

A Formal Safety Characterization of Advanced Driver Assist Systems in the Car-Following Regime with Scenario-Sampling^{*}

Bowen Weng^{+*} Minghao Zhu^{+*} Keith Redmill^{*}

** Department of Electrical and Computer Engineering at The Ohio
State University, OH, 43210 USA (e-mail: weng.172@osu.edu,
zhu.1385@osu.edu, redmill.1@osu.edu).*

Abstract: The capability to follow a lead-vehicle and avoid rear-end collisions is one of the most important functionalities for human drivers and various Advanced Driver Assist Systems (ADAS). Existing safety performance justifications of car-following systems either rely on simple concrete scenarios with biased surrogate metrics or require a significantly long driving distance for risk observation and inference. In this paper, we propose a guaranteed unbiased and sampling efficient scenario-based safety evaluation framework inspired by previous work on $\epsilon\delta$ -almost safe set quantification. The proposal characterizes the complete safety performance of the test subject vehicle in the car-following regime. The performance of the proposed method is also demonstrated in challenging cases including some widely adopted car-following decision-making modules and the commercially available Openpilot driving stack by CommaAI.

Keywords: Test and Validation, Scenario Sampling, Set Invariance, Advanced Driver Assist Systems.

1. INTRODUCTION

The car-to-car rear-end collision has been the most common crash type in the U.S. for decades. Various Advanced Driver Assist Systems (ADAS) have been developed and deployed to help mitigate the rear-end collision risk, including crash-imminent braking (CIB), autonomous emergency braking (AEB), traffic jam assist (TJA), adaptive cruise control (ACC), and pedestrian crash avoidance mitigation (PCAM). In this paper, we are primarily interested in vehicle following ADAS implementations, which cover a large portion of the currently available ADAS. We assume the Subject Vehicle (SV) performs sufficiently well in other operational modules, such as lane-keeping. This is a common assumption and is feasible to achieve in the practice of ADAS tests. We also emphasize that the proposed approach is applicable to evaluate other ADAS modules, such as the Lane-Keeping Assist System (LKAS), but details are beyond the scope of this paper.

The safety evaluation of an ADAS-equipped SV in the car-following and rear-end collision avoidance regime seeks to characterize the SV's safety performance against stationary/moving vehicles in front of the SV within the same lane or along the SV's current trajectory. One common testing approach is to observe the SV's performance in the real-world or a simulated naturalistic driving environment for a sufficiently long driving distance. One then observes

or infers the collision rate estimate. This is formally known as Monte-Carlo sampling, with other importance sampling based variants from Zhao et al. (2017) that help improve the sampling efficiency. However, the required testing effort is still too expensive and time consuming to be widely applicable in practice. The naturalistic driving environment is not necessarily static and may vary significantly from time to time. For those importance sampling based variants, the importance function estimate was developed with various heuristics, making it difficult to justify its accuracy. Also, as reported by Weng et al. (2022), such a statistical inference method occurs in an implicitly defined operable domain with the tendency to over-estimate the risk. Finally, a simple scalar measure of risk is not necessarily sufficient to cover and justify the complete safety performance of an SV.

The dominant approach adopted by most existing regulations and standards follows the scenario-based test where the SV is deployed as a black-box system (uncontrollable and partially observable) in a testing case with the lead vehicle following a certain prescribed behavior control policy. The common practice in this case presents a finite set of concrete scenarios and analyzes the testing outcome through an independent safety metric (i.e. the metric is computed independently from the test execution and data acquisition, and the testing data is presented as it stands). Some commonly observed concrete scenarios in the rear-end collision avoidance regime include the car-to-car lead vehicle braking in Forkenbrock and Snyder (2015), the suddenly revealed stationary vehicle (SRSV) and the lead vehicle lane change and brake (LVLCB) in Rao et al. (2019),

^{*} This work was supported in part by the U.S. Department of Transportation (USDOT) under grant 69A3552047138 for the CARMEN University Transportation Center (UTC).

⁺ These authors contributed equally.

also known as the frontal cut-in scenario. The testing is mostly performed in a real-world proving grounds with a certain strikable target that emulates the motion and the appearance of a lead vehicle. Some also execute the test in a hardware-in-the-loop fashion such as the augmented scenes by Feng et al. (2020). The results are then analyzed using an extrinsic metric, such as the observed collision rate, time-to-collision violation (TTCV) by Wishart et al. (2020), and other surrogate measures summarized in Wang et al. (2021). Note that this is also the testing approach adopted by many regulatory standards such as the Europe NCAP by EuroNCAP (2019). However, as reported by Weng et al. (2021b), the set of concrete scenarios has very poor coverage of the SV's operational domain and is not of sufficient risk. The safety metrics are mostly biased and fail to arrive at a consensus agreement and make a fair comparison among various SVs as shown in Weng (2021). The approach is also fundamentally problematic if the underlying system is stochastic which is a common phenomena in practice, especially with the increased involvement of learning-based methods in perception and decision-making modules.

