Measuring the Safety of ADS: How Safe is Safe Enough?

December 2023 | Final Report

SAFE-D
SAFETY THROUGH DISRUPTION

VIRGINIA TECH TRANSPORTATION INSTITUTE
VIRGINIA TECH

Texas A&M Transportation Institute

San Diego State University
Leadership Starts Here

Disclaimer

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No. 06-014	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle Measuring the Safety of ADS: How Safe is Safe Enough?		5. Report Date December 2023 6. Performing Organization Code:
7. Author(s) Eileen Herbers Andy Petersen Jon Hankey Zachary Doerzaph Elizabeth White		8. Performing Organization Report No.
9. Performing Organization Name and Address: Safe-D National UTC Virginia Tech Transportation Institute		10. Work Unit No. 11. Contract or Grant No. 69A3551747115/[Project 06-014]
12. Sponsoring Agency Name and Address Office of the Secretary of Transportation (OST) U.S. Department of Transportation (US DOT)		13. Type of Report and Period Final Research Report 7/2022-12/2023 14. Sponsoring Agency Code

15. Supplementary Notes

This project was funded by the Safety through Disruption (Safe-D) National University Transportation Center, a grant from the U.S. Department of Transportation – Office of the Assistant Secretary for Research and Technology, University Transportation Centers Program.

16. Abstract

This research aims to determine how to measure the potential safety benefit of automated vehicle technologies and automated driving systems (ADS), and where current technology practices may not deliver on projected safety promises in order to answer a portion of the question "how safe is safe enough." A missing piece that could allow for more cohesion and safer implementation of ADS is the knowledge of what type of data is needed for the refinement and further development of these systems, as well as which scenarios may not be able to be addressed by the currently implemented technology. The purpose of this project was to use naturalistic driving data to inform the scenario selection that is used to measure how ADS will perform in these scenarios. This project highlighted key areas where ADS may fail, such as in high-speed turns, blind turns and hills, lane-change events with other vehicles, and scenarios in which there is significant visual obstruction and occlusion.

17. Key Words Automated Vehicles, Connected Vehicles,	18. Distribution Statement No restrictions. This document is available to the public through the Safe-D National UTC website, as well as the following repositories: VTechWorks, The National Transportation Library, The Transportation Library, Volpe National Transportation Systems Center, Federal Highway Administration Research Library, and the National Technical Reports Library.

19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 20	22. Price $0

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

Abstract

This research aims to determine how to measure the potential safety benefit of automated vehicle technologies and automated driving systems (ADS), and where current technology practices may not deliver on projected safety promises in order to answer a portion of the question "how safe is safe enough." A missing piece that could allow for more cohesion and safer implementation of ADS is the knowledge of what type of data is needed for the refinement and further development of these systems, as well as which scenarios may not be able to be addressed by the currently implemented technology. The purpose of this project was to use naturalistic driving data to inform the scenario selection that is used to measure how ADS will perform in these scenarios. This project highlighted key areas where ADS may fail, such as in high-speed turns, blind turns and hills, lane-change events with other vehicles, and scenarios in which there is significant visual obstruction and occlusion.

Acknowledgements

This project was funded by the Safety through Disruption (Safe-D) National University Transportation Center, a grant from the U.S. Department of Transportation – Office of the Assistant Secretary for Research and Technology, University Transportation Centers Program. The team would like to thank Kevin Kefauver as the subject matter expert for this report.

Table of Contents

List of Figures

List of Tables

Introduction

When looking to the future deployment of automated driving systems (ADS), a common question is how safe is safe enough? One potential answer is that as long as the system is safer than a human driver, then it is safe enough to deploy. This answer is much more complicated than it initially appears to be because it depends on understanding how to quantify human safety, which in turn relies on understanding human driver behavior. For example, the average human driver is not very safe. In 2022, it is projected that 42,795 deaths occurred on US roadways [1], which represents a small decrease (0.3%) compared to 2021. A survey collected by the National Highway Traffic Safety Administration in 2005–2007 found that the reason for the critical pre-crash event for 94% of crashes in the US can be assigned to the driver [2]. Although this statistic does not directly assign blame to the driver, it indicates that the critical reasons for crashes are primarily driver-related, which can be further categorized into driver recognition error, decision error, performance error, non-performance error, and other (Figure 1).

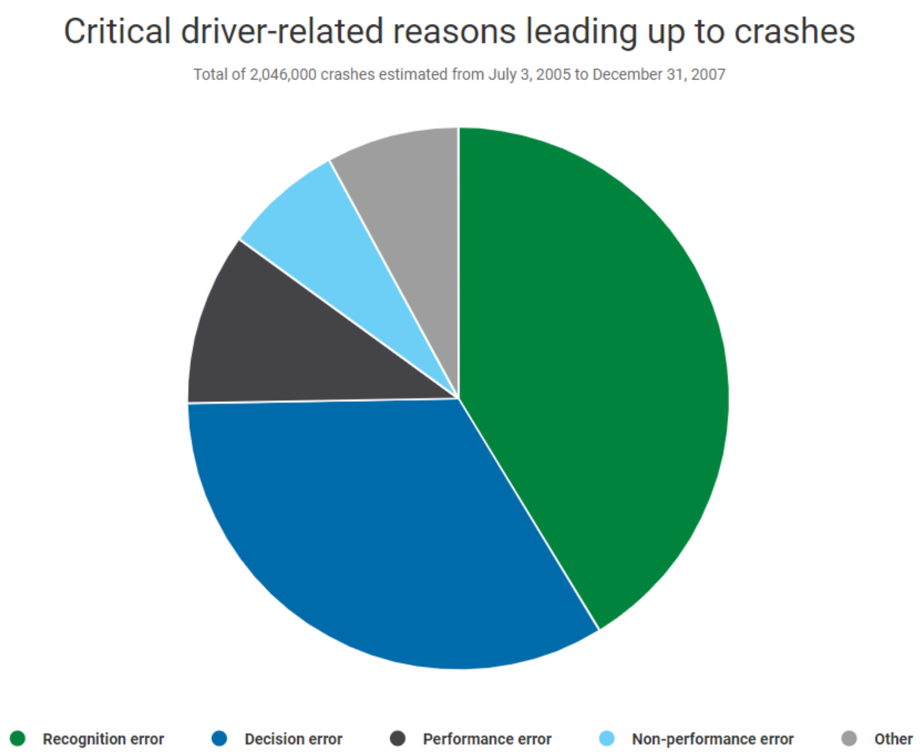

Figure 1. Chart. Critical driver-related reasons leading to 2,046,000 crashes estimated in 2005-2007 [2].

