Collision Avoidance Effectiveness of an Automated Driving System Using a Human Driver Behavior Reference Model in Reconstructed Fatal Collisions

John M. Scanlon, Kristofer D. Kusano, Johan Engström, Trent Victor

Waymo, LLC

Abstract

Avoiding and mitigating any potential collision is dependent on (1) road user ability to avoid entering into a conflict, the (*conflict avoidance effect*), and (2) road user response should a conflict be entered, the (*collision avoidance effect*). This study examined the *collision avoidance effect* of an automated driving system (ADS) using a human behavior reference model. The potential benefits of ADS technology have been widely anticipated, and reliable performance benchmarking methodologies for assessing ADS performance is an essential component of determining system readiness. The reference model used in this study reflects the response time and evasive action (referred to as collision avoidance) of a human driver that is **n**on-**i**mpaired, with **e**yes **on** the conflict (NIEON) - a consistently performing, always-attentive driver that does not exist in the human population. Notably, the NIEON model is a tool for evaluating *collision avoidance* only, as it inherits the pre-conflict behaviors (includes conflict avoidance) of the ADS that is being evaluated. In this way, the testing regime isolates the analysis to focus only on the *collision avoidance effect*.

Counterfactual simulations of the Waymo Driver and the NIEON model were run on potentially avoidable (excludes rear-end struck) responder collision scenarios (ADS is responding to some potential collision partners sudden, unexpected actions). All of these simulated events were reconstructed fatal crashes that occurred during a 10-year period in the Operational Design Domain of the Waymo ADS in Chandler, Arizona (Scanlon et al., 2021). In this area, users of the commercial Waymo One service can hail an ADS without a human in the driver's seat.

A conflict must be entered to examine the collision avoidance effect using the NIEON model (or else there is no potential collision to examine performance). 16 total conflicts were identified for this study from a prior analysis by Scanlon et al. (2021). Regarding the *collision avoidance effect*, of 16 conflicts entered, 12 (75%) were prevented by the Waymo Driver, and 10 (62.5%) were prevented by the NIEON model. The NIEON Model mitigated an additional 5 collisions and did not mitigate 1 collision. In these 16 conflicts entered, 93% of serious injury risk was reduced by the Waymo Driver, whereas 84% of serious injury risk was reduced by the NIEON model. Further, in a case-by-case evaluation, the Waymo Driver's collision avoidance led to reduced serious injury risk when compared to the NIEON model in every simulated event.

The results of this paper demonstrate that a reference model like NIEON can be used to benchmark ADS responder performance in response to high-risk initiating behaviors performed by the current driving population. Outperforming established benchmarks, like the Waymo Driver did versus the NIEON model in these simulations, could be used to demonstrate that fatal and serious-injury crashes will be reduced. The analysis also helps demonstrate the safety benefits of an ideal state of human driving in Chandler, Arizona, where many of these fatal collision outcomes would have been preventable.

1. Introduction

1.1 Counterfactual Simulation in Autonomous Driving Verification

There were over 6.7 million police-reported collisions in the United States in 2019 (NCSA, 2021). Although collisions regularly occur across the population, they are rarely experienced in an individual's lifetime. Police-reported collisions in the United States have been estimated to occur about once per every one-half million driving miles (Blanco et al., 2016). Of those police-reported collisions, an injury occurred approximately 28% of the time and at least one fatality occurred in 0.05% of those collisions.

Collision and injury rates have historically been used to provide empirical evidence on the efficacy of a potential technology after the technology is deployed widely in the vehicle fleet. For example, Blower and Woodrooffe (2013) examined a trucking carrier's loss data, containing over 112,000 records, to demonstrate the efficacy of heavy vehicle rollover stability control. Reixinger et al. (2019) used a quasi-induced exposure model on over 1.3 million passenger vehicle rollover crashes in the United States to demonstrate the efficacy of electronic stability control (ESC) in reducing rollover propensity. Cicchino (2017) and Isaksson-Helma & Lindman (2015) each used insurance loss data to demonstrate the efficacy of forward collision warning (FCW) and automatic emergency braking (AEB) data in reducing passenger vehicle front-to-rear collision rates.

The relatively low rate of collisions resulting in serious injury or fatalities compared to those collisions of lesser severities restricts evaluators from making definitive safety claims about Automated Driving Systems (ADS; specifically referring to SAE levels 3 through 5 technologies) (SAE J3016, 2021)[1], which have not yet been widely deployed. ADS technology, including the Waymo Driver, is operating in limited deployments on public roads. Waymo has previously shared details about collisions that occurred during the Waymo One autonomous operations (Schwall et. al, 2020). This crash data from calendar year 2019 included over 6.1 million miles of autonomous driving and resulted in 47 contact events consisting of 18 actual and 29 simulated contact events. Each of these actual and simulated events had low injury risk. Because serious injury and fatal collisions so rarely occur, and therefore require extreme mileage for statistical power, prospective methodologies are required to predict safety performance for an ADS that is not yet widely deployed.

One technique to generate a prospective benefit estimate is to rely on simulation of reconstructed collisions to evaluate collision and injury prevention capabilities. Previously, the authors published a simulation study that replaced the human drivers with Waymo's ADS for all applicable (see section 3.1.1 for more details) fatal collisions in its Operational Design Domain (ODD) in Chandler, AZ (the location of Waymo's commercial ADS ride-hailing service, Waymo One) (Scanlon et al., 2021). This previous simulation study utilized crash reconstructions of the original human-driver-only events and compared the outcome of the actual collisions with the simulated performance of the Waymo ADS. The results were promising: The ADS was able to avoid 100% of the collisions when replacing the initiator role (the conflict partner that performed an unexpected maneuver that resulted in a conflict) and 82% of collisions when replacing the responder role.

Scanlon et al. (2021) and other prospective simulation studies that used crashes as their basis (Sander & Lubbe, 2018, Kusano & Gabler, 2012; Rexinger et al., 2019; Haus et al., 2019; Scanlon et al., 2016; Scanlon et al., 2017) generally analyzed to what extent a simulated technology was able to prevent or mitigate the actual human crashes. This analysis can be useful for an estimate of potential safety benefits if a technology were to be widely deployed (Najm & daSilva, 2000). This type of analysis, among others, is covered extensively in ISO T/R 21934-1, which

[1] The hardware and software that are collectively capable of performing the entire DDT on a sustained basis, regardless of whether it is limited to a specific operational design domain (ODD); this term is used specifically to describe a Level 3, 4, or 5 driving automation system.

details “prospective assessment of traffic safety for vehicle-integrated technologies acting in the pre-crash phase by means of virtual simulation” (ISO, 2021a). The ISO document focuses on active safety and advanced driver assistance systems (ADAS), but many of the principles for relying on scenario-based testing for prospective assessment can be applied - at least in part - to level 3-5 ADS technology. Beyond estimating the prospective safety benefits, it is useful to relate these benefits to well-defined benchmarks. Such benchmarks can be defined in terms of *behavior reference models* against which ADS performance and associated estimated safety benefits can be compared.

1.2 Benchmarking ADS Safety Performance

When considering ADS-equipped vehicles that will perform the entire driving task without the expectation that a driver will respond to a request to intervene (SAE level 4-5), there are two main approaches to assessment of the safety of a system: (1) a positive risk balance (PRB) at the aggregate system performance level and (2) the absence of unreasonable risk (AUR) at the event (or hazard) level (Favarò, 2021). There are opportunities to apply PRB at an event-level, but its utility as a tool for measuring aggregate performance (namely in safety benefits analyses) is more widely accepted. When defining the acceptance criteria for these two levels of assessment, data from the performance of the current human driving population is often referenced to define acceptance criteria. These acceptance criteria can often be from different driving populations, and the relationship between AUR benchmarks and the overall system performance assessed by PRB is not clear (Favarò, 2021).

For aggregate system performance benchmarks for PRB, it is difficult to translate human data (which often includes undesirable behaviors such as reckless driving, or driving while intoxicated, drowsy, or distracted) into a desired ADS safety acceptance criteria. Goodall (2021) proposed a method to derive benchmark crash rates under nominal (or model) driving that is free from risky behavior from naturalistic driving data. The assumption that preventing a certain risky behavior (e.g., interacting with a cell phone) will lead to this model driving in human drivers is potentially flawed, due to human driver behavioral adaptation (Wege et al., 2013), whereby the propensity to do other activities needs to be accounted for. Flannagan et al. (2019) used naturalistic driving data to create propensity models to estimate weighted odds ratios to estimate the reduction in crash risk due to the absence of cell phone use. They found that drivers that were not using cell phones were more likely to engage in other distracting activities, and thus the realized crash reduction was lower than comparing cell phone use to model driving.

In parallel or in combination with PRB, event level evaluations of AUR can also draw upon human performance data to set acceptance criteria. For example, a behavior reference model provides a repeatable framework for examining how some defined actor may behave if it were to encounter some scenario. This study presents an implementation of a reference behavior model for response time and evasive maneuvers (i.e., braking and steering) in collision avoidance situations that is based on the performance of a Non-Impaired Eyes ON the conflict (NIEON) human driver. By applying this reference behavior model to a set of reconstructed real-world fatal collisions, in the way of counterfactual simulations discussed in the previous section, this study is the first known to the authors to investigate the relationship between an event-level acceptance criteria defined as a reference behavior model and estimated safety benefits of an ADS on a set of existing human crashes. A similar analysis was described in Rothoff et al. (2019) in a feasibility study, where four scenarios (Same direction, rear-end frontal; Rear end, Rear; Car-to-pedestrian SCP, pedestrian from right; Opponent vehicle lane change) were simulated with an ADS and compared against “an attentive, skilled and experienced Reference Driver.”

1.3 Avoiding the Need for Collision Avoidance Action

Reference behavior models are designed to function within limited driving contexts. This study focuses on a model designed for representing collision avoidance actions of a non-impaired human with their eyes on the forward roadway within multiple potential collision types. Accordingly, it is important to consider the contexts for which a driver may find themselves in a scenario requiring some level of collision avoidance action.

Consider the potential ADS collision space. There are those collisions that involve the subject vehicle in the initiator role, i.e., where the subject vehicle's driver takes some initial surprising action. Surprise in this context refers to the deviation of expectations from outcomes for a given road user's behavior (Engström et al., 2022; Itti and Baldi, 2009). Specifically, the initiator road user behaved unexpectedly (e.g., failed to yield for relevant traffic signals or failed to maintain lane positioning), which led to the emergence of a potential collision event. There are also scenarios where the driver is in a responder role, whereby they must respond to some initiator's actions. In both initiator and responder collisions, human and ADS capabilities are dictated by conflict avoidance and collision avoidance evasive maneuvering. Conflict avoidance refers to the proactive actions that a driver may take to avoid getting into conflicts in the first place and thus mitigate or prevent potential collision and injury risk (e.g., keeping a sufficient following distance, managing occlusions, adhering to the posted speed limit, following traffic laws). If a conflict does occur, collision avoidance action is required. In this current study, collision avoidance is referring to "any action performed by any conflict partner to change its trajectory or speed in an attempt to avoid or reduce the severity of a potential crash, avoid or reduce the severity of a road departure, or regain vehicular control after a loss of control" (ISO TR21974-1, 2018). Specifically, collision avoidance refers to the urgent, high acceleration magnitude evasive action required immediately before an impending potential collision event (e.g., braking and steering action). A breakdown of the differentiation can be found in Figure 1 below.

