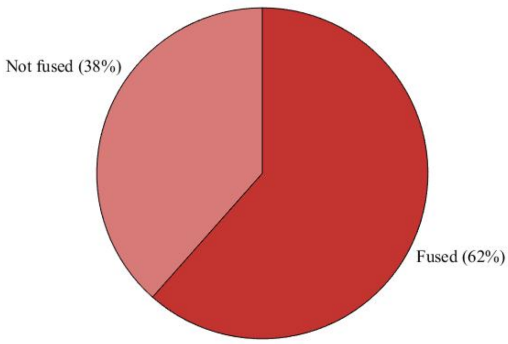

(d) *False interventions in data set B*

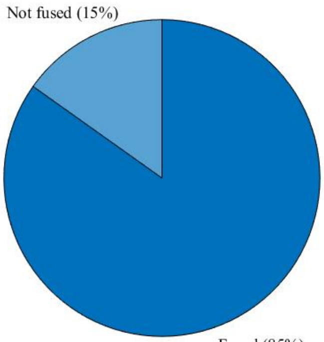

(e) *True interventions in data set C*

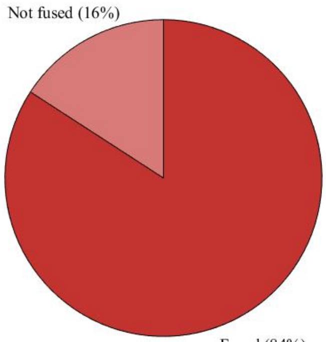

(f) *False interventions in data set C*

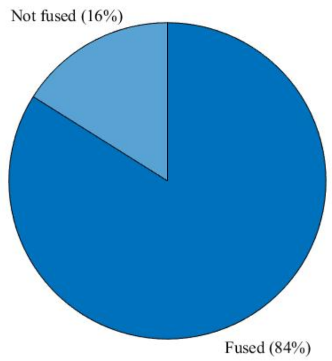

(g) *True interventions in data set D*

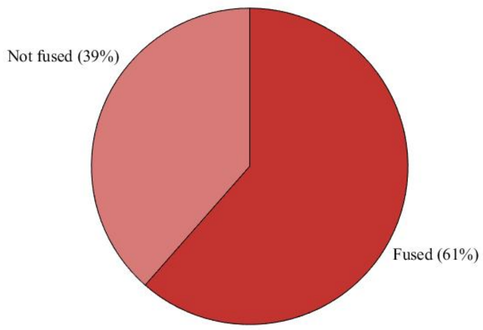

(h) *False interventions in data set D*

Figure 4.6: *Pie charts showing proportion of fused targets at the start of intervention.*

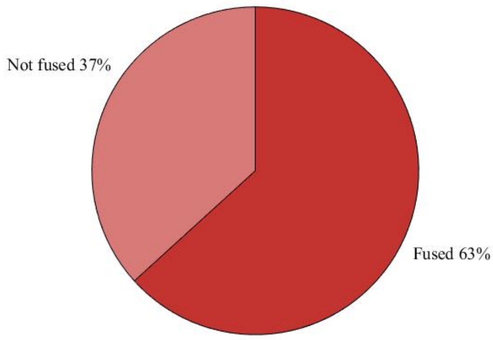

Figure 4.7: *Pie chart showing proportion of fused targets at start of false interventions in data set A, excluding interventions classified as false by FCC4.*

Figure 4.6 shows the fraction of targets that are fused and not fused at the start of the intervention. In all four data sets, the majority of true interventions occurred with a fused target. The fraction of fused targets among the false interventions differs for the different data sets, but is lower than among the true interventions. The fraction of fused and not fused targets in the false interventions in data set A where FCC4 was excluded is shown in Figure 4.7. This was found to be similar to the distributions among the false interventions in data sets B and D.

4.7 Longitudinal distance to target

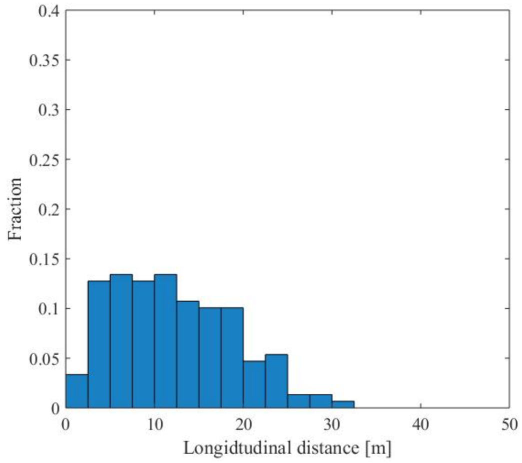

(a) *True interventions in data set A*

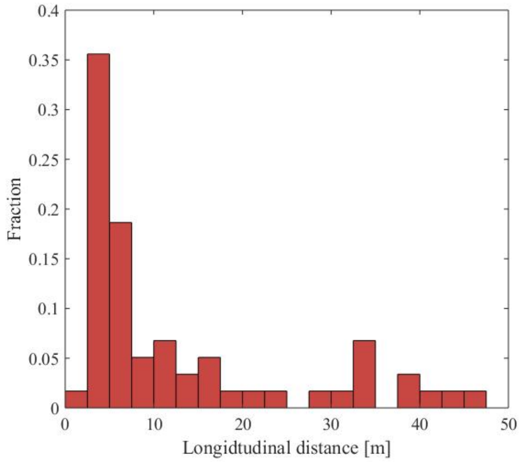

(b) *False interventions in data set A*

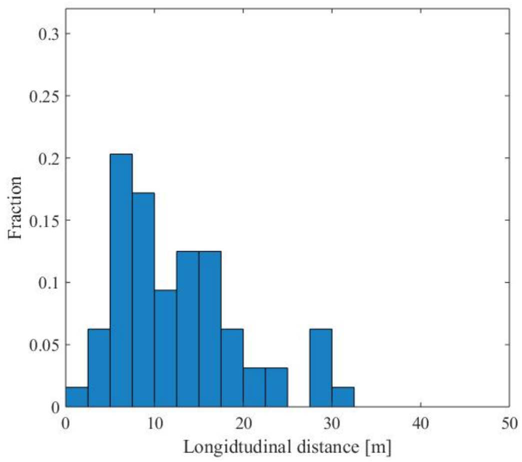

(c) *True interventions in data set B*

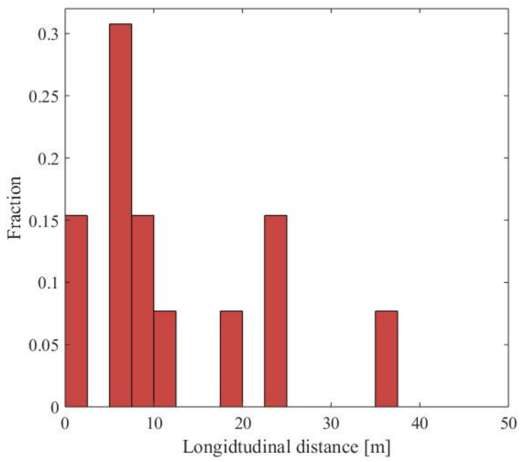

(d) *False interventions in data set B*

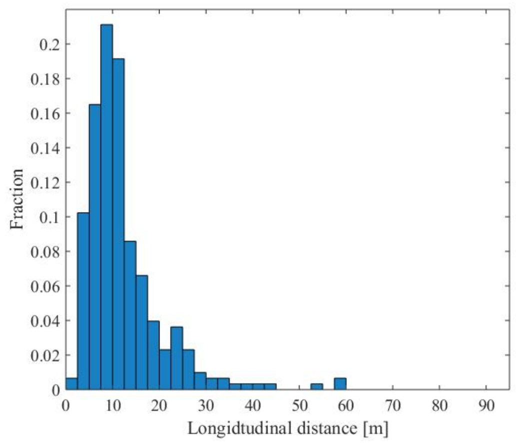

(e) *True interventions in data set C*

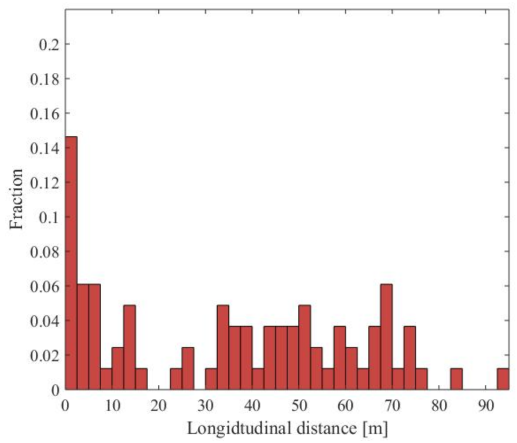

(f) *False interventions in data set C*

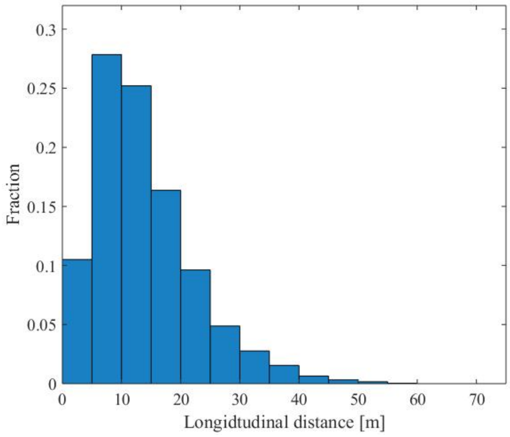

(g) *True interventions in data set D*

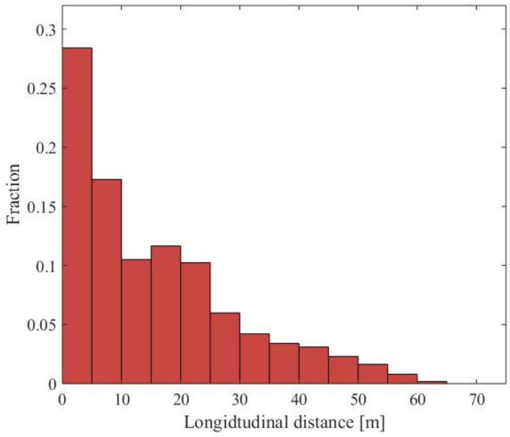

(h) *False interventions in data set D*

Figure 4.8: *Histograms showing distribution of longitudinal distance to target at start of intervention.*

In the true interventions in data sets A and B, the longitudinal distance from the host to the target is below 35 meters, most often below 20 meters. In data set D, there are true interventions with higher longitudinal distances, but also here the distance tend to be below 20 meters. In the false interventions, the distance is most often below 10 meters in the four data sets. In the true interventions in data set C, the longitudinal distance tends to be below 35 meters as well, but also reaches higher values (see Figure 4.8e). In the false interventions, a peak is obtained at distances below 5 meters, but is otherwise pretty evenly distributed between 5-75 meters.

