THESIS FOR THE DEGREE OF LICENTIATE OF ENGINEERING

MACHINE AND VEHICLE SYSTEMS GRADUATE SCHOOL

Modelling driver behaviour in run-off-road crashes:
Applications in safety system development and safety benefit estimation

DANIEL NILSSON

Department of Mechanics and Maritime Sciences

CHALMERS UNIVERSITY OF TECHNOLOGY

Gothenburg, Sweden 2017

Modelling driver behaviour in run-off-road crashes: Applications in safety system development and safety benefit estimation

DANIEL NILSSON

THESIS FOR LICENTIATE OF ENGINEERING no 2017:12
Department of Mechanics and Maritime Sciences
Vehicle Safety Division
Chalmers University of Technology
SE-412 96 Gothenburg, Sweden
Telephone + 46 (0)31-772 1000

Chalmers Reproservice
Gothenburg, Sweden 2017

And there is no map, and a compass wouldn't help at all

— Björk, *Human behaviour* (1993)

Modelling driver behaviour in run-off-road crashes:
Applications in safety system development and safety benefit estimation
DANIEL NILSSON
Department of Mechanics and Maritime Sciences
Vehicle Safety Division
Chalmers University of Technology

Abstract

Run-off-road crashes have been identified as a major concern for automobile safety and several advanced driver assistance systems (ADASs) targeting run-off-road crashes are on the market today. Assessment of ADASs require relevant test scenarios and valid computational models of driver behaviour. Therefore, the objectives of this thesis has been: (A) define run-off-road test scenarios, and (B) identify a conceptual framework suitable for modelling relevant behavioural mechanisms for crash causation.

Cluster analysis was applied to run-off-road crashes from representative in-depth crash data from the German GIDAS database. Nine different clusters were identified, forming a basis for test scenarios. The two largest clusters included crashes relevant for current lane support ADASs (i.e. *drift during daytime/night-time*), while other clusters suggested that drivers may need support in judging the physical limits of the vehicle (e.g. on *snowy rural roads*). However, a need for more detailed driver behaviour data was identified. Indeed, naturalistic data, which include more information about driver behaviour in critical situations, may help the definition of test scenarios by linking them to the behavioural mechanisms contributing to the crash causation.

This thesis also shows that modelling of driver behaviour may be supported by a framework based on new findings in contemporary neurocognitive science and, specifically, on predictive processing. This new framework improved the interpretation of the clusters and facilitated the formulation of plausible behavioural mechanisms leading to run-off-road crashes.

Keywords: run-off-road crashes; advanced driver assistance systems; driver behaviour; predictive processing; effectiveness assessment; cluster analysis

Acknowledgements

First and foremost, I would like to express my gratitude towards my supervisors (both current and former) for their valuable support and contributions: to Marco, for structure and direction; to Trent, for visions and perspectives; to Johan, for knowledge and the value of theory; and to all of them, for persistence and forbearance. I am also grateful for the presence of Gustav and Ola, upon whose shoulders I stand.

I would also like to thank my wonderful colleagues at Volvo Cars: all of 91440, for valuable feedback and Friday afternoons; Magda, for constantly reminding me about the excitement in our work; Malin, for your support and structure; and the rest of you, for whenever you were there when needed.

Neither last, nor least, I am forever thankful for my fellow colleagues at SAFER; I cannot mention you all, but I will mention Christian-Nils and Alberto, for sharing the experience. And Morgan.

And to a few, for spiritual support. To my family, for being there. Whenever. Wherever. To my parents; forgive when I forget. To my wonderful sister and my brother, the eternal sunshine. To those who catch me when I fall.

Och till dig, Moa. För att du fanns där, även när du inte fanns här.

The research work was carried out jointly at the Accident Prevention group at the Division of Vehicle Safety, Department of Mechanics and Maritime Sciences, Chalmers University of Technology, and at the Active Safety Analysis and Verification group, Volvo Cars Safety Centre, Volvo Car Group. The work was conducted at SAFER within the QUADRÆ project, sponsored by FFI/Vinnova (grant number 2015-04863).

Table of Contents

1 Background

Safe mobility—to move unharmed between the current position and a desired destination—is at the very core of human activities. Yet, 6 million crashes involving light vehicles were reported by the police during 2003—in the US alone. These crashes consisted mainly of rear-end crashes (29%), run-off-road crashes (22%), and lane-change crashes (9%) (Najm & Smith, 2007). While exact definitions of crash types vary within the literature, it is clear that run-off-road crashes—where the vehicle inadvertently departs from the roadway and collides with an off-road object, rolls over or otherwise suffers a breakdown—pose a serious societal challenge. In a study of four US crash databases, it was found that single-vehicle crashes with a road departure involved 10 percent of the occupants, but accounted for 31 percent of the fatalities. Together with single vehicle crashes with control loss (with an unknown fraction leading to a road departure), these two crash types accounted for 42 percent of fatalities (Kusano & Gabler, 2014). In a study of insurance claims for crashes in Germany, 29 percent of all insurance claims for crashes concerned crashes resulting from a lane departure. Out of these, 22 percent (6.3% of total insurance claims) were due to a road departure (Kuehn, Hummel, & Bende, 2009). Also in Sweden, crashes resulting from lane departures are of high importance, accounting for 49 out of 154 car occupant fatalities in 2010 (Strandroth, 2015). Regardless of the exact definition, effective countermeasures addressing run-off-road crashes should be devised, evaluated, and deployed.

In the context of automobile development, countermeasures that try to avoid collisions or reduce severity of impacts by actuating vehicle systems preemptively are typically referred to as *active safety* (Eskandarian, 2012, p. 9), or synonymously as *advanced driver assistance systems* (ADASs) (Winner, 2016). Over the last decades, there has been a clear increase both in prevalence and in capabilities of active safety technologies. In particular, there has been an increase in ADASs that uses a multitude of sensors in order to detect potential threats and urgent conflict situations. This is done using e.g. camera-based computer vision, LIDAR, and RADAR systems, that enables environmental mapping and vehicle positioning (Lundquist, Schön, & Gustafsson, 2012), or sensors measuring current vehicular states such as steering wheel angle or wheel slip (Mörbe, 2016). An ever-growing subset of ADASs are designed to address run-off-road crashes by supporting the driver in the monitoring and control aspects of lane keeping and to initiate warnings or interventions in order to resolve the situation before it evolves into a crash. These systems can be classified as lane support systems (LSS). LSS currently include lane departure warning (LDW) systems, and lane keeping assist (LKA) systems (Euro NCAP, 2015). In addition, emerging countermeasures include always-on emergency LKA (ELK) systems that target drift events towards road edge or into oncoming or overtaking traffic (Euro NCAP, 2017). LDW systems support the driver's monitoring task by issuing a warning when the vehicle is about to depart from the lane, relying on the driver to initiate corrective steering manoeuvres in order to avoid a road departure. LKA systems, on the other hand, support the driver in the control part of the lane keeping task by providing lateral control interventions (automatic heading corrections) when an imminent lane departure is detected (Bartels, Rohlfs, Hamel, Saust, & Klauske, 2016). Warning strategies for LDW may include tactile (e.g. vibrations in steering wheel or seat), auditory (e.g. beep sound or rumble strip imitation), and/or visual (e.g. heads-up display lights) modalities (Gayko, 2012). The most commonly used actuation strategy for the LKA functionality has been a system-generated steering wheel torque, either in a loose fashion engaging close to lane markings or road edge, or in a lane-centring manner (Gayko, 2012), but other actuation strategies for steering assist are possible, for example using single-sided braking or torque vectoring (Bartels et al., 2016, p. 1224; Dang et al., 2012). Recently, more advanced features have also emerged on the market that, when active, continuously detects the forward roadway and automatically tries to position the vehicle within the lane by proactive steering interventions[1]. It is clear is that there are many different run-off-road countermeasure strategies that are based on supporting the driver in the driving task. For ongoing and future work to improve traffic safety, it is important to understand the effectiveness of these systems, and what crash-relevant situations current and future systems are able to address, as well as the situations that remain to be addressed by future systems.

1.1 Effectiveness assessment during development and evaluation of run-off-road crash countermeasures using virtual simulation

Naturally, the degree to which ADASs are able to prevent road departures is a topic of high interest, both to vehicle manufacturers (wanting to improve their systems), to rating organisations (wanting to recommend specific

[1] For example Volvo's Pilot Assist, Mercedes-Benz's Active Steering Assist, and Tesla's Autopilot.

cars/systems to customers concerned with safety), and to policy makers (wanting to know if a system should be made a legal requirement in new cars). The answer to this question can be estimated in early stages of development or deployment, without having to resort to retrospective analysis, by the use of prospective effectiveness methods. Methods for prospective effectiveness assessment aim to evaluate future and current systems, prior to market introduction, and are typically quantified in terms of the estimated safety benefit. The safety benefit is generally defined as the reduction in fatalities or injuries—for a given market penetration of an ADAS—as compared to a baseline where no system is present (Alvarez et al., 2017).

In the P.E.A.R.S. (Prospective Effectiveness Assessment for Road Safety) initiative, data on methodological approaches to prospective safety benefit estimation have been gathered from over 30 (mainly European) organisations—ranging from car manufacturers and insurance companies to research institutes and universities—in an effort to harmonise prospective effectiveness assessments among different stakeholders (Alvarez et al., 2017; Page et al., 2015). *Virtual simulation* was chosen as the basis for a harmonised general approach to effectiveness assessment and safety benefit estimation within the initiative (Page et al., 2015, p. 7). Virtual simulation has emerged as a promising method for safety benefit estimation, capable of evaluating system effectiveness in a large number of cases, while offering flexibility, reproducibility, and experimental control (Alvarez et al., 2017, p. 3). It constitutes a model-based approach that typically takes the form of *counterfactual simulations*—sometimes referred to as "what if"-simulations—that try to answer the question: could a (specific) crash have been avoided or mitigated, had the vehicle been equipped with a specific ADAS? (Alvarez et al., 2017; Bärgman, Boda, & Dozza, 2017; Bärgman, Lisovskaja, Victor, Flannagan, & Dozza, 2015; McLaughlin, Hankey, & Dingus, 2008; Scanlon, Kusano, Sherony, & Gabler, 2015). By simulating the same situation both with and without an ADAS, the results can be compared and the impact of the system can be assessed. Safety benefit estimation thus consists of simulating a target crash population—the *baseline situations*, consisting of either specific real-world crashes or synthesised crashes—in a counterfactual manner to assess the effectiveness of the ADAS in that crash population (Alvarez et al., 2017). However, the validity of a safety benefit estimation based on virtual simulation depends on the extent to which the input data is representative for the real-world crash scenarios, and the validity of the models used in the simulation process. In particular, Bärgman et al. (2017) showed how safety benefits for ADAS strongly depend on the driver model used in the counterfactual simulation and the potential of naturalistic data to support such simulations both in terms of test scenarios and to develop accurate driver models.

