Thesis for the Degree of Doctor of Philosophy

Computational Verification Methods for Automotive Safety Systems

Jonas Nilsson

Department of Signals and Systems
Chalmers University of Technology
Göteborg, Sweden 2014

Computational Verification Methods for Automotive Safety Systems
JONAS NILSSON
ISBN 978-91-7385-973-8

Doktorsavhandlingar vid Chalmers tekniska högskola
Ny serie nr 3654
ISSN 0346-718X

Department of Signals and Systems
Mechatronics Group
CHALMERS UNIVERSITY OF TECHNOLOGY
SE–412 96 Göteborg
Sweden
Telephone: +46 (0)31 – 772 1000

Jonas Nilsson
Telephone: +46 (0)70 – 266 5397
Email: jonas.nilsson@volvocars.com
jonas.nilsson@chalmers.se

Typeset by the author using LaTeX.

Chalmers Reproservice
Göteborg, Sweden 2014

to Monica

Abstract

This thesis considers computational methods for analysis and verification of the class of automotive safety systems which support the driver by monitoring the vehicle and its surroundings, identifying hazardous situations and actively intervening to prevent or mitigate consequences of accidents. Verification of these systems poses a major challenge, since system decisions are based on remote sensing of the surrounding environment and incorrect decisions are only rarely accepted by the driver. Thus, the system must make correct decisions, in a wide variety of traffic scenarios. There are two main contributions of this thesis. First, theoretical analysis and verification methods are presented which investigate in what scenarios, and for what sensor errors, the absence of incorrect system decisions may be guaranteed. Furthermore, methods are proposed for analyzing the frequency of incorrect decisions, including the sensitivity to sensor errors, using experimental data. The second major contribution is a novel computational framework for determining the errors of mobile computer vision systems, which is one of the most widely used sensor technologies in automotive safety systems. Augmented photo-realistic images, generated by rendering virtual objects onto a real image background, are used as input to the computer vision system to be tested. Since the objects are virtual, ground truth is readily available and varying the image content by adding different virtual objects is straightforward, making the proposed framework flexible and efficient. The framework is used for both performance evaluation and for training object classifiers.

Keywords: Automotive, Active Safety, Semi-Autonomous Vehicles, Verification, Performance Evaluation, Decision Making, Augmented Reality.

Acknowledgments

Writing a Ph.D. thesis is a bit like hiking in the mountains. The view is beautiful and inspiring but as one strives to the top, there is some real physical pain involved. You pin your hope on the summit resting right behind the next crest, only to find yet another crest. Anyhow, for a person appreciating cardiovascular exercise, the journey has been truly enjoyable.

Many people have accompanied me on this journey and deserve my deepest gratitude. My academic supervisor, Dr. Jonas Fredriksson, has always been available for discussions (often on topics related to cross-country skiing), and has helped me not to loose sight of the bigger picture. Thank you for your commitment. My industrial supervisor, Dr. Anders Ödblom, has continuously supported me using his impeccable eye for details. For this I thank you. Together I think we have finally understood what this project is really about. Thank you also to Prof. Jonas Sjöberg for welcoming me into his research group and occasionally allowing me to beat him in a race.

From the many colleagues at Volvo Cars and Chalmers contributing to a stimulating work environment, I would like to mention a few. My current and former managers, Georgios Minos and Peter Janevik respectively, for giving me the freedom needed to work efficiently in between academia and industry. Dr. Mohammad Ali, for fruitful collaboration and enriching discussions. The room got awfully quiet without you. Co-authors Adeel Zafar, Patrik Andersson and Prof. Irene Gu, I hope we can collaborate again in the future. Dr. Erik Coelingh, Dr. Mattias Brännström and Jonas Ekmark, for indirect contributions through your confidence and support. Dr. Fredrik Persson, Feng Liu and David Hultberg, for your committed work on VCAV. Thank you also to Vinnova, SAFER and Volvo Cars for funding this project.

Last but most important I thank my family, starting with my siblings Fredrik and Pernilla for proofreading this thesis. Monica, for your everlasting support and your complementing talents, bringing balance to our family. Ludvig and Valdemar, for showing me there is joy in every detail in life and for being inexhaustible sources of ideas on new forms of transportation.

Jonas Nilsson
Göteborg, January 2014

Contents

II Included papers

Part I

Introductory Chapters

Chapter 1

Introduction

Road traffic accidents are a global problem of epidemic proportions. According to the World Health Organization (WHO), road traffic injuries are the leading cause of death globally for young people aged $15 - 29$, and the eight leading cause of death in total, [1]. In the developed countries primarily, road traffic accidents have been on the agenda in the past few decades. Governments have invested in infrastructure and passed laws to improve road safety. The automotive industry has put emphasis on designing systems that protect the occupants of the vehicle in case of a crash, so called *passive safety* systems. Passive safety innovations include seat belts, crumple zones and airbags.

In the 1970s, the introduction of Anti-lock Braking Systems (ABS) marked a first milestone for *active safety* systems, i.e. systems which actively intervene to prevent or mitigate consequences of accidents. In recent years, active safety systems which monitor the surrounding environment, using remote sensing technologies, have been introduced to the market. By using information on the surrounding traffic environment, systems can identify hazardous situations, e.g. when the driver has failed to observe a crossing pedestrian and a collision is imminent. If and when hazardous situations are detected, the system can actively intervene to prevent an accident either by informing the driver of the upcoming danger or by autonomously performing an evasive maneuver such as Automatic Emergency Braking (AEB).

This thesis concerns the problem of verifying that a given active safety system acts correctly in the wide variety of possible traffic scenarios. There are two major reasons why this is a challenging task. First, the variations in operating conditions are essentially unlimited, a fact easily acknowledged when reflecting and comparing a snowy country road in northern Sweden to downtown Tokyo. Second, incorrect decisions by highly intrusive systems, like AEB, can only be accepted on very rare occasions.

1.1 Aims and Objectives

The aim of the work presented in this thesis is to develop computational methods for efficient verification of automotive safety systems. In this context, computational verification methods are defined as methods which predict system performance by performing computations with recorded experimental data and/or mathematical models as input.

In active safety systems, decision functions use input from sensors to decide how to appropriately support the driver. A vital part of active safety system performance is the ability to make correct decisions, also in the presence of sensor measurement errors. Consequently, three objectives are formulated, namely to develop methods that

I. For a given active safety decision function, identify traffic scenarios where the function makes incorrect decisions

II. For a given active safety decision function, quantify the robustness to input errors

III. Generate virtual sensor data with sufficient quality for analysis and verification

The first two objectives are addressed by Papers 1-3, while the third objective is treated in Papers 4-6.

1.2 Delimitations

This thesis is concerned with semi-autonomous vehicles where active safety systems monitor the traffic situation and intervene if needed to ensure safety. Objectives I and II are delimited to evaluating the correctness of the intervention decision as opposed to the choice and execution of the intervention. With regards to the same two objectives, only traffic scenarios with single moving objects are considered. In Objective II, we primarily consider input errors which are bounded and systematic, where systematic means that they depend on the specific traffic situation. Objective III is concerned with efficiently determining said input errors and is delimited to computer vision sensors, which is one of the dominating technologies used in active safety applications.

1.3 Thesis Outline

The thesis is divided into two parts. Part I serves as an introduction to Part II by presenting background information and related work. Part II contains six scientific papers that constitute the base of the thesis.

Part I provides context to the appended papers and is organized as follows. In Chapter 1, the topic of the thesis is introduced and aims, objectives and delimitations are described. Chapter 2 gives an overview of in-vehicle safety systems with a strong emphasis on active safety systems. In Chapter 3, an overview of methods for system verification is provided. Chapter 4 briefly summarizes the papers included in Part II while Chapter 5 presents the main scientific contributions and gives suggestions for future research.

1.4 List of Publications

This thesis is based on the following publications:

Paper 1

J. Nilsson, A. Ödblom and J. Fredriksson, Worst Case Analysis of Automotive Collision Avoidance Systems, submitted for possible journal publication.

Paper 2

J. Nilsson, J. Fredriksson and A. Ödblom, Verification of Collision Avoidance Systems using Reachability Analysis, submitted as invited paper to *the 19th IFAC World Congress*, Cape Town, South Africa, 2014.

Paper 3

J. Nilsson and M. Ali, Sensitivity Analysis and Tuning for Active Safety Systems, in *Proceedings of the 13th International IEEE Conference on Intelligent Transportation Systems*, 2010, pages 161-167, Madeira Island, Portugal.

Paper 4

J. Nilsson, A. Ödblom, J. Fredriksson, A. Zafar and F. Ahmed, Performance Evaluation Method for Mobile Computer Vision Systems using Augmented Reality, in *Proceedings of the IEEE Virtual Reality Conference*, 2010, pages 19-22, Waltham, Massachusetts, USA.

Paper 5

J. Nilsson, J. Fredriksson and A. Ödblom, Reliable Vehicle Pose Estimation using Vision and Single-Track Model, submitted for possible journal publication.

Paper 6

J. Nilsson, P. Andersson, I. Gu and J. Fredriksson, Pedestrian Detection using Augmented Training Data, submitted to *the 22nd International Conference on Pattern Recognition*, Stockholm, Sweden, 2014.

Other Publications

In addition to the publications above, the following publications by the thesis author are related to the topic of this thesis:

J. Nilsson, J. Fredriksson, and A. Ödblom, Bundle Adjustment using Single-Track Vehicle Model, in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2013, pp. 2888-2893.

J. Nilsson, A. Ödblom, J. Fredriksson, and A. Zafar, Using Augmentation Techniques for Performance Evaluation in Automotive Safety, in *Handbook of Augmented Reality*, 1st ed., B. Furht, Ed. Springer, 2011, pp. 631-649.

J. Nilsson, On Performance Evaluation of Automotive Active Safety Systems, Licentiate Thesis R014/2010, ISSN 1403-266X, Chalmers University of Technology, Göteborg, Sweden, 2010.

J. Nilsson and A. Ödblom, On Worst Case Performance of Collision Avoidance Systems, in *Proceedings of the IEEE Intelligent Vehicles Symposium*, 2010, pages 1084-1091, San Diego, California, USA.

A. Ödblom and J. Nilsson, Augmented Vision in Image Sequence Generated from a Moving Vehicle, Patent pending, EP2639771, European Patent Office, 2012.

J. Nilsson, Operating Method and System for Supporting Lane Keeping of a Vehicle, Patent granted, US8428821, U.S. Patent and Trademark Office, 2007.

J. Nilsson, Operating Method and System for Supporting Lane Keeping of a Vehicle, Patent granted, EP2188168, European Patent Office, 2007.

Chapter 2

Automotive Safety Systems

The over 1 million annual fatalities caused by road traffic accidents are merely the tip of the iceberg, e.g. the WHO estimates that road traffic accidents also lead to between 20 and 50 million non-fatal injuries each year, [1]. On top of that, the economic burden linked to road traffic accidents is significant. In 1998, a crude estimate of the annual global cost was found to be in the order of US$500 billion, [2].