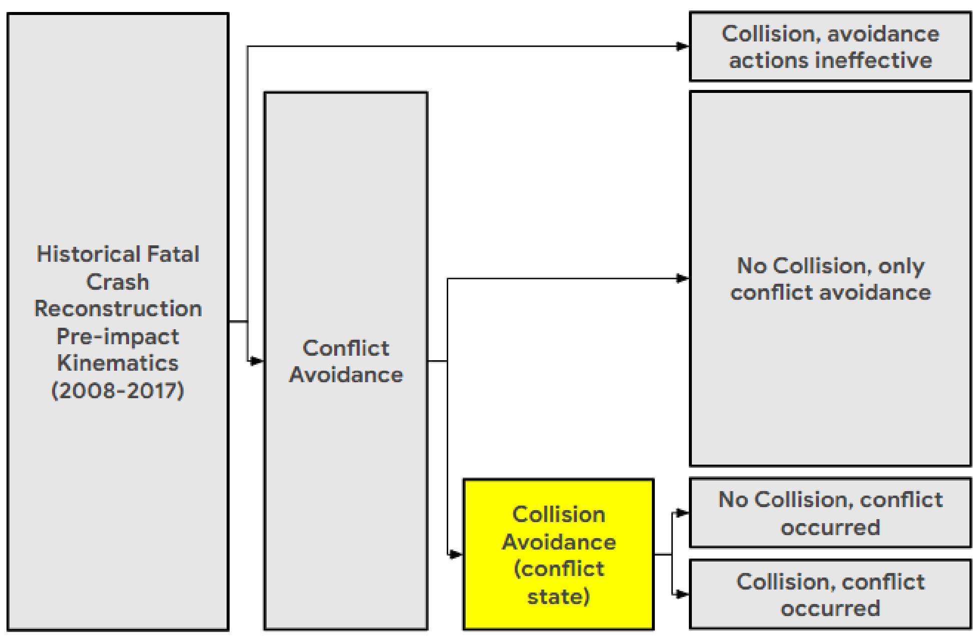

Figure 1: A depiction of the relationship of conflict avoidance and collision avoidance in mitigating and preventing real-world collisions. This study focused on evaluating ADS collision avoidance performance (yellow highlighted box). The uppermost path in this depiction ("Collision, avoidance actions ineffective") is intended to capture crash types where neither conflict nor collision avoidance have any effect on the potential severity outcome. In the current study, all of the ADV simulated events labeled as "Collision, avoidance actions ineffective" were in a lead vehicle position when the vehicle was struck in a front-to-rear crash type.

As a potential initiator, avoidance of a conflict or crash is dictated first by the capacity to not perform said initial, unexpected action. These initiating actions can take a variety of context-specific forms. An initiating vehicle may have violated some traffic control device, stopped unexpectedly in the roadway, failed to yield the right of way, or failed to maintain control of the vehicle. The most effective drivers who do not initiate conflicts or crashes are, implicitly, those that do not perform these precursor initiator actions. More simply, the most effective drivers behave as other road users expect them to behave.

In responder collisions the subject vehicle is required to take some responsive actions following some initial, surprising action by another actor (Engström et al., 2022). The evolution of a responder role event varies by the scenario. A scenario represents a unique set of conditions - or in the case of this counterfactual simulation study - a set of reconstructed actor behaviors. Some scenarios emerge abruptly with little to no notice - e.g., a red light runner emerges from a perpendicular road into the path of a straight traveling ADS with a green light. In these scenarios, a driver may be required to take collision avoidance actions to mitigate or prevent the collision. Other scenarios are more influenced by conflict avoidance. For example, a vehicle starts to inch out into an intersection to make an unprotected left turn, and the ADS proactively slows in anticipation of a potential crossing action. Still, there remains those collisions where neither conflict nor collision avoidance are effective at mitigating or preventing a collision. Consider, for example, the common collision mode where some subject vehicle is stopped at a queue of vehicles when it is struck by some trailing vehicle.

The most effective drivers in the responder roles have several complementary traits. First, they are fully aware of their surroundings and actively anticipate the potential for a collision scenario to emerge. Second, they actively execute conflict avoidance to limit the opportunity for conflicts to emerge. Third, when a conflict does emerge, these drivers are effective at collision avoidance through timely actuation of driving controls to mitigate and prevent potential harm. This study focuses on the use of a behavior reference model to examine the performance of this third trait: how effective is the Waymo Driver relative to the reference model at collision avoidance compared to this reference behavior model? As will be discussed, to independently examine the collision avoidance effect, both the Waymo Driver and the reference model receive the Waymo Driver conflict avoidance, i.e., the pre-conflict driving behavior of the Waymo Driver, such as travel speeds leading up to the conflict.

1.4 Reference Models of Collision Avoidance Behavior

Human driving data is one source for developing reference behavior models. Several reference models have been proposed in the literature. The Responsibility Sensitive Safety (RSS) model presents a mathematical formulation of boundary or envelope conditions that an ADS should adhere to (Shalev-Shwartz et al., 2017). The RSS model is a set of constraints for a road user to follow in order to avoid entering into a scenario where a collision is unavoidable. RSS can be considered a conflict avoidance modeling approach, whereby adhering fully to these sets of guidelines will enable the vehicle to (a) not initiate collision scenarios and (b) proactively respond to some scenarios where conflict avoidance is effective at avoiding or mitigating potential collisions in the set of scenarios covered by the model.

Another class of reference models to determine the avoidability of collisions rely on simplified perception, cognition, and action models of human behavior. For example, consider a rear-end brake avoidance model for a human driver in response to some lead vehicle. A brake avoidance maneuver model could assume some response time (or distribution of response times) followed by a braking response dictated by some brake application model (Kusano & Gabler, 2012). Such fixed response time models are also used in regulation as a means for specifying the performance requirements for highway automated lane keeping system (ALKS) (UN/ECE 2021). One drawback of such reaction time and response models is that (a) reaction time is highly dependent on the nature of the events with various factors, including vehicle kinematics and looming of an approaching vehicle, (b) braking jerk is also dependent on kinematic urgency at braking onset (Markkula et al., 2016), and (c) determining the stimulus onset time, that is, when to start measuring the reaction time, varies by scenario. Due to these scenario-specific factors that influence response time, a fixed response or a distribution of response times observed from real-world events may not accurately predict response time, especially when applied to scenarios that have different kinematic properties (see Engström et al., 2022, for further discussion).

To address the shortcomings of fixed response and reaction models, Waymo has developed a framework for modeling road user response timing in naturalistic traffic conflicts (Engström et al., 2022). This model represents human avoidance responses taking into account the dependence of response time on the situation at hand. The model can be fit to naturalistic driving crashes and near-crashes representing the targeted reference population, for example human drivers with their eyes on the forward roadway during the onset of the conflict (Engström et al., 2022). This study utilized this response timing model alongside an evasive maneuver model to represent braking and steering avoidance performance of a non-impaired and attentive driver in response to an imminent collision scenario, as further described in section 2.4.

In this paper, we apply a reference human collision avoidance behavior model to the collision imminent situations reconstructed from fatal collisions in a given ADS ODD. This evaluation was performed on the fatal collision dataset from Scanlon et al. (2021). Our analysis focused on the simulated cases from this study where the Waymo Driver entered into a conflict, where urgent, evasive maneuvering (collision avoidance) was required to prevent or mitigate a potential collision. Within these conflict events, the goal was to compare the simulated ADS crash prevention and mitigation performance in the reconstructed scenarios to a reference model that represents response and evasive maneuvering performance representative of a non-impaired human with their eyes directed toward the conflict.

As discussed previously, high severity collisions like the ones examined in this study are rare and often have multiple contributing factors that lead to such an outcome. In this dataset, the police reported that many of the involved drivers were impaired and it is possible some of these drivers were drowsy or distracted, which would not be captured in the police accident report. As an alternative to comparing against the original human driver that may have been distracted or impaired, this study introduces the NIEON model, a benchmark representing a non-impaired, eyes on conflict human response. This NIEON model directly incorporates the novel response time modeling presented by Engström et al. (2022). The performance of this benchmark is juxtaposed against the performance of the Waymo Driver and the original human driving outcome. The NIEON model represents a repeatable modeling approach that notably also enables performance evaluations to be performed on collision scenarios that were artificially created, e.g., during scenario-based testing, with the goal to establish a collision avoidance benchmark against which the performance of the ADS can be compared.

1.5 Project Scope and Research Question

The current study extends prior work by Scanlon et al. (2021), where the Waymo Driver's ability to respond to, and not initiate, a set of real-world fatal crash scenarios was evaluated. The current study implements a reference behavior model of collision avoidance performance on this same representative set of fatal collisions. The key contribution of the paper is to demonstrate how such a behavior reference model can be used to establish a benchmark for crash prevention and mitigation performance in scenario-based testing that the performance of an ADS can be compared against. In this analysis, we specifically investigated how the Waymo Driver's collision avoidance performance compared to a reference model representing a non-impaired human driver with their eyes on the conflict in the reconstructed fatal collision scenarios. This focused research evaluates only how the Waymo Driver behaves when a conflict has been encountered (collision avoidance effect), whereas overall injury prevention and mitigation performance also has dependencies on the driver's ability to avoid conflicts altogether (conflict avoidance effect).

2. Methodology

2.1 Data Source

This study relied upon a dataset of identified real-world collisions that occurred in Chandler, Arizona. All collisions were captured within the Arizona Department of Transportation's (ADOT) publicly accessible crash database. Chandler, Arizona is a part of the commercial Waymo One service, where users can hail an ADS without a human behind the wheel. All fatal collisions that occurred in Chandler, Arizona from 2008 to 2017 (a 10-year period) were identified (a total of 107 events). This dataset is identical to that previously used in Scanlon et al. (2021).

2.1.1 Reconstruction Materials

This study required reconstructed collision sequences be generated for all identified events. To complete this exercise, materials required for reconstruction were requested for identified cases to the Arizona Department of Transportation, Arizona Department of Public Safety, and the Chandler Police Department. A complete list of requested materials can be found below:

- Police report(s)
- Scene diagrams
- Photographs
- Witness statements
- Event data recorder (EDR) reports
- Other miscellaneous reconstruction-relevant materials (e.g. surveillance footage)

As noted in Scanlon et al. (2021), every case varied in the amount of available materials for reconstruction. Downstream, this affects the scope of any derived conclusions. The impact of this material constraint, and the utility of a reference model in addressing this limitation, are discussed in the Discussion section.

2.1.2 Inclusion Criteria

The inclusion criteria for this study was identical to that performed in Scanlon et al. (2021). First, only the first two collision partners in the collision sequence were considered. Specifically, the first collision event was considered, and the total partners in that event (one or two) was considered. Second, class-1 and class-2 vehicles not towing any vehicles or objects were exclusively considered in this study to be approximately consistent with dynamics considerations of the Waymo Chrysler Pacifica vehicle platform. Third, a valid map for operating the Waymo must have been available. A valid map is required to have the vehicle navigate on that portion of the roadway. This requirement was three-part. The first map requirement was that the road configuration at the time of the collision must represent the current mapped configuration. Any modification from the map was excluded, for example, if a roadway had major construction that changed the number of lanes or intersection configuration. The second map requirement was that the Waymo ADS operate only within a drivable location (i.e., within its ODD). This requirement would exclude areas like private roads and driveways inaccessible to the Waymo Driver. The requirement also entailed that the subject vehicle being replaced could not be operating on a roadway with a speed limit greater than 45 mph, given that those locations would be outside of the Waymo Driver ODD in that region. In addition, crashes that occurred on highway on-ramps or off-ramps were excluded as the current Waymo One service does not travel on limited access highways.

