CHALMERS
UNIVERSITY OF TECHNOLOGY

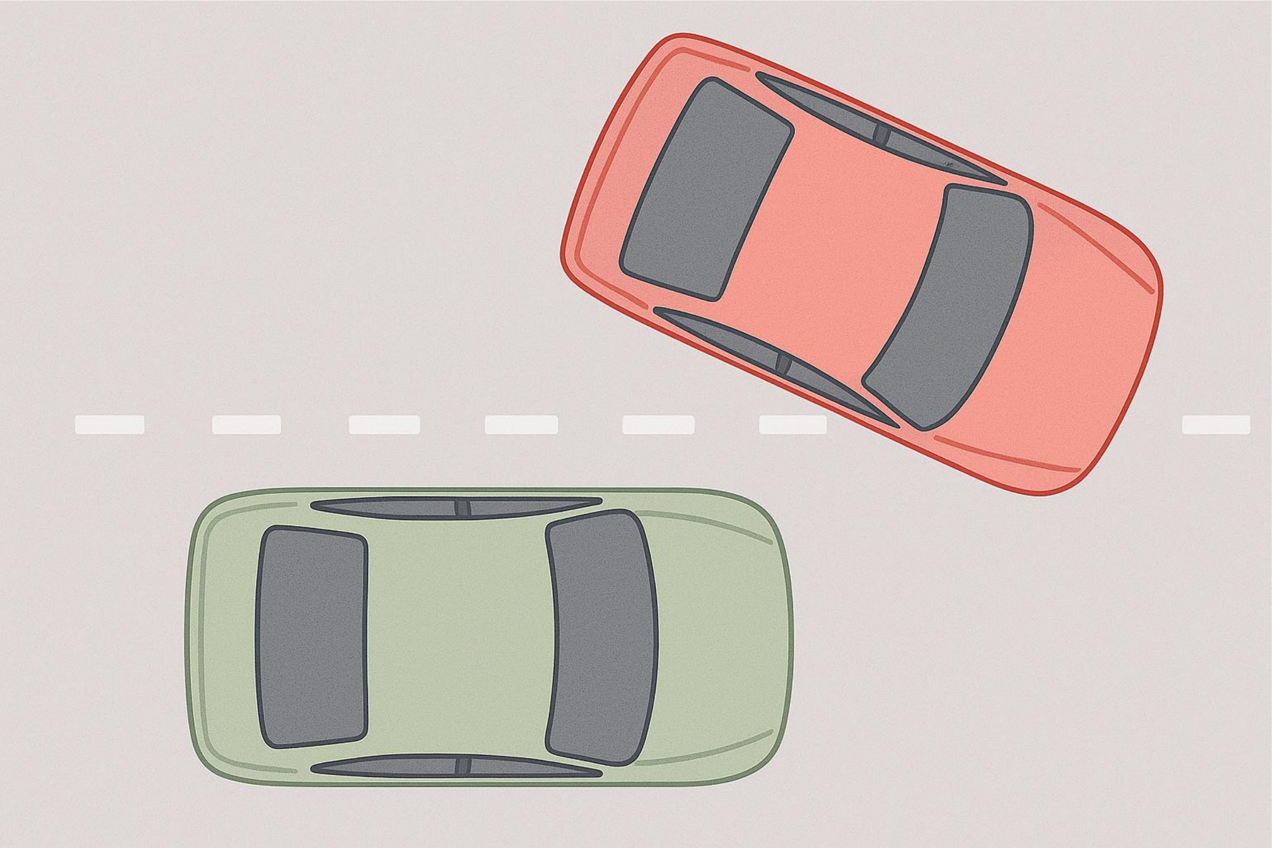

Assessing the crash avoidance potential of cut-in crashes

Master's thesis in Complex Adaptive Systems

Jonatan Solin

Department of Mechanics and Maritime Sciences

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2025
www.chalmers.se

MASTER’S THESIS IN COMPLEX ADAPTIVE SYSTEMS

Assessing the crash avoidance potential of cut-in crashes

JONATAN SOLIN

Department of Mechanics and Maritime Sciences
Division of Vehicle Safety
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2025

Assessing the crash avoidance potential of cut-in crashes
JONATAN SOLIN

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

Cover:
Top-down illustration of a cut-in scenario, where one vehicle merges into the lane of another. Image generated with assistance from ChatGPT.

Printed by Chalmers Digitaltryck
Department of Mechanics and Maritime Sciences
Göteborg, Sweden 2025-06-30

Assessing the crash avoidance potential of cut-in crashes
Master's thesis in Complex Adaptive Systems
JONATAN SOLIN
Department of Mechanics and Maritime Sciences
Division of Vehicle Safety
Chalmers University of Technology

Abstract

Road traffic crashes remain a major safety concern. To reduce their occurrence, automated driving functions (ADFs) are being developed—but these systems must be thoroughly validated. Virtual simulations have been used to assess their performance, often comparing results between different systems without establishing a benchmark for what is theoretically achievable. This study focuses on cut-in crashes, a relatively common and challenging scenario for automated driving systems. It aims to assess the proportion of such crashes that could theoretically be avoided through braking alone and to compare this benchmark with the performance of two reference driver models and an ADF. This was achieved by developing an idealized model that reacts earlier and brakes harder than realistically possible, ensuring that no other model should be capable of outperforming it. This ideal model as well as the reference driver models and the ADF were then applied in virtual counterfactual simulations to estimate the proportion of crashes they could avoid. The cut-in scenarios simulated were categorized as either frontal or non-frontal cut-ins. The study was able to establish an upper limit for the non-frontal cut-in crashes but not for the frontal ones, as the ideal model avoided all frontal collisions. The two reference driver models avoided 38.5% and 81.8% of the frontal crashes, respectively, illustrating a large discrepancy. It remains unclear whether this reflects differences in the modeled driver's behavior or limitations in how well the models represent real human drivers.

Key words: counterfactual simulations, cut-in, crashes, vehicle safety

Contents

Acknowledgement

I would like to thank my supervisor, Jordanka Kovaceva, and my examiner, Jonas Bärgman, for their guidance and support throughout this thesis. I am also grateful to the Hi-Drive project for granting access to the ADF model prior to its publication, and to Volvo Cars for allowing the use of data from the VCTAD dataset in this study.

Large language models (LLMs) were used to assist with improving the grammar, formality, and flow of the report text. They were also used to help generate code for some of the figures included in the report. All research, analysis, and original writing content are my own.

Göteborg June 2025-06-30

JONATAN SOLIN

Abbreviations

Abbreviation	Definition
ADF	Autonomous driving function
ADAS	Advanced driver-assistance systems
AEB	Automatic emergency brakes
ALKS	Automated lane-keeping system
CCDM	Competent and Careful Driver Model
CI	Confidence interval
FSM	Fuzzy Safety Model
ODD	Operational design domain
TTC	Time to collision
UNECE	United Nations Economic Commission for Europe
VCTAD	Volco Cars's traffic accident database

1 Introduction

Traffic accidents result in the death of 1.19 million people worldwide every year and are the leading cause of death globally for individuals aged 5-29 years (World Health, 2023). In recent years, considerable progress has been made in the development of advanced driver assistance systems (ADAS), aimed at reducing the number of traffic fatalities and serious injuries. These systems, such as automated emergency braking (AEB), can successfully prevent many types of crashes.

However, not all situations are equally manageable for automated systems. Sudden maneuvers, such as abrupt lane changes, still pose challenges due to the limited time available to react. One scenario that has proven particularly challenging for ADAS is the highway cut-in — a lane change maneuver discussed further in Section 1.1. Crashes in such scenarios are relatively difficult for ADAS to avoid (Li et al., 2024), largely due to the short time window available to detect and respond to other vehicles cutting in.

1.1 Cut-in crashes

In the context of traffic, a cut-in is a common maneuver where one vehicle changes lane to end up in front of another vehicle. This occurs frequently on highways when faster vehicles overtake slower ones. Most of the time, this maneuver is unproblematic, but if performed without sufficient attention to surrounding traffic, it can lead to accidents.

Studies have suggested that between 4% (U.S.) and 5% (Germany) of all accidents in traffic occur during lane changes (Chovan et al., 1994; Habenicht et al., 2011) and 7% of all fatalities on the German Autobahn occur during such maneuvers (Rodemerk et al., 2012). A naturalistic driving study conducted in Washington, U.S., identified 'Lane Change Without Sufficient Gap' as the most frequent incident type when heavy vehicle drivers were at fault in heavy–light vehicle interactions, accounting for 26.6% of such cases. It was also the second most frequent incident type for light vehicle drivers at fault in the same type of interaction (Hanowski et al., 2006).

In this work, the term ‘cut-in’ refers to any two-vehicle collision that occurred during a lane change, based on the criteria used to identify cut-ins when selecting scenarios from Volvo Cars’ Traffic Accident Database (VCTAD) (Isaksson-Hellman & Norin, 2005) for the L3Pilot project (Bjorvatn et al., 2021). See Section 3.2 for more details on the data selection.