1.2 Virtual simulation for safety benefit estimation requires relevant test scenarios

While virtual simulation for ADAS evaluation can be done for multiple purposes, simulations are most commonly done with the objective to estimate the safety benefit of the system, based on comparisons between results from simulation conditions with and without the system (Alvarez et al., 2017; Markkula, 2015). However, as indicated above, a mapping is required between the simulated baseline situations and the targeted real-world crash situations in order to assure validity between the simulation results and the potential real-world safety benefit. This issue is typically addressed by basing the baseline situations on test scenarios that are representative of, or has a known relation to, a relevant real-world crash population of interest (Alvarez et al., 2017; Page et al., 2015). In addition, computational and analytical efforts can be saved by defining the test scenarios in a way that ensures that the system has some sort of impact in the simulated situations (Page et al., 2015, p. 4). Within the P.E.A.R.S. initiative, three[2] general strategies to obtain baseline situations (or *reference* situations) for simulation were identified, hierarchically categorised by the involvement of simulation in the scenario generation process (Alvarez et al., 2017):

1. At the first level, baseline situations are generated by directly selecting reconstructed crashes from crash databases.
2. At the second level, baseline situations are generated through simulation, for example by using marginal distributions from real-world crashes to generate synthetic scenarios through Monte-Carlo simulations.
3. At the third level, a large number of situations—based on processes and contributing factors identified in real-world crashes—are simulated in a Monte-Carlo fashion, in order to generate a small(er)

[2] The zeroth level is excluded here, as it only consists of a crude effectiveness assessment of the system based on expert assessments of the test scenarios, and no simulation is used.

number of crashes. Baseline situations are then based on these crashes, rather than on (statistical representations of) real-world crashes directly, as in the previous methods.

At all three levels, the initial step—before the baseline situations are established—is to identify the relevant set of test scenarios through analysis of crash databases. While the P.E.A.R.S. initiative identified the urgent need for sound input data "in terms of quality, representativeness, scalability, and real-world relevance" (Page et al., 2015, p. 4) in order to identify a relevant basis for test scenarios, a process for test scenario identification has not been thoroughly outlined. A relevant question for anyone who would like to enable safety benefit estimation of ADASs using virtual simulation is then: how should run-off-road test scenarios be defined to ensure relevance? The term relevance should be understood both in terms of relevance to real-world crash populations, (representativeness), and in terms of relevance for ADAS development and evaluation.

1.3 The state and intended role of the driver influences the effectiveness of run-off-road crash countermeasures

Notably, the effectiveness of a system is not only a function of the system itself, but of the cooperative effort of the vehicle (including the ADAS) and the driver—together referred to as the *joint driver-vehicle system* (JDVS) (Engström & Hollnagel, 2007). This cooperative effort can be conceptualised as the *situational control* exerted by the JDVS. Situational control is defined as the JDVS's ability to manipulate its trajectory within a *driver, vehicle, and environment* (DVE) state space (Ljung Aust & Engström, 2011). In the vast majority of run-off-road crashes, the critical reason for the crash is attributed to the driver (Kusano & Gabler, 2014; Liu & Ye, 2011; Pomerleau et al., 1999, pp. 7–8). Hence, situational control would be described in a space where one or several parameters are influenced by the behaviours of the driver, such as choice of speed or level of attentional effort. To exemplify, insufficient attentional effort on behalf of the driver decreases the monitoring performance of the driver, which reduces the situational control of the JDVS and eventually results in an imminent road departure. This can be visualised as a JDVS trajectory in the DVE space (Figure 1). Figure 1 also visualises how different run-off-road crash countermeasures are, conceptually, intended to enhance the level of situational control available to the JDVS so that the crash can be avoided or mitigated. LDW and rumble strips are visualised as providing an increase in the driver's attentional effort, followed by driver-initiated corrective steering manoeuvres to reduce the urgency of the situation. LKA and ELK, on the other hand, are intended to induce a corrective yaw rate of the vehicle, which can be visualised as a reduction in the urgency of the situation followed by an increase in driver attentional effort. LKA activate when the vehicle is approaching the lane markings, while ELK activate in relation to the road edge. Thus, their corresponding trajectories will start to deviate from a run-off-road crash trajectory at a different time to road departure. It should be noted that despite their analogous relation, ELK is an always-on approach whereas LKA can be turned off, and LDW is typically available on all roads with visible lane markings whereas rumble strips are only present on certain roads.

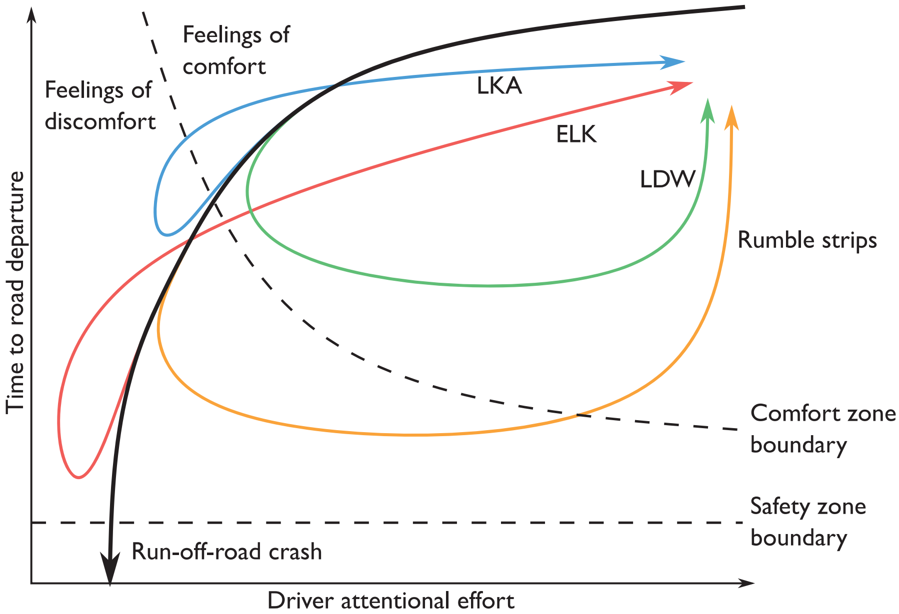

Figure 1: A run-off-road crash, visualised as a JDVS trajectory in the DVE space at a conceptual level. Decreased driver attentional effort leads to an imminent road departure. The intended effect of LDW and rumble strips is visualised as an increase in driver attentional effort, followed by corrective steering manoeuvres, whereas the intended effect of LKA and ELK (road-edge LKA) is visualised as a decrease in urgency (longer time to road departure).

However, it is important to capture the cooperative effort of the JDVS during ADAS effectiveness assessment beyond the conceptual stage. This can be exemplified by considering that one common implementation of LKA functionality is the use of steering wheel torque actuators. However, several studies have reported that the steering wheel adjustment resulting from a steering wheel torque intervention depends heavily on the (neuromuscular) response from the driver, as well as the implementation of the haptic intervention (Abbink, Cleij, Mulder, & Van Paassen, 2012; Benderius, 2014; Navarro, Mars, & Young, 2011; Petermeijer, Abbink, & de Winter, 2014). Thus, the intended role of LKA/ELK, as visualised in Figure 1, may not automatically hold true for all implementations of such systems.

To summarise, ADASs targeting run-off-road crashes can be described in terms of the cooperative effort of the JDVS, where the (re-)actions of the driver to a warning or intervention will, to a varying degree, determine the outcome of a system intervention in a given conflict situation. Therefore, it is important to acknowledge that the effectiveness of an ADAS depends on the state of the driver, the driver's response to, and intended role in relation to, the ADAS (i.e. how the shared control is distributed), as well as the actions of the driver.

1.4 Virtual simulation requires computational models of driver behaviour representative of targeted test scenarios

Since the effectiveness of an ADAS depends on the actions of the driver, a representation of the interactions between the JDVS and the DVE states is required for successful effectiveness assessment (Markkula, 2015). In their proposed framework for evaluation of active safety functionality, Ljung Aust and Engström (2011) conceptualises this interaction in terms of *satisficing*[3] control on behalf of the driver, in relation to a *goal state* in the DVE space (see Figure 1). This description suggests that drivers will adapt to changing circumstances by controlling the vehicle in order to stay within a comfort zone, and maintain a satisfying and sufficient distance (in DVE space) to situations that may induce feelings of discomfort (Summala, 2007), rather than to optimally reach the goal state (such as always being in the centre of the lane). However, the model of driver behaviour presented in Ljung Aust and Engström's framework is an example of a *conceptual model*, and even if such a qualitative description can support understanding of mechanisms and causal factors, it is insufficient for formative evaluation of ADASs based on virtual simulation. This is a topic discussed in more depth by Markkula (2015), where it is noted that virtual simulations require a more quantitative account of driver behaviour that, mathematically, describes the JDVSs effort to navigate the DVE space. Markkula further states that such quantitative driver behaviour models should, at least, be able to describe relevant aspects

[3] *Satisfice* is a portmanteau of *satisfy* and *suffice*.

of driver control in the designated target scenario—both with and without the ADAS. Such models typically take the form of a *statistical model*, such as a distribution of reaction times, or of a *process model* that, in a closed-loop manner, describe a reaction based on observed quantities. Many quantitative process models exists, capturing various aspects of driver lateral control (Gordon & Magnuski, 2006; Markkula, Boer, Romano, & Merat, 2017; Najm & Smith, 2004; Salvucci & Gray, 2004; Weir & McRuer, 1970). However, it has been shown that the outcome of counterfactual simulations is sensitive to small variations in the design of the driver behaviour model (e.g. Bärgman et al., 2017), so careful model selection is crucial in the evaluation procedure. Depending on the test scenario, existing models may therefore prove insufficient or unable to capture relevant aspects, requiring development of additional models. Specifically, they may not have been validated for crash-relevant situations. Therefore, it is important to understand what aspects that has to be described in various run-off-road situations, and how these aspects should be modelled.

In reference to the development of driver behaviour models, Carsten notes that "*the variety of models of the driving task is almost as numerous as the number of authors who have contributed the models*" (Carsten, 2007, p. 105). This is an observation that appears to hold true for computational models of driver control in critical situations as well (see e.g. Markkula, Benderius, Wolff, & Wahde, 2012). Harmonisation in the development of (computational) driver behaviour models is therefore desired, to create synergies and enable knowledge transfer. This can be obtained through the use a unifying conceptual framework that enables identification of relevant aspects, and supports the process of developing models of these aspects. Such a framework can preferably be adopted from contemporary neuroscience, cognitive science, or psychology—in order to increase generalisability and reusability of models (Markkula, Benderius, Wolff, & Wahde, 2012, p. 1137).

1.5 There is a need for more specific explanations of driver behaviour in run-off-road crashes

In the above sections, some of the requirements for virtual simulation-based effectiveness assessment of ADASs targeting run-off-road crashes have been described. Specifically, focus was put on the need for representative and relevant run-off-road test scenarios, and the need for quantitative, computational models of driver behaviour with validity in the selected target situations. These two aspects have in common that they stress the need for more specific explanations of driver behaviour in run-off-road situations;

relevant test scenarios rely on an account of driver behaviour in the target crash population, and computational models should quantify such specific driver behaviours in virtual simulations of these test scenarios. However, in order to reach specific explanations, there is a need for detailed, relevant, and reliable data able to represent behaviours exhibited in real-world situations. Notably, no data set is currently able to satisfy all of these conditions, meaning that each set of data has to be analysed and used in a way that best exploits their corresponding strengths. Large national crash databases are able to ensure representativeness and could be used to describe relevant target populations, while data from naturalistic driving studies (NDSs) or field operational tests (FOTs) are able provide sufficient detail to support the development of driver behaviour models.

2 Objectives

The overall aim of this PhD work (of which this licentiate thesis is a part) is to progress the development of driver behaviour models, from general explanations of driver behaviour to quantitative explanations of specific behaviours in safety critical situations potentially leading to run-off-road crashes (see Figure 2). Conceptual models of driver behaviour can support general understanding of processes or mechanisms in automobile driving, but they tell us less about the specific behaviours exhibited in real-world crash situations. To form hypotheses about specific explanations, analyses of driver behaviour in more detailed and specific situations should be done. However, an increase in detail and specificity is often coupled with scarcity and a decrease in the representativeness of the data. It is therefore important to connect the use of representative, but low-detail crash data, with the use of naturalistic driving study data, and the collection of experimental data with higher levels of detail and experimental control. Once more specific explanations of driver behaviour are obtained, modelling efforts will be able to inform development and evaluation of ADASs—such as LDW and LKA/ELK—by supporting virtual simulation of ADASs and enabling safety benefit analysis.

