

Simulator Ensembles for Trustworthy Autonomous Driving Systems Testing

Lev Sorokin ¹  · Matteo Biagiola ²  ·
Andrea Stocco ^{1,3} 

Received: date / Accepted: date

Abstract Scenario-based testing with driving simulators is extensively used to identify failing conditions of automated driving assistance systems (ADAS) and reduce the amount of in-field road testing. However, existing studies have shown that repeated test execution in the same as well as in distinct simulators can yield different outcomes, which can be attributed to sources of flakiness or different implementations of the physics, among other factors.

In this paper, we present **MultiSim**, a novel approach to multi-simulation ADAS testing based on a search-based testing approach that leverages an ensemble of simulators to identify failure-inducing, simulator-agnostic test scenarios. During the search, each scenario is evaluated jointly on multiple simulators. Scenarios that produce consistent results across simulators are prioritized for further exploration, while those that fail on only a subset of simulators are given less priority, as they may reflect simulator-specific issues rather than generalizable failures.

Our empirical study, which involves testing three lane-keeping ADAS with increasing complexity on different pairs of three widely used simulators, demonstrates that **MultiSim** outperforms single-simulator testing by achieving, on average, a higher rate of simulator-agnostic failures by 66%. Compared to a state-of-the-art multi-simulator approach that combines the outcome of independent test generation campaigns obtained in different simulators, **MultiSim** identifies on average up to 3.4× more simulator-agnostic failing tests and higher failure rates.

To avoid the costly execution of test inputs on which simulators disagree, we propose an enhancement of **MultiSim** that leverages surrogate models to

¹ Technical University of Munich, Munich, Germany. Email: lev.sorokin@tum.de, andrea.stocco@tum.de

² University of St. Gallen and Università della Svizzera italiana, St. Gallen / Lugano, Switzerland. Email: matteo.biagiola@unisg.usi.ch

³ fortiss GmbH, Munich, Germany. Email: stocco@fortiss.org

predict simulator disagreements and bypass test executions. Our results show that utilizing a surrogate model during the search does not only retain the average number of valid failures but also improves its efficiency in finding the first valid failure. These findings indicate that combining an ensemble of simulators during the search is a promising approach for the automated cross-replication in ADAS testing.

Keywords. Search-based software testing, scenario-based testing, autonomous driving, testing deep learning systems, simulator-agnostic

1 Introduction

Advanced Driver Assistance Systems (ADAS) rely on perception systems like cameras and LiDAR, powered by deep neural networks (DNNs), for real-time tasks such as lane-keeping, object detection and image segmentation. While effective, these systems must operate reliably across diverse environments, yet training data cannot cover all possible scenarios [59, 87]. Consequently, ADAS may encounter unseen inputs post-deployment, making DNNs highly sensitive to variations in road shapes, lighting, and noise. These discrepancies can lead to ADAS prediction errors, misclassifications, and inaccurate segmentations, which impact vehicle decision-making and compromise safety [24, 50].

Validating the safety of ADAS through virtual testing with scenario-based simulation is the default option for companies [24, 87]. Simulators enable developers to quickly prototype ADAS and evaluate them across a wide range of challenging scenarios. In the literature, researchers have proposed automated testing techniques to expose failing conditions and corner cases [1, 33, 47, 63, 64, 70, 76, 98], using various open-source ADAS simulators, such as CARLA [30], BeamNG [12], and Udacity [94], or commercial close-source solutions, such as Siemens PreScan [78], ESI Pro-SiVIC [36], and PTV VISSIM [95].

However, the result of a test execution through a simulation environment only approximates the actual test outcome in the real world. Indeed, multiple studies have shown that simulation-based testing might produce inconsistent results both within simulator, i.e., when the same test scenario is executed multiple times on the same simulator [5, 18], and across simulators [16, 20]. This phenomenon may be due to test flakiness, and it is characterized by multiple runs of the same test exhibiting non-deterministic behavior (i.e., tests pass or fail non-deterministically). Test flakiness has been studied extensively in the software testing literature [73], including ADAS testing. A study by Amini et al. [5] shows that test flakiness is quite common also for simulation-based testing of ADAS, potentially leading to a distrust in virtual testing, as also reported in the field of robotic simulation [3]. Possible causes of flakiness stem from uncertainties in the simulation environment to timing and synchronization in the interaction between the ADAS and the simulator [5]. In general, such sources of non-determinism are unknown and difficult to control, making the results of automated testing techniques, i.e., the failure-inducing test cases, unreliable.

Despite the significance of the problem, very few solutions have been proposed to mitigate test flakiness in simulation-based testing (beyond discarding flaky tests [18]). Amini et al. [5] used machine-learning classifiers to predict the flakiness of an ADAS test using few test runs. Another way to mitigate test flakiness is by combining the outcome of multiple simulators when performing a testing campaign. In particular, Biagiola et al. [16], proposed a multi-simulator approach to approximate the outcome of a test case on a digital twin. Their framework named *digital siblings* (DSS henceforth) executes two independent test generation algorithms on two simulators, generating two feature maps that characterize test cases executed on such simulators. Through the operations of migration and merge, the framework outputs a combined feature map that prioritizes *agreements* across simulators, i.e., when the same test case has the same outcome on different simulators. The intuition is that we can expect test cases where the two simulators agree to be more reliable, and hence less flaky, than those where the two simulators disagree. Although their results show that the combination of the two simulators is able to predict the failures of the digital twin better than each individual one, the drawback of the digital siblings framework is that test case execution outcomes on the two simulators are merged as a post-processing step. This way, the test generation algorithm runs the risk of evolving flaky tests as each test is evaluated only on a single simulator during the search, making the whole process inefficient.

In this paper we present a novel approach named **MultiSim**, which incorporates multiple simulators as an ensemble directly as part of the test generation process. In particular, we cast the testing problem as a multi-objective optimization problem, where we define, for each simulator, a fitness function that evaluates the quality of the ADAS under test. In this way, we treat simulator disagreement as a first-class citizen during the optimization process. By minimizing the fitness values obtained from executing each test case in multiple simulators and their distance, we direct the search towards failure-inducing tests that are *simulator-agnostic*, i.e., reliable, tests since the respective outcomes on the two simulators are close to each other.

We evaluate our **MultiSim** approach with three lane-keeping ADAS based on convolutional DNNs [19], transformers [11] and end-to-end trajectory/control prediction [97], on three different and widely used simulators, namely Udacity, Donkey and BeamNG [5, 33, 42, 76, 86, 102]. In particular, we compare **MultiSim** with a single simulator approach, as well as with the DSS framework in terms of effectiveness, i.e., the number of simulator-agnostic failures each approach triggers given a fixed search budget, and efficiency, i.e., how quickly each approach generates the first simulator-agnostic failure.

Our results show that **MultiSim** identifies on average 70% valid failures across all simulator configurations, outperforming DSS with 65% and single simulator-based testing with 47%. Regarding the rate of simulator-agnostic failures, BD (i.e., BeamNG combined with Donkey) outperforms both single-simulator testing and DSS-based testing, except for DSS-BD, where it has a similar validity rate but identifies a significantly higher number of failures. Using an ensemble of simulators during search-based testing demonstrates

efficiency comparable to testing with only one simulator. Compared to **DSS**, in the **BD** configuration, it achieves significantly better results, while in another comparison, i.e., **BU** (BeamNG combined with Udacity) is outperformed by **DSS-BU**. Finally, integrating a Random Forest classifier to predict simulator-related disagreements into **MultiSim** improves the effectiveness, leading to a higher median of simulator-agnostic failures, as well as a lower standard deviation, compared to running the search without the classifier. Similarly, in terms of efficiency, we can identify a smaller median and variation in the search budget required for finding the first valid failure.

Our paper makes the following contributions:

- **Approach.** A novel search-based test generation approach called **MultiSim** that integrates multiple simulators in an ensemble to identify simulator-agnostic failures for ADAS.
- **Evaluation.** An extensive empirical study using three simulators and three different approaches showing that **MultiSim** is effective at generating simulator-agnostic failures for three state-of-the-art lane-keeping ADAS. To encourage open research, our test generation approach and experimental data are available [91].

The outline of the paper is as follows: In Section 2, we present the preliminaries. In Section 3, we introduce the problem statement and the notion of simulator-agnostic failures. In Section 4, we present our test generation approach, followed by our validation strategy. In Section 5, we evaluate our testing approach in comparison with single-simulator based testing and a default multi-simulator based testing approach. In Section 6, we discuss the insights and clarifications regarding our approach. In Section 7, we outline the most important threats regarding the validity of our results. In Section 8, we present related work. We finally conclude in Section 9, with a summary and insights into future work.

2 Background

2.1 Simulation-based ADAS Testing

We consider Level 2 ADAS, as classified by the National Highway Traffic Safety Administration (NHTSA) [92], which perform vision-based perception tasks using data gathered by camera sensors of a vehicle. Despite the adoption of Level 2 ADAS in many commercial vehicles, their reliability remains a concern, as evidenced by numerous recent crash reports [93] and large scale real-world validation experiments [71]. Although Levels 3 and 4 ADAS have been proposed [10], their real-world deployment remains highly constrained. Consequently, addressing the limitations of Level 2 systems is crucial for advancing to higher levels of autonomy. Particularly, in this paper we focus on ADAS that learn the *lane-keeping* functionality, a critical component for the safe operation of self-driving vehicles, from human-labeled driving data.

In the early stages of development, ADAS undergo model-level testing [75]. This involves evaluating performance metrics such as accuracy, mean squared error (MSE), and mean absolute error (MAE) on an unseen test set, i.e., a dataset not used during training. This form of testing is analogous to unit testing in traditional software development, helping test engineers identify inadequately trained models [38].

Following model-level testing, ADAS are subjected to system-level testing [27, 37, 84]. This phase assesses the impact on the entire decision-making process of an ADAS, ensuring that its predictions align with expected vehicle behavior. A typical system-level test, i.e., test scenario, involves generating a one-lane or two-lane road, each defined by a starting and an endpoint, with varying length, curvature, and number of turns [4]. While driving, the ADAS processes input images and generates steering commands. System failures occur when the vehicle deviates from system requirements, such as violating safety constraints, e.g., driving off the road or causing harm to other vehicles, the environment, or pedestrians [86]. These failures often stem from errors in the perception component [50, 51]. The testing objective is therefore to identify road topologies where the ADAS fails to maintain the vehicle in lane, either by driving off-road (for one-lane scenarios) or crossing the opposite lane (for two-lane scenarios).

System-level ADAS testing is primarily conducted in simulation environments using software-in-the-loop testing to ensure safety and minimize costs. This approach enables safe measurement, analysis, characterization, and reproduction of driving failures. Well-established ADAS simulation platforms include industrial platforms such as Siemens PreScan [78] or ESI Pro-SiVIC [36], whereas open-source solutions widely adopted by researchers include Udacity [94], Donkey Car’s sdsandbox [29] (hereafter referred to as Donkey, for simplicity), and BeamNG [12].

2.2 Search-based Software Testing

Our framework uses the concepts behind search-based testing (SBST), a testing technique where the testing problem is modeled as an optimization problem to be solved with metaheuristic optimization techniques [65]. SBST is defined as follows:

Definition 1 A search-based testing problem P is defined as a tuple $P = (T, D, F, O)$, where

- T is the system under test.
- $D \subseteq \mathbb{R}^n$ is the search domain, where n is the *dimension* of the search space. The vector $\mathbf{x} = (x_1, \dots, x_n) \in D$ is called test input.
- F is the vector-valued fitness function defined as $F : D \mapsto \mathbb{R}^m$, $F(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_m(\mathbf{x}))$, where f_i is a scalar fitness function (or fitness function, for short) that assigns a quantitative value to each test input and \mathbb{R}^m is the *objective space*, and m corresponds to the number of fitness functions.

A fitness function evaluates how *fit* a test input is, assigning a *fitness value* to it.

- O is the oracle function, $O : \mathbb{R}^m \mapsto \{0, 1\}$, which evaluates, given the objective space of the fitness functions, whether a test passes or fails. A test that fails is called *failure-inducing*.

For instance, a search based testing problem for the ADAS system from Section 2.1 is defined as follows: roads represent test inputs passed to the ADAS and the execution of the ADAS on a road is performed in a simulator yielding location traces. Location traces are used by the fitness function to calculate to each executed test input the maximal distance of the vehicle to the center line. The oracle function specifies that an execution is failing when the vehicle drives off the lane when the fitness function value is above a predefined threshold.

In our approach, solving such a search-based testing problem requires the definition of multiple fitness functions. We use the concepts of multi-objective optimization to solve the problem.

Definition 2 In search-based testing, a multi-objective optimization (MOO) problem is defined as:

$$\min_{\mathbf{x} \in X} F(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_m(\mathbf{x}))$$

where f_i is a scalar fitness function and $X \subseteq D$ is called the feasible solution set. In general, a set of equality and inequality constraints is defined, which have to be satisfied by solutions in X . A solution \mathbf{x} with $F(\mathbf{x}) = (v_1, \dots, v_m)$ is said to *dominate* another solution \mathbf{x}' with $F(\mathbf{x}') = (v'_1, \dots, v'_m)$ $\iff \exists v_i. (v_i < v'_i) \wedge \forall v_j. (v_j \leq v'_j)$, i.e., \mathbf{x} is superior to \mathbf{x}' in at least one fitness value and at least as good in all other fitness values. Consequently, a solution \mathbf{x} is called Pareto optimal if no solution exists that dominates it. The set of all Pareto optimal solutions is called Pareto set PS .

3 Problem Definition

3.1 Non-determinism in Virtual Testing

Recent studies have shown that there are several simulation platforms available for system-level testing, both commercial and open-source [49, 57], with no consolidated omni-comprehensive solution. As multiple simulation platforms exist, researchers have started performing cross-replication studies of ADAS in different simulation platforms [5, 20], to confirm the testing results obtained on a driving simulator on another, possibly independent, simulator. Ideally, when we execute the test input on the same ADAS in the different simulators, the pass/fail verdict should not change, providing trustworthiness to simulation-based testing. However, the results of these studies highlight the negative aspects of simulation-based testing for ADAS. In most cases, the

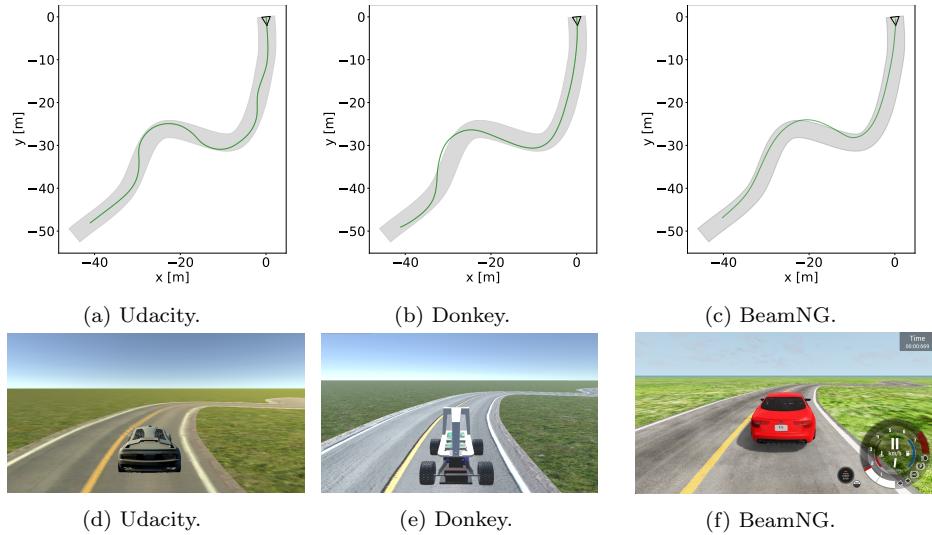


Fig. 1: Difference in executing a lane-keeping ADAS on the same road in three different simulators, along with their rendering. In Udacity the trajectory of the vehicle (visualized in green, starting with a triangle) is within the lane’s bounds while in Donkey and BeamNG the vehicle is departing off lane.

results obtained in one simulator (e.g., the failure conditions) cannot be reproduced by another simulator [5, 16, 20]. These discrepancies can lead to a distrust in simulation-based testing, as reported by recent surveys [3, 34].