1.2 Virtual counterfactual simulations

Counterfactual simulations are simulations where a scenario from reality is reconstructed in a controlled environment and then replayed with some modification, often called the treatment, to try to see if that modification changes the outcome of the scenario (ISO, 2021, 2024). The simulated scenario with the treatment can then be compared to the untreated scenario, known as the baseline (ISO, 2021, 2024).

In the context of traffic, these scenarios are often crashes or near-crashes and the treatment comes in terms of some behavior model or safety system, for example automated emergency braking (AEB), being added to at least one of the vehicles.

Counterfactual traffic simulations are typically conducted in virtual platforms, where driver models can be applied to vehicles to modify their behavior through changes in acceleration and steering (Alvarez et al., 2017; Page et al., 2015; Wimmer et al., 2023; Wimmer et al., 2019). If only a single vehicle has modifications in its behavior, that vehicle is typically called the ego vehicle. In this study, the term ego vehicle will be used to denote the vehicle following its lane in cut-in scenarios, while the vehicle doing the lane change will be referred to as the cut-in vehicle.

Several previous studies have been conducted where virtual counterfactual simulations have been used (Bärgman et al., 2015; Kovaceva et al., 2022; Sander, 2018). One example of particular interest for this work is Olleja et al. (2025), where near-crash scenarios were reconstructed and then used to assess the performance of reference driver models. However, to the author's knowledge, no previous study has used counterfactual simulation on reconstructed crashes either to assess the performance of reference driver models or to estimate the proportion of cut-in crashes that are theoretically avoidable.

1.3 Reference driver models

A regulation from United Nations Economic Commission for Europe (UNECE) declares two different driver models as reference driver models (UNECE, 2023) that can be used to determine whether a crash is preventable or not by human standards. These two models are the Competent and Careful Driver Model (CCDM) and the Fuzzy Safety Model (FSM).

1.3.1 Competent and Careful Driver Model

The Competent and Careful Driver Model (CCDM) (JAMA, 2022; UNECE, 2023) is a driver model that can be used as a reference model for human drivers.

The way it looks for cut-ins is by defining a lateral wandering zone in surrounding lanes that vehicles can move within. The width of this wandering zone is defined as 0.375 m, meaning that if a vehicle inside this wandering zone moves more than 0.375 m away from its center, it is considered to have left its wandering zone. If a vehicle goes outside of this zone, it is considered a cut-in vehicle, and a safety maneuver will start.

This maneuver starts with some perception time, to model how the driver needs some time to detect and perceive a vehicle starting a cut-in (Green, 2000). The perception time used is 0.4 seconds. After the perception time, the braking is delayed by another 0.75 seconds to account for decision making time, accelerator release time, foot transfer time and time before braking. After this delay, braking starts. The braking linearly increases deceleration from 0 to $0.774g$ in 0.6 seconds, which implies a constant jerk of $1.29\frac{\mathrm{g}}{\mathrm{s}}$. If at any point during this maneuver the time to collision (TTC) goes below a threshold, the AEB will trigger and increase the jerk to $1.42\frac{g}{s}$ and a maximum deceleration of $0.85g$. These values would mean maximum deceleration is reached in 0.6 seconds if starting from 0 deceleration, however it will often be the case that deceleration has already started when the AEB is triggered resulting in a shorter time needed to reach maximum deceleration. A graph showing this can be seen in Figure 1.

An open source implementation of this model is available in Esmini (Esmini, 2024). This implementation requires the cut-in vehicle to be longitudinally ahead of the ego vehicle for any reaction to occur.

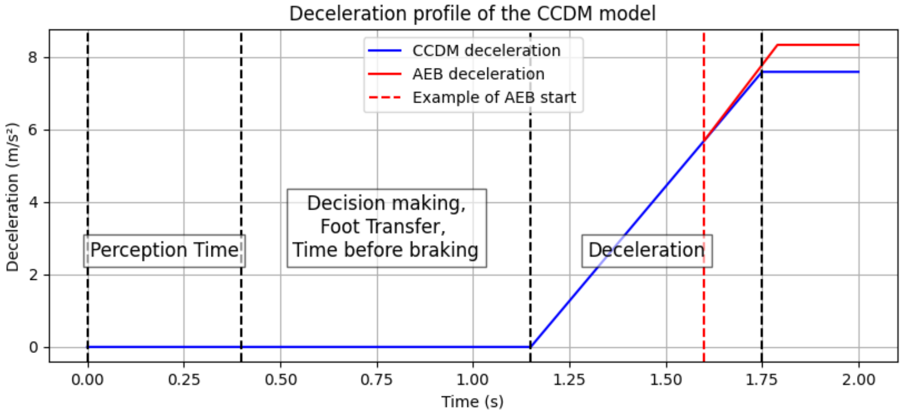

Figure 1: The deceleration for the CCDM model is shown. Time 0.0 here represents the point in time where the cut-in vehicle left its wandering zone. The time point where AEB activates here is an example and can vary depending on the scenario.

1.3.2 Fuzzy Safety Model

The second reference driver model used in this study is the Fuzzy Safety Model (FSM) (Mattas et al., 2022; Mattas et al., 2020; UNECE, 2023). While the CCDM only has two different levels of deceleration, either with AEB or without, the FSM attempts to judge the criticality of the situation and decelerate accordingly. For cut-ins, the model does a lateral check to determine whether a lateral threat is apparent or not and then, if such a threat exists, it determines how much to decelerate by looking at longitudinal dynamics.

The lateral check has four criteria that all need to be satisfied for braking to be considered.

- The cut-in vehicle must be ahead longitudinally of the ego vehicle.
- The cut-in vehicle is moving laterally towards the ego vehicle.
- The cut-in vehicle must be moving slower longitudinally than the ego vehicle.
- If velocities are maintained, the vehicles will overlap laterally and longitudinally simultaneously.

If all these criteria are fulfilled, calculations are made on the longitudinal dynamics of the vehicles.

The longitudinal check determines the criticality of the situation based on the distance to the cut-in vehicle and both vehicles' velocities by using two fuzzy surrogate safety metrics. The details on how this is done can be read in the UNECE regulation and the paper where is was proposed (Mattas et al., 2022;

UNECE, 2023). This model uses a reaction time for the driver of 0.75 seconds and uses a maximum deceleration of $0.6\,\frac{\mathrm{m}}{\mathrm{s}^2}$ $(0.061g)$ with a jerk of $12.65\,\frac{\mathrm{m}}{\mathrm{s}^3}$ $\left(1.29\,\frac{\mathrm{g}}{\mathrm{s}}\right)$.

1.4 Hi-Drive’s autonomous driving functions

Hi-Drive (Hi-Drive, 2020) is an European Commission-funded project involving numerous partners across Europe. It aims to extensively test automated vehicles and ADAS in a wide range of traffic scenarios to evaluate their impact on safety, efficiency, the environment, mobility, transport systems, and society (Sintonen et al., 2023).

As part of the project, driver models for autonomous driving functions (ADFs) were developed (Vater et al., 2023). An ADF is a functionality of an automated driving system (ADS) that controls the vehicle's movement and behavior within a specific operational design domain (ODD). Hi-Drive's ADF models have been used in various types of virtual safety assessments, including both counterfactual simulations and scenario-based traffic simulations (Hi-Drive, 2025), to determine the effects of autonomous driving on driving behavior (Vater et al., 2023).

One of Hi-Drive's ADFs, the rural ADF, is of particular interest for this work as cut-ins on rural roads are part of its ODD. In the remainder of this thesis, this model is referred to as Hi-Drive’s ADF. As the Hi-Drive project deliverable describing this model has not been published yet (the intent is, as far as the author of this work has understood, to make the ADF code open-source), it will not be described in detail in this work. The maximum deceleration of the model is higher than that of the FSM and CCDM-models and just like those models, it only considers vehicles that are longitudinally in front when looking in other lanes for potential cut-ins.

2 Aim and scope

The aim of this thesis is to assess the crash avoidance potential of ego vehicle braking maneuvers in cut-in scenarios. To guide this assessment, the work addresses two research questions.

Cut-in crashes are often difficult to avoid, especially if the vehicle cutting-in is longitudinally close to, or even overlapping with, the ego vehicle. This led to the formulation of the first research question: "What proportion of cut-in crashes are avoidable solely by braking, under ideal conditions?"

A second research question was formulated related to the performance of the UNECE reference models and the Hi-Drive's automated driving function in cut-ins: "How do some of the existing reference driver and ADF models compare to each other and to the ideal cut-in avoidance rates?"

One way to avoid accidents is to minimize the risk of being in a dangerous situation to start with. For the case of cut-ins, this could for example include decelerating whenever there is another vehicle beside the ego vehicle. As this type of maneuver is hard to evaluate in counterfactual simulations (further explained in 3.4) and greatly increases the number of maneuvers that would need to be considered, the scope of this work was set to not include this type of maneuver.

A second restriction in the scope of this thesis work was that turn signals were not considered. The reason for this is that data about turn signals were not available in the data set used.