The following five objectives have been set to ensure meeting the overall aim.

1. Define a basis for test scenarios for run-off-road crashes—based on analysis of in-depth cases from crash databases—that can inform the design of ADASs to prevent these crashes.
2. Identify a conceptual framework suitable for describing human behaviour, and apply it to modelling of the behavioural mechanisms that lead to run-off-road crashes.
3. Identify and describe the behavioural mechanisms leading to run-off-road crashes from naturalistic data.
4. Develop driver models of the behavioural mechanisms leading to run-off-road crashes, to support virtual simulation and validation of ADASs.
5. Implement the driver models in a virtual simulation environment to evaluate ADASs.

The first two objectives have been achieved for this licentiate thesis; while all five objectives are intended to be achieved for the PhD thesis.

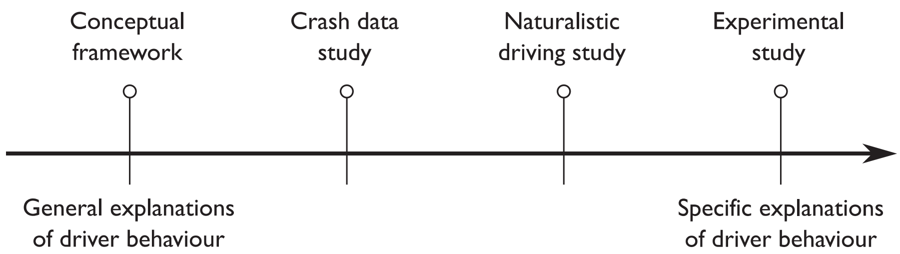

Figure 2: Development and evaluation of ADAS progresses as explanations of driver behaviour becomes more specific. The research work covered in this thesis supports progression of our explanations of driver behaviour in run-off-road situations, from general towards specific, by presenting a unifying conceptual framework of driver behaviour, and by using real-world crash data to define behaviourally relevant run-off-road test scenarios for ADASs.

3 Defining test scenarios relevant for development and evaluation of ADASs targeting run-off-road crashes

An explicit objective of this thesis work is to define a basis for test scenarios for run-off-road crashes, in a way that can inform the design of ADASs to prevent these crashes. Previously, test scenarios for effectiveness assessment and safety benefit estimation have, to a large extent, used a top-down (or rule-based) approach to analysis and target situation selection from crash databases, based on categorisation of crashes according to a set of *conflict situations* (e.g. Najm, Koopmann, Boyle, & Smith, 2002). This chapter introduces the work covered in Paper I, wherein categorisation of crashes instead was done using the statistical method of cluster analysis. First, a brief argument will be given as to why it may be of interest to go beyond a pre-defined categorisation of crashes, followed by an argument for the use of cluster analysis. Last, some important considerations when using cluster analysis for test scenario definition will be presented.

3.1 Beyond our best understanding of conflict situations

In several previous efforts to define run-off-road test scenarios, the starting point has been to define a set of conflict situations that describe, at high levels of abstraction, the events preceding the crash. Two examples of such high-level descriptions—obtained from the pre-crash scenario typology proposed by Najm, Smith, and Yanagisawa (2007)—are "road edge departure without prior vehicle manoeuvre" and "control loss with prior vehicle action". The next step has then been to subdivide the crashes in a crash database according to the chosen conflict situation typology, and derive test scenarios using the distribution of pre-crash factors, e.g. vehicle manoeuvres, speed, or number of lanes. This general approach has previously been applied to define run-off-road test scenarios (Najm et al., 2002). Similarly, run-off-road crashes were categorised according to so-called driving scenarios in the ACAT project (Gordon et al., 2010), where the driving scenarios were obtained through analysis of causal factors and crash sequences found in crash data. For the simulations in the ACAT project, various data sets (including crash databases and in-depth crash studies) were sampled for the given driving scenario, in order to obtain specific values for required parameters (Gordon et al., 2010, Chapter 4). While Najm et al. (2002) and Gordon et al. (2010) generate their baseline situations at levels 2 and 3 (as described in Chapter 1.2), it is also possible to use the selected crashes directly, corresponding to level 1 (Alvarez et al., 2017). This was done by Kusano, Gabler, and Gorman (2014), who estimated the safety benefit of

LDW based on crashes from the National Automotive Sampling Systems, Crashworthiness Data Systems (NASS/CDS) database, where the crashes were selected according to previously defined pre-crash scenarios (Kusano & Gabler, 2013).

One issue with the top-down approach is that crashes are divided according to the predefined conflict situations, forcing that structure upon the data and implicitly claiming that it is the most reasonable way to divide the data. In this process, similarities and patterns existing in the data may be disregarded. To exemplify, crashes triggered by driver fatigue and sleepiness may result both in drift-off-the-road events and in loss-of-control events, potentially classifying them as belonging to different conflict situations. Still, both crashes may have been avoided using the same countermeasure strategy (e.g. a driver state monitoring system with a warning). In some sense, the predefined conflict situations are based on the currently best understanding about the structures in the data, and if a certain pattern is not already accounted for in the existing conflict situations, it will be ignored. Specifically, as in the example with fatigue-related crashes above, structures relevant for ADAS development or crash mechanism analysis may be overlooked. Therefore, alternative approaches should be examined, based on the requirements imposed by desired properties of test scenarios. Desired properties of a test scenario could include relevance for system development and evaluation, namely that the crashes represented by a test scenario should—to the largest extent possible—be addressed by a specific countermeasure, and be defined in parameters available to a system. Test scenarios should also be based on a minimal number of assumptions about how the underlying data is best categorised. This speaks for the use of a statistical method that is data-driven and exploratory, while still allowing some level of control over the relation between input and desired output. One potential approach is to use cluster analysis, which is an objective statistical method that can support researchers in identifying structures and segmenting of crashes according to their relative similarities (Depaire, Wets, & Vanhoof, 2008).

3.2 Cluster analysis provides a data-driven approach to defining test scenarios

The cluster analysis approach used in Paper I can be described as bottom-up or data-driven, and is part of a family of methods that finds clusters in data according to homogeneity, similarity, or relationship between observations, in some statistical sense. Due to the objective statistical nature, (unsupervised) cluster analysis has previously been used in crash data analysis, but only to a limited extent for defining test scenarios. Cluster analysis as a method for crash data analysis has been used, for example, to find "hotspots", where geographical distances between crashes have been used to identified areas with elevated crash rates (Anderson, 2009; Bíl, Andrášik, & Janoška, 2013; Kim & Yamashita, 2007). It has also been used as a way to improve understanding of patterns in road traffic crashes, and for identification of relevant pre-crash factors, for a large number of purposes (Berg, Gregersen, & Laflamme, 2004; De Oña, López, Mujalli, & Calvo, 2013; Depaire, Wets, & Vanhoof, 2008; Kaplan & Prato, 2013; Nowakowska, 2012; Sasidharan, Wu, & Menendez, 2015; Theofilatos & Efthymiou, 2012; Weiss, Kaplan, & Prato, 2016). Cluster analysis thus provides a general method able to support specific purposes. For the explicit purpose of defining test scenarios, cluster analysis of in-depth crash data has been used on a smaller scale. For example, cluster analysis was applied to the STATS19-OTS database in order to establish test scenarios for autonomous emergency braking in car-to-pedestrian situations (Lenard, Badea-Romero, & Danton, 2014). By basing the cluster analysis on representative real-world crash data, the cluster solutions will automatically carry relevance for, and be representative of, the real-world crash population upon which the analysis was carried out, which is a desirable property of test scenarios.

There are other arguments suggesting that cluster analysis could be a suitable method for defining test scenarios. Cluster analysis is essentially an exploratory method, meaning that it does not focus on testing existing hypotheses, but rather enables identification of structures in the data (regardless of the analysist's previous knowledge of these structures). This can be used to facilitate formulation of hypotheses regarding causes for the observed features in the results. As such, cluster analysis is a method that is driven by the underlying structures in the data, and that makes minimal assumptions about the nature of the data. This means that, by using cluster analysis, it may be possible to capture patterns in the crashes that can facilitate our understanding of their causal mechanisms. The argument for cluster analysis as a method to define run-off-road test scenarios is thus an idea is that crashes that have a similar set of pre-crash factors may also share similar causal structures. Furthermore, by basing the cluster analysis on a set of variables that represent properties important for ADAS development (e.g. within the limitations of sensors), the results will be coupled to system relevance.

In summary, the argument is that an approach based on cluster analysis may be better, compared to previously used top-down approaches, at capturing structures created by similar causation factors and crash mechanisms. This property is of elevated importance if the test scenarios should support development of effective countermeasures (e.g. conceptualised in terms of situational control), and also provide a better understanding of the driver's role in the crash, for it to be modelled properly. In addition, cluster analysis encompasses means to, simultaneously, maintain relevance for ADAS development.

3.3 Care should be taken when selecting cluster analysis method

As explained in the previous section, an exploratory, unsupervised cluster analysis method would be well suited for the purpose of defining test scenarios. Furthermore, the method should provide some way to determine the number of clusters, as the total number of clusters in a set of crashes is not known beforehand. For this purpose, *hierarchical agglomerative clustering* offers a suitable alternative (Kaufman & Rousseeuw, 2005, Chapter 5). In hierarchical agglomerative clustering, each observation (crash) is initially considered a separate cluster. In the next step, the two clusters that are—in some sense—closest to one another are merged. Distance, in this case, could be any suitable measure of (dis-)similarity, such as Euclidean distance or Manhattan distance. Proximity between clusters is decided by an agglomeration (linkage) criterion. Since a test scenario is supposed to be representative of a distinct set of crashes, an agglomeration criterion that yields high within-cluster similarity is preferable. For that purpose, *complete linkage*[4] provides a good alternative, as it tends to create relatively compact clusters with high within-cluster similarity (Kaufman & Rousseeuw, 2005, p. 227). The agglomeration step is then repeated until all crashes has been merged into one cluster, creating a cluster hierarchy (Kaufman & Rousseeuw, 2005, Chapter 3). The final step is then to determine the number of clusters from the cluster hierarchy by examining how a clustering criterion changes at each merging step (Kaufman & Rousseeuw, 2005, p. 208). The hierarchical merging is often visualised as a tree structure, going from individual clusters to a single cluster. Hence, the last step consists of using a statistical measure

[4] Complete linkage is also known as *furthest neighbour*, as the dissimilarity between the most remote observations of two clusters is used as the between-cluster dissimilarity measure.

that describes the most likely number of clusters present in the data to decide where to cut the hierarchical tree structure.

When performing cluster analysis, selection of variables to include in the analysis should be done with care. First of all, variable selection has strong influence on the outcome of the cluster analysis (Milligan & Cooper, 1987; Punj & Stewart, 1983). Using run-off-road crashes as an example, the results are highly likely to be different depending on whether the analysis uses information on road surface conditions or not, as sufficient road friction is essential for maintaining lateral control. Likewise, one needs to select variables that reflect the purpose of the analysis. In the case of ADAS development and evaluation, variables should be chosen that carry relevance for system specification (Conquest, Spyridakis, Haselkorn, & Barfield, 1993; Lenard et al., 2014). Consequently, analysis of crash severity and injury levels would require a different set of variables (De Oña et al., 2013; De Oña, Mujalli, & Calvo, 2011). For example, personal characteristics of the driver, such as gender, age and weight, could be important for injury mechanism analysis, but may be less interesting for ADAS development. In the context of ADAS development, for age and gender to be parameters available to the system, installation of sophisticated sensory equipment not currently available in vehicles would be required, which would likely not provide a cost-efficient strategy.