If humans are already not very safe, then how do we determine what is acceptably safe? Primarily, "acceptably safe" can be viewed from a variety of lenses. For example, current drivers view driving risk differently from the engineers who develop the vehicles, from insurance companies, and from regulatory agencies. ADS has the potential to reduce the risk of driving a vehicle by removing the driver from the equation. "Safe enough" means that everyone involved believes that the benefits of ADS deployment outweigh the risks. The research in this report looks at specific real-world events when the actual driving risk results in a safety-critical event (SCE), and how ADS could have avoided or mitigated this event. By calculating safety surrogate measures, this research can be used to quantitatively show how the current driving risk relates to the potential risk of ADS in order to begin to answer the question of "how safe is safe enough?"

The SCEs were extracted from naturalistic driving studies (NDS) and analyzed to determine how ADS could have affected the outcome of these crash and near-crash events, and how additional sensors and technology, like vehicle communication, could support the ADS during these events to ensure safe operation. Vehicle communication in this context refers to vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) technology, which allows for the exchange of information between vehicles on the road, and between vehicles and the infrastructure. Combined, these can provide vehicles with real-time information about road conditions, traffic congestion, crashes, construction zones, and even other vehicles' trajectories.

Including V2V and V2I technologies in ADS is a possible solution to addressing many vehicle safety problems. However, a common problem in getting ADS on the road today is the ever-changing trade-off between what features are necessary for the public sector and which are necessary for the private sector. The uncertainty of which technology will be available in the infrastructure becomes a problem when developing ADS because these vehicles must then develop a solution for every possible scenario.

This research highlights scenarios that cannot be simply addressed by ADS alone, and how more vehicular communication could reduce or mitigate these specific SCEs in the hopes that this can inform decisions for determining what is required for V2I communication.

The purpose of this research is to expand the understanding of ADS capability using exemplar cases from naturalistic driving data that demonstrate situations in which ADS may not provide the intended benefits. This may help safety researchers understand potential challenges to wide-scale ADS deployment that simulation or closed-track testing cannot and thus address an important aspect of the question "how safe is safe enough?"

Background

Levels of Automation

In order to discuss the potential safety impact of ADS, we first define what ADS is and how it fits into the levels of automation. ADS combine technological advances in a vehicle that can be classified within one of the six levels of driving automation defined by the Society of Automotive Engineers (SAE) in J3016 [3]. These levels of driving automation are used to clarify the capacity in which a vehicle is fully autonomous, and when a driver is still expected to be fully engaged in the driving task. Additionally, by having these levels of automation clearly stated, it makes it easier to determine what is expected of each vehicle, and thus makes it easier to grade the performance of the vehicle. From Figure 3, we can see that Levels 0–2 (L0, L1, and L2) expect the human driver to continually monitor the road. In Levels 3–5 (L3, L4, and L5), the vehicle itself is in charge of monitoring the road. However, in cars with L3 and L4, the driver can be the fallback when automation fails or when the vehicle is outside of its operational design domain (ODD). It is important to make sure that the system correctly identifies when this occurs so that it can return functionality to the human driver at the correct moment, and that the human driver is prepared to do so.

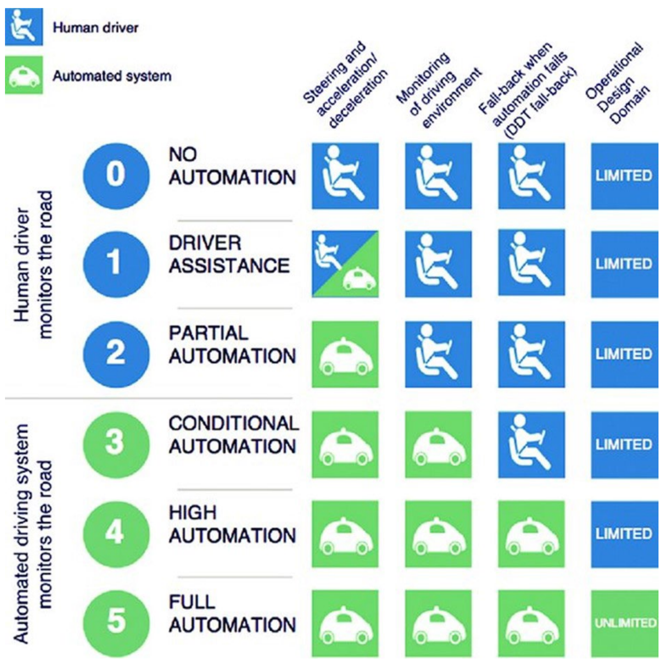

Figure 2. Illustration. SAE J3016 levels of driving automation [3].

ADS generally signifies cars with L3 and above, whereas advanced driver assistance systems (ADAS) are characterized as L1 and L2 automation. Currently, most cars available to the public are those with L2 capabilities and below. There are some vehicles now available with higher levels of autonomy, but they are still rare in 2023. In 2022, 46.5% of vehicle sales were cars with L2, and 3.9% of vehicle sales were vehicles marketed as L2+, which are essentially L2 vehicles with more advanced ADAS features that still require drivers to monitor the driving environment at all times [4]. Therefore, current crash databases do not have data on crashes involving ADS, and different methods are required to predict the potential safety benefit of ADS technology. This project aims to look at how the potential safety benefit of ADS can be predicted, and how to assess the accuracy of the predictions.

How to Measure the Potential Impact of ADS

The purpose of this research is to ensure that the potential safety benefit of ADS is realized. First, we must understand the current approaches to measure this theorized impact. This includes looking at the impact of ADAS features and how current technologies on the road have reduced crashes, or how safety testing of these features has been done in the past. This also includes understanding the simulation and field testing that ADS developers and researchers use to assess how ADS might perform in certain safety-critical scenarios.

Impact of ADAS (L1 and L2)

Before we look at the impact of ADS on roadway safety, it makes sense to first look at the theorized and actual impact of ADAS technologies on the road, since ADAS essentially form the basis of ADS [5]. ADAS are primarily focused on collision avoidance technologies (like automated emergency braking [AEB], lane-keeping assist [LKA], and blind spot warning [BSW]), whereas ADS could be classified as a driver convenience feature since the purpose is to take on more of the driving task. Advances in ADAS have the potential to "reduce roadway crashes, fatalities, and injuries and assist the USDOT in managing safety risks along the path" to ADS [5]. It is anticipated that increasing the overall market penetration of these systems and improvements in their functionality and performance will contribute to overall improvements in traffic safety.

Crash Statistics

There are many studies that capture how ADAS may reduce the number of fatalities, injuries, and crashes in general. Research done by the AAA Foundation for Traffic Safety sought to quantify the number of crashes, injuries, and deaths that occurred in the US in 2016 that theoretically might have been avoided or reduced in severity had the involved vehicles been equipped with AEB, LKA, BSW, forward collision warning (FCW), and lane-departure warning (LDW). They found that, if installed on all vehicles, these ADAS would have had the potential to prevent or mitigate roughly 40% of all crashes involving passenger vehicles, 37% of all injuries, and 29% of all fatalities that occurred in those crashes [6]. A study at the University of North Carolina developed a model that estimated that ADAS technologies currently available to US consumers could prevent approximately 16% of crashes and injuries, and 22% of deaths between 2012 and 2050 [7]. The National Security Council used data on the predicted availability of ADAS features from the Highway Loss Data Institute and the crash population for crash avoidance technologies to estimate that ADAS technologies could have prevented about 62% of total traffic deaths between 2011 and 2015 [8, 9].