2.1.3 Initiator and Responder Roles

Scanlon et al. (2021) introduced a novel classification scheme for identifying the role of actors in a collision sequence. This methodology designates the “initiator” and “responder” roles, which represent who performed the initial surprising action that led to the collision events and who was responding to those initiator actions. This concept is orthogonal to other role designation schemes, such as responsibility, right-of-way, and fault assignments, and serves as a technique for differentiating the type of action required to avoid or mitigate a potential collision. For example, in the initiator role, preventing associated collisions requires that the actor not perform the set of surprising actions that were performed by the original actor, which would create the same conflict state. If the actor fails to prevent this conflict state, evasive maneuvering must be performed in order to prevent or mitigate a potential collision. Similarly for the responder role, evaluating performance within this role indicates the capacity to perceive and respond in a timely fashion to some potential conflict state created by an initiator.

Occasionally, there was uncertainty in the associated police reporting as to who was the initiator or responder (e.g., due to conflicting reporting by witnesses). Our approach in the current study was to model the Waymo Driver as the responder in all simulated events where it was not explicitly clear whether the replacing vehicle was in an initiator or responder role. This notably only occurred in this study in situations that involved a red light runner, where it was unclear which actor violated the traffic control device. This study took the approach of modeling all these vehicles as responding to the other vehicle running the red light runner.

2.2 Collision Reconstruction

The reconstructed methods relied upon in the current analysis are detailed in Scanlon et al. (2021). For more details please refer to this prior publication. A third-party engineering firm reconstructed each collision using the requested materials. Trained experts reconstructed the pre-crash events of each case using a variety of techniques, each of which reflects industry best practices, according to the available materials. When the reconstructions were performed, the trained experts were not aware of the intended use case or identity of the simulating party (Waymo).

On a high-level, the following reconstruction materials were required to create simulations:

1) Pre-crash kinematics of involved actors,
2) Pre-crash kinematics of actors and objects influencing the causation or severity of the collision sequence,
3) Actor and object dimensions and inertial attributes,
4) Traffic signal phase timings, and
5) A driving environment (generated from scaled, orthonormal aerial images) that could be transposed to the Waymo driving environment (the map).

2.3 Waymo Simulation

The same alignment procedure used in the previous study, Scanlon et al. (2021), was used in the current study. To briefly summarize, this alignment procedure took the human involved crash reconstructions, replaced one of the actors with the Waymo Driver, then aligned the trajectories such that the original reconstructed party and the Waymo Driver simulation were synchronized in their approach to the potential collision configuration. The methodologies enabled the Waymo Driver to drive in a manner consistent with its programmed behavior throughout its approach to the potential collision, i.e., the Waymo Driver was not forced to speed or disobey traffic controls. This alignment procedure enabled a comparison to the original actor, and a representation of what the Waymo Driver would have done had it organically experienced the scenario during normal operations.

Waymo's simulation platform was utilized to simulate the Waymo performance and that of the reference human behavior model. As discussed in our previous study (Scanlon et al., 2021), Waymo's simulation platform aims to recreate the actual driving environment through simulation of the sensors (including noise and latency), perception system, and planner/behavior systems. The system software and map data used in this study were the same versions that are being used in the commercial Waymo One service at the time of writing.

Waymo's simulation platform creates a driving environment where the Waymo Driver can be evaluated. This virtual platform is intended to recreate the performance of the system as if it was encountering the scenario in real-world conditions. As previously described, there are several key features of this platform that enable this recreation. First, a three-dimensional driving environment is used, which is referred to as the map. The map captures environmental conditions, such as road geometry or sight obstruction, that can influence driving capabilities. Second, the simulator recreates the vehicle's on-board sensing capacity. This sensor simulation captures nuances of sensor mechanics, such as range, field of view, sweeping characteristics, and measurement noise that may contribute to uncertainty downstream processes such as object classification. Additionally, the simulation approach models interpretation and processing of those sensor measurements. These sensing simulations were developed and validated to recreate measured performance from an actual vehicle in the virtual world. Third, the onboard behavior level is being simulated. This logic considers the environment and perception feedback, and plans the path of the vehicle accordingly.

Latency in the system is considered on top of the aforementioned simulation features. A number of modules, including perception-related and planning-related, have potential latency, which cause a delay in any action taken by the platform in real-world conditions. In the current study, our approach was to create simulations using the 95th percentile execution time of every simulated module. The 95th percentile latency values are unlikely to occur in all modules concurrently. Accordingly, the end-to-end latency of the simulated platform can be considered to be conservative.

2.4 Reference Model Simulation

This study presents an implementation of a behavior reference model to evaluate the Waymo's collision avoidance performance within this reconstructed Chandler dataset. On a high level, this reference model (NIEON) aims to represent the response of a **n**on-**i**mpaired, with **e**yes always **on** the conflict human driver that is presented with some critical situation that can be avoided or mitigated through either braking and/or steering action. In this current study, the NIEON model inherits the pre-conflict movements of the ADS (including conflict avoidance) and independently models collision avoidance actions. In this way, the NIEON model only provides a benchmark for the collision avoidance performance effect (i.e., how well did the Waymo Driver respond when a conflict was encountered?). It is also important to note that, because the pre-conflict movements of the Waymo Driver are being adopted, the same alignment procedures are followed for the NIEON model as with the Waymo Driver (described in the prior section and more extensively in the prior publication by Scanlon et al., 2022). It is also important to note that this study was not meant to provide a comprehensive technical summary of the NIEON model. Rather, the study focused on the performance of the NIEON model in a series of real-world, representative collision scenarios.

The NIEON model used in the current study leveraged a model of human response timing in response to the perception of the critical threat based on the framework outlined in Engström et al. (2022). In this framework, response timing is conceptualized as a belief update process where an initial (prior) hypothesis is updated to an alternative (posterior) hypothesis through the accumulation of surprising evidence. Specifically, the timing of the evasive action is dictated by (a) the timing of the initial surprising moment and (b) the rate of accumulation of evidence for the need to take evasive action.. Engström et al. (2022) presents a method for fitting a statistical model to predict the response time in a given scenario based on the difference between the stimulus end and onset, called ramp-up time. The annotation of the stimulus onset and stimulus end is based on manual annotation guided by procedural heuristics. The example scenario presented in Engström et al. (2022) was a rear-end conflict scenario. Using similar principles, Waymo has developed a library of scenarios that defines the initial hypothesis (H_i), alternative (surprising) hypothesis (H_a), stimulus onset (T_1), and stimulus end (T_2). Using these definitions, Waymo has used additional naturalistic driving data sources to create statistical models of response time for a wide range of scenarios. The NIEON model uses a single point estimate response time from this statistical model.

In addition to the model of response timing, this study used a model of typical avoidance maneuvers (i.e., steering and braking) magnitudes observed from human naturalistic near crashes. For braking action, the reference model employs a constant jerk magnitude of -30 m/s^3 up to a maximum deceleration of -8.3 m/s^2. These values were derived from fitting a piecewise linear model of acceleration to human crash and near-crash data using a method similar to the one described in Markkula et al. (2016). If the vehicle is traveling up or down a slope, the maximum deceleration is adjusted by the component of gravity perpendicular to the slope of the road.

For the swerving model, a pure pursuit controller was used to control the nominal agent toward a point 4.0 m laterally and 2.0 meters plus 1.0 seconds times the vehicle speed longitudinally along the pre-maneuver trajectory. The controller's change in acceleration and path curvature were constrained by maximum lateral jerk (16 m/s^3), maximum lateral acceleration (6.4 m/s^2), maximum time pinch (the time derivative of the path curvature; 0.35 rad/m/s), and maximum curvature (0.2 rad/m). The maximum lateral jerk and lateral acceleration values were based on human driver avoidance from naturalistic driving near-crash events. The maximum curvature was based on the steering geometry of the Chrysler Pacifica (i.e., the same physical constraints apply for the Waymo driver and NIEON model). Based on this maximum curvature and steering geometry of the vehicle, the pinch time was derived assuming a steering wheel angle rate of 720 degrees per second, which is used in NHTSA's fishhook rollover stability testing for the New Car Assessment Program as an urgent maneuver most drivers can achieve (Forkenbrock et al., 2003; National Highway Traffic Safety Administration 2013). The reference model's swerving maneuver included both steering described here in addition to the braking described in the previous paragraph. Together, the scenario description library, the response time model and the evasive maneuver model create a general purpose collision avoidance reference behavior model that can be applied consistently to all the crash scenarios in the current study.

For each role simulation where urgent avoidance maneuvers were required by the ADS (i.e., steering or braking as described above), three reference model responses were simulated: swerving and braking to the left and right and braking without swerving. Swerving maneuvers that departed the roadway or impacted fixed roadside objects (e.g., signs, raised curbs) were excluded as invalid avoidance maneuvers. Furthermore, swerving avoidance maneuvers were restricted in scenarios where the other vehicle crossed the path of the ego vehicle, for example, another vehicle that turns left across the path from the opposite direction of the ego vehicle. In these crossing path scenarios, the stopping distance of the other vehicle was estimated at the time of the swerving maneuver start, assuming an instantaneous 4 m/s^2 deceleration. The swerving maneuver in the direction opposite to the other vehicle's direction was excluded from consideration if the hypothetical 4 m/s^2 braking at the swerve maneuver start time would have resulted in the other vehicle stopping in the path of the ADS's trajectory. In the left turn across path example, the restriction would exclude the left swerving maneuver reference behavior model if the other agent's stopping distance met the criteria above. This restriction was introduced to account for the well-documented phenomenon that, in human crashes and near-crashes, the driver in the responder role, in crossing path conflicts, usually swerve in the same direction as the crossing vehicle (Hu and Li, 2017; Li et al., 2019; Malaterre et al., 1988; Mazzae et al, 1999; Scanlon, 2017; Weber et al., 2015). However, drivers also swerve in the opposite direction relatively often and the decision in which direction to swerve seems to be strongly scenario-dependent (Li et al., 2019; Hu and Li., 2017; Scanlon, 2017). The rationale behind this exclusion criteria is the assumption that human drivers are less likely to swerve to the side opposite of the other vehicle's direction if there is some uncertainty whether the crossing vehicle will come to a stop in the ADS's path. For the scenarios considered in this study that required urgent collision avoidance maneuvers, this swerving constraint was not present, and therefore did not affect the results of the current study.

2.5 Event Severity Assessment

Every original and simulated collision outcome had an event severity assessment performed. Three levels of injury prevention potential were evaluated: Waymo conflict avoidance only, the reference behavior model, and the Waymo Driver. For this study, the "Waymo conflict avoidance only" measures the safety benefits due to fully avoiding collisions using conflict avoidance and does not measure the added severity mitigation effect of conflict avoidance when collision avoidance would have been also required. As indicated earlier, the reference behavior model simulations only simulates collision avoidance behavior (the urgent evasive action before a collision), so any conflict avoidance action taken by the Waymo is also awarded to the reference behavior model.

The goal of this assessment was to evaluate serious injury risk for persons potentially involved. Serious injury designation at the occupant-level and associated risk calculations, p(MAIS3+), were computed according to the Abbreviated Injury Scale (AAAM, 2015). Only the first potential collision event was considered. A dose-response model was used to estimate the count of serious injury events (Kullgren, 2008). Our approach was to consider the maximum serious injury risk across all modeled occupants. For the current study, vehicles with potential for multiple occupancy (e.g., passenger vehicles, heavy vehicles, motorcycles, etc.) were modeled as being occupied by a single person. This single person modeling technique was performed by considering that the occupant could have been in either the driver or right front passenger seat, and the maximum of those two risks was extracted. This was done for the original reconstruction, the Waymo Driver simulation, and the reference model simulation to control for any potential biases due to occupancy. The maximum risk at the collision level was then aggregated to form a single estimate, n_{MAIS3+}, which is representative of overall serious injury risk across events. The formulation of this metric can be found below.

$$n_{MAIS3+} = \sum_{i=1}^{\#\,cases} max(p(MAIS3\ +))$$

The injury benefits estimate was then calculated using the following equation:

$$\text{Injury Benefits} = (1 - \frac{n_{MAIS3+}\ in\ simulated\ events}{n_{MAIS3+}\ in\ actual\ fatal\ crashes})\%$$

This "Injury Benefits" metric is intended to be an estimation of the percentage of serious injury events that would be achieved if the Waymo Driver or reference behavior model had experienced the original collision scenarios. It should be noted that an absolute estimation of injuries occurring before and after is not being modeled here. To achieve this, the approach would need to consider occupancy and a number of occupant-specific characteristics that are not necessarily readily available or easily modeled. Rather, this simplified modeling approach allowed us to approximate injury prevention benefits given the outcomes of the counterfactual simulations.