Furthermore, police reported crashes and in-depth case studies both suffer from the fact that different variables have different reliability. For example, each crash is often ascribed a cause, which is typically the result of a subjective assessment process, while other parameters such as the road width can be determined by objective measures. In the arguments for cluster analysis presented above, the idea is presented that a similar set of pre-crash factors may imply similarities in causal structures. Therefore, one promising approach is to exclude the (relatively) unreliable variables in which crash causation is attributed from the cluster analysis, and use them as a means to evaluate the soundness and usefulness of the resulting cluster solution. In addition, cluster analysis provides a unique opportunity to compare different cluster solutions, based on changes in agglomeration criterion or input variables. Thus, it is possible to, for example, iteratively include or exclude variables and evaluate the usefulness of the resulting cluster solutions.

To summarise, cluster analysis as a method enables definition of test scenarios from identified structures in the data, obtained using a multivariate measure of crash similarity in an objective and scalable manner. This has the benefit of allowing researchers to specify a set of parameters with relevance to the current objective, without having to put limitations on the potential outcomes beforehand. The method can also be seen as exploratory, in that it allows for variations in how the clustering is done (e.g. inclusion/exclusion of parameters) in order to iteratively produce and interpret results to maximise the usefulness of the obtained solution. The usefulness can, for example, be evaluated using attributed causes for the crash, if such variables are purposefully excluded from the cluster analysis itself.

4 A conceptual framework for automobile driving and the predictive processing account of lateral control

Throughout the history of traffic safety research, the role of the driver has attracted a great deal of attention (Evans & Schwing, 1985). By adopting theories developed in a number of fields, such as psychology, cognitive science, and human factors, researchers have been trying to understand how humans carry out the driving task and, perhaps more urgently, the safety aspects of driver behaviour. This chapter will deal with the theoretical background to the contents of Paper II, in which the neurocognitive framework often labelled as *predictive processing* is applied to automobile driving. This work covered in Paper II was carried out in order to identify and establish a conceptual framework able to support and unify explanations of driver behaviour, and to harmonise the development of driver behaviour models. Furthermore, a conceptual framework can support generalisation of specific explanations by putting findings in a larger context.

In Section 4.1 below, a brief summary will be given to explain key concepts of the predictive processing framework. In Section 4.2, an expansion beyond what is covered in Paper II will be provided, where the aim is to elaborate on how the predictive processing account can be used to understand lateral control in automobile driving.

4.1 Active inference: infer the state of the world through predicted sensory stimuli, and act to realise predictions

Predictive processing can be seen as an umbrella term, capturing a range of models for the human brain functions (Clark, 2013). The brain is viewed as a statistical entity that tries to infer the state of the world, in a Bayesian fashion, in order to predict its own sensory input by tapping into regularities in the environment. By maintaining a *hierarchical generative model* that describes the relation between causal structures in the world and the resulting sensory stimuli, the brain can predict its sensory stimuli by inferring the state of the world (Friston, 2010). If the actual sensory stimuli agree with the predicted sensory stimuli, it can be seen as confirmation that the *external* state that causes the *sensory* state was correctly inferred. If there is a disagreement between the actual and the predicted sensory state, the brain typically tries to cancel out or minimise this *prediction error* in two ways: either by *updating the predictions given the sensory stimuli*, or by *taking actions that will align the sensory stimuli with the predicted stimuli*. This prediction error minimisation is the core mechanism of the predictive processing account.

Perception, in this sense, is the process of minimising prediction errors by updating one's beliefs about the world, while *action* is seen as manipulation of external states in order to bring about the predicted sensory state. Or, as Clark puts it: "you treat the desired (goal) state as observed and perform Bayesian inference to find the actions that get you there." (Clark, 2013, p. 6). It is noteworthy that actions themselves are explained using the same description, namely that motor control strive to fulfil predictions about proprioceptive sensations (Friston, Daunizeau, Kilner, & Kiebel, 2010). What the predictive processing account hence brings is a new take on the classical notion of perception and action, in which they are unified under the label of so-called *active inference*. That is, perception and action can be seen as manifestations of the same thing—they are the means by which prediction errors are minimised. Specifically, action is the consequence of unexpected sensory stimuli.

One important aspect of the predictive processing account is the hierarchical nature of the generative model. Persistent prediction errors at a lower level will propagate upwards to generate prediction errors at higher levels, and new predictions will be sent downwards (formally, this can be thought of as a new *priors*) (Friston, Shiner, et al., 2012). This is a process similar to the concept of *predictive coding* of visual information in the visual cortex, as proposed by Rao and Ballard (1999). As predictions are generated at all levels of the hierarchy, it is the overall prediction error that should be minimised. This process is weighted by the reliability of the predictions, namely by their precision (inverse variance), in a process similar to that ascribed to attention or biased competition (Friston, 2010, pp. 132–133).

The most comprehensive formal treatments of the predictive processing account is likely provided by Karl Friston and colleagues. They propose the *free energy formulation*, which is a Bayesian description of the brain functions, working under the assumption that brain activity strive to minimise the *variational free energy* (Friston et al., 2010; Friston, Kilner, & Harrison, 2006). The contents of the formal treatment will not be discussed in depth within this thesis, but interested readers are encouraged to explore published texts, providing accounts of e.g. visual processes and perception (Friston, Adams, Perrinet, & Breakspear, 2012), motor control (Adams, Shipp, & Friston, 2013; Friston, 2011), and active inference and learning (Friston, FitzGerald, Rigoli, Schwartenbeck, & Pezzulo, 2017).

In summary, the predictive processing account describes the brain as a statistical entity that tries to predict its sensory states based on its beliefs about the relation between causes and sensations—formed by the statistical regularities in the environment. If predicted and actual sensory stimuli is in disagreement, the brain either updates its beliefs based on the prediction error (perception), or tries to bring about the predicted state (action), in order to minimise the overall prediction error.

4.2 A predictive processing account of lateral control in automobile driving

To exemplify the predictive processing account, let us imagine a driver travelling on a curvy road. Under the assumption that the driver's generative model is properly tuned, it will provide meaningful predictions of the state of the visual stimuli when the future path coincides with the desired trajectory. At the lower level of the hierarchy, this prediction is compared to the actual visual stimuli. While the detailed content of the prediction is not known, it could, for example, encode a stationary state of near and far points, as in the two-level model of steering proposed by Salvucci and Gray (2004). If a prediction error exists, that is, if there is a deviation from the predicted state in the actual stimuli—such as a movement of the far point—the driver will seek out actions believed to realise or fulfil the predicted state. That is, by taking corrective steering actions, the driver will stabilise the far point and bring about the predicted sensory stimuli.

In order for this process to be successful, so that the driver can stay on the road, the brain has to model the relationship between visual stimuli from the forward roadway (sensory states) and the corresponding required steering manoeuvres (active states)—phrased more strictly, the brain should model the way external states (the vehicle's trajectory in the road environment) give rise to sensory states (optical flow patterns), as well as the way actions (steering corrections) affect the external states (a change of the vehicle's trajectory). However, external states are not directly available to the brain, in a statistical sense, but have to be inferred from the agent's internal, sensory, and active states (Friston, 2013). In the same way, the brain does not possess direct influence over external states *per se*, but may indirectly influence external states by changing its active states. If the active states influence the same external states that cause the current sensory states a dependency is created, which Friston describes as "a circular causality that is reminiscent of the action-perception cycle […] external states cause changes in internal states, via sensory states, while the internal states couple back to the external states through active states—such that internal and external states cause each other in a reciprocal fashion" (Friston, 2013, p. 2). In other words, the driver should both be able to predict the visual stimuli that indicates following the road, and be able to predict the required steering manoeuvre in cases where the visual stimuli indicate the driver is not following the road ahead. However, given the nature of the predictive processing account, it is possible to find explanations to a number of phenomenon observed in relation to the driver's ability to maintain and exert lateral control. Below, three examples will be given, namely the aspect of satisficing control, gaze target selection, and driver behaviour in relation to rare events.

4.2.1 Satisficing control as an optimising strategy

One implication of the variational free energy formulation of the predictive processing account is that the generative model will be maximising its accuracy, while at the same time minimising the model complexity (Friston, FitzGerald, Rigoli, Schwartenbeck, & Pezzulo, 2017). Implicitly, this means that causes and sensory effects can be understood as a few-to-many relationship (Trappenberg & Hollensen, 2013), and favours the use of simple models of the causal structures in the world that, at the same time, provides sufficiently good predictions about the sensory states (Pezzulo, 2014). To understand why this is important in relation to lateral control, it should be coupled with the notions that: (A) the brain makes (steering) decisions based on noisy accumulation of evidence that a steering manoeuvre is required, resulting in an inherent intermittency in control application (e.g. Gold & Shadlen, 2007; Markkula, Boer, Romano, & Merat, 2017) ; and (B) steering manoeuvres can be understood in terms of execution of motor primitives (Benderius & Markkula, 2014; Martínez-García, Zhang, & Gordon, 2016). Thus, feedback from a corrective manoeuvre is not instantaneous, as natural variability from the road environment and vehicular response has to be decoupled from the effect of the manoeuvre in the noisy evidence signal. This yields a level of uncertainty, or variability, in the monitoring and control task, meaning that the performance of the generative model is bound by the statistical regularities in its environment—no predictions can be made with any certainty in an environment that is inherently irregular. Furthermore, as noted above, error minimisation occurs in a precision-weighted manner for all predictions made by the generative model simultaneously, at all levels in its hierarchy. It may therefore be part of an overall optimisation strategy to adopt a simpler, but less precise, model for low-level lateral control, that may reduce the precision (and thus weighting) of deviations in lane centring without compromising the performance of higher level predictions (i.e. to keep the vehicle on the road). This lends support to the notion that the lane keeping task, in itself, can be understood in terms of satisficing rather than optimising control (Summala, 2007).

4.2.2 Gaze target selection should minimise overall uncertainties about predictions

The second aspect of lateral control where predictive processing presents intuitive explanations is the role of attention and visual target selection in automobile driving. Lane keeping is only one task among many in automobile driving that requires visual monitoring to generate and correct predictions, and competes with tasks such as monitoring of the speedometer or the rear-view mirrors. In addition, the driver may have higher motives, creating additional predictions about e.g. the state of the climate control or the location of a song in a playlist. Now, given that the lane-keeping task is considered satisficing, resolving the contradictions of multi-tasking through active inference will prioritise the tasks that—at any given time—produces the highest precision-weighted prediction errors (or, in some sense, features the highest uncertainties about the correctness of the predictions). Gaze target selection then corresponds to the driver's beliefs on where visual sampling will result in the largest overall reduction in uncertainty (Friston, Adams, Perrinet, & Breakspear, 2012). Thus, if a driver predicts, with high certainty, that a road departure is not imminent, it will become more important to allocate visual attention in a way that reduces uncertainties for other tasks than the lane-keeping task. In other words, the more predictable a traffic environment is, the more likely it is that the driver will engage in "secondary" tasks. This notion finds support in studies of driver glance length patterns in naturalistic driving studies (Tivesten & Dozza, 2014).