As a last restriction, it was decided to not consider surrounding traffic as the aim is to avoid the crash with the cut-in vehicle itself. While excessive braking can lead to a crash with a vehicle coming from behind, such crashes were not considered.

3 Method

A key method used in this work has been to perform counterfactual virtual simulations in the environment simulator Esmini (Esmini, 2024). The simulated scenarios were reconstructed crashes from the Volvo Cars Traffic Accident Database (VCTAD) (Bjorvatn et al., 2021). In these reconstructed crashes, the vehicle following its lane was considered the ego vehicle and had various driver models applied to it while the cut-in vehicle was set to just follow its original trajectory.

An ideal driver model was developed in this study to establish an upper limit on the number of crashes that can be avoided within the defined scope. It is considered 'ideal' in the sense that it is designed to outperform what would be feasible in a real-world scenario, by reacting earlier and with stronger braking capabilities than a regular car.

In addition to the model implemented, three models from other studies were used to get a comparison between the ideal avoidances and the avoidances from those models. Two of these models, the FSM (Mattas et al., 2022) and the CCDM (JAMA, 2022), are included in UNECE regulations as reference driver models for the evaluation of automated lane-keeping systems (ALKS). The final model used in this study is an autonomous driving function developed by the Hi-Drive project, used with explicit permission for this study. Details of the model are expected to be published in a forthcoming project deliverable (Hi-Drive, 2025).

An overview of the method is illustrated in Figure 2. The methods used are explained in more detail in this chapter.

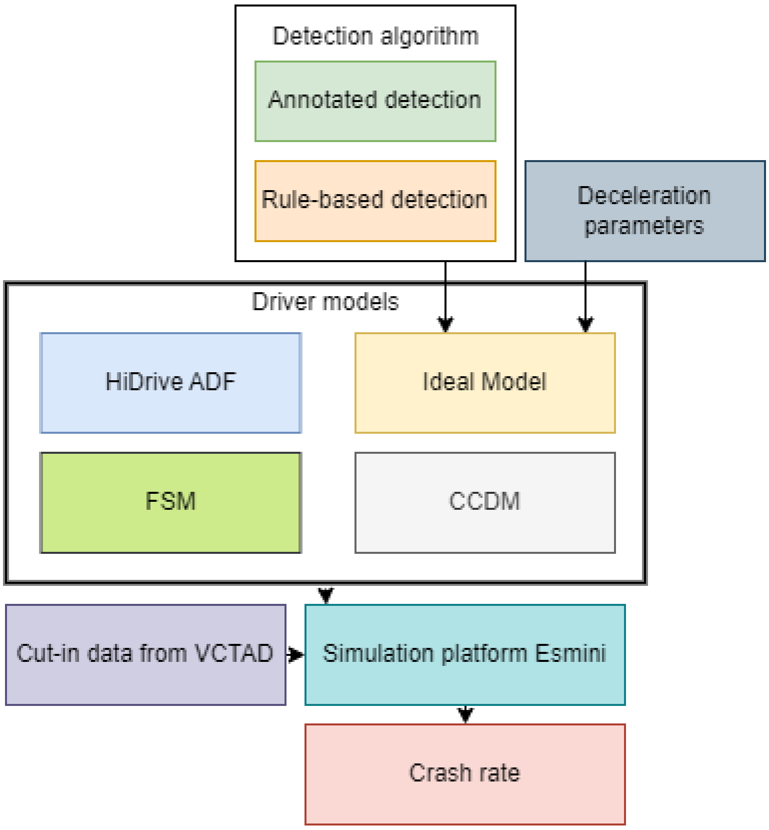

Figure 2: Overview of the simulation methodology used in this study. The flowchart shows how detection algorithms, deceleration parameters, and driver models were used in the Esmini simulation platform to evaluate crash rates using cut-in data from Volvo Cars Traffic Accident Database (VCTAD).

3.1 Virtual simulations in Esmini

To determine what proportion of cut-in crashes that are avoidable under ideal conditions, it was decided to run virtual simulations in the environment simulator Esmini. Esmini was selected for the virtual simulations mainly because of our prior experience with the platform. Esmini is an open-source virtual environment simulator that plays OpenScenario XML files to simulate traffic (Esmini, 2024). One of its features is the ability to add controllers to vehicles, allowing them to be controlled based on various human or technology driver models.

3.2 Reconstructed crashes

The simulated scenarios were reconstructed cut-in crashes obtained from the Volvo Cars Traffic Accident Database (VCTAD). A total of 51 crashes, previously selected as cut-in scenarios for the L3Pilot project (Bjorvatn et al., 2021), were included in this study. Out of these 51 crashes, three were excluded as they did not sufficiently resemble cut-in crashes.

The first removed crash involved the ego vehicle being stationary when the cut-in vehicle drove into it. This scenario was deemed not representative enough of typical cut-in crashes to be included.

The second excluded case involved the cut-in vehicle approaching from an orthogonal direction relative to the ego vehicle, resembling a crossing scenario rather than a cut-in.

The final excluded data point was a crash in a slow-moving traffic queue in which a cut-in occurred. This case was removed because it differed substantially from the other scenarios, and designing a model to avoid it would not be representative of general cut-in crash behavior.

The remaining 48 crashes were categorized as either frontal or non-frontal crashes, based on two criteria. A crash was classified as frontal if it met both of the following conditions:

- The entirety of the cut-in vehicle was ahead of the ego vehicle at the start of the cut-in maneuver.
- The front side of the ego vehicle collided with the cut-in vehicle.

Crashes that met both criteria were generally easier to avoid, and results are presented separately for each category (i.e., when both criteria were met on the one hand, and when at least one criterion was not fulfilled, on the other). Of the 48 analyzed crashes, 13 were classified as frontal and 35 as non-frontal.

3.3 Ideal model

This section explains in detail how the ideal model was used to identify an upper limit for the number of avoidable cut-in crashes.

3.3.1 Overview

To estimate the proportion of cut-in crashes that are theoretically avoidable within the defined scope, a driver model was developed. The model initially follows the trajectory that led to the original crash. Upon detecting a cut-in, it begins decelerating at a high rate. Two detection methods and three deceleration settings were evaluated.

The first detection method was a rule-based algorithm that evaluates lateral velocity, lateral acceleration, lateral distance and longitudinal distance between the cut-in and ego vehicles. When those values matched certain rules (further explained in Section 3.3.2), the maneuver was considered a cut-in and braking would start.

For the second method, each crash was manually reviewed in the simulation platform, and the start time of each cut-in maneuver was annotated. These time stamps were then used as the detection times. The two detection methods were applied separately and are presented independently in the results.

3.3.2 Rule-based detection

The first (i.e., rule-based) method for determining when a cut-in begins continuously evaluates whether another vehicle's movement satisfies predefined conditions.

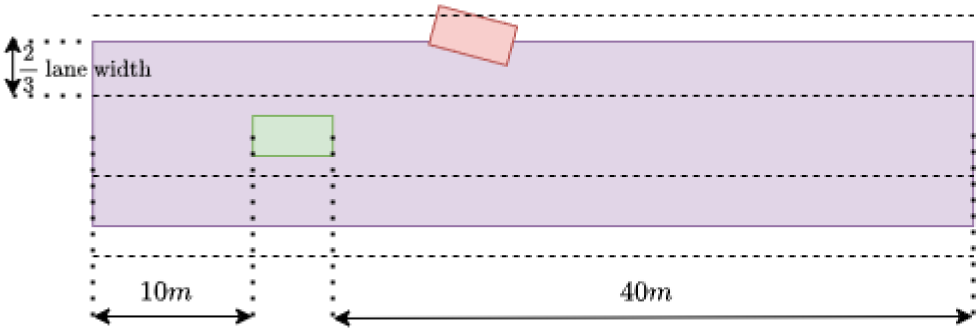

Figure 3 – The ego vehicle (green box) is moving toward the right. Dashed lines indicate lane markings. Any vehicle within the purple area is considered a potential cut-in. It is sufficient for only part of the vehicle to be within the area.

An initial check was performed to determine whether the cut-in vehicle was sufficiently close to the ego vehicle to be considered. The evaluated area extended longitudinally from 10 meters behind to 40 meters in front of the ego vehicle. Laterally, it covered two-thirds of the width of the adjacent lanes. A vehicle was considered as soon as any part of it was within this area. Figure 3 illustrates this region.

The longitudinal threshold was chosen to exclude distant cut-ins that posed no threat to the ego vehicle. These events were typically followed by a subsequent cut-in occurring closer to the ego vehicle. The lateral constraint was intended to prevent vehicles from more distant lanes from being misclassified as cut-ins into the ego vehicle's lane.