In this paper, we propose a scenario-sampling framework built on the Synchronous Pruning and Exploration (SPE) for safe set quantification in Weng et al. (2021b) with various improvements dedicated to the practical car-following testing regime. The basic idea of the proposed framework seeks to characterize the safe operational design domain (ODD) of the SV in the car-following regime through repeatedly sampling runs of scenarios in a guided manner. With a certain desired confidence level, one can then claim at *what* states the SV is potentially safe and *how* safe the SV is within the derived set of states. The proposed method is further demonstrated in Section 4, where it is shown capable of capturing various subtle safety properties and insights of widely adopted car-following models in both academic research as well as commercially available ADAS products. The studied ADAS are more realistic and difficult to evaluate than some of the previous work by Fan et al. (2017) and Zhao et al. (2016). To the best of our knowledge, many of the obtained properties have never been captured by existing work in the field.

Notation: The set of real and positive real numbers are denoted by \mathbb{R} and $\mathbb{R}_{>0}$ respectively. \mathbb{Z} denotes the set of all positive integers and $\mathbb{Z}_N = \{1, \dots, N\}$. $|\mathcal{X}|$ is the cardinality of the set \mathcal{X} .

2. PRELIMINARIES AND PROBLEM FORMULATION

Consider the general discrete-time system dynamics

$$\mathbf{s}(t+1) = f(\mathbf{s}(t); \omega(t)) \quad (1)$$

with state $\mathbf{s} \in \mathcal{S} \subseteq \mathbb{R}^n$ and uncertainties and disturbances $\omega \in \mathcal{W} \in \mathbb{R}^w$ for some $n, w \in \mathbb{Z}$. Let $\mathcal{C} \subset \mathcal{S}$ denote the set of failure states. Intuitively, for the system (1) to remain statistically safe, there should exist an $\mathcal{S}^* \subset \mathcal{S}$ with $\mathcal{S}^* \cap \mathcal{C} = \emptyset$ such that all trajectories initialized in \mathcal{S}^* remain inside \mathcal{S}^* with high probability. The safety performance justification then seeks to characterize the set \mathcal{S}^* . In practice, \mathcal{S}^* could be non-convex, non-unique, and of other complex structures, leading to various challenges for accurate characterization, statistically or determinis-

tically. In this paper, we adopt the $\epsilon\delta$ -almost safe set based methods from Weng et al. (2021a). Some important definitions and theorems are revisited in the following subsection.

2.1 $\epsilon\delta$ -Almost Safe Set

The following definition is adapted from Weng et al. (2021b,a).

Definition 1. (δ -Covering Set) Given a compact set $\mathcal{X} \subset \mathbb{R}^n$ for some $n \in \mathbb{Z}$ and $\delta \in \mathbb{R}_{>0}$. For any $\mathbf{x} \in \mathcal{X}$, let $\mathcal{N}_\delta(\mathbf{x})$ be the δ -neighbourhood of \mathbf{x} , i.e., $\forall \mathbf{x}' \in \mathcal{N}_\delta(\mathbf{x}), |\mathbf{x} - \mathbf{x}'| \leq \delta$. We claim that $\Phi_\delta^\mathcal{X}$ is a δ -covering set of \mathcal{X} if for some $k \in \mathbb{Z}$ and $\mathbf{x}_i \in \mathcal{X}, i = 1, \dots, k$, we have $\Phi_\delta^\mathcal{X} = \bigcup_{i \in \{1, \dots, k\}} \mathcal{N}_\delta(\mathbf{x}_i) \supseteq \mathcal{X}$ and $\Phi_\sigma^\mathcal{X} = \{\mathbf{x}_i\}_{i \in \{1, \dots, k\}} \subseteq \mathcal{X}$. Furthermore, $\Phi_\sigma^\mathcal{X}$ are centroids of $\Phi_\delta^\mathcal{X}$.

Recall \mathcal{C} is the set of failure states (e.g. collisions). The following definition formally characterizes the notion of the SV being ‘‘almost’’ safe in a certain set.

Definition 2. ($\epsilon\delta$ -Almost Safe Set) Given the system dynamics (1), $\epsilon \in (0, 1]$, $\delta \in \mathbb{R}_{>0}$, $\Phi \subseteq \mathcal{S}$. The set Φ is $\epsilon\delta$ -almost safe for the system (1) if there exists a δ -covering set Φ_δ of Φ with Φ_σ such that $\Phi_\delta \cap \mathcal{C} = \emptyset$ and

$$\mathbb{P}\left(\left\{\forall \mathbf{s} \in \Phi_\sigma, \forall \omega \in \mathcal{W} : f(\mathbf{s}; \omega) \notin \Phi_\delta\right\}\right) \leq \epsilon. \quad (2)$$

It is immediate from the above definition that $\lim_{\delta \rightarrow 0} \Phi_\delta^\mathcal{X} = \mathcal{X}$. Also note that as ϵ tends to zero, the $\epsilon\delta$ -almost safe set becomes an absolutely safe δ -covering set. To adapt the above definitions to the application of car-following regime safety analysis, we shall first characterize the car-following scenario in the form of (1).