4.2.3 False certainty in the perception-action cycle dictates driver behaviour in rare events

As noted above, the perception-action cycle can be seen as driven by the brain's ability to predict how actions influences external states, and how these external states in turn are predicted to cause future sensory states. The predictions, on the other hand, emerge from the brain's beliefs about the world, which in turn are based on previously experienced regularities. In other words, the more frequent and reliable a situation is, the more likely it is that the brain will make correct predictions about current and future states, and how they are affected by specific actions. Conversely, if the brain encounters a situation that rarely occurs (such as imminent road departures), it is likely that the brain will be less capable to make correct and high-precision predictions about the effects of a specific action. This would then especially hold true if the results of an action do not adhere to regularities exhibited in situations the driver more frequently experiences (such as low road surface friction). However, due to lack of experience, the brain runs the risk of having a false certainty about the rarely encountered situation, resulting in a failure to perceive the irregular external state (e.g. the generative model does not provide relevant predictions for sensory stimuli related to a reduction of road surface friction). In other words, when encountering rare situations, the brain's beliefs about the state of the world will not necessarily account for the rare state, resulting in a false certainty that the perception-action cycle valid in regular situations is also valid for the irregular situation.

4.3 Summary

In Section 4.1 above, a brief introduction to the predictive processing framework was given, introducing the concept of active inference, through which perception and action can be understood as ways to reduce uncertainties about the current state of the world by minimising overall deviations from predicted sensory inputs. These concepts were then applied to automobile driving in Section 4.2, where focus was put on investigating concepts relevant to lane keeping and run-off-road crashes beyond what was covered in Paper II, bringing up concepts such as satisficing control, gaze target selection, and false certainties about causal structures.

5 Summary of papers

A summary of the two papers included in the thesis will be given in the next sections. Paper I deals with defining run-off-road test scenarios using cluster analysis of in-depth crash data. Paper II investigates how the (neuro-)cognitive framework of predictive processing can be applied to automobile driving.

Paper I

Nilsson, D., Lindman, M., Dozza, M., & Victor, T. (2017). *Definition of run-off-road crash clusters—for safety benefit estimation and driver assistance development.* Manuscript submitted for publication.

The study was designed by Nilsson, Lindman, Victor and Dozza. Nilsson carried out data reduction and statistical analyses. Nilsson authored the paper, with minor contributions from Victor. Essential feedback was provided by Lindman and Dozza.

Paper II

Engström, J., Bärgman, J., Nilsson, D., Seppelt, B., Markkula, G., Piccinini, G. B., & Victor, T. (2017). Great expectations: a predictive processing account of automobile driving. *Theoretical Issues in Ergonomics Science*, (April), 1–39. http://doi.org/10.1080/1463922X.2017.1306148

The paper was authored by Engström. Nilsson participated in the literature study, partook in discussions of the topic, and provided feedback on the text during workshop sessions and in personal communication in between sessions.

Paper I. Definition of run-off-road crash clusters—for safety benefit estimation and driver assistance development

Introduction

To ensure adequate real-world crash prevention performance of ADASs, information on causal factors in real-world crashes should be included in test scenarios used in the design and evaluation of the ADAS. Previously, test scenarios for run-off-road crashes have been based on dividing real-world crashes according to a conflict situation typology, hence running the risk of disregarding underlying structures in the data.

Aim

The goal of this study was to create a basis for test scenarios for run-off-road crashes, by dividing real-world crashes based on a multi-variate similarity measure, rather than on the more top-down conflict situation typology classification.

Method

The study was performed by applying hierarchical, agglomerative cluster analysis to run-off-road crashes from the German In-Depth Accident Study database, collected during the years 2008-2014. The cluster analysis was based on a number of variables deemed relevant for characterising run-off-road crashes. In addition to analysing the distribution of these variables for each cluster, the distributions of attributed causes were also analysed.

Results

The cluster analysis yielded a solution with 13 clusters, out of which four had a very low number of crashes (1-3) and were considered clusters of outliers. The remaining nine clusters were found to consist of crashes characterised by: drift during daytime, drift during night, high speeds, high departure angles on narrow roads, highways, snowy roads, loss-of-control on wet roadways, sharp curves, and high speeds in severe road conditions (Figure 3).

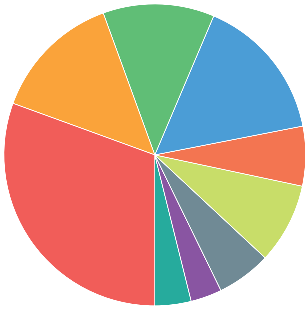

Figure 3: Distribution of the run-off-road crashes among the nine clusters along with the identified characteristics of each cluster.

Discussion

Nine groups of run-off-road crashes were found, showing both that run-off-road crashes occur in different ways, but also that there are patterns to be found in the data. These patterns provides a basis for improved understanding of crash causation mechanisms. Cluster analysis was thus found to provide an exploratory and statistically sound tool relevant for dividing real-world crashes according to underlying structures in the data. Analysis of the obtained cluster solution provided valuable insights, specifically highlighting scenarios with relevance for the development of ADASs to prevent run-off-road crashes. The two largest clusters were considered targets for current lane support systems, including lane departure warnings and lane keeping assist. However, limitations in the data on driver behaviour from retrospective crash databases highlight a need for additional data to specify causation mechanisms in higher detail.

Paper II. Great expectations: a predictive processing account of automobile driving

Introduction

Recently, predictive processing has emerged as a promising framework for describing brain functions in terms of minimisation of prediction errors. Neural activity can be understood in terms of a hierarchical generative model that generates predictions of incoming sensory stimuli. Thus, the brain only needs to deal actively with errors between actual and predicted stimuli. The brain, then, tries to minimise any prediction errors, either by a change of predictions, or by taking actions that bring about the predicted state—a process referred to as active inference. In this process, prediction errors are weighted by their respective precision, where precision describes the certainty of the prediction. Precision weighting thus serves as a way of controlling behaviour and select the most relevant action, correcting errors in high precision predictions. Learning, in this context, is thought of as model tuning, where the generative model tries to capture regularities in the environment to better infer the hidden states of the worlds that causes observed sensory stimuli.

Aim

This study aimed to survey the predictive processing account of human brain functions in order to provide a unifying framework for understanding human behaviour in an automobile driving context, creating synergies in ongoing driver behaviour modelling efforts.

Method

The study was conducted as a collaborative literature survey, organised through a series of workshops, where relevant literature was discussed with a focus on implications for safety research and human factors in automobile driving.

Results

By applying predictive processing concepts to automobile driving, new perspectives were obtained for a number of driving phenomena, such as driver brake response, countersteering, and driver interaction with automated driving (AD) functions.

Responses to braking lead vehicles have previously been described as triggered by (accumulated) optical expansion on the retina (looming). The predictive processing account suggests that it may instead be a response to *unexpected* looming, such that drivers may—in certain situations—predict the looming cues, thus inhibiting the brake response.

Similarly, a predictive processing approach suggests that it may be possible to cue drivers prior to autonomous steering interventions from ADASs, in order to inhibit countersteering behaviour.

Furthermore, the predictive processing account has important implications for the design of AD functionality. If the AD functionality exhibits timely and appropriate responses—providing consistent visual and vestibular cues to the driver—drivers will generate higher precision predictions, facilitating detection of functional limitations. However, high precision predictions and consistent AD responses may result in overreliance—where the driver's trust in the AD functionality exceeds its capabilities, typically in rare events—while at the same time it may reduce the driver's monitoring efforts.

Discussion

Predictive processing presents a potent framework, able to explain several elusive phenomena observed in automobile driving and provide guidance for future designs of ADASs and AD functionality. Rather than presenting a radical alternative to previous models and ideas in human factors, predictive processing could present a way to unify them, and further development and applications of these concepts are recommended.

6 Discussion

6.1 Test scenarios based on cluster analysis: exploratory but limited and purpose-specific

The first objective of the thesis, presented in Chapter 0, was to obtain a basis for run-off-road test scenarios that can inform the development of ADAS. In Paper I, a method to achieve this was described, wherein cluster analysis was used to reveal structures in run-off-road crash data, enabling formulation of plausible hypotheses about crash causation mechanisms. By exploiting similarities in pre-crash parameters to segment crashes, cluster analysis presented a promising approach for definition of ADAS test scenarios, where the resulting clusters define sets of unique real-world crashes that are similar to one another and could potentially be addressed using the same countermeasure strategy.

The use of cluster analysis to define test scenarios should be seen in contrast to the approach where crashes instead are divided according to an *a priori* understanding of how crashes are best categorised (e.g. Najm, Smith, & Yanagisawa, 2007). Cluster analysis has the benefit of being an exploratory method, providing a bottom-up approach where less assumptions has to be made regarding the underlying structures in the data. As such, it enables identification of patterns in the data that may otherwise have been overlooked. However, this is a benefit that is not free of costs. In some previous efforts (e.g. Kusano & Gabler, 2013), a more specific procedure have been used to select crashes relevant for ADAS effectiveness assessment. By specifying requirements for the data beforehand, based on properties that are relevant to the specific ADAS under assessment, selection could be done in a top-down manner. This specificity in data selection is lost with cluster analysis. Instead, careful and iterative selection of clustering method and input variables—coupled with interpretation of obtained clusters—can be used to find sensible cluster solutions that are relevant for specific ADASs. However, doing so removes some of the generality of the process, making it more of a data-driven top-down process. Thus, a cluster solution obtained using parameters relevant to ADAS development and analysis of crash causation mechanisms is not necessarily relevant in categorisation of crashes for other purposes, such as injury mechanism analysis, making it more of a purpose-specific approach.

Although Paper I provides a solid foundation for test scenarios based on crash data, the overall goal in which this thesis plays a part—the progression towards specific explanations of driver behaviour in run-off-road crashes—suggests a need for additional data sources in order to establish a comprehensive test bench for ADAS development. Parameters describing driver behaviour found in retrospective in-depth crash case studies contain limited information, as data on driver behaviour are often summarised as a stated cause for the crash, obtained by analysing the crash site and the interviews conducted with the driver according to some given causation analysis method (Larsen, 2004; Ljung Aust et al., 2010; Otte, Pund, & Jänsch, 2009). Consequently, the use of retrospective in-depth crash data imposes a severe limitation on the specificity of explanations of driver behaviour, both in terms of reliability (e.g. it is based on memories of subjective experiences) and in detail (e.g. due to a lack time series data). To exemplify, even if the cause of a run-off-road crash is attributed to driver distraction, there will be no detailed time-series data on the nature and character of the distraction. Therefore, test scenarios based on cluster analysis of crash data should not be taken as a complete description of the topography of driver behaviours in the metaphorical run-off-road crash landscape. It should also be noted that the resulting test scenarios from the cluster analysis may have been different if more detailed data on driver behaviour would have been available, as the method is inherently data-driven. Instead, the current cluster solution provides a crude map, pointing out general directions and areas of interest for further study. Specifically, the results presented in Paper I suggests that a primary target to investigate further is driver behaviour and performance in relation to distraction, intoxication, and fatigue. This can be done using data from naturalistic driving studies or field operational tests in which instrumented vehicles were used to collect detailed and objective data on driver behaviour with high validity (*FESTA Handbook Version 6*, 2016; Victor et al., 2015).

6.2 The use of a conceptual framework in analysis of crash data and for explanations of driver behaviour

The second objective of this thesis was to identify a conceptual framework suitable for describing human behaviour, and to apply it to modelling of the behavioural mechanisms that lead to run-off-road crashes. In Paper II, the predictive processing framework was presented as a promising candidate when applied to a number of safety-critical aspects of automobile driving, also able to unify multiple theories in human factors. In addition, Chapter 4 provided a specific account for some concepts relevant for driver lateral control, and subsequent run-off-road crashes caused by loss of lateral control.