For vehicle collisions, the Kudlich-Slibar impact model was used to assess collision dynamics (Brach et al., 2011; Kudlich, 1966). The vehicle inertial parameters of the ego vehicle in both the original and counterfactual simulation were always taken to be the 2017-present Chrysler Pacifica Minivan. This method allowed behavioral competencies of the Waymo Driver to be assessed independently of any inertial considerations. Two collision dynamics parameters were extracted to evaluate injury risk in vehicle collisions: delta-V and the principal direction of force (PDOF). The injury risk model, a continuous, omnidirectional injury risk curve model developed by McMurray et al. (2021), was then used to assign the probability of a serious injury, p(MAIS 3+), given the computed delta-v and PDOF.

For motorcycle, pedestrian, and cyclists severity estimation, collision kinematics were assessed at the point of impact. Models developed by Nie et al. (2013) were used to assess serious injury risk in pedestrian and cyclist events. For motorcycle events, serious injury risk curves developed by Tominaga et al. (2002) were used.

3. Results

3.1 Review of Prior Study

3.1.1 Simulation Case Set

The Waymo simulation results shown in the current study have been previously published (Scanlon et al., 2021). This section will provide a high-level overview of those prior results. For a more detailed assessment, please refer to the prior study.

As noted above, all fatal collisions (107) from 2008 to 2017 had public records requested. Materials were received for 92 of those requests. One of the fatal cases was not reconstructed due to substantial changes to the road following the time of the crash.

Reconstructions were performed for all 91 fatal cases. From these collisions, 166 collision partners (90 initiators and 76 responders) were identified as the initial collision partners that would be potentially eligible for simulation in the current study.

As noted, several inclusion criteria were established for case inclusion. First, an important inclusion criteria was that the simulated Waymo Driver should only replace class-1 or class-2 vehicles not towing any objects or vehicles. A total of 56 actors were removed due to this criteria. Second, an additional 19 actors were removed due to the aforementioned incompatible map roadgraph criteria or due to operating on roadways with a speed limit greater than 45 mph. The final dataset consisted of 72 crashes with 91 actors (52 initiators and 39 responders).

3.1.2 Waymo Driver Simulated Outcomes

Every simulated perspective from this prior study was categorized by the urgency and preventability of the scenario. As previously shown in Figure 1, the outcomes were dictated by the ADS's ability to prevent or mitigate by conflict avoidance and collision avoidance. The current study determined that the ADS took collision avoidance action if urgent evasive braking (deceleration greater than 0.5 g) or steering action (yaw rate greater than 8.3 degrees/s) was executed, which is consistent with magnitudes previously observed prior to real-world collision events (Scanlon et al., 2015). These outcome designations were as follows:

1) **Rear-end Struck (trailing vehicle is initiator)**: Conflict avoidance and collision avoidance are ineffective at mitigating or preventing these potential collisions, and a collision occurred regardless of any avoidance action. This classification does not include scenarios where the lead vehicle may have been the initiator (e.g., unexpected stopping behavior by the lead vehicle).
2) **No Collision, only conflict avoidance**: A potential collision was prevented by conflict avoidance without any evasive maneuvering action
3) **No collision, conflict occurred:** Prevented with some degree of collision avoidance.
4) **Collision, conflict occurred:** The collision occurred and some degree of collision avoidance was required.

A full breakdown of the Waymo simulated events within the simulation case set by these four designations can be seen in Figure 2. These results are the same as those produced in Scanlon et al. (2021). In total, 100% of initiator cases and 82% of responder cases were avoided in the current simulation case set. The only simulated collisions involving the Waymo Driver were in intersection scenarios.

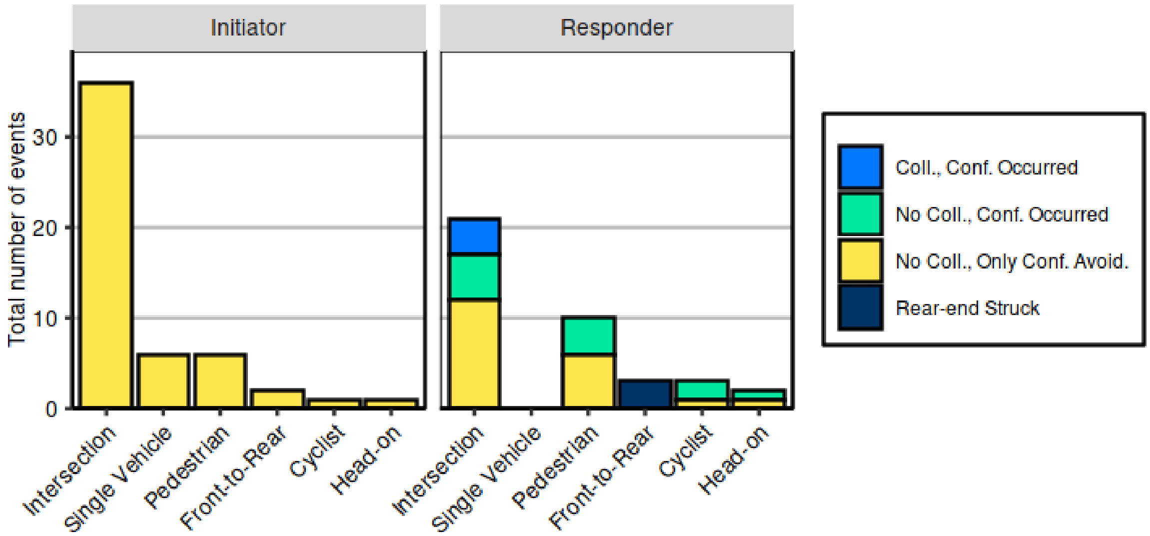

Figure 2: A breakdown of the collision outcomes for all Waymo simulated events.

The majority (79%) of events were avoided with conflict avoidance alone without need for collision avoidance action. All (100%) initiator events were avoided with conflict avoidance alone. This result indicated that, when put into a similar scenario as the original human driver, the simulated Waymo Driver did not perform the initiating behavior that led to a conflict that might require some urgent evasive action. The relevant conflict avoidance behaviors included, but were not limited to, maintaining proper spatiotemporal boundaries, obeying traffic rules, and proper gap selection.

Responder collisions were avoided 51% of the time with conflict avoidance only. Avoidance in these cases are a product of the Waymo Driver perceiving the initiator's behavior early enough that it could take some non-urgent preventive action. A variety of conflict types fall into this category. Some responder collisions (3 scenarios; 8% of responder role) were found to be unaffected by conflict avoidance and collision avoidance. These scenarios were all front-to-rear collisions where some trailing actor struck a stationary or moving lead vehicle (replaced by ADS) intending on traveling straight. The remaining cases (16 scenarios, 41% of responder role) represent those cases in which conflict avoidance alone was insufficient in enabling the Waymo Driver to avoid the collision. The Waymo Driver was able to take collision avoidance (i.e., evasive steering and/or braking action) for all remaining scenarios, and a simulated collision was avoided for all but four of these scenarios.

3.2 Reference Behavior Model Results

The NIEON model results are presented for all potentially avoidable responder role scenarios. These scenarios exclude the three front-to-rear struck events where conflict avoidance and collision avoidance were ineffective at mitigating or preventing the collisions. As discussed, our analysis focused on those scenarios where a conflict was entered and collision avoidance action was required to mitigate or prevent injury risk. As stated in the methodology section, three comparison points were examined: Waymo Conflict Avoidance Only, the NIEON reference behavior model, and the Waymo Driver. The total breakdown of collision prevention and injury prevention benefits can be found in Figure 3 below.

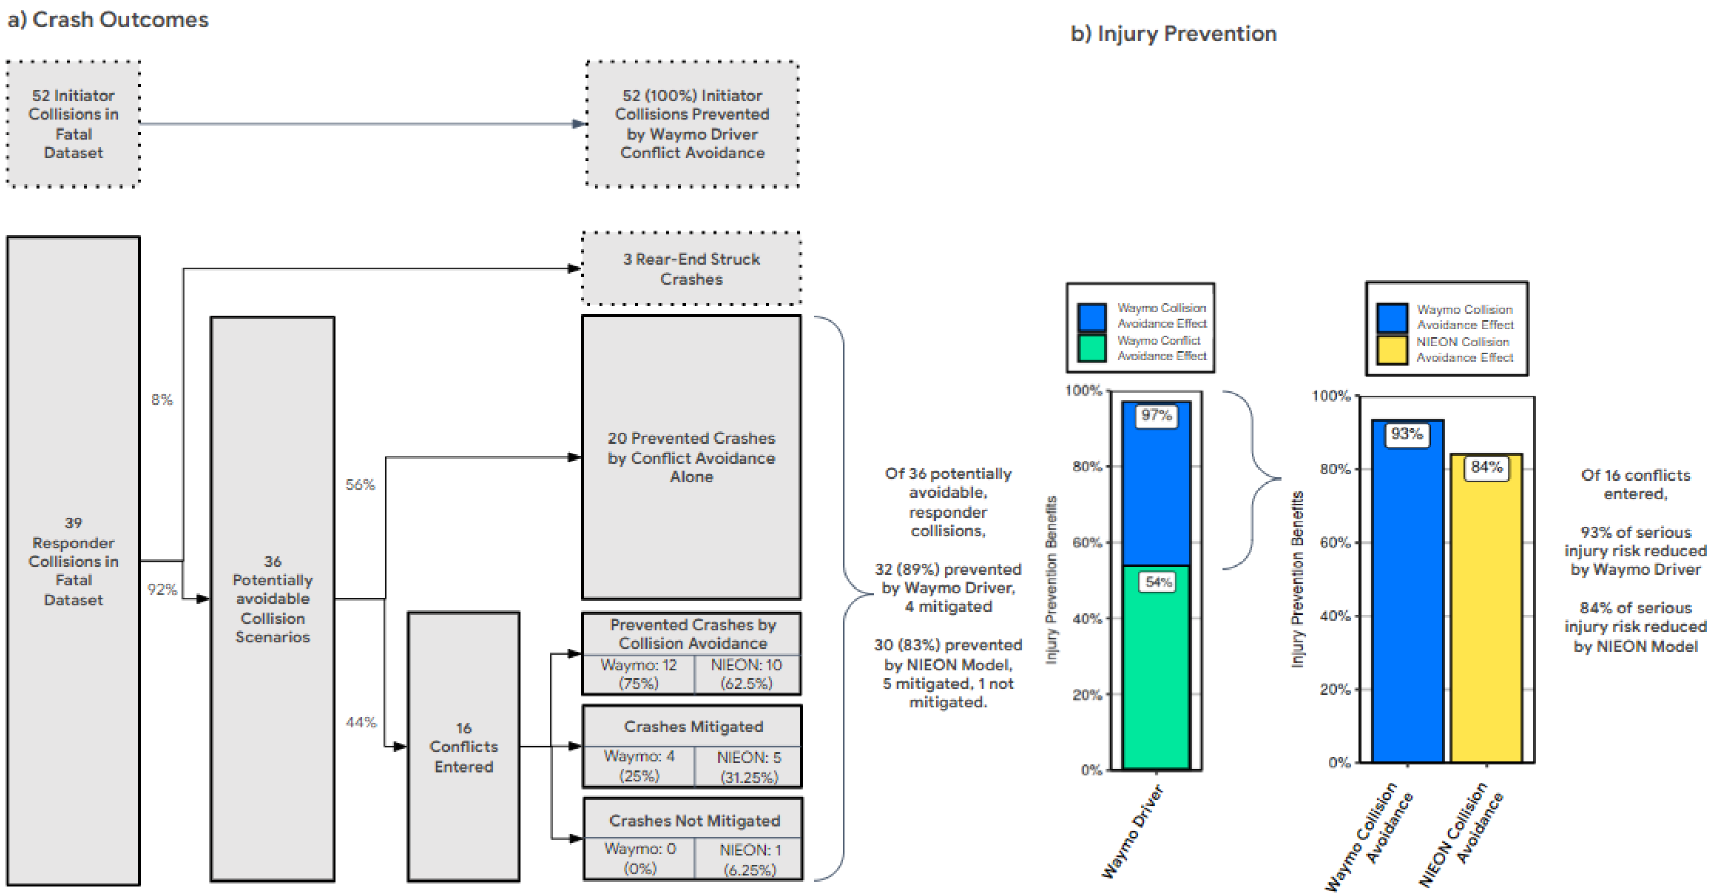

Figure 3: A complete breakdown of (a) crash outcomes and (b) injury risk reduction by both the Waymo Driver and NIEON model.