4.8 Longitudinal speed of host

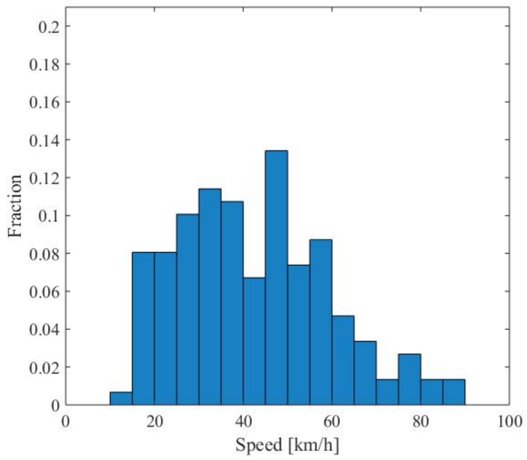

(a) *True interventions in data set A*

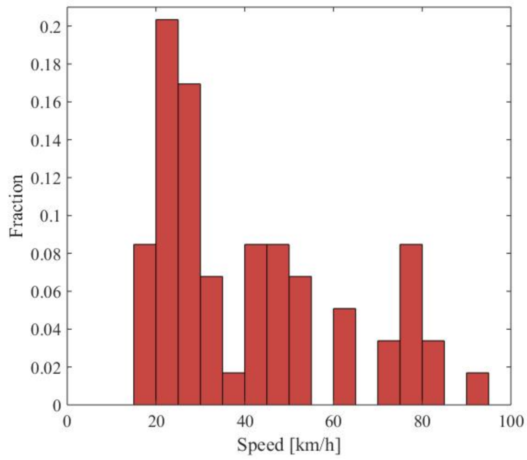

(b) *False interventions in data set A*

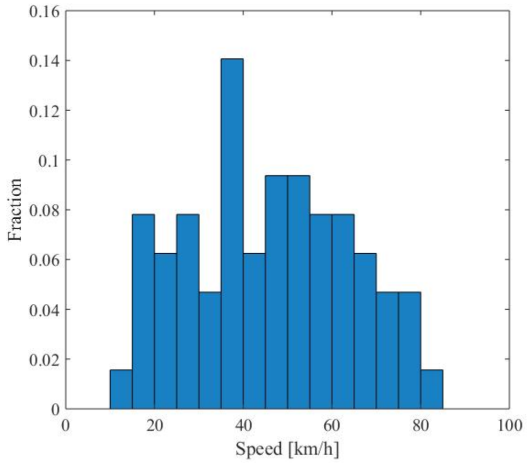

(c) *True interventions in data set B*

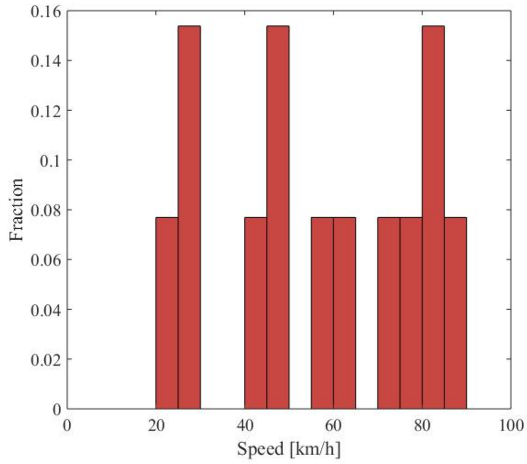

(d) *False interventions in data set B*

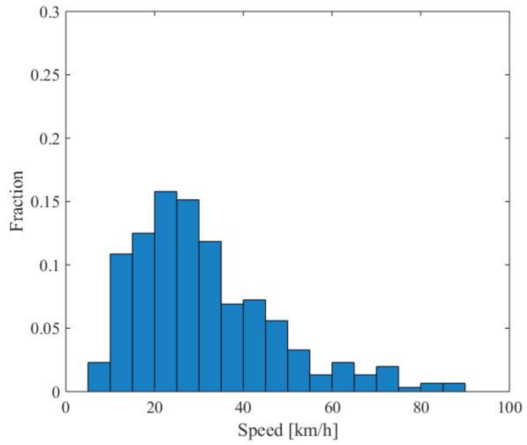

(e) *True interventions in data set C*

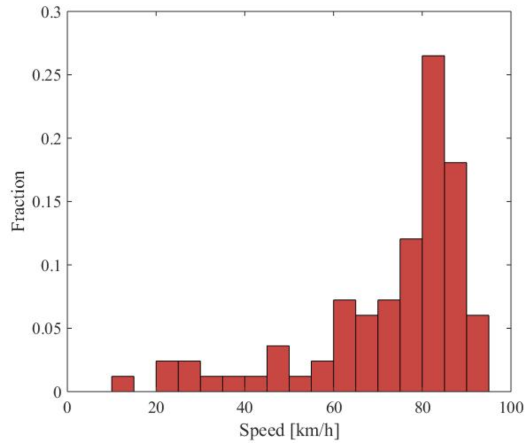

(f) *False interventions in data set C*

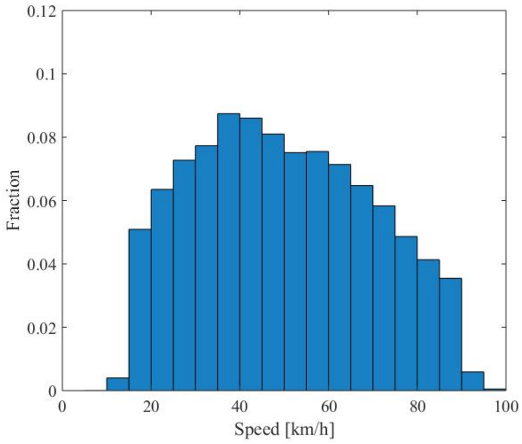

(g) *True interventions in data set D*

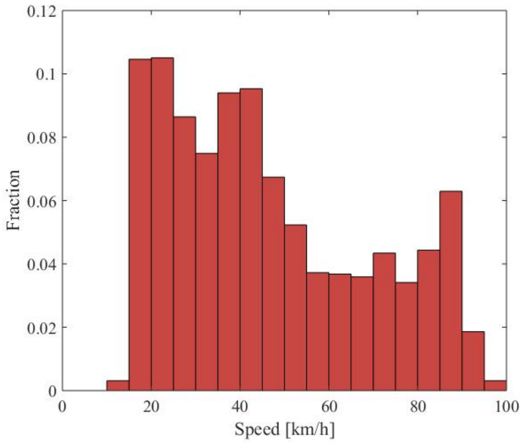

(h) *False interventions in data set D*

Figure 4.9: *Histograms showing distribution of longitudinal speed of host at start of intervention.*

The speed of the host in the true interventions is between 10 and 90 km/h in both data set A and B. In data set A, there is a peak at 45-50 km/h, and in data set B at 35-40 km/h. In data set D, the distribution is similar and peaks at 35-40 km/h. The false interventions in data set A tend to occur at lower speeds (15-30 km/h) than the true ones, but there are also false interventions at high speeds. Also for the false interventions, the distribution is similar for data set D. However, for data set B, the speed om the false interventions tend to vary more and is evenly distributed between 20-90 km/h.

Data set C differs from data sets A and B. For the true interventions, the most common host speed is 15-35 km/h, but many false interventions occurred at host speeds of 80-95 km/h. In data set D, a smoother distribution is shown and all velocities between 15 and 90 are clearly represented. for both true and false interventions. The true interventions show a wide peak around 40 km/h, whereas the false interventions has sharper peaks around 20, 40 and 85 km/h and fewer interventions around 50-75 km/h.