In calculating these percentages, different inclusion and exclusion criteria were used to determine the populations of crashes that would have been or could be avoided with ADAS. Inclusion criteria, as seen in Table 1, include factors from the crash that indicate that an ADAS-equipped vehicle could have avoided or mitigated the crash.

Table 1. Examples of the target population of crashes for each ADAS feature.

ADAS Feature	Target Population
AEB, FCW	Crashes in which a vehicle or pedestrian is rear-ended by a forward-moving passenger vehicle.
LDW, LKA	Crashes in which the vehicle left its lane prior to the initial event, but intentionally did so.
BSW	Crashes in which a vehicle left its travel lane prior to the initial event and was not leaving the travel lane on purpose.

These inclusion criteria do not take into account other factors that would indicate if an ADAS feature could have successfully mitigated or prevented a crash. For example, most ADAS features operate at specific speed ranges, and so the ADAS feature may not be active at very slow or very fast speeds [10]. BSW systems may not be able to warn drivers of vehicles travelling much faster in an adjacent lane [6]. Additionally, tight curves or steep inclines may affect how well the vehicle could brake in AEB situations [11]. The crash population percentages are a theoretical upper bound for the population of crashes that could have been affected by having ADAS available for all vehicles on the road. Therefore, there are many crashes within this population that would still most likely not be able to be addressed by ADAS, but more information about the specifics of the crash would need to be known.

It is also important to look at the exclusion criteria, or the populations of crashes that were deemed to be not addressable by ADAS. This included crashes in which the police indicated that the driver was asleep, ill, or impaired by alcohol or drugs. Additionally, any events in which sensors would be blocked or dirty, or lane markings would be absent, degraded, or obscured by rain or snow would not be addressable by ADAS. Crashes like these have the potential to be addressable with ADS if there is additional vehicular communication between the infrastructure and the ADS (V2I). However, any crashes in which the road conditions are not ideal (snow/ice) or if there is any off-roading involved would not be simply addressed with the inclusion of advanced technology.

Field Testing

The crash statistics are a top-down approach to estimate how many crashes could have been avoided with ADAS features. Field tests use a bottom-up approach by testing ADAS features and calculating which metrics can be used to assess the safety of each feature. For example, the National Highway Traffic Safety Administration (NHTSA) completed an objective test for automatic crash imminent braking (CIB) systems that defines minimum performance requirements and objective tests for CIB [12]. Results from tests like these are used to develop standards and requirements for ADAS features, such as the new requirement for AEB systems for light vehicles [13].

The variables used in these tests are similar to the inclusion and exclusion criteria in the crash statistics. For example, the crash configuration, crash object, speed, acceleration, braking maneuver, range, impact location, and time-to-collision were all standardized for each test.

Insurance Claims

Another way to determine if ADAS has contributed to fewer or less severe crashes is by evaluating how insurance claims have changed since ADAS-equipped vehicles have been on the road. The Highway Loss Data Institute estimates that, overall, many collision avoidance technologies have reduced the frequency of claims: "front automatic emergency braking (AEB) and rear AEB, technologies that automatically take action for the driver in a crash-imminent situation, were associated with larger reductions in the frequency of claims than technologies that rely on the driver to respond to warnings" [14]. Another analysis of 11 million vehicles for model years 2014–2019 showed a notable reduction in loss cost for bodily injury, property damage, and collision claims [15].

Impact of ADS (L3–L5)

When it comes to measuring the potential safety impact of ADS, there are no crash statistics or insurance claim changes available yet for ADS since they are not on the road. Therefore, only field testing or simulation studies can effectively measure the potential safety impact of ADS.

Simulation

Simulation studies have also been used to measure the safety impact of ADAS features, but for the purposes of this research, we are only focusing on simulation studies that have looked at measuring ADS safety. Similar to the initial design of this research, much ADS simulation testing revolves around scenario testing. Because testing every possible scenario that ADS may come across can be very time-consuming, one study was able to transform single highway scenarios into multiple different tests to evaluate the safety benefit of ADS [16]. Another study used deep learning to develop a range of factors and simulations to test scenarios that would inherently put ADS at risk [17]. They used safety surrogate measures to assess the safety impact of ADS within each of these scenarios. A combination of micro and macro traffic simulations can also be used to predict how ADS would affect the overall traffic landscape and safety outcomes of potential conflicts. A study using surrogate safety models and traffic microstimulators found that automated vehicles (AVs) could reduce the number of conflicts by 20% in signalized intersections and 29% in roundabouts [18].

These are all very similar to the research that this project proposes. However, for this project, we used real-world crash and near-crash events to simulate if ADS could have avoided or mitigated these SCEs. This ability to use naturalistic driving data allows for more insight into unique situations that are not available from simulating post-crash events that do not have the available pre-crash data.

Potential Problems

There is also a question of how ADS may develop new problems. According to a study that used reaction times and accident statistics to predict vulnerable road user (VRU) interaction with ADS, AVs "might very well give rise to new types of crashes (in particular sensor failure, cyber risks or misinterpreted behavior)" [19]. Even if ADS are designed perfectly to interact with other vehicles and VRUs, there is still a share of accidents in which VRU action is a major contributing accident cause.

Research Gap

When it comes to determining "how safe is safe enough" for ADS deployment, the approaches to determine what is acceptably safe can become convoluted. One way to assess ADS safety can be generally organized into three categories: measurement, process, and thresholds [20-22]. Measurement establishes evidence to determine acceptable safety. Process includes activities associated with the verification and validation of safety measures. Thresholds establish the level of safety that an AV must achieve to be considered acceptably safe. This research focuses on using leading measurements (like safety surrogate measures) to highlight scenarios and processes that might not be addressable by ADS alone.

From a review of current processes, there is a population of SCEs that are predicted to be mitigated or avoided by the introduction of ADS but might actually be unavoidable. "How safe is safe enough" cannot be answered accurately if there are some crashes and near-crashes that are embedded into some safety threshold that should in fact be addressed separately. This research uses current crash populations to determine which scenarios have not been addressed. Figure 3 is an illustration that shows an example percentage of the population of crashes that could be avoided by ADAS and ADS. The large rectangle represents the total population of crashes, and the color blocks show the relative percentage of SCEs that could be avoided with different automation levels. This research focuses on real-world events collected through NDS that may not necessarily be avoided by ADS alone; these events may highlight where the proposed safety benefit of ADS may not be realized.

