Paracosm: A Test Framework for Autonomous Driving Simulations

Rupak Majumdar¹ , Aman Mathur¹ , Marcus Pirron¹ , Laura Stegner² , and Damien Zufferey¹

¹ MPI-SWS, Kaiserslautern, Germany {rupak, mathur, mpirron, zufferey}@mpi-sws.org

² University of Wisconsin, Madison, USA stegner@wisc.edu

Abstract. Systematic testing of autonomous vehicles operating in complex real-world scenarios is a difficult and expensive problem. We present PARACOSM, a framework for writing systematic test scenarios for autonomous driving simulations. PARACOSM allows users to programmatically describe complex driving situations with specific features, e.g., road layouts and environmental conditions, as well as reactive temporal behaviors of other cars and pedestrians. A systematic exploration of the state space, both for visual features and for reactive interactions with the environment is made possible. We define a notion of test coverage for parameter configurations based on combinatorial testing and low dispersion sequences. Using fuzzing on parameter configurations, our automatic test generator can maximize coverage of various behaviors and find problematic cases. Through empirical evaluations, we demonstrate the capabilities of PARACOSM in programmatically modeling parameterized test environments, and in finding problematic scenarios.

Keywords: Autonomous driving · Testing · Reactive programming.

1 Introduction

Building autonomous driving systems requires complex and intricate engineering effort. At the same time, ensuring their reliability and safety is an extremely difficult task. There are serious public safety and trust concerns [63], aggravated by recent accidents involving autonomous cars [48]. Software in such vehicles combine well-defined tasks such as trajectory planning, steering, acceleration and braking, with underspecified tasks such as building a semantic model of the environment from raw sensor data and making decisions using this model. Unfortunately, these underspecified tasks are critical to the safe operation of autonomous vehicles. Therefore, testing in large varieties of realistic scenarios is the only way to build confidence in the correctness of the overall system.

Running real tests is a necessary, but slow and costly process. It is difficult to reproduce corner cases due to infrastructure and safety issues; one can neither run over pedestrians to demonstrate a failing test case, nor wait for specific weather and road conditions. Therefore, the automotive industry tests

Fig. 1: A PARACOSM program consists of parameterized reactive components such as the test vehicle, the environment, road networks, other actors and their behaviors, and monitors. The test input generation scheme guarantees good coverage over the parameter space. The test scenario depicted here shows a test vehicle stopping for a jaywalking pedestrian.

autonomous systems in virtual simulation environments [21, 26, 53, 61, 68, 72]. Simulation reduces the cost per test, and more importantly, gives precise control over all aspects of the environment, so as to test corner cases.

A major limitation of current tools is the lack of customizability: they either provide a GUI-based interface to design an environment piece-by-piece, or focus on bespoke pre-made environments. This makes the setup of varied scenarios difficult and time consuming. Though exploiting parametricity in simulation is useful and effective [10, 23, 31, 67], the cost of environment setup, and navigating large parameter spaces, is quite high [31]. Prior works have used bespoke environments with limited parametricity. More recently, programmatic interfaces have been proposed [27] to make such test procedures more systematic. However, the simulated environments are largely still fixed, with no dynamic behavior.

In this work, we present PARACOSM, a programmatic interface that enables the design of *parameterized environments* and *test cases*. Test parameters control the environment and the behaviors of the actors involved. PARACOSM supports various test input generation strategies, and we provide a notion of coverage for these. Rather than computing coverage over intrinsic properties of the system under test (which is not yet understood for neural networks [39]), our coverage criteria is over the space of test parameters. Figure 1 depicts the various parts of a PARACOSM test. A PARACOSM program represents a family of tests, where each instantiation of the program’s parameters is a concrete test case.