4.9 Longitudinal speed of target

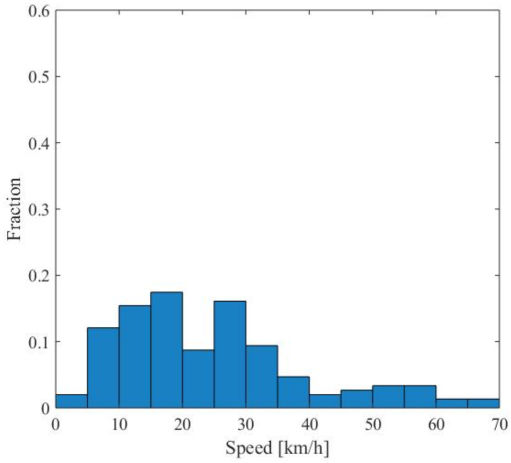

(a) *True interventions in data set A*

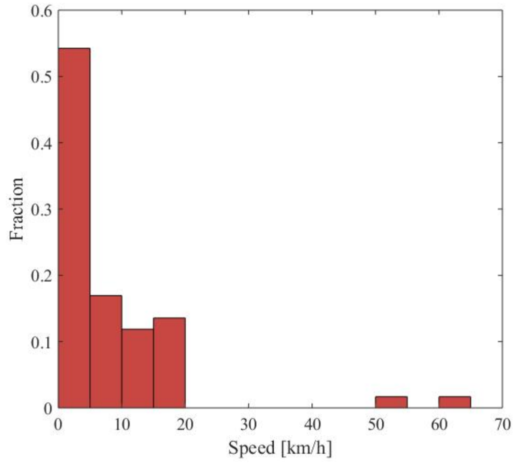

(b) *False interventions in data set A*

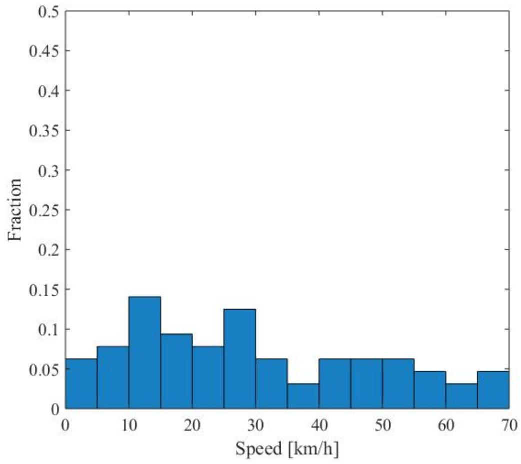

(c) *True interventions in data set B*

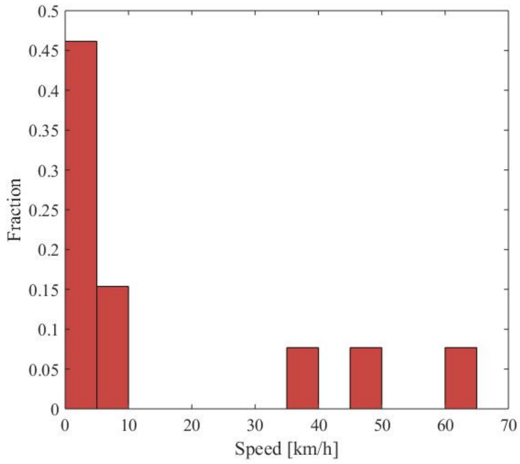

(d) *False interventions in data set B*

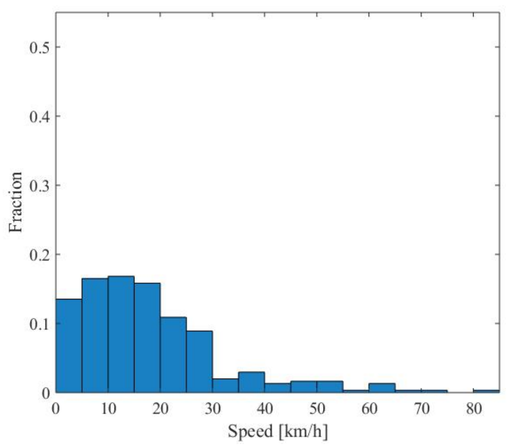

(e) *True interventions in data set C*

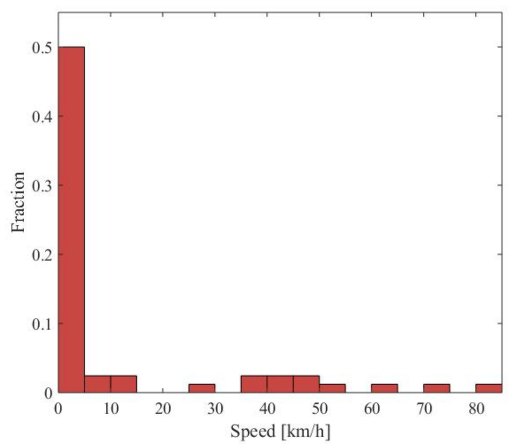

(f) *False interventions in data set C*

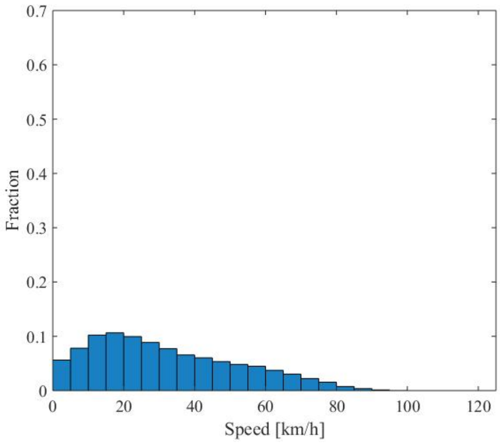

(g) *True interventions in data set D*

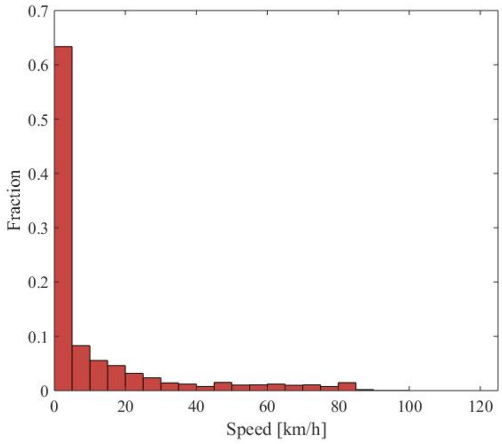

(h) *False interventions in data set D*

Figure 4.10: *Histograms showing distribution of longitudinal speed of target at start of intervention.*

The speed of the target in the true interventions is varying but has somewhat higher density between 10-30 km/h in all data sets. In the false interventions, there is a clear at 0-5 km/h in all data sets, i.e. mainly stationary targets.

4.10 Dynamics of target

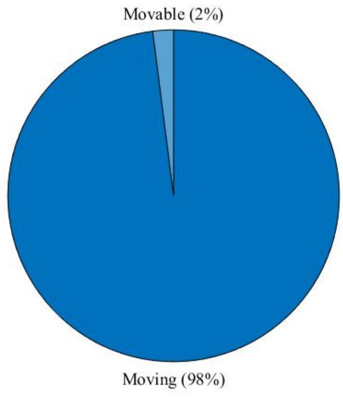

(a) *True interventions in data set A*

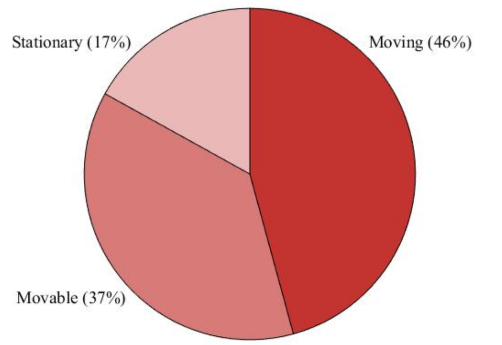

(b) *False interventions in data set A*

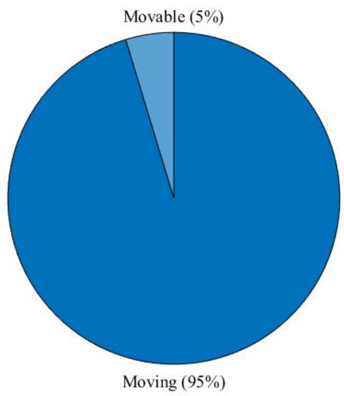

(c) *True interventions in data set B*

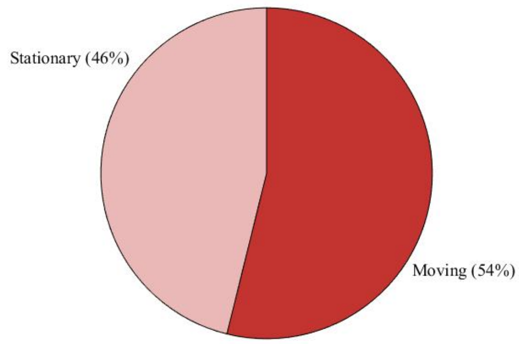

(d) *False interventions in data set B*

Figure 4.11: *Pie charts showing distribution of the dynamics of the target causing the brake intervention.*

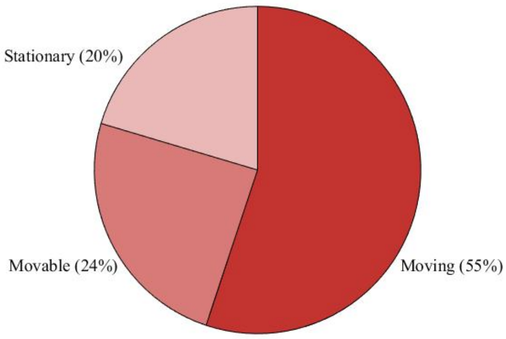

Figure 4.12: *Pie chart showing distribution of the dynamics of the target in false interventions in data set A, excluding interventions classified as false by FCC4.*

In data sets A and B, there is a variable defining if the target is considered stationary, movable (i.e. has previously been seen moving) or moving. In both data sets, the true interventions almost exclusively are due to moving targets. The distribution in the false interventions differs from the trues ones. Also here, the greatest proportion is moving targets, but there is also a noticeable proportion of stationary and/or movable targets. Figure 4.12 shows the distribution of the dynamics in the false interventions in data set A, but without the interventions considered to be due to misclassified targets, i.e. the ones fulfilling FCC4. Here, the proportion of moving targets is similar to the proportion of moving targets in the false interventions in data set B.