An example is shown in Figure 1, in which we replicated the execution of a logical test scenario for a lane-keeping ADAS model on three simulators, namely Udacity [94], Donkey [88] and BeamNG [12]. The logical test scenario includes a two-lane 70 meters long road with three curves, one on the right with curvature 0.24, followed by one on the left with curvature 0.22, and one on the right with curvature 0.08. The objective for the ADAS model is to keep the car within the right lane from the start of the road until the end. The figure depicts only the right lane and shows that the driving behavior of the lane-keeping ADAS model (shown by the green line trajectory) is quite different across simulators. While the car reaches its destination in all three examples, it yields different simulation traces, specifically failing in the Donkey and BeamNG simulators but not in Udacity.

Potential root causes are related to flakiness in the simulation environments, bugs in the simulator, or synchronization issues in the communication between the simulator and the testing framework [5, 20, 49, 57]. As the simulation platform is used as-is, as a runtime testbed, it is usually not possible for developers to debug such issues, even less try to fix the simulator. This motivates the development of approaches to increase the reliability of simulation-based testing of ADAS, without requiring access to the simulation environments or its code, which is treated it as a black-box.

3.2 Simulator-agnostic Failures

In this paper, we focus on identifying test scenarios where failures are not caused by the limitations of the simulation environment but rather by deficiencies in the system under test. To achieve this, we introduce the concept of *simulator-agnostic failures*.

Let us consider an ADAS \mathcal{A} and a failing test input \mathbf{x} that can be executed in a simulation environment \mathcal{S} . We define \mathbf{x} as *simulator-agnostic* failure-inducing test input if \mathcal{A} also fails on \mathbf{x} when executed in a different simulator \mathcal{S}' with equivalent capabilities. In other words, the execution outcome of \mathbf{x} should not depend on compatibility issues between \mathcal{A} and the simulators, i.e., \mathcal{S} and \mathcal{S}' . This concept can naturally be extended to multiple simulators. In this study, we consider three simulation environments and assume that if a failure is consistently observed across all three, it indicates a true defect in the ADAS, as it is independent of any single simulator. Conversely, failures that occur only in specific simulators are regarded as *simulator-dependent*.

Following the established terminology from the deep learning testing domain [77], we refer to simulator-agnostic failures as *valid* failures, while we refer to simulator-dependent failures as *invalid* failures throughout this paper.

4 Approach

4.1 Ensemble-based Test Case Generation

Our proposed **MultiSim** approach incorporates the evaluation capabilities of multiple simulators to identify simulator-agnostic failures. The key idea is to apply ensemble learning within search-based ADAS testing, harnessing the strengths and compatibilities of individual simulators within a unified processing unit. We reformulated the multi-objective optimization problem to incorporate simulator-specific agreement and disagreement objectives, allowing the search to explicitly reason about the consistency of failures across heterogeneous simulators. Instead of evaluating a test input in one simulator, **MultiSim** combines the test input evaluation results of multiple simulators in one evaluation vector. Our approach aims to obtain more reliable, and hence more valid, failure-inducing test inputs than a single-simulator approach that only evaluates a test input in a single simulator.

Algorithm 1 details the steps of the **MultiSim** approach. The first step is to initialize the system under test T , the set of all failure-inducing test inputs C as well as the archive A (Line 1). Then, a random set of test inputs is sampled within the defined search domain (Line 2, P_0 being the initial population). In the next step, test inputs are evaluated based on the **MultiSim** specific evaluation function (see Algorithm 2). The main loop of the algorithm begins by applying default genetic operators such as selection, crossover and mutation (Lines 5–6). After the operators have been applied, a set of candidate test inputs to form the next population is defined. In particular, when all newly

Algorithm 1: Test case generation in **MultiSim**.

```

Input :  $(T, D, F, O)$ : An SBST Problem.
         $S = \{S_1, \dots, S_j\}$ : A set of simulators, with  $j > 1$ .
Output:  $C$ : Set of all failure-inducing test inputs.

1 INIT( $T$ );  $C \leftarrow \emptyset$ ;  $A \leftarrow \emptyset$ 
2  $P_0 \leftarrow \text{RANDOMSAMPLING}(T, D)$ 
3  $P_0 \leftarrow \text{EVALUATEMSIM}(P_0, S, F, T, A)$  /* Algorithm 2. */
4 while  $\text{budgetAvailable}$  do
5    $Par \leftarrow \text{TOURNAMENTSELECTION}(P_i)$ 
6    $P_i \leftarrow \text{CrossoverMutate}(Par)$ 
7    $P_i \leftarrow \text{EVALUATEMSIM}(P_i, S, F, T, A)$  /* Algorithm 2. */
8   for  $\mathbf{x} \in P_i$  do
9      $T \leftarrow T \cup \{\mathbf{x}\}$  /* Store all evaluated tests. */
10   $P_i \leftarrow \text{SURVIVAL}(P_i, O)$ 
11   $P_i \leftarrow \text{REPOPULATE}(P_i, D, T)$ 
12 for  $\mathbf{x} \in T$  do
13    $allFail \leftarrow \bigwedge_{k=1}^j O(f_k(\mathbf{x}))$ 
14   if  $allFail$  then
15      $C \leftarrow C \cup \{\mathbf{x}\}$ 
16 return  $C$ 

```

generated test inputs have been evaluated in the **MultiSim** specific evaluation function (Line 3), the survival operator is applied to rank the test inputs based on dominance (Line 10). Then, the re-population operator is used to replace a portion of dominated tests of each population by randomly generated ones to diversify the search (Line 11). When the search budget is exhausted, each test input is marked as failing, when all simulators agree on the failure of such test (Line 12–15). Practically, the oracle function O returns 1 for each simulator-specific fitness function $f_k(\mathbf{x})$, where $k : (i, \dots, j) < m$ indicates the fitness function for each specific simulator (Line 13).

4.1.1 Multi-Simulator Evaluation

As detailed in Algorithm 2, in **MultiSim** the evaluation of a test input is defined as the fitness vector:

$$F(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_j(\mathbf{x}), f_d(\mathbf{x}), f_a(\mathbf{x}, A)) \quad (1)$$

where (1) $f_1(\mathbf{x}) = v_1, \dots, f_j(\mathbf{x}) = v_j$ represent the fitness functions for the execution of \mathbf{x} in simulator S_1, \dots, S_j (Lines 3–4), (2) $f_d(\mathbf{x}) = v_d$ is defined as the average distance between the fitness values $f_1(\mathbf{x}), \dots, f_j(\mathbf{x})$ (Line 5), and (3) $f_a(\mathbf{x}, A) = v_a$ is the distance of the test input \mathbf{x} to previously found test inputs in the archive A (Line 6).

While optimizing the functions $f_1(\mathbf{x}), \dots, f_j(\mathbf{x})$ prioritizes failure-revealing test inputs, minimizing $f_d(\mathbf{x})$ drives the search towards test inputs on which all simulators agree. The fitness function $f_a(\mathbf{x}, A)$ aims to diversify the search

Algorithm 2: Test input evaluation in MultiSim.

Input :	T : System under test.
	$S = \{S_1, \dots, S_j\}$: A set of simulators, $j > 1$.
	δ : distance threshold
	P : Not evaluated test inputs.
	A : Archive of test inputs.
Output:	Q : Evaluated test inputs.

```

1  $Q \leftarrow \emptyset$ 
2 for  $\mathbf{x} \in P_i$  do
3   for  $S_k \in S$  do
4      $v_k \leftarrow \text{SIMULATE}(\mathbf{x}, T, S_k)$       /* Evaluate fitness for each simulator. */
5      $v_d \leftarrow \frac{\sum_{k=1}^{j-1} \sum_{l=k+1}^j |v_k - v_l|}{\binom{j}{2}}$   /* Compute average distance between fitness values. */
6     */
7      $v_a \leftarrow \text{GETDISTANCETOARCHIVE}(\mathbf{x}, A)$  /* Compute distance w.r.t. archived test
inputs. */
8      $\mathbf{x}.F \leftarrow (v_1, \dots, v_j, v_d, v_a)$       /* Assign fitness values to test input. */
9      $\text{ADDTOARCHIVE}(\mathbf{x}, A, \delta)$ 
10     $Q \leftarrow Q \cup \{\mathbf{x}\}$ 
11 return  $Q$ 

```

by calculating the distance of a test input to previously generated tests in the archive A [76]. As distance metric we use the Euclidean distance between normalized test inputs. In particular, we apply Min-Max normalization, scaling values in the test input representation to a range from 0 to 1. By maximizing $f_a(\mathbf{x}, A)$ the goal is to identify more diverse test cases. After the evaluation of a test input is completed, it is added to the archive in case its distance to the closest test in the archive exceeds a predefined threshold δ (the function `ADDTOARCHIVE` in Line 8, takes care of comparing the current test input \mathbf{x} with those in the archive A).

4.1.2 Test Representation

We adopt a model-based representation of test inputs. In line with previous studies [4, 76], each road scenario is defined through a sequence of control points. These points are interpolated using Catmull–Rom splines to generate smooth and continuous trajectories. To be able to generate diverse and challenging roads we parameterize the road definition using the vector $\mathbf{x} = (\alpha_1, \dots, \alpha_n, l_1, \dots, l_n)$, where $\alpha_i \in [-180; 180]$ denotes the clockwise angle between the i -th segment and the horizontal axis, while $l_i \in \mathbb{R}^+$ represents the length of the i -th segment between the control points c_{i+1} and c_{i+2} . For instance, the road in Figure 2a has 7 control points, 5 segments and is represented by the vector $(-90, -110, 173, -120, -140, 10, 20, 18, 17, 15)$. Note that this choice of input representation is not exclusive of our approach and equivalent parameterized representations can be defined, following similar principles [25].

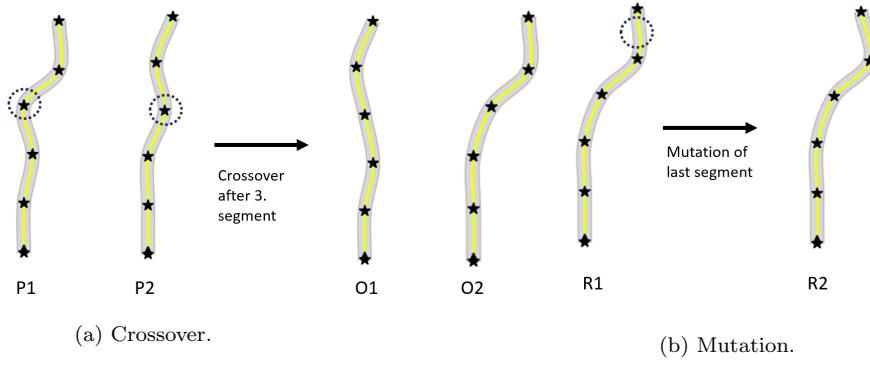


Fig. 2: Illustration of crossover (a) and mutation (b) of roads in MultiSim including control points which are placed within the road (visualized as stars). For sake of simplicity only angles are modified. In Figure a) tails of roads after the third segment are exchanged. In Figure b) the angle of the last segment is increased by 10 degrees.

4.1.3 Mutation/Crossover

Our implementation of the crossover operator on roads follows that of existing studies [40, 66]. Specifically, we use a one-point crossover operator, where for two given roads $\mathbf{p}_1, \mathbf{p}_2$ (i.e., the parents), we first randomly select an index, followed by exchanging the tail of the road \mathbf{p}_2 , including segments with a lower index, with the tail of road \mathbf{p}_1 , including segments with a higher index. Figure 2a illustrates the crossover of the road \mathbf{p}_1 with the road \mathbf{p}_2 resulting in the roads \mathbf{o}_1 and \mathbf{o}_2 (i.e., the offsprings), by exchanging the tails of the roads after the third segment.

The mutation of a road \mathbf{x} is performed by randomly selecting a segment, followed by increasing or decreasing the angle, or the segment length, by a given amount. A mutated road can be invalid because of intersecting segments being placed out of the given map, or when violating a threshold which defines the maximum angle between segments. If a road is invalid, we randomly generate a new road by applying the same operator as used to generate the initial population for MultiSim. As an example, Figure 2b shows the mutation of a road \mathbf{r}_1 by increasing the angle of the last segment from 90° to 100° .

4.1.4 Fitness Function

For the evaluation of a single test input within one simulator, we define a fitness function that captures the system-level behavior of the SUT. In the case of lane-keeping ADAS, the chosen fitness function is the cross-track error, denoted as f_{XTE} . This function computes the maximum lateral deviation between the vehicle's center and the lane center over the entire duration of the test execution, providing a direct quantitative measure of lane-keeping

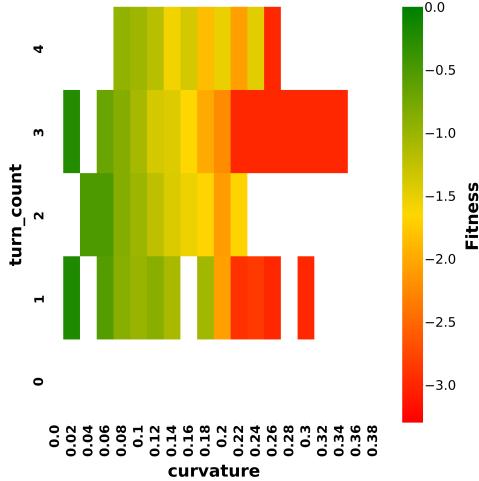


Fig. 3: Example of a feature map with two features, namely number of turns (i.e., `turn_count`) and curvature. The color of a cell is defined based on the worst XTE value of test inputs stored in a cell (colorbar on the right-hand side of the map). A cell is green when the XTE value is 0, red for the maximum, in absolute value, XTE value of -3. A cell is white if there is no test input that *covers* it. The fitness is negative, because the fitness function is to be minimized. In total 412 tests are stored in the feature map, which has 13 failing-cells and 35 non-failing cells.

performance. We selected this function because it has been shown to be an effective fitness function for triggering failure-inducing test inputs [16, 40, 76]. As oracle function we use $O = f_{XTE} > d$ which evaluates a scenario execution as a failure when the maximal XTE value exceeds d .