PARACOSM is based on a synchronous reactive programming model [13, 35, 40, 70]. Components, such as road segments or cars, receive streams of inputs and produce streams of outputs over time. In addition, components have graphical assets to describe their appearance for an underlying visual rendering engine and physical properties for an underlying physics simulator. For example, a vehicle in PARACOSM not only has code that reads in sensor feeds and outputs steering angle or braking, but also has a textured mesh representing its shape, position

and orientation in 3D space, and a physics model for its dynamical behavior. A PARACOSM configuration consists of a composition of several components. Using a set of system-defined components (road segments, cars, pedestrians, etc.) combined using expressive operations from the underlying reactive programming model, users can set up complex temporally varying driving scenarios. For example, one can build an urban road network with intersections, pedestrians and vehicular traffic, and parameterize both, environment conditions (lighting, fog), and behaviors (when a pedestrian crosses a street).

Streams in the world description can be left “open” and, during testing, PARACOSM automatically generates sequences of values for these streams. We use a coverage strategy based on *k-wise combinatorial coverage* [14, 38] for discrete variables and *dispersion* for continuous variables. Intuitively, *k-wise coverage* ensures that, for a programmer-specified parameter *k*, all possible combinations of values of any *k* discrete parameters are covered by tests. Low dispersion [57] ensures that there are no “large empty holes” left in the continuous parameter space. PARACOSM uses an automatic test generation strategy that offers high coverage based on random sampling over discrete parameters and *deterministic* quasi-Monte Carlo methods for continuous parameters [49, 57].

Like many of the projects referenced before, our implementation performs simulations inside a game engine. However, PARACOSM configurations can also be output to the OPENDRIVE format [7] for use with other simulators, which is more in-line with the current industry standard. We demonstrate through various case studies how PARACOSM can be an effective testing framework for both qualitative properties (crash) and quantitative properties (distance maintained while following a car, or image misclassification).

Our main contributions are the following: (I) We present a programmable and expressive framework for programmatically modeling complex and parameterized scenarios to test autonomous driving systems. Using PARACOSM one can specify the environment’s layout, behaviors of actors, and expose parameters to a systematic testing infrastructure. (II) We define a notion of test coverage based on combinatorial *k-wise coverage* in discrete space and low dispersion in continuous space. We show a test generation strategy based on fuzzing that theoretically guarantees good coverage. (III) We demonstrate empirically that our system is able to express complex scenarios and automatically test autonomous driving agents and find incorrect behaviors or degraded performance.

2 Paracosm through Examples

We now provide a walkthrough of PARACOSM through a testing example. Suppose we have an autonomous vehicle to test. Its implementation is wrapped into a parameterized class:

```
AutonomousVehicle(start, model, controller) {
    void run(...) { ... } }
```

where the `model` ranges over possible car models (appearance, physics), and the `controller` implements an autonomous controller. The goal is to test this class in

many different driving scenarios, including different road networks, weather and light conditions, and other car and pedestrian traffic. We show how PARACOSM enables writing such tests as well as generate test inputs automatically.

A *test configuration* consists of a composition of *reactive objects*. The following is an outline of a test configuration in PARACOSM, in which the autonomous vehicle drives on a road with a pedestrian wanting to cross. We have simplified the API syntax for the sake of clarity and omit the enclosing `Test` class. In the code segments, we use ‘:’ for named arguments.

```

1 // Test parameters
2 light = VarInterval(0.2, 1.0) // value in [0.2, 1.0]
3 nlanes = VarEnum({2,4,6}) // value is 2, 4 or 6
4 // Description of environment
5 w = World(light:light, fog:0)
6 // Create a road segment
7 r = StraightRoadSegment(len:100, nlanes:nlanes)
8 // The autonomous vehicle controlled by the SUT
9 v = AutonomousVehicle(start:...,model:...,controller:...)
10 // Some other actor(s)
11 p = Pedestrian(start:..., model:..., ...)
12 // Monitor to check some property
13 c = CollisionMonitor(v)
14 // Place elements in the world
15 run_test(env: {w, r, v, p}, test_params: {light, nlanes},
           monitors: {c}, iterations: 100)

```

An instantiation of the reactive objects in the test configuration gives a *scene*—all the visual elements present in the simulated world. A *test case* provides concrete inputs to each “open” input stream in a scene. A test case determines how the scene evolves over time: how the cars and pedestrians move and how environment conditions change. We go through each part of the test configuration in detail below.