4.11 Relative longitudinal speed

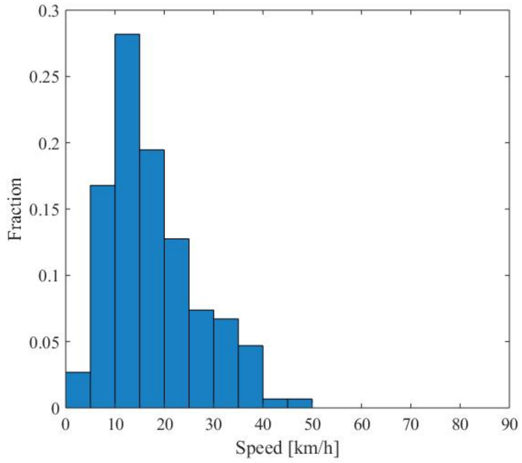

(a) *True interventions in data set A*

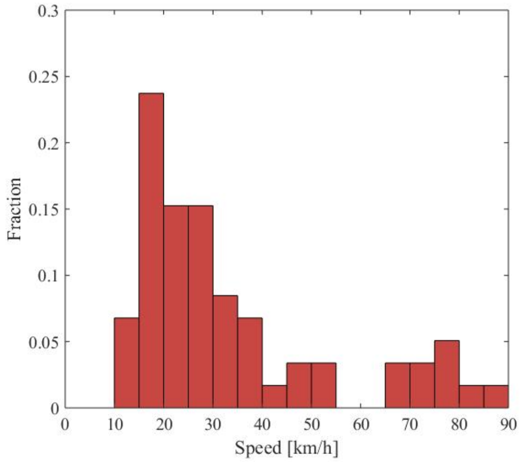

(b) *False interventions in data set A*

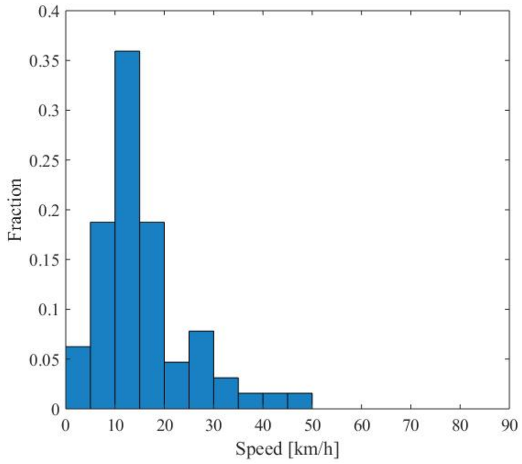

(c) *True interventions in data set B*

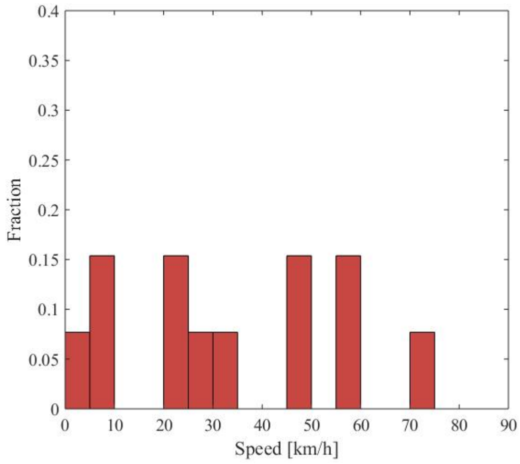

(d) *False interventions in data set B*

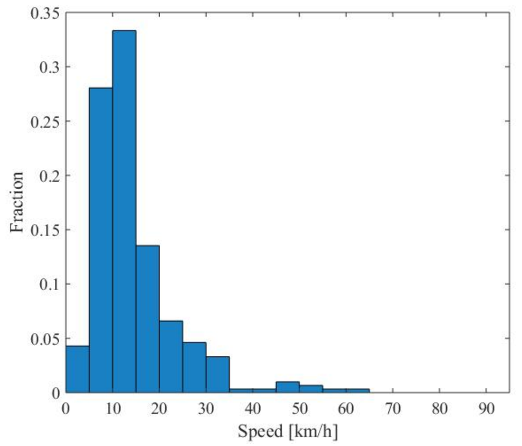

(e) *True interventions in data set C*

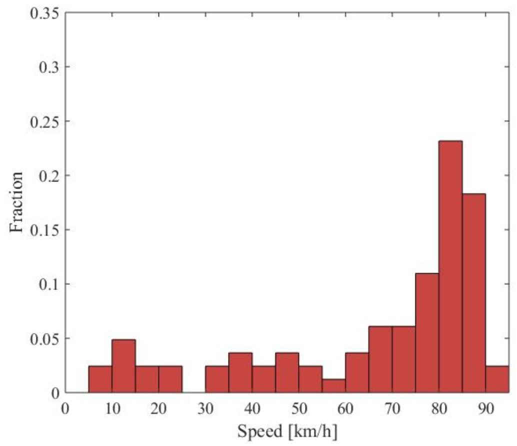

(f) *False interventions in data set C*

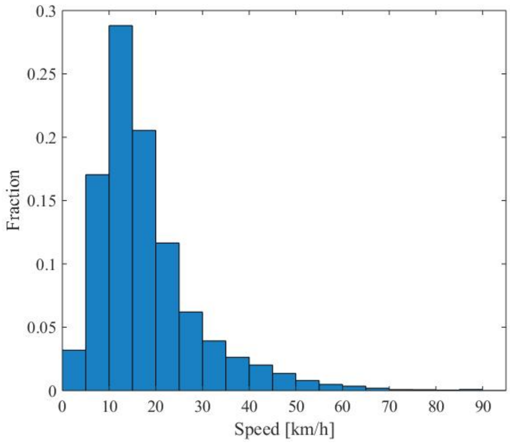

(g) *True interventions in data set D*

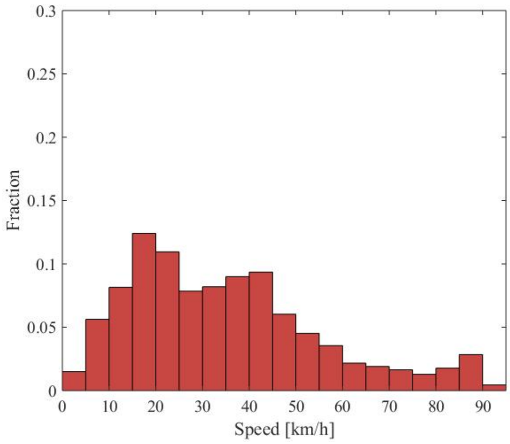

(h) *False interventions in data set D*

Figure 4.13: *Histogram showing distribution of relative longitudinal speed between host and target at start of intervention.*

The relative speed between the host and the target is most often 10-15 km/h in the four data sets, and always below 50 km/h in data sets A and B. In the false interventions in data set A, the relative speed is most often at 15-20 km/h which is similar to the true interventions. However, there is also some interventions at high relative speeds. The distribution among the false interventions in data set D looks similar, but with a lower peak at 15-20 km/h and somewhat higher proportion of interventions with higher relative speeds. The false interventions in data set B are evenly distributed. In data set C, the relative speed in the false interventions peaks at 80-85 km/h, which is considerably higher than in the other data sets. Note that the speed of the host tends to be high and the speed of the target tends to be low in this selection (see Figures 4.9f and 4.10f), which motivates this result.

4.12 Time of day

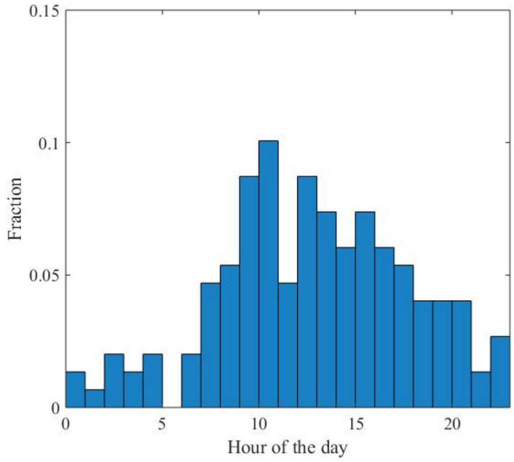

(a) *True interventions in data set A*

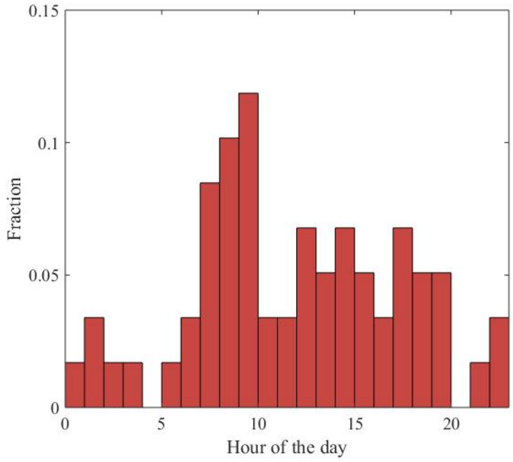

(b) *False interventions in data set A*

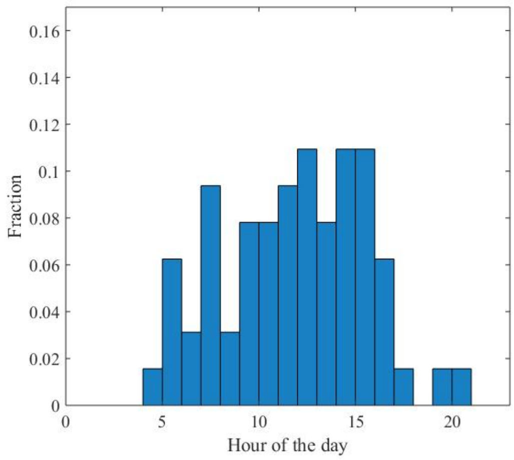

(c) *True interventions in data set B*

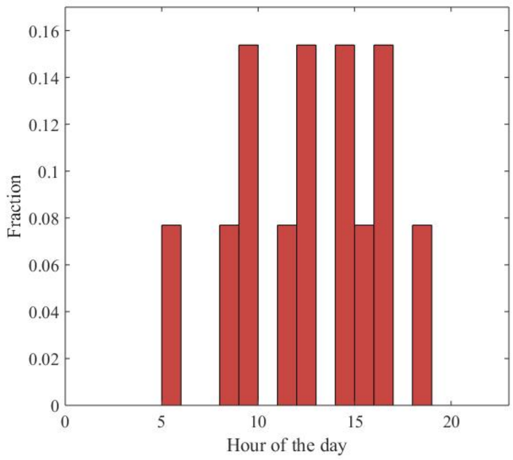

(d) *False interventions in data set B*

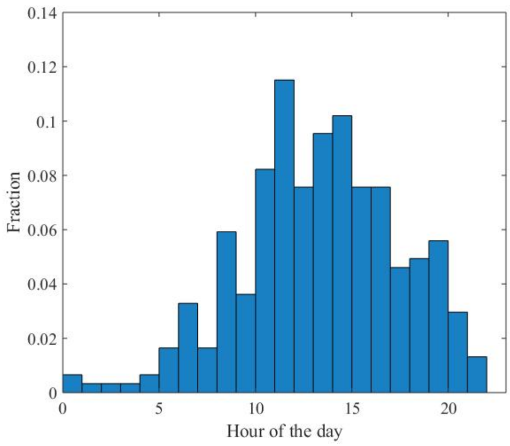

(e) *True interventions in data set C*

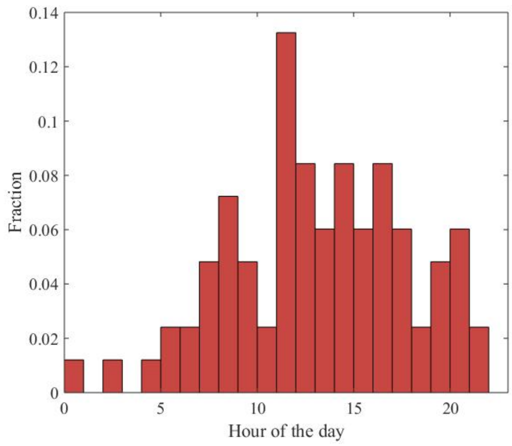

(f) *False interventions in data set C*

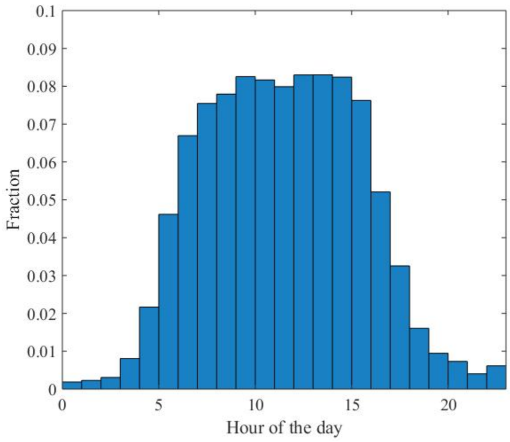

(g) *True interventions in data set D*

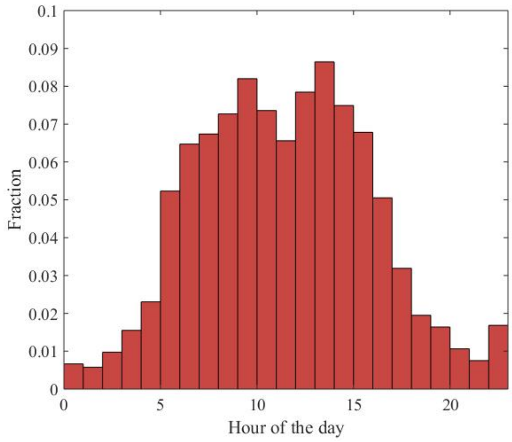

(h) *False interventions in data set D*

Figure 4.14: *Histograms showing distribution of which hour of the day the intervention occurred.*

During what hour of day the interventions occur was investigated to find whether the interventions tend to occur during day or night. The distributions shown in Figure 4.14, show that the majority of interventions occurs during the day.