2.2 The Scenario-based Car-Following System

In this paper, we consider the following system to formulate the interactive motion between a Subject Vehicle (SV) follower and a leading Principal Other Vehicle (POV) in the front sharing the same lane with the SV:

$$\mathbf{s}(t+1) = f_s(\mathbf{s}(t), \mathbf{u}(t); \omega_s(t)). \quad (3)$$

The state $\mathbf{s} = [d, v_0, v_1] \in \mathcal{S} \subset \mathbb{R}_{\geq 0}^3$, where $d \in [0, \infty)$ denotes the distance headway (simplified as headway or DHW in this paper) between the two vehicles and $v_0 \in [0, v_{\max}]$ and $v_1 \in [0, v_{\max}]$ denote the longitudinal velocity of the SV follower and the lead POV, respectively. In practice, a significantly large value of d is not a safety concern, hence the upper bound of d is often replaced with a sufficiently large value $d_{\max} \in \mathbb{R}_{>0}$. Other disturbances and uncertainties are denoted as $\omega_s \in \mathcal{W}_s$, which could involve environmental features (e.g. weather condition and road surface friction), infrastructure information (e.g. road curvature, road gradient, and speed limit), other kinematic and dynamic features (e.g. lateral offset between the vehicles and acceleration status of vehicles), other road users (e.g. pedestrian, cyclist, and other vehicles), planning parameters (e.g., free-traffic speed), and measurement error, to name a few. As also discussed by Weng et al. (2022), elements of the state \mathbf{s} and some of the uncertainties ω_s may be interchangeable depending on a particular feature's observability and how important it is in determining safety related properties. For example, EuroNCAP

(2019) considers the lateral offset between vehicles as an important feature that affects the performance of the SV, leading to an extra dimension added to the state \mathbf{s} . The action $\mathbf{u} \in \mathcal{U} \subset \mathbb{R}$ represents the control input of the lead POV, such as the desired velocity and the commanded acceleration. Note that the SV is the test subject in the testing content, thus it is an uncontrollable and (partially) observable black-box system (see Remark 3 in Weng et al. (2021b)).

Furthermore, the action \mathbf{u} is typically determined by a certain feedback control policy

$$\mathbf{u} = \pi(\mathbf{s}, \omega_s; \omega_u), \quad (4)$$

with \mathbf{s} , ω_s as defined above, and the uncertainties $\omega_u \in \mathcal{W}_u$. Intuitively, the policy π describes the lead POV driving behavior. In the scenario-based safety evaluation regime, the testing policy is a given function.

As a result, composing (3) with (4) we have the exact system dynamics of (1) with $n = 3$. The disturbances and uncertainties $\omega \in \mathcal{W}$ is jointly affected by \mathbf{s} and ω_s in (3) and ω_u in (4). In practice, the scenario system may not necessarily exhibit the Markov Decision Process (MDP) nature induced by (1) as the next-step state may be dependent upon not only the current state, but also a series of historic observations. One can extend the state space to involve those observations, yet the state space complexity will also increase significantly. In the particular car-following domain studied by this paper, we argue that the capability of SV to take advantage of historical information, if applicable, would only enhance safety performance. As a result, the safety property obtained from system (1) still remains as the worst-case justification.

A run of a test scenario, $\mathcal{RS}(\mathbf{s}_0, K)$ ($K \in \mathbb{Z}, K \geq 2$), thus starts from a certain state initialization $\mathbf{s}_0 \in \mathcal{S}$, consecutively collects a set of states admitting the system dynamics (1), and terminates either when encountering a failure event (e.g., collision) or when the K -th step of observation is reached. If f is explicitly known or approximately characterized, one can execute the test scenario and collect data through computer simulations. On the other hand, the scenario-based test can also be performed in a real-world testing proving ground with f implicitly induced.

The standard scenario-based safety evaluation methods (e.g. NCAP EuroNCAP (2019) and NHTSA guidelines in Forkenbrock and Snyder (2015); Rao et al. (2019)) specify the \mathbf{s}_0 based on expert-knowledge and real-world crash databases. The test policy $\pi(\cdot)$ is typically presented as a deterministic function with constant deceleration magnitude (e.g. $\pi(\mathbf{s}) = -6\text{m/s}^2, \forall \mathbf{s} \in \mathcal{S}$ in some of the car-to-car AEB cases). In this paper, we adopt a similar design of $\pi(\cdot)$ used by the above mentioned standardized tests (i.e., the lead POV executes the braking maneuver at a constant deceleration rate). This evaluates the SV's safety performance in a more adversarial environment than the naturalistic driving environment. We also emphasize that the proposed method does not rely on a particular testing policy, and will generalize easily to other testing policies such as those emulating naturalistic driving behaviors in Zhao et al. (2016).

2.3 The Almost Safe Set Quantification Problem

Let a scenario-sampling algorithm consecutively sample runs of scenarios on \mathcal{S} following the system dynamics (1). We are now ready to present the car-following safe set quantification problem as follows.

Problem 1. Given $\delta \in \mathbb{R}^n$, $\epsilon \in (0, 1]$, and $\beta \in (0, 1]$, a testing policy $\pi(\cdot)$ in the form of (4), and the corresponding car-following scenario system in the form of (1). Let $\mathcal{S}_0 \subseteq \mathcal{S}$ be the sup-set of all safe sub-sets in \mathcal{S} . The *car-following safe set quantification problem* seeks to find a scenario-sampling algorithm $\mathcal{ALG} : \mathcal{S} \times (0, 1] \times (0, 1] \times \mathbb{R}^n \rightarrow \mathcal{S}$, such that with a confidence level of at least $1 - \beta$, the result of $\mathcal{ALG}(\mathcal{S}_0, \epsilon, \delta, \beta)$ is an $\epsilon\delta$ -almost safe set for (1).

The previous work by Weng et al. (2021b) has already presented various algorithms that provably solve the above problem with a primary focus on completeness and asymptotic optimality properties. Such properties occur as the number of samples tends to infinity which leads to a significant number of samples being required in practice. In this paper, we propose a modified version of the Synchronous Pruning and Exploration for safe set quantification by Weng et al. (2021b) with a specific focus on the car-following regime. This leads to a theoretically sound and practically feasible safe set quantification solution as we shall see in the next two sections.