There are large regional differences across the world as the variations in vehicle safety, infrastructure and driver education are substantial. Remarkable progress has been made in the developed countries during the last decades, as can be seen in Figure 2.1. Improved vehicle design, road infrastructure investments and road safety policies have contributed to reducing the risk of getting killed in traffic, in most developed countries, by more than 40% since 1990, [3].

Success in reducing fatalities has spurred stakeholders in road safety to set more and more ambitious goals, as described in [6]. The most ambitious goal possible, i.e. a vision of zero fatalities in road traffic, has been expressed in road safety policies in Sweden and the Netherlands. The current and future automotive safety systems discussed in this chapter have the potential to contribute significantly to this goal.

We categorize automotive safety systems into *passive safety* systems, which protect the vehicle occupants when collision has occurred, and three types of *active safety* systems, which are designed to prevent accidents. The first category of active safety systems, *vehicle dynamics control* systems, prevent unwanted dynamical behaviours such as instability. *Driver Assistance* (DA) systems monitor the vehicle surroundings to assist the driver. In a not too distant future, *Autonomous Driving* (AD) systems may take complete responsibility for the driving task. The line between these categories is by no means sharp, as exemplified by the Roadway Departure Prevention Assist (RDPA) system described in [7] which incorporate both

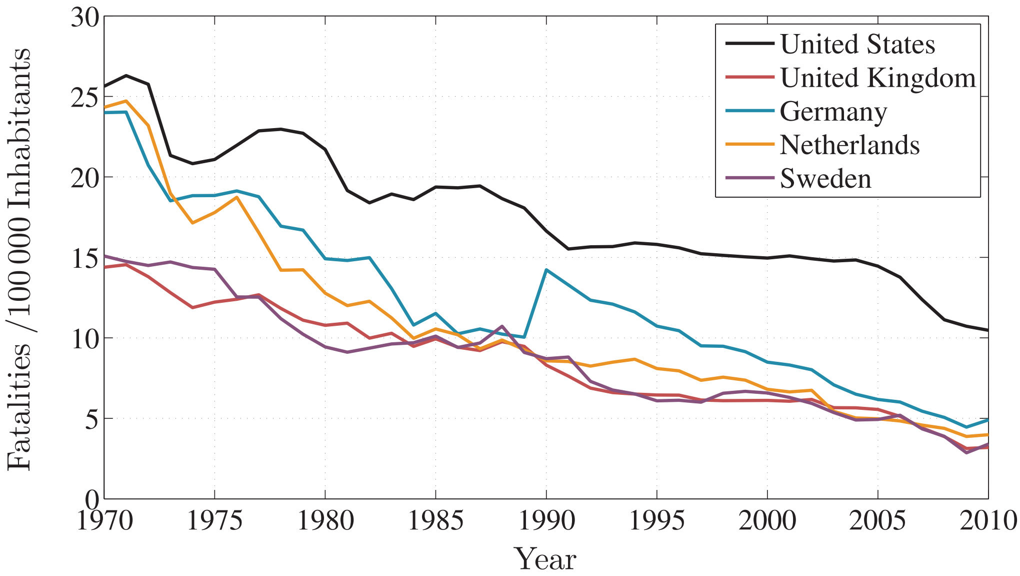

Figure 2.1: Historical road traffic fatalities, obtained from [4], for some of the developed countries. As a reference, low- and middle-income countries have annual road traffic fatalities of 18.3 and 20.1 per 100 000 inhabitants respectively, [1]. The sharp increase in fatalities for Germany in 1990 is an effect of the reunification of Germany, [5].

stability control and collision avoidance in a single framework.

In this chapter, the different categories of active safety systems are described, after first providing some context by briefly discussing the causes of accidents. In the final part of the chapter, the effectiveness of these systems is reviewed followed by a discussion on the challenges associated with system verification, which is the core problem addressed in this thesis.

2.1 Traffic Accident Causation

To efficiently prevent accidents, the causes of accidents need to be understood. A common approach for identifying accident causes is to study accident statistics. Figure 2.2 shows the accident distribution in terms of major crash types, obtained from [8]. There are numerous ways to classify accidents, e.g. by gender, age, type of vehicle, time of day or weather conditions. Extensive reports with accident classifications based on national accident statistics are published continuously, see e.g. [9] for the U.S. or [10] for Sweden.

Human error plays a major role in a majority of accidents. In an in depth

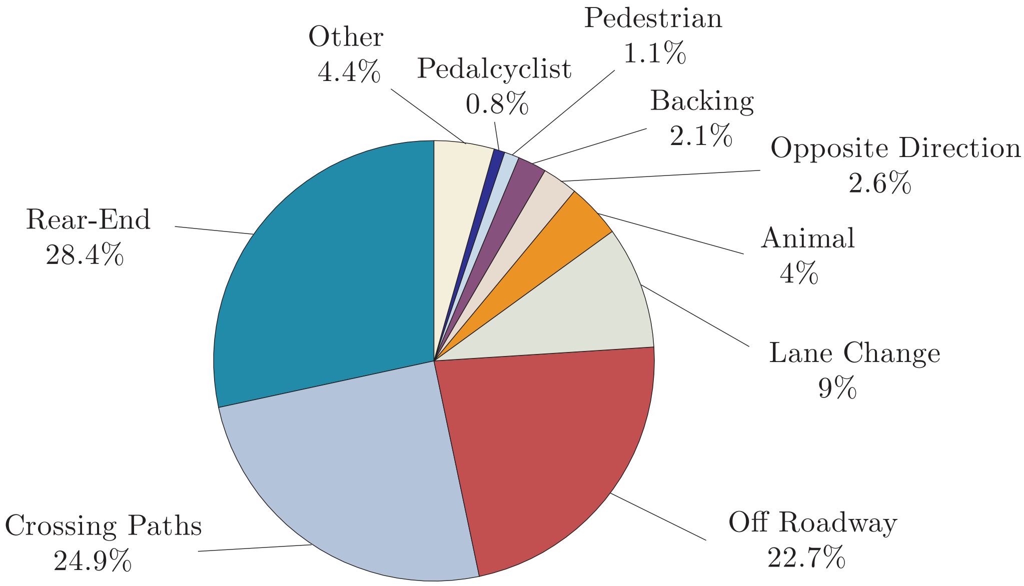

Figure 2.2: Distribution in terms of major crash types for all 6 394 000 police-reported motor vehicle crashes in the U.S. which resulted in 3 189 000 injured people and a total of 41 821 fatalities, [8]. The figure is based on statistics from the 2000 National Automotive Sampling System (NASS)/General Estimates System (GES) crash database.

study of real accidents in the 1970s, [11], including on-scene investigations, it was concluded that human participants were solely or partly to blame in 92.6% of the investigated accidents. The corresponding numbers for environmental and vehicular factors were 33.8% and 12.6% respectively. Common human errors were e.g. excessive speed, improper evasive action and driver inattention or distraction. Environmental factors were e.g. view obstructions and slippery road surfaces while vehicular factors included brake failures and inadequate tyre tread depth.

More recently, in 2005, a Field Operational Test (FOT) known as the 100-Car Study, [12, 13], was completed. 100 cars were equipped with unobtrusive data collection instrumentation to collect naturalistic data from normal driving. The study reaffirms that drivers are often to blame for accidents as nearly 80% of all crashes involved the driver looking away from the forward roadway just prior to the collision. Driver inattention or distraction, e.g. using a mobile phone while driving, does not necessarily lead to an accident but if coinciding with another unfortunate event, e.g. the vehicle in front suddenly braking, the probability of an accident increases significantly. Multiple accident causes mean that there are multiple possible preventive measures. As accidents are very diverse, preventing a majority of accidents requires the deployment of a large number of preventive measures.

2.2 Vehicle Dynamics Control

Following advances in electronics technology, mass production of ABS started on road vehicles in the 1970s but the innovation had been present in the railway and aviation industries decades before that. ABS monitors the rotational speed of the wheels and automatically reduce the brake force if the wheels cease to rotate, thus preventing brake lock-up. This enables steering of the vehicle while simultaneously braking hard.

In the 1990s, Electronic Stability Control (ESC) was introduced to handle problems with vehicle instability. ESC detects when the vehicle starts to skid and counteracts this by automatically braking the wheels individually, as illustrated in Figure 2.3. A natural evolution of ESC is to also prevent the vehicle from rolling over, as presented in [14]. Roll Stability Control (RSC) is mostly relevant for vehicles with high center of gravity, such as Sport Utility Vehicles (SUVs) and trucks, and was first introduced in 2002, [15].

The interested reader is referred to e.g. [16, 17], for more comprehensive treatments of vehicle dynamics control systems.

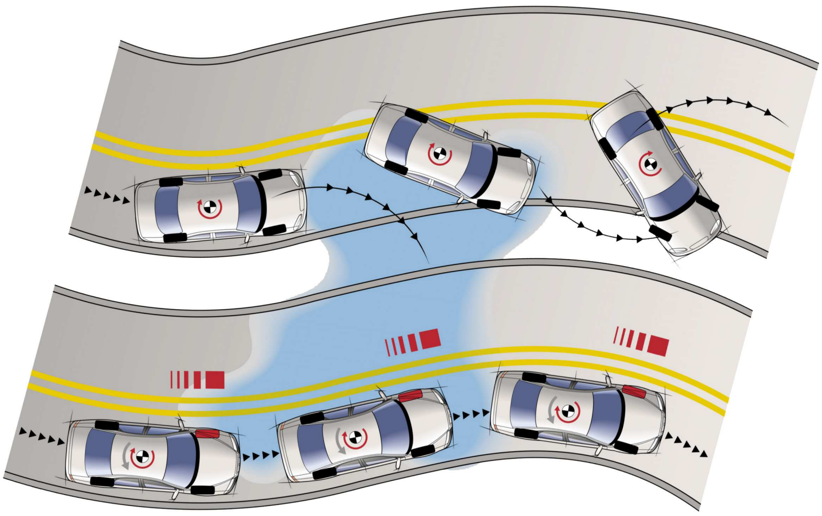

Figure 2.3: A vehicle drives onto an ice patch in a curve. Without ESC the vehicle becomes unstable and starts spinning. With ESC the left front wheel is braked, thereby counteracting the rotation, ensuring that stability is maintained.

2.3 Driver Assistance

Recent advances in remote sensing technology have led to the introduction of several DA systems, see e.g [17, 18] for extensive overviews. One of the first examples, launched in 1995, is an extension of the cruise control which automatically maintains a constant vehicle speed set by the driver. Adaptive Cruise Control (ACC), thoroughly described in [16,19], uses information from a forward looking sensor, e.g. a radar, to maintain a constant distance or time gap, set by the driver, to the vehicle in front of the host vehicle, see Figure 2.4. ACC contributes to safe driving by assuring that a safe distance is kept to the vehicle ahead. Also, ACC can reduce fuel consumption and congestion through smooth control of the brakes and throttle, thereby contributing to a cleaner environment.