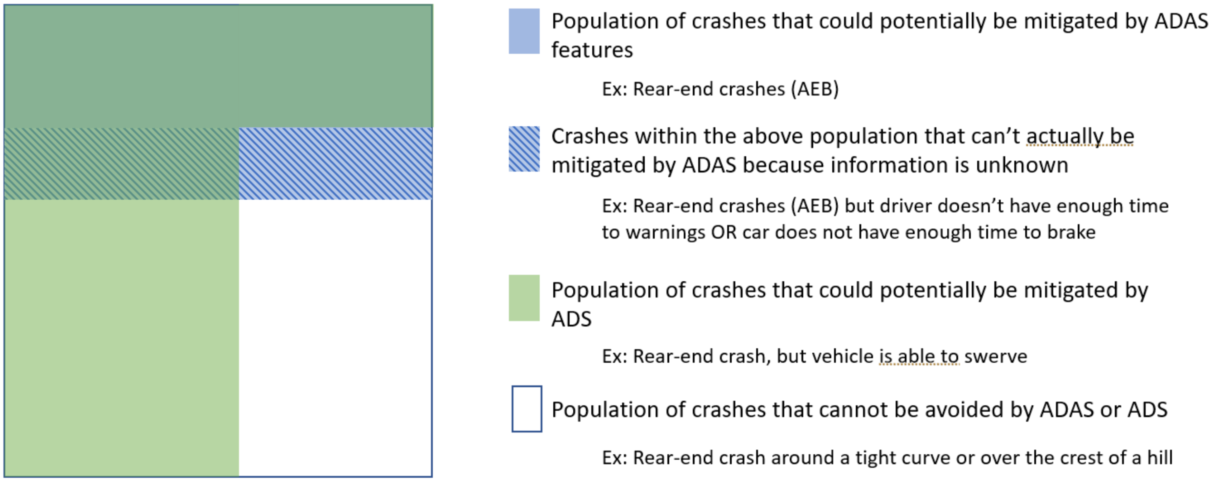

Figure 3. Diagram. Example of the percentage of SCEs that could be avoided by ADAS and ADS.

Methods

Dataset

NDS

The data for this research came from a set of NDS: the Second Strategic Highway Research Program (SHRP 2) NDS, and the Canada NDS. NDS capture real-world driving conditions and real-time driver behavior and vehicle kinematics by the continuous recording of driving information using advanced instrumentation. The SHRP 2 database consists of over 5.5 million trips driven by 3,542 drivers across six collection sites in the continental United States [23]. The Canada NDS contained 149 participants and was conducted at one Canadian site that was comparable to the U.S. sites that were part of SHRP 2 [24]. Both NDS include vehicles with data acquisition systems that collect a variety of data under normal driving conditions. This includes time-series data such as speed, acceleration, yaw, and radar data such as range, range rate, and heading. Through data reduction, events are chosen that can be broken up into four case type levels: crash, near-crash, baseline, and epoch. This project looked at SCEs, which encompass crashes and near-crashes. Each event then has other coded features such as driver evasive maneuver, event nature, impact or proximity time, driver reaction time, locality, driver behavior, etc.

Scenario Selection

Events were first chosen based on the difficulty for them to be addressed by ADS. Based on the inclusion and exclusion criteria mentioned in the background section, they were broken into two categories:

- Operator Factor: Fault of the other driver
- Visual Obstruction: Present

From initial video review and reviewing the literature, it was determined that ADS might have trouble mitigating or reducing crashes and near-crashes in these two categories. Many crash and near-crash events can be predicted when the current driver is at fault because their following distance is too close, their speed is too high, or the driver does not notice another vehicle in its conflict path. However, other drivers on the road are much more difficult to predict. Drivers may drive erratically or unpredictably, which makes it much harder for ADS to perform the correct evasive maneuver. Additionally, any visual obstruction makes it much harder to accurately predict any conflict events in which there are surprise events, or the conflict cannot be identified by the sensors on the ADS vehicle.

The selection of these scenarios led to 3,009 events. After some quality assurance, as described below, we were left with 1,780 events with reliable time-series data. However, the original 3,009 events were still used to initially assess some driver behaviors and scenario classifications.

Data Analysis

Timepoints

For each of the events, some initial timepoints were chosen to calculate different safety surrogate measures. These timepoints are shown in Table 2.

Table 2. Important event timepoints and their definitions.

Variable	Definition
T0	The timepoint that the conflict object was identified by the radar.
T1	The conflict begin timepoint. This is the timepoint at which the sequence of events defining the crash or near-crash begins; for example, the lead vehicle brake lights activate, the vehicle enters the intersection, or the vehicle loses control.
T2	The subject reaction start. This is the timepoint at which the driver initially reacts to the conflict.
T3	Impact or proximity frame. This is the timepoint at which the vehicle first makes contact with the conflict object, or when they are at the closest proximity.

Conflict Object

The first necessary calculation was to determine which radar object corresponded to the conflict object. Although the datasets for each NDS were similar, the available time-series data varied depending on the dataset. Some datasets had validated data that included the actor type of each radar object. These were simple to extract, since any actor labeled as Actor Type 2 corresponded to the conflict object. However, older datasets did not have the actor type variable available. Therefore, the conflict object was determined to be the closest detected object at timepoint T3 (the impact or proximity frame). To validate this, a check was done to see if the conflict object was also the lead vehicle as long as the event included a conflict with the lead vehicle. If the check did not match, then the video was manually reviewed to determine if the correct radar object was identified.

After this quality assurance, 764 events were removed because there was no radar data available, 91 events were removed because the subject vehicle was struck from behind, and another 365 events were removed because the conflict object was not detected correctly by the radar. Therefore, we were left with 1,780 events.

Safety Surrogate Measures

Safety surrogate measures are mathematical measures used to predict the safety implications of new practices, which are important to use since crashes are rare events [25]. The safety surrogate measures chosen in this research were used to establish how to potentially measure the safety impact of ADS. These include the relative velocity, the time-to-collision (TTC), and the minimum required deceleration. These measures were used to compare how the human drivers performed in these events to how ADS could have performed.

In order to do this, we must also determine at what timepoint the ADS would have begun to perform an evasive maneuver (or identified a potential collision). To simplify this calculation, the conflict begin timepoint (T1) is used as the moment in time in which the ADS would have identified a conflict and made a decision whether an evasive maneuver was necessary. This is a large simplification, and it would be beneficial to predict the trajectory of the subject vehicle and conflict object to determine a timepoint that better reflects current ADS development. However, since ADS development is so dependent on the manufacturer, this simplification is sufficient for this project.

Relative Velocity

The relative velocity is the difference in velocity between the subject vehicle and the host vehicle. If the vehicles are moving towards each other, then this value is negative. If the vehicles are separating, then this value is positive. The range rate from the radar was used, and was validated by using the range of the conflict object over a period of 1 second and the velocity of the host vehicle. The relative velocity at each timepoint was an average of these two values.