We conclude this section by addressing the following assumption and justifying its practical feasibility.

Assumption 1. Given the state space \mathcal{S} , the set of failure states \mathcal{C} , and the system (1), we assume that the run of scenario can be initialized from any $\mathbf{s} \in \mathcal{S} \setminus \mathcal{C}$.

In practice, if one can control the engagement of the subject ADAS sufficiently accurately, the above assumption is naturally feasible, such as in the test protocol by EuroNCAP (2019). On the other hand, if the ADAS is expected to engage before triggering the test, accurate initialization becomes more difficult for some states. In this situation, the above assumption is most easy to achieve at the control equilibrium sub-set of \mathcal{S} . For example, $v_0 = v_1$ for some $d \in \mathbb{R}_{>0}$, which denotes the steady-state car-following condition. This is also the initialization condition adopted by Forkenbrock and Snyder (2015). Some non-control equilibrium states can be initialized through customized procedures. For example, in the LVLCB test from the NHTSA report by Rao et al. (2019), the lead-vehicle on the side lane can choose to perform a lane change at any speed with any headway, which has the potential to initialize some non-control equilibrium states such as when $v_0 \gg v_1$. Note that even with the above techniques, some states are still difficult to initialize, such as $v_0 \gg v_1, d = 0$. However, those difficult-to-achieve initialization states are typically of obvious high-risk, thus they may not need to be tested anyway, as we shall see in Section 4.

3. MAIN METHOD

To solve Problem 1, the overall algorithm follows a two-step procedure. First, one continuously constructs a candidate set as more runs of scenarios are collected through scenario sampling. Second, as the constructed set becomes close to the actual almost safe set, one should observe a

sufficiently large number of runs of scenarios that start from and remain inside the candidate set. For the second step, the sampling sufficiency is justified by the following theorem.

Theorem 1. ($\epsilon\delta$ -Almost Safe Set Validation) Given the system dynamics (1), $\epsilon \in (0, 1]$, $\beta \in (0, 1]$, $\delta \in \mathbb{R}^n$, $\Phi \subseteq \mathcal{S}$, and the corresponding δ -covering set Φ_δ with centroids Φ_σ defined by Definition 1. Consider N runs of scenarios, $\{\mathcal{RS}_i(\mathbf{s}_0, K)\}_{i=1, \dots, N}$ ($K \in \mathbb{Z}, K \geq 2$), with the state initialization of each run being i.i.d. w.r.t. the underlying distribution on Φ_σ . The set Φ is the $\epsilon\delta$ -almost safe set for (1) with confidence level at least $1 - \beta$ if $\bigcup_{i=1}^N \mathcal{RS}_i(\mathbf{s}_0, K) \subseteq \Phi_\delta \cap \mathcal{C} = \emptyset$ and $N \geq \frac{\ln \beta}{\ln(1-\epsilon)}$.

That is, under the given conditions, if one consecutively observes $\frac{\ln \beta}{\ln(1-\epsilon)}$ runs of scenarios all remaining inside Φ_δ , one then can claim with a confidence level at least $1 - \beta$ that the probability of any trajectory starting from Φ_σ leaving Φ_δ is less than ϵ , i.e., the SV is $\epsilon\delta$ -almost safe in the set Φ . One can refer to Weng et al. (2021b) for the proof of Theorem 1.

The proposed algorithm to solve Problem 1 is presented as Algorithm 1 which relies on Theorem 1. Note that `pop`, `reachable`, `nearest`, `remove`, and `append` are all notional functions. $\mathcal{X}.\text{pop}()$ returns a point $\mathbf{x} \in \mathcal{X}$ and removes it from the set. `reachable(s, G)` returns all vertices on the graph G that connects, directly and indirectly, to the point \mathbf{s} through a depth-first-search routine (see Weng (2022)). $\mathcal{X}.\text{nearest}(\mathbf{x})$ returns the nearest point to \mathbf{x} in \mathcal{X} in terms of ℓ_2 -norm distance. The commands `remove` and `append` simply remove a point from or add a point to the given set, respectively.

Overall, Algorithm 1 consists of four major steps. The initialization step (line 2) configures two graphs, G_σ and G_u , that are intended to contain potentially safe and observed unsafe states and transitions, respectively, through scenario-sampling. The sampling step (line 4-7) takes a i.i.d. sample by Theorem 1 if the prioritized replay buffer \mathcal{B} is empty. Otherwise, i.e. when some unsafe states have been observed and added to \mathcal{B} at line 12, it prioritizes sampling points in Φ_σ that are close to the points in \mathcal{B} as they are intuitively of higher-risk. Such a sampling heuristic will not jeopardize the claimed property in Theorem 1 for set validation, as \mathcal{B} will be empty eventually, but will accelerate the convergence to a sufficiently almost safe set as unsafe points are removed more frequently. The third important stage happens at line 10-19. When a sampled run of a scenario is observed to converge to \mathcal{C} , any reachable states to the points in the collected run are removed from Φ_σ . On the other hand (line 21-32), one either adds an uncovered point to the covering set (line 23-25) or consecutively observes N runs of scenarios that remain inside Φ_δ to claim the $\epsilon\delta$ -almost safe property.