Several points of comparison are made. First, the effect of conflict avoidance alone is shown. For these avoidable responder collisions, conflict avoidance alone (without urgent evasive action) prevented 20 potential collisions, a crash rate reduction of approximately 56%. Second, the effect of collision avoidance was examined. The NIEON model adopts the conflict avoidance of the Waymo Driver, which includes elements like proactive slowing and adhering to the posted speed limit. The NIEON model then simulates - independent of the Waymo Driver - collision avoidance (observed as evasive action response) to some imminent collision threat. When comparing the NIEON model and Waymo Driver collision avoidance responses, the NIEON model prevented an additional 10 collisions, while the Waymo Driver was found to prevent an additional 12 collisions.

Injury prevention benefits account for additional mitigation benefits that are achieved even if a collision cannot be completely avoided. Rather than relying on a quantification of some actual injury outcome, the injury prevention estimates the potential injury risk reduction for individual collision events. This modeling technique is used because each fatal collision scenario poses some unique injury risk magnitude. One advantage of this modeling strategy for injury risk is that events with higher injury risk potential contribute more to the overall injury prevention estimate.

Of 36 potentially avoidable **responder** collisions, the *combined effect of conflict avoidance and collision avoidance* of the Waymo Driver prevented 32 (89%), and mitigated 4 collisions. The combined effect of conflict avoidance and collision avoidance also led to a serious injury risk reduction for these 36 responder collisions: 97% for the Waymo Driver.

Waymo conflict avoidance was effective - by itself - at preventing 56% of collision events. As stated previously, because the Waymo collision avoidance cannot be disabled, the Waymo Conflict Avoidance Only results assume that there are no mitigation benefits for scenarios where collision avoidance was required. Specifically, if a scenario required collision avoidance, the injury risk for that scenario was modeled as being the same as the original fatal collision. This result shows the importance of conflict avoidance in preventing fatal collisions in this ODD.

Regarding the effect of *collision avoidance only*, of 16 conflicts entered, 12 (75%) were prevented by the Waymo Driver, and 10 (62.5%) were prevented by the NIEON model. The NIEON Model mitigated an additional 5 collisions and did not mitigate 1 collision. In these 16 conflicts entered, 93% of serious injury risk was reduced by the Waymo Driver, whereas 84% of serious injury risk was reduced by the NIEON model.

A breakdown of the results for all potentially avoidable responder collisions (36) by collision type can be found in Figure 4. Of those collisions not prevented by conflict avoidance alone, both the Waymo Driver and the NIEON model were able to avoid all the remaining collisions in the pedestrian, cyclist, and head-on scenarios (7 collisions in these scenarios). Of the 9 intersection scenarios not prevented by conflict avoidance alone, the Waymo driver was able to completely avoid 5 collisions compared to 3 avoided by the NIEON model.

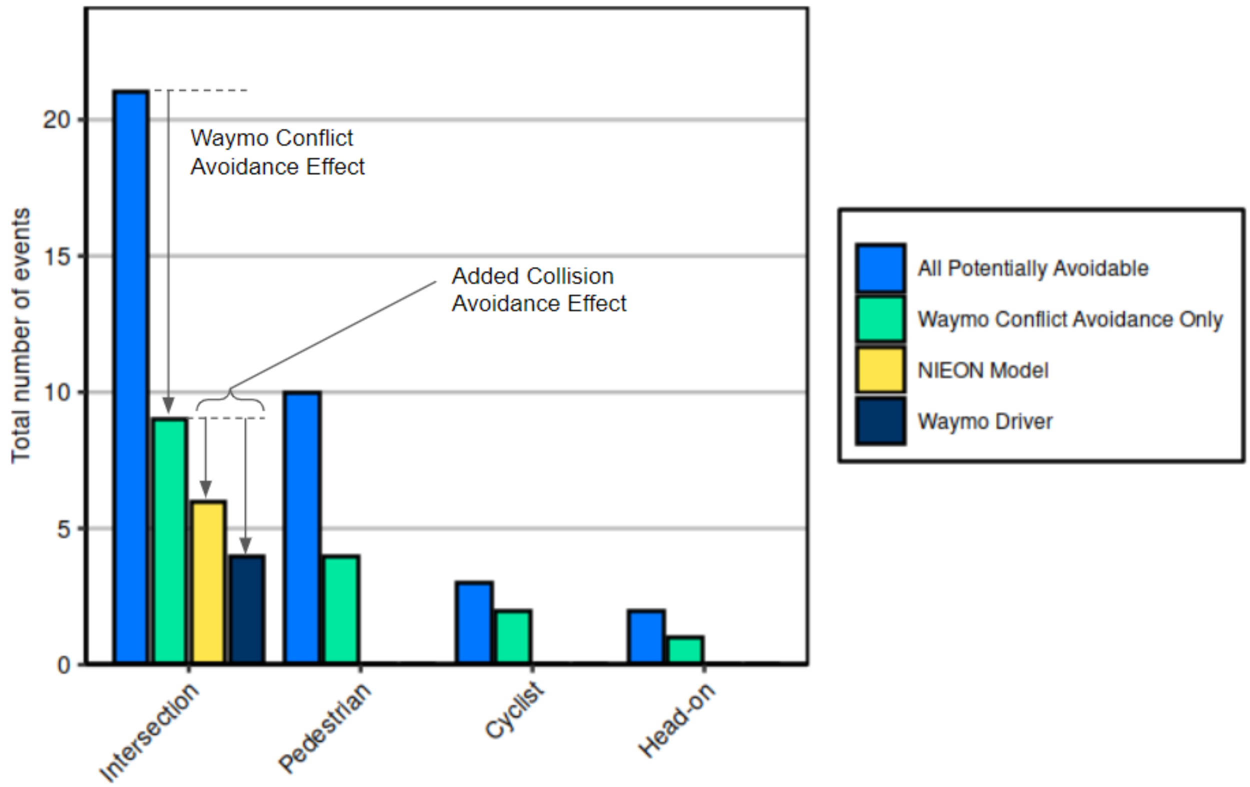

Figure 4: Number of simulated collision events by collision type. The "Waymo Conflict Avoidance Only" bar depicts the residual collisions as a result of Waymo's pre-conflict state driving alone, i.e., if the Waymo driver was not able to take evasive steering or braking action (collision avoidance), how many collisions would be remaining? The "Waymo Driver" bar reflects the residual collisions when simulating on-road driving (capturing the added benefit of the Waymo Driver taking collision avoidance action). The "NIEON model" bar reflects residual collisions when the NIEON model takes collision avoidance action following some initial Waymo Driver control of the vehicle (i.e., following any Waymo Driver conflict avoidance).

Reference behavior models also enable an analysis of driving behavior within individual cases. Figure 5 shows computed serious injury risk for each of the six residual collisions from the NIEON model comparison analysis. Each of the collisions is numbered according to the severity risk of the original collision. All six of the simulated residual scenarios involved an unprotected left turning initiator actor across the path of another approaching actor (twice from the lateral direction, LTAP/LD, and four times from the opposite direction, LTAP/OD). The Waymo Driver was found to outperform the NIEON model in all 6 simulated residual collisions. The Waymo Driver mitigated all simulated residual collisions with respect to the severity of the original fatal collision. The NIEON model simulations mitigated injury risk with respect to the original fatal collision for five out the six residual events. The Waymo Driver took some evasive action prior to the collision point in all of simulated residual collisions. The NIEON model did not begin its evasive action before the collision for three of its six residual collisions. In all scenarios, earlier response times were the primary factor enabling the Waymo driver to outperform the NIEON model.

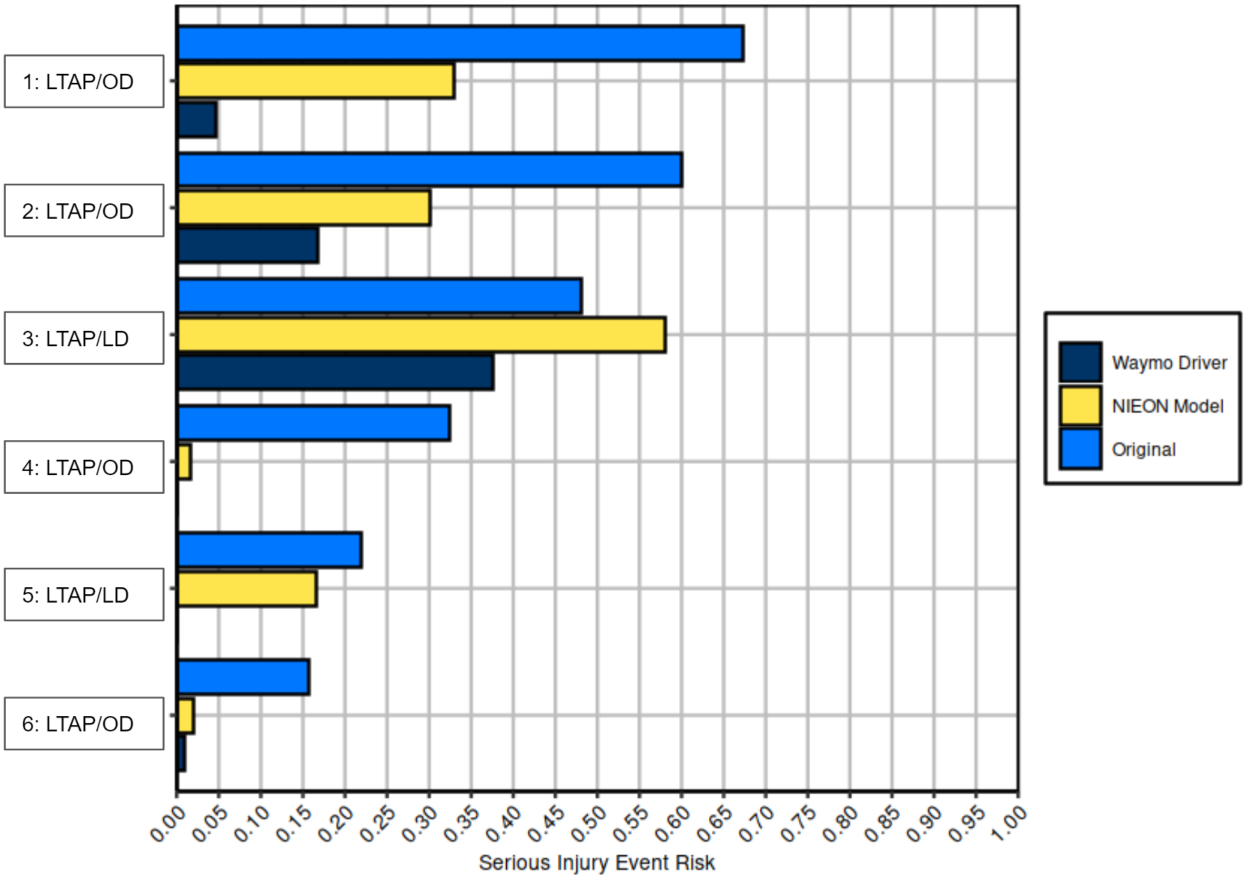

Figure 5: Probability of serious injury event (AIS3+) for simulated cases where either the Waymo Driver or Reference Model had a collision.[2]

The initiating turning actor in five of the six simulated residual collisions was a vehicle. In the sixth case, the initiating actor was a turning motorcycle (case 3). Cases 1, 2, 3, and 6 were all previously discussed in Scanlon et al. (2021). It is worth noting that the computed injury risk differs slightly for cases 1, 2, and 6 from this prior publication due to the adoption of the novel injury risk model for vehicle occupants developed by McMurray et al. (2021).