In particular, the notion that drivers will be minimising the overall uncertainty of predictions has important implications for our understanding of everyday driving tasks. To exemplify, the concept of satisficing can be framed as an increased uncertainty in predictions about control demands, driven by changes in top-down predictions on required performance. By decreasing control demands, the driver enables reallocation of attention towards other sources of uncertainty, such as monitoring of the traffic environment, or a conversation with a passenger. Hence, strategies for gaze target selection and optimality in control should minimise uncertainties about all tasks in a precision-weighted manner and not just the driving task itself. This insight enables thinking about safety interventions at multiple levels of control (Michon, 1985). Current LSS systems support and intervene at the control-level, by redistributing control within the JDVS to address situations where a given allocation of attentional resources on behalf of the driver leads to loss of situational control. In that perspective, LSS systems deals with the symptoms of insufficient allocation of attentional resources towards the forward roadway, in relation to situational demands, and does not necessarily address the underlying cause. Naturally, strategies can be imagined that would increase precision at lower levels by encouraging high-level predictions about lane-keeping performance, for example through gamification of the driving task or increasing the driver's awareness of safety implications of long off-road glances, resulting in adaptive driver behaviour at tactical or strategic levels.

Furthermore, by adopting a predictive processing mind-set, it is possible to inspect the cluster solutions from Paper I in order to formulate hypotheses on explanations of driver behaviour, connecting specific explanations to useful generalisations. For example, the cluster with the highest proportion of crashes where driver distraction is the given cause also features fair driving conditions (daylight, dry and mostly straight roads). This is in agreement with the reasoning presented in Chapter 4.2.2, where it was hypothesised that high levels of predictability in the driving task should result in increased engagement in e.g. secondary tasks. Another example is the cluster characterised by crashes that occurred in severe road surface conditions (icy or snowy roads), where the most commonly stated cause for the crash is speeding (or rather, transgression of physical limits), which can be interpreted in terms of erroneous predictions of the physical limits of the JDVS in the current driving situation, as a direct consequence of limited experience from driving in adverse conditions. This explanation, also outlined in Chapter 4.2.3, is strengthened by the fact that this cluster typically features younger drivers[5], where age can be seen as a proxy for inexperience.

It should be noted that while application of a predictive processing mind-set allows for intuitive explanations of observed patterns, there is room for alternative explanations. For example, proneness to risk-taking behaviours (e.g. driving in severe conditions) that is typically present in a younger population of drivers would provide another plausible explanation of the age distribution in the latter cluster, even if the strength of this alternative explanation is potentially clouded by the fact that most drivers were travelling at a speed much lower than the posted speed limit[5]. Furthermore, explanations obtained from analysis of crash data should be understood as representing driver behaviour in crashes, rather than driver behaviour in general.

Still, the argument seems strong for predictive processing as a useful tool able to support conceptual understanding and explanations of behaviour in automobile driving in general, and for run-off-road crashes in particular, and its full potential remains to be seen in future research. What has already been seen, is that some of the aspects of the framework have had synergetic effects on driver behaviour modelling efforts (Markkula, Boer, Romano, & Merat, 2017; Svärd, Markkula, Engström, Bärgman, & Granum, 2017). Furthermore, an interesting future direction is to investigate the potential for a variational free energy formulation, as proposed by Friston and colleagues, in modelling of automobile driving, and it is yet to be explored.

6.3 Implications for development and effectiveness assessment of ADASs targeting run-off-road crashes

In collision avoidance, the role of an ADAS is to support the driver in particular aspects of the driving task in such a way that targeted crashes can be avoided (or mitigated). However, as clear from the analysis in Paper I, run-off-road crashes occur in many different ways and for many different reasons. To define a particular and well-specified aspect of the driving task in which an ADAS could provide sufficient support to avoid run-off-road crashes may therefore prove a laborious task, if even possible. Due to this, run-off-road crashes should perhaps be considered a misguided classification in relation to driver behaviour and ADAS development, albeit the "run-off-road" label carries merit for specific contexts of road traffic

[5] Unpublished data from the study presented in Paper I.

safety (e.g. Jakobsson et al., 2014). Fortunately, the results in Paper I also suggests that it is possible to subdivide run-off-road crashes into a limited number of sub-categories, for which specific countermeasure strategies, as well as driver behaviour models, can be developed. Naturally, to perform effectiveness assessment of a specific ADAS, only those run-off-road test scenarios wherein the ADAS is believed to make a difference should be selected.

Analysis of the results in Paper I suggests applicability for a range of potential countermeasure strategies, such as improved headlights, curve overspeed warnings, traction control systems, and making information on road surface conditions (friction) readily accessible to the driver. Specifically, the two largest clusters, totalling 46% of the analysed run-off-road crashes (see Figure 3, p. 27), can be understood as relevant targets for LSS systems, based on the high proportion of crashes related to distraction, intoxication, and fatigue. As such, they constitute a primary target for ADAS development and effectiveness assessment, and subsequently driver behaviour model development. However, the data used for the cluster analysis was taken from the German In-Depth Accident Study (GIDAS), years 2008-2014, and the database is considered representative for the German crash population (Hautzinger, Pfeiffer, & Schmidt, 2004). Hence, effectiveness estimation based on the obtained test scenarios should not be considered representative of non-German markets.

To support virtual simulation-based effectiveness assessment of LSS functionality in prioritised test scenarios, driver behaviour models must account for all processes relevant to the JDVS in those scenarios. This includes reactions to warnings, threat assessment, reaction to detected threats, reaction to system control interventions, and, if applicable, behaviours that may override interventions. In further specifying these aspects, the conceptual framework presented in Paper II can provide valuable support. To exemplify, as driver reaction to stimuli should be understood as updated predictions about the state of the world, based on noisy evidence accumulation—as briefly mentioned in Chapter 4.2.1—this implies that reaction times to a warning will depend on the current deviation from predicted (visual) stimuli, rather than being fixed values or reaction time distributions (Gold & Shadlen, 2007; Markkula, Boer, Romano, & Merat, 2017; Markkula, Engström, Lodin, Bärgman, & Victor, 2016). Hence, several previous studies on reaction times to LDW interventions (e.g. Stanley, 2006; Suzuki & Jansson, 2003) would have to be re-analysed and re-interpreted, or new studies should be conducted, in order to separate the re-allocation of attentional resources from the threat assessment (i.e. accumulation of unexpected visual stimuli suggesting a steering manoeuvre is required). Potentially, the use of peripheral vision must also be considered (Summala, Nieminen, & Punto, 1996). Recently, promising steps have been taken to describe driver steering control along these lines of thought, in particular by Markkula et al., who has been adapting the often cited model of steering control proposed by Salvucci & Gray (2004) to "concepts that are well established in contemporary neuroscience: *motor primitives*, neuronal *evidence accumulation*, and *prediction of sensory consequences of motor actions*" (2017, p. 3). However, to the author's knowledge, that model—albeit promising—has neither been validated and parameterised for driver behaviours exhibited in imminent road-departure situations, nor does it account for off-road glances. Furthermore, the potential for priming drivers prior to LKA control interventions in order to avoid countersteering and increase driver acceptance to steering wheel torques, as proposed in Paper II, has to be investigated further, coupled with an in-depth description of neuromuscular dynamics (Benderius, 2014; Cole, 2012). Hence, some work remains before successful development and effectiveness of ADASs targeting run-off-road crashes can be done using virtual simulation, but the contents of this thesis has provided a foundation for the work to come by extending our knowledge on relevant processes, and how they can be described in a unified way.

7 Conclusions

The aim of this PhD work, as described in Chapter 0, is to progress the development of driver behaviour models relevant to run-off-road situations, to enable timely ADAS effectiveness assessment using virtual simulation. To reach this goal, the two initial objectives were to: (1) define relevant test scenarios using representative in-depth crash data; and (2) identify a conceptual framework that can support modelling and understanding of behavioural mechanisms present in run-off-road crashes.

Paper I showed that while the run-off-road crashes occurred in different ways, and for different reasons, it was possible to find a smaller number of clusters wherein crashes showed high levels of similarity. Furthermore, by using the predictive processing framework proposed in Paper II, it was possible to analyse the contents of the clusters in Paper I to formulate hypotheses about crash causation mechanisms. Specifically, a predictive processing mind set suggests that distraction-related crashes are to be expected in less demanding traffic situations, as a driver's relative uncertainty about secondary tasks would increase when the primary (driving) task gets more predictable, meaning that more attentional resources will be directed towards the secondary tasks. This notion was reflected in the fact that the largest cluster, which had a high proportion of distraction-related crashes, was characterised by crashes that occurred on straight and dry roads in daytime, suggesting fair driving conditions.

By basing test scenarios on the results of exploratory cluster analysis, it was possible to find relevant targets both for ADASs currently on the market and for future systems. Specifically, the two largest clusters were found to be relevant for existing lane support systems, such as LDW and LKA/ELK, as the clusters were characterised by distracted, intoxicated, and fatigued drivers—suggesting that the crashes could potentially have been avoided had the driver been supported in the monitoring and control aspects of lane keeping. Furthermore, it was shown that while a conceptual framework of driver behaviour improves interpretation of the results from data-driven crash analysis, the analysis is hampered by limited reliability and detail in the descriptions of driver behaviour found in retrospective crash data. Hence, the work also highlights a need for additional studies in order to fulfil the overall aim—to develop quantitative, computational models of driver behaviour. Furthermore, Paper II will have implications for the use of previous findings in the field of human factors and driver behaviour modelling, as certain results must be re-interpreted in order to adhere to the story told by the predictive processing framework.

The work discussed in this thesis has been able to progress explanations of driver behaviour in run-off-road crashes in a way that is useful for development and evaluation of ADASs, ensuring the fulfilment of the first two objectives presented in Chapter 0, but has also paved the way towards fulfilment of the remaining three objectives. The results also have implications for the design of effective ADASs targeting run-off-road crashes, as an understanding of crash causation mechanisms is essential in the development of successful crash avoidance strategies.

8 Future work

The future work should help fulfil the overall aim presented in Chapter 0, namely to progress the development of driver behaviour models, from general explanations of driver behaviour to quantitative explanations of specific behaviours in safety critical situations potentially leading to run-off-road crashes. Specifically, driver behaviour models valid for effectiveness assessment of real-world crashes targeted by LSS are of primary interest. To reach this goal, additional analyses are required, targeting for example the use of visual cues and corresponding lateral control manoeuvres.

As outlined in Chapter 0, detailed data from on-road studies with high levels of validity, such as naturalistic driving studies or field operational tests, should be used to formulate more specific hypotheses about crash causation mechanisms. Specifically, it is of interest to look at threat assessment and control actuation, e.g. regarding the explicit nature of the visual stimuli predicted and perceived by the driver in situations relevant to run-off-road crashes, and the connection to steering wheel corrections, to understand driver behaviour in critical situation. This analysis should be based on existing knowledge on driver steering behaviour (e.g. Benderius & Markkula, 2014; Markkula et al., 2017; Martínez-García, Zhang, & Gordon, 2016) and required visual information (e.g. Land & Lee, 1994; Lappi, Lehtonen, Pekkanen, & Itkonen, 2013; Lappi, Rinkkala, & Pekkanen, 2017; Salvucci & Gray, 2004; Wilkie & Wann, 2003) and be framed by the predictive processing account presented in Paper II.

Specific hypotheses should then be tested and specified in further detail using controlled studies on test tracks or in driving simulators with sufficient validity (Mullen, Charlton, Devlin, & Bedard, 2011) so that the controlled variables (i.e. specific visual cues) can be reliably and readily manipulated. This could include manipulation of the availability of visual information, for example by experimentally induce (partial) visual occlusion during driving, while controlling the extent and/or duration of the occlusions. Once specific and quantifiable explanations are obtained, they should be incorporated into virtual simulation environments—either through relevant parameterisation of existing computational models of driver behaviour, or through development of new models—to enable timely ADAS effectiveness assessment with high validity.