The proposed algorithm differs from the SPE for safe set quantification in Weng et al. (2021b) in two main ways, (1) the use of prioritized sampling with a replay buffer and (2) the removed stage of $\epsilon\delta$ decay. The prioritized sampling with a replay buffer is a heuristic approach that improves the convergence rate to a potentially almost safe set. The fixed choice of δ and ϵ compromises the probabilistic completeness of the algorithm in return for

practical feasibility with improved sampling efficiency (as we shall also see empirically in Section 4). One can always re-obtain the completeness and optimality properties, or at least achieve an appropriate level of compromisation, by configuring δ and ϵ to be arbitrarily close to zero, yet the number of required samples might also increase dramatically.

Algorithm 1 Car-following Safe Set Quantification

```

1: Input: Initial set  $\mathcal{S}_0 \subseteq \mathcal{S}$ , collision set  $\mathcal{C}$ ,  $\epsilon \in (0, 1]$ ,
    $\beta \in (0, 1]$ , trajectory horizon  $K$ .
2: Initialize: The  $\delta$ -covering set of  $\mathcal{S}_0$ ,  $\Phi_\delta$ , and cen-
   troids  $\Phi_\sigma$  by Definition 1, the state graph  $G_\sigma =$ 
    $(\Phi_\sigma, E_\sigma), E_\sigma = \emptyset \subset \mathcal{S}^2$ , the unsafe state graph  $G_u =$ 
    $(\mathcal{D}_u, E_u), \mathcal{D}_u = \emptyset \subset \mathcal{S}, E_u = \emptyset \subset \mathcal{S}^2$ , prioritized replay
   buffer  $\mathcal{B} = \emptyset, N=0$ .
3: While  $N < \frac{\ln \beta}{\ln(1-\epsilon)}$ :
4:   If  $\mathcal{B} = \emptyset$ 
5:      $\mathbf{s}_0 \sim P(\Phi_\sigma)$ 
6:   Else
7:      $\mathbf{s}_b = \mathcal{B}.\text{pop}(), \mathbf{s}_0 = \Phi_\sigma.\text{nearest}(\mathbf{s}_b)$ 
8:   End If
9:   Get  $\mathcal{T} = \mathcal{RS}(\mathbf{s}_0, K)$ 
10:  If  $\mathcal{T} \cap \mathcal{C} \neq \emptyset$ 
11:    For  $i$  in  $\mathbb{Z}_{|\mathcal{T}|-1}$  do
12:       $\mathcal{B}.\text{append}(\mathcal{T}[i])$ 
13:      For  $\mathbf{s}$  in  $\text{Reachable}(\mathcal{T}[i], G_\sigma)$  do
14:         $\Phi_\sigma.\text{remove}(\mathbf{s})$ 
15:         $E_u.\text{append}((\mathcal{T}[i], \mathcal{T}[i+1]))$ 
16:      End For
17:       $\mathcal{B}.\text{append}(\mathcal{T}[i+1])$ 
18:    End For
19:     $N = 0$ 
20:  Else
21:     $\bar{\mathbf{s}} = \mathbf{s}_0, N_s = |\Phi_\sigma|$ 
22:    For  $i$  in  $\{2, \dots, |\mathcal{T}|\}$  do
23:      If  $\mathcal{T}[i] \notin \Phi_\delta$ 
24:         $E_\sigma.\text{append}((\bar{\mathbf{s}}, \mathcal{T}[i]))$ 
25:         $\bar{\mathbf{s}} = \mathcal{T}[i]$ 
26:      End If
27:    End For
28:    If  $N_s = |\Phi_\sigma|$  and  $\mathcal{B} = \emptyset$ 
29:       $N_s = N_s + 1$ 
30:    Else
31:       $N = 0$ 
32:    End If
33:  End If
34: Output:  $\Phi_\delta$ 

```

4. CASE STUDIES

To demonstrate the performance of the proposed Algorithm 1, we start with examples of safety evaluations of deterministic decision-making systems where the perception and the control modules are both sufficiently accurate. We then move to an end-to-end case study taking the CommaAI's Openpilot by Shihadeh et al. (2018) as an example which involves a neural-network based perception module, camera-radar sensor fusion, model-based decision-making, and control modules. The source code for Algorithm 1 in Python can be found at Weng (2022).




Fig. 1. Some $\epsilon\delta$ -almost safe sets, presented as slices, obtained for the car-following case study with various decision-making modules ($\epsilon = 0.01, \beta = 0.001$): (a) ACC-AEB with $\delta = [10, 2, 2]$, (b) ACC-AEB with $\delta = [10, 6, 6]$, (c) N_IDM with $\delta = [10, 6, 6]$, (d) M_IDM with $\delta = [10, 6, 6]$, (e) H_IDM with $\delta = [10, 6, 6]$.