Reactive Objects. The core abstraction of PARACOSM is a *reactive object*. Reactive objects capture geometric and graphical features of a physical object, as well as their behavior over time. The behavioral interface for each reactive object has a set of *input* streams and a set of *output* streams. The evolution of the world is computed in steps of fixed duration which corresponds to events in a predefined *tick* stream. For streams that correspond to physical quantities updated by the physics simulator, such as position and speeds of cars, etc., appropriate events are generated by the underlying physics simulator.

Input streams provide input values from the environment over time; output streams represent output values computed by the object. The object’s constructor sets up the internal state of the object. An object is updated by event triggered computations. PARACOSM provides a set of assets as base classes. Autonomous driving systems naturally fit reactive programming models. They consume sensor input streams and produce actuator streams for the vehicle model. We differentiate between static *environment* reactive objects (subclassing

Fig. 2: Reactive streams represented by a marble diagram. A change in the value of test parameters `nlanes` or `light` changes the environment, and triggers a change in the corresponding sensor (output) stream `camera`.

Geometric) and dynamic *actor* reactive objects (subclassing **Physical**). Environment reactive objects represent “static” components of the world, such as road segments, intersections, buildings or trees, and a special component called the *world*. Actor reactive objects represent components with “dynamic” behavior: vehicles or pedestrians. The world object is used to model features of the world such as lighting or weather conditions. Reactive objects can be *composed* to generate complex assemblies from simple objects. The composition process can be used to connect static components structurally—such as two road segments connecting at an intersection. Composition also connects the behavior of an object to another by binding output streams to input streams. At run time, the values on that input stream of the second object are obtained from the output values of the first. Composition must respect geometric properties—the runtime system ensures that a composition maintains invariants such as no intersection of geometric components. We now describe the main features in PARACOSM, centered around the test configuration above.

Test Parameters. Using test variables, we can have general, but constrained streams of values passed into objects [59]. Our automatic test generator can then pick values for these variables, thereby leading to different test cases (see Figure 2). There are two types of parameters: continuous (**VarInterval**) and discrete (**VarEnum**). In the example presented, `light` (light intensity) is a continuous test parameter and `nlanes` (number of lanes) is discrete.

World. The **World** is a pre-defined reactive object in PARACOSM with a visual representation responsible for atmospheric conditions like the light intensity, direction and color, fog density, etc. The code segment

```
w = World(light:light, fog:0)
```

parameterizes the world using a test variable for light and sets the fog density to a constant (0).

Road Segments. In our example, **StraightRoadSegment** was parameterized with the number of lanes. In general, PARACOSM provides the ability to build complex road networks by connecting primitives of individual road segments and intersections. (A detailed example is presented in our Technical Report [43].)

It may seem surprising that we model static scene components such as roads as reactive objects. This serves two purposes. First, we can treat the number of lanes in a road segment as a constant input stream that is set by the test case, allowing parameterized test cases. Second, certain features of static objects can also change over time. For example, the coefficient of friction on a road segment may depend on the weather condition, which can be a function of time.