4.2 Test Validation

This step aims to validate that the identified failing tests are not failing due to simulator-specific behaviors. To characterize the ADAS unique failures, we adopt the feature maps by Zohdinasab et al. [68, 102].

Feature Maps. In the first step of our validation approach, we map test inputs generated by `MultiSim` into a feature space [68, 102], as it allows to represent important characteristics of test inputs in a human-interpretable way. In particular, with a feature space we can discretize the multi-dimensional search space, which is the space of roads (see Section 4.1.2), for our lane-keeping ADAS, and consider diverse test cases for the evaluation of our approach. Our feature space is characterized by two static road attributes, i.e., *num of turns* and *curvature*, i.e., the reciprocal of the minimal value of the radius over all

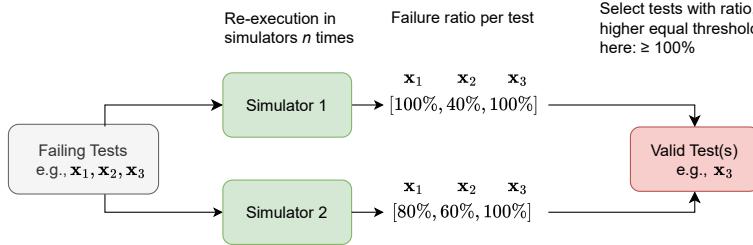


Fig. 4: Validation of failure-inducing test inputs, i.e., x_1, x_2, x_3 . Each test is re-executed in multiple simulators, i.e., two in this case S_1 and S_2 , and the failure rates for each test are computed. Then the failure rates on S_1 and S_2 are compared, and the test inputs are filtered according to a failure rate threshold. (in this case 100%). In this example, the only simulator-agnostic/valid failure-inducing test input is x_3 .

circles that can be placed through three consecutive control points of the road. Similarly to the fitness function, such choices of features were also found to be effective at characterizing the search space of road generators [102].

An example of a feature map is shown in Figure 3. A cell with the coordinates (x, y) contains test inputs whose curvature is in the range $[x - 0.02, x + 0.02]$, where 0.02 is the granularity of the map, and the number of turns is y . For instance, the road in Figure 1a is to be placed in the cell $(0.24, 3)$ as it has a maximum of curvature 0.23 and three turns. A cell of a feature map is colored based on the road with the best fitness placed in the cells, which is ranging from green (worst, $f_{XTE} = 0$) to red (best, $f_{XTE} = -3$) fitness value. We deem a cell as *failing* cell, when at least one test in the cell is failing w.r.t. the given oracle O . For instance, the cell $(0.32, 3)$ is a failing cell, while the cell at the position $(0.04, 2)$ is non-failing cell.

Re-execution. In the second step, we select failing cells from the feature map, extract a portion of failing and duplicate-free test inputs per cell, and re-execute each of these selected failure-inducing test inputs multiple times, to account for the flakiness of each specific simulation environment. Specifically, we execute such test inputs on simulators that have been not used by MultiSim during its search process, to confirm whether the failure-inducing test inputs generated by MultiSim are simulator-agnostic or not. For the validation we consider the notion of *hard* flakiness as defined by Amini et al. [5], where a test input is considered hard flaky when re-executing the test yields different pass or fail verdicts. Figure 4 illustrates the validation approach, when employing two simulators for the validation. First, each failure-inducing test input is executed on each simulator for n times. For each execution the test input is evaluated based on the fitness and oracle functions (the same used during the test case generation process of MultiSim). Afterwards, a failure rate is computed. Finally, the failure rates are compared between the simulators

for each test input to decide if the test is failing according to a user-defined threshold. For instance, in Figure 4 we have selected a 100% threshold, i.e., a failure-inducing test input is considered as simulator-agnostic if for every re-execution and for every simulator the test fails. In the example in Figure 4 only the test \mathbf{x}_3 is a simulator-agnostic/valid failure-inducing test input, as both simulators, i.e., simulator S_1 and S_2 , report a 100% failure rate. For \mathbf{x}_1 a re-execution in S_1 yields a failure rate below 100%. Similarly, for \mathbf{x}_2 re-executions in both simulators S_1 and S_2 , achieve failure rates below the threshold.

5 Empirical Study

In this section, we present the evaluation of our approach for identifying simulator-agnostic failures for a lane-keeping ADAS case study. We describe the research questions, our evaluation technique, evaluation experiments, implementation and results.

5.1 Research Questions

We consider the following research questions:

RQ₁ (effectiveness). *How effective is MultiSim in generating simulator-agnostic failure-inducing test inputs compared to single-simulator testing and a state-of-the-art approach?*

The first question evaluates the effectiveness of MultiSim in finding simulator-agnostic failures. After the test case generation phase, we validate a portion of failure-inducing test scenarios based on the approach described in Section 4.2.

RQ₂ (efficiency). *How efficient is MultiSim in generating simulator-agnostic failures compared to single-simulator testing and a state-of-the-art approach?*

The second research question evaluates whether MultiSim is efficient in identifying failures. Test efficiency is particularly relevant when the testing budget is limited [13].

RQ₃ (prediction). *To what extent can machine learning predictors increase the effectiveness and the efficiency of MultiSim in finding simulator-agnostic failures?*

Since simulator-specific failures, i.e., disagreements between simulators, might impact the identification of simulator-agnostic failures during the search, we use machine learning (ML) classifiers to predict whether a test input is leading to a disagreement between pass/fail verdicts on different simulators. We investigate whether this approach can be considered as a calibration technique to fine-tune the original MultiSim approach to mitigate disagreements during the search caused by simulator-specific failures.

5.2 Baselines

We compare the performance of **MultiSim** to a search-based testing approach which only employs one simulator (**SingleSim**) during the search. We also compare with the recently proposed **DSS** approach [16], which employs multiple simulators by re-executing test inputs generated for one simulator in a different simulator and vice-versa. **DSS** is, to the best of our knowledge, the only existing approach in the literature that employs multiple simulators for automated test case generation in the ADAS domain.

5.3 Procedure and Metrics

To answer all three research questions for our experiment, we first validate the failures as described in Section 4.2. Specifically, we assign all tests to cells of the feature map, by extracting curvature and number of turns of the corresponding roads, using the same feature space granularity as proposed by Biagiola et al. [16]. We randomly select three failure-inducing test inputs from each failing cell. Preliminary experiments with a higher number of selected tests per failing cell, i.e., five, did not affect the results significantly.

To answer RQ₁, to validate failures found by **MultiSim**, we use the third simulator that was not used during the search. We then re-execute each test five times to account for the flakiness of the simulation environment. As failure rate threshold we use 100%, i.e., a failure-inducing test input found by **MultiSim** is only valid if it always fails in the third validation simulator.

To validate failures found by **SingleSim**, we use the same validation approach, but we use the two simulators that are not used during the search. For **DSS** we also use only one simulator for validation, as two of the three simulators available for our study are already used in the search. Based on the number of selected failures and the validation results, we assess the performance of **MultiSim**, **SingleSim** and **DSS** using two metrics: (1) *n_valid*, which is the number of valid failures, and (2) *valid_rate*, which is defined as:

$$\text{valid_rate} = \frac{\# \text{ valid failures}}{\# \text{ failures selected for validation}}$$

To answer RQ₂, after validating all the selected failures, we identify at which time during the search each valid failure is encountered. Thus, we define efficiency as the proportion of the search budget required to identify the first valid failure. It measures how early in the search process the approach discovers a fault-revealing test relative to the total number of generations or evaluations. This metric reflects the practical usefulness of the search procedure in time- and resource-constrained testing contexts, where faster identification of valid failures is beneficial.

For **DSS**, as it consists of two independent search executions and two migration steps, we first identify in which simulation run the failure is detected.

Based on the order of the run executions and the evaluation number we compute the total run time to find the failure. That is, when a failure is detected in simulator DS1 with in total A test evaluations at the evaluation number a , while simulator DS2 performed in total B test evaluations, we assign to the test the final evaluation count as $A + B + a$. In case the valid failure is detected during the execution of B at evaluation number b , we assign to the test the evaluation count $A + B + A + b$, respectively. We evaluate efficiency using this methodology because DSS consists of two independent testing experiments, and unlike single-simulator-based testing, it is not possible to incrementally order the test cases generated by DSS for a single run.

To answer RQ₃, we select the configuration of **MultiSim** with the *best* results in terms of validity rate and number of failures, and collect test cases from past executions where the simulators both agree or disagree. In addition, we execute a modified version of **MultiSim** with five different seeds where the fitness function prioritizes disagreements. Specifically, we reuse Algorithm 2, but maximize $f_d(\mathbf{x})$, i.e., the distance between maximum XTE evaluations. We pass the roads as inputs and the labels representing an agreement or disagreement to five different and widely used ML classification models for training. Specifically, we use Decision Trees [22], Random Forests [21], Support Vector Machines [28], Logistic Regression [23] and Stochastic Gradient Boosting [31]. For the training of the classification models we use default hyperparameter configurations available in the Sklearn library.

We then evaluate the ML models based on F-1-score and AUC-score. The best performing ML model on all metrics is then integrated into **MultiSim** and executed with the same search configuration as **MultiSim**.

In terms of metrics, we measure the number of valid failures (n_valid) and validity rate ($valid_rate$) as in RQ₁, and compare **MultiSim** with and without the integration of a disagreement predictor. We also measure the first valid failure as the efficiency metric for both configurations.

5.4 Experimental Setup

5.4.1 Simulation Environments

We have selected the following simulators for our study: BeamNG, Udacity and Donkey Car. We have chosen such simulators for the following reasons: 1) they seamlessly integrate with our lane-keeping ADAS case study, as the simulation environment, i.e., asphalt road on green grass and sunny weather, has been aligned [16] to be able to replicate the same simulation conditions across different simulators (see Figure 1); 2) they are complementary as they differ in terms of the physics implementation or the rendering engine; 3) they are open-source or available under permissive academic-licenses (BeamNG). A brief description of the simulators provided is as follows:

- BeamNG is based on Unreal engine, implements soft-body physics, serves as the simulator in the SBFT competition [15, 32, 58], and it has been used

extensively in the literature to benchmark testing approaches [12, 40, 76, 102].

- The Udacity simulator is a cross-platform and rigid-body physics developed with Unity3D [94]. While originally developed for educational purposes, it has been widely applied for ADAS testing [11, 16, 35, 41, 42, 82, 85].
- The Donkey Car simulator (Donkey, hereafter) is part of the Donkey Car framework that is used to build and test small-scale self-driving physical vehicles [29]. Similarly to Udacity, it implements rigid-body physics, and is developed with Unity 3D. The simulator is actively used in the literature for the validation of testing approaches [16, 52, 53, 67, 83, 84].

Unlike previous work [16], we make no assumptions about simulator fidelity and treat all simulators as black-box environments. In our experiments, we execute **MultiSim** in three different configurations: in the configuration BD we use BeamNG and Donkey as simulator ensemble, in the configuration BU we use BeamNG and Udacity, and lastly in the configuration UD we use Udacity and Donkey. We also evaluate DSS in such configurations; in the tables and figures, we prefix such configurations with DSS- (e.g., DSS-BU indicates the configuration BU with the DSS approach, while BU indicates the same configuration with the **MultiSim** approach). It is noteworthy that in the original DSS paper [16] only the configuration DSS-BU was proposed and evaluated. The extension of the approach to other simulators constitutes a novel experimental contribution of our work.

5.4.2 Study Subjects

To assess the generality and robustness of our approach, we evaluate it on three lane-keeping ADAS that differ in their design paradigms, from convolutional models, to transformer-based architectures leveraging global attention, and hybrid vision–trajectory policies that integrate planning information, described next:

- **CNN-based ADAS (DAVE-2):** DAVE-2 is a convolutional neural network developed for multi-output regression tasks based on imitation learning [19]. The model architecture includes three convolutional layers for feature extraction, followed by five fully connected layers. The model takes as input an image representing a road scene, and it is trained to predict vehicle’s actuators commands. We use a pre-trained DAVE-2 model from previous work [16] that has been trained on more than 60,000 images and steering commands automatically by a PID-controlled autopilot driving the car on roads with different road complexity in the three simulators used in this work. DAVE-2 has been extensively used in a variety of ADAS testing studies [16, 17, 42, 50, 51, 83, 89, 99].
- **Vision Transformer-based ADAS (ViT):** ViT [48] is a DNN based on the transformer architecture to predict steering commands. While Transformers were originally developed for natural language generation and understanding, they have since been widely adopted for visual tasks [6], in-

cluding lane-keeping, where ViT achieves results competitive with systems such as DAVE-2 [11]. ViT receives an image of the road scene and predicts the steering command. Similarly to DAVE-2, we trained the ViT model using images and steering commands collected by the PID controller driving the car on 100 tracks for each simulator. Differently from DAVE-2, we trained ViT using approximately twice as many images (i.e., 120,000).

- **Trajectory-Control Policy ADAS (TCP):** TCP [97] is a hybrid deep neural network that combines vision-based perception with trajectory-conditioned control to predict steering and throttle commands. The model adopts a two-branch architecture: a CNN-based vision encoder that extracts high-level spatial features from the input image, and a trajectory encoder (an MLP) that processes the planned waypoints provided by the controller. The two representations are fused in a joint latent space and passed through fully connected layers to produce continuous control outputs. Training follows a procedure similar to DAVE-2 and ViT. We collected synchronized driving data from 40 tracks per environment by running a PID controller to generate reference control signals. Each sample includes the RGB scene image, the corresponding trajectory points, and the ground-truth steering and throttle commands. The final dataset comprises over 62,000 samples, enabling TCP to learn a context-aware mapping between the visual scene, planned path, and control actions.

5.4.3 Search Setup

Search Algorithm. For the test case generation with **MultiSim** we use Algorithm 1. For **SingleSim**, we employ only one simulator for the evaluation of the fitness value of a test case. We use the same hyper-parameter configuration for all algorithms, such as mutation, crossover rate and population size. To select the archive threshold δ that controls which test inputs shall be included in the archive during search, we follow the guidelines of Riccio et al. [76] and perform preliminary experiments with different thresholds, ranging from 4 to 6. We assess the diversity of the corresponding feature map after the search, by identifying how many cells of the map are covered. As the archive threshold affects the diversity of the search, we use the coverage of the feature map to identify an appropriate threshold, which is 4.5 in our case. Further, we configure a maximum search budget of six hours for all testing approaches with DAVE-2 and ViT and 3.5 hours with TCP, which proved sufficient to explore the search space and obtain an adequate number of failures.