4.1 Decision-Making Safety Evaluation

We consider two classes of decision making systems in this section. The first is a combination of ACC and AEB (ACC-AEB) first introduced by Zhao et al. (2016). When the perceived time-to-collision value is greater than a pre-determined threshold, the ACC module is engaged as a discrete Proportional-Integral (PI) controller to achieve a desired time headway. Otherwise, the AEB module extracted from a 2011 Volvo V60 is active. Our ACC-AEB module takes the same hyper-parameters and configuration as Zhao et al. (2016), having a maximum braking capability of -10m/s^2 subject to a deceleration change rate limit of -16m/s^3 . The second decision-making module studied by this section is the Intelligent Driving Model (IDM) in Treiber and Kesting (2013), which is a widely adopted car-following model in the field. Note that we have created three IDM variants based on the maximum brake control capability. In particular, we have the normal-brake IDM (N_IDM) with -5m/s^2 , the mild-brake IDM (M_IDM) with -3m/s^2 , and the hard-brake IDM (H_IDM) with -7m/s^2 . Other IDM parameters include the minimum safe distance (2 m), maximum acceleration (0.73m/s^2), comfortable deceleration (1.67m/s^2), safe time headway (2 s), exponent of acceleration (4), and vehicle length (4 m). Unless mentioned otherwise, we consider the state space \mathcal{S} with headway $d \in [0, 100]$ m, SV speed $v_0 \in [0, 30]$ m/s, and lead POV speed $v_1 \in [0, 30]$ m/s. Note that the collected run of a scenario might leave \mathcal{S} with a large headway value that is greater than the given upper bound (100 m), in which case one shall either truncate the trajectory or saturate the headway value at the given upper bound before proceeding to line 10 of Algorithm 1. The simulation of each run of scenario operates at 10 Hz with $K = 300$. The testing policy admits the form of $\pi(\mathbf{s}) = -5\text{m/s}^2, \forall \mathbf{s} \in \mathcal{S}$. We also assume the free-traffic speed to be 30 m/s.

We execute Algorithm 1 for 10 times with 10 different random seeds. The set of 10 seeds remains the same among different SVs. Some of the obtained almost safe sets for $\epsilon = 0.01, \beta = 0.001$ are illustrated in Fig 1 for the same seed. The three-dimensional safe set is illustrated with a series of subplots on the (v_0, v_1) domain, each representing a subspace slicing of a certain headway value. Intuitively, one expects the size of the safe set to increase as the lead-POV becomes further away since the state is of lower-risk as the lead-POV operates at a higher speed than the SV follower. This is mostly correct if one observes the IDM cases where M_IDM has the smallest almost safe set and H_IDM has the largest almost safe set, which aligns




Fig. 2. M_IDM's commanded acceleration inputs for a group of (v_0, v_1) pairs at 40-meter headway in the car-following scenario.

with the underlying configurations of M_IDM having the lowest braking capability and H_IDM having the strongest braking capability among all tested IDMs.

However, for most of the subplots in the ACC-AEB case, especially those with large headway values, one exhibits a non-convex almost safe set with a white notch, which indicates some unsafe states even when the headway is sufficiently large. This is mainly due to the ACC design nature where one tends to reach the free-traffic speed aggressively when the headway value is high, v_0 thus increases, ending up in a certain unsafe state. For a similar cause, ACC-AEB also fails all of the CCRb and CCRm tests in Fig 6. As a result, if one considers the free-traffic speed as an observable state and expands the \mathcal{S} to be of dimension four, the corresponding almost safe set will also change w.r.t. the desired velocity. A detailed analysis regarding this variant, and possibly other variants considering different added features, are of future interest.

Returning to the notch observation, why isn't a similar shape showing up on any of the IDM variants in Fig 1? This is because the IDM is primarily a car-following model and may not necessarily exhibit expected behaviors outside the normal car-following work domain. For example, Fig 2 illustrates the M_IDM's acceleration outputs for a group of (v_0, v_1) pairs with 40-meter headway. Note that at $v_0 = 12$ m/s, $v_1 = 25$ m/s, the M_IDM decides to execute maximum brake maneuver, rather than to accelerate to track the desired speed. This leads to a utility performance degradation in terms of velocity tracking, but on the other hand, improves the safety performance against potential rear-end collisions. Fundamentally speaking, the observed phenomena is caused by a squared term associated with the $(v_0 - v_1)$ term in the IDM formulation, the details are beyond the scope of this paper.

Table 1. Safety evaluation results for various decision-making modules in the car-following case ($\beta = 0.001, \delta = [10, 6, 6]$) of Section 4.1 and Openpilot presented in Section 4.2 ($\beta = 0.001, \delta = [3, 3, 3]$).

SV	\mathcal{S}_0	ϵ	scenario runs	collision runs	IoU
ACC-AEB	\mathcal{S}	0.1	867.5 ± 281.2	268.3 ± 34.5	0.915
	\mathcal{S}	0.01	1912.6 ± 146.4	185.4 ± 1.4	1.000
H_IDM	\mathcal{S}	0.1	194.2 ± 14.6	40.7 ± 3.2	0.965
	\mathcal{S}	0.01	1376.0 ± 182.1	49.0 ± 0.0	1.000
N_IDM	\mathcal{S}	0.1	368.5 ± 95.0	69.6 ± 4.4	0.952
	\mathcal{S}	0.01	1628.8 ± 266.3	74.6 ± 0.8	0.998
	Fig 1e	0.01	1578.6 ± 220.1	26 ± 0.0	1.000
M_IDM	\mathcal{S}	0.1	830.9 ± 88.3	155.6 ± 3.7	0.956
	\mathcal{S}	0.01	1892.6 ± 237.5	161.0 ± 0.0	1.000
	Fig 1e	0.01	1731.4 ± 125.5	112.0 ± 0.0	1.000
Openpilot	\mathcal{S}	0.1	704.2 ± 54.3	141.8 ± 3.4	0.897

Moreover, comparing Fig 1(a) and Fig 1(d), the ACC-AEB has a relatively larger safe set than M.IDM when the headway value is small. As the headway value increases, the safe set of M.IDM enlarges significantly and eventually out-performs ACC-AEB in terms of the safe set size. That is, the simplistic notion of “one vehicle being safer than the other” can be problematic as it is essentially a multi-dimensional comparison. A similar point was also made by Weng et al. (2022) through observing real-world car-following performance in the naturalistic driving environment. Such a subtle safety characterization is difficult to obtain by existing concrete scenario-based testing strategies such as the NCAP AEB testing shown in Fig 6.