4. Discussion

The current study demonstrated a novel methodology for evaluating ADS technologies using a collection of fatal, real-world collisions from Waymo's current commercial ODD and a behavior reference model defining a certain level of collision avoidance performance in these scenarios. Counterfactual simulation was used to determine the capacity of the Waymo Driver to avoid or mitigate these collisions through either conflict avoidance or collision avoidance. The collision avoidance behavior reference model, representing a non-impaired driver with eyes on the

[2] It is notable that case 2011-118246 was the only simulated scenario where serious injury risk in the NIEON model simulation exceeded the original case. In this LTAP/LD collision, the original responder driver in the crash reconstruction was traveling at just below the posted speed limit when the left turning initiator crossed the vehicle's path. The NIEON model inherits the speed of the Waymo Driver, which in this case was the posted speed limit. The NIEON model is not able to reduce its speed below the impact speed of the original human driver, which results in a slightly higher injury risk.

conflict (NIEON) model, was then used to evaluate the Waymo Driver's performance in situations where evasive maneuvering can serve to mitigate or prevent a collision.

When any driver is facing a potential collision scenario, they must perceive the potential threat, determine the potential actions of that threat, and enact a physical decision on what actions need to be taken. These events can take many forms, and can come with little, or no, notice. The current study emphasized, again, the importance of conflict avoidance as a primary means for mitigating and avoiding high severity collisions. The study also emphasized the importance of collision avoidance. Scenarios requiring collision avoidance in addition to conflict avoidance are observed during cases with unprotected left turning vehicles that fail to detect an oncoming vehicle and force responding drivers to respond urgently to the potential threat. In another example, the threat may also emerge suddenly, such as a vulnerable road user from an occluded area or a perpendicular crossing vehicle traveling that is violating a traffic signal. These real events occur on roadways within the intended ODD of ADS technologies. Following rules of the road, alone, is insufficient for maximizing safety benefits. To be the most effective driver, of any form, the capacity to respond to these reasonably foreseeable scenarios is paramount to the safety benefits of that driver.

4.1 Reference Model as a Potential Target for ADS Performance

Reference behavior models enable a method for evaluating the competency of driver actions. They are intended to represent some potential target by which system performance can be evaluated. This study focused on using reference behavior models as a means for evaluating collision avoidance behavior. The goal of this current study was not to establish exact specifications for a reference behavior model-derived safety target (what ADS performance should be). Rather, the current study shows a method for how a behavior reference model benchmark can be used to evaluate ADS performance.

Translating any behavior reference model to an actionable performance benchmark requires many layers of consideration. Public acceptance is one key component of this consideration. Generally, public acceptance will be driven, in part, by the opportunity for ADS to substantially outperform the current driving population in many crash types. For example, inherently avoiding some crashes with human-related pre-crash behaviors as main contributing factors (e.g. speeding, reckless driving, impairment, drowsiness distraction). These pre-crash factors were associated with collision-initiating behaviors or poor responder actions (Mueller et al., 2020). Additionally, the benchmarks established should consider the continual improvement of the technology over time. Another dimension is the potential safety impact of an ADS. Kalra & Groves (2017) found that for a hypothetical ADS, deploying an ADS that was marginally better than human drivers sooner rather than waiting to deploy a system substantially better than humans but at a later date would have an overall benefit to society.

The NIEON model used in this study was intended to reflect collision avoidance (response time and evasive maneuvering) performance of a non-impaired human with their eyes on the conflict. To help illustrate this concept, Figure 6 provides a conceptual overview of how to interpret the NIEON model within a set of reference points. There is an inherent relationship between driving behavior performance during a conflict (i.e., response time and evasive maneuvering) and safety benefits (i.e., reduction in injuries). As presented, there are other important contributory factors, such as conflict avoidance, that also dictate injury reduction benefits.

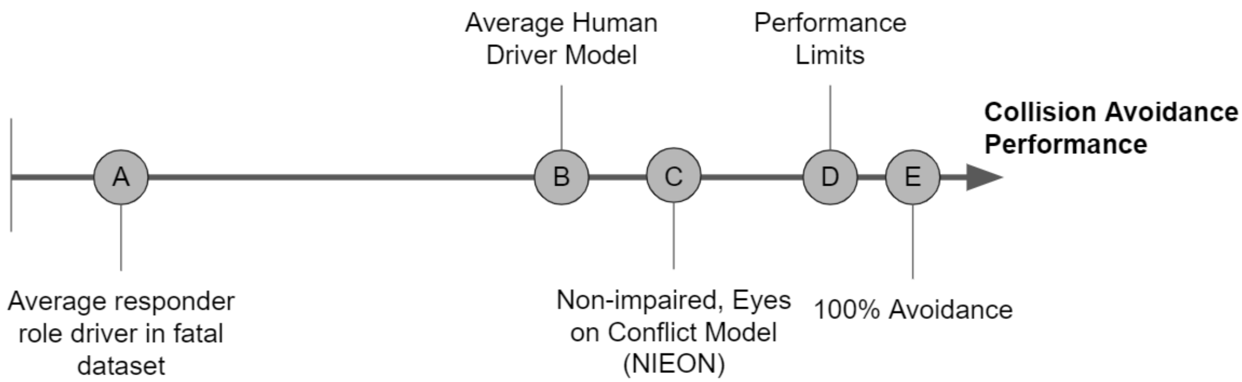

Figure 6: Conceptual illustration of collision avoidance performance on a continuum with reference points, including the original human drivers in the Chandler fatal dataset and the NIEON model.

When examining driver behavior during conflicts, the responses of the drivers in the actual reconstructed injury-causing crashes can be used as one reference point (A). The depicted data point (A) is intended to reflect an average response performance in the fatal crashes in the current dataset. Although the collision avoidance of the original actor was not explicitly evaluated (independent of conflict avoidance), this study's results suggest that many of the responders executed insufficient evasive action or none at all, which is unsurprising given that the dataset was selected from collisions with the most extreme outcomes, that is, fatal injury. With that being said, these responses exist on a spectrum, and being involved in a fatal collision does not necessarily imply a poor response time or poor evasive maneuvering (since there could have been other factors contributing to that fatal outcome). This can be more explicitly evaluated in future studies by placing the NIEON model in the position of the original driver and then allowing the NIEON to take over control from the original driver, which would allow for a one-to-one evaluation of that original driver's collision avoidance action. A second behavioral reference is the general driving population, from which we have depicted a single point, an average human driver response (B). While drivers in the overall driving population exhibit a wide spectrum of response times and evasive maneuvers should they organically experience a potential fatal collision event, this depicted reference point only represents a single instance (the average) of these. Thus, the average collision avoidance performance represented by the single reference point (B) is a fundamentally different concept than average human crash rate; a hypothetical driver always performing like B when encountering a conflict would likely have a significantly lower crash rate than a driver exhibiting an average human crash rate whose performance varies over time, i.e. due to distraction, impairment etc.). The NIEON model (C) represents a single collision avoidance performance reference point for a driver that is non-impaired and always has their eyes on the conflict, hence representing an idealized consistently performing driver that does not exist in the current population. In the current study, the NIEON model (with Waymo conflict avoidance) far outperformed the original human responses exhibited in the crashes in the vast majority of the original crashes. Using a statistical model fit to data from always eyes-on-conflict human drivers, the current study shows that ADS achieves the same or superior performance compared to this reference model in high-severity collision events. Note that choosing to compare ADS performance with this ideal, consistently performing reference behavior model will continue to drive performance improvements, which we believe will translate into substantial safety benefits. Additionally, by being modeled on response time data from actors with their eyes on the conflict, the NIEON model is inherently representing higher performing collision avoidance actions than most of the drivers within the human driving population that may be inattentive to or distracted from detecting an imminent threat or simply execute poor evasive maneuvering. It is also important to note that the NIEON model captures a single modeled response time and set of discrete evasive maneuvers; additional performance specifications (e.g., varying response time intercepts; Engstrom et al., 2022) for the NIEON model may be considered that might represent various collision avoidance performance levels. Beyond this current NIEON model's performance, one can envision a reference model operating at performance limits (D) in terms of perception, processing, actuation, and other physical limits that restrain the collision avoidance capabilities of any road user. Such performance limits exist for both humans and ADS technology. and avoidance of 100% of potential collisions (E) is currently infeasible given these performance limits. As ADS technology is further deployed, reducing and eliminating injury-causing driver initiator behaviors will further boost ADS safety benefits. Whereas the current study compared the Waymo Driver to the NIEON model, future studies might investigate the safety benefits of different reference behavior models and model parameters to aid in selecting an acceptable performance level to hold an ADS to.

Outside the scope of the current study was an evaluation of the *conflict* avoidance capacity of the Waymo Driver. The NIEON model used in the present paper does not account for this often nuanced behavior, such as managing occlusions or proactive mitigation measures when faced with uncertain, unexpected actor behaviors. Despite this, conflict avoidance affects all of the simulations in this study, including the NIEON model which relies on a handoff from the Waymo Driver before executing collision avoidance. Accordingly, the NIEON model comparison results should be viewed only as an evaluation of when and how the Waymo Driver took collision avoidance action once a conflict had already emerged - rather than an evaluation of the Waymo Driver to avoid these conflicts altogether. As noted in Najm and daSilva (2000), quantifying the ability to limit conflict exposure is a critical aspect of quantifying safety benefits. An additional reference behavior model of conflict avoidance would enable an additional ADS evaluation.

4.2 Case-Level Versus Aggregate-Level Evaluation

4.2.1 Evaluating ADS Performance on a Case-by-Case Basis

This study presented two different levels to consider the results of a reference behavior model approach. The first consideration is on the single case-level. This approach evaluated the performance of the ADS within the context of each specific collision scenario. There are multiple benefits to performing the analysis on a single case. It allows an assessment of the desirability of the behavior. With regards to the NIEON model used in the current study, it supports answering the question: how would a non-impaired driver with their eyes on the conflict behaved under identical crash avoidance conditions? One might also imagine using this type of approach to evaluate the progress of the ADS collision avoidance performance over time.