In the **DSS** approach, we also use Algorithm 1 for test generation. Test cases are produced through two independent search executions, each employing a different simulator, i.e., S_1 and S_2 . The test cases generated by S_1 are subsequently re-executed in S_2 , and vice-versa. A test case is classified as critical only if it has been evaluated as critical by both simulators. The results from both simulators, including all generated and migrated test cases, are then aggregated to represent the final output.

Table 1: Test generation configuration for the approaches under test **MultiSim**, **SingleSim** and **DSS**. All configurations are equal across methods, beside the number of fitness functions used for **MultiSim**.

Parameter	SingleSim	MultiSim	DSS
Population size	20	20	20
Time budget	6 hours	6 hours	6 hours
Mutation rate	1/10	1/10	1/10
Crossover rate	0.6	0.6	0.6
Archive distance	4.5	4.5	4.5
# Fitness functions	2	4	2
# Search variables	10	10	10

For **DSS**, we adapted the search budget for the search in two selected simulators before the test migration: we set it to a quarter of the total search budget, because first, two independent search runs are executed, and in the second step results have to be migrated in the respective simulator. The remaining parameters are set as for **SingleSim** and **MultiSim**. An overview of the complete configuration for all search methods can be found in Table 1.

Initial Population. To generate the initial population of roads, we use the test generator implemented in the study from Riccio et al. [76], where a road is iteratively generated while persevering the user-defined constraints such as segment length and maximum angle between consecutive segments. In particular the algorithm samples segments randomly based on the given maximum angle and segment length.

Mutation/Crossover. The mutation extent is randomly selected in the range of $[-8, 8]$, while the segment length is randomly set in the range $[10, 20]$.

Fitness/Oracle Function. As fitness and oracle functions we use the function f_{XTE} as defined in the motivation example. As oracle function we use $O = f_{XTE} > 2.2$ that evaluates a scenario execution as a failure when the vehicle leaves the driving lane (similar to existing case studies [16, 39]). We stop the simulation on a given road when the current XTE value exceeds a threshold of 3 mt, or the car reaches the end of the last segment.

5.5 ADAS Preliminary Evaluation

For DAVE-2, we considered an ADAS model from previous work [16], to mitigate the threats to validity of training our own model for testing. Vision perception models like DAVE-2, are negatively affected by changes of image distributions [50–52]. Thus, we assess the driving performance of the ADAS from previous work in our environment prior to the empirical evaluation. Specifically, we have defined 16 different types of roads with a road-wise maximum curvature ranging from 0 to 0.19, while the number of turns was in the range 1 to 4, and the road length was 100 mt. The evaluation allows us, on the one hand to investigate the flakiness of the simulation environments, and on the

other hand to verify whether the testing setup including the hardware and software configurations is suitable for our study.

For ViT, we selected an existing model architecture [11] and trained the model on collected images with corresponding PID-based steering commands by letting the model drive on 100 diverse automatically generated roads.

For TCP, we reused the model architecture from the original work [97], which includes a ResNet model for visual feature extraction trained in the CARLA simulator. Because the pre-trained ADAS exhibited consistent failures across our simulators, we retrained it on a unified dataset collected from all three environments through PID-controlled driving over 40 randomly generated tracks, ensuring balanced exposure to the visual and physical characteristics of each simulator. The collected data was composed of images, steering commands, vehicle position information, and target waypoint positions. The TCP model defines a parameter *alpha* that specifies the weight of the *trajectory branch* vs. the *predict control branch* for combined steering value prediction. We selected the alpha score, i.e., 0.1, based on preliminary experiments using validation roads. We executed all ADAS in every simulator on each of the predefined roads for 10 times to account for the flakiness, and assessed the quality of the driving using the XTE fitness function.

The results show that all ADAS exhibit low XTE variation (below 0.7m) across all roads without failing, even on roads with a high number of turns and high curvature. These observations suggest that while the vehicle’s trajectories and driving styles are not harmonized across simulators, the ADAS’s driving performance is not impacted by hardware limitations or latencies, as the variation is minimal and the ADAS does not fail. Thus, we conclude that the evaluated ADAS under test are sufficiently robust for exhaustive testing.

The results of all preliminary experiments with a detailed description of the roads used can be found in our replication package [91].

5.6 Implementation

We have implemented **MultiSim** using the testing framework OpenSBT [80]. OpenSBT is a modular framework that eases the execution, as well as the extension of SBT components for conducting testing experiments. OpenSBT has been recently applied for evaluating testing techniques and replication studies [70, 79, 81]. To implement our test generator in OpenSBT, we have extended the problem definition and test execution interfaces by evaluating a test input in multiple simulators and merging the result using a migration function. Our implementation allows to use any simulator and evaluation migration strategy, which is integrated into the framework. The code of **MultiSim** is available in our replication package [91].

All experiments have been executed on two computing devices, one PC with a Ryzen 5 5600X CPU, 32 GB RAM, and a 4090 Ti GPU, and on a Linux Laptop with an i7 CPU, 32 GB RAM, and a 790 GX GPU. To accelerate DNN computations we used CUDA 11.8 with cuDNN 8.9. Note that due to

simulator-specific requirements a considerable effort was required to identify a hardware configuration which was able to support our case study. BeamNG exhibited initially a very low FPS number when running on a device with the GPU 4090 Ti. In addition, BeamNG does not support the Linux, but runs on Windows. Udacity, in contrast, could not run on our Windows machine, so that we had to employ a second computing device, i.e., a Linux laptop, to perform the experiments for the BU/DSS-BU configuration. Finally, we made sure that all simulators run at 20 FPS.

Overall, our experiments required approximately 1,395 hours i.e., ≈ 58 days (3 systems \times 6/10 repetitions \times 9 methods \times 3.5/6 hours) for all compared test generation approaches. Additional more than 25 hours were required to validate more than 6,000 failing tests.

5.7 Results

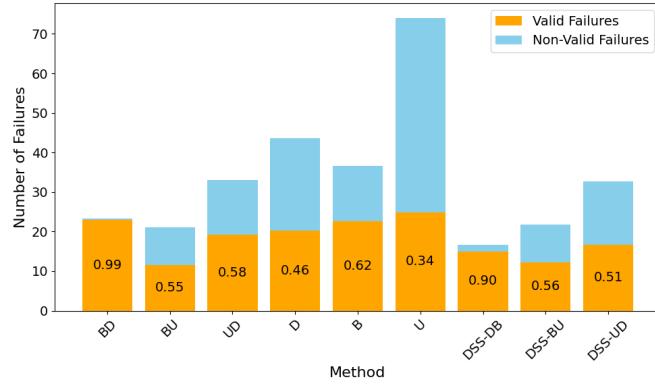
5.7.1 Effectiveness (RQ_1)

The results for RQ_1 are shown in Figure 5. The figure depicts the average number of failures and the number of *valid* failures, identified from 10 runs for **MultiSim**, **SingleSim** and **DSS** for the system DAVE-2, and from 6 runs for the systems ViT and TCP. In addition, in each bar plot the validity rate is annotated in the orange bar.

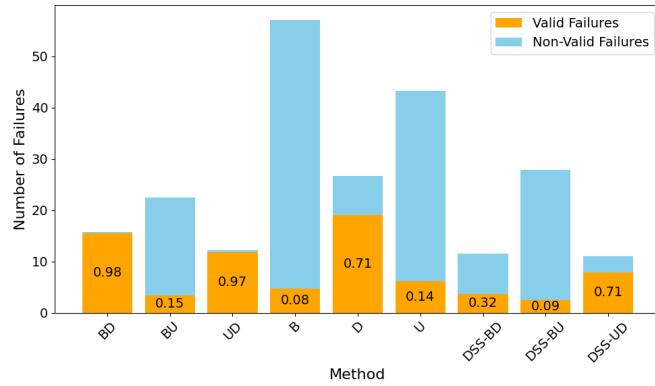
Across all configurations and systems under test, single-simulator (**SingleSim**) approaches generally detect more total failures than multi-simulator ones, but many of these are invalid. The highest number of valid failures is achieved by **SingleSim** on Udacity (U), which is comparable to the best multi-simulator configuration, BD for DAVE-2. While **DSS** yields in general fewer total failures overall, the number of valid failures remains for DAVE-2 similar to **MultiSim**. However, for ViT and TCP the number of valid failures is lower for **DSS**.

Simulator-wise, for DAVE-2 and ViT, BD produces fewer total failures than its single-simulator counterparts (B and D) but maintains a comparable number of valid ones. In contrast, BU identifies a similar total number of failures as B and U but significantly fewer valid cases, a pattern also observed for DSS-BU and UD. For TCP, the configurations BU and UD reach a significantly higher number of valid failures compared the single simulator-based counterparts, as well as to **DSS** and the multi-simulator configuration BD.

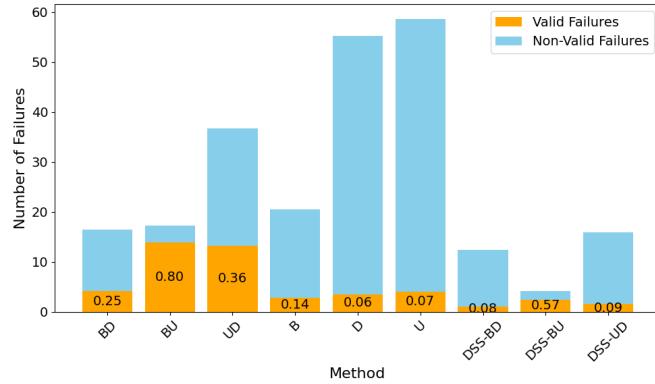
The results in Figure 5 show that for DAVE-2 and ViT, BD achieves the highest validity rate across all compared approaches and configurations, (99% for DAVE-2, 98% for ViT). For DAVE-2, BU has a validity rate of 55%, which is surpassed by B, being close to that of DSS-BU. For UD the validity rate is higher compared to those of D, U, and DSS-UD. For TCP, BU has the highest validity rate (80%), which is higher than the corresponding **DSS** based counterpart. Also, we can observe that all **SingleSim** approaches achieve lower validity rates than the worst **MultiSim** combination BD (0.25).



(a) DAVE-2



(b) ViT



(c) TCP

Fig. 5: Validity rate (*valid_rate*) and number of valid failures (*n_valid*) identified for MultiSim, DSS, and SingleSim averaged for DAVE-2, ViT, and TCP. The average validity rate is shown in each bar plot.

Table 2: Effectiveness: Statistical tests (Wilcoxon and Vargha-Delaney) for the comparison between **MultiSim** and **SingleSim**, as well as between **MultiSim** and **DSS** for the metrics validity rate and number of valid failures. The letter L represents a large effect size magnitude, M represents a medium effect, and S represents a small effect (repeated comparisons are excluded (N/A)). Effect values are annotated with a star if the first approach of the comparison yields lower values than the compared approach.

	U p (effect)	B p (effect)	D p (effect)	BU p (effect)	UD p (effect)	DSS-BD p (effect)	DSS-BU p (effect)	DSS-UD p (effect)
DAVE-2								
BD								
valid_rate	~ 0 (L)	~ 0 (L)	~ 0 (L)	0.01 (L)	~ 0 (L)	0.03 (L)	0.03 (L)	0.14 (-)
n_valid	1.00 (-)	0.85 (-)	0.83 (-)	0.03 (L)	0.77 (-)	0.03 (L)	0.03 (L)	0.06 (S)
BU								
valid_rate	0.11 (-)	0.77 (-)	0.43 (-)	0.77 (-)	0.02 (L)	N/A	0.56 (L*)	0.77 (-) (-)
n_valid	0.04 (L)	0.01 (L)	N/A	0.16 (-)	0.11 (-)	1.00 (M)	0.31 (M)	0.08 (-)
UD								
valid_rate	0.03 (L)	0.92 (-)	0.32 (-)	N/A	N/A	0.16 (M)	0.03 (L)	0.14 (-)
n_valid	0.85 (-)	0.77 (-)	0.68 (-)	N/A	N/A	0.16 (L)	0.09 (M)	0.44 (S)
ViT								
BD								
valid_rate	0.03 (L)	0.03 (L)	0.03 (L)	0.03 (L)	0.59 (S)	0.04 (L)	0.03 (L)	0.14 (-)
n_valid	0.09 (M)	0.03 (L)	0.56 (S)	0.06 (L)	0.31 (M)	0.03 (L)	0.03 (L)	0.06 (S)
UD								
valid_rate	0.03 (L)	0.03 (L)	0.03 (L)	0.03 (L)	N/A	0.16 (M)	0.03 (L)	0.14 (-)
n_valid	0.22 (M)	0.08 (M)	0.31 (M)	0.16 (M)	N/A	0.16 (L)	0.09 (M)	0.44 (S)
BU								
valid_rate	0.84 (-)	0.03 (S)	0.03 (L*)	N/A	N/A	1.00 (M)	0.31 (M)	0.06 (-)
n_valid	0.44 (S)	0.31 (M)	0.04 (L*)	N/A	N/A	1.00 (M)	0.50 (-)	0.08 (-)
TCP								
BD								
valid_rate	0.04 (L)	0.69 (M)	0.04 (L)	0.56 (S*)	0.06 (L*)	0.04 (L)	0.14 (L)	0.08 (L*)
n_valid	0.78 (L)	0.59 (S)	1.00 (-)	0.09 (L*)	0.06 (L*)	0.07 (L*)	0.04 (L)	0.36 (M*)
UD								
valid_rate	0.06 (L*)	0.16 (L*)	0.06 (L*)	0.09 (L*)	N/A	0.09 (L)	0.04 (L)	0.56 (S)
n_valid	0.09 (L*)	0.06 (L*)	0.22 (L*)	1.00 (-)	N/A	0.06 (L)	0.04 (L)	0.06 (L*)
BU								
valid_rate	0.04 (L)	0.06 (L*)	0.04 (L)	N/A	N/A	0.04 (L)	0.04 (L)	0.29 (L*)
n_valid	0.04 (L)	0.06 (L*)	0.04 (L)	N/A	N/A	0.04 (L)	0.04 (L)	0.06 (L*)

We further evaluated the convergence behavior of **MultiSim** and **SingleSim** using the Hypervolume (HV) quality indicator, a widely adopted metric in search-based software engineering for assessing both the optimality and spread of Pareto fronts [55, 56]. The HV analysis (available in our replication package [91]) shows that **SingleSim** configurations quickly reach stable HV values, whereas **MultiSim** achieves comparable values that continue to improve slightly over time. This suggests that extending the execution time of **SingleSim** would likely not yield further gains in effectiveness, while **MultiSim** maintains gradual improvements as the search progresses.