If any part of the cut-in vehicle was within the detection area, the check proceeded to a second stage, where four conditions were evaluated:

- $d_{lat} \cdot v_{lat} < 0$
- $|v_{lat}| > 0.2\frac{\mathrm{m}}{\mathrm{s}}$
- $d_{lat} \cdot a_{lat} < 0$
- $|a_{lat}| > 0.5\frac{\mathrm{m}}{\mathrm{s}^2}$

Here, a_{lat} and v_{lat} represent the lateral acceleration and lateral velocity of the cut-in vehicle, respectively. d_{lat} is the lateral distance, measured from the ego vehicle. Figure 4 illustrates how these variables are defined. The first condition requires the lateral velocity of the cut-in vehicle to be directed toward the ego vehicle. This ensures that no reaction is triggered by vehicles moving away from the ego vehicle. A similar criterion was also used in the FSM. The second condition specifies a minimum lateral velocity required for a maneuver to be considered a cut-in. A similar threshold was used in a previous study (Jokhio et al., 2024), although the limit was slightly increased from $0.15\,\frac{\mathrm{m}}{\mathrm{s}^2}$ to $0.2\,\frac{\mathrm{m}}{\mathrm{s}^2}$ in this study. This adjustment was made because the lower threshold often triggered detections in situations where no actual cut-in was occurring in the data used for this study.

The third and fourth conditions were similar in form to the first and second but were based on lateral acceleration rather than velocity. The third condition required the lateral acceleration to be directed toward the ego vehicle. This helped distinguish whether a cut-in vehicle was initiating or completing a lane change. Without this condition, vehicles performing a single lane change from two lanes away could be misclassified as cut-ins.

The fourth condition required the lateral acceleration to exceed $0.5\,\frac{\mathrm{m}}{\mathrm{s}^2}$. This threshold helped distinguish cut-ins from regular driving in curves, as there were several instances where all other conditions were met during curved driving. Including this requirement reduced false detections that occurred before the actual cut-in.

If all these conditions were satisfied, a cut-in was considered detected, and the ego vehicle began decelerating as described in Section 3.3.4.

3.3.3 Detection using manually annotated time stamps

While the rule-based detection algorithm was able to correctly identify most cut-ins shortly after they began, there were cases where it failed to detect them. Since this study aims to estimate an upper limit on avoidable collisions, failing to detect some cut-ins properly is problematic. To address this, an alternative detection method was introduced. This method used manually annotated time stamps indicating when the cut-in began to determine when deceleration should

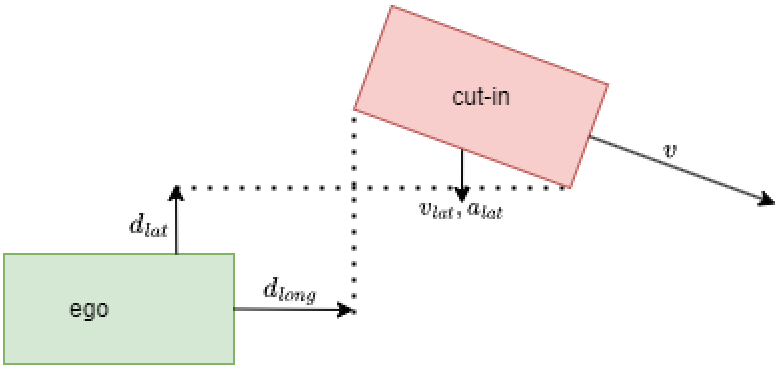

Figure 4 – The directions of variables are shown. Velocity and acceleration are from the cut-in vehicle will distance is measured from the ego vehicle. The lateral distance, velocity and acceleration can all be positive or negative depending on direction. Which direction is positive, and which is negative does not matter for the model.

start. As soon as the simulation time matched the time stamp, braking was initiated, as explained in Section 3.3.4.

To annotate those time stamps for each scenario, simulations were run without any treatment model in place, making both vehicles follow their original trajectories leading to the collision. These simulations were visualized and the trajectories for both vehicles were shown to the author of this work. The simulation was paused frequently when approaching the start of a cut-in to get an accurate time of when the cut-in started.

The annotated-cut-in time stamp was recorded as soon as the cut-in vehicle stopped following its lane and the trajectory showed that a cut-in was imminent. Note however that in some instances the movement was so subtle such that it may not be possible to already at that time judge that this would be a cut-in, it was anyway decided to use the earliest identifiable time, despite hindsight bias. This is further discussed as a limitation in Section 5.5. How the visualization looked when annotating these time stamps can be seen in Figure 5.

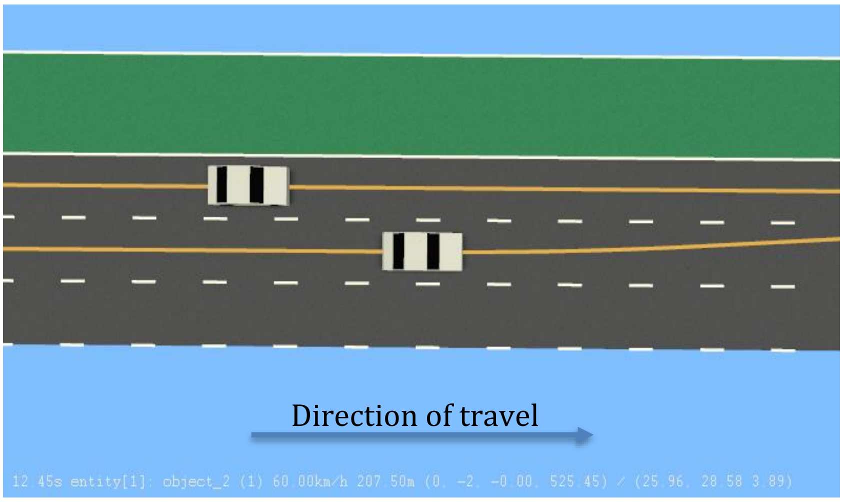

```
12.45s entity[1]: object_2 (1) 60.00km/h 207.50m (0, -2, -0.00, 525.45) / (25.96, 28.58 3.89)
```

Figure 5 – A figure showing how the visualization of the simulations looked like when annotating the start times of the cut-ins. Trajectories were shown and the simulations was paused frequently when approaching the start of the cut-in. The arrow and the "Direction of travel" text were not part of the original visualization but have been added here to clarify the direction of travel.

3.3.4 Model braking

As soon as a cut-in is detected, the ideal model will initiate braking. To get an upper limit on the number of avoidable collisions, the braking must be as strong as realistically possible for a real vehicle. Since AEB systems are designed to apply strong emergency braking in critical situations, the model was designed to have similar braking capabilities as AEB systems.

Many AEB systems work in two stages (Alsuwian et al., 2022), the first one where deceleration is lower which acts as a warning to the driver while lowering the velocity in a comfortable way. If the situation becomes more critical, it goes to a second stage where it maximizes braking.

In a study where various vehicles with AEB systems were tested (Miholjcic et al., 2019) at velocities 20 $\frac{\text{km}}{\text{h}}$ and 40 $\frac{\text{km}}{\text{h}}$, the highest deceleration any model reached was 0.92 g.

As part of a sensitivity analysis, three sets of braking parameters were defined for this study: very strong braking, strong braking, and medium braking.

The very strong braking profile was designed to exceed realistic braking capabilities, ensuring that no practical system would outperform it. In this profile, braking began immediately upon cut-in detection with a constant deceleration of 1 g (9.81 $\frac{\text{m}}{\text{s}^2}$).

The strong braking profile incorporated an initial mild braking phase. It began with 0.3 g deceleration for 0.5 s, after which it ramped up to 1 G. This profile aimed to remain within the upper bounds of realistic braking behavior.

The medium braking profile reduced braking intensity and initial responsiveness. It applied 0.2 g deceleration for 0.75 s, followed by 0.9 g deceleration. This setting was intended to represent a realistic braking strategy.

All braking profiles assumed instantaneous transitions in deceleration, neglecting jerk. The resulting deceleration curves are shown in Figure 6.

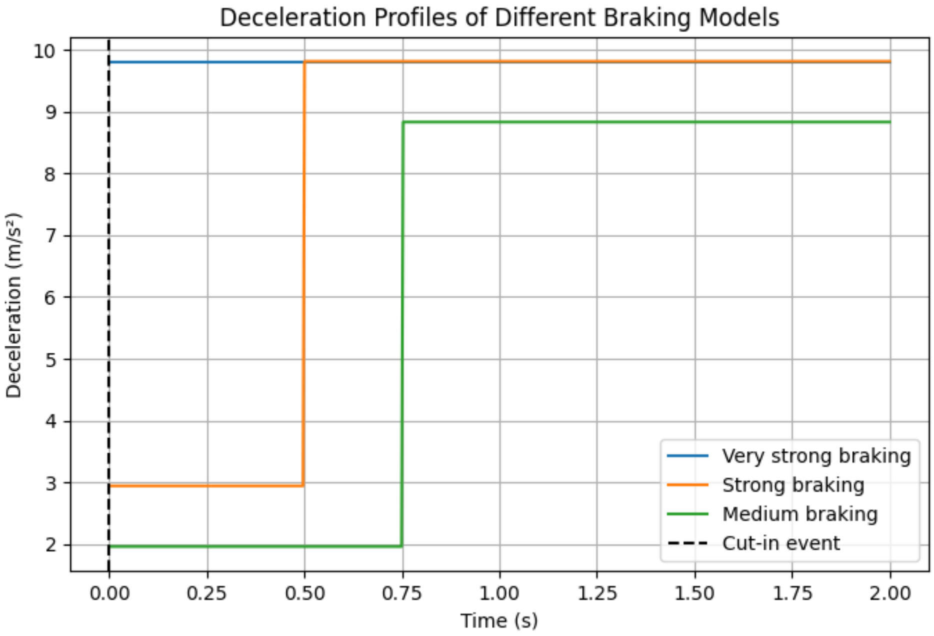

Figure 6: Braking profiles for the different parameter settings. After reaching the higher level of deceleration, the vehicle will continue braking until fully stopping or colliding. Note that the very strong braking is ramping up directly at the cut-in event (i.e., it is hidden behind the dashed line).