Autonomous Vehicles & System Under Test (SUT). `AutonomousVehicle`, as well as other actors, extends the `Physical` class (which in turn subclasses `Geometric`). This means that these objects have a visual as well as a physical model. The visual model is essentially a textured 3D mesh. The physical model contains properties such as mass, moments of inertia of separate bodies in the vehicle, joints, etc. This is used by the physics simulator to compute the vehicle's motion in response to external forces and control input. In the following code segment, we instantiate and place our test vehicle on the road:

```
v = AutonomousVehicle(start:r.onLane(1, 0.1), model:
    CarAsset(...), controller:MyController(...))
```

The `start` parameter “places” the vehicle in the world (in relative coordinates). The `model` parameter provides the implementation of the geometric and physical model of the vehicle. The `controller` parameter implements the autonomous controller under test. The internals of the controller implementation are not important; what is important is its interface (sensor inputs and the actuator outputs). These determine the input and output streams that are passed to the controller during simulation. For example, a typical controller can take sensor streams such as image streams from a camera as input and produce throttle and steering angles as outputs. The PARACOSM framework “wires” these streams appropriately. For example, the rendering engine determines the camera images based on the geometry of the scene and the position of the camera and the controller outputs are fed to the physics engine to determine the updated scene. Though simpler systems like OPENPILOT [15] use only a dashboard-mounted camera, autonomous vehicles can, in general, mix cameras at various mount points, LiDARs, radars, and GPS. PARACOSM can emulate many common types of sensors which produce streams of data. It is also possible to integrate new sensors, which are not supported out-of-the-box, by implementing them using the game engine's API.

Other Actors. A test often involves many actors such as pedestrians, and other (non-test) vehicles. Apart from the standard geometric (optionally physical) properties, these can also have some pre-programmed behavior. Behaviors can either be only dependent on the starting position (say, a car driving straight on the same lane), or be dynamic and reactive, depending on test parameters and behaviors of other actors. In general, the reactive nature of objects enables complex scenarios to be built. For example, here, we specify a simple behavior of a pedestrian crossing a road. The pedestrian starts crossing the road when a car is a certain distance away. In the code segments below, we use ‘_’ as shorthand for a lambda expression, i.e., “`f(_)`” is the same as “`x => f(x)`”.

```

Pedestrian(value start, value target, carPos, value dist,
  value speed) extends Geometric {
  ... // Initialization
  // Generate an event when the car gets close
  trigger = carPos.Filter( abs(_ - start) < dist )
  // target location reached
  done = pos.Filter( _ == target )
  // Walk to the target after trigger fires
  tick.SkipUntil(trigger).TakeUntil(done).foreach( ... /*
    walk with given speed */ )
}

```

Monitors and Test Oracles. PARACOSM provides an API to provide qualitative and quantitative temporal specifications. For instance, in the following example, we check that there is no collision and ensure that the collision was not trivially avoided because our vehicle did not move at all.

```

// no collision
CollisionMonitor(AutonomousVehicle v) extends Monitor {
  assert(v.collider.IsEmpty()) }
// cannot trivially pass the test by staying put
DistanceMonitor(AutonomousVehicle v, value minD) extends
  Monitor {
    pOld = v.pos.Take(1).Concat(v.pos)
    D = v.pos.Zip(pOld).Map( abs(_ - _) ).Sum()
    assert(D >= minD)
  }

```

The ability to write monitors which read streams of system-generated events provides an expressive framework to write temporal properties, something that has been identified as a major limitation of prior tools [31]. Monitors for metric and signal temporal logic specifications can be encoded in the usual way [18,33].

3 Systematic Testing of Paracosm Worlds

3.1 Test Inputs and Coverage

Worlds in PARACOSM directly describe a parameterized family of tests. The testing framework allows users to specify various strategies to generate input streams for both, static, and dynamic reactive objects in the world.

Test Cases. A *test* of *duration* T executes a configuration of reactive objects by providing inputs to every open input stream in the configuration for T ticks. The inputs for each stream must satisfy `const` parameters and respect the range constraints from `VarInterval` and `VarEnum`. The runtime system manages the scheduling of inputs and pushing input streams to the reactive objects. Let In denote the set of all input streams, and $\text{In} = \text{In}_D \cup \text{In}_C$ denote the partition of In into *discrete* streams and *continuous* streams respectively. Discrete streams take

their value over a finite, discrete range; for example, the color of a car, the number of lanes on a road segment, or the position of the next pedestrian (left/right) are discrete streams. Continuous streams take their values in a continuous (bounded) interval. For example, the fog density or the speed of a vehicle are examples of continuous streams.