4.13 Investigation of ghost targets

In the literature review about automotive radar, it was found that certain shapes cause strong reflections which might be misinterpreted to derive from a vehicle. Even though the radar antenna uses beamforming to suppress signals not deriving from objects ahead of the HGV, the side lobes and secondary reflections might still be problematic. This means that in rare occasions, the radar might receive radiation reflected at a surface that is not an object.

The radar is able to measure the azimuth angle (the angle in the horizontal plane), which is required for the fusion, but the accuracy is not very good. The ability to measure the elevation angle is however often even worse, making it difficult to distinguish the relevant vehicles from reflections from the ground or above the host vehicle.

Another important aspect is how the noise level is affected by the environment. Sometimes, for example when driving in a tunnel or close to a metal side barrier, there will be lots of small reflections that are weak enough to distinguish from other vehicles, but still contributing to a higher noise level, and thus reducing the accuracy of the actual measurements. It is possible to adapt the threshold to the current noise level to reduce the risk for misinterpretations, but that would also increase the risk that a real target is missed.

One of the most important findings in the literature study is that the road environment contains a large number of structures that have corners. These corners have extremely large radar cross sections, even when they are not perfectly aimed towards the radar. This makes them highly visible to the sensor. In addition, since some sensors use the radar cross section of the object to classify its object type (car, truck, pedestrian etc), the corner shapes might easily be misclassified as a vehicle relevant for AEBS.

Using the GPS positions of locations where false interventions have occurred, the environments causing the ghost targets were analysed. In data set A, there were in total 6 locations with multiple interventions, of which 5 had 2 interventions and 1 location with 6 interventions. Thus, in total there were 16 false interventions on these 6 locations in data set A. At the majority of these locations, large metal structures can be found, for example tunnels or bridges for elevated roads or other flat metal object such as side barriers and drain covers. In data set B and C, there were no locations with multiple interventions. This is believed to be because of the small number of interventions collected from a large area, making it unlikely that different HGVs have driven in the same areas. In data set D, there were several locations with multiple false interventions, however most of them were from test tracks and airfields, which is believed to derive from test runs of the system. Therefore, these were excluded from this study. The environments at the found locations are shown in Figures 4.15, 4.16 and 4.17. As these images are from Google Street View, they might be taken from a different lane than the intervention occurred. This might lead to an image taken in a direction against the traffic in the current lane, to show the environment in the direction the HGV was driving when the intervention occurred. As there occasionally are roads on different vertical levels, it can be uncertain which of the roads the interventions occurred on. This is discussed further in Section 5.5.

Figure 4.15: *A category or road environment where multiple false interventions have occurred. The metal bridge above the road, in combination with the pillar and beam holding it up, seems to be the common factor. Image: partly censored screenshots of Google Street View[21].*

Figure 4.16: *Another category or road environment where multiple false interventions have occurred. The metal walls beside the road, in combination with metal seams or drain covers, seems to be the common factor. Image: partly censored screenshots of Google Street View[21].*

From the environments studied, it is clear that a few characteristic road structures are frequently occurring. As it seems, there are three main categories that might have caused the ghost targets: below bridges for elevated roads, on the elevated road with metal walls on the sides and inside tunnels. All these categories are quite similar, as a large fraction of the field of view is covered in metal structures with many corners. This might both result in very strong reflections and an increased noise level due to the many corners acting as great reflectors. Some examples of corner shapes and flat surfaces perpendicular to the direction of travel found at the investigated locations are shown in Figure 4.18. These types of shapes exist both on the ground, on walls or barriers beside the road and on the underside of bridges above the road.

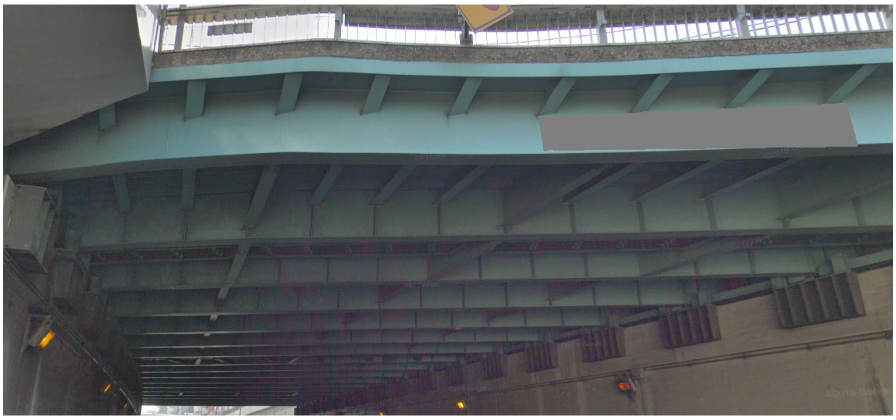

Figure 4.17: *A third category or road environment where multiple false interventions have occurred. The ceiling of the tunnel has a metal grid structure which can cause many unwanted reflections. Image: partly censored screenshot of Google Street View[21].*

Figure 4.18: *Zoomed in view of objects that might cause unwanted reflections that resulted in a false intervention.Image: partly censored screenshots of Google Street View[21].*

In these false interventions, the target either appeared close ahead of the host (4-18 m) and slowly drifted forward (5 m/s), or appeared as an almost stationary object at a large distance (above 45 m). In almost all cases, the target drifted towards the lateral center of the HGV path and disappears shortly before ”impact”. Most of the locations were found to be one-way roads. Since the number of lanes seems to vary greatly, it does probably not have a big impact on the performance.

5 Discussion

To estimate the benefits from AEBS is not trivial. In this thesis project, data in the form of intervention logs has been used for analysis. This data is in general believed to reflect how AEBS is used in the real world, but there might be factors that influence the validity of the data. This is described further in Section 5.1, followed by a discussion about the limitations in the analysis in Section 5.2, and then also about the TFC program and how well it performed in Section 5.3.

The analysis of different variables yielded many histograms and pie charts. How these should be interpreted and their validity, together with the most important sources of error and how they were handled is discussed in Section 5.4. The analysis of false interventions was linked together with the principles of a radar to find potential weaknesses, which is described in Section 5.5. Finally, some ideas that were generated during the project and areas that would be interesting to investigate further are described in Section 5.6.

5.1 The data

The intervention logs used in this study cover a very broad range of near-collision scenarios. The fact that all interventions in the data used are unique necessitates a manual analysis to understand the scenarios, which is sensitive to misinterpretations due to the lack of hindsight. However, one of the data sets had video and sound recordings of the interventions, which could be used together with the TFC program for verification purposes. Altogether, the verification showed that the TFC program performed very well (above 95% correct classifications, see more in 5.3), which would not have been possible if the misinterpretations in the manual analysis were severe, and thus, this data was appropriate for this analysis.

The data used is retrieved from HGVs and then transferred to Volvo. Since there is a risk that errors occur during the retrieval or transfer of data, the data has been checked for duplicated logs and erroneous logs (e.g. missing values or corrupt files). To cope with this, a pre-processing was performed, that removed all duplicated logs and logs with apparent errors. However, there might have been other errors that went through the error check unnoticed, e.g. if a log for some reason would contain reasonable but incorrect values.

Even data that had been correctly transferred could be unsuitable for analyses. For example, if an HGV has been used to test the AEBS and a large amount of brake interventions were caused on purpose, it might give a different picture of how well the system performs in real traffic. If the system is tested several times at the same location, this location can be found in the comparison of GPS coordinates. In this project, a few of these locations where the AEBS system has been used in testing purposes were found, but only in data set D. Since these interventions were very few compared to the entire data set, they were not excluded from the analysis in this project. If the programs used in this project are reused in other projects, it is recommended to check the number of interventions from tests and make a decision whether they need to be removed. There is also a risk that the AEBS has been tested at different locations, which not can be found by only comparing the GPS coordinates. The risk of this was assumed to be low, and therefore, no action was taken to cope with it.

There is a risk that HGVs involved in a collision are taken to a scrap yard without retrieving its intervention logs. Therefore, there is a risk that the data set do not cover the interventions where collisions occur. Apart from that, the data sets are believed to be a good representation of how the AEBS is used in the real world.

The strategy for the data analysis was to use the smaller data sets A and B to develop the TFC program, use the medium sized data set C to verify the script and finally use it on the large data set D for the deeper analysis of true and false interventions. Since data set B is small, and in particular the number of false interventions are very few, normal variations will appear very clearly, and it is thus difficult to draw conclusions from the histograms and pie charts. On the other side, data set D is very large and the histograms show smooth shapes. Moreover, this data set contains more recently occurring interventions, and thus later versions of AEBS. Therefore, data set D is assumed to reflect how AEBS works in the real world. The conclusions about true and false interventions are mainly based on this data set, whereas the analysis of the radar is based on observations from all data sets, and in particular data set A due to the infrastructure of this data set.

In the beginning of the project, only data sets A and B were available. As the structure of data set C and D was unknown at the time, the data extraction could not be designed to also handle these. When data set C and D later became available, some changes had to be made to adapt the data extraction to their format and structure. Luckily, data set D only required a change in encoding, and was thus easily implemented. The data sets contained different variables depending on whether logging type α or β was used. Some variables did not exist in data set C (using logging type β), but had to be calculated using other variables. The irregularity in the measuring frequency in data set C was one of the main issues of the project. To be able to use data set C without spending too much time on data extraction or re-writing the TFC, a decision was made to convert it to the same format as the other data sets by interpolating to the same time vector as in data set A, B and D (i.e. the structure used in logging type α). The interpolation of the logs resulted in a large fraction of the information in data set C being left out of the analysis. However, the information that remained after the interpolation was considered to be sufficient to verify the TFC (see Section 5.3) and using the data for further analysis.

5.2 Limitations of the data analysis

Using simulations or driving on test tracks can give an indication of the extent of the benefits from AEBS, but the complexity of real world driving makes it impossible to account for every detail, and thus the results from this kind of tests will not be fully accurate.

Another approach to estimate the benefits from a system is case analysis of the system when it is used in the real world, i.e. in the form of a field test. This approach is very useful for determining the actual benefit from a system with high accuracy, but is not as useful for testing of specific scenarios since it is not possible to decide which factors and variables to consider or exclude. In addition, the analysis cannot be performed until the system has been implemented and used for a period of time, and the analysis is thus delayed.

From the time when a safety system is launched on the market, there is also a delay until a sufficiently high number of HGVs equipped with the system are on the roads to achieve a significant difference. For systems as AEBS, which are not activated very often, it might take some time for the HGVs to gather sufficient data for a statistical analysis. Altogether, there is a long time delay from the development of the system until its effects actually can be measured. Since AEBS is a relatively new system, there is not much information about the actual performance of the system in terms of how often it intervenes and how many collisions that actually are avoided, mitigated or even caused by the system. In addition, the constant development of the system makes it difficult to draw conclusions about the performance of the system, since there are different versions of the system in different HGVs.