TTC

The TTC is the most common safety surrogate measure. It measures the time that the host vehicle would collide with the conflict object based on the current range and range rate [26]. This assumes that the host vehicle and the conflict object will continue on their current trajectories at the current velocities. This is commonly used to measure the potential severity of collisions even if there is not a collision (near-crash). It is simply the range divided by the range rate.

Required Deceleration

The minimum required deceleration is the smallest deceleration value required to avoid the collision if the subject vehicle and conflict object do not change trajectories or velocities. Although this value is already inherent in the TTC calculation, the minimum required deceleration can be more easily compared to deceleration rates already set in ADAS features (like AEB). The lower the value, the less braking that is required to avoid the crash. The required deceleration is calculated by taking the relative velocity divided by the TTC.

Video Review
Throughout the scenario selection, data analysis, and validation, video review was constant. If any specific events stood out, notes were taken to highlight these events to review later. After the data analysis, any outliers were reviewed as well to determine what that could possibly forecast in terms of ADS involvement. This was more of a qualitative analysis, but as patterns emerged, it was easier to take note of cases that would be important for ADS testing and deployment.

Results and Discussion

The results are separated into a few categories based on how ADS might be the most effective. These categories include the driver reaction or evasive maneuver, as well as the event configuration. The event configuration significantly impacts how the safety surrogate measures were calculated, since the direction of the velocity was important for these calculations.

Driver Reaction
Since ADS has the potential to remove the driver from the driving equation, it was important to look at the driver's reactions for the two scenario categories chosen. The diagram in Figure 4 is meant to represent the driver's primary (and potential secondary) reaction. The overlap in the diagrams represents the combination of both reactions for a primary and secondary reaction. This diagram was included to show that the majority of drivers react by braking only, or by braking and steering. Generally, braking is easily addressed by ADAS features because only one direction of motion is required, whereas steering requires two directions of motion. Thus, if we were to compare human drivers to ADS, an evasive braking and steering maneuver requires more complex decision-making from the ADS.

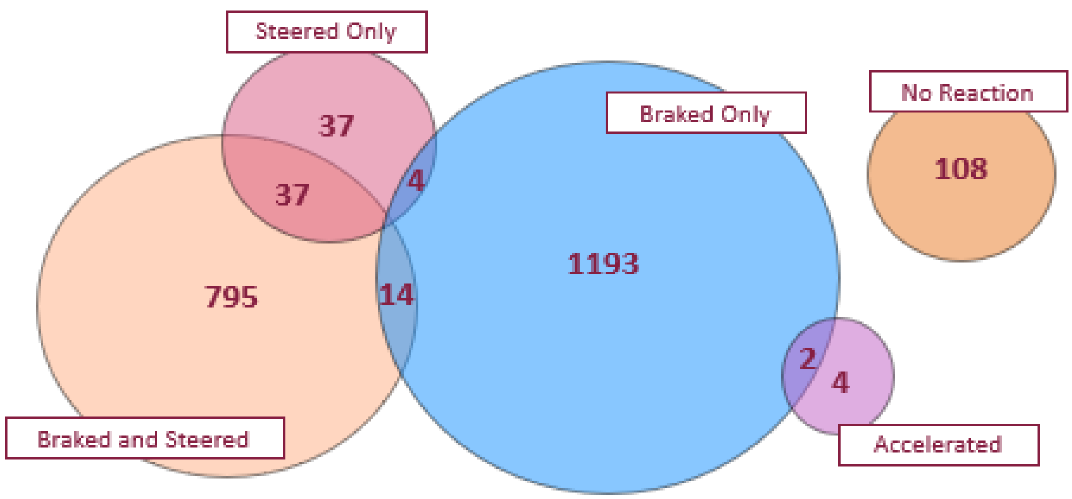

Figure 4. Diagram. The number of crashes and near-crashes broken up by the driver's initial (and secondary) reactions.

No Reaction

It is also important to note the 108 events in which the driver had no reaction. These events could be easily avoided or mitigated by introducing ADS because the driver did not perform any sort of evasive maneuver in these events. Within these scenarios, 102 of them (93.58%) were with vehicles in the adjacent lane in which the subject vehicle was cut-off. With additional blind-spot detection, ADS would easily be able to mitigate these cases.

TTC

The TTC was calculated for the moment when the subject reacted (began braking, steering, accelerating, did not react, or a combination thereof). Figure 5 shows the average TTC for seven event configurations, which were broken up into angle, sideswipe, merge, cut-in; head on (initial opposite directions); forward impact; perpendicular; backing up; and roadside departure. The backing up and roadside departure categories were more of a catch-all for events that did not fit into the other four categories. The line in the center of each box represents the median TTC for that category, the box holds the 25th–75th percentiles, the whiskers represent the 10th and 90th percentiles, and the dots represent outliers. Each category is then split between the subject reaction TTC and the ADS TTC.