3.4 Collision synchronization

In counterfactual simulations, real-world crash scenarios are recreated and replayed in a virtual environment with specific modifications, often referred to as treatments when used to assess the safety impact of safety systems. In this study, the treatment involved adding various controller systems (i.e., human or automated behavior models) to the ego vehicle to evaluate their potential impact on crash avoidance for a specific dataset of cut-in crashes.

Since the outcome of a crash can be highly sensitive to timing and positioning, it is crucial to preserve the key characteristics of the scenario to ensure it remains representative of the original crash event. To achieve this, vehicles in the simulations followed the reconstructed trajectories leading to the original crash, until the ego vehicle's equipped system detected a critical situation and took over longitudinal control.

Since the crashes occurred shortly after a known critical event—the lane change—even minor braking before that point could shift the scenario enough to prevent the collision, without any meaningful reaction to the cut-in. For example, if the ego vehicle reduced its speed by just $1\,\frac{\mathrm{m}}{\mathrm{s}}$ at the start of the scenario and the original crash occurred 10 seconds later, it would be positioned 10 meters further back—potentially avoiding the crash without any actual reaction to the cut-in.

To avoid misclassifying controller reactions that occurred too early to be plausible responses to the cut-in, the concept of early triggers was introduced. These are defined as cases where the controller initiated braking well before the cut-in maneuver began. Since such responses are unlikely to be related to the actual cut-in, they were excluded from the main evaluation and categorized separately.

Manually annotated time stamps from Section 3.3.3 were used as reference points for when each cut-in began. Due to the uncertainty in these annotations, a one-second window prior to each time stamp was permitted for valid responses. If a controller acted more than one second before the annotated cut-in time, the case was classified as an early trigger rather than an avoidance or a collision. In practice, most early triggers occurred several seconds before the cut-in, so the exact choice of threshold had limited impact on the results.

These early triggers could occur for a variety of reasons, most commonly because the system responds to behavior that resembles a cut-in according to the model's logic, even though no actual cut-in is taking place.

3.5 Statistical analysis

To account for the limited amount of data available, 95% confidence intervals were used to estimate the range of plausible avoidance rates. A 95% confidence interval represents the range of values that, under repeated sampling, would be expected to contain the true value in 95% of cases.

Due to the small sample size in this study, the standard Wald confidence interval was deemed unsuitable, as it is known to perform poorly with limited data. Instead, the Wilson score interval was used, as it provides more accurate coverage probabilities, particularly for small samples (Wallis, 2013).

The equation used to compute the confidence intervals was:

$$w^{\pm} = \frac{\hat{p} + \frac{z^2}{2n} \pm z\sqrt{\frac{\hat{p}(1-\hat{p})}{n} + \frac{z^2}{4n^2}}}{1 + \frac{z^2}{n}}$$

Where n is the total number of samples — in this case, the number of avoidances plus the number of collisions. z is the z-score, a value that corresponds to the desired confidence level. For the 95% confidence intervals used in this study, $z \approx 1.96$. $\hat{p}$ is the observed success rate, defined here as the proportion of avoidances out of total events.

4 Results

In this section the results of the simulations will be presented. In Section 4.1, the driver models used in the study will be compared in terms of how well they avoided collisions. Due to major differences in avoidance rates between frontal crashes (as described in Section 3.2) and non-frontal crashes, results will be presented separately for those two as well as together. In Section 4.2, the rule-based detection algorithm will be evaluated by comparing it to the model with annotated time stamps.

4.1 Collision avoidance

To evaluate how well the different models performed, we look at the avoidance rates for the different models for frontal crashes, non-frontal crashes and all crashes. Scenarios where driver models took control of vehicles prior to the cut-in event have been considered as 'early triggers' here and are counted as neither a collision nor an avoidance (described further in Section 3.4). If the ego vehicle and the cut-in vehicle touch at any point during a simulated scenario, that scenario is considered a collision.

4.1.1 Frontal crash scenarios

Table 1 and Figure 7 shows early triggers, avoidance rates and confidence intervals for the frontal crash scenarios for each of the models. As expected, all models showed a higher avoidance rate in this category of scenarios than they did in the non-frontal scenarios. A total of 13 scenarios were played in this category.

There was a single scenario which the rule-based ideal model with medium braking failed to avoid. That scenario ended up in a collision for the CCDM- and FSM-model as well and caused an early trigger by Hi-Drive's ADF-model so neither of the non-ideal models avoided the collision in that scenario either. That collision was avoided by all other settings of the ideal model. The rule-based detection model detected this cut-in 0.19 seconds after the annotated time stamp, which made the difference to cause a collision. A combination of a short distance between the vehicles and a large speed difference when the cut-in started made this scenario particularly difficult to avoid. The baseline collision for this scenario occurred 2.14 seconds after the annotated time stamp.

Table 1: Early triggers, avoidance rates and 95% confidence interval for each model for the frontal crashes.

	Early triggers	Avoidance rate	95% CI
Hi-Drive's autonomous driving function	3	100.0% (10)	72.2%-100.0%
Competent and Careful Driver Model	0	38.5% (13)	17.7%-64.5%
Fuzzy Safety Model	2	81.8% (11)	52.3%-94.9%
Ideal, rule-based time, very strong braking	0	100.0% (13)	77.2%-100.0%
Ideal, rule-based time, strong braking	0	100.0% (13)	77.2%-100.0%
Ideal, rule-based time, medium braking	0	92.3% (13)	66.7%-98.6%
Ideal, annotated time, very strong braking	0	100.0% (13)	77.2%-100.0%
Ideal, annotated time, strong braking	0	100.0% (13)	77.2%-100.0%
Ideal, annotated time, medium braking	0	100.0% (13)	77.2%-100.0%

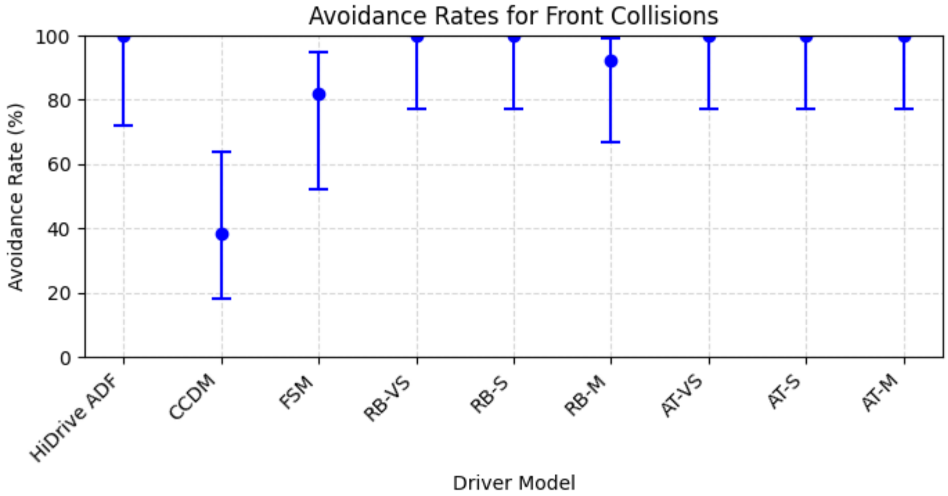

Figure 7: Confidence intervals and mean values for the avoidance rate for each model for the frontal scenarios. RB is the rule-based model, AT is the annotated time model, VS, S and M mean very strong, strong and medium braking, respectively.

4.1.2 Non-frontal crash scenarios

Table 2 and Figure 8 shows early triggers, avoidance rates and 95% confidence intervals for non-frontal crashes. A total of 35 scenarios were played for this category of crashes.

The Hi-Drive's ADF-, CCDM- and FSM-model were all unable to avoid every collision for this category. It is important to note that all the models except the ideal ones only reacted to vehicles ahead of them which made these models unable to avoid most of these collisions by design. The definition used for frontal-crashes in this study did leave some room for scenarios where these models could have acted, but with the limited data used in this study, no collisions in this category were avoided.

Table 2: Early triggers, avoidance rates and 95% confidence interval for each model for the non-frontal crashes. While it may seem wrong that the rule-based algorithm with strong and medium braking have the same number of avoidances, those values have been double-checked.