Coverage. In the setting of autonomous vehicle testing, one often wants to explore the state space of a parameterized world to check “how well” an autonomous vehicle works under various situations, both qualitatively and quantitatively. Thus, we now introduce a notion of coverage. Instead of structural coverage criteria such as line or branch coverage, our goal is to cover the parameter space. In the following, for simplicity of notation, we assume that all discrete streams take values from $\{0, 1\}$, and all continuous streams take values in the real interval $[0, 1]$. Any input stream over bounded intervals—discrete or continuous—can be encoded into such streams. For discrete streams, there are finitely many tests, since each co-ordinate is Boolean and there is a fixed number of co-ordinates. One can define the coverage as the fraction of the number of vectors tested to the total number of vectors. Unfortunately, the total number of vectors is very high: if each stream is constant, then there are already 2^n tests for n streams. Instead, we consider the notion of *k-wise testing* from combinatorial testing [38]. In *k-wise testing*, we fix a parameter k , and ask that every interaction between every k elements is tested. Let us be more precise. Suppose that a test vector has N co-ordinates, where each co-ordinate can get the value 0 or 1. A set of tests A is a *k-wise covering family* if for every subset $\{i_1, i_2, \dots, i_k\} \subseteq \{1, \dots, N\}$ of co-ordinates and every vector $v \in \{0, 1\}^k$, there is a test $t \in A$ whose restriction to the $i_1, \dots, i_k$ is precisely v .

For continuous streams, the situation is more complex: since any continuous interval has infinitely many points, each corresponding to a different test case, we cannot directly define coverage as a ratio (the denominator will be infinite). Instead, we define coverage using the notion of *dispersion* [49, 57]. Intuitively, dispersion measures the largest empty space left by a set of tests. We assume a (continuous) test is a vector in $[0, 1]^N$: each entry is picked from the interval $[0, 1]$ and there are N co-ordinates. Dispersion over $[0, 1]^N$ can be defined relative to sets of neighborhoods, such as N -dimensional balls or axis-parallel rectangles. Let us define $\mathcal{B}$ to be the family of N -dimensional axis-parallel rectangles in $[0, 1]^N$, our results also hold for other notions of neighborhoods such as balls or ellipsoids. For a neighborhood $B \in \mathcal{B}$, let $\text{vol}(B)$ denote the volume of B . Given a set $A \subseteq [0, 1]^N$ of tests, we define the *dispersion* as the largest volume neighborhood in $\mathcal{B}$ without any test:

$$\text{dispersion}(A) = \sup \{ \text{vol}(B) \mid B \in \mathcal{B} \text{ and } A \cap B = \emptyset \}$$

A lower dispersion means better coverage.

Let us summarize. Suppose that a test vector consists of N_D discrete co-ordinates and N_C continuous co-ordinates; that is, a test is a vector (t_D, t_C) in $\{0, 1\}^{N_D} \times [0, 1]^{N_C}$. We say a set of tests A is (k, ε) -covering if

1. for each set of k co-ordinates $\{i_1, \dots, i_k\} \subseteq \{1, \dots, N_D\}$ and each vector $v \in \{0, 1\}^k$, there is a test $(t_D, t_C) \in \{0, 1\}^{N_D} \times [0, 1]^{N_C}$ such that the restriction of t_D to the co-ordinates $i_1, \dots, i_k$ is v ; and
2. for each $(t_D, t_C) \in A$, the set $\{t_C \mid (t_D, t_C) \in A\}$ has dispersion at most ϵ .

3.2 Test Generation

The goal of our default test generator is to maximize (k, ϵ) for programmer-specified number of test iterations or `ticks`.

k-Wise Covering Family. One can use explicit construction results from combinatorial testing to generate k -wise covering families [14]. However, a simple way to generate such families with high probability is random testing. The proof is by the probabilistic method [4] (see also [44]). Let A be a set of $2^k(k \log N - \log \delta)$ uniformly randomly generated $\{0, 1\}^N$ vectors. Then A is a k -wise covering family with probability at least $1 - \delta$.