References

Abbink, D. A., Cleij, D., Mulder, M., & Van Paassen, M. M. (2012). The importance of including knowledge of neuromuscular behaviour in haptic shared control. *Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics*, 3350–3355. http://doi.org/10.1109/ICSMC.2012.6378309

Adams, R. A., Shipp, S., & Friston, K. (2013). Predictions not commands: Active inference in the motor system. *Brain Structure and Function*, *218*(3), 611–643. http://doi.org/10.1007/s00429-012-0475-5

Alvarez, S., Page, Y., Sander, U., Fahrenkrog, F., Helmer, T., Jung, O., … Op den Camp, O. (2017). Prospective effectiveness assessment of ADAS and active safety systems via virtual simulation: A review of the current practices. In *The 25th International Technical Conference on the Enhanced Safety of Vehicles (ESV)*.

Anderson, T. K. (2009). Kernel density estimation and K-means clustering to profile road accident hotspots. *Accident Analysis and Prevention*, *41*(3), 359–364. http://doi.org/10.1016/j.aap.2008.12.014

Bartels, A., Rohlfs, M., Hamel, S., Saust, F., & Klauske, L. K. (2016). Lateral Guidance Assistance. In *Handbook of Driver Assistance Systems* (pp. 1207–1233). Cham: Springer International Publishing. http://doi.org/10.1007/978-3-319-12352-3_49

Benderius, O. (2014). *Modelling driver steering and neuromuscular behaviour*. Retrieved from http://linkinghub.elsevier.com/retrieve/pii/S2214241X13000047

Benderius, O., & Markkula, G. (2014). Evidence for a fundamental property of steering. *Human Factors and Ergonomics Annual Meeting*, 884–888. http://doi.org/10.1177/1541931214581186

Berg, H. Y., Gregersen, N. P., & Laflamme, L. (2004). Typical patterns in road-traffic accidents during driver training: An explorative Swedish national study. *Accident Analysis and Prevention*, *36*(4), 603–608. http://doi.org/10.1016/S0001-4575(03)00068-X

Bíl, M., Andrášik, R., & Janoška, Z. (2013). Identification of hazardous road locations of traffic accidents by means of kernel density estimation and cluster significance evaluation. *Accident Analysis and Prevention*, *55*, 265–273. http://doi.org/10.1016/j.aap.2013.03.003

Bärgman, J., Boda, C.-N., & Dozza, M. (2017). Counterfactual simulations applied to SHRP2 crashes: The effect of driver behavior models on safety benefit estimations of intelligent safety systems. *Accident Analysis and Prevention*, *102*, 165–180. http://doi.org/10.1016/j.aap.2017.03.003

Bärgman, J., Lisovskaja, V., Victor, T., Flannagan, C., & Dozza, M. (2015). How does glance behavior influence crash and injury risk? A "what-if" counterfactual simulation using crashes and near-crashes from SHRP2. *Transportation Research Part F: Traffic Psychology and Behaviour*, *35*, 152–169. http://doi.org/10.1016/j.trf.2015.10.011

Carsten, O. (2007). From driver models to modelling the driver: What do we really need to know about the driver? In *Modelling Driver Behaviour in Automotive Environments: Critical Issues in Driver Interactions with Intelligent Transport Systems* (pp. 105–120). http://doi.org/10.1007/978-1-84628-618-6_6

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*(3), 181–204. http://doi.org/10.1017/S0140525X12000477

Cole, D. J. (2012). A path-following driver–vehicle model with neuromuscular dynamics, including measured and simulated responses to a step in steering angle overlay. *Vehicle System Dynamics*, *50*(4), 573–596. http://doi.org/10.1080/00423114.2011.606370

Conquest, L., Spyridakis, J., Haselkorn, M., & Barfield, W. (1993). The effect of motorist information on commuter behavior: Classification of drivers into commuter groups. *Transportation Research Part C*, *1*(2), 183–201. http://doi.org/10.1016/0968-090X(93)90013-6

Dang, T., Desens, J., Franke, U., Gavrila, D., Schäfers, L., & Ziegler, W. (2012). Steering and Evasion Assist. In *Handbook of Intelligent Vehicles* (pp. 759–782). London: Springer London. http://doi.org/10.1007/978-0-85729-085-4_29

De Oña, J., López, G., Mujalli, R., & Calvo, F. J. (2013). Analysis of traffic accidents on rural highways using Latent Class Clustering and Bayesian Networks. *Accident Analysis and Prevention*, *51*, 1–10. http://doi.org/10.1016/j.aap.2012.10.016

De Oña, J., Mujalli, R. O., & Calvo, F. J. (2011). Analysis of traffic accident injury severity on Spanish rural highways using Bayesian networks. *Accident Analysis and Prevention*, *43*(1), 402–411. http://doi.org/10.1016/j.aap.2010.09.010

Depaire, B., Wets, G., & Vanhoof, K. (2008). Traffic accident segmentation by means of latent class clustering. *Accident Analysis and Prevention*, *40*(4), 1257–1266. http://doi.org/10.1016/j.aap.2008.01.007

Engström, J., & Hollnagel, E. (2007). A general conceptual framework for modelling behavioural effects of driver support functions. In *Modelling Driver Behaviour in Automotive Environments: Critical Issues in Driver Interactions with Intelligent Transport Systems* (pp. 61–84). http://doi.org/10.1007/978-1-84628-618-6_4

Eskandarian, A. (2012). Introduction to Intelligent Vehicles. In *Handbook of Intelligent*

Vehicles (pp. 1–13). London: Springer London. http://doi.org/10.1007/978-0-85729-085-4_1

Euro NCAP. (2015). European New Car Assessment Programme - Test Protocol - Lane Support Systems v.1.0, (November). Retrieved from http://www.euroncap.com/en/for-engineers/protocols/safety-assist/

Euro NCAP. (2017). European New Car Assessment Programme - Test Protocol - Lane Support Systems v.2.0, (June). Retrieved from http://www.euroncap.com/en/for-engineers/protocols/safety-assist/

Evans, L., & Schwing, R. C. (Eds.). (1985). *Human Behavior and Traffic Safety*. *Human Behavior and Traffic Safety*. Boston, MA: Springer US. http://doi.org/10.1007/978-1-4613-2173-6

FESTA Handbook Version 6. (2016). Retrieved from http://wiki.fot-net.eu/index.php/FESTA_Handbook

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, *11*(2), 127–138. http://doi.org/10.1038/nrn2787

Friston, K. (2011). What is optimal about motor control? *Neuron*, *72*(3), 488–498. http://doi.org/10.1016/j.neuron.2011.10.018

Friston, K. (2013). Life as we know it. *Journal of The Royal Society Interface*, *10*(86), 20130475–20130475. http://doi.org/10.1098/rsif.2013.0475

Friston, K., Adams, R. A., Perrinet, L., & Breakspear, M. (2012). Perceptions as hypotheses: Saccades as experiments. *Frontiers in Psychology*, *3*(MAY), 1–20. http://doi.org/10.3389/fpsyg.2012.00151

Friston, K., Daunizeau, J., Kilner, J., & Kiebel, S. J. (2010). Action and behavior: A free-energy formulation. *Biological Cybernetics*, *102*(3), 227–260. http://doi.org/10.1007/s00422-010-0364-z

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active Inference: A Process Theory. *Neural Computation*, *29*(1), 1–49. http://doi.org/10.1162/NECO_a_00912

Friston, K., Kilner, J., & Harrison, L. (2006). A free energy principle for the brain. *Journal of Physiology Paris*, *100*(1–3), 70–87. http://doi.org/10.1016/j.jphysparis.2006.10.001

Friston, K., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R. A., Brown, H., … Bestmann, S. (2012). Dopamine, affordance and active inference. *PLoS Computational Biology*, *8*(1). http://doi.org/10.1371/journal.pcbi.1002327

Gayko, J. E. (2012). Lane Departure and Lane Keeping. In *Handbook of Intelligent*

Vehicles (pp. 689–708). London: Springer London. http://doi.org/10.1007/978-0-85729-085-4_26

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscienceeuroscience*, *30*, 535–574. http://doi.org/10.1146/annurev.neuro.29.051605.113038

Gordon, T. J., & Magnuski, N. (2006). Modeling Normal Driving as a Collision Avoidance Process. In *Proceedings of the 8th International Symposium on Advanced Vehicle Control*.

Gordon, T. J., Sardar, H., Blower, D., Ljung Aust, M., Bareket, Z., Barnes, M., … Theander, H. (2010). Advanced Crash Avoidance Technologies (ACAT) Program – Final Report of the Volvo-Ford- UMTRI Project : Safety Impact Methodology for Lane Departure Warning – Method Development And Estimation of Benefits, (October), 218.

Hautzinger, H., Pfeiffer, M., & Schmidt, J. (2004). Expansion of GIDAS Sample Data to the Regional Level: Statistical Methodology and Practical Experiences. In *1st International Conference on ESAR "Expert Symposium on Accident Research* (pp. 38–43).

Jakobsson, L., Lindman, M., Axelson, A., Lokensgard, B., Petersson, M., Svanberg, B., & Kovaceva, J. (2014). Addressing Run Off Road Safety. *SAE International Journal of Passenger Cars - Mechanical Systems*, *7*(1), 2014-01–0554. http://doi.org/10.4271/2014-01-0554

Kaplan, S., & Prato, C. G. (2013). Cyclist–Motorist Crash Patterns in Denmark: A Latent Class Clustering Approach. *Traffic Injury Prevention*, *14*(7), 725–733. http://doi.org/10.1080/15389588.2012.759654

Kaufman, L., & Rousseeuw, P. J. (2005). *Finding Groups in Data: An Introduction to Cluster Analysis*. *Wiley-Interscience* (Vol. 33). http://doi.org/10.1007/s00134-006-0431-z

Kim, K., & Yamashita, E. Y. (2007). Using a k-means clustering algorithm to examine patterns of pedestrian involved crashes in Honolulu, Hawaii. *Journal of Advanced Transportation*, *41*(1), 69–89. http://doi.org/10.1002/atr.5670410106

Kuehn, M., Hummel, T., & Bende, J. (2009). Benefit Estimation of Advanced Driver Assistance Systems for Cars Derived from real-live accidents. *Proceedings of the 21st International Technical Conference of Th Enhanced Safety of Vehicles Conference (EVS), Stuttgart, Germany, June 15-18*, 1–10.

Kusano, K. D., & Gabler, H. C. (2013). Pre-crash scenarios for determining target populations of active safety systems. *Proceedings of the 23rd International Technical Conference on the Enhanced Safety of Vehicles (ESV)*. Retrieved from http://wbldb.lievers.net/10092130.html

Kusano, K. D., & Gabler, H. C. (2014). Comprehensive Target Populations for Current Active Safety Systems Using National Crash Databases. *Traffic Injury Prevention*, *15*(7), pp 753-761. http://doi.org/10.1080/15389588.2013.871003

Kusano, K. D., Gabler, H. C., & Gorman, T. I. (2014). Fleetwide Safety Benefits of Production Forward Collision and Lane Departure Warning Systems. *SAE International Journal of Passenger Cars - Mechanical Systems*, *7*(2), 2014-01–0166. http://doi.org/10.4271/2014-01-0166

Land, M., & Lee, D. N. (1994). Where we look when we steer. *Nature*, *369*(6483), 742–744. http://doi.org/10.1038/369742a0

Lappi, O., Lehtonen, E., Pekkanen, J., & Itkonen, T. H. (2013). Beyond the tangent point: gaze targets in naturalistic driving. *Journal of Vision*, *13*(2013), 11. http://doi.org/10.1167/13.13.11

Lappi, O., Rinkkala, P., & Pekkanen, J. (2017). Systematic Observation of an Expert Driver's Gaze Strategy—An On-Road Case Study. *Frontiers in Psychology*, *8*(April), 1–15. http://doi.org/10.3389/fpsyg.2017.00620

Larsen, L. (2004). Methods of multidisciplinary in-depth analyses of road traffic accidents. *Journal of Hazardous Materials*, *111*(1–3), 115–122. http://doi.org/10.1016/j.jhazmat.2004.02.019

Lenard, J., Badea-Romero, A., & Danton, R. (2014). Typical pedestrian accident scenarios for the development of autonomous emergency braking test protocols. *Accident Analysis and Prevention*, *73*, 73–80. http://doi.org/10.1016/j.aap.2014.08.012

Liu, C., & Ye, T. J. (2011). *Run-Off-Road Crashes : An On-Scene Perspective* (No. DOT HS 811 500).