	Early triggers	Avoidance rate	95% CI
Hi-Drive's autonomous driving function	5	0.0% (30)	0.0%-11.4%
Competent and Careful Driver Model	2	0.0% (33)	0.0%-10.4%
Fuzzy Safety Model	0	0.0% (35)	0.0%-9.9%
Ideal, rule-based, very strong braking	0	60.0% (35)	43.6%-74.4%
Ideal, rule-based, strong braking	0	42.9% (35)	28.0%-59.1%
Ideal, rule-based, medium braking	0	42.9% (35)	28.0%-59.1%
Ideal, annotated time, very strong braking	0	74.3% (35)	57.9%-85.8%
Ideal, annotated time, strong braking	0	62.9% (35)	46.3%-76.8%
Ideal, annotated time, medium braking	0	48.6% (35)	33.0%-64.4%

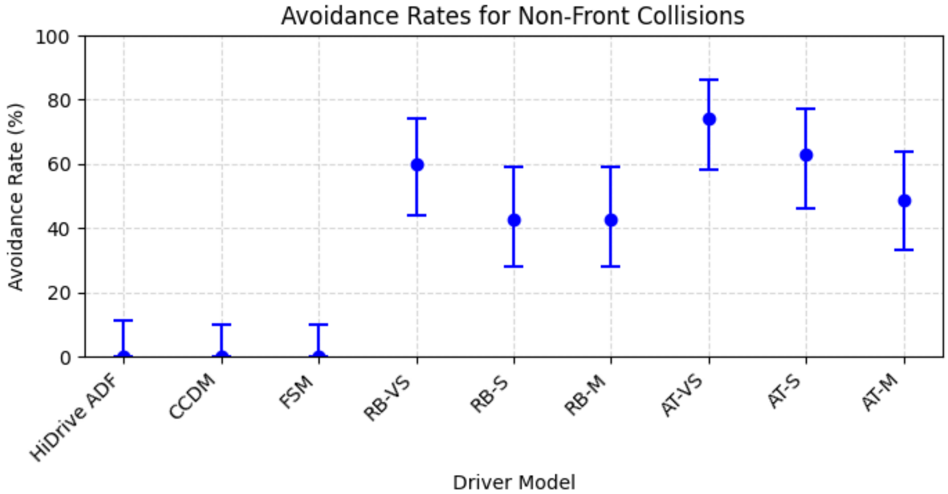

Figure 8: Confidence intervals and mean values for the avoidance rate for each model for the non-frontal scenarios. RB is the rule-based model, AT is the annotated time model, VS, S and M mean very strong, strong and medium braking, respectively.

4.1.3 All scenarios

Last, Table 3 and Figure 9 shows the results for the 13 frontal and 35 non-frontal crashes combined, for a total of 48 scenarios.

As stated in the previous subsection, it is important to note that the models other than the ideal one by design do not react to cut-ins from vehicles that are not in front.

Table 3: Number of early triggers, avoidance rates, and 95% CI for avoidance for each of the different controller systems. This includes all cases (both frontal and non-frontal).

	Early triggers	Avoidance rate	95% CI
Hi-Drive's autonomous driving function	8	25.0% (40)	14.2%-40.2%
Competent and Careful Driver Model	2	10.9% (46)	4.7%-23.0%
Fuzzy Safety Model	2	19.6% (46)	10.7%-33.2%
Ideal, rule-based, very strong braking	0	70.8% (48)	56.8%-81.8%
Ideal, rule-based, strong braking	0	58.3% (48)	44.3%-71.2%
Ideal, rule-based, medium braking	0	56.3% (48)	42.3%-69.3%
Ideal, annotated time, very strong braking	0	81.3% (48)	68.1%-89.8%
Ideal, annotated time, strong braking	0	72.9% (48)	59.0%-83.4%
Ideal, annotated time, medium braking	0	62.3% (48)	48.4%-74.8%

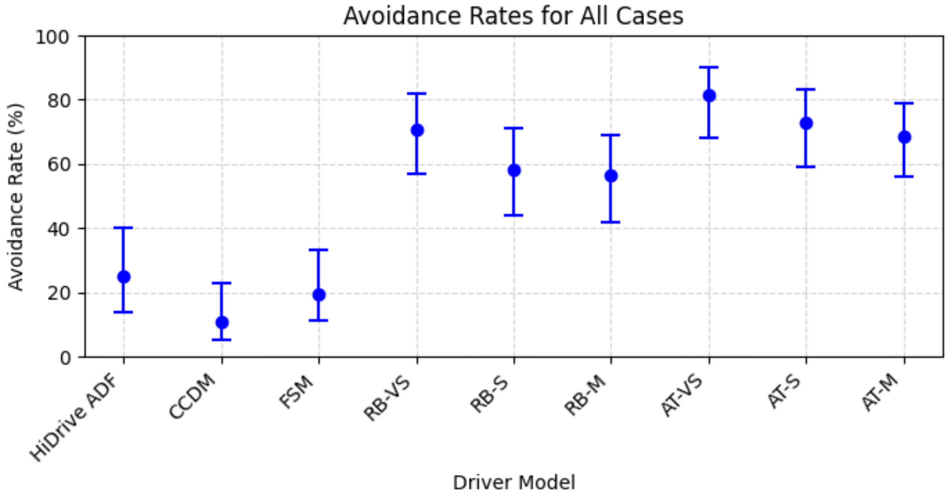

Figure 9: Confidence intervals and mean values for the avoidance rate for each model for the frontal and non-frontal scenarios combined. RB is the rule-based model, AT is the annotated times model, VS, S and M mean very strong, strong and medium braking, respectively.

4.2 Evaluation of the rule-based detection algorithm

To evaluate the performance of the rule-based detection algorithm, Figure 10 shows the difference in detection time between the rule-based model and the annotated time stamps. A positive detection error indicates that the rule-based model reacted after the annotated time stamp. The mean detection delay was 0.247 seconds. In addition to the cases shown in the figure, there were two scenarios in which the rule-based algorithm failed to detect the cut-in entirely. Although this delay is called an "error" here, it should be noted that even an extremely skilled human or the best possible realistic technology-based system would likely in several of these scenarios be unable to with any confidence judge the lateral movement to be a cut-in.

The histogram uses color to indicate which scenarios resulted in additional collisions for the rule-based detection model. Specifically, red bars represent cases where at least one of the braking profiles led to a collision using the rule-based model, while the same braking profile did not result in a collision under the annotated detection model. Blue bars indicate that both models produced the same outcome across all braking profiles. There were no cases where the rule-based model avoided a collision and the model using annotated times collided.

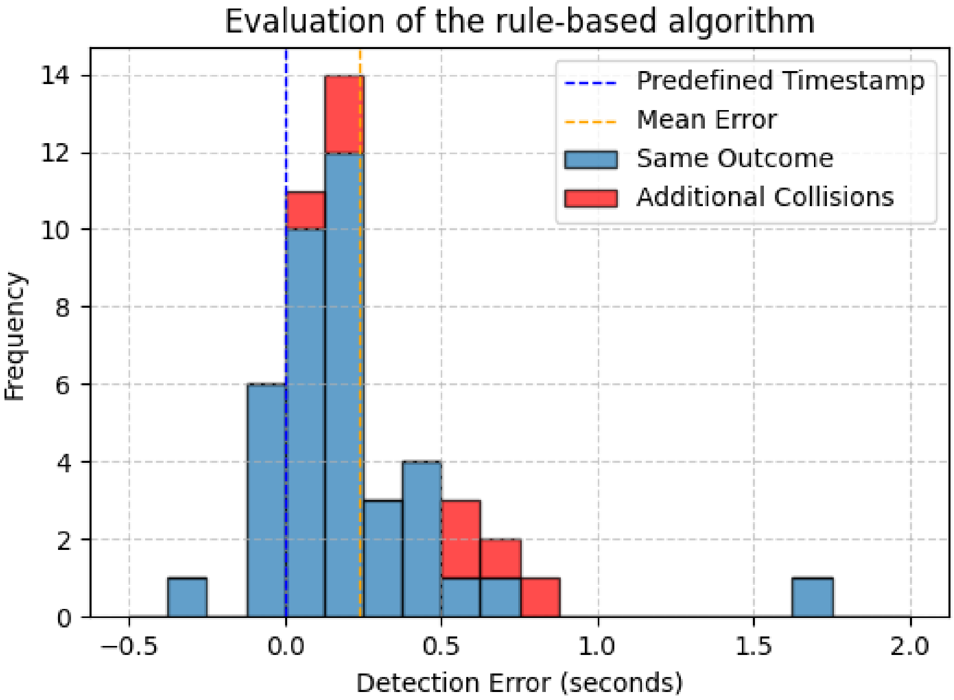

Figure 10: This histogram shows the difference in time of detection between the annotated timestamps and the rule-based detection algorithm's detection time. A positive time difference means that the rule-based algorithm reacted later than the one with annotated times. The color of the bar shows whether the change in reaction time caused any additional collisions for any of the models, that is, if any set of braking parameters led to a collision with the rule-based model but not with the annotated times model. There were no cases where the rule-based model avoided a collision and the model using annotated times collided. The mean error is 0.247 seconds.