Low Dispersion Sequences. It is tempting to think that uniformly generating vectors from $[0, 1]^N$ would similarly give low dispersion sequences. Indeed, as the number of tests goes to infinity, the set of randomly generated tests has dispersion 0 almost surely. However, when we fix the number of tests, it is well known that uniform random sampling can lead to high dispersion [49, 57]; in fact, one can show that the dispersion of n uniformly randomly generated tests grows asymptotically as $O((\log \log n/n)^{\frac{1}{2}})$ almost surely. Our test generation strategy is based on *deterministic quasi-Monte Carlo sequences*, which have much better dispersion properties, asymptotically of the order of $O(1/n)$, than the dispersion behavior of uniformly random tests. There are many different algorithms for generating quasi-Monte Carlo sequences deterministically (see, e.g., [49, 57]). We use *Halton sequences*. For a given ϵ , we need to generate $O(\frac{1}{\epsilon})$ inputs via Halton sampling. In Section 4.2, we compare uniform random and Halton sampling.

Cost Functions and Local Search. In many situations, testers want to optimize parameter values for a specific function. A simple example of this is finding higher-speed collisions, which intuitively, can be found in the vicinity of test parameters that already result in high-speed collisions. Another, slightly different case is (greybox) fuzzing [5, 55], for example, finding new collisions using small mutations on parameter values that result in the vehicle narrowly avoiding a collision. Our test generator supports such *quantitative* objectives and *local search*. A quantitative monitor evaluates a cost function on a run of a test case. Our test generation tool generates an initial, randomly chosen, set of test inputs. Then, it considers the scores returned by the Monitor on these samples, and performs a local search on samples with the highest/lowest scores to find local optima of the cost function.

4 Implementation and Tests

4.1 Runtime System and Implementation

PARACOSM uses the Unity game engine [69] to render visuals, do runtime checks and simulate physics (via PhysX [16]). Reactive objects are built on top of UniRx [36], an implementation of the popular Reactive Extensions framework [56]. The game engine manages geometric transformations of 3D objects and offers easy to use abstractions for generating realistic simulations. Encoding behaviors and monitors, management of 3D geometry and dynamic checks are implemented using the game engine interface. The project code is available at: <https://gitlab.mpi-sws.org/mathur/paracosm>.

A simulation in PARACOSM proceeds as follows. A test configuration is specified as a subclass of the `EnvironmentProgramBaseClass`. Tests are run by invoking the `run_test` method, which receives as input the reactive objects that should be instantiated in the world as well as additional parameters relating to the test. The `run_test` method runs the tests by first initializing and placing the reactive objects in the scene using their 3D mesh (if they have one) and then invoking a reactive engine to start the simulation. The system under test is run in a separate process and connects to the simulation. The simulation then proceeds until the simulation completion criteria is met (a time-out or some monitor event).

Output to Standardized Testing Formats. There have been recent efforts to create standardized descriptions of tests in the automotive industry. The most relevant formats are OPENDRIVE [7] and OPENSCENARIO (only recently finalized) [8]. OPENDRIVE describes road structures, and OPENSCENARIO describes actors and their behavior. PARACOSM currently supports outputs to OPENDRIVE. Due to the static nature of the specification format, a different file is generated for each test iteration/configuration.

4.2 Evaluation

We evaluate PARACOSM with respect to the following research questions (**RQs**):

RQ 1: Does PARACOSM’s programmatic interface enable the easy design of test environments and worlds?

RQ 2: Do the test input generation strategies discussed in Section 3 effectively explore the parameter space?

RQ 3: Can PARACOSM help uncover poor performance or bad behavior of the SUT in common autonomous driving tasks?