Utilizing the same forward-looking sensor, Forward Collision Warning (FCW) indicates to the driver, as exemplified in Figure 2.5a, when imminent action is needed to avoid a collision, e.g. when the vehicle ahead suddenly brakes. If there is insufficient time or if the driver fails to respond to warnings, a Collision Avoidance (CA) system can autonomously control the vehicle to avoid the impending collision. A common action for CA systems is to automatically apply the brakes in situations where a collision is imminent, so called AEB, illustrated in Figure 2.5b. If the collision is unavoidable, AEB may still be triggered to reduce impact speed, so called Collision Mitigation (CM).

There are also numerous DA systems which support the lateral control of the vehicle, as illustrated in Figure 2.6. If the vehicle crosses a lane marking a Lane Departure Warning (LDW), [20], may be issued to the driver. A lane guidance system closely related to LDW is Lane Keeping Assistance (LKA), [16], where the driver is supported by a torque on the steering wheel to stay in the current lane. In [7] the problem of road or lane departures and vehicle stability are addressed in a common framework,

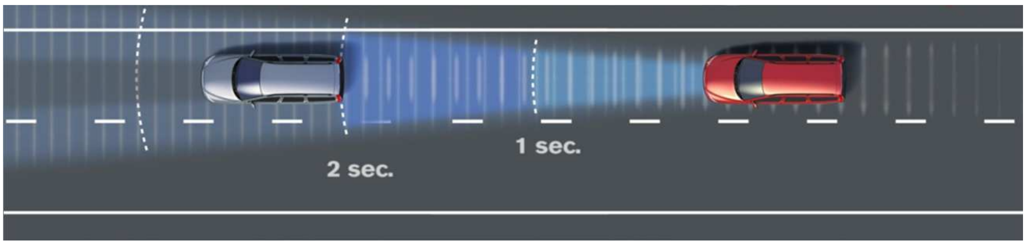

Figure 2.4: ACC automatically maintains a driver set time gap to the vehicle in front.

(a)

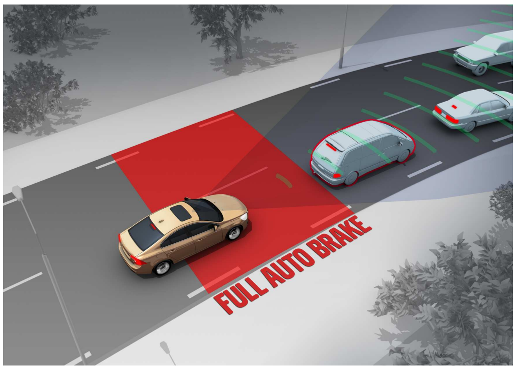

(b)

Figure 2.5: (a) FCW displayed in a Head Up Display (HUD). The red light displayed to the driver in the windshield is designed to resemble the appearance of vehicle brake lights. (b) When the host vehicle enters the red zone, an imminent collision is detected and an autonomous brake intervention is initiated.

thereby combining and enhancing the functionality of lane guidance systems and ESC. There are also systems that support the driver when performing lane change maneuvers. Lane Change Aid (LCA) systems, [21], monitor adjacent lanes and inform the driver when an obstacle is present in the blind spot of the rear view mirrors, see Figure 2.6b. In some situations there is very little, if any, time to warn the driver of a potential hazard, making it justified for a CA system to control the *steering* of the host vehicle to avoid accidents. A system designed to avoid collisions with oncoming traffic using steering interventions, referred to as Emergency Lane Assist (ELA), is presented in [22].

Information of host vehicle motion and road geometry can also be used to assess the present state of the driver. If a driver is fatigued, distracted or even impaired by drugs, this will affect the driver's ability to maneuver the vehicle smoothly in the current road lane. [23] presents a method for detecting inadequate driving behaviour, which can be used by systems to e.g. inform the driver when about to fall asleep.

The underlying technology for DA systems is discussed in the following subsections. DA systems are mechatronic systems and consist of three basic layers, namely the *perception*, *decision* and *action* layers. The architecture for a DA system performing autonomous interventions is illustrated in Figure 2.7.

(a) (b)

Figure 2.6: (a) A lane guidance system detects the lane markings and warns the driver (LDW), or applies a steering wheel torque (LKA), when crossing the lane boundary. (b) The colored zones visualize the blind spots, i.e. the zones not visible to the driver through the rear view mirrors. LCA indicates that an obstacle is present in the blind spot by lighting a small lamp close to the rear view mirror.

2.3.1 Sensor Technology

A key enabler for DA systems is reliable remote sensing technology. In the perception layer, see Figure 2.7, sensors collect observations from the environment, driver and host vehicle. Depending on the requirements imposed by the system, various technologies can be chosen to deliver an interpretation of the surrounding environment.

A frequently used sensor technology is computer vision, which detects and classifies objects in the environment using image data collected by cameras. Computer vision is the dominant technology to retrieve information on the road geometry and the relative position of the host vehicle to the road, which is done by detecting the lane markings or the edge of the road.

Active sensors such as radar, laser or ultrasonic sensors transmit radio, optical or sound signals and evaluate object attributes by interpreting the reflected response of the transmitted signal. Also, observations from digital maps and sensors mounted on other vehicles or infrastructure can be made available to the safety system through a communication device.

In many applications, system requirements cannot be fulfilled by a single sensor. Sensor observations from multiple sensors are combined, or fused, to provide an enhanced view of the environment. Also, objects observed by sensors are tracked over time to reduce the influence of noise. General frameworks for sensor data fusion and tracking are described in [24, 25] while [26–29] describe work tailored to DA systems.

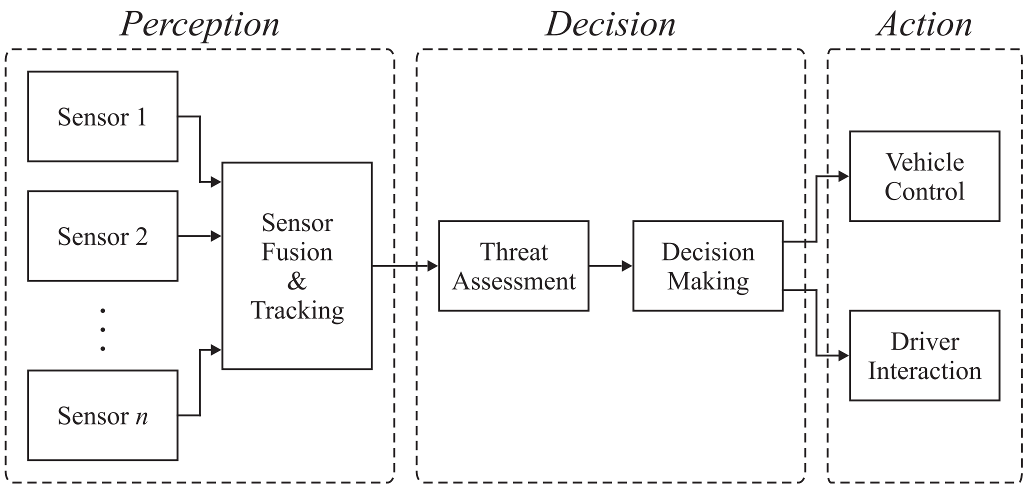

Figure 2.7: System architecture for an active safety system designed to intervene in case a critical situation arises. The *perception* layer provides information used for decision making in the *decision* layer. The decision is executed in the *action* layer via one or multiple actuators, e.g. brake system or driver information displays.

2.3.2 Decision-Making and Interventions

In the decision layer, see Figure 2.7, input from the perception layer is used to decide if and how to intervene. This *decision function* consists of two parts. The process of converting state estimations, e.g. object positions, into measures describing whether or not the host vehicle is in a hazardous situation, i.e. if surrounding road users and objects constitute a threat of collision, is termed *threat assessment*. Based on the threat measures, a *decision-making* algorithm chooses what, if any, action should be taken by the system.

The earlier, relative to the potential accident, the system intervenes, the more likely it is to prevent the accident. Also, the earlier the system intervenes, the more likely it is that the driver is well aware of the hazard and thus perfectly capable of preventing the accident. If the latter is true then the driver would consider the intervention *unnecessary*. Therefore, the aim of the decision function is usually to intervene at the latest point in time when the intervention type is still likely to succeed, where success is defined as e.g. preventing or mitigating the consequences of an accident.

A CA system aims to avoid all potential collisions. For lane guidance systems, the aim is not as straightforward to define since a lane departure not necessarily leads to a dangerous situation. Most LDW systems aim at issuing warnings exclusively when lane departures are unintentional. In situations when the driver intentionally deviates from the current lane, it is assumed that the driver can manage the situation.

There is a range of possible actions, or intervention types, which can be applied when a hazardous situation is detected. If the situation is detected early, the system, e.g. FCW or LDW, can warn the driver by for instance audible, visual or haptic feedback. In certain situations, there is no time for the driver to react to the feedback and perform a driving maneuver to avoid the impending accident. In those situations the system can, to avoid the accident, autonomously control the brakes or the steering.

System interventions are sometimes perceived as *intrusive* by the driver. The level of intrusiveness varies between intervention types where warnings or information to the driver are generally less intrusive than autonomous vehicle control. The amplitude of the intervention also has an influence as e.g. a loud warning signal is often considered more intrusive than a subtle warning signal. The possibility for the driver to override an intervention also affects the level of intrusiveness.

2.4 Autonomous Driving

Automotive safety systems which intervene autonomously to prevent accidents are currently commercially available from a large number of vehicle manufacturers. The systems are evolving to handle more and more operating scenarios such as intersections and night-time driving, and this trend is likely to continue, see Figure 2.8. An enabler for this evolution is the availability of more accurate, affordable remote sensors.

The research community has for quite some time focused on the next major step in automotive safety, namely *Autonomous Driving* systems. These are systems which takes full responsibility for the driving task as opposed to DA systems which still require the driver to monitor the system. In the 2007 *DARPA Urban Challenge*, [30], 35 teams formed from collaborations between industry and academia competed with driverless vehicles in an urban environment. A total of six self-driving vehicles completed the course which included tasks such as negotiating intersections, parking and avoiding vehicles stalled on the road.

In many ways AD systems are a natural evolution of DA systems and a number of companies, vehicle manufacturers and others, have communicated their aim to commercialize this technology. The potential benefit of AD systems is undoubtedly huge, not only in terms of safety, but also in terms of reduced fuel consumption, reduced congestion and added driver convenience.