The data sets A and B were collected some time ago, and with the older hardware version than data set C, while data set D is a combination of both hardware versions. The infrastructure in the geographic location in data set A was known to yield problems for the system at the time of the data collection. The collection of data set C occurred when the system was in a development state. Software releases have been updated since the collection of both of these data sets. Therefore, the number of false interventions would most likely be lower if an analysis was carried out with a later software version. Since the largest part of data set D was collected later than A, B and C, it is believed to best represent the general performance of the different AEBS versions combined. Data set D is also much larger than A, B and C, and therefore the performance in this study is believed to be a good measure of the actual performance.

5.3 True false classification

The task to develop a program that automatically classifies interventions as true or false was complicated. To identify all false interventions would require many and very specific criteria since traffic scenarios are very complex. These criteria would most likely not be applicable to different data sets, which is the point of implementing the TFC. Thus, to make the TFC applicable on different data sets, it cannot classify the logs 100% correct. There are always borderline cases, and also, the sensors have some limitations and inaccuracies. In the TFC, four false classification criteria were implemented. In each criterion, limits of different variables were defined and tuned based on manual analysis of the logs. These limits are not trivially defined, and can vary from case to case. The decision to implement a grey zone ("possibly false") made borderline cases easier to handle. Then, to further validate the performance of the script, it was run on a data set that contained videos of the interventions. The videos had been manually analysed by a Volvo employee, and the interventions were then classified as true or false. The TFC classification was in agreement with the manual analysis in 97 out of 100 of these interventions. Thus, the performance of the TFC was satisfactory.

During the manual analysis, some trade-offs between classifying an intervention as true or false had to be made. There are scenarios in which the driver can see what is going to happen while the system do not, which means that the intervention is considered true from the system's point of view, but the driver might not want to brake. The TFC is mainly based on the system's functionality, and not on the driver's perspective. However, in these borderline cases, the intervention tends to be short and the velocity reduction not that high, meaning that the false intervention will not cause any serious consequences.

The false classification criteria have a few minor flaws. The first criterion, i.e. if the object causing the intervention is only selected as target for a short time, is not 100% accurate. This criterion only considers for how long the object causing the intervention has been selected as target, and not for how long the object has been tracked by the system in total. Thus, it can still have been a critical object before it was selected as target, but then another object was considered even more critical. This was found to be the case in one of the three faulty classified logs among the 100 investigated logs mentioned above.

5.4 Variable analysis

Based on the analysed variables in Sections 4.2 - 4.12, conclusions about the performance of AEBS can be drawn, both based on the individual variables but also in relation to each other. Below follows a discussion regarding these variables.

Regarding the TTC visualised in Figure 4.2, it turned out that almost all true interventions are initiated when the TTC is between 0.5 and just above 2 seconds. In this period of time, it can be hard for a driver to react and brake the HGV enough to avoid a collision, even if the driver is attentive. Thus, the system seems to be able to intervene at the latest possible moment, as it should. For the false interventions, there were cases with TTC close to 0 in data set C and D. This means that the target must have appeared very close to the host, and the accelerations required for objects to suddenly "jump" in to the path is not physically possible to achieve. Therefore, these are mainly ghost targets. There were also a few cases with a TTC above 2.5 seconds, which indicates that the relative velocity with respect to the target was high.

In Section 4.3, the duration of the interventions is showed. The false interventions tend to have a short duration. This is no surprise, since one of the false classification criterion was "Short time as target", and if the target is lost, the intervention will in most cases be terminated. An exception from this is if the host speed has been reduced to below 15 km/h, in which case the intervention will continue until zero speed is achieved. One possible explanation of the short duration among false intervention is that they might derive from ghost targets. Since ghost targets not necessarily derive from reflections from large surfaces the same way as reflections from vehicles, the ghost targets are likely to be more sensitive to angular differences, and thus more unstable. This means they have a higher tendency of disappearing shortly after their appearance. False interventions might also be due to a temporarily erroneous measurement or incorrect fusion, which reverts after a short period of time.

The duration of the true interventions in the four data sets varies more than for the false interventions, and tends to be higher. For some true interventions, there might be sufficient with a short intervention to avoid an imminent collision, e.g. if the target switches lane and a minor speed reduction can be sufficient for the target to reach a sufficient lateral distance to come out of the future path of the host.

The speed reduction, shown in Figure 4.4, showed characteristics similar to the duration of the intervention. In a large fraction of the false interventions, the host had a speed reduction below 5 km/h. This is assumed to be related to the often short duration of the brake intervention, and also the fact that AEBS often only reaches the pre-brake state. The combination of low speed reduction and low deceleration due to only initiating a full brake, is interpreted as the false interventions generally not being severe. This, and the fact that only a low fraction of interventions is false, motivates the assumption that the false interventions are very unlikely to cause accidents. A significant fraction of the true interventions, somewhere around 40%, had a low speed reduction as well. This indicates that HGV drivers occasionally drive with a short distance to the preceding vehicle, or brake late enough for the AEBS to initiate an intervention. The intervention might in other words be considered unnecessary by the driver, but still be correct from a system point of view, due to a risky behaviour of the driver. There are also interventions with higher speed reduction of the HGV, in which the brake intervention most likely is crucial to avoid a collision.

From Figure 4.5, it is clear that most true interventions are due to targets with a high life length, while among the false interventions, most targets had a life length shorter than 3 seconds. This indicates that the targets causing true interventions have been tracked for a long time and are likely to be tracked with high confidence (since they otherwise might have been dropped and the life length reset), while false interventions are often caused by targets that were detected shortly before the start of the intervention.

In data set C, the different way to measure the target life length makes it difficult to draw any conclusions. Therefore, the analysis was mainly based on the other data sets. For the false interventions in data set A, almost half of the targets have a life count below 10 frames, corresponding to 0.66 s, at the initiation of the brake interventions. This distinguishes data set A from the other data sets, and is assumed to derive from a higher likelihood of ghost targets. Apart from this difference, the data sets A, B and D have similar characteristics. It seemed like data set B had a slightly lower fraction of true interventions with maximum life length than data sets A and D, but the reason for this was not found. Since data set B is small, normal variations will appear more dominant, which could explain the behaviour. There were also some false interventions of targets with a maximal life length. These interventions might for example derive from a bad measurement of the lateral distance of a vehicle that is overtaken by the host. This can happen if the vehicle is outside of the field of view of the camera but still tracked by the radar, which has a low angular accuracy. Such an intervention would likely be cancelled as soon as the vehicle is outside the field of view of the radar. In data set D, around 13% of the false interventions showed a maximum target life length.

When an object has been tracked by both the radar and camera for some time, it is likely to become fused. In general, targets that are fused are measured with a higher confidence in position and velocity. Ghost targets and inaccurate measurements are less likely to be fused. This is clearly shown in Figure 4.6, where most true interventions have a fused target while the fraction of fused targets is lower among the false. The proportion of fused targets varies between the data sets. In data set A, which is assumed to contain more ghost targets, the proportion of targets that are not fused is significantly higher than in the other data sets. This result supports the assumption of ghost targets, since ghost targets often are radar only targets. As data set A also differed in having a higher fraction of logs classified as false by FCC4, which considers the fusion status, a pie chart was made of the fusion status for all interventions fulfilling any of FCC1, FCC2 and FCC3 (i.e. excluding false logs that only fulfil FCC4), which is shown in Figure 4.7. The fraction of not fused target is then 39%, which is closer to the other data sets. From this, it can be concluded that the higher fraction of false interventions in data set A derives from ghost targets that are not fused.

In Figure 4.8, there is a visible tendency of low distance (5-20 m) to the target at the start of the true interventions, which seems accurate since a brake intervention should only by initiated when a collision is considered inevitable, and thus the distance should not be very high. In the false interventions, the distance is most often below 10 meters in all of the four data sets. The tendency of short distance could be correlated to the short life length in the false interventions (see Figures 4.5b, 4.5d, 4.5f and 4.5h), which both are typical behaviours for ghost targets. In other words, the false targets are detected when they are close and the brake intervention is initiated immediately. There are also many false interventions with high distances (above 40 meters), in contrast to the true interventions. These probably derive from almost stationary ghost targets detected at a high host speed, i.e. a high relative speed.

The longitudinal speed of the host varies greatly for both true and false interventions. This probably indicates that collision-critical scenarios occur no matter what speed the HGV is driven in, and thereby most road types are probably represented. The fact that few interventions occur at higher speeds than 90 km/h is probably due to that HGVs are not allowed to travel faster than that. Very few interventions occurred below 15 km/h, which not is so surprising since the first two versions of Volvo's AEBS are limited to 15 km/h while only the most recent version can brakes at host speeds down to 5 km/h.

That the speed of the target in the true interventions is most frequently 10-30 km/h seems reasonable since the host is approaching the target when an intervention is initiated, and since the system aims to prevent rear-end collisions. In the false interventions, the speed tends to be at 0-5 km/h, which could for example be due to inaccurate measurements of stationary objects, i.e. objects that the HGV would pass and that the system thus should not brake for.

In the true interventions, the relative speed between the host and the target is most often in the interval 10-15 km/h. This seems accurate since the host should be approaching the target at the initiation of the brake intervention. Among the false interventions, there is higher tendency of high relative speeds. This could be due to stationary targets that are interpreted to be moving. In general, it seems harder for the system to correctly interpret imminent collisions with stationary targets.

Most true and false interventions occur at daytime, which most likely is a consequence of a higher number of HGVs driving during daytime, but also that there is more traffic on the roads. Since the lighting condition at different times of the day depends on the time of the year, it is difficult to draw conclusions about whether or not the lighting conditions affects the frequency of interventions.

Since radar ghost targets caused by the road environment possibly are independent of the amount of traffic or lighting conditions, the likelihood of these might reflect how many HGVs are driving at each time. If they do, the distribution of false interventions should correspond to which hours HGVs are driving. Then, the relative difference between the histograms for true and false interventions in Figure 4.14, i.e. the difference in fraction of true and false interventions at each hour of the day, divided by the fraction of false, should give an indication of which hours the likelihood for true interventions is higher or lower. For example, the true interventions have a lower value than the false in the night, early morning, and in the evening but a higher value around 07-08 in the morning, at lunchtime and around 15-16. Thus, the likelihood of being close to colliding might be lower or higher, respectively, at these hours of the day. One possible explanation to the increased risk is that drivers are hungry and have driven a long time without break before lunch, and thus are more tired or inattentive. It could also be a consequence of HGVs driving off the highways to smaller roads in order to reach a lunch restaurant, and that this kind of road environment possess a higher risk of critical situations. In the night time, there is probably fewer vehicles on the roads and it is easier to keep the distance to nearby vehicles.

A simple collision detection script was run on all data sets, which compared the TTC and the time required to decelerate the relative velocity when assuming a maximum deceleration of 8 m/s^2. Logs that fulfilled this criterion in multiple frames were analysed closer, but only a handful seemed to possibly contain an actual collision, and none of them seemed serious. Therefore, the conclusion can be made that either the system does avoid collisions rather than mitigate them, or that the true interventions which not fully avoid the collisions unintentionally were excluded from the data sets.