Statistical Test. To assess whether the differences observed are statistical significant, we use the non-parametric pairwise Wilcoxon rank sum test (significance level 0.05) and the Vargha Delaney's \hat{A}_{12} effect size to compare the *valid_rate* and the *n_valid* values shown in Figure 5. We adopt the following standard classification for effect size values: an effect size e is small, when $0.56 \leq e < 0.64$ or $0.36 < e \leq 0.44$, it is medium when $0.64 \leq e < 0.71$ or $0.29 < e \leq 0.36$ and large when $e \geq 0.71$ or $e \leq 0.29$. Otherwise, the effect size is negligible. The results are shown in Table 2.

For DAVE-2, the statistical test results show that BD yields significantly higher validity rates with large effect sizes than **SingleSim**, **MultiSim** and DSS approaches DSS-BU and DSS-UD. As for BU, the validity rate results for **MultiSim** are in general not better or worse than **SingleSim** or **MultiSim**. For UD, the results are only significant compared to **SingleSim** with Udacity (i.e., U). Regarding the number of valid failures, we see that BD outperforms DSS-BD with a medium effect size, while having a similar validity rate. Regarding BU, **SingleSim** B and U are significantly better than BU with a large effect size, what is consistent with the results in Figure 5. Concerning UD, we do not observe a statistical significant difference w.r.t. **SingleSim** and DSS.

For ViT, similarly there is a statistically significant difference with large effect sizes for the valid rate between **MultiSim**-based and **SingleSim** based approaches as well as the corresponding DSS-BD baseline. Also for UD there is a statistically relevant difference compared to **SingleSim** approaches which is inline with the results in Figure 5. For the number of valid failures, we see statistically significant higher values for BD compared to the single simulator approach B as well to corresponding DSS counterparts.

For TCP, both BD and BU configurations exhibit significantly higher validity rates and numbers of valid failures than the single-simulator (U, D) and corresponding DSS configurations, with large effect sizes. UD also shows statistically significant improvements over most **SingleSim** configurations, though with smaller effects. Overall, results consistently confirm that combining complementary simulators enhances the discovery of valid failures across architectures, with BD emerging as the most effective configuration.

RQ₁ (effectiveness). Across all ADAS, **MultiSim** achieves higher validity rates and numbers of valid failures than both **SingleSim** and DSS (70% vs. 65% and 47%). Among multi-simulator settings, BD (BeamNG–Donkey) performs best for DAVE-2 and ViT (up to 99% validity), while BU (BeamNG–Udacity) excels for TCP with up to 80%. This indicates that the most effective simulator pairing depends on the SUT: BD benefits vision- and control-based systems, whereas BU better captures perception-driven failures. Statistical tests confirm that **MultiSim** significantly outperforms both **SingleSim** and DSS, except for DSS-BD, where it achieves comparable validity but a higher number of failures.

Table 3: Efficiency results, i.e., time to first valid failure for DAVE-2, ViT, and TCP. Mean and standard deviation (std) are in percent.

Method	DAVE-2		ViT		TCP	
	Mean	Std	Mean	Std	Mean	Std
BD	39.7	22.9	34.2	9.4	40.6	31.2
UD	35.3	24.0	56.5	22.2	31.0	33.3
BU	55.6	29.4	60.1	29.4	30.6	31.2
U	22.8	16.1	44.5	37.3	59.0	31.1
B	25.5	11.3	51.4	31.1	79.5	26.9
D	22.4	14.1	37.0	20.9	74.4	34.4
DSS-BD	56.6	16.3	37.1	2.2	61.1	12.5
DSS-UD	39.9	13.6	43.4	6.5	45.8	63.5
DSS-BU	33.5	13.2	34.6	5.2	51.6	10.2

5.7.2 Efficiency (RQ_2)

Table 3 shows the efficiency results, specifically the ratio of the search budget and its standard deviation when the first valid failure is detected.

For DAVE-2 we observe that all **MultiSim** configurations exhibit higher percentages for the first valid failure compared to **SingleSim**, as expected. Compared to **DSS**, we see that the configuration BD achieves a better result than **DSS-BD**. However, across all configurations, the standard deviations for **MultiSim** are higher than those of **SingleSim** or **DSS**. The combination UD is surpassed by **SingleSim** and **DSS-BU**, while **DSS-UD** and **DSS-BD** yield higher search times. For BU, we observe a similar behavior; here, **DSS-UD** additionally remains more efficient.

The statistical test results are shown in Table 4. While for DAVE-2 D significantly outperforms BD, BD significantly performs better than **DSS-BD**. Regarding BU, we can only observe that it is outperformed by **DSS-BU** and D, while there is no significant difference with the other approaches. The **SingleSim** configurations U and D yield significantly lower search times than UD, with large effect sizes. However, we do not observe a statistical significant difference of UD w.r.t. **DSS-UD**, which confirms our observations from Table 4.

For ViT, we can see that, similarly as for DAVE-2, **MultiSim** combinations achieve on average higher values than **SingleSim**-based search and **DSS**.

For TCP, **MultiSim** achieves lower scores with similar variation compared to **SingleSim**. However, the difference is not statistically significant except for the comparison between BU and B for TCP, where **MultiSim** achieves significantly better results.

Figure 6 shows the study results considering both effectiveness and efficiency (validity rate vs. search budget required for first valid failure) for DAVE-2, ViT and TCP for **MultiSim**, **SingleSim** and **DSS**. Particularly, for DAVE-2 the methods whose values are non-dominated either in effectiveness nor efficiency are BD, B, U and D (Pareto-optimal solutions). **SingleSim**, when using BeamNG and Donkey, is more efficient in finding the first valid

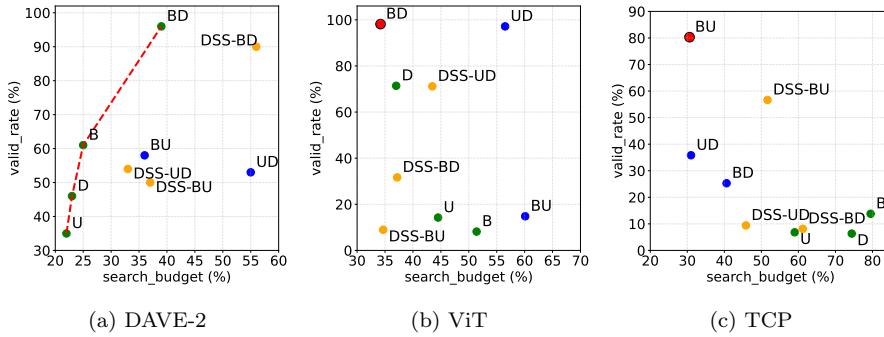


Fig. 6: Visualization of the tradeoff between efficiency (first valid failure) and effectiveness (*valid_rate*) for **MultiSim**, **DSS** and **SingleSim** for different configurations across three case studies. The Pareto front is visualized in red.

failure, while based on the results of the **MultiSim** combination BD, **MultiSim** is significantly more effective. For ViT/TCP, we see that the **MultiSim** combination BD/BU is dominating all remaining configurations yielding both a high validity rate with a low search effort for the first valid failure.

RQ₂ (efficiency). Across all configurations, **MultiSim** achieves efficiency comparable to **SingleSim**, requiring a similar search budget to find the first valid failure. Among multi-simulator settings, BD and UD are the most efficient for DAVE-2 and ViT, while BU performs best for TCP. Compared to DSS, in general **MultiSim** shows better efficiency, whereas DSS-BU remains in two out of three comparisons slightly more time-efficient than BU. Overall, **MultiSim** maintains competitive efficiency while yielding substantially higher failure validity.

5.7.3 Prediction (*RQ₃*)

To reduce *unnecessary* simulator evaluations, we extend **MultiSim** with a *surrogate disagreement predictor*. This surrogate is a classifier, trained using **labeled outcomes** of previous runs, where each instance is annotated as an *agreement* or *disagreement*. The surrogate outputs the probability of a *disagreement event* between simulator pairs. Specifically, we gathered all test cases from five completed runs using BD, along with five additional runs using a modified version of **MultiSim** focused on identifying disagreements (Section 5.3). In total, we collected a dataset comprising 580 agreements and 572 disagreements for DAVE-2, 581 agreements and 581 disagreements for ViT, and 383 agreements and 383 disagreements for TCP. These datasets were then used to train five classification models aimed at predicting disagreement occurrences. We train and validate multiple surrogate classifiers (Decision Tree, Random Forest, SVM, Logistic Regression, and Gradient Boosting) using 5-

Table 4: (DAVE-2) Efficiency. Statistical tests (Wilcoxon and Vargha-Delaney) for identifying the first valid failure (repeated comparisons are excluded (N/A)) for the ADAS DAVE-2, ViT and TCP. Effect size magnitudes are annotated with a star, if the first approach (rows) of the comparison yields lower values than the second approach (columns).

	U <i>p</i> (effect)	B <i>p</i> (effect)	D <i>p</i> (effect)	UD <i>p</i> (effect)	BU <i>p</i> (effect)	DSS-BD <i>p</i> (effect)	DSS-BU <i>p</i> (effect)	DSS-UD <i>p</i> (effect)
DAVE-2								
BD								
BD	0.11 (-)	0.06 (-)	0.03 (L*)	0.70 (-)	0.16 (-)	0.01 (L*)	0.38 (-)	1.00 (-)
BU	0.16 (-)	0.28 (-)	0.04 (M*)	0.19 (-)	N/A	0.85 (-)	0.03 (L)	0.16 (-)
UD	0.03 (L*)	0.02 (L*)	0.01 (L*)	N/A	N/A	0.01 (L*)	0.92 (-)	0.49 (-)
ViT								
BD								
BD	1.00 (L)	0.31 (S)	0.84 (S)	0.16 (L*)	0.16 (L)	0.06 (L)	0.84 (-)	0.09 (L)
BU	0.56 (M)	0.56 (S)	0.44 (L)	N/A	0.69 (-)	0.09 (L)	0.16 (L)	0.31 (M)
UD	0.56 (M)	0.84 (S)	0.31 (L)	N/A	N/A	0.16 (L)	0.16 (L)	0.44 (S)
TCP								
BD								
BD	0.22 (M*)	0.16 (L*)	0.09 (L*)	0.31 (L)	0.44 (L)	0.31 (L*)	0.56 (L*)	0.56 (L*)
BU	0.16 (L*)	0.04 (L*)	0.16 (L*)	N/A	N/A	0.09 (L*)	0.31 (L*)	0.44 (L*)
UD	0.22 (L*)	0.16 (L*)	0.16 (L*)	N/A	1.00 (-)	0.06 (L*)	0.44 (L*)	0.44 (L*)

fold cross-validation with an 80:20 split. As reported in Table 5, the Random Forest model achieved the best performance for every system under test, with F-1 scores close to or above 80%. We therefore integrate this model into `MultiSim` as the surrogate predictor.

During the test generation phase, each newly generated candidate scenario is first evaluated by the surrogate model. If the predicted probability of simulator disagreement exceeds a threshold ($\tau = 0.7$), the candidate is discarded before running the simulation. Otherwise, it is executed across both simulators. This filtering mechanism allows `MultiSim` to focus simulation budget on the exploration on high-value, disagreement-prone regions of the search space.

As Table 6 shows, across all three case studies, integrating the surrogate disagreement predictor into `MultiSim` (BD-P) improved the overall test generation effectiveness. The surrogate-guided runs produced a higher number of valid and fault-revealing test cases compared to the baseline configuration (BD), increasing the mean *n_valid* and *valid_rate* by approximately 38% on average. This confirms that the surrogate model not only reduces redundant

Table 5: Classifier results after 5-fold cross-validation (in %) for predicting disagreements in three simulator configurations. The best results are marked in bold.

	DAVE-2		ViT		TCP	
	F-1	AUC-ROC	F-1	AUC-ROC	F-1	AUC-ROC
DT	79	78	58	65	60	65
RF	83	89	74	81	78	86
SVM	69	71	65	77	58	80
LREG	57	62	58	65	60	65
GBOOST	80	87	73	80	75	82

Table 6: Mean and standard deviation of effectiveness and efficiency metrics for BD-P and BD approaches across three use cases, averaged over 10 runs for DAVE-2, 6 runs for ViT, and 6 runs for TCP.

	DAVE-2		ViT		TCP	
	BD-P	BD	BD-P	BD	BD-P	BD
n_valid	25.3 ± 7.3	23.0 ± 13.1	15.8 ± 2.0	15.5 ± 8.3	8.3 ± 3.6	4.1 ± 4.8
valid_rate (%)	92.9 ± 7.9	98.5 ± 3.0	97.9 ± 3.0	98.1 ± 4.1	44.3 ± 8.3	25.9 ± 18.7
first_fail (%)	34.0 ± 14.1	39.7 ± 22.9	$43. \pm 16.7$	34.2 ± 9.4	39.8 ± 25.9	40.6 ± 31.2

simulator queries but also helps focusing the search on higher-yield regions of the input space, maintaining an average validity rate above 90% across runs.

The use of the surrogate model also leads to efficiency improvements in two out of the three case studies (Table 6). For DAVE-2 it yields an improvement of 14% (first failure found at 34% vs. 39.7% of the search budget), and 2% (first failure found at 39.8% vs. 40.6% of the search budget) for TCP. These results indicate that surrogate-based filtering can accelerate simulation-based testing while maintaining the ability to identify a broad range of failure cases.

RQ₃ (prediction). Among all trained classifiers, the Random Forest achieved the best disagreement-prediction performance (AUC: 0.81–0.89, F-1: 0.74–0.83). Integrating this surrogate into **MultiSim** enabled the framework to bypass low-value simulations, improving both effectiveness and efficiency by increasing the number of valid failures and reducing variability compared to baseline approaches.

5.8 Qualitative Analysis

In the following, we describe a qualitative analysis concerning our evaluation results. In the first part, we highlight disagreement scenarios encountered with **MultiSim**. In the second part, we report on the observations we made regarding results using the disagreement classifier from RQ₃.

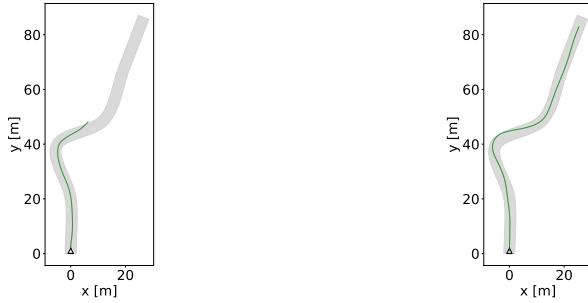


Fig. 7: Example of a disagreement scenario. Left: Vehicle simulated in BeamNG leaves short segment heading towards following segment. Right: Vehicle simulated in Donkey stays within the lane.

5.8.1 Disagreements

We manually analyzed the disagreements found by `MultiSim` across different configurations. We made following observations.

In the BeamNG simulator for DAVE-2, occasionally the vehicle switches to the left lane whenever the underlying segment is short, i.e., $< 15\text{m}$, and the following segment is a left turn (Figure 7). The car continues driving on the left lane towards the next segment, incurring in a failure.