Although the NIEON reference behavior model used in this study is novel in its composition, the reliance on a benchmark derived from the current fleet of vehicles in operation to evaluate system performance is not a new vehicle safety concept. For example, Insurance Institute for Highway Safety (IIHS, 2020) utilized existing seat and head restraint performance data to determine the appropriate criteria for testing safety performance in a simulated rear impact test. Likewise, Schram et al. (2015) described EuroNCAP's determined scoring procedures as considering "expected performance of current systems and accident priorities." The goal of these prior benchmarks is to establish some performance level based on something that both advances safety and is achievable. The NIEON model used in the current study follows this same philosophy and represents a level of safety that, if achieved by every driver on the roadway for a given scenario, might advance traffic safety by reducing injury risk.

4.2.2 Evaluating ADS Performance at the Aggregate-Level

This study also presented the potential to use reference behavior models in aggregate-level evaluation. To our knowledge, this study is the first presentation of aggregate-level real-world crash performance using a reference behavior model. The implementation in the current study is similar in nature to what may be observed in prior safety benefits study. A commonly executed approach for ADS benefits study is to take some representative set of crashes, and then perform counterfactual simulations to establish the anticipated potential benefits relative to the original outcomes. For level-4 ADS technology, this benefits analysis requires accounting for exposures to potential collision scenarios (i.e., modeling the likelihood of experiencing some set of conflict scenarios). Although comparison to some original outcome achieves some performance relative to the current fleet, a reference behavior model, like the NIEON model, unlocks additional aggregate-level analysis for evaluators and, in particular, establishes interpretable benchmarks for crash prevention and mitigation in scenario-based testing. This study presents one implementation of aggregate-level assessment using the NIEON model. This section discusses the multiple areas that the NIEON model can be applied and utilized for performance assessment.

Relying solely on reconstruction data to benchmark aggregate performance allows only a comparison to some original actor that was involved. The aggregate performance level of some original actors involved is not necessarily something to aspire to surpass only marginally, given that it is not uncommon for drivers in high severity collisions to be distracted, drowsy, or intoxicated. Merely avoiding those frequent crash causes would, of course, provide aggregate-level benefits that may constitute risk reduction. As this study demonstrates, there are also opportunities to avoid and mitigate collisions when the driver is not a contributing party in the pre-conflict phase. A reference model allows evaluators to implement a custom, repeatable avoidance performance-level that may represent superior behavior to the overall driving population (that includes some undesirable behavior).

One key potential for using reference models is to serve as a tool to help generate fleetwide safety benefits estimates for some conflict types. As ADS developers look to deploy their technologies into various ODDs, there may be a need to establish what level of benefits - as measured by propensity to reduce injuries - the technology may achieve. As Najm and daSilva (2000) indicate, one way to arrive at this estimate is to utilize both (a) the propensity to enter into this sort of conflict and (b) avoidance action given that a potential conflict has occurred. For evaluating avoidance, a critical piece of this formulation is to generate a relative performance estimate to some driving population. There is potential for leveraging reference behavior models as an independent tool for establishing relative avoidance performance with respect to some benchmark road user other than the avoidance performance exhibited in the human crash data. Using some modeled exposure to the conflict and the relative avoidance performance, an aggregate crash and injury reduction for a given set of existing crashes can be estimated. This is an important area of research and development that we believe ADS evaluators should consider as designers strive to establish “positive risk balance” (Favarò, 2021).

4.3 Implications for Scenario-Based Testing

The current study evaluated the Waymo Driver within a replicated version of a collision that previously occurred. One alternative to this approach would be to take some previously observed scenario (a seed event), vary the conditions of that scenario, and evaluate the performance of the ADS under those set of conditions. Alternatively, one may consider building scenarios using some novel mechanism that may not necessarily have been previously observed, but that represents something that is reasonably foreseeable given known driving behaviors and failure modes.

A variety of scenario creation methods exist beyond replaying collisions that previously occurred. Zhao et al. (2016) examined inserting adversarial actors into naturalistic driving logs in order to create collision avoidance scenarios. Feng et al. (2021) did similar work, but as an alternative to modifying existing naturalistic driving data logs, proposed creating synthetic naturalistic driving environments using traffic simulation, and allowing some actors to take on adversarial behaviors. In a review by Riedmaier et al. (2020), additional approaches are discussed, but it is clear that there are multiple competing methodologies to create sufficient coverage of the potential collision space, although best practices such as ISO/DIS 34502 are starting to harmonize these methodologies (ISO, 2021b). As presented in Webb et al. (2020), Waymo leverages driving data to simulate variations of those historical logs. Additionally, closed-course testing, and variations of those tests, can also be used to create synthetic, realistic collision avoidance scenarios. Waymo implements a Collision Avoidance Testing program using a large collection of seeded events for evaluating performance. Beyond being used daily in internal development testing at Waymo, reference behavior models have utility in regulatory and evaluation agencies that are looking to enact some reasonable safety benchmarks.

In any sort of these synthetic simulation conditions, the question exists: how adequate was the performance of the ADS in that scenario? This study's reference model implementation is intended to be a demonstration of one potential methodology for evaluating and benchmarking that performance.

The utility of reference behavior models goes beyond scenario-based testing. One might consider evaluating the reasonableness of some real-world action by retrospectively evaluating the capacity of the reference model to avoid some known collision modality. This has potential utility in both the human driving population evaluation and in the continued evaluation of deployed ADS technologies.

4.4 Limitations and Future Challenges

There are several important limitations and challenges to note from the current study. First, the current dataset consists entirely of collisions that involve only human actors. This means that the failure modes that led to the initiation of and responding to the collision sequence were present in these human actors' behaviors. It is, of course, important to consider how an ADS may respond to these known human-initiated collision modalities and to test the system's capacity to not perform the same initiating behaviors. In parallel to these known collision scenarios, ADS vehicles may have their own unique exposure to conflict events, or unique pre-crash contributing factors. The difference in exposure is not being tested in this current study, but plays a key role in the safety impact achieved by the ADS technology. Additional testing regimes are required in order to uncover and verify the larger spectrum of potential conflict scenarios (Webb et al., 2020). This particular study focuses exclusively on the effect of collision avoidance, which is also a key factor in achieving safety benefits.

Second, this study was based on reconstructions of police-reported collisions. There was variability in the amount of crash-related documentation that was available for reconstructing these collisions. This variability translates to uncertainty in the reconstructions themselves. Our approach in this study was to rely on a single representation of the collision reconstruction for our counterfactual simulation, which has been done in prior studies (Gruber et al., 2019; Hamdane et al., 2015). Alternatively, probabilistic reconstructions may help to better capture the actual collision scenario with some degree of precision (Bareiss et al., 2019; Scanlon et al., 2017; Schachner et al., 2020). Additionally, other traffic for uninvolved vehicles was largely not captured in the current reconstructions. This can affect the available opportunities for avoidance action and increase the potential for secondary collision events. As has been done in prior works (Feng et al., 2021; Zhao et al., 2016), incorporating traffic into the counterfactual simulations may provide more robust testing conditions.

Third, the current study implemented a single reference model, the NIEON model, for evaluating the performance of the Waymo Driver. The NIEON model executed a single driving response using a point estimate of response time (Engstrom et al., 2022). This presentation of performance was intended as a demonstration of using reference models in scenario-based testing. Looking toward the future, reference models could be built using a different set of parameter inputs, or with a different or larger scope in mind. For example, incorporating reference conflict avoidance behaviors would further enable evaluation of safe driving behavior also in pre-conflict scenarios. It is also important to note that this study was not meant to provide a comprehensive technical summary of the NIEON model. Rather, the study focused on the performance of the NIEON model in a series of real-world, representative collision scenarios.

Fourth, system performance was tested in a simulated environment with a distinct set of conditions. Variability in those conditions, such as due to small variations in the behavior of other actors or some module-specific challenges that may arise, were not rigorously tested using this approach. Module-specific failures were not explored in the current study. Our methods were intended to focus our evaluation on the behavioral competencies of the Waymo Driver given the conditions of the real-world collision scenarios. Rather than test our expected median response to those conditions, our latency modeling technique did replicate the 95th percentile of system modules to conservatively account for latency effects. This latency modeling technique covers expected operational behavior of the ADS, but the aggregate safety impact of the ADS will also be affected by extreme, long-tail latencies that could occur. As discussed above, Waymo utilizes multiple methodologies to ensure safety, including platform requirements related to latency that attempt to mitigate the likelihood of failure, like extreme latency. Additionally, some non-determinism is inherent in ADS algorithms and the simulation environment. We previously tested this non-determinism and found no notable injury risk change over five repeated tests, which suggested that the effect of non-determinism on the current study was negligible (Scanlon et al., 2021).

5. Conclusion

This study evaluated the capacity of the Waymo Driver to prevent and mitigate a set of reconstructed, real-world fatal collisions within Chandler, Arizona in comparison to the NIEON collision avoidance model. A counterfactual simulation methodology was used to evaluate the Waymo Driver in a scenario similar to the original crash. The Waymo Driver's performance within collision scenarios that required collision avoidance was evaluated using the NIEON model. The present analysis extended upon the previous results in Scanlon, et al. (2022), which showed that 100% of initiator collisions were avoided in simulation using conflict avoidance alone, and 82% of responder collisions were prevented by conflict and collision avoidance. Three responder scenarios (8% of responder role) involved the Waymo Driver being rear-end struck, and were labeled as being unaffected by conflict avoidance and collision avoidance in preventing or mitigating the potential collision outcomes.

Of 36 potentially avoidable (excludes rear-end struck) responder collisions, the *combined effect of conflict avoidance and collision avoidance* of the Waymo Driver prevented 32 (89%), and mitigated 4 collisions. Waymo conflict avoidance was effective at preventing 56% of the 36 events. The combined effect of conflict avoidance and collision avoidance also led to a serious injury risk reduction for these 36 responder collisions: a 97% reduction for the Waymo Driver. Regarding the effect of *collision avoidance only* (i.e., a direct evaluation of the Waymo Driver and NIEON Model collision avoidance performance), of 16 conflicts entered, 12 (75%) were prevented by the Waymo Driver, and 10 (62.5%) were prevented by the NIEON model. The NIEON Model mitigated an additional 5 collisions and did not mitigate 1 collision. In these 16 conflicts entered, 93% of serious injury risk was reduced by the Waymo Driver, whereas 84% of serious injury risk was reduced by the NIEON model.

Further, results were also examined on a case-by-case basis to evaluate relative performance and the degree of effectiveness for each scenario. The only residual collision events for both the Waymo Driver and the reference behavior model were left turn across path (LTAP) intersection collision types, where the driver was responding to an initiator making an unprotected left in either an LTAP/LD or LTAP/OD configuration.

The current study is intended to showcase how reference behavior models can be used to help establish interpretable safety performance benchmarks in scenario-based safety benefit analysis. With regards to the specific model showcased, the results suggest that performing at or beyond this collision avoidance benchmark will translate to safer driving performance. This is demonstrated through effective crash and injury mitigation performance of the NIEON model within this fatal collision dataset. It is, therefore, proposed that using this model to benchmark collision avoidance driving behavior will also stand to promote safer driving behavior. Future evaluations may consider additional reference models of desirable driving behaviors that also promote safety benefits. Additionally, the utility of these reference models extend beyond simulating crash reconstructions, and have utility for establishing performance benchmarks in many forms of scenario-based testing and evaluation of on-road performance. We believe that these methods have important implications for technology developers, researchers, and evaluators interested in establishing ADS benchmarks.

6. References

Association for the Advancement of Automotive Medicine. (2015). The Abbreviated Injury Scale: 2015 Revision. Chicago, IL: AAAM.

Blanco, M., Atwood, J., Russell, S. M., Trimble, T., McClafferty, J. A., & Perez, M. A. (2016). Automated vehicle crash rate comparison using naturalistic data. Virginia Tech Transportation Institute.

Blower, D., & Woodrooffe, J. (2013). *Real-World Safety Effect of Roll Stability Control* (No. 2013-01-2392). SAE Technical Paper.