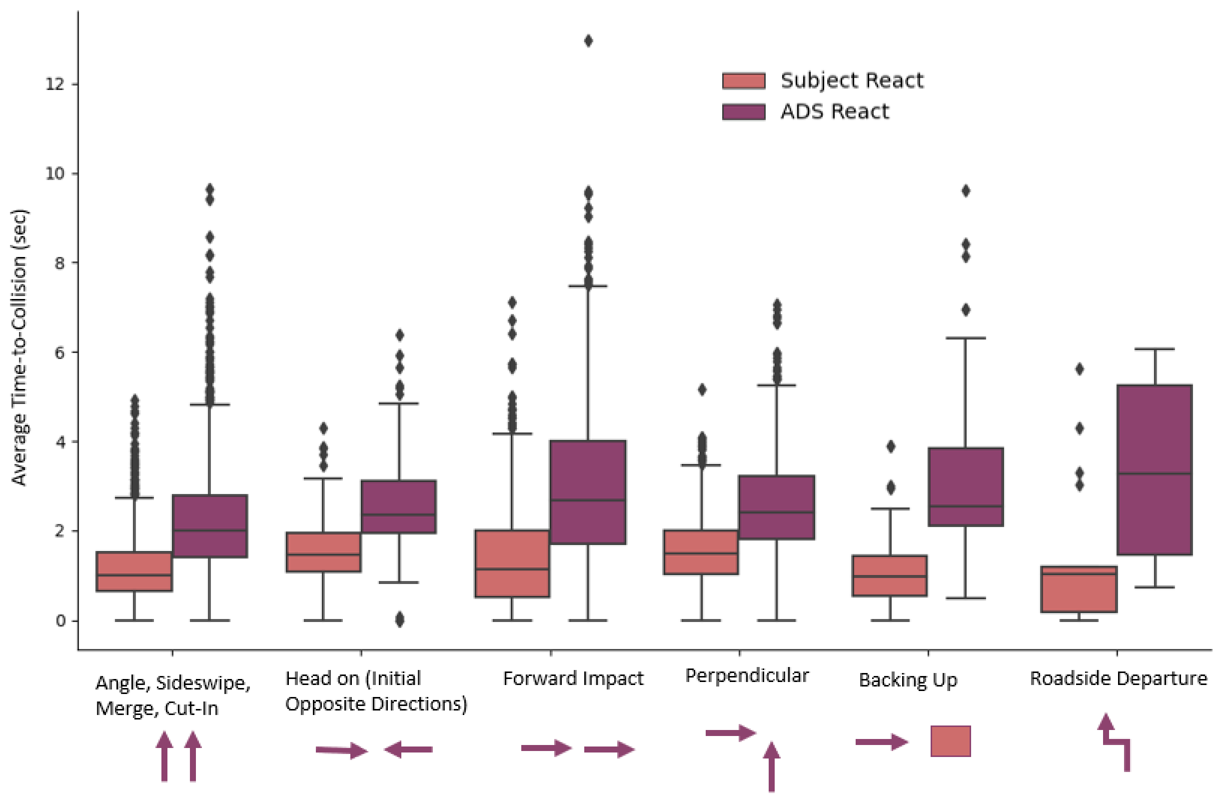

Figure 5. Boxplot. Average difference in TTC (s) between when the subject reacted and when the ADS could have begun braking for seven event configurations.

We can see that the ADS provided a high TTC on average for all seven categories. On average, ADS provided 2.23 additional seconds of TTC across all seven categories. This is an expected result, since we expect that any automated technology will reduce any required reaction time from drivers. However, it is promising that ADS has the potential to provide almost 2 more seconds to make the correct decision. In high-speed scenarios, two seconds provide crucial time to begin braking and reduce the potential energy in a crash.

Minimum Required Deceleration

Although TTC provided an expected result, it is necessary to investigate what that additional time could provide to ADS. The minimum required deceleration to avoid a crash uses the current vehicle speed and the TTC to determine the braking force necessary to avoid a collision between the subject vehicle and the conflict object. Figure 6 shows the number of events for each binned minimum required deceleration value. According to a study by AAA that looked at maximum deceleration values in AEB scenarios, most AEB systems had an abrupt deceleration rate that was around 1 g [27]. The red dashed line in Figure 6 represents the 1-g threshold. Therefore, it could be stated that any events that require a minimum required deceleration above this threshold of 1 g would not be feasible by ADS. In this case, 206 (11.57%) of the events could not be avoided by braking alone, even with the added TTC from an ADS system. Since less than 1% of all the analyzed events actually resulted in a crash, it indicates that human drivers in these events are pretty good at performing the correct evasive maneuver. In fact, in the 206 events that are over the 1-g threshold, the driver both braked and steered in 92.43% of them. Therefore, it is important to look into the combination of the braking and steering maneuvers for situations that ADS might not be able to address. These scenarios would most likely be impacted the most with the implementation of V2I communication.

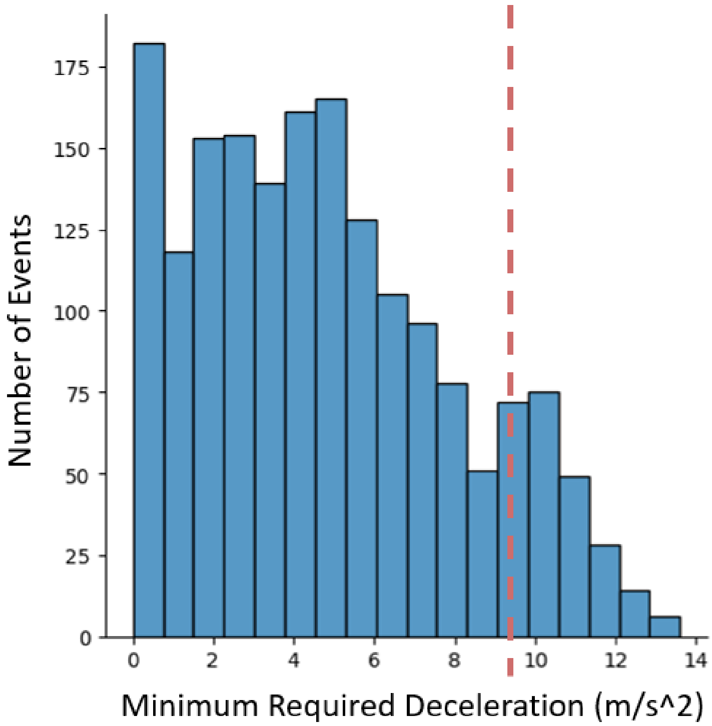

Figure 6. Chart. Minimum required deceleration to avoid a crash if the subject vehicle was equipped with ADS.

To look further into the steering maneuvers of the events, each dot in Figure 7 represents the corresponding yaw and acceleration for each event in which the driver braked and steered. A lower deceleration rate and lower yaw would correspond to a less severe evasive maneuver. The purple dots represent the events that had a minimum required deceleration value of greater than 1 g if the vehicle was equipped with ADS. We can see that the drivers of these events also had higher yaw and deceleration rates.

A traction circle is generally used in physics to depict how fast vehicles can take turns without losing traction. The circle in Figure 7 represents the concept of using a traction circle to see how these events could potentially be evaluated in future analyses. The application of the traction circle could be used to show how vehicle dynamics may influence how effective ADS evasive maneuvers are compared to human drivers (i.e., any events outside of the traction circle would lose traction on the road). The actual application of this concept was considered out of scope for this project, and so Figure 7 would need to include more in-depth analysis into the correct parameters chosen.