A similar behavior was observed for TCP in Donkey and Udacity for sharp and short left turns followed by a sharp right turn, where the vehicle did not continue driving to the left, but proceeded straight to align with the subsequent path after the right turn. This behaviour could be explained by the fact that the vehicle observes the upcoming segment, and considers this shortcut as the immediate straight continuation of its path. However, in the remaining simulators and system under test configurations, the vehicle stays in the lane, even when driving on short segments.

Specifically to TCP, we observed that the model exhibits partially in all simulations environments a slightly higher XTE variation compared to DAVE-2 and ViT even when the single-simulator based failure count is less than for DAVE-2 or ViT. At the same time we could see that the model steers in general more extreme in curves compared to DAVE-2 or ViT which could explain its robust performance and the lower failure rate compared to the other system under tests.

Across all system under tests, we observed also that in the BeamNG simulator, the vehicle switches often to the left lane in the last but one segment after meeting a sharp curve. This behaviour might be related to the fact that the road terminates after the last segment. However, we did not observe a similar behaviour for the vehicle in the Donkey or Udacity simulators, which might be attributed to the visualization/rendering differences in the simulators.

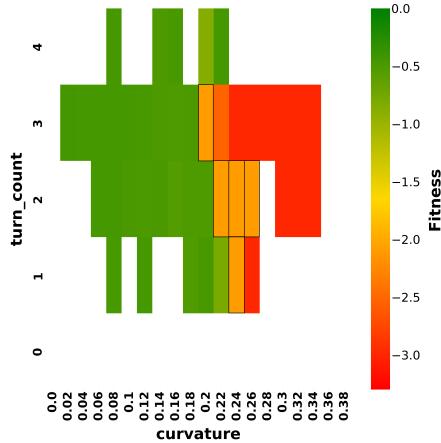


Fig. 8: Feature map of tests found in one run with **MultiSim** employing disagreement classifier (BD-P) for DAVE-2. Cells containing roads for which disagreements are predicted are marked in orange with a black box.

5.8.2 Feature Maps of Prediction-based Search

When evaluating the results of **MultiSim** augmented with the prediction classifier, we inspected the feature maps before failure validation. We could observe that disagreements are predicted in regions which are between failing and non-failing cells of the map. An example for DAVE-2 is shown in Figure 8. Cells that contain tests that are predicted as disagreements by the classifier are marked as orange. This observation highlights the fact that our classifier is likely to perform accurately in predicting disagreements, as disagreement include test cases which lie on the boundary between scenarios/roads that pass and fail. Similar situations were observed in the other simulator ensembles.

6 Discussion

6.1 Approach

While our study utilizes two simulators in the ensemble, **MultiSim** can be extended to incorporate a larger number of simulators. To reach consensus on the criticality, we suggest to use an odd number of simulators. Further, the feasibility of the approach with multiple simulators is basically constrained by the underlying computational resources. To mitigate this, test executions in **MultiSim** with multiple simulators might be parallelized, independent of the underlying test algorithm used.

6.2 Lessons Learned

Our case study shows that employing multiple simulators in the search does improve the effectiveness of finding simulator-agnostic failures compared to single-simulator testing. However, in the literature, test case generation approaches in general report the evaluation results which are based on single-simulator executions. We suggest that failures found *during* the test generation or at least the final identified failures, should be evaluated in a second simulator to assess whether the failures are likely to be simulator-agnostic.

6.3 Root Cause Analysis

Our paper proposes a testing technique to identify failing test cases where simulators agree on. It is out of the scope of this study to investigate the underlying root causes of disagreements between simulator executions, as it would likely require to inspect, debug, and profile the simulator’s code (Section 5.8 we qualitatively analyzed two disagreement scenarios). Additionally, as outlined in RQ₃, the **MultiSim** approach can be configured to search for tests on which simulators disagree to support the identification of the root causes behind failures. As shown by Jodat et al. [44] and Ben Abdessalem et al. [1], decision trees or decision rules can be then applied on underlying roads and simulation traces to derive conditions on input variables and analyze the root causes behind inconclusive and simulator-specific failures.

Among the possible explanations for the disagreements, we conjecture that variations in vehicle dynamics may have influenced our results. Specifically, we observed that the vehicle in Udacity exhibited reduced steering capability compared to those in Donkey and BeamNG. This could account for the high number of simulator-specific failures in Udacity and explain why Udacity appeared more reliable for independent validation (in BD) than during the multi-simulator search, where it introduced spurious failures. As actionable feedback, we recommend incorporating or explicitly detailing vehicle characteristics [72] in scenario definitions for future scenario-based ADAS testing.

We observed that BD performs best for DAVE-2 and ViT, whereas BU excels for TCP. This can be explained by the differing sensitivities of the ADAS architectures to physical versus perceptual variability. BD combines simulators with heterogeneous vehicle dynamics (soft- vs. rigid-body physics), which better exposes control-related faults, while BU combines simulators with distinct rendering and camera models, thus better revealing perception-related inconsistencies. Consequently, it appears that the optimal simulator ensemble depends on the dominant failure modes of the SUT.

6.4 Validation

In our study, the validation was dependent on three hyperparameters: the number of re-executions, the failure rate threshold, and the number of failures

selected per cell. We set the failure rate threshold to 100%. However, in practice, the threshold for the validation can be configured based on the number of simulators used. We expect that the more diverse simulation environments are employed, the more difficult would be to achieve failure rates of 100%. Regarding the number of re-executions we suggest to perform preliminary experiments with a different number of re-executions and compare the results, as conducted in our study. To select the number of selected failing tests per cell, we carried out preliminary experiments with a higher number of failures, i.e., 5. However, we did not observe significantly different evaluation results. Moreover, we did not select more tests per cell because the feature map already discretizes the search domain, making the diversity of roads within a single cell being likely low.

6.5 Generalization of Our Approach

While our empirical evaluation focuses on lane-keeping assistance systems, the **MultiSim** framework is not limited to this specific ADAS function. The proposed architecture and workflow, based on simulator coupling, multi-objective search, and cross-simulator agreement analysis, can be generalized to other ADAS and ADS functionalities.

First, the representation of tests is model-based and modular: road geometry is encoded through control points and spline interpolation, but this abstraction can be extended to other types of driving scenarios. For instance, intersections, pedestrian crossings, or obstacle avoidance can be represented using higher-level scene graphs or semantic maps instead of simple lane curves. Similarly, perception-heavy ADAS such as automatic emergency braking (AEB), adaptive cruise control (ACC), or lane-change assist could be evaluated by defining task-specific fitness functions (e.g., time-to-collision, minimum distance to obstacles, or stability of steering trajectories).

Second, the optimization procedure is independent of the specific fitness formulation. Our search and repopulation strategies can operate over any domain where simulation feedback is measurable, making it possible to reuse the same search backbone for tasks with different control policies or sensing modalities (e.g., radar-, LiDAR-, or camera-based). This in particular is enabled by the modular architecture of the extended framework OpenSBT as shown in various case studies [80].

Finally, the **MultiSim** framework could be extended to support system-level ADS validation by incorporating full autonomy stacks like Autoware [8] or Apollo [9]. In this context, cross-simulator disagreement analysis could help identify environment- or simulator-induced discrepancies in perception, planning, and control, offering a systematic means of analyzing simulation fidelity and testing robustness of the end-to-end driving system. However, since no existing autonomous driving stack is truly simulator-agnostic, extending **MultiSim** to full ADS validation remains an open engineering challenge that

requires standardized interfaces and consistent scene representations across simulators.

7 Threats to Validity

External validity concerns the extent to which our results can be generalized beyond the studied settings. We evaluated **MultiSim** on a well-established case study involving three lane-keeping models of increasing complexity, each tested across different simulator pairs drawn from a diverse set of environments, BeamNG, DonkeyCar, and Udacity. It would be interesting to investigate whether the results generalize to other ADAS functions, such as automated emergency braking, or to higher levels of autonomy using full driving stacks like Autoware [8] or Apollo [9]. This would require testing in more complex urban simulation environments, such as AWSIM [90] or CARLA [30]. The main difficulty lies in the technical integration of modern autonomous driving stacks with multiple simulators, as these frameworks are typically coupled with a single specific simulator. Achieving cross-simulator consistency in terms of scenarios, visual appearance, simulation conditions, and interactions with other actors poses significantly greater challenges than in our current study.

Internal validity risks refer to confounding factors that could affect the interpretation of results. To mitigate these, we adopted several strategies. First, to avoid biased evaluations due to unaligned deployment conditions (e.g., communication delays between the SUT and the simulator), we conducted a preliminary validation of XTE variation across manually defined driving scenarios. Second, we aligned the simulated roads across all simulators, building upon existing research [16]. Third, we defined validity constraints for generated roads and regenerated invalid samples after mutation. For the BU setup, simulators were executed on different hardware configurations due to compatibility constraints (BeamNG requires running on a Windows operating system). This may have influenced BU results, although BU did not outperform other configurations in our experiments. To ensure that detected failures were simulator-agnostic, we validated all failures identified by **MultiSim**, **DSS**, and **SingleSim** on one or two simulators left out during the search. We also re-ran failing tests multiple times to filter out failures caused by transient or flaky simulation conditions.

Construct validity threats relate to the adequacy of the chosen metrics. To evaluate performance, we used two primary metrics: the number of simulator-agnostic failures and the validity rate, i.e., the ratio between simulator-agnostic failures and validated failures. The first metric is a standard measure in test case generation studies [47], while the second captures the proportion of failures that remain consistent across simulators, reflecting robustness and transferability. Additionally, we analyzed the convergence of the Hypervolume (HV) indicator for **MultiSim** and **SingleSim**, showing that **SingleSim** converges earlier, while **MultiSim** continues improving slightly over time. We excluded **DSS** from HV analysis as it is not Pareto-based.

To foster *replicability*, we publicly release the implementation of **MultiSim**, along with all experimental data for **MultiSim**, **SingleSim**, and **DSS**. The search algorithms and result analyses are implemented using the modular and open-source framework OpenSBT [80], which supports easy integration of additional case studies for future research.

8 Related Work

In this section, we outline two categories of related work. The first category targets ADAS testing approaches that use one simulator for test generation. The second category reports studies using multiple simulators for testing.

8.1 Single-Simulator Approaches for ADAS Testing

Most current test generation methods use search-based approaches to automatically create test cases for DNN-driven ADAS systems [1, 2, 14, 45, 66, 74, 76, 89, 100]. In this domain, test cases consist of individual driving images or road topologies, which are rendered through a driving simulator. Abdessalem et al. [1, 2, 14] integrate genetic algorithms with machine learning techniques to test a pedestrian detection system. Mullins et al. [69] apply Gaussian processes to guide search-based test generation toward unexplored regions within the input space. Gambi et al. [33] leverage procedural content generation to propose a search-based test generation approach for ADAS. Riccio and Tonella [76] introduce, a model-based test generator that leverages Catmull-Rom splines, the test representation as used in our approach, to define road shapes, producing test cases at the behavioral frontier of self-driving car models. Arrieta et al. [7] apply a genetic algorithm to generate tests for cyber-physical systems, optimizing them across three dimensions: requirement coverage, test case similarity, and execution time. Lastly, Lu et al. [61] employ reinforcement learning to discover environmental configurations that induce crashes. DeepQTest [62] is a testing approach that uses reinforcement learning to learn environment configurations with a high chance of revealing ADAS misbehaviors. In another work [60], epigenetics algorithms are used to test ADAS in dynamically changing environments.

Among the fuzzing domain, DriveFuzz [46] leverages the physical state of the vehicle, along with oracles grounded in real-world traffic rules, to guide the fuzzer toward uncovering potential misbehaviors. AutoFuzz [101], on the other hand, focuses on fuzzing the test scenario specifications. Before initiating the fuzzing process, it employs a seed selection mechanism using a binary classifier that identifies seeds with a higher likelihood of violating traffic rules. AV-Fuzzer [54] applies a genetic algorithm, informed by the positioning of globally monitored non-player characters (NPCs) within each driving scenario. NPCs deemed to have a higher likelihood of safety violations [43] are prioritized for evolution. Cheng et al. [26] introduce BehaviorMiner, an unsupervised model that extracts temporal features from predefined scenarios and

employs clustering-based abstraction to group behaviors with similar features into abstract states.

All these approaches use a single-simulator approach, with some works validating their propositions on multiple simulators [26], for example, because different ADS are compatible/integrated only with specific simulation platforms [72]. Our approach differs from these solutions because it uses an ensemble of simulations during search-based testing to retrieve more accurate fitness signals, based on the consensus among the simulators.

8.2 Multi-Simulator Approaches for ADAS Testing

A study by Borg et al. [20] investigates the comparability of multiple simulators for testing a pedestrian vision detection system. The study evaluates a large set of test scenarios on both PreScan [78] and Pro-SiVIC [36]. The study reports inconsistent results in terms of safety violations and behaviors across these simulators. Consequently, the authors suggest that a single-simulator approach for ADAS testing might be unreliable, especially when failures are highly dependent on the chosen simulator. Moreover, a recent study by Amini et al. [5] has analyzed the degree of flakiness affecting ADAS testing. The study evaluates several simulators and ADAS showing that test flakiness is common and can significantly impact the test results. The authors propose the usage of machine learning classifiers to identify flaky ADAS tests. Another study by Wagner et al. [96] evaluates the translation of real-world driving scenarios to executed scenarios in a simulator. Their results show that a reprocessing error exists, which can be basically attributed to sensor model offsets and can be tackled by employing scenario-based sensor models.

To address simulator disagreements, Biagiola et al. [16] involve search-based testing across multiple simulators, provided that the same test scenario and ADAS under test can be consistently represented. Their method combines the predicted failure probabilities from each simulator, reporting a failure only when there is consensus among the simulators. In this method, the search process is conducted separately for each simulator, which can lead to numerous simulator-specific failures. Since discrepancies between simulators are only addressed after the searches terminate, this can lead to unnecessary consumption of the testing budget when such failures occur.

Following the approach of Biagiola et al. [16], we also base our methodology on multi-simulator, search-based testing, leveraging simulators capable of accommodating analogous configurations in terms of test scenarios and ADAS. However, our approach differs by performing a joint test evaluation *during the search*, ensuring that scenarios leading to simulator-dependent failures (i.e., simulator disagreements) are filtered out before progressing further.

9 Conclusion and Future Work

In this paper, we presented the approach **MultiSim**, to mitigate simulator-specific failures when testing ADAS using an ensemble of simulators. Our approach leverages search-based testing and executes test inputs in multiple simulators during the search to identify generalizable failures.