Overall, with a fused, moving target that has been tracked for at least a few seconds, the system seems to work as it should, and thus prevent many rear-end collisions. This conclusion seems to hold irrespective of the longitudinal velocity of the host, but interventions seem to be most common at a host velocity of around 40 km/h. The reason for this could be the circumstances when driving at this speed, e.g. the type of road and how heavy traffic there is.

In the false positives, the target was often newly detected and had a low speed, and the proportion of not fused targets was higher than among the true interventions. These attributes indicate the existence of ghost targets among the false interventions. There is also a higher tendency of low host speed in the false interventions. In these cases, the target was most often found to be stationary. In general, the system seems to have a harder time to properly identify stationary targets (see Figure 4.11).

Data set D, i.e. the data set most suitable for the analysis, was found to have a low proportion of false interventions. However, it is still important to not only consider the proportion of false interventions, but also focus on the actual number of true interventions since these have potentially prevented or mitigated a rear-end collision. Also, the false interventions tend to be short and the speed reduction of the host tend to be low, and thus, the false intervention will most likely not cause any serious consequences or affect the surrounding traffic noticeably. It could cause nuisance for the driver, but the benefits from the true interventions will probably be of more weight than this potential nuisance.

5.5 Radar analysis

From the literature study, some possible weaknesses of the radar were derived. These were then kept in mind when analysing the locations where false interventions had occurred. During the analysis it turned out that the logged GPS position might deviate from the actual position, and possibly with a larger deviation in areas without open sky (in tunnels, underneath bridges etc.). The deviation was assumed to be smaller than 20 meters.

To verify this assumption, the logged GPS coordinates were checked in maps, and in most cases they corresponded to positions on the roadway. The path travelled by the host was visualised (as described in Figure 3.2) to compare with the road characteristics shown in the maps. For example, where the intervention occurred relative to a curve in the path visualisation could be compared with the logged GPS coordinates appear relative to the same curve in the maps. From this verification, it was clear that the GPS coordinate and the actual position of the intervention seem to match well. Thus, it is reasonable to assume that the estimation of a maximal GPS position deviation of 20 meters is correct. Moreover, the infrastructure in the geographic area of data set A sometimes contained flyover roads, i.e. a road on a bridge passing above another road, which means both the roads have the same GPS latitude and longitude. In addition, the distance threshold for two interventions to be considered to take place at the same location was set to above 100 m. Altogether, these factors result in difficulties to state which roadway the host travelled on at the moment of the intervention (if there are two or more roads next to each other, which was the case for most of the locations). The possible position deviation had to be accounted for. Sometimes the heading direction could be used to determine the correct roadway. In other cases the nearby roadways were also analysed, which is the reason why there were more images in Figures 4.15, 4.16 and 4.17 in this report than locations with multiple interventions. Besides, not all false interventions are due to ghost targets, but rather an inaccurate measurement of an actual vehicle. With that in mind, the analysis of the environment proceeded.

Three different categories of road environments were identified to probably be problematic for the radar, as described in Section 4.13. Since the environments in each category look similar, it is reasonable to assume that the road infrastructure is the main reason behind the false interventions. However, since multiple roads were analysed, it might be that only one of category 1 and 2 caused the interventions. If that is the case, the first category is assumed to have a higher likelihood of causing problems due to the larger fraction of the field of view containing metal objects, and the large number of corners found on the underside of the elevated roads.

To strengthen the conclusion that the road environment causes false intervention, the environment of 50 true interventions was also investigated. Most of the analysed true interventions had occurred on a rural road near an intersection or on a highway, without the elements found in the case of false interventions, i.e. no elevated roads, tunnels or flat metal walls beside the road. Therefore, these metal structures are assumed to be the cause of ghost targets.

The large metal structures in the road environment that cause ghost targets are in general stationary, and the AEBS system used in Volvo HGVs only brakes for stationary objects if they are fused (i.e. detected by both the camera and radar). A radar-only target (i.e. a target not detected by the camera) will only cause a brake intervention if it has been seen moving. When analysing the 6 interventions that had occurred at the same location, none of them were caused by a fused target and they were all interpreted as moving forward. Therefore, we came up with the idea that there might be a way that ghost targets appear as moving even though they derive from stationary structures. One way that this could occur was found, as described below.

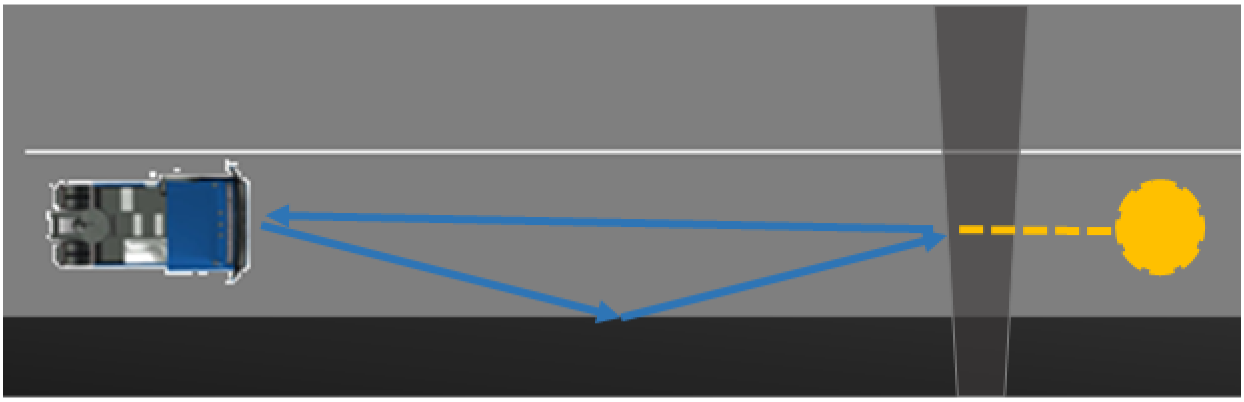

Figure 5.1: *A schematic image how ghost targets can be created due to secondary reflections and side lobes. The blue arrows show how the radiation is reflected at a wall and an object almost perpendicular to the road, whereas the yellow dotted line and circle show the radar interpretation of the radiation and reflection object. The interpreted object appears as further away than the actual reflecting object.*

A secondary reflection can occur if there for example is a parallel wall next to the road and a metal beam close to perpendicular to the road but angled a little towards the wall. A schematic figure of this is shown in Figure 5.1. This type of environment was found in the two lower images in Figure 4.15 and also in the tunnel in Figure 4.17, with the perpendicular object placed above the roadway. There is a risk that the radiation first is reflected at the parallel object at the side of the road and then reflected at the perpendicular object, and finally returning to the radar again. This is the type of secondary reflection that is shown in the right part of Figure 2.4. The radar has no way to determine that a secondary reflection has occurred, and can only determine the angle from which the radiation is received and estimated the distance travelled by the radiation (based on the assumption of a primary reflection). Since the secondary reflection has travelled a somewhat longer path than a primary reflection would, the radar interprets it as an object at the same angle as the perpendicular object (i.e. straight ahead), but at a larger distance, denoted as a *position discrepancy.*

The position discrepancy, denoted as $\hat{d}$ derives from the radiation travelling hypotenuses while the object causing the latest reflection is straight ahead of the HGV. The hypotenuses derive from both lateral and vertical distances between the HGV and the parallel wall and perpendicular reflection. If the longitudinal distance between the HGV and the perpendicular object is denoted d_{long}, the lateral distance between the center of the HGV and the parallel wall is denoted as d_{lat}, the vertical distance between the radar and the perpendicular object is denoted as d_{vert} and the longitudinal distance interpreted by the radar is denoted as d_{int}, their relation can be written

$$2d_{int} = \sqrt{d_{long}^2 + 2d_{lat}^2 + d_{vert}^2} + \sqrt{d_{long}^2 + d_{vert}^2}. \tag{5.1}$$

Dividing with d_{long} gives

$$2\frac{d_{int}}{d_{long}} = \sqrt{1 + 2\frac{d_{lat}^2}{d_{long}^2} + \frac{d_{vert}^2}{d_{long}^2}} + \sqrt{1 + \frac{d_{vert}^2}{d_{long}^2}}. \tag{5.2}$$

Taylor expansion can now be used for the square roots, which yields

$$2\frac{d_{int}}{d_{long}} = 1 + \frac{1}{2}\left(2\frac{d_{lat}^2}{d_{long}^2} + \frac{d_{vert}^2}{d_{long}^2}\right) + 1 + \frac{1}{2}\frac{d_{vert}^2}{d_{long}^2} + \mathcal{O}(d^4), \tag{5.3}$$

where $\mathcal{O}(d^4)$ is an error term that in this case can be neglected in as long as d_{long} is larger than or equal to d_{vert} and d_{lat}. The remaining terms can be written as

$$2\frac{d_{int}}{d_{long}} = 2 + \frac{d_{lat}^2}{d_{long}^2} + \frac{d_{vert}^2}{d_{long}^2} = 2 + \frac{d_{lat}^2 + d_{vert}^2}{d_{long}^2}. \tag{5.4}$$

The sum $d_{offset}^2 = d_{lat}^2 + d_{vert}^2$ is the squared hypotenuse created from the lateral and vertical distances, i.e. if the HGV was positioned at the perpendicular object, d_{offset} would be the distance between the radar and the corner where the parallel wall and the perpendicular beam are connected. Using this, the equation can now be written as

$$d_{int} = d_{long}\left(1 + \frac{d_{offset}^2}{2d_{long}^2}\right). \tag{5.5}$$

The position discrepancy, defined as $\hat{d} = d_{int} - d_{long}$, is the difference between the interpreted distance and the actual distance, i.e.

$$\hat{d} = \frac{d_{offset}^2}{2d_{long}}. \tag{5.6}$$

From equation 5.6, one can directly see that $\hat{d}$ is proportional to the inverse of the longitudinal distance, and thus will increase as the host approaches the perpendicular object. Important to note is that this conclusion is made based on the assumption that the d_{long} is larger than d_{offset}. When the host is moving forward, d_{long} will reduce according to the speed of the host v_{long} whereas d_{offset} will remain constant. Using the time derivative, one can find the interpreted velocity $\hat{v}$ of the target as

$$\hat{v} = \frac{\partial \hat{d}}{\partial t} = \frac{\partial \hat{d}}{\partial d_{long}} \frac{\partial d_{long}}{\partial t} = -\frac{d_{offset}^2}{2d_{long}^2}(-v_{long)} = v_{long}\frac{d_{offset}^2}{2d_{long}^2}. \tag{5.7}$$

Equation 5.7 shows a *velocity discrepancy* effect that increases when the longitudinal distance to the target decreases. As Equations 5.1 - 5.7 are based on our own assumptions and calculations, an attempt was made to assess their reasonability. Using the values $v_{long} = 15$ m/s and $d_{long} = 2d_{offset}$ yields a velocity discrepancy of 1.875 m/s, which could be considered as a moving object. Thus, the host speed might cause ghost targets to appear as moving away from the host even though they are stationary.