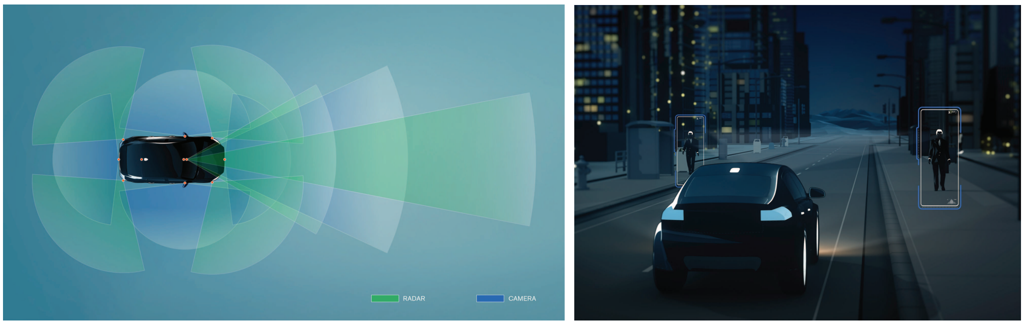

(a) (b)

Figure 2.8: (a) Possible sensor setup for future vehicles: 360° field of view with cameras and radars. (b) Autonomous vehicles and future driver assistance systems must handle more traffic scenarios, e.g. night conditions.

2.5 System Effectiveness

In the last decades, passive safety systems have made a major contribution to road traffic safety through innovations such as the safety belt, crumple zones and airbags, see Figure 2.9. Their effectiveness has been extensively studied using accident statistics. In the U.S. during 2008, according to [31], seat belts saved 13 250 lives, frontal airbags 2 546 and child restraints 244. In [32], it is shown that passive safety improvements have contributed to a significant decrease in injury severity between the 1970s and the 1990s, also when ignoring effects from seat belts and airbags.

The effectiveness of passive safety systems is assessed by governments around the world. In Europe, EuroNCAP has since 1997 assessed cars, by e.g. crash tests, in order to provide consumers with an independent rating of safety performance. Active safety systems such as ESC are included in this rating and in 2014 AEB will also be included, [33]. These ratings are important selling arguments for vehicle manufacturers and thus they encourage rapid development of new safety technology.

Vehicle dynamics control systems have been widely deployed in the market for many years, making it possible to assess their effectiveness in improving road safety using accident statistics. In [34], multiple studies investigating the safety impact of ABS are reviewed. A majority of the studies indicate that equipping vehicles with ABS significantly reduces the occurrence of accidents involving multiple vehicles. Some studies also indicate that ABS increases the occurrence of run-off-road accidents. Possible contributing factors to this increase include inappropriate use of ABS and driver behaviour adaption, e.g. when the driver decreases driving safety margins

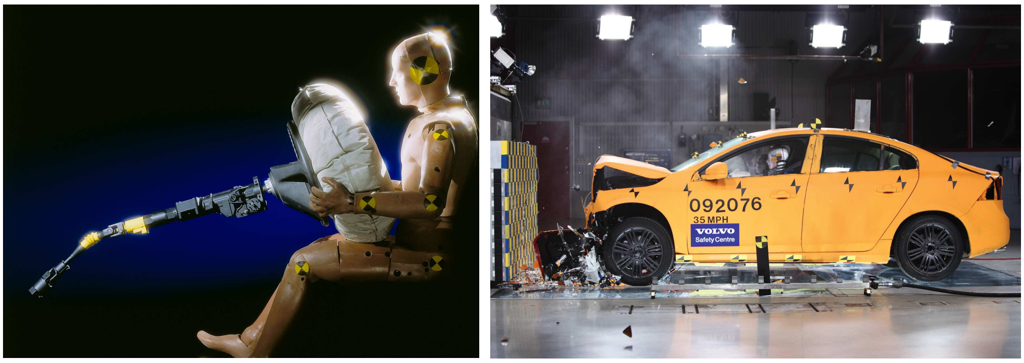

(a) (b)

Figure 2.9: (a) The airbag is an example of passive safety technology. (b) Crash tests are used to assess passive safety effectiveness.

due to awareness of the positive safety effects of ABS.

According to [35], use of ESC reduces fatal single-vehicle accidents involving cars and Sport Utility Vehicles (SUVs) by 30-50% and 50-70% respectively. Considering that single-vehicle crashes stand for 60% of all fatal crashes in the U.S., [31], the potential safety impact of ESC is significant. Additionally, the reduction of rollover accident fatalities, related to the use of ESC, is in [35] estimated to 70-90%, regardless of vehicle type.

DA systems have, if at all, been introduced to the market relatively recently which explains why their effectiveness has not been studied to the same extent. An overview on the subject is given in [17] which concludes that the safety impacts of DA systems are expected to be considerable. Due to the lack of data, several approaches have been proposed to predict system effectiveness such as reconstructing real-world accidents from accident databases and using simulations to determine if a given system could prevent these accidents. Using this method [36] predicts that a newly introduced CA system could prevent up to 24% of pedestrian fatalities and [37] predicts that a similar system could reduce driver fatalities in rear-end crashes by up to 50%.

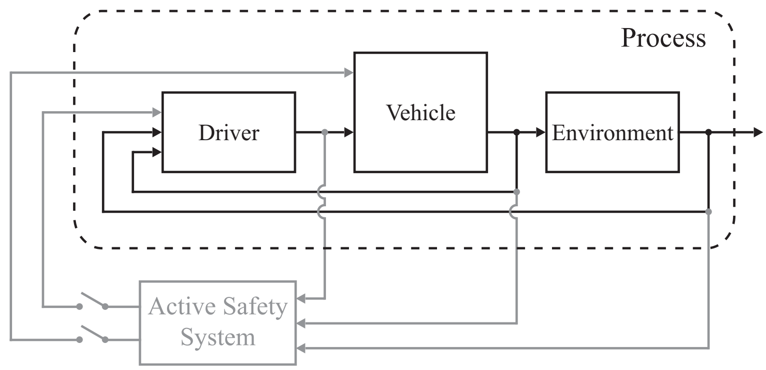

Figure 2.10: The driver monitors both the vehicle and the surrounding environment to control the vehicle. The active safety system monitors the complete process and control the vehicle either directly, or indirectly through driver interaction, see Figure 2.7. The switches determine if the system is executed in open- or closed-loop, see Section 3.2.

2.6 Verification Challenges

This section introduces terminology and discusses the challenges associated with system verification. Consider a *process*, as illustrated by the top part of Figure 2.10 consisting of a *vehicle*, a *driver*, and an *environment*. The environment has in general both static and dynamic content where static content is e.g. roads, trees and traffic signs and dynamic content is e.g. road users such as cars, bicycles and pedestrians. The *active safety system*, interact with the process according to Figure 2.10. The system monitor and control the process to ensure that the host vehicle is operated safely.

The purpose of *system verification* is to ensure that the system performance meets the system requirements. This must be addressed for the complete set of *operating scenarios*, defined by the variations in the process, i.e. the variations in vehicle, driver and environment behaviour. The set of operating scenarios is essentially unlimited in size as combinations of e.g. weather conditions, road user types, appearance and motion patterns are infinite, see Figure 2.11a. Defining the boundaries of this set is a challenge in itself since the system is mobile and travels in an environment which is completely or partially unknown to the system a priori.

The system relies on real-time remote sensing of e.g. road users and road geometry to make decisions on when and how to intervene. The sensing performance depends on variations in the environment, as illustrated by Figure 2.11b. For instance, a camera subjected to direct sunlight will exhibit poor object detection performance, much like the human eye.

The more intrusive an intervention type is, see Section 2.3.2, the less

(a) (b)

Figure 2.11: Illustration of system verification challenges as seen by a vision sensor. (a) One of the many possible complex traffic situations. (b) The sensor is partially blinded when exiting the tunnel.

likely the driver is to accept an unnecessary intervention. Consequently, the acceptable rate of unnecessary interventions is very low for systems performing intrusive interventions. The large quantity of operating scenarios makes verification of this requirement, on a low rate of unnecessary interventions, especially challenging.

Chapter 3

Verification Methods

The goal of *system performance evaluation*, in the context of this thesis, is to determine the performance of an active safety system in a given set of operating scenarios. If system performance evaluation is used for *system verification*, the performance estimate is compared to a set of *system requirements*, which specify the acceptable level of system performance. Accurate and efficient methods for system performance evaluation are needed for several purposes, e.g. system verification, system tuning or analysing the system sensitivity to disturbances. For verification purposes it is usually sufficient to derive or estimate a bound on performance, to show that the system requirements are fulfilled. As a consequence, some methods focus on performance bounds while some focus on performance estimates. This chapter provides an overview on verification methods used in an active safety context.

3.1 Performance Metrics

In this section, performance metrics describing the ability of the system to make correct decisions are presented. A commonly used terminology for describing the nature of incorrect decisions comes from *statistical hypothesis testing*, extensively covered in [38], and was first discussed in [39]. A hypothesis test is classified with regards to the test outcome, i.e. the decision on what hypothesis to accept, and the true hypothesis, see Figure 3.1. The default decision, often a decision *not* to perform an action, is in statistical hypothesis testing represented by the *null hypothesis*. A test outcome is said to be *negative* if the null hypothesis is accepted and *positive* in the opposite case.

In an active safety context, a test would be e.g. to decide whether or not to initiate an autonomous brake intervention, the true hypothesis would

	System does *not* intervene (Null hypothesis is *accepted)*	System intervenes (Null hypothesis is *rejected)*
Intervention is *not* needed (Null hypothesis is *true)*	True Negative	**Unnecessary Intervention** *Type I Error* False Positive
Intervention is needed (Null hypothesis is *false)*	**Missed Intervention** *Type II Error* False Negative	True Positive

Figure 3.1: Error types for a system deciding on whether or not to intervene.

represent the correct decision and the null hypothesis would represent the decision not to intervene.

Linked to this, there are two types of errors, commonly referred to as *Type I* and *Type II* errors. If the null hypothesis is true and is rejected by the test, the error is Type I or *false positive*. If instead the null hypothesis is false and is accepted by the test, the error is Type II or *false negative*. False positives and false negatives are in this thesis referred to as *unnecessary* and *missed* interventions respectively, since these terms are more descriptive for active safety applications.

3.2 Method Properties

Different methods have different properties and each property contributes to the overall strength or weakness of the method. Below, relevant method properties are defined and discussed.

Coverage

Coverage is a measure used to describe the degree to which the set of operating scenarios is evaluated. A major benefit of theoretical methods is that full coverage is possible to attain.

By conducting experiments, i.e. tests, system performance can be evaluated in a chosen set of scenarios. For a complex process, generally, the set of operating scenarios can be described by an unbounded number of parameters. As the number of operating scenarios grows exponentially with the number of scenario parameters, this set is very large. This effect, known as the *curse of dimensionality*, makes full coverage of the set of operating scenarios unrealistic.

Methods for selecting a set of scenarios to evaluate are generally referred to as *experimental design* or *Design of Experiments* (DoE), see e.g. [40] for a wide treatment of the subject. The scenario parameter space may for instance be covered by drawing random samples, or using a more systematic approach, the samples may be chosen such that the coverage is evenly spread while minimizing the number of evaluated scenarios.

Online/Offline

Systems are evaluated either online or offline, where these terms are used according to the following definitions.

Definition 1 *A system is* online *when forced to execute in real-time.*

Definition 2 *A system is* offline *when not forced to execute in real-time.*

Online experiments evaluate if the system comply with real-time requirements but have the obvious disadvantages of not being able to execute slower, or faster, than real-time.