In our empirical study, we evaluated our approach on testing three DNN-enabled ADAS trained for lane-keeping. We compared our approach in terms of effectiveness and efficiency in identifying simulator-agnostic/valid failures, to single-simulator based testing as well as state-of-the-art testing approach which employs multiple simulators. The results of the study show that combining evaluation results from multiple simulators during testing outperforms single-simulator testing, identifying nearly 100% valid failures. Compared to existing multi-simulator approaches, **MultiSim** identifies, on average, more valid failures. In terms of efficiency, **MultiSim** performs similarly to other approaches. Additionally, the study demonstrates that machine learning can be used within **MultiSim** to predict disagreements between test outcomes across different simulators, increasing efficiency and reducing the variation in the number of valid failures.

Our future work is to extend our study to other use cases including more complex systems under tests such as full-stack ADAS and to investigate the root causes behind the disagreements of evaluation outcomes. Another future direction could consist in the integration of the proposed multi-simulator concept within other existing test generators, which can act as a complementary layer that analyzes simulator-induced flakiness and verifies the transferability of detected faults across environments.

10 Declarations

10.1 Funding

This research was funded by the Bavarian Ministry of Economic Affairs, Regional Development and Energy. Matteo Biagiola is partially supported by Fondo Istituzionale per la Ricerca granted by Università della Svizzera italiana (USI).

10.2 Ethical Approval

Not applicable.

10.3 Informed Consent

Not applicable.

10.4 Author Contributions

Lev Sorokin: conceptualization, methodology, implementation, evaluation, writing, review, editing. **Matteo Biagiola:** conceptualization, methodology, review, editing. **Andrea Stocco:** conceptualization, methodology, review, editing.

10.5 Data Availability Statement

All our results, the source code, and the simulator are accessible and can be reproduced [91].

10.6 Conflict of Interest

The authors declare no conflict of interest.

10.7 Clinical Trial Registration

Clinical trial number: Not applicable.

10.8 Acknowledgements

We acknowledge the support of Adnan Fidan in integrating the TCP system for the evaluation of our approach. We also thank the reviewers for their valuable feedback.

References

1. Abdessalem, R.B., Nejati, S., Briand, L.C., Stifter, T.: Testing vision-based control systems using learnable evolutionary algorithms. In: M. Chaudron, I. Crnkovic, M. Chechik, M. Harman (eds.) Proceedings of the 40th International Conference on Software Engineering, ICSE 2018, Gothenburg, Sweden, May 27 - June 03, 2018, pp. 1016–1026. ACM (2018). DOI 10.1145/3180155.3180160. URL <https://doi.org/10.1145/3180155.3180160>
2. Abdessalem, R.B., Panichella, A., Nejati, S., Briand, L.C., Stifter, T.: Testing autonomous cars for feature interaction failures using many-objective search. In: Proceedings of ASE '18, ASE 2018, pp. 143–154. ACM (2018). DOI 10.1145/3238147.3238192. URL <http://doi.acm.org/10.1145/3238147.3238192>
3. Afzal, A., Katz, D.S., Le Goues, C., Timperley, C.S.: Simulation for robotics test automation: Developer perspectives. In: 2021 14th IEEE Conference on Software Testing, Verification and Validation (ICST), pp. 263–274. IEEE (2021). DOI 10.1109/ICST49551.2021.00036
4. Ali, Q., Stocco, A., Mariani, L., Riganelli, O.: OpenCat: Improving Interoperability of ADS Testing. In: Proceedings of 47th International Conference on Software Engineering Workshops, ICSEW '24, p. 10 pages. IEEE (2025)

5. Amini, M.H., Naseri, S., Nejati, S.: Evaluating the impact of flaky simulators on testing autonomous driving systems. *Empirical Softw. Engg.* **29**(2) (2024). DOI 10.1007/s10664-023-10433-5. URL <https://doi.org/10.1007/s10664-023-10433-5>
6. Ando, A., Gidaris, S., Bursuc, A., Puy, G., Boulch, A., Marlet, R.: Rangevit: Towards vision transformers for 3d semantic segmentation in autonomous driving. In: CVPR (2023)
7. Arrieta, A., Wang, S., Markiegi, U., Sagardui, G., Etxeberria, L.: Search-based test case generation for cyber-physical systems. In: 2017 IEEE Congress on Evolutionary Computation (CEC), pp. 688–697 (2017). DOI 10.1109/CEC.2017.7969377
8. Autoware. <https://autoware.org> (2024)
9. Baidu: Baidu Apollo. <https://github.com/ApolloAuto/apollo/> (2024). URL <https://github.com/ApolloAuto/apollo>
10. Baidu Inc.: Baidu Apolloscapes Dataset. <https://apolloscape.auto/index.html> (2018). Accessed: [2024-01-15]
11. Baresi, L., Hu, D.Y.X., Stocco, A., Tonella, P.: Efficient domain augmentation for autonomous driving testing using diffusion models. In: Proceedings of 47th International Conference on Software Engineering, ICSE '25. IEEE (2025)
12. BeamNG.research: BeamNG GmbH. <https://www.beamng.gmbh/research> (2022). Online; accessed 18/08/2019
13. Ben Abdessalem, R., Nejati, S., Briand, L.C., Stifter, T.: Testing advanced driver assistance systems using multi-objective search and neural networks. ASE, p. 63–74 (2016). DOI 10.1145/2970276.2970311. URL <https://doi.org/10.1145/2970276.2970311>
14. Ben Abdessalem, R., Nejati, S., Briand, L.C., Stifter, T.: Testing advanced driver assistance systems using multi-objective search and neural networks. In: 2016 31st IEEE/ACM International Conference on Automated Software Engineering (ASE), pp. 63–74 (2016)
15. Biagiola, M., Klikovits, S.: Sbft tool competition 2024 - cyber-physical systems track. In: Proceedings of the 17th ACM/IEEE International Workshop on Search-Based and Fuzz Testing, SBFT '24, p. 33–36. Association for Computing Machinery, New York, NY, USA (2024). DOI 10.1145/3643659.3643932. URL <https://doi.org/10.1145/3643659.3643932>
16. Biagiola, M., Stocco, A., Riccio, V., Tonella, P.: Two is better than one: digital siblings to improve autonomous driving testing. *Empirical Softw. Engg.* **29**(4) (2024). DOI 10.1007/s10664-024-10458-4. URL <https://doi.org/10.1007/s10664-024-10458-4>
17. Biagiola, M., Tonella, P.: Boundary state generation for testing and improvement of autonomous driving systems (2023). DOI 10.1109/TSE.2024.3420816. URL <https://doi.org/10.1109/TSE.2024.3420816>
18. Birchler, C., Khatiri, S., Bosshard, B., Gambi, A., Panichella, S.: Machine learning-based test selection for simulation-based testing of self-driving cars software. *Empir. Softw. Eng.* **28**(3), 71 (2023). DOI 10.1007/S10664-023-10286-Y. URL <https://doi.org/10.1007/s10664-023-10286-y>
19. Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L.D., Monfort, M., Muller, U., Zhang, J., Zhang, X., Zhao, J., Zieba, K.: End to end learning for self-driving cars. *CoRR* **abs/1604.07316** (2016). URL <http://arxiv.org/abs/1604.07316>
20. Borg, M., Abdessalem, R., Nejati, S., Jegeden, F., Shin, D.: Digital twins are not monozygotic – cross-replicating adas testing in two industry-grade automotive simulators pp. 383–393 (2021). DOI 10.1109/ICST49551.2021.00050. URL <https://doi.ieeecomputersociety.org/10.1109/ICST49551.2021.00050>
21. Breiman, L.: Random forests. *Mach. Learn.* **45**(1), 5–32 (2001). DOI 10.1023/A:1010933404324. URL <https://doi.org/10.1023/A:1010933404324>
22. Breiman, L., Friedman, J., Stone, C., Olshen, R.: Classification and Regression Trees. Taylor & Francis (1984). URL <https://books.google.de/books?id=JwQx-W0mSyQC>
23. Cabrera, A.: Logistic regression analysis in higher education: An applied perspective. Higher education: Handbook of theory and research X/Agathon Press (1994)
24. Cerf, V.G.: A comprehensive self-driving car test. *Commun. ACM* **61**(2), 7–7 (2018). DOI 10.1145/3177753. URL <http://doi.acm.org/10.1145/3177753>
25. Chen, X., Biagiola, M., Riccio, V., d'Amorim, M., Stocco, A.: XMutant: XAI-based Fuzzing for Deep Learning Systems (2025). URL <https://arxiv.org/abs/2503.07222>

26. Cheng, M., Zhou, Y., Xie, X.: Behavexplor: Behavior diversity guided testing for autonomous driving systems. In: Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2023, p. 488–500. Association for Computing Machinery, New York, NY, USA (2023). DOI 10.1145/3597926.3598072. URL <https://doi.org/10.1145/3597926.3598072>
27. Codevilla, F., López, A.M., Koltun, V., Dosovitskiy, A.: On offline evaluation of vision-based driving models. CoRR **abs/1809.04843** (2018). URL <http://arxiv.org/abs/1809.04843>
28. Cortes, C., Vapnik, V.: Support-vector networks. *Mach. Learn.* **20**(3), 273–297 (1995). DOI 10.1023/A:1022627411411. URL <https://doi.org/10.1023/A:1022627411411>
29. Donkey Car. <https://www.donkeycar.com/> (2021)
30. Dosovitskiy, A., Ros, G., Codevilla, F., López, A., Koltun, V.: CARLA: an open urban driving simulator. CoRR **abs/1711.03938** (2017). URL <http://arxiv.org/abs/1711.03938>
31. Friedman, J.H.: Stochastic gradient boosting. *Computational Statistics & Data Analysis* **38**(4), 367–378 (2002). DOI [https://doi.org/10.1016/S0167-9473\(01\)00065-2](https://doi.org/10.1016/S0167-9473(01)00065-2). URL <https://www.sciencedirect.com/science/article/pii/S0167947301000652>. Nonlinear Methods and Data Mining
32. Gambi, A., Jahangirova, G., Riccio, V., Zampetti, F.: Sbst tool competition 2022. In: Proceedings of the 15th Workshop on Search-Based Software Testing, SBST '22, p. 25–32. Association for Computing Machinery, New York, NY, USA (2023). DOI 10.1145/3526072.3527538. URL <https://doi.org/10.1145/3526072.3527538>
33. Gambi, A., Mueller, M., Fraser, G.: Automatically testing self-driving cars with search-based procedural content generation. In: Proceedings of ISSTA '19 (2019)
34. García, S., Strüber, D., Brugali, D., Berger, T., Pelliccione, P.: Robotics software engineering: A perspective from the service robotics domain. In: Proceedings of ESEC/FSE '20, ESEC/FSE 2020, pp. 593–604. ACM, USA (2020). DOI 10.1145/3368089.3409743
35. Grewal, R., Tonella, P., Stocco, A.: Predicting Safety Misbehaviours in Autonomous Driving Systems using Uncertainty Quantification. In: Proceedings of 17th IEEE International Conference on Software Testing, Verification and Validation, ICST '24, p. 12 pages. IEEE (2024)
36. Group, E.: Esi prosivic. <https://myesi.esi-group.com/downloads/software-downloads/pro-sivic-2021.0> (2021)
37. Haq, F.U., Shin, D., Nejati, S., Briand, L.: Can offline testing of deep neural networks replace their online testing? a case study of automated driving systems. *Empirical Softw. Engng.* **26**(5) (2021). DOI 10.1007/s10664-021-09982-4. URL <https://doi.org/10.1007/s10664-021-09982-4>
38. Humbatova, N., Jahangirova, G., Bavota, G., Riccio, V., Stocco, A., Tonella, P.: Taxonomy of Real Faults in Deep Learning Systems. In: Proceedings of 42nd International Conference on Software Engineering, ICSE'20, p. 12 pages. ACM, New York, NY, USA (2020). DOI 10.1145/3377811.3380395
39. Humeniuk, D., Khomh, F., Antoniol, G.: A search-based framework for automatic generation of testing environments for cyber-physical systems. *Inf. Softw. Technol.* **149**, 106936 (2022). DOI <https://doi.org/10.1016/j.infsof.2022.106936>. URL <https://www.sciencedirect.com/science/article/pii/S0950584922000866>
40. Humeniuk, D., Khomh, F., Antoniol, G.: Reinforcement learning informed evolutionary search for autonomous systems testing. *ACM Trans. Softw. Eng. Methodol.* (2024). DOI 10.1145/3680468. URL <https://doi.org/10.1145/3680468>. Just Accepted
41. Hussain, M., Ali, N., Hong, J.E.: Deepguard: A framework for safeguarding autonomous driving systems from inconsistent behaviour. *Automated Software Engng.* **29**(1) (2022). DOI 10.1007/s10515-021-00310-0. URL <https://doi.org/10.1007/s10515-021-00310-0>
42. Jahangirova, G., Stocco, A., Tonella, P.: Quality metrics and oracles for autonomous vehicles testing. In: Proceedings of 14th IEEE International Conference on Software Testing, Verification and Validation, ICST '21. IEEE (2021)
43. Jha, S., Banerjee, S.S., Tsai, T., Hari, S.K.S., Sullivan, M.B., Kalbarczyk, Z.T., Keckler, S.W., Iyer, R.K.: ML-based fault injection for autonomous vehicles: A case

for bayesian fault injection. 2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) pp. 112–124 (2019). URL <https://api.semanticscholar.org/CorpusID:195776612>