5 Discussion

This section discusses the results of the study, along with its limitations and potential directions for future work. The discussion in this section focuses on how the results relate to the aim of the study. Of particular interest is whether the ideal model established an upper limit on the number of avoidable crashes, and how the other models compared to this limit and to each other.

5.1 Frontal crash scenarios

For the frontal crashes in this study, all models except the CCDM avoided most collisions. Since both the reference driver model's from the UNECE regulation (UNECE, 2023), FSM and CCDM, are intended to model human driving behavior, the large discrepancy in their performance is concerning. While human drivers naturally exhibit a wide range of behaviors, a model designed to represent a competent and careful driver should perform closer to the optimal human level. In contrast, the CCDM performs considerably worse than the FSM in this study.

These findings are consistent with those of Olleja et al. (2025), where the models were evaluated in near-crash scenarios. That study found that the CCDM tends to react later than typical human drivers, occasionally turning a near-crash into a crash. The FSM, on the other hand, was overly cautious, responding earlier than most humans. Another study that evaluated the reference driver models (Mattas et al., 2022) also found that the FSM was able to avoid many more collisions than the CCDM. Together, these results raise questions about the validity of current reference driver models and highlight the need for further validation and refinement.

The only case in which the ideal model failed to avoid a frontal collision was when the rule-based detection method was combined with a medium braking level. This highlights that some frontal scenarios are close to the boundary of avoidability and may remain unavoidable depending on the vehicle's braking capabilities. Since all frontal crashes were avoided by at least one model, it is not possible to establish an upper bound for their avoidability. The results are consistent with the possibility that all such crashes could, in theory, be avoided.

Hi-Drive's ADF triggered early in three frontal scenarios, including the one where the ideal model collided, making the outcome in those cases uncertain. However, in all scenarios without early triggers, the ADF avoided the crash. This suggests that even without idealized assumptions, most frontal collisions could still be prevented by a well-performing ADF with perfect sensing.

5.2 Non-frontal crash scenarios

For the non-frontal crash scenarios, only the ideal model was able to avoid collisions. This is largely because the other models by design do not react to vehicles that are not in front of them. The definition of frontal crashes used in this study does leave some room for scenarios where these models could have reacted, but it is expected that those cases are rare and with the small data set used in this study, no cases where the collision was avoided by these models were encountered.

The ideal model's success in avoiding non-frontal crashes varied depending on the braking parameters and the detection method used. Across all combinations, the highest avoidance rate was achieved using the detection method based on annotated timestamps combined with the strongest braking capability, resulting in a 74.3% avoidance rate with a 95% confidence interval ranging from 57.9% to 85.8%. This corresponds to a confidence interval for the rate of unavoidable crashes ranging from 14.2% to 42.1%, even under the strongest braking profile evaluated.

This result defines a lower bound on the proportion of non-frontal cut-in crashes that are theoretically unavoidable under idealized conditions. Based on this lower bound, at least 14.2% of non-frontal crashes are unavoidable by this model. Since the model is idealized to outperform both human drivers and real-world automated systems, this limitation is expected to apply to those systems as well.

This estimate is based on the boundary of the confidence interval and the very strong braking profile. If, instead, the mean avoidance rate under the strong braking setting is used, the corresponding proportion of unavoidable non-frontal crashes would be 38.1%.

Theoretical boundaries like this can be useful as a benchmark for other studies that evaluates different methods for avoiding cut-ins, like Raiyn & Weidl (2025), but also for manufacturers (BMW, 2025; Mercedes, 2025; Volvo, 2025) of ADAS and autonomous driving systems to aid in the assessment of those systems. It should, however, again be noted that these are theoretical boundaries and many aspects of this assessment are ideal (e.g., perfect sensing without uncertainty, and no other road users present).

5.3 Frontal vs. non-frontal crashes

The results of this study show that all the models performed better by a large margin in the frontal crash scenarios than in the non-frontal crashes. While this result was expected, it highlights the importance of separating these two types of crashes when performing counterfactual simulations of cut-in events. Treating them as the same type of scenario can lead to misleading conclusions.

5.4 The rule-based detection algorithm

As the rule-based algorithm was developed and tested using the same small set of data points, it is possible that that the breakpoints of the algorithm were to some extent fitted to the data and that its results would be worse with another data set. This is not an issue for this study though as the model aimed to be ideal and react as early as possible, but it does mean that one needs to be careful when using this model in other studies.

The rule-based detection algorithm failed to detect the cut-ins entirely in two scenarios. For the rest of the scenarios, it detected the cut-in on average 0.247 seconds after the annotated times. Given the limited time available for a master's thesis, it was not feasible to make significant improvements to the model beyond this. As it would theoretically be possible for a well-developed autonomous system to react faster to cut-ins than this rule-based model, relying on it when attempting to establish an upper bound to the number of avoidable crashes is problematic. For this reason, primarily the results from the model using annotated timestamps were used in the discussion to draw conclusions regarding upper avoidance limits.

5.5 Limitations

There were several limitations in this study that may have influenced the results.

A major limitation is the small dataset, consisting of 13 frontal cut-in crashes and 35 non-frontal ones (and 3 crashes that were excluded due to not resembling cut-ins). This limited number of data points results in wide confidence intervals, which makes it difficult to draw firm conclusions from the findings. While a 100% avoidance rate was observed for the frontal crashes with some models, the author considers it unlikely that this result would generalize to a larger dataset.

In most non-frontal crashes, the cut-in vehicle never moved in front of the ego vehicle. Since the ADF, CCDM, and FSM models are designed to not respond when the cut-in vehicle remains beside or behind the ego vehicle, their poor performance is expected. A more informative comparison could have been achieved by extending them to attempt collision avoidance in non-frontal situations.

While the ADF used in this study offers some insight into how an autonomous system might perform in cut-in scenarios, it is a simplified representation of real-world autonomous driving systems. However, it is unlikely that a complete commercial system would be made available to researchers for direct evaluation.

Another key limitation is the exclusion of surrounding traffic in the simulations. The high deceleration levels used by the ego vehicle could lead to it being struck from behind by a following vehicle, which would constrain the amount of safe deceleration. Including surrounding traffic would increase simulation complexity, as it would require modeling how other vehicles respond to the ego vehicle’s actions.

The annotated time stamps used in this study were manually annotated with trajectories visible, eliminating the need to distinguish between within-lane fluctuations and actual lane changes. This level of clarity would not be available to a driver in real time, introducing a hindsight bias. However, as the model aims to set an upper limit to avoidance rate, this was not considered problematic for the aims of this study.

Finally, another limitation of this work is the exclusion of lateral avoidance maneuvers. Some of the collisions observed in this study might have been preventable through lateral movement. Including such maneuvers would also increase the importance of modeling surrounding traffic in the scenario. Still, when evaluating systems that rely solely on longitudinal control, an upper limit based only on braking may be more appropriate than one that also considers lateral movement.

5.6 Future work

While this work establishes an upper bound on the number of avoidable cut-in crashes, this boundary could likely be refined with additional data, more accurate representations of ideal braking, and further assumptions about how quickly an autonomous vehicle can react. Lateral motion and surrounding traffic are other factors that could improve the accuracy of this estimate.

This boundary can serve as a benchmark for safety systems operating under similar constraints. With continued research, the gap between the number of crashes that theoretically could be avoided and those that practically are avoided should gradually shrink.

This study also reveals a large discrepancy between the two reference driver models. While it is expected that different models may reflect different types of human behavior, the model intended to represent a competent and careful driver would be expected to perform close to the optimal human level. Its considerably lower performance compared to the Fuzzy Safety Model raises questions about the validity and consistency of these models and highlights the need for further work to validate them.

6 Conclusions

In this study, virtual counterfactual simulations were run to set an upper limit on the number of cut-in crashes that can be avoided solely by braking, and to assess how reference driver models and an ADF compare to this theoretical benchmark.

The results show that even under idealized conditions, with fast reactions and high deceleration capabilities, at least 14.2% of all non-frontal cut-in crashes appear to be unavoidable. This estimate is based on the strongest braking profile evaluated and corresponds to the boundary of a 95% confidence interval. As such, it represents a lower bound on the true proportion of unavoidable crashes, suggesting that the actual proportion is likely higher. The full confidence interval for the unavoidable crash rate for the strongest braking profile ranges from 14.2% to 42.1%. In contrast, for frontal cut-in crashes, no such limit could be established because both the ideal model and Hi-Drive's ADF were able to avoid all the collisions.

None of the non-ideal models were able to avoid any of the non-frontal crashes. This is primarily due to the design of these models, which only decelerate in response to vehicles executing a cut-in if the cut-in vehicle is ahead of the ego vehicle.

For the frontal cut-ins, the ADF from Hi-Drive either avoided or reacted too early to all the cut-ins. The reference driver models, CCDM and FSM, avoided 38.5% and 81.8% of the crashes, respectively. While differences in their performance are not necessarily unexpected—since the models may reflect different types of human drivers—the relatively low avoidance rate of the CCDM is notable, given that it is designed to represent a careful and competent driver. It remains unclear whether the CCDM underperforms, the FSM overperforms, or both, relative to real human drivers. Additionally, the confidence intervals for these results are quite wide, which may reduce the reliability of this comparison.