Ljung Aust, M., & Engström, J. (2011). A conceptual framework for requirement specification and evaluation of active safety functions. *Theoretical Issues in Ergonomics Science*, *12*(1), 44–65. http://doi.org/10.1080/14639220903470213

Ljung Aust, M., Habibovic, A., Tivesten, E., Sander, U., Bärgman, J., & Engström, J. (2010). Manual for DREAM version 3.2. Retrieved November 1, 2017, from http://publications.lib.chalmers.se/records/fulltext/204828/204828.pdf

Lundquist, C., Schön, T. B., & Gustafsson, F. (2012). Situational Awareness and Road Prediction for Trajectory Control Applications. In *Handbook of Intelligent Vehicles* (Vol. 1–2, pp. 365–396). London: Springer London. http://doi.org/10.1007/978-0-85729-085-4_15

Markkula, G. (2015). *Driver behavior models for evaluating automotive active safety From* *neural dynamics to vehicle dynamics*. Chalmers University of Technology, Gothenburg, Sweden.

Markkula, G., Benderius, O., Wolff, K., & Wahde, M. (2012). A Review of Near-Collision Driver Behavior Models. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *54*(6), 1117–1143. http://doi.org/10.1177/0018720812448474

Markkula, G., Boer, E., Romano, R., & Merat, N. (2017). Sustained sensorimotor control as intermittent decisions about prediction errors: Computational framework and application to ground vehicle steering. Retrieved from http://arxiv.org/abs/1703.03030

Markkula, G., Engström, J., Lodin, J., Bärgman, J., & Victor, T. (2016). A farewell to brake reaction times? Kinematics-dependent brake response in naturalistic rear-end emergencies. *Accident Analysis & Prevention*, *95*, 209–226. http://doi.org/10.1016/j.aap.2016.07.007

Martínez-García, M., Zhang, Y., & Gordon, T. (2016). Modeling Lane Keeping by a Hybrid Open-Closed-Loop Pulse Control Scheme. *IEEE Transactions on Industrial Informatics*, *12*(6), 2256–2265. http://doi.org/10.1109/TII.2016.2619064

McLaughlin, S. B., Hankey, J. M., & Dingus, T. A. (2008). A method for evaluating collision avoidance systems using naturalistic driving data. *Accident Analysis and Prevention*, *40*(1), 8–16. http://doi.org/10.1016/j.aap.2007.03.016

Michon, J. A. (1985). A critical view of driver behavior models: What do we know, what should we do? *Human Behavior and Traffic Safety*, (New York: Plenum Press), 485–520. http://doi.org/10.1007/978-1-4613-2173-6_19

Milligan, G. W., & Cooper, M. C. (1987). Methodology Review: Clustering Methods. *Applied Psychological Measurement*, *11*(4), 329–354. http://doi.org/10.1177/014662168701100401

Mullen, N., Charlton, J., Devlin, A., & Bedard, M. (2011). Simulator Validity: Behaviros Observed on the Simualtor and on the Road. In *Handbook of Driving Simulation for Engineering, Medicine, and Psychology* (pp. 13-1-13–18).

Mörbe, M. (2016). Vehicle Dynamics Sensors for DAS. In *Handbook of Driver Assistance Systems* (pp. 279–302). Cham: Springer International Publishing. http://doi.org/10.1007/978-3-319-12352-3_15

Najm, W. G., Koopmann, J., Boyle, L., & Smith, D. L. (2002). *Development of Test Scenarios for Off-Roadway Crash Countermeasures Based on Crash Statistics* (No. DOT HS 809 505). Cambridge.

Najm, W. G., & Smith, D. L. (2004). Modeling driver response to lead vehicle decelerating. *SAE 2004 World Congress, Paper*, No. 2004-01, 0171. http://doi.org/10.4271/2004-01-0171

Najm, W. G., & Smith, J. D. (2007). *Development of Crash Imminent Test Scenarios for Integrated Vehicle-Based Safety Systems* (No. DOT HS 810 757). Retrieved from http://www-nrd.nhtsa.dot.gov/pdf/nrd-12/DOT-HS-810757_IVBSS.pdf

Najm, W. G., Smith, J. D., & Yanagisawa, M. (2007). *Pre-Crash Scenario Typology for Crash Avoidance Research* (No. DOT HS 810 767).

Navarro, J., Mars, F., & Young, M. S. (2011). Lateral control assistance in car driving: classification, review and future prospects. *IET Intelligent Transport Systems*, *5*(3), 207. http://doi.org/10.1049/iet-its.2010.0087

Nowakowska, M. (2012). Road Traffic Accident Patterns: A Conceptual Grouping Approach to Evaluate Crash Clusters. *Archives of Transport*, *24*(1), 73–98. http://doi.org/10.2478/v10174-012-0006-4

Otte, D., Pund, B., & Jänsch, M. (2009). A New Approach of Accident Causation Analysis by Seven Steps. In *The 21th International Technical Conference on the Enhanced Safety of Vehicles (ESV)*.

Page, Y., Fahrenkrog, F., Fiorentino, A., Gwehenberger, J., Helmer, T., Lindman, M., … Wimmer, P. (2015). A Comprehensive and Harmonized Method for Assessing the Effectiveness of Advance Driver Assistance Systems by Virtual Simulation. In *The 24th International Technical Conference on the Enhanced Safety of Vehicles (ESV)*. Retrieved from http://www-esv.nhtsa.dot.gov/Proceedings/24/isv7/main.htm

Petermeijer, S. M., Abbink, D. A., & de Winter, J. C. F. (2014). Should Drivers Be Operating Within an Automation-Free Bandwidth? Evaluating Haptic Steering Support Systems With Different Levels of Authority. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *57*(1), 5–20. http://doi.org/10.1177/0018720814563602

Pezzulo, G. (2014). Why do you fear the bogeyman? An embodied predictive coding model of perceptual inference. *Cognitive, Affective & Behavioral Neuroscience*, *14*, 902–911. http://doi.org/10.3758/s13415-013-0227-x

Pomerleau, D., Jochem, T., Batavia, P., Pape, D., Hadden, J., McMillan, N., & Brown, N. (1999). Run-off-road collision avoidance using IVHS countermeasures, (December), 134. Retrieved from http://ntl.bts.gov/lib/jpodocs/repts_te/13342.pdf

Punj, G., & Stewart, D. W. (1983). Cluster Analysis in Marketing Research: Review and Suggestions for Application. *Journal of Marketing Research*, *20*(2), 134–148. http://doi.org/10.2307/3151680

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87. http://doi.org/10.1038/4580

Salvucci, D. D., & Gray, R. (2004). A two-point visual control model of steering. *Perception*, *33*(10), 1233–1248. http://doi.org/10.1068/p5343

Sasidharan, L., Wu, K.-F., & Menendez, M. (2015). Exploring the application of latent class cluster analysis for investigating pedestrian crash injury severities in Switzerland. *Accident Analysis & Prevention*, *85*, 219–228. http://doi.org/10.1016/j.aap.2015.09.020

Scanlon, J. M., Kusano, K. D., Sherony, R., & Gabler, H. C. (2015). Potential Safety Benefits of Lane Departure Warning and Prevention Systems in the U.S. Vehicle Fleet. *Proceedings of the 24th Enhanced Safety of Vehicles Conference*, 15–0080.

Stanley, L. M. (2006). Haptic and auditory cues for lane departure warnings. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *50*(22), 2405–2408. http://doi.org/10.1177/154193120605002212

Strandroth, J. (2015). Validation of a method to evaluate future impact of road safety interventions, a comparison between fatal passenger car crashes in Sweden 2000 and 2010. *Accident Analysis and Prevention*, *76*, 133–140. http://doi.org/10.1016/j.aap.2015.01.001

Summala, H. (2007). Towards understanding motivational and emotional factors in driver behaviour: Comfort through satisficing. In *Modelling Driver Behaviour in Automotive Environments: Critical Issues in Driver Interactions with Intelligent Transport Systems* (pp. 189–207). http://doi.org/10.1007/978-1-84628-618-6_11

Summala, H., Nieminen, T., & Punto, M. (1996). Maintaining lane position with peripheral vision during in-vehicle tasks. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *38*(3), 442–451. http://doi.org/10.1518/001872096778701944

Suzuki, K., & Jansson, H. (2003). An analysis of driver's steering behaviour during auditory or haptic warnings for the designing of lane departure warning system. *JSAE Review*, *24*(1), 65–70. http://doi.org/10.1016/S0389-4304(02)00247-3

Svärd, M., Markkula, G., Engström, J., Bärgman, J., & Granum, F. (2017). A quantitative driver model of pre-crash brake onset and control. In *Human Factors and Ergonomics Society Annual Meeting Proceedings*.

Theofilatos, A., & Efthymiou, D. (2012). Investigation of Pedestrians' Accident Patterns in Greater Athens Area. *Procedia - Social and Behavioral Sciences*, *48*, 1897–1906. http://doi.org/10.1016/j.sbspro.2012.06.1164

Tivesten, E., & Dozza, M. (2014). Driving context and visual-manual phone tasks influence glance behavior in naturalistic driving. *Transportation Research Part F: Traffic Psychology and Behaviour*, *26*(PA), 258–272. http://doi.org/10.1016/j.trf.2014.08.004

Trappenberg, T., & Hollensen, P. (2013). Sparse coding and challenges for Bayesian models of the brain. *Behavioral and Brain Sciences*, *36*(3), 232–233. http://doi.org/10.1017/S0140525X12002300

Victor, T., Dozza, M., Bärgman, J., Boda, C.-N., Engström, J., Flannagan, C., … Markkula, G. (2015). *Analysis of Naturalistic Driving Study Data: Safer Glances, Driver Inattention, and Crash Risk*. http://doi.org/10.17226/22297

Weir, D. H., & McRuer, D. T. (1970). Dynamics of driver vehicle steering control. *Automatica*, *6*(1), 87–98. http://doi.org/10.1016/0005-1098(70)90077-4

Weiss, H. B., Kaplan, S., & Prato, C. G. (2016). Fatal and serious road crashes involving young New Zealand drivers: a latent class clustering approach. *International Journal of Injury Control and Safety Promotion*, *23*(4), 427–443. http://doi.org/10.1080/17457300.2015.1056807

Wilkie, R. M., & Wann, J. P. (2003). Controlling steering and judging heading: retinal flow, visual direction, and extraretinal information. *Journal of Experimental Psychology. Human Perception and Performance*, *29*(2), 363–378. http://doi.org/10.1037/0096-1523.29.2.363

Winner, H. (2016). *Handbook of Driver Assistance Systems*. http://doi.org/10.1007/978-3-319-12352-3