Open/Closed-Loop

An active safety system monitors a process and use this information to influence said process, as shown in Figure 2.10. In some experiments the scenario is partially or completely fixed, meaning that the system has limited or no influence on the process. As a consequence, the following definitions are useful.

Definition 3 *A system is executed in* closed-loop *when the control loop between the system and the process is closed.*

Definition 4 *A system is executed in* open-loop *when the control loop between the system and the process is open.*

Note that open-loop execution does not equal open-loop control, which commonly refers to a control system operating without feedback. Open-loop execution means that the system cannot influence the process during execution. When evaluating the correctness of decisions, it is in many cases sufficient to execute the system in open-loop. This is valid when the system does not perform any action prior to the decision to intervene, and will consequently not influence the process prior to said decision.

Efficiency

The cost in terms of time and money are measures of the method *efficiency*. Online methods are time consuming as real-time execution is required and, in general, methods involving real world experiments have higher financial cost than theoretical analysis and computer experiments. Also, the cost associated with method development vary between different methods.

Repeatability and Reproducibility

Repeatability and reproducibility are statistical terms associated with accuracy. An experiment is *repeatable* if it can be performed on two different occasions with no substantial change between measured quantities. Repeatability only requires this to be possible using the same personnel and equipment. An experiment is also *reproducible* if it is repeatable using different personnel and equipment when performed on two different occasions. Repeatability and reproducibility ensures that the experiment results are not significantly affected by temporary factors.

Ground Truth Data

Ground truth data refers to information that is confirmed in an actual field check at a location, as opposed to information acquired from a distance. In remote sensing, the term is commonly used to describe information considered accurate, relative to information acquired from the remote sensing system being evaluated. Ground truth data describes the true scenario, e.g. how objects move in the scene, and aids in evaluating sensor and control system performance.

Model Accuracy

The *model accuracy* is the ability of the model to generate output equivalent to output from the real system. Model accuracy should not be confused with system accuracy which is the ability of the system to generate output with small errors, e.g. a sensor delivering accurate measurements. If the input to an accurate model is equivalent to the real operating conditions, the output is *realistic*. Output collected from real systems in the real operating environment is realistic per definition. Achieving high output realism from models often induces high cost, in the form of time, money or both.

Scenario Representativeness

For a scenario set to be *representative*, it must correctly reflect the set of operating scenarios in terms of system performance. Sampling scenarios using the real system in the real operating environment is the most obvious way to collect data with high representativeness. When modeling or recreating real scenarios, e.g. in computer simulations or real world test environments, limitations imposed by process models or test equipment make the scenarios less representative to a varying degree.

Process Controllability

The ability to control the process during evaluation is referred to as *process controllability*. Lack of process controllability is primarily an issue in real world experiments, where control of the process related to for instance weather or multiple object dynamics is challenging. Also, safety-critical situations such as collisions and near-collisions are difficult to realize in real world experiments where they are potentially hazardous for involved personnel and destructive to the equipment used.

3.3 Models

Many verification methods use mathematical models to describe the active safety system, the process and their interaction, see Figure 2.10. The complexity of the models vary significantly and also depends on the interfaces to other models, e.g. interfaces to perception or action layer models might require more or less complexity in the corresponding process models.

3.3.1 System Models

Active safety systems consist of three layers, see Figure 2.7. Commonly, the decision layer is an available software interacting only with the perception and action layers. The latter two layers interact physically with the process, e.g. the surrounding environment, and modeling of these layers are discussed below.

Perception Layer Models

As described in Section 2.3, the perception layer provides input data to the decision layer, based on sensor observations of a process. In the perception layer, observations from one or several sensors are generally passed through multiple layers of advanced signal processing, fusing sensor observations into estimated states such as positions and velocities of detected objects. Sensor models describe how the process is perceived by the sensors and can be formulated on many different abstraction levels.

Low-level sensor models describe the transformation between the process and the unprocessed sensor observations whereas *high-level* sensor models describes the transformation between the process and the estimated states. Modeling the physics of remote sensing technologies such as cameras, lidars and radars, is a complex task, especially when considering situations where the sensor is observing a complex environment, see e.g. [41] for an overview of low-level radar models. This is why high-level empirical models are often used.

A common high-level approach is to model state estimates, e.g. object position or velocity, as the true estimate influenced by a noise model. If noise is ignored, the models represent *ideal* or *perfect* sensors, as used in e.g. [42]. A common noise model is additive Gaussian noise, as used in e.g. [43–45] and Paper 3. High-level models are in many cases a major simplification of the sensor and incorporate very limited information on how the sensor errors depend on the observed process. Nevertheless, they are useful when studying systems in limited scenario sets, systems with very accurate sensors, or the aspects of system performance not affected by sensor errors.

For a computer vision system, a low-level sensor model describes how a camera perceives the process, i.e. generates a sequence of images. Techniques for generating images with computers, known as *rendering*, are studied in *computer graphics*. Rendering imagery requires process models which describe e.g. the 3D structure of objects in the environment. Rendered imagery based on 3D models, in contrast to real imagery collected from cameras, is denoted *virtual imagery* while rendered imagery where virtual objects are superimposed on real imagery is denoted *augmented imagery*.

In [46], published in 1995, it is argued that the realism of virtual imagery is sufficient for evaluation of mobile computer vision systems. Since then computer graphics has evolved rapidly, as can be observed in for instance the video gaming and movie industries. Nonetheless, photo-realism in virtual imagery is not easily achieved and an overview of the multidisciplinary challenges of rendering is found in [47]. Paper 4 explores the possibility of rendering *augmented imagery* for offline evaluation of computer vision systems.

Online evaluation methods require image rendering in real-time, making it more challenging to attain high realism in rendered images. The traffic simulation environments described in [48, 49] have software modules available for rendering of virtual imagery in real-time.

Action Layer Models

Models of e.g. braking and steering systems are needed to describe how the driver and vehicle are influenced by system decisions. Descriptions and models of automotive systems and components, including active safety actuators, are thoroughly described in [50]. Note that when the system is executed in open-loop, modeling the action layer is unnecessary.

3.3.2 Process Models

Modeling of the process, i.e. the driver, vehicle and surrounding traffic environment, is discussed in the following sections.

Driver Models

For evaluation methods based on real or augmented data, the behaviour of the driver is incorporated in the data. Therefore, only purely model-based methods require a driver model to generate the driver input to the vehicle, e.g. steering and braking, based on feedback from the vehicle, environment and active safety system. Driver modeling is a wide field of research and models are often more or less application specific. A collection of papers treating driver models in the automotive domain from a variety of perspectives is found in [51].

Vehicle Models

Vehicle motion models are needed to describe both the motion of the host vehicle as well as vehicles in the surrounding environment. Vehicle motion is studied within the field of *vehicle dynamics* which is the topic of several books, e.g. [52].

Environment Models

When modeling a dynamic traffic environment, each object in the environment, e.g. cars, roads and pedestrians, are described by individual models. Depending on the interface to the active safety system, e.g. sensor and actuator models, the environment models need to include different aspects. If low-level sensor models are used, the level of detail of the environment models is usually higher compared to when high-level sensor models are used. If for instance virtual imagery is generated by a sensor model, a complete 3D structure of the environment is required.

Presently, there exist several simulations environments for simulating traffic environments including active safety systems such as PreScan, [48], v-TRAFFIC, [49], or the Volvo Cars Traffic Simulator (VCTS), [27]. These softwares include models of driver, vehicle and environment.

3.4 Methods

This section describes different types of analysis and verification methods. The methods evaluate real physical systems, mathematical models, or a combination thereof.

3.4.1 Real Driving

Online experiments using real vehicles are performed both in real traffic and on dedicated test tracks. They are repeatable to some degree at test tracks but to a minor degree in real traffic. If the system is online, it can be evaluated in closed-loop and sensor data often have the advantage of being realistic.

Real Traffic

Real traffic experiments are primarily used to estimate the probability of an unnecessary intervention from a set of randomly sampled scenarios. Also, experiments are conducted to estimate the probability of a missed intervention, given that the tested system has relatively frequent and non-intrusive interventions, which is valid for e.g. an LDW system.

Variations between different vehicles and system components are handled by using multiple vehicles and components in testing. For a randomly sampled scenario set to be representative, the scenarios available for sampling must also be representative. In [27, 53, 54], a Real World User Profile (RWUP) is used to ensure that a representative scenario set is sampled, taking into account for instance different driving styles, weather and driving environments.

As discussed in Section 2.6, the acceptable rate of unnecessary interventions for highly intrusive systems is very low, meaning that a large amount of driving data needs to be collected to ensure that the requirement is fulfilled. The obvious drawback is that real traffic experiments are both expensive and time consuming. Also, ground truth data is challenging to obtain since the environment is uncontrolled.

Test Track

On test tracks, specific types of scenarios are tested in a more controlled setting. Compared to real traffic experiments, test track experiments of-

(a) (b) (c)

Figure 3.2: Non-destructive tests in collision and near-collision scenarios where (a) shows stationary pedestrian dummies of both adult and child size, (b) shows an inflatable moving object representing a moving vehicle and (c) shows an artificial object representing a moose.

fer a higher degree of process controllability, repeatability and reproducibility. Motions of involved objects can be controlled to create desired scenarios. Also, ground truth can be obtained by e.g. positioning involved traffic participants and objects with an accurate positioning system such as Differential Global Positioning System (DGPS). Process controllability on test tracks is better but not without limitations as for instance weather, e.g. snow or rain, and animals crossing the road are still difficult to reproduce on demand.

When recreating collision and near-collision situations on test tracks, non-destructive tests are preferred to ensure safety. Therefore, collisions are conducted between the host vehicle and low-mass objects such as inflatable cars, see Figure 3.2. This creates limitations on scenarios possible to recreate as even state-of-the-art inflatable car or pedestrian systems cannot recreate all motions possible for real cars or pedestrians. It also degrades the representativeness of the scenario since an inflatable object might not be perceived by the sensors as would an equivalent real object.

For highly intrusive systems such as AEB, the scenarios in which the system should intervene are very rare, meaning that estimating the probability of a missed intervention would require an unrealistic amount of driving data from real traffic conditions. Consequently, the probability of a missed intervention is often, e.g. [27, 53, 54], assessed by replicating collision situations on test tracks. Such tests are also used to estimate the effectiveness of the system, for instance impact speed reduction.

For verification purposes, scenarios in which an incorrect system decision is most likely, i.e. the *worst case* scenarios, are often replicated on test tracks, thus complementing tests in real traffic. If the system behaves correctly in these worst case scenarios it can be argued that less challenging scenarios are not likely to pose a problem. Paper 1 presents a theoretical method for identifying the worst case scenarios for a CA system. Examples of scenarios most likely to cause unnecessary interventions are near-collision situations, e.g. evasive maneuvers where the time or distance margins to a potential collision are small.