7 References

Alsuwian, T., Saeed, R. B., & Amin, A. A. (2022). Autonomous Vehicle with Emergency Braking Algorithm Based on Multi-Sensor Fusion and Super Twisting Speed Controller. *Applied Sciences*, *12*(17), 8458. https://www.mdpi.com/2076-3417/12/17/8458

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

Bjorvatn, A., Page, Y., Fahrenkrog, F., Weber, H., Aittoniemi, E., Heum, P., Lehtonen, E., Silla, A., Bärgman, J., Borrack, M., Innamaa, S., Itkonen, T., Malin, F., Pedersen, K., Schuldes, M., Sintonen, H., Streubel, T., Hagleitner, W., Hermitte, T.,…Torrao, G. (2021). *L3Pilot Deliverable D7.4: Impact evaluation results*.

BMW. (2025). *BMW front page*. https://www.bmw.com/

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. https://doi.org/https://doi.org/10.1016/j.trf.2015.10.011

Chovan, J. D., Tijerina, L., Alexander, G., & Hendricks, D. L. (1994). Examination of Lane Change Crashes and Potential IVHS Countermeasures [Tech Report]. https://rosap.ntl.bts.gov/view/dot/2978

Esmini. (2024). *Esmini (Github)*. https://github.com/esmini/esmini

Green, M. (2000). "How Long Does It Take to Stop?" Methodological Analysis of Driver Perception-Brake Times. *Transportation Human Factors*, *2*(3), 195-216. https://doi.org/10.1207/STHF0203_1

Habenicht, S., Winner, H., Bone, S., Sasse, F., & Korzenietz, P. (2011, 5-9 June 2011). A maneuver-based lane change assistance system. 2011 IEEE Intelligent Vehicles Symposium (IV),

Hanowski, R. J., Olson, R. L., Hickman, J. S., & Dingus, T. A. (2006). The 100-Car Naturalistic Driving Study: A Descriptive Analysis of Light Vehicle-Heavy Vehicle Interactions from the Light Vehicle Driver's Perspective [Research Paper]. https://doi.org/https://doi.org/10.21949/1502657

Hi-Drive. (2020). *Hi-Drive (front web page)*. https://www.hi-drive.eu/

Hi-Drive. (2025). *Deliverable 7.3: Effects on Safety (unpublished project deliverable)*.

Isaksson-Hellman, I., & Norin, H. (2005). How Thirty Years of Focused Safety Development has Influenced Injury Outcome in Volvo Cars. *Annual proceedings / Association for the Advancement of Automotive Medicine. Association for the Advancement of Automotive Medicine*, *49*, 63-77.

ISO. (2021). Road vehicles — Prospective safety performance assessment of pre-crash technology by virtual simulation — Part 1: State-of-the-art and general method overview. *ISO/TR 21934-1:2021(E)*. https://www.iso.org/obp/ui/en/#iso:std:iso:tr:21934:-1:ed-1:v1:en

ISO. (2024). Road vehicles - Prospective safety performance assessment of pre-crash technology by virtual simulation, Part 2: Guidelines and requirements for application. *ISO/TS 21934-2:2024*. https://www.iso.org/standard/81790.html

JAMA. (2022). *Automated Driving Safety Evaluation Framework*.

Jokhio, S., Olleja, P., Bärgman, J., Yan, F., & Baumann, M. (2024). Analysis of Time-to-Lane-Change-Initiation Using Realistic Driving Data. *IEEE Transactions on Intelligent Transportation Systems*, *25*(5), 4620-4633. https://doi.org/10.1109/TITS.2023.3329690

Kovaceva, J., Bärgman, J., & Dozza, M. (2022). On the importance of driver models for the development and assessment of active safety: a new collision warning system to make overtaking cyclists safer. *Accident Analysis & Prevention*, *165*, 106513.

Li, P., Zhu, X., Ren, Y., Tan, Z., Hu, W., Zhang, Y., & Xu, C. (2024). Generalization of cut-in pre-crash scenarios for autonomous vehicles based on accident data. *Scientific Reports*, *14*(1), 17664. https://doi.org/10.1038/s41598-024-68263-4

Mattas, K., Albano, G., Donà, R., Galassi, M. C., Suarez-Bertoa, R., Vass, S., & Ciuffo, B. (2022). Driver models for the definition of safety requirements of automated vehicles in international regulations. Application to motorway driving conditions. *Accident Analysis & Prevention*, *174*, 106743. https://doi.org/10.1016/j.aap.2022.106743

Mattas, K., Makridis, M., Botzoris, G., Kriston, A., Minarini, F., Papadopoulos, B., Re, F., Rognelund, G., & Ciuffo, B. (2020). Fuzzy Surrogate Safety Metrics for real-time assessment of rear-end collision risk. A study based on empirical observations. *Accident Analysis & Prevention*, *148*, 105794. https://doi.org/https://doi.org/10.1016/j.aap.2020.105794

Mercedes. (2025). *Mercedes-Benz Front Page*. https://www.mercedes-benz.com/

Miholjcic, D., Fabbroni, M., & Robinson, R. (2019). *A Study of the Performance of Automatic Emergency Braking (AEB) Systems Equipped on Passenger Vehicles for Model Years 2013 to 2018* https://doi.org/10.4271/2019-01-0416

Olleja, P., Markkula, G., & Bärgman, J. (2025). Validation of human benchmark models for automated driving system approval: How competent and careful are they really? *Accident Analysis & Prevention*, *213*, 107922. https://doi.org/https://doi.org/10.1016/j.aap.2025.107922

Page, Y., Fahrenkrog, F., Fiorentino, A., Gwehenberger, J., Helmer, T., Lindman, M., Camp, O., Rooij, L., Puch, S., Fränzle, M., Sander, U., & Wimmer, P. (2015). *A comprehensive and harmonized method for assessing the effectiveness of advanced driver assistance systems by virtual simulation: The P.E.A.R.S. initiative*.

Raiyn, J., & Weidl, G. (2025). Improvement in Collision Avoidance in Cut-In Maneuvers Using Time-to-Collision Metrics. *Smart Cities*, *8*(1), 15. https://www.mdpi.com/2624-6511/8/1/15

Rodemerk, C., Habenicht, S., Weitzel, A., Winner, H., & Schmitt, T. (2012, 3-7 June 2012). Development of a general criticality criterion for the risk estimation of driving situations and its application to a maneuver-based lane change assistance system. 2012 IEEE Intelligent Vehicles Symposium,

Sander, U. (2018). *Predicting Safety Benefits of Automated Emergency Braking at Intersections-Virtual simulations based on real-world accident data*. Chalmers University of Technology.

Sintonen, H., Bolovinou, A., Guelsen, B., Hogema, J., Innamaa, S., Lee, Y. M., Mallada, J. L., Metz, B., Reimer, F., & Tango, F. (2023). Hi-Drive D4. 3 Experimental procedure: Version 1.1.

UNECE. (2023). UN Regulation No. 157 - Automated Lane-Keeping Systems (ALKS). https://unece.org/transport/documents/2023/03/standards/un-regulation-no-157-amend4

Vater, L., Aittoniemi, E., Bjorvatn, A., Fahrenkrog, F., Flechon, C., Innamaa, S., Lehtonen, E., Rahmani, S., Weber, H., Anagnostopoulou, C., Bärgman, J., Bolovinou, A., Calvert, S., Eckstein, L., Georgiou, N., Heum, P., Hogema, J., Itkonen, T., Klunder, G.,...Jie, Z. (2023). *Hi-Drive D4.5 Effects evaluation methods*.

Volvo. (2025). *Volvo Cars Corporation Front Page*. https://www.volvocars.com/

Wallis, S. (2013). Binomial Confidence Intervals and Contingency Tests: Mathematical Fundamentals and the Evaluation of Alternative Methods. *Journal of Quantitative Linguistics*, *20*(3), 178-208. https://doi.org/10.1080/09296174.2013.799918

Wimmer, P., Camp, O., Weber, H., Chajmowicz, H., Wagner, M., Mallada, J., Fahrenkrog, F., & Denk, F. (2023). *Harmonized Approaches for Baseline Creation in Prospective Safety Performance Assessment of Driving Automation Systems*.

Wimmer, P., Michael, D., Henri, C., Fredrik, G., Julian, K., Harald, K., Olaf, O. d. C., Paolo, S., & and Wagner, M. (2019). Toward harmonizing prospective effectiveness assessment for road safety: Comparing tools in standard test case simulations. *Traffic Injury Prevention*, *20*(sup1), S139-S145. https://doi.org/10.1080/15389588.2019.1616086

World Health, O. (2023, 2023/12). *Road Traffic Injuries*. World Health Organization. https://www.who.int/news-room/fact-sheets/detail/road-traffic-injuries

Department of Mechanics and Maritime Sciences
Division of Vehicle Safety
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2025
www.chalmers.se