More detailed results regarding this case are listed in Table 1. The “IoU” denotes the intersection-over-union ratio of all obtained safe sets from different seeds w.r.t. the same SV. It is clear that the higher the IoU value, the more similar the obtained sets are among different seeds. Considering that the studied decision-making modules in this section are both deterministic, the IoU value should converge to one for sufficiently small ϵ and β . This has been validated empirically by row 2, 4, 7, 9, and 10 in Table 1. We emphasize that even for the cases with IoU values less than one, the results are not wrong, as the $\epsilon\delta$ -almost safe set is simply not unique for the studied system. Also, note that if the set initialisation is not S_0 but is another set that is closer to the final almost safe set, one should expect a smaller number of runs of scenarios and, more importantly, a smaller number of runs of scenarios with collisions, to converge to the desired outcome (e.g. comparing row 6 with row 7, and comparing row 9 with row 10 in Table 1).

Overall, the total number of runs of scenarios varies w.r.t. the SV, the selected hyper-parameters (e.g. ϵ, β) and the random seed but remains below 2000, which is less than 17-hours of actual scenario-running time (2000 runs of scenarios with at most 30 seconds for each run) excluding the testing preparation and scenario restoration time. This is a slightly higher testing burden than the existing standards for the car-following regime but should still be considered feasible in practice. Moreover, the testing effort may be further reduced for a smaller K , and the exploration regarding this direction is of future interest. More importantly, among the methods that are capable of providing similar theoretical guarantees, the proposed solution appears to be the most practical and is capable of capturing the subtle differences among various SVs. For comparison, the importance sampling and Monte-Carlo sampling based methods reported by Zhao et al. (2017) require hundreds of millions of test runs in simulation for safety evaluation with car-following maneuvers and only generate a risk estimate.

4.2 End-to-End Safety Evaluation

For an end-to-end case study, we evaluate the CommaAI Openpilot’s safety performance in the car-following regime through simulation using the Carla simulator. To run the Openpilot in Carla, we use the Openpilot-Carla bridge provided by CommaAI as a foundation to which we have added clustered radar results for radar-camera fusion needed to enable the ACC in Openpilot. The radar points




Fig. 3. Some $\epsilon\delta$ -almost safe sets, presented as slices, obtained for the car-following case study with Openpilot for three different random seeds ($\epsilon = 0.1, \beta = 0.001$)




Fig. 4. The $\epsilon\delta$ -almost safe sets ($\epsilon = 0.01, \beta = 0.001$) obtained for ACC-AEB and Openpilot in a lead-obstacle scene where the lead-POV remains stationary for all time.




Fig. 5. The trajectories on the (v_0, d) domain of Openpilot tested in two standard NCAP Car-to-Car Rear moving scenarios. For both scenarios, the lead POV remains at 20 km/h (5.56 m/s). All other parameters and environmental configurations remain identical among all test runs for the same initialization condition. Within each subplot, Openpilot is enabled at the illustrated initialization state and both vehicles, unless specified otherwise by the testing procedure, remain at the steady-state stage with zero acceleration.

clustering configuration is identical to the work by Zhong et al. (2021). The detailed implementation can be found at Zhu (2022). The state space S takes the configuration $d \in [0, 30]$ m, $v_0 \in [0, 15]$ m/s, and $v_1 \in [0, 15]$ m/s. The simulation of each run of scenario operates at 100 Hz with $K = 500$. The free-traffic speed is 11.176 m/s (25 mph) if $v_0(0) < 11.176$ and $v_0(0)$ otherwise, which is the default configuration of Openpilot.

Note that Openpilot is not designed for emergency collision avoidance as indicated by CommaAI at Shihadeh et al. (2018). It is primarily a car-following model. As a result, an adversarial testing policy such as the one adopted for the decision-making case could lead to a very limited safe set. For example, as shown in Fig 4, if the lead vehicle remains stationary (similar to the CCRs case by EuroNCAP (2019) and also included in Fig 6), the Openpilot SV rarely avoids any rear-end collisions if $v_0 \geq 4.5$ m/s. The Openpilot’s almost safe set is also significantly smaller than a regular almost safe set in cases such as the one shown for ACC-AEB in Fig 4a. In this section, we admit the testing policy as $\pi(s) = 0$ m/s², which emulates the steady-state car-following situation.