3.4.2 Closed-Loop Simulations

If the process is mathematically modeled, the system behaviour can be simulated in closed-loop with computer generated inputs. *Model-In-the-Loop* (MIL) simulations use a system model while *Software-In-the-Loop* (SIL) simulations use an actual system implementation, which not necessarily is executed on the production hardware. The border between MIL and SIL is sometimes hard to define but examples of one or the other is found in [27, 43, 44, 55].

MIL/SIL offers many benefits over real driving when comparing for instance efficiency and process controllability. Experiments are repeatable and reproducible and these are important properties when comparing different system configurations. MIL/SIL are offline methods, meaning there are no real-time constraints, making it possible to simulate scenarios with speeds limited only by the computational power available. In addition, systems can be tested before deployment at early stages in development, without the need of functioning hardware.

If system hardware components are available, their performance can be tested online with computer generated inputs as *Hardware-In-the-Loop* (HIL). The benefit of HIL, compared to MIL/SIL, is that the hardware is also evaluated. The drawback is the online property, constraining HIL simulations to real-time execution. In [56] a test facility where a complete vehicle is set up on a chassis dynamometer, with robot vehicles representing the surrounding environment, is described and referred to as *Vehicle Hardware-In-the-Loop* (VeHIL).

3.4.3 Data Replay

In *data replay* methods, real recorded data is used to evaluate the system offline. If real data is used exclusively, different system software configurations can be evaluated by computer simulations without any loss in input data realism, as done in [57], but often without the impeccable ground truth data accessible using in-the-loop methods. Limited ground truth, e.g. improved estimates of object motion, can be obtained by offline processing of real measurement data, as done in [43, 57]. Data replay methods are restricted to open-loop, since the scenario is fixed by the collected data.

Another option is to combine real data with model-based methods thus generating *mixed* or *augmented* data, for example by adding new objects or errors in recorded sensor data. This is exemplified in [43] where three FCW algorithms are simulated with input data consisting of accurate lead vehicle motion, obtained from real data, and noise, from a radar model. *Augmented data replay* has the potential to pick the best out of two worlds but also risk picking the worst. The modeling effort compared to purely model-based methods is limited and many of the advantages are partly preserved, e.g. process controllability, or completely preserved, e.g. repeatability and lack of real-time constraints. The downside is that simulation is limited to open-loop, as some real data is used, and that the data realism is now dependent on a model-based data augmentation method which then requires validation.

In Paper 4, an augmented data replay framework is formulated, used for computer vision systems. This framework uses a low-level sensor model, discussed in Section 3.3. If instead high-level sensor models are available, Paper 3 presents efficient data replay methods for decision function tuning and sensitivity analysis with regards to input perturbations, which can be applied to real, model-based or augmented data.

3.4.4 Theoretical Methods

The ultimate goal of system verification is to *prove* that the system meets the system requirements. Methods for proving system properties, such as requirement compliance, are known as *formal methods*, see [58] for an extensive survey.

If the active safety system and the set of operating scenarios are described mathematically it is sometimes possible to derive analytical expressions describing system performance, as done in Paper 1. Generally, this is only possible when making quite significant simplifications.

For dynamical systems, guarantees of not entering an undesired system state may be obtained by computing the set of reachable states. Paper 2 explores the use of reachability analysis, [59], and viability theory, [60], to formally verify a collision avoidance system.

3.5 Method Comparison

This section provides a brief comparison of the methods presented in Section 3.4 with regards to the properties discussed in Section 3.2. An overview of the more or less discrete properties is presented in Table 3.1. Figures 3.3 and 3.4 compare process controllability, sensor data realism and efficiency

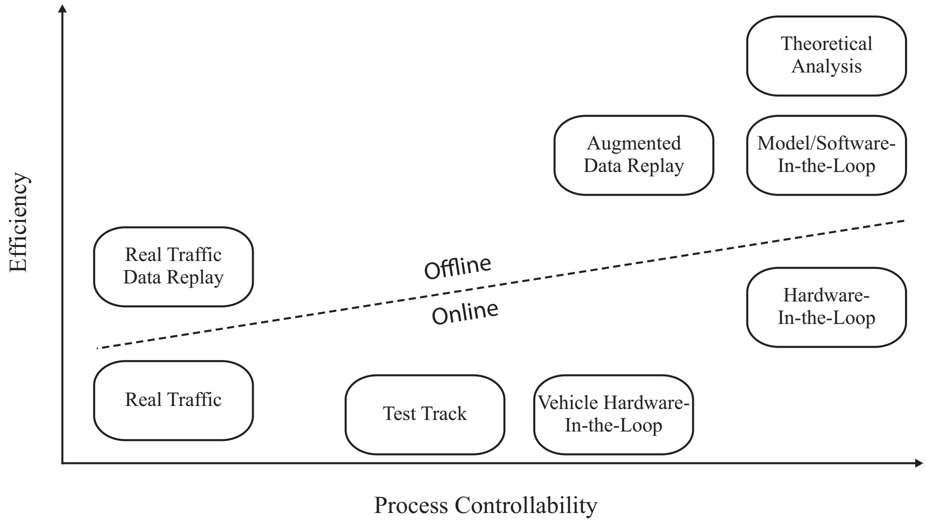

Figure 3.3: A qualitative sketch for the relation between efficiency and process controllability for different evaluation methods.

for different methods. It should be noted that these properties are application dependent, meaning that the figures should not be considered absolute truths.

In Figure 3.3 it can be noted how the model-based methods are superior

	Vehicle	System hardware	Online	Closed-loop	Ground truth	Repeatability	Reproducibility
Theoretical analysis				x	x	x	x
Real traffic	x	x	x	x			
Test track	x	x	x	x	x	(x)	(x)
Model/software-in-the-loop				x	x	x	x
Hardware-in-the-loop		x	x	x	x	x	x
Vehicle hardware-in-the-loop	x	x	x	x	x	x	x
Real traffic data replay	(x)	(x)			(x)	x	
Augmented data replay	(x)	(x)			x	x	

Table 3.1: Overview of method properties for different methods. The fact that data replay methods use vehicles and system hardware indirectly, i.e. for initial data collection, is represented by a tentative "(x)". The same notation indicates that, on test tracks, some aspects are repeatable and reproducible while others are not. Ground truth for real data replay is also tentatively marked as offline processing can offer limited ground truth data.

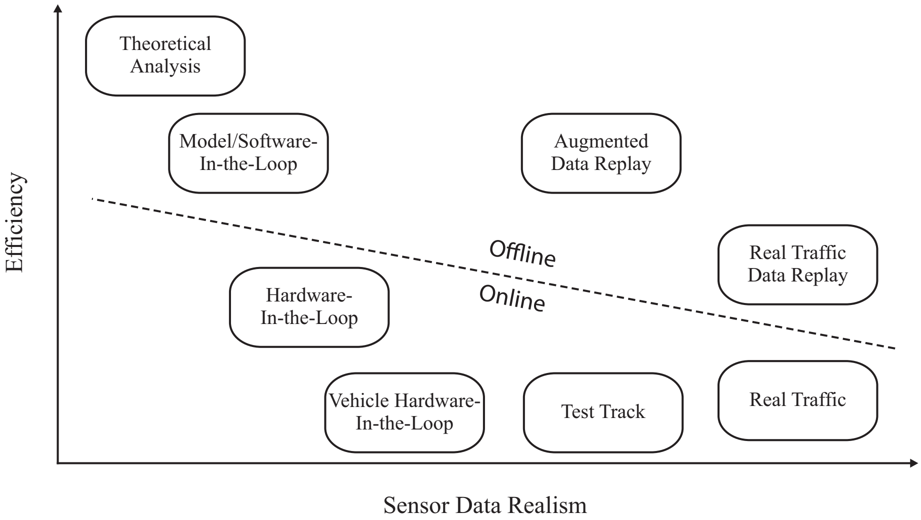

Figure 3.4: A qualitative sketch for the relation between efficiency and sensor data realism for different evaluation methods.

in terms of process controllability and also in many cases are very efficient, largely due to the offline property. The major challenge for the model-based methods is related to the sensor model accuracy, as visualized in Figure 3.4. For the purely model-based methods, e.g. closed-loop simulations, sensor models generating realistic data are either unexisting or resource demanding. Methods using exclusively real traffic data have, by definition, realistic sensor data. Augmented data is relatively realistic but with the drawback that augmented data replay is limited to open-loop execution of the system, as shown in Table 3.1.

The comparisons in Table 3.1, Figure 3.3 and Figure 3.4 clearly show the complementary nature of the presented methods. Thus, verification is often carried out using a variety of methods, as exemplified in [27, 53, 54]. Methods which require complete vehicles or system hardware are constrained to use in the later stages of the development process. Alternatively, they may be employed with decreased prediction accuracy using early hardware prototypes.

Chapter 4

Summary of Included Papers

This chapter provides a brief summary of the papers included in the thesis and also describes the contributions to each paper by the author of this thesis. Full versions of the papers are included in Part II.

Paper 1

> J. Nilsson, A. Ödblom and J. Fredriksson, Worst Case Analysis of Automotive Collision Avoidance Systems, submitted for possible journal publication.

As discussed in Section 2.6, the set of traffic scenarios which generates the input to an active safety decision function is very large. This paper theoretically identifies scenarios with a high risk of incorrect system decisions, i.e. the worst case scenarios. The main challenge with this approach, as discussed in Chapter 3, is to model system and scenarios in such a way that performance can be described analytically while still including the key factors affecting performance, e.g. sensor errors or object motion.

The key idea of this paper is to theoretically investigate the fundamental limitations of a collision avoidance system, subject to systematic measurement errors and unexpected future object motion, in terms of early and unnecessary interventions. Specifically, we include effects of sensor and actuator delays, and derive closed-form expressions for the worst case performance, with regards to longitudinal or lateral prediction and measurement errors. For a system example, numerical results show how decision timing and robustness depend on scenario and system parameters. The method can be used for system verification, tuning or sensitivity analysis with regards to scenario variations and sensor errors. Also, scenarios with inadequate performance can be identified, thus improving existing test methods by directing testing and analysis efforts towards relevant scenarios.

The thesis author was responsible for the problem formulation, derivation of the closed-form expressions, implementation and writing the paper.

Paper 2

J. Nilsson, J. Fredriksson and A. Ödblom, Verification of Collision Avoidance Systems using Reachability Analysis, submitted as invited paper to *the 19th IFAC World Congress*, Cape Town, South Africa, 2014.

The closed-form expressions for performance derived in Paper 1 are very useful from a verification perspective but for many complex active safety decision functions, they are not possible to derive. The alternative of evaluating state trajectories, as done in traditional simulations and real vehicle tests, does not provide guarantees for system performance for all possible state trajectories.

To address these limitations, Paper 2 describes a novel set-based framework for analyzing under what conditions the absence of incorrect decisions may be guaranteed for a given collision avoidance decision function. Reachability analysis and viability theory are used to compute unsafe and safe sets, i.e. sets where an ideal system should or should not intervene respectively. In these sets, incorrect decisions for a given decision function are identified using optimization techniques. By separating the dynamics of the input space from the decision function, non-linear and ad-hoc decision functions are efficiently handled in the proposed framework.