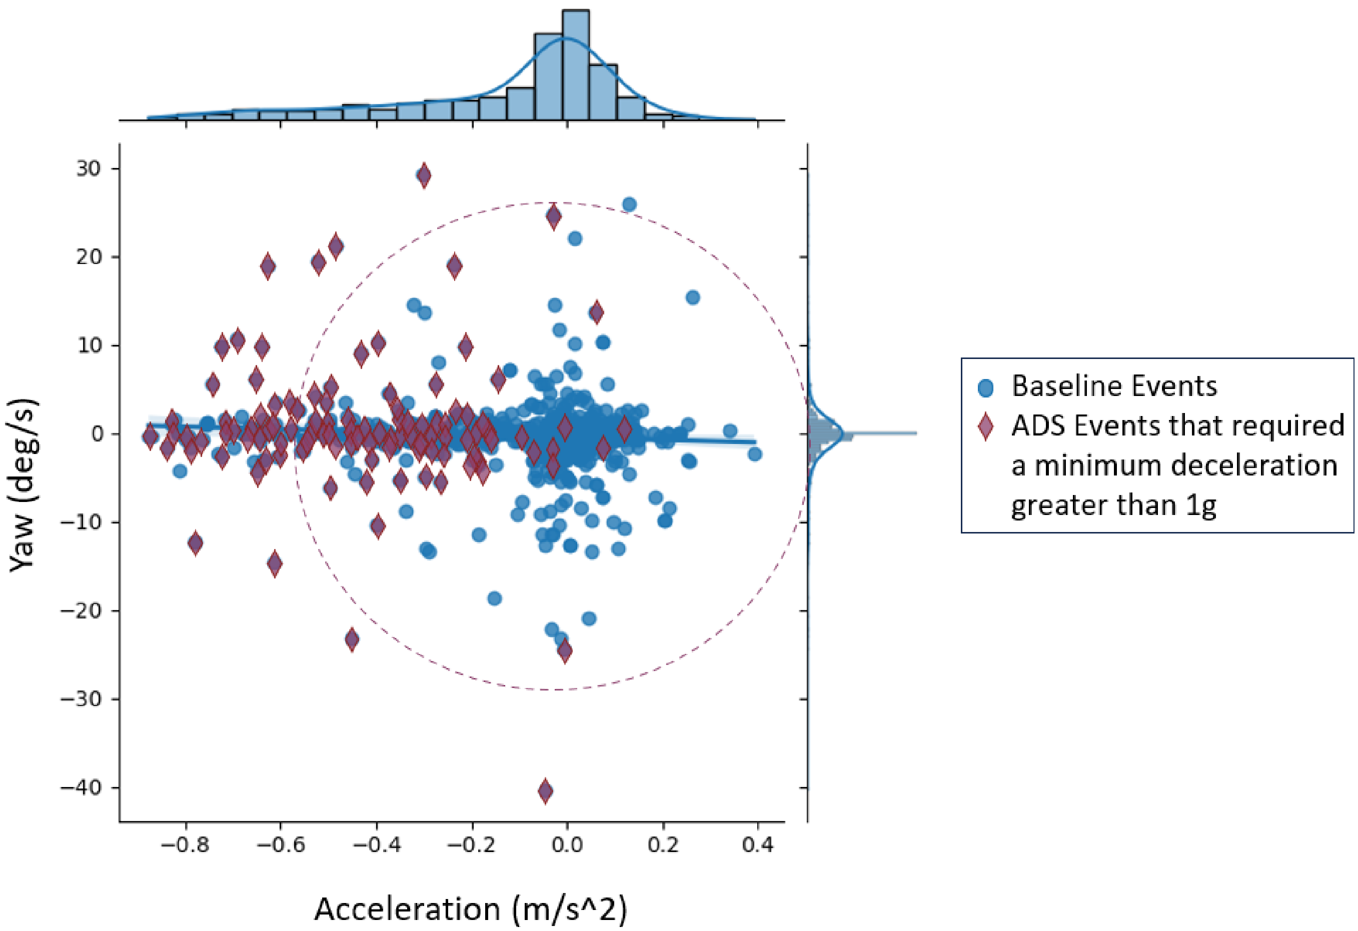

Figure 7. Chart. Yaw (deg/s) and acceleration (m/s^2) for each event with a generalized traction circle. The purple dots represent the events that required a minimum deceleration greater than 1 g for ADS.

Scenario Review

By using the data above, specific scenarios were able to be easily selected to review independently. This is a more anecdotal but important part of the project, as it highlighted key areas in which ADS may not be able to perform as expected. These key areas included high-speed turns, blind turns and hills, lane-change events with other vehicles, and scenarios in which there is significant visual obstruction and occlusion. The most common visual-obstruction cases were those in which the vehicle was making a left turn across stopped traffic and another vehicle came unexpectedly down the outside lane. Although these are very specific examples, they may indicate additional testing that would be required before full ADS deployment.

Conclusions and Recommendations

The answer to "how safe is safe enough" is complicated and ever-changing. The purpose of this research was to expand our understanding of the actual capabilities of ADS when applied to real-world data. Using a small set of naturalistic data has the potential to convey important information to wide-scale ADS deployment that simulation or closed-track testing cannot.

This project looked specifically at events in which ADS may have the potential to not perform well. These scenarios included those in which the "other" driver was at fault, and when there was a visual obstruction leading to the cause of the SCE. Safety surrogate measures such as the relative velocity, TTC, and required deceleration were calculated to compare baseline SCEs to these same events as if the vehicle was equipped with ADS. It was found that although ADS could provide more reaction time before an imminent collision, drivers are generally good at performing evasive maneuvers that require braking and steering, which requires a complex set of decisions for ADS. Although this may be counterintuitive to the original point that "safer than a human driver" might not be safe enough, it is important to note that human drivers can perform better than ADS in certain scenarios.

A large portion of this project was spent reviewing video and looking at videos that specifically were pointed out as potentially being important for ADS deployment from the safety surrogate measures. It was found that key areas in which ADS may not perform as expected include high-speed turns, blind turns and hills, lane-change events with other vehicles, and scenarios in which there is significant visual obstruction and occlusion. One goal of this project is to provide more testing scenarios for developing ADS in order to improve the algorithms and systems within the vehicle, which are highlighted by these key events. These events can also be used to show the importance of V2V or V2I implementation when even the most advanced ADS technology cannot prevent a SCE. This has the possibility of reducing crashes for future ADS, or preventing secondary events caused by the incorrect action of ADS.

The results from this project supplement techniques that have previously been used to determine the safety impact of ADAS technology. This research uses more physics-based methods to determine how an ADS could react in different scenarios and point to real events to determine scenario selection. This can inform future tests that determine the safety impact of ADS to support "how safe is safe enough" arguments.

It would be beneficial to look specifically into how ADS developers use sensor data to predict potential conflicts, as well as to look into more of the physics-based methods of how ADS plan any evasive maneuvers. Additionally, this project would greatly benefit from looking at how additional vehicular communication could enhance ADS safety with V2V or V2I technology.

Although this research plays a role in answering "how safe is safe enough," the answer to this question requires a lot more collaboration between all perspectives of what acceptably safe means. It requires partnership between manufacturers, engineers, researchers, drivers, policy makers, and everyone in-between to determine how safe is safe enough.

Additional Products

Safe-D Project Website:

https://safed.vtti.vt.edu/projects/measuring-the-safety-of-ads-how-safe-is-safe-enough/

In order to comply with participant informed consent and IRB requirements, use of this data set is governed by a data use license (DUL). Please send data use license inquiries to datasharing@vtti.vt.edu.

Education and Workforce Development Products

Course Module

Course Module developed for graduate-level course on advanced transportation systems: BMES 5234 – Advanced Vehicle Safety Systems

This course teaches the emerging challenges, applications, technology, and tools used in advanced vehicles to graduate students at Virginia Tech. This research showed how to use available data to measure the potential safety impact of advanced driving assistant systems and automated driving systems. It highlights some key areas that may be neglected when this technology is being developed and allows the students to find gaps to apply their own research and develop technology that could address these issues.