44. Jodat, B.A., Chandar, A., Nejati, S., Sabetzadeh, M.: Test generation strategies for building failure models and explaining spurious failures. ACM Trans. Softw. Eng. Methodol. **33**(4) (2024). DOI 10.1145/3638246. URL <https://doi.org/10.1145/3638246>
45. Kim, J., Feldt, R., Yoo, S.: Guiding deep learning system testing using surprise adequacy. In: Proceedings of the 41st International Conference on Software Engineering, ICSE '19, pp. 1039–1049. IEEE Press, Piscataway, NJ, USA (2019). DOI 10.1109/ICSE.2019.00108. URL <https://doi.org/10.1109/ICSE.2019.00108>
46. Kim, S., Liu, M., Rhee, J.J., Jeon, Y., Kwon, Y., Kim, C.H.: DriveFuzz. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security. ACM (2022). DOI 10.1145/3548606.3560558. URL <https://doi.org/10.1145/3548606.3560558>
47. Klück, F., Zimmermann, M., Wotawa, F., Nica, M.: Genetic algorithm-based test parameter optimization for adas system testing. In: QRS, pp. 418–425 (2019). DOI 10.1109/QRS.2019.00058
48. Kolesnikov, A., Dosovitskiy, A., Weissenborn, D., Heigold, G., Uszkoreit, J., Beyer, L., Minderer, M., Dehghani, M., Houlsby, N., Gelly, S., Unterthiner, T., Zhai, X.: An image is worth 16x16 words: Transformers for image recognition at scale (2021)
49. Koroglu, Y., Wotawa, F.: Towards a review on simulated adas/ad testing. In: 2023 IEEE/ACM International Conference on Automation of Software Test (AST), pp. 112–122 (2023). DOI 10.1109/AST58925.2023.00015
50. Lambertenghi, S.C., Leonhard, H., Stocco, A.: Benchmarking image perturbations for testing automated driving assistance systems. In: Proceedings of the IEEE 18th IEEE International Conference on Software Testing, Verification and Validation, ICST '25, p. 12 pages. IEEE (2025)
51. Lambertenghi, S.C., Stocco, A.: Assessing quality metrics for neural reality gap input mitigation in autonomous driving testing. In: Proceedings of 17th IEEE International Conference on Software Testing, Verification and Validation, ICST '24, p. 12 pages. IEEE (2024)
52. Lambertenghi, S.C., Valdez, M.F., Stocco, A.: A multi-modality evaluation of the reality gap in autonomous driving systems. In: Proceedings of the 40th IEEE/ACM International Conference on Automated Software Engineering, ASE '25, p. 12 pages. IEEE (2025)
53. Li, D., Auerbach, P., Okhrin, O.: Autonomous driving small-scale cars: A survey of recent development (2024). URL <https://arxiv.org/abs/2404.06229>
54. Li, G., Li, Y., Jha, S., Tsai, T., Sullivan, M., Hari, S.K.S., Kalbarczyk, Z., Iyer, R.: Av-fuzzer: Finding safety violations in autonomous driving systems. In: 2020 IEEE 31st International Symposium on Software Reliability Engineering (ISSRE), pp. 25–36 (2020). DOI 10.1109/ISSRE5003.2020.00012
55. Li, M., Chen, T., Yao, X.: How to evaluate solutions in pareto-based search-based software engineering: A critical review and methodological guidance. IEEE Transactions on Software Engineering **48**(5), 1771–1799 (2022). DOI 10.1109/TSE.2020.3036108
56. Li, M., Yao, X.: Quality evaluation of solution sets in multiobjective optimisation: A survey. ACM Comput. Surv. **52**(2), 1–38 (2019). DOI 10.1145/3300148. URL <https://doi.org/10.1145/3300148>
57. Li, Y., Yuan, W., Zhang, S., Yan, W., Shen, Q., Wang, C., Yang, M.: Choose your simulator wisely: A review on open-source simulators for autonomous driving. IEEE Transactions on Intelligent Vehicles **9**(5), 4861–4876 (2024). DOI 10.1109/tiv.2024.3374044
58. Liu, D., Metzman, J., Bohme, M., Chang, O., Arya, A.: SBFT Tool Competition 2023 - Fuzzing Track. In: 2023 IEEE/ACM International Workshop on Search-Based and Fuzz Testing (SBFT), pp. 51–54. IEEE Computer Society, Los Alamitos, CA, USA (2023). DOI 10.1109/SBFT59156.2023.00016. URL <https://doi.ieee.org/10.1109/SBFT59156.2023.00016>

59. Lou, G., Deng, Y., Zheng, X., Zhang, M., Zhang, T.: Testing of autonomous driving systems: where are we and where should we go? In: Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2022, p. 31–43. Association for Computing Machinery, New York, NY, USA (2022). DOI 10.1145/3540250.3549111. URL <https://doi.org/10.1145/3540250.3549111>
60. Lu, C., Ali, S., Yue, T.: Epitester: Testing autonomous vehicles with epigenetic algorithm and attention mechanism. *IEEE Transactions on Software Engineering* pp. 1–19 (2024). DOI 10.1109/TSE.2024.3449429
61. Lu, C., Shi, Y., Zhang, H., Zhang, M., Wang, T., Yue, T., Ali, S.: Learning configurations of operating environment of autonomous vehicles to maximize their collisions. *IEEE Transactions on Software Engineering* **49**(1), 384–402 (2023). DOI 10.1109/TSE.2022.3150788
62. Lu, C., Yue, T., Zhang, M., Ali, S.: Deepqtest: Testing autonomous driving systems with reinforcement learning and real-world weather data (2023). URL <https://arxiv.org/abs/2310.05170>
63. Marko, N., Ruebsam, J., Biehn, A., Schneider, H.: Scenario-based testing of adas - integration of the open simulation interface into co-simulation for function validation. In: Proceedings of the 9th International Conference on Simulation and Modeling Methodologies, Technologies and Applications - SIMULTECH,, pp. 255–262. IN-STICC, SciTePress (2019). DOI 10.5220/0007838302550262
64. Matinnejad, R., Nejati, S., Briand, L., Bruckmann, T., Poull, C.: Search-based automated testing of continuous controllers: Framework, tool support, and case studies. *Information and Software Technology* **57**, 705–722 (2015). DOI <https://doi.org/10.1016/j.infsof.2014.05.007>. URL <https://www.sciencedirect.com/science/article/pii/S0950584914001244>
65. McMinn, P.: Search-based software testing: Past, present and future. In: 2011 IEEE Fourth International Conference on Software Testing, Verification and Validation Workshops, pp. 153–163 (2011). DOI 10.1109/ICSTW.2011.100
66. Moghadam, M.H., Borg, M., Saadatmand, M., Mousavirad, S.J., Bohlin, M., Lisper, B.: Machine learning testing in an adas case study using simulation-integrated bio-inspired sbt. *J. Softw.: Evol. Process* p. e2591
67. Mokhtarian, A., Scheffe, P., Kloock, M., Schäfer, S., Heeseung Bang, Viet-Anh Le, Sangeet Ulhas, Betz, J., Wilson, S., Berman, S., Prorok, A., Alrifae, B.: A survey on small-scale testbeds for connected and automated vehicles and robot swarms (2024). DOI 10.13140/RG.2.2.16176.74248/1. URL <https://rgdoi.net/10.13140/RG.2.2.16176.74248/1>
68. Mouret, J.B., Clune, J.: Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015)
69. Mullins, G.E., Stankiewicz, P.G., Hawthorne, R.C., Gupta, S.K.: Adaptive generation of challenging scenarios for testing and evaluation of autonomous vehicles. *Journal of Systems and Software* **137**, 197–215 (2018). DOI <https://doi.org/10.1016/j.jss.2017.10.031>. URL <https://www.sciencedirect.com/science/article/pii/S0164121217302546>
70. Nejati, S., Sorokin, L., Safin, D., Formica, F., Mahboob, M.M., Menghi, C.: Reflections on surrogate-assisted search-based testing: A taxonomy and two replication studies based on industrial adas and simulink models. *Information and Software Technology* **163**, 107286 (2023)
71. Opletal, J.: China's massive adas test: 36 cars, 15 hazard scenarios, 216 crashes (2025). URL <https://carnewschina.com/2025/07/24/chinas-massive-adas-test-36-cars-15-hazard-scenarios-216-crashes/>
72. Pan, Q., Wang, T., Ma, J., Arcaini, P., Yue, T.: Simulation-based safety assessment of vehicle characteristics variations in autonomous driving systems. *ACM Trans. Softw. Eng. Methodol.* (2025). DOI 10.1145/3743673. URL <https://doi.org/10.1145/3743673>
73. Parry, O., Kapfhammer, G.M., Hilton, M., McMinn, P.: A survey of flaky tests. *ACM Trans. Softw. Eng. Methodol.* **31**(1), 17:1–17:74 (2022). DOI 10.1145/3476105. URL <https://doi.org/10.1145/3476105>

74. Pei, K., Cao, Y., Yang, J., Jana, S.: Deepxplore: Automated whitebox testing of deep learning systems. In: Proceedings of the 26th Symposium on Operating Systems Principles, SOSP '17, pp. 1–18. ACM, New York, NY, USA (2017). DOI 10.1145/3132747.3132785. URL <http://doi.acm.org/10.1145/3132747.3132785>
75. Riccio, V., Jahangirova, G., Stocco, A., Humbatova, N., Weiss, M., Tonella, P.: Testing Machine Learning based Systems: A Systematic Mapping. *Empirical Software Engineering* (2020)
76. Riccio, V., Tonella, P.: Model-Based Exploration of the Frontier of Behaviours for Deep Learning System Testing. In: Proceedings of ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE '20 (2020)
77. Riccio, V., Tonella, P.: When and why test generators for deep learning produce invalid inputs: an empirical study. In: 2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE), ICSE '23, pp. 1161–1173. IEEE Press (2023). DOI 10.1109/ICSE48619.2023.00104. URL <https://doi.org/10.1109/ICSE48619.2023.00104>
78. Software, S.D.I.: Simcenter prescan. <https://www.plm.automation.siemens.com/global/en/products/simcenter/prescan.html> (2022). URL <https://plm.sw.siemens.com/de-DE/simcenter/autonomous-vehicle-solutions/prescan/>
79. Sorokin, L., Kerscher, N.: Guiding the search towards failure-inducing test inputs using support vector machines. In: Proceedings of the 5th IEEE/ACM International Workshop on Deep Learning for Testing and Testing for Deep Learning, DeepTest '24, p. 9–12. Association for Computing Machinery, New York, NY, USA (2024). DOI 10.1145/3643786.3648023. URL <https://doi.org/10.1145/3643786.3648023>
80. Sorokin, L., Munaro, T., Safin, D., Liao, B.H.C., Molin, A.: OpenSBT: A modular framework for search-based testing of automated driving systems (2023). DOI 10.1145/3639478.3640027. URL <https://github.com/opensbt/opensbt-core>
81. Sorokin, L., Safin, D., Nejati, S.: Can search-based testing with pareto optimization effectively cover failure-revealing test inputs? *Empirical Software Engineering* **30**(1), 26 (2024). DOI 10.1007/s10664-024-10564-3. URL <https://doi.org/10.1007/s10664-024-10564-3>
82. Stocco, A., Nunes, P.J., d'Amorim, M., Tonella, P.: ThirdEye: Attention Maps for Safe Autonomous Driving Systems. In: Proceedings of 37th IEEE/ACM International Conference on Automated Software Engineering, ASE '22. IEEE/ACM (2022). DOI 10.1145/3551349.3556968
83. Stocco, A., Pulfer, B., Tonella, P.: Mind the Gap! A Study on the Transferability of Virtual Versus Physical-World Testing of Autonomous Driving Systems. *IEEE Transactions on Software Engineering* **49**(04), 1928–1940 (2023). DOI 10.1109/TSE.2022.3202311
84. Stocco, A., Pulfer, B., Tonella, P.: Model vs system level testing of autonomous driving systems: A replication and extension study. *Empirical Softw. Engng.* **28**(3), 73 (2023). DOI 10.1007/s10664-023-10306-x. URL <https://doi.org/10.1007/s10664-023-10306-x>
85. Stocco, A., Tonella, P.: Confidence-driven weighted retraining for predicting safety-critical failures in autonomous driving systems. *Journal of Software: Evolution and Process* (2021). DOI 10.1002/sm.2386. URL <https://doi.org/10.1002/sm.2386>
86. Stocco, A., Weiss, M., Calzana, M., Tonella, P.: Misbehaviour prediction for autonomous driving systems. In: Proceedings of 42nd International Conference on Software Engineering, ICSE '20, p. 12 pages. ACM (2020)
87. Tang, S., Zhang, Z., Zhang, Y., Zhou, J., Guo, Y., Liu, S., Guo, S., Li, Y., Ma, L., Xue, Y., Liu, Y.: A survey on automated driving system testing: Landscapes and trends. *ACM Trans. Softw. Eng. Methodol.* **32**(5), 124:1–124:62 (2023). DOI 10.1145/3579642. URL <https://doi.org/10.1145/3579642>
88. Tawn Kramer, M.E., contributors: Donkeycar. <https://www.donkeycar.com/> (2022)
89. Tian, Y., Pei, K., Jana, S., Ray, B.: Deptest: Automated testing of deep-neural-network-driven autonomous cars. In: Proceedings of ICSE '18, ICSE '18, pp. 303–314. ACM, New York, NY, USA (2018). DOI 10.1145/3180155.3180220. URL <http://doi.acm.org/10.1145/3180155.3180220>
90. Tier IV, Inc.: Awsim: Autonomous driving simulator by tier iv. <https://github.com/tier4/AWSIM> (2023). Accessed: 2025-11-03

91. Replication package. (2025). URL <https://github.com/ast-fortiss-tum/MultiSim>
92. of Transportation, U.D.: A framework for automated driving system testable cases and scenarios. https://rosap.ntl.bts.gov/view/dot/38824/dot_38824_DS1.pdf (2018)
93. of Transportation, U.D.: Standing general order on crash reporting for level 2 advanced driver assistance systems. <https://www.nhtsa.gov/sites/nhtsa.gov/files/2022-06/ADAS-L2-SGO-Report-June-2022.pdf> (2022)
94. Udacity: A self-driving car simulator built with Unity. <https://github.com/udacity/self-driving-car-sim> (2017). Online; accessed 18 August 2019
95. VISSIM: VISSIM website. <https://www.ptvgroup.com/en-us/products/ptv-vissim> (2023)
96. Wagner, S., Groh, K., Kühbeck, T., Knoll, A.: Towards cross-verification and use of simulation in the assessment of automated driving. In: 2019 IEEE Intelligent Vehicles Symposium (IV), pp. 1589–1596 (2019). DOI 10.1109/IVS.2019.8814268
97. Wu, P., Jia, X., Chen, L., Yan, J., Li, H., Qiao, Y.: Trajectory-guided control prediction for end-to-end autonomous driving: a simple yet strong baseline. In: Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22. Curran Associates Inc., Red Hook, NY, USA (2022)
98. Zeller, A.: Search-based testing and system testing: A marriage in heaven. SBST, pp. 49–50 (2017). DOI 10.1109/SBST.2017.3
99. Zhang, M., Zhang, Y., Zhang, L., Liu, C., Khurshid, S.: DeepRoad: GAN-based metamorphic testing and input validation framework for autonomous driving systems. In: Proceedings of the 33rd ACM/IEEE ASE, ASE '18, p. 132–142. ACM (2018). DOI 10.1145/3238147.3238187
100. Zhang, M., Zhang, Y., Zhang, L., Liu, C., Khurshid, S.: Deeproad: Gan-based metamorphic testing and input validation framework for autonomous driving systems. In: Proceedings of ASE '18, ASE 2018, pp. 132–142. ACM, New York, NY, USA (2018). DOI 10.1145/3238147.3238187. URL <http://doi.acm.org/10.1145/3238147.3238187>
101. Zhong, Z., Kaiser, G., Ray, B.: Neural network guided evolutionary fuzzing for finding traffic violations of autonomous vehicles. IEEE Trans. Softw. Eng. **49**(4), 1860–1875 (2023). DOI 10.1109/TSE.2022.3195640. URL <https://doi.org/10.1109/TSE.2022.3195640>
102. Zohdinasab, T., Riccio, V., Gambi, A., Tonella, P.: Deephyperion: exploring the feature space of deep learning-based systems through illumination search. In: Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2021, p. 79–90. Association for Computing Machinery, New York, NY, USA (2021). DOI 10.1145/3460319.3464811. URL <https://doi.org/10.1145/3460319.3464811>