The method is demonstrated on a collision avoidance system example and, given the models used and absence of measurements errors, we show that the system does not make incorrect decisions. Furthermore, we describe and demonstrate how to evaluate the robustness to measurement errors, using the proposed framework.

The thesis author was responsible for the problem formulation, development of the proposed methods, implementation and writing the paper.

Paper 3

J. Nilsson and M. Ali, Sensitivity Analysis and Tuning for Active Safety Systems, in *Proceedings of the 13th International IEEE Conference on Intelligent Transportation Systems*, 2010, pages 161-167, Madeira Island, Portugal.

Papers 1 and 2 are full coverage methods, i.e. are concerned with verification of the complete scenario parameter space. Full coverage methods are desirable but set limitations on the complexity of the involved mathematical models. In contrast, Paper 3 considers verification given that a representative experimental data set is available.

The design and tuning of an active safety decision function, e.g. how thresholds are placed, will decide how sensitive the system performance is to input errors. Investigating the interplay between input errors, decision function and system performance gives rise to three relevant questions:

i. Given a decision function and input errors, what is the system performance?

ii. Given a decision function and system performance requirements, what are the input requirements?

iii. Given input errors and system performance requirements, how should the decision function be tuned?

This paper proposes a framework for open-loop analysis of decision functions, with regards to the above mentioned questions. By introducing a robustness measure, describing the robustness to input errors for the decision function, efficient offline methods are formulated. The robustness measure is independent of the input errors, meaning that it needs to be estimated only once for each decision function and data set. This allows for efficient evaluation of the system performance as combinations of decision function and input errors can be processed without evaluating the decision function output for each combination. The framework is applied to data collected in an experimental setting. Also, it is demonstrated how it can be used for setting input requirements and tuning the decision function.

The formulation of the presented framework and writing the paper were jointly conducted by both authors of the paper. The author of this thesis is responsible for the demonstration of the framework while the second author is responsible for the collection of experimental data and development of the decision function example.

Paper 4

J. Nilsson, A. Ödblom, J. Fredriksson, A. Zafar and F. Ahmed, Performance Evaluation Method for Mobile Computer Vision Systems using Augmented Reality, in *Proceedings of the IEEE Virtual Reality Conference*, 2010, pages 19-22, Waltham, Massachusetts, USA.

The methods for analyzing decision functions in Papers 1-3, all rely on accurate modeling of sensor errors. In Paper 4, a novel framework using augmented imagery is proposed for determining sensor errors of computer vision systems, which are widely used in active safety systems. The proposed framework exploits the possibility to add virtual agents into a real data sequence collected in an unknown environment, thus making it possible to efficiently create augmented data sequences, including ground truth, to be used for performance evaluation. Varying the content in the data sequence by adding different virtual agents is straightforward, making the proposed framework very flexible.

The method has been implemented and tested on a pedestrian detection system used for collision avoidance. Preliminary results show that the method has the potential to replace and complement physical testing, for instance by creating collision scenarios, which are difficult to test in reality.

The formulation of the novel framework was jointly conducted by the first two authors of the paper. The author of this thesis was also responsible for writing the paper and supervising the case study implementation done by authors four and five.

Paper 5

> J. Nilsson, J. Fredriksson and A. Ödblom, Reliable Vehicle Pose Estimation using Vision and Single-Track Model, submitted for possible journal publication.

The method in Paper 4 relies on an accurate 3D reconstruction of the camera motion in six Degrees of Freedom (6-DoF). Extensive use of this method requires this to be done without adding additional expensive sensors to the vehicle. The core idea of Paper 5 is to use a single-track vehicle model in a local bundle adjustment framework to improve the pose estimates obtained from a standard vehicle sensor setup, i.e. a forward looking monocular camera, wheel speed, yaw rate and steering wheel angle sensors. This means pose estimates are optimized not only with regards to observed image features, but also with respect to a single-track vehicle model and standard in-vehicle sensors.

The described method has been tested experimentally on challenging data sets at both low and high vehicle speeds as well as on a data set with moving objects. The vehicle motion model in combination with in-vehicle sensors exhibit good accuracy in estimating planar vehicle motion. Results show that this property is preserved when combining these information sources with vision. Furthermore, the accuracy obtained from vision-only in direction estimation is improved, primarily in situations where the matched visual features are few.

The thesis author was responsible for the problem formulation, development of algorithms, implementation, experimental validation and writing the paper.

Paper 6

J. Nilsson, P. Andersson, I. Gu and J. Fredriksson, Augmented Training Data for Pedestrian Detection, submitted to *the 22nd International Conference on Pattern Recognition*, Stockholm, Sweden, 2014.

Machine learning techniques are widely used in computer vision to train object classifiers. In many applications, e.g. pedestrian detection, the dominating approach in literature is to use supervised learning, e.g. Support Vector Machines (SVM), to train a classifier using labelled data. This labelled data is chosen such that it represents the environment where the classifier will be used. Thus, for a mobile system operating in a complex and uncontrolled environment, e.g. a car, the training data set must contain a great amount of variation. Collecting and manually labelling large amounts of data is an expensive and time consuming process.

In Paper 6, we propose to replace or complement real data with augmented data, using the method presented in Paper 4. Augmented data can be automatically labelled while still exhibiting a real, and consequently realistic, background. The proposed solution is evaluated by training pedestrian classifiers using one of the gold-standard methods in pedestrian classification, specifically a linear SVM and the Histogram of Oriented Gradients (HOG), [61]. Experimental validation is performed on real data sets and the results are compared to performance obtained using real training data.

The thesis author was responsible for the problem formulation and writing the paper. The design of experiments was conducted jointly by the author of this thesis and the second author of the paper. Note that the development and implementation of algorithms were primarily the responsibility of the second author, and not the author of this thesis.

Chapter 5

Concluding Remarks

This chapter states the most important contributions and provides recommendations for future research.

5.1 Contributions

System verification of an automotive safety system must assess the correctness of system decisions in a vast array of traffic scenarios. These decisions are based on remote sensing of the surrounding environment and consequently, including sensors in the analysis and verification methods is crucial. Computational methods have the potential to significantly improve the verification process in terms of e.g. efficiency and coverage. This thesis focus on computational methods for both decision function analysis, including the dependence on sensor errors, and methods for determining these sensor errors.

Related to decision function analysis and verification, the main contributions of this thesis are:

- Derivation of closed-form expressions for the worst case decision timing, in the presence of prediction and measurement errors, for a collision avoidance system example. Also, closed-form expressions are derived for robust avoidance scenarios, i.e. scenarios which are guaranteed not to exhibit an unnecessary intervention. These results are presented in Paper 1.

- A novel set-based framework for analyzing under what conditions the absence of incorrect decisions may be guaranteed for a given active safety decision function. In contrast to evaluating state trajectories, reachability analysis and viability theory are used to compute unsafe and safe sets, in which absence of incorrect decisions and robustness to sensor errors may be guaranteed using optimization techniques. This framework is presented in Paper 2 and forms a generalization of the work shown in Paper 1.

- A framework for active safety decision function analysis using recorded or simulated data. Efficient methods for system performance evaluation are derived and these can be used to analyze the decision function sensitivity to input errors, or for decision function tuning. This framework is presented in Paper 3.

Related to performance evaluation of computer vision systems, the main contributions of this thesis are:

- A novel performance evaluation approach using augmented imagery for evaluation of mobile computer vision systems. Performance is evaluated in collision and near-collision scenarios, safely and non-destructively, while still using a real image background from recorded data. This concept is presented in Paper 4 and the use of augmented data is extended from performance evaluation to training of a pedestrian classifier in Paper 6.

- An approach for 6-DoF vehicle pose estimation using a single vehicle-based standard camera. Visual features are complemented by standard in-vehicle sensors and a single track vehicle model in a bundle adjustment framework. The method has been validated experimentally in challenging situations at both low and high vehicle speeds. This method is presented in Paper 5 and is an important module needed for the framework introduced in Paper 4.

5.2 Directions of Future Research

There is a great need for more efficient verification methods to handle the challenges associated with future automotive safety systems. The work presented in this thesis has inspired multiple ideas on this topic.

Sensor error models

To make full use of the theoretical methods for performance estimation, presented in Papers 1-3, accurate sensor error models are needed. This requires acquiring and processing large amounts of sensor data, with associated ground truth, but also proper choices of model structures. The presented framework for sensor evaluation using augmented data may prove to be a valuable resource.

Extending reachability methods

The dynamical models used in Paper 2 are linear and low-dimensional, handling only a single moving object. Applying existing methods for reachabality analysis of more complex systems is an interesting approach. This could enable the analysis of the same problem with more complex vehicle dynamics models and/or multiple objects.

Augmenting other sensors

The augmentation framework in Paper 4 has been applied primarily on image data. Many safety systems fuse information from different sensor technologies, e.g. radar, laser. Thus, a natural extension would be to extend the concept to include also other sensor types. This requires in-depth knowledge of the sensor technology to be added and also accurate and detailed modeling of the specific sensor used.

References

[1] "Global status report on road safety 2013: supporting a decade of action," World Health Organization (WHO), Geneva, Switzerland, 2013.

[2] G. Jacobs, A. Aeron-Thomas, and A. Astrop, "Estimating global road fatalities," Transport Research Laboratory, Global Road Safety Partnership (GRSP), TRL445, 2000.

[3] "IRTAD Annual Report 2009," International Traffic Safety Data & Analysis Group (IRTAD), International Transport Forum, 2010.

[4] (2013, Dec.) OECD Statistics. OECD. [Online]. Available: http://stats.oecd.org/

[5] F. Winston and R. Menon, "Traffic Mortality in Germany Before, During, and After Reunification," *Annual Proceedings / Association for the Advancement of Automotive Medicine*, pp. 239–250, 1999.

[6] "Towards Zero: Ambitious Road Safety Targets and the Safe System Approach," OECD, International Transport Forum, OECD Publishing, 2008.

[7] M. Ali, "Decision Making and Control for Automotive Safety," PhD Thesis, No 3413, ISSN 0346-718X, Chalmers University of Technology, Göteborg, Sweden, 2012.

[8] W. Najm, B. Sen, J. Smith, and B. Campbell, "Analysis of Light Vehicle Crashes and Pre-Crash Scenarios Based on the 2000 General Estimates System," U.S. Department of Transportation, National Highway Traffic Safety Administration, John A. Volpe National Transportation Systems Center Cambridge, DOT HS 809 573, 2003.

[9] "Traffic Safety Facts 2011 - A Compilation of Motor Vehicle Crash Data from the Fatality Analysis Reporting System and the General Estimates System," National Highway Traffic Safety Administration, National Center for Statistics and Analysis, U.S. Department of Transportation, DOT HS 811 754, Washington, DC, USA, 2012.

[10] "Vägtrafikskador 2012 (Road traffic injuries 2012)," Trafikanalys, Statistik 2013:9, Stockholm, Sweden, 2013.