A pilot test was performed where the values of d_{long} and v_{long} were taken from logs of false interventions at the locations where multiple interventions had occurred, and d_{offset} was estimated from the Street View images. Only a handful of logs were analysed and the velocity discrepancy was mostly in the same magnitude as the target velocity in the logs, i.e. $\hat{v}$, but did not match perfectly. Therefore, it is possible that the hypothesis of secondary reflections causing a speed discrepancy is correct, but not the only factor. The assumptions used in this simple model might deviate somewhat from real world scenarios. Besides, the details in the velocity calculations of the radar used were not available, and thus, it is not possible to exclude that more factors are taken into account when the target velocity is determined. The speed discrepancy phenomenon might possibly be avoided if the velocity is measured using only the Doppler effect, but whether this is the case or not is unknown.

Normally the radar antenna has a low gain for high elevation angles, which would make it unlikely that reflections from objects above the roadway cause false interventions. However, there is a risk that side lobes affect both the transmission and reception of radiation. Since the secondary reflection causes the angle of the received radiation to be different than the transmitted (in contrast to primary reflections), different side lobes might coincide.

One possible way to reduce false brake interventions due to ghost targets, is to use a large amount of naturalistic driving data to determine what causes the ghost targets. With this knowledge, it might be possible to make the system recognise unwanted reflections and learn to distinguish these from relevant reflections. It could also be possible to use this data for an improved beamforming that reduces the antenna gain in the directions where many ghost targets are found.

As described in Section 2.3, the weather might influence the performance of the radar in terms of range capacity and reduced angular accuracy. The data used in this project did not contain variables that can be used for analysis of the impact of the weather. No other findings were made regarding temporarily reduced performance of the radar.

5.6 Future work

A number of variables were chosen to be used in this analysis. However, there are more variables and combinations of variables that could be interesting to analyse, to be able to draw further conclusions about the performance of AEBS, i.e. when the systems intervenes correctly and incorrectly respectively. A couple of examples of possibly interesting areas are the road type, the speed limit of the road and the infrastructure where the interventions occur. Another interesting aspect could be to analyse the performance of AEBS in combination with other safety systems, e.g. ACC or Driver Alert Support (DAS), to find if AEBS intervenes more or less frequently when the other system is activated.

However, some variables can be difficult to analyse and draw conclusions about. This was for example the case for the variable time of day. Clearly, there are more interventions during day time since there are both more HGVs driving during the day and also more traffic on the roads in general, and thereby a higher probability of an intervention. This was also the result from the analysis (see Figure 4.14). However, there seems to be a drop at 12 pm among the false interventions, and the cause for this could be of interest to investigate further. Also, the false interventions tend to have a higher proportion during the late night than the true interventions, which could be due to darkness. If this is the case is also suggested as future work.

The output of the TFC is useful for further analysis of the system. The logs in the data set used as input in the TFC can easily be filtered and sorted by classification, what FCC is fulfilled as well as different variable values included in the output. Since the fraction of false interventions was found to be low, the program is also useful to detect the false interventions. These logs can then be manually analysed to find what caused the false intervention, which is very useful information to further develop the AEBS system.

Another interesting analysis would have been to investigate the circumstances of the interventions fulfilling each false classification criterion, to be able to find differences among the false interventions. Unfortunately, the data sets A, B and C were too small to use for this purpose, and data set D was obtained too late to perform this analysis within the scope of this thesis. Since it can give a deeper insight in the scenarios where AEBS intervenes, it is recommended as future work.

Finally, it would be interesting to experimentally verify if the hypotheses of how secondary reflection can cause stationary targets to appear as moving is correct or not. This could be done on a test track with a metal wall beside the road and a flat metal surface above the road, close to perpendicular to the direction of travel.

6 Conclusions

This thesis project consisted of data analysis of actual interventions with Volvo's AEBS system. The thesis was divided into three main parts. The first part was to develop a program that can determine if an intervention was true or false based on logged information from the intervention. The second part consisted of an analysis of data from a large number of logs, to better understand the characteristics of true and false interventions, and in what circumstances the true and false interventions occur respectively. The third part was to study the principles of a radar and analyse the road environments to investigate if radar ghost targets are a main contributor to the false interventions and determine what causes the ghost targets.

The classification of logs resulted in a program with four different criteria which are described in Table 4.1. The program performed well, and the proportion of incorrect classifications is assumed to be below 3%. From the output of the program, fractions of true and false AEBS interventions could be derived, and a grey zone was used for logs that were difficult to classify. Due to confidentiality, the fractions of true and false interventions cannot be disclosed, but the false rate was in general found to be low. The false rate depends on the road infrastructure, and will therefore differ depending on the geographic location.

In the analysis of true interventions, conclusions were made that they vary in duration and thus generate different speed reductions, but a large fraction generate a significant speed reduction. This gives a strong indication that AEBS can avoid many collisions. In contrast, most of the false interventions have a short duration and a vast majority generate a speed reduction less than 5 km/h, and are thus not severe brakings. Conclusions were drawn that the false interventions are very unlikely to cause a collision. Most true interventions are due to the host getting too close to a vehicle that has been tracked by the sensors for a long time, whereas most false interventions were due to objects that had often been tracked for less than 3 seconds. The false interventions were in most cases caused by stationary or slowly moving targets, which also resulted in very high relative velocities. The true interventions are more likely to occur with a higher target velocity, and the relative velocity is most often between 5 and 20 km/h. Both true and false intervention mainly occur at daytime. Most true interventions are initiated when the time to collision is between 0.5 and 2 seconds, but for false interventions it varies more.

The radar analysis resulted in three road environment categories that are likely to cause false interventions: roads below bridges, roads with flat metal walls and roads inside a tunnel. The first and the last category, i.e. roads below bridges and in tunnels, are believed to be the most likely to cause false interventions since in those environments there is a higher likelihood that secondary reflections cause stationary objects in the road environment appear as moving.

References

[1] European Commission. *Heavy goods vehicles*. 2019. URL: https://ec.europa.eu/transport/road_safety/specialist/knowledge/vehicle/safety_design_needs/heavy_goods_vehicles_en (visited on 02/07/2019).

[2] World Health Organization. *The top 10 causes of death*. URL: https://www.who.int/news-room/fact-sheets/detail/the-top-10-causes-of-death (visited on 06/13/2019).

[3] European Road Safety Observatory. *Annual Accident Report 2018*. 2018. URL: https://ec.europa.eu/transport/road_safety/sites/roadsafety/files/pdf/statistics/dacota/asr2018.pdf (visited on 03/20/2019).

[4] European Road Safety Observatory. *Traffic Safety Basic Facts 2018 - Heavy Goods Vehicles and Buses*. 2018. URL: https://ec.europa.eu/transport/road_safety/sites/roadsafety/files/pdf/statistics/dacota/bfs2018_hgvs.pdf (visited on 03/20/2019).

[5] European Commission. *Commission Regulation (EU) No 347/2012 of 16 April 2012 implementing Regulation (EC) No 661/2009 of the European Parliament and of the Council with respect to type-approval requirements for certain categories of motor vehicles with regard to advanced emergency braking systems*. 2012. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32012R0347 (visited on 02/22/2019).

[6] Insurance Information Institute. *Facts + Statistics: Highway safety*. 2019. URL: https://www.iii.org/fact-statistic/facts-statistics-highway-safety (visited on 03/26/2019).

[7] National Highway Traffic Safety Administration. Analyses of Rear-End Crashes and Near-Crashes in the 100-Car Naturalistic Driving Study to Support Rear-Signaling Countermeasure Development (Oct. 2007).

[8] N. Fecher et al. Analysis of driver behaviour in autonomous emergency hazard braking situations (2008), 1–10.

[9] S. Kockum et al. *Volvo Trucks Safety Report 2017*. May 2017.

[10] AB Volvo. *Emergency Brake - a system that saves lives*. 2017. URL: https://www.volvotrucks.com/en-en/news/volvo-trucks-magazine/2017/jul/tech-focus-emergency-brake.html (visited on 02/15/2019).

[11] L. Decoster. *Collision Warning with Emergency Braking*. PowerPoint (internal source).

[12] Volvo Trucks. *Volvo Trucks - Emergency braking at its best!* Video. URL: https://www.youtube.com/watch?v=ridS396W2BY (visited on 05/30/2019).

[13] K. Grove et al. Field study of heavy-vehicle crash avoidance systems. *National Highway Traffic Safety Administration* (June 2016).

[14] I. Isaksson-Hellman and M. Lindman. Evaluation of the crash mitigation effect of low-speed automated emergency braking systems based on insurance claims data. *Traffic Injury Prevention* (2016). ISSN: 1538-9588. DOI: 10.1080/15389588.2016.1186802. URL: https://www-tandfonline-com.proxy.lib.chalmers.se/doi/pdf/10.1080/15389588.2016.1186802.

[15] I. Isaksson-Hellman and M. Lindman. Using Insurance Claims Data to Evaluate the Collision-Avoidance and Crash-Mitigating Effects of Collision Warning and Brake Support Combined with Adaptive Cruise Control. *2016 IEEE Intelligent Vehicles Symposium (IV)* (2016). URL: https://ieeexplore.ieee.org/document/7535538.

[16] B. Fildes et al. Effectiveness of low speed autonomous emergency braking in real-world rear-end crashes. *Accident Analysis and Prevention* (2015).

[17] J. Saadé. *Autonomous Emergency Braking AEB (city, inter-urban)*. 2017. URL: www.roadsafety-dss.eu (visited on 01/24/2019).

[18] Y. Jammes et al. Emergency braking is affected by the use of cruise control. *Traffic Injury Prevention* **18**.6 (2017), 636–641. ISSN: 1538-957X. DOI: 10.1080/15389588.2016.1274978. URL: https://doi.org/10.1080/15389588.2016.1274978.

[19] H. Winner. Automotive RADAR. *Handbook of Driver Assistance Systems* (2016), 325. ISSN: 9783319123516. URL: http://proxy.lib.chalmers.se/login?url=http://search.ebscohost.com.proxy.lib.chalmers.se/login.aspx?direct=true&db=edb&AN=112579552&site=eds-live&scope=site.

[20] radartutorial.eu. *Radar Basics*. URL: http://www.radartutorial.eu/01.basics/The%20Radar%20Range%20Equation.en.html (visited on 05/13/2019).

[21] Google Maps. URL: https://www.google.com/maps/ (visited on 05/29/2019).