Brach, R. M., & Brach, M. (2011). Vehicle Accident Analysis and Reconstruction Methods, Second Edition. SAE International.

Chai, C., Zeng, X., Wu, X., & Wang, X. (2019). Safety evaluation of responsibility-sensitive safety (RSS) on autonomous car-following maneuvers based on surrogate safety measurements. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)* (pp. 175-180). IEEE.

Cicchino, J. B. (2017). Effectiveness of forward collision warning and autonomous emergency braking systems in reducing front-to-rear crash rates. *Accident Analysis & Prevention*, *99*, 142-152.

Engström, J., Liu, S. Y., Dinparastdjadid, A., & Simoiu, C. (2022). Modeling road user response timing in naturalistic settings: a surprise-based framework. *arXiv preprint arXiv:2208.08651*.

Favarò, F. (2021). Exploring the Relationship Between" Positive Risk Balance" and" Absence of Unreasonable Risk". *arXiv preprint arXiv:2110.10566*. https://arxiv.org/ftp/arxiv/papers/2110/2110.10566.pdf

Feng, S., Yan, X., Sun, H., Feng, Y., & Liu, H. X. (2021). Intelligent driving intelligence test for autonomous vehicles with naturalistic and adversarial environment. Nature communications, 12(1), 1-14.

Flannagan, C., Bärgman, J., & Bálint, A. (2019). Replacement of distractions with other distractions: A propensity-based approach to estimating realistic crash odds ratios for driver engagement in secondary tasks. *Transportation research part F: traffic psychology and behaviour*, *63*, 186-192.

Forkenbrock, G. J., Garrott, W. R., Heitz, M., & O'Harra, B. C. (2003). An experimental examination of J-turn and fishhook maneuvers that may induce on-road, untripped, light vehicle rollover. *SAE transactions*, 1112-1127.

Goodall, N. J. (2021). Potential Crash Rate Benchmarks for Automated Vehicles. *Transportation Research Record*. https://doi.org/10.1177/03611981211009878

Haus, S. H., Sherony, R., & Gabler, H. C. (2019). Estimated benefit of automated emergency braking systems for vehicle–pedestrian crashes in the United States. *Traffic injury prevention*, *20*(sup1), S171-S176.

Hu, M., & Li, Y. (2017). Drivers' avoidance patterns in near-collision intersection conflicts. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)* (pp. 1-6).

Insurance Institute for Highway Safety (2020). Vehicle Seat/Head Restraint Evaluation Protocol Dynamic Criteria (Version VI). Ruckersville, VA: IIHS.

Isaksson-Hellman, I., & Lindman, M. (2015). Real-world performance of city safety based on Swedish insurance data. In *24th international technical conference on the enhanced safety of vehicles (ESV)* (paper number 15-0121).

ISO 21934. (2021a). 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.

ISO 34502 (2021b). Road vehicles — Scenario-based safety evaluation framework for Automated Driving Systems. ISO/DIS 34502:2021.

ISO/TR 21974–1 (2018). Naturalistic driving studies – Vocabulary – Part 1.

Itti, L., & Baldi, P. (2009). Bayesian surprise attracts human attention. Vision research, 49(10), 1295-1306.

Kalra, N., & Paddock, S. M. (2016). Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?, RAND Corporation, Santa Monica, Calif. RR-1478-RC, 2016. As of February 02, 2021: https://www.rand.org/pubs/research_reports/RR1478.html

Kalra, N., & Groves, D. G. (2017). *The enemy of good: Estimating the cost of waiting for nearly perfect automated vehicles*. Rand Corporation. Accessed from: https://www.rand.org/pubs/research_reports/RR2150.html

Kudlich, H. (1966). Beitrag zur Mechanik des Kraftfahrzeug-Verkehrsunfalls. Technische Hochschule Wien.

Kullgren, A. (2008). Dose-response models and EDR data for assessment of injury risk and effectiveness of safety systems. In Proc of Int. IRCOBI Conf., Bern, Switzerland (pp. 3-14).

Kusano, K. D., & Gabler, H. C. (2012). Safety benefits of forward collision warning, brake assist, and autonomous braking systems in rear-end collisions. *IEEE Transactions on Intelligent Transportation Systems*, *13*(4), 1546-1555.

Li, X., Rakotonirainy, A., & Yan, X. (2019). How do drivers avoid collisions? A driving simulator-based study. *Journal of safety research*, *70*, 89-96.

Lindman, M., Isaksson-Hellman, I., & Strandroth, J. (2017, September). Basic numbers needed to understand the traffic safety effect of automated cars. In *Proceedings of the IRCOBI Conference*.

Liu, P., Yang, R., & Xu, Z. (2018). How safe is safe enough for self-driving vehicles?. *Risk analysis*, *39*(2), 315-325.

Malaterre, G., Ferrandez, F., Fleury, D., & Lechner, D. (1988). Decision making in emergency situations. *Ergonomics*, *31*(4), 643-655.

Markkula, G., Engström, J., Lodin, J., Bärgman, J., & Victor, T. (2016). A farewell to brake reaction times? Kinematics-dependent brake response in naturalistic rear-end emergencies. *Accident Analysis & Prevention*, *95*, 209-226.

Mazzae, E. N., Baldwin, G. S., & McGehee, D. V. (1999). Driver crash avoidance behavior with ABS in an intersection incursion scenario on the Iowa driving simulator. *SAE transactions*, 2353-2363.

McMurry, T. L., Cormier, J. M., Daniel, T., Scanlon, J. M., & Crandall, J. R. (2021). An omni-directional model of injury risk in planar crashes with application for autonomous vehicles. Traffic injury prevention, 1-6.

Mueller, A. S., Cicchino, J. B., & Zuby, D. S. (2020). What humanlike errors do autonomous vehicles need to avoid to maximize safety?. Journal of safety research, 75, 310-318.

Najm, W. G., & daSilva, M. P. (2000). Benefits estimation methodology for intelligent vehicle safety systems based on encounters with critical driving conflicts. In ITS America 10th Annual Meeting and Exposition: Revolutionary Thinking, Real ResultsIntelligent Transportation Society of America (ITS America).

National Center for Statistics and Analysis. (2021). Traffic safety facts 2019: A compilation of motor vehicle crash data (Report No. DOT HS 813 141). National Highway Traffic Safety Administration.

National Highway Traffic Safety Administration (2013). Laboratory test procedure for dynamic rollover: the fishhook maneuver test procedure. *Washington (DC): US, Department of Transportation*.

Nees, M. A. (2019). Safer than the average human driver (who is less safe than me)? Examining a popular safety benchmark for self-driving cars. *Journal of safety research*, *69*, 61-68.

Nie, J., Li, G., Yang, J., Zhou, X., Zhang, C., Yu, X., ... & Wang, M. (2013, September). A study of injury risk of bicyclist and pedestrian in traffic accidents in Changsha of China. In Proceedings of The 5th Expert Symposium on Accident Research (ESAR). Hanover, Germany.

Nyberg, T., Pek, C., Dal Col, L., Norén, C., & Tumova, J. (2021). Risk-aware Motion Planning for Autonomous Vehicles with Safety Specifications. In *IEEE Intelligent Vehicles Symposium*.

Riedmaier, S., Ponn, T., Ludwig, D., Schick, B., & Diermeyer, F. (2020). Survey on Scenario-Based Safety Assessment of Automated Vehicles. IEEE Access, 8, 87456-87477.

Riexinger, L., Sherony, R., & Gabler, H. (2019). *Has Electronic Stability Control Reduced Rollover Crashes?* (No. 2019-01-1022). SAE Technical Paper.

Rothoff, M., Zarghampour, H., Tivesten, E., Victor, T., Ödblom, A., Bergh, T., & Lindenberg, B. (2019). ADEST Autonomous Driving Effects on Sustainable Transportation & ADFE Autonomous Driving Fuel Economy.

Sander, U., & Lubbe, N. (2018). Market penetration of intersection AEB: Characterizing avoided and residual straight crossing path accidents. Accident Analysis & Prevention, 115, 178-188.

Scanlon, J. M. (2017). Evaluating the Potential of an Intersection Driver Assistance System to Prevent US Intersection Crashes (Doctoral dissertation, Virginia Tech).

Scanlon, J. M., Kusano, K. D., & Gabler, H. C. (2015). Analysis of driver evasive maneuvering prior to intersection crashes using event data recorders. Traffic injury prevention, 16(sup2), S182-S189.

Scanlon, J. M., Kusano, K. D., & Gabler, H. C. (2016). Influence of roadway characteristics on potential safety benefits of lane departure warning and prevention systems in the us vehicle fleet. In Transportation Research Board 95th Annual Meeting (No. 16-1893).

Scanlon, J. M., Sherony, R., & Gabler, H. C. (2017). Injury mitigation estimates for an intersection driver assistance system in straight crossing path crashes in the United States. Traffic injury prevention, 18(sup1), S9-S17.9.

Scanlon, J. M., Kusano, K. D., Daniel, T., Alderson, C., Ogle, A., & Victor, T. (2021). Waymo Simulated Driving Behavior in Reconstructed Fatal Crashes within an Autonomous Vehicle Operating Domain. https://waymo.com/safety/simulated-reconstruction

Schwall, M., Daniel, T., Victor, T., Favaro, F., & Hohnhold, H. (2020). Waymo public road safety performance data. *arXiv preprint arXiv:2011.00038*.

Schram, R., Williams, A., van Ratingen, M., Ryrberg, S., & Sferco, R. (2015). Euro NCAP's first step to assess Autonomous Emergency braking (AEB) for Vulnerable Road Users. In Proceedings of 24th Enhanced Safety of Vehicles (ESV) conference.

SAE J3016: Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles. (2021). https://doi.org/10.4271/J3016_202104

Shalev-Shwartz, S., Shammah, S., & Shashua, A. (2017). On a formal model of safe and scalable self-driving cars. *arXiv preprint arXiv:1708.06374*.

Tominaga, S., & Sakurai, M. (2002). Fundamental Analysis of Motorcyclist Injury Risk Using A Statistical Model Based on Real-world Crashes (No. 2002-32-1830). SAE Technical Paper.

UN/ECE (2021). *Uniform provisions concerning the approval of vehicles with regard to Automated Lane Keeping Systems.* Economic Commission for Europe, Inland Transport Committee, World Forum for Harmonization of Vehicle Regulations. E/ECE/TRANS/505/Rev.3/Add.156

Webb, N., Smith, D., Ludwick, C., Victor, T., Hommes, Q., Favaro, F., ... & Daniel, T. (2020). Waymo's Safety Methodologies and Safety Readiness Determinations. arXiv preprint arXiv:2011.00054. https://arxiv.org/abs/2011.00054

Weber, S., Blum, K., Ernstberger, A., & Färber, B. (2015). Standard reactions–driver reactions in critical driving situations. *Procedia Manufacturing*, *3*, 2489-2496.

Wege, C. A., Pereira, M., Victor, T. W., Krems, J. F., Stevens, A., & Brusque, C. (2014). Behavioural adaptation in response to driving assistance technologies: A literature review. *Driver adaptation to information and assistance systems* (S. 13–34). London: The Institution of Engineering and Technology.

Xu, X., Wang, X., Wu, X., Hassanin, O., & Chai, C. (2021). Calibration and evaluation of the Responsibility-Sensitive Safety model of autonomous car-following maneuvers using naturalistic driving study data. *Transportation research part C: emerging technologies*, *123*, 102988.

Zhao, D., Lam, H., Peng, H., Bao, S., LeBlanc, D. J., Nobukawa, K., & Pan, C. S. (2016). Accelerated evaluation of automated vehicles safety in lane-change scenarios based on importance sampling techniques. IEEE transactions on intelligent transportation systems, 18(3), 595-607.