[11] J. Treat, N. Tumbas, S. McDonald, D. Shinar, R. Hume, R. Mayer, R. Stansifer, and N. Castellan, "Tri-level Study of the Causes of Traffic Accidents: Executive Summary," U.S. Department of Transportation, National Highway Traffic Safety Administration (Contract No. DOT HS 034-3-535), DOT HS 805 099, Washington, DC, 1979.

[12] V. L. Neale, T. A. Dingus, S. G. Klauer, J. Sudweeks, and M. Goodman, "An overview of the 100-car naturalistic study and findings," in *Proceedings of the 19th International Technical Conference on the Enhanced Safety of Vehicles*, 2005, pp. paper 05–0400.

[13] T. A. Dingus, S. Klauer, V. L. Neale, A. Petersen, S. E. Lee, J. Sudweeks, M. A. Perez, J. Hankey, D. Ramsey, S. Gupta, C. Bucher, Z. R. Doerzaph, J. Jermeland, and R. Knipling, "The 100-Car naturalistic driving study: Phase II - Results of the 100-Car field experiment," U.S. Department of Transportation, National Highway Traffic Safety Administration, DOT HS 810 593, Washington, DC, USA, 2006.

[14] L. Palkovics, A. Semsey, and E. Gerum, "Roll-Over Prevention System for Commercial Vehicles - Additional Sensorless Function of the Electronic Brake System," *Vehicle System Dynamics*, vol. 32, no. 4-5, pp. 285–297, Nov. 1999.

[15] J. Lu, D. Messih, and A. Salib, "Roll Rate Based Stability Control-The Roll Stability Control System," in *Proceedings of the 20th Enhanced Safety of Vehicles Conference (ESV)*, paper 07-0136, Lyon, France, 2007.

[16] R. Rajamani, *Vehicle Dynamics and Control.* Springer, 2006.

[17] M. Bayly, B. Fildes, M. Regan, and K. Young, "Review of crash effectiveness of Intelligent Transport Systems," TRACE Project, No. 027763, Deliverable D4.1.1 - D6.2, 2007.

[18] R. Bishop, *Intelligent Vehicle Technology and Trends.* Norwood, MA, USA: Artech House, 2005.

[19] R. K. Jurgen, *Adaptive Cruise Control.* SAE International, 2006.

[20] D. LeBlanc, G. Johnson, P. Venhovens, G. Gerber, R. DeSonia, R. Ervin, A. Ulsoy, and T. Pilutti, "CAPC: A Road-Departure Prevention System," *IEEE Control Systems Magazine*, vol. 16, no. 6, pp. 43–60, 1996.

[21] M. Ruder, W. Enkelmann, and R. Garnitz, "Highway Lane Change Assistant," in *IEEE Intelligent Vehicle Symposium.* IEEE, 2002, pp. 240–244.

[22] A. Eidehall, J. Pohl, F. Gustafsson, and J. Ekmark, "Toward Autonomous Collision Avoidance by Steering," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 1, pp. 84–94, Mar. 2007.

[23] W. Birk, M. Brännström, and D. Levin, "Method for determining a measure for evaluating the behaviour of a driver of a vehicle," Patent, EP1674375, European Patent Office, 2006.

[24] S. S. Blackman and R. Popoli, *Design and analysis of modern tracking systems.* Norwood, MA: Artech House, 1999.

[25] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman filter : particle filters for tracking applications.* Boston, Mass.: Artech House, 2004.

[26] M. Darms and H. Winner, "A Modular System Architecture for Sensor Data Processing of ADAS Applications," in *IEEE Intelligent Vehicles Symposium.* IEEE, 2005, pp. 729–734.

[27] E. Coelingh, H. Lind, W. Birk, and D. Wetterberg, "Collision Warning with Auto Brake," in *FISITA World Congress*, Yokohama, Japan, 2006.

[28] A. Polychronopoulos, M. Tsogas, A. Amditis, and L. Andreone, "Sensor Fusion for Predicting Vehicles' Path for Collision Avoidance Systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 3, pp. 549–562, Sep. 2007.

[29] F. Bengtsson and L. Danielsson, "A design architecture for sensor data fusion systems with application to automotive safety," in *Intelligent Transportation Systems World Congress*, New York, USA, 2008.

[30] (2014, Jan.) DARPA Urban Challenge. Defense Advanced Research Projects Agency (DARPA). [Online]. Available: http://archive.darpa.mil/grandchallenge/index.asp

[31] "Traffic Safety Facts 2008 - A Compilation of Motor Vehicle Crash Data from the Fatality Analysis Reporting System and the General Estimates System," National Highway Traffic Safety Administration, National Center for Statistics and Analysis, U.S. Department of Transportation, DOT HS 811 170, Washington, DC, USA, 2009.

[32] M. Richter, H.-C. Pape, D. Otte, and C. Krettek, "Improvements in passive car safety led to decreased injury severity - a comparison between the 1970s and 1990s." *Injury*, vol. 36, no. 4, pp. 484–8, Apr. 2005.

[33] (2014, Jan.) The European New Car Assessment Programme (EuroNCAP). [Online]. Available: http://www.euroncap.com/

[34] D. Burton, A. Delaney, S. Newstead, D. Logan, and B. Fildes, "Effectiveness of ABS and vehicle stability control systems," Report No. 04/01, Royal Automobile Club Of Victoria (RACV), 2004.

[35] S. A. Ferguson, "The effectiveness of electronic stability control in reducing real-world crashes: a literature review." *Traffic injury prevention*, vol. 8, no. 4, pp. 329–38, Dec. 2007.

[36] M. Lindman, A. Ödblom, E. Bergvall, A. Eidehall, B. Svanberg, and T. Lukaszewicz, "Benefit estimation model for pedestrian auto brake functionality," in *Proceedings of the ESAR Conference*, no. 300961, 2010.

[37] K. Kusano and H. Gabler, "Safety benefits of forward collision warning, brake assist, and autonomous braking systems in rear-end collisions," *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 4, pp. 1546–1555, Dec. 2012.

[38] E. L. Lehmann, *Testing statistical hypotheses*, 2nd ed. Pacific Grove, Calif.: Wadsworth & Brooks/Cole Advanced Books & Software, 1991.

[39] J. Neyman and E. S. Pearson, "On the Use and Interpretation of Certain Test Criteria for Purposes of Statistical Inference: Part I," *Biometrika*, vol. 20A, no. 1/2, pp. 175–240, 1928.

[40] R. O. Kuehl, *Design of Experiments: Statistical Principles of Research Design and Analysis*, 2nd ed. Pacific Grove : Duxbury-Thomson Learning, 2000.

[41] L. Danielsson, "Tracking and radar sensor modelling for automotive safety systems," PhD Thesis, No 3064, ISSN 0346-718X, Chalmers University of Technology, Göteborg, Sweden, 2010.

[42] J. Hillenbrand, A. M. Spieker, and K. Kroschel, "A Multilevel Collision Mitigation Approach - Its Situation Assessment, Decision Making, and Performance Tradeoffs," *IEEE Transactions on Intelligent Transportation Systems*, vol. 7, no. 4, pp. 528–540, Dec. 2006.

[43] P. Zheng and M. McDonald, "The effect of sensor errors on the performance of collision warning systems," in *IEEE Intelligent Transportation Systems Conference*, vol. 1. Piscataway, NJ, USA: IEEE, 2003, pp. 469–474.

[44] L. Yang, J. Yang, E. Feron, and V. Kulkarni, "Development of a performance-based approach for a rear-end collision warning and avoidance system for automobiles," in *IEEE Intelligent Vehicles Symposium*, 2003, pp. 316–321.

[45] O. Gietelink, B. De Schutter, and M. Verhaegen, "Probabilistic validation of advanced driver assistance systems," in *Proceedings of the 16th IFAC World Congress*, vol. 19, Prague, Czech Republic, 2005.

[46] W. Burger, M. Barth, and W. Sturzlinger, "Immersive Simulation for Computer Vision," in *joint 19th AGM and 1st SDRV workshop Visual Modules*. Maribor, Slovenia: Oldenbourg Press, 1995, pp. 160–168.

[47] A. S. Glassner, *Principles of Digital Image Synthesis*, 1st ed. Morgan Kaufmann, 1995.

[48] (2014, Jan.) PreScan. TASS International. [Online]. Available: https://www.tassinternational.com/prescan

[49] (2014, Jan.) v-TRAFFIC. VIRES Simulationstechnologie GmbH. [Online]. Available: http://www.vires.com/Products_ToolChain.htm

[50] *Bosch Automotive Handbook*, 7th ed. Robert Bosch GmbH, 2007.

[51] P. C. Cacciabue, Ed., *Modelling Driver Behaviour in Automotive Environments*. Springer, 2007.

[52] H. Pacejka, *Tyre and Vehicle Dynamics*, 2nd ed. Elsevier Ltd, 2006.

[53] R. Kiefer, D. Leblance, M. Palmer, J. Salinger, R. Deering, and M. Shulman, "Development and Validation of Functional Definitions and Evaluation Procedures for Collision Warning/Avoidance System," National Highway Traffic Safety Administration, DOT HS 808 964, Washington, DC, USA, 1999.

[54] M. Distner, M. Bengtsson, T. Broberg, and L. Jakobsson, "City Safety - A system addessing rear-end collisions at low speeds," in *Proceedings of the 21st International Technical Conference on the Enhanced Safety of Vehicles (ESV)*, Stuttgart, Germany, 2009.

[55] J. Hillenbrand and K. Kroschel, "A Study on the Performance of Uncooperative Collision Mitigation Systems at Intersection-like Traffic Situations," in *IEEE Conference on Cybernetics and Intelligent Systems*. IEEE, Jun. 2006, pp. 1–6.

[56] O. Gietelink, J. Ploeg, B. De Schutter, and M. Verhaegen, "Development of advanced driver assistance systems with vehicle hardware-in-the-loop simulations," *Vehicle System Dynamics*, vol. 44, no. 7, pp. 569–590, Jul. 2006.

[57] K. Lee and H. Peng, "Evaluation of automotive forward collision warning and collision avoidance algorithms," *Vehicle System Dynamics*, vol. 43, no. 10, pp. 735–751, Oct. 2005.

[58] J. Woodcock, P. G. Larsen, J. Bicarregui, and J. Fitzgerald, "Formal methods: Practice and Experience," *ACM Computing Surveys*, vol. 41, no. 4, pp. 1–36, Oct. 2009.

[59] I. Mitchell, "Comparing forward and backward reachability as tools for safety analysis," in *10th International Workshop, Hybrid systems: computation and control*. Pisa, Italy: Springer Berlin Heidelberg, 2007, pp. 428–443.

[60] J.-P. Aubin, A. M. Bayen, and P. Saint-Pierre, *Viability Theory - New Directions*, 2nd ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011.

[61] N. Dalal, "Finding people in images and videos," PhD Thesis, Institut National Polytechnique de Grenoble / INRIA Rhône-Alpes, Grenoble, 2006.