We execute Algorithm 1 for 5 times with 5 different random seeds. Some of the obtained almost safe sets for $\epsilon = 0.1, \beta = 0.001$ with three different seeds are illustrated




Fig. 6. The testing outcomes of all studied SVs in Section 4 with the standard NCAP car-to-car AEB testing procedure discussed in EuroNCAP (2019). The procedure specifies 48 different scenario configurations from three categories including the Car-to-Car Rear stationary (CCRs), Car-to-Car Rear moving (CCRm), and Car-to-Car Rear braking (CCRb), where the lower-case letter after CCR induces the lead-POV’s driving behavior (staying stationary, moving at a constant velocity, or braking to stop). Each deterministic decision-making module is only tested once. The Openpilot enabled SV is tested with the same set of 48 scenarios for 10 times. The detailed parameters related to the order of all testing cases can be found in “ncap_bridge.py” at Zhu (2022).

in Fig 3. Other statistical properties are summarized in the last row of Table 1. Note that the IoU rate in Table 1 is slightly smaller than the presented cases in Section 4.1. This is mainly due to the fact that Openpilot is fundamentally stochastic, as also illustrated by Fig 5 and Fig 6 where, starting from the same s_0 , the Openpilot enabled SV is shown capable of generating both safe and collision outcomes. As a result, the almost safe set for Openpilot in the studied domain is fundamentally non-unique, making it a particularly challenging case for many existing scenario-based techniques and surrogate safety metrics. As for the proposed method, the obtained safe set aligns with the claimed operational domain by CommaAI. The SV remains safe with high probability when $v_0 \geq v_1$ regardless of the following distance. The size of the almost safe set also increases as the headway value becomes larger.

5. CONCLUSION

In this paper, we have presented a theoretically sound and sampling efficient scenario-sampling framework for the safety performance evaluation of various car-following and rear-end collision avoidance systems. The performance of the proposed method has been demonstrated empirically through a series of challenging cases. It is of future interest to improve the completeness of the formulated scenario state space and develop more sampling-efficient safe set quantification algorithms. The proposed method is also expected to generalize to the safety evaluation of other cooperative car-following systems and human drivers within the same operable domain.

REFERENCES

EuroNCAP (2019). European new car assessment programme (euro ncap) test protocol – AEB car-to-car systems. Technical report, The European New Car Assessment Programme.

Fan, C., Qi, B., Mitra, S., and Viswanathan, M. (2017). D ry vr: data-driven verification and compositional reasoning for automotive systems. In *International Conference on Computer Aided Verification*, 441–461. Springer.

Feng, S., Feng, Y., Yan, X., Shen, S., Xu, S., and Liu, H.X. (2020). Safety assessment of highly automated driving systems in test tracks: a new framework. *Accident Analysis & Prevention*, 144, 105664.

Forkenbrock, G.J. and Snyder, A.S. (2015). NHTSA’s 2014 automatic emergency braking test track evaluations. Technical report, National Highway Traffic Safety Administration.

Rao, S.J., Forkenbrock, G.J., et al. (2019). Test procedures traffic jam assist test development considerations. Technical report, United States. Department of Transportation. National Highway Traffic Safety Administration.

Shihadeh, A. et al. (2018). openpilot. <https://github.com/commaai/openpilot>.

Treiber, M. and Kesting, A. (2013). Traffic flow dynamics. *Traffic Flow Dynamics: Data, Models and Simulation*, Springer-Verlag Berlin Heidelberg.

Wang, C., Xie, Y., Huang, H., and Liu, P. (2021). A review of surrogate safety measures and their applications in connected and automated vehicles safety modeling. *Accident Analysis & Prevention*, 157, 106157.

Weng, B. (2021). A class of model predictive safety performance metrics for driving behavior evaluation. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, 180–187. doi:10.1109/ITSC48978.2021.9565013.

Weng, B. (2022). SDQ tools. https://gitlab.com/Bobeye/sdq_tools.

Weng, B., Capito, L., Ozguner, U., and Redmill, K. (2021a). A formal characterization of black-box system safety performance with scenario sampling. *IEEE Robotics and Automation Letters*. doi:10.1109/LRA.2021.3122517.

Weng, B., Capito, L., Ozguner, U., and Redmill, K. (2022). A finite-sampling, operational domain specific, and provably unbiased connected and automated vehicle safety metric. *IEEE Transactions on Intelligent Transportation Systems*. doi:10.1109/TITS.2022.3164358.

Weng, B., Capito Ruiz, L.J., Ozguner, U., and Redmill, K. (2021b). Towards guaranteed safety assurance of automated driving systems with scenario sampling: An invariant set perspective. *IEEE Transactions on Intelligent Vehicles*. doi:10.1109/TIV.2021.3117049.

Wishart, J., Como, S., Elli, M., Russo, B., Weast, J., Altekar, N., James, E., and Chen, Y. (2020). Driving safety performance assessment metrics for ads-equipped vehicles. *SAE Technical Paper*, 2(2020-01-1206).

Zhao, D., Huang, X., Peng, H., Lam, H., and LeBlanc, D.J. (2017). Accelerated evaluation of automated vehicles in car-following maneuvers. *IEEE Transactions on Intelligent Transportation Systems*, 19(3), 733–744.

Zhao, D., Lam, H., Peng, H., Bao, S., LeBlanc, D.J., Nobukawa, K., and Pan, C.S. (2016). Accelerated evaluation of automated vehicles safety in lane-change scenarios based on importance sampling techniques. In *IEEE Transactions on Intelligent Transportation Systems*, volume 18, 595–607. IEEE.

Zhong, Z., Hu, Z., Guo, S., Zhang, X., Zhong, Z., and Ray, B. (2021). Detecting safety problems of multi-sensor fusion in autonomous driving. *arXiv preprint arXiv:2109.06404*.

Zhu, M. (2022). Openpilot in Carla. https://github.com/pgchui/openpilot_in_carla.