Women in Transportation Webinar Outreach

This webinar is a panel of women in transportation around the Commonwealth of Virginia that talk to high school girls interested in STEM fields.

Radford High School Career Fair; Radford High School, Radford, VA; April 20, 2023

This career fair had about 150 high school sophomores that rotated through different booths to get exposure to diverse career paths. Various VTTI personnel discussed different research projects they had been a part of and showed some technology on an instrumented research vehicle.

C-Tech^2: Computers and Technology Summer Camp at Virginia Tech; VTTI: Blacksburg, VA; June 26, 2023

A 15-minute presentation about this project to 16 high-school women during a larger tour of VTTI. Led to a great discussion about the future deployment of ADS and questions that the general public may have.

Radford University: Young Women in STEM Summer Camp; Radford University, Radford, VA; July 18, 2023

A 30-minute presentation about this project to about 30 high-school women during a session at their STEM summer camp. Again, this led to a great discussion about their interaction with ADAS features and how ADS may affect how they interact with vehicles in the future.

Technology Transfer Products

Chapter of Dissertation

Herbers, E. (2023). Using Naturalistic Driving Data to Measure the Safety Impact of Automated Driving Systems [Doctoral Dissertation, Virginia Tech].

Conference Presentation

Herbers, E. (2023, September 13-14). Using Naturalistic Driving Data to Measure the Safety Impact of Automated Driving Systems. Distracted Driving Summit, Blacksburg, VA, United States.

(Potential) Herbers, E. (2024, April 7-9). Using Naturalistic Driving Data to Measure the Safety Impact of Automated Driving Systems. Lifesavers Conference on Roadway Safety, Denver, CO, United States.

Webinar Presentation

Herbers, E. (2024) Measuring the Safety of ADS: How Safe is Safe Enough? Safe-D UTC Webinar.

References

1. NHTSA, *Early Estimate of Motor Vehicle Traffic Fatalities in 2022*, in *National Traffic Safey Facts: Crash Stats*. 2023.
2. Singh, S., *Critical Reasons for Crashes Investigated in the National Motor Vehicle Crash Causation Survey*, in *Traffic Safety Facts*. 2018.
3. SAE, *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*, in *Society of Automotive Engineers (SAE) International*. 2021.
4. Sharma, M. and B. Wang, *USA Autonomous Vehicle Tracker Q2 2022*, in *Counterpoint Research*. 2022.
5. USDOT *Ensuring American Leadership in Automated Vehicle Technologies - Automated Vehicles 4.0*, in *National Science & Technology Council United States Department of Transportation (USDOT)*. 2020. p. 56.
6. Benson, A.J., et al., *Potential Reductions in Crashes, Injuries, and deaths from Lrage-Scale Deployment of Advanced Driver Assistance Systems*. 2018, AAA Foundation for Traffic Safety: Washington, D.C.
7. Naumann, R.B., et al., *Examining the Safety Benefits of Partial Vehicle Automation Technologies in an Uncertain Future*, in *AAA Foundation for Traffic Safety*. 2023.
8. Wang, J.-S., *Target Crash Population for Crash Avoidance Technologies in Passenger Vehicles*, in *National Highway Traffic Safety Administration*. 2019.
9. HLDI, *Predicted Availability of Safety Features on Registered Vehicles - A 2022 Update.* Highway Loss Data Institute, 2022. **39**(2).
10. Ayoub, J., et al., *Cause-and-Effect Analysis of ADAS: A Comparison Study between Literature Review and Complaint Data*, in *14th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. 2022: Seoul, South Korea. p. 17.
11. Zhang, L., et al., *Research on Automatic Emergency Braking System Based on Target Recognition and Fusion Control Strategy in Curved Road.* Electronics, 2023. **12**: p. 3490.
12. NHTSA, *Objective Tests for Automatic Crash Imminent Braking (CIB) Systems*, USDOT, Editor. 2011, USDOT: USDOT. p. 150.
13. NHTSA, *NPRM: Automatic Emergency Braking Systems for Light Vehicles*, in *49 CFR Parts 571 and 596* F.M.V.S. Standards, Editor. 2023.
14. HLDI, *Compendium of HLDI collision avoidance research*, in *HLDI*, HLDI, Editor. 20018.
15. Kanet, J. and P. Kohli *The Impact of Advanced Driver Assistance Systems (ADAS) on Insurance Claims*. Lexis Nexis Risk Solutions, 2020.
16. Kim, I., et al., *Method of Evaluating Multiple Scenarios in a Single Simulation Run for Automated Vehicle Assessment.* Sensors (Basel), 2023. **23**(19).
17. Formosa, N., et al., *A New Modeling Approach for Predicting Vehicle-Based Safety Threats.* IEEE Transactions on Intelligent Transportation Systems, 2022. **23**(10): p. 18175-18185.
18. Morando, M.M., et al., *Studying the Safety Impact of Autonomous Vehicles Using Simulation-Based Surrogate Safety Measures.* Journal of Advanced Transportation, 2018. **2018**: p. 6135183.
19. Hula, A., et al., *Using reaction times and accident statistics for safety impact prediction of automated vehicles on road safety of vulnerable road users.* Safety Science, 2023. **162**: p. 106091.

20. Koopman, P., *How Safe Is Safe Enough? Measuring and Predicting Autonomous Vehicle Safety,* 2022: Independently Published. 368. ISBN: 9798846251243
21. Blumenthal, M.S., et al., *Safe Enough: Approaches to Assessing Acceptable Safety for Automated Vehicles*. 2020, Santa Monica, CA: RAND Corporation.
22. Fraade-Blanar, L., et al., *Measuring Automated Vehicle Safety: Forging a Framework*. 2018, Santa Monica, CA: RAND Corporation.
23. Campbell, K.L., *The SHRP2 Naturalistic Driving Study: Addressing Driver Performance and Behavior in Traffic Safety*, in *Transportation Research Board*. 2012, Transportation Research Board of the National Academies.
24. *The Canada Naturalistic Driving Study (CNDS)*, in *Council of Deputy Ministers Responsible for Transportation and Highway Safety*. 2015: University of Saskatchewan.
25. Wang, C., et al., *A review of surrogate safety measures and their applications in connected and automated vehicles safety modeling.* Accident Analysis & Prevention, 2021. **157**: p. 106157.
26. Brown, T.L., *Adjusted Minimum Time-to-Collision (TTC): A Robust Approach to Evaluating Crash Scenarios*, in *DSC 2005 North America*. 2005: Orlando.
27. AAA, *Automatic Emergency Braking Performance in the Context of Common Crash Scenarios*, in *American Automobile Association*. 2022.