Methodology. To answer **RQ 1**, we develop three independent environments rich with visual features and other actors, and use the variety generated with just a few lines of code as a proxy for ease of design. To answer **RQ 2**, we use coverage maximizing strategies for test inputs to all the three environments/case studies. We also use and evaluate cost functions and local search based methods. To answer **RQ 3**, we test various neural network based systems and demonstrate

Table 1: An overview of our case studies. Note that even though the Adaptive Cruise Control study has 2 discrete parameters, we calculate k-wise coverage for 3 as the 2 parameters require 3 bits for representation.

	Road segmentation	Jaywalking pedestrian	Adaptive Cruise Control
SUT	VGGNet CNN [62]	NVIDIA CNN [12]	NVIDIA CNN [12]
Training	191 images	403 image & car control samples	1034 image & car control samples
Test	3 discrete	2 continuous	3 continuous & 2 discrete
Test iters	100	100, 15s timeout	100, 15s timeout
Monitor	Ground truth	Scored Collision	Collision & Distance
Coverage	$k = 3$ with probability ~ 1	$\epsilon = 0.041$	$\epsilon = 0.043$, $k = 3$ with probability ~ 1

(a) A good test with all parameter values same as the training set (true positive: 89%, false positive: 0%).

(b) A bad test with all parameter values different from the training set (true positive: 9%, false positive: 1%).

Fig. 3: Example results from the road segmentation case study. Pixels with a green mask are segmented by the SUT as a road.

how PARACOSM can help uncover problematic scenarios. A summary of the case studies presented here is available in Table 1. In our Technical Report [43], we present more case studies, specifically experiments on many pre-trained neural networks, busy urban environments and studies exploiting specific testing features of PARACOSM.

4.3 Case Studies

Road segmentation Using PARACOSM’s programmatic interface, we design a long road segment with several vehicles. The vehicular behavior is to drive on their respective lanes with a fixed maximum velocity. The test parameters are the number of lanes ($\{2, 4\}$), number of cars in the environment ($\{0, 5\}$) and light conditions ($\{Noon, Evening\}$). Noon lighting is much brighter than the evening. The direction of lighting is also the opposite. We test a deep CNN called VGGNet [62], that is known to perform well on several image segmentation benchmarks. The task is road segmentation, i.e., given a camera image, identifying which pixels correspond to the road. The network is trained on 191 dashcam images

Table 2: Summary of results of the road segmentation case study. Each combination of parameter values is presented separately, with the parameter values used for training in bold. We report the SUT’s average true positive rate (% of pixels corresponding to the road that are correctly classified) and false positive rate (% of pixels that are not road, but incorrectly classified as road).

# lanes	# cars	Lighting	# test iters	True positive (%)	False positive (%)
2	5	Noon	12	70%	5.1%
2	5	Evening	14	53.4%	22.4%
2	0	Evening	12	51.4%	18.9%
2	0	Noon	12	71.3%	6%
4	5	Evening	10	60.4%	7.1%
4	5	Noon	16	68.5%	20.2%
4	0	Evening	13	51.5%	7.1%
4	0	Noon	11	83.3%	21%

Table 3: Results for the jaywalking pedestrian case study.

Testing strategy	Dispersion (ϵ)	% fail	Max. collision
Random	0.092	7%	10.5 m/s
Halton	0.041	10%	11.3 m/s
Random+opt/collision	0.109	13%	11.1 m/s
Halton+opt/collision	0.043	20%	11.9 m/s
Random+opt/almost failing	0.126	13%	10.5 m/s
Halton+opt/almost failing	0.043	13%	11.4 m/s

captured in the test environment with fixed parameters (2 lanes, 5 cars, and *Noon* lighting), recorded at the rate of one image every $1/10^{th}$ second, while manually driving the vehicle around (using a keyboard). We test on 100 images generated using PARACOSM’s default test generation strategy (uniform random sampling for discrete parameters). Table 2 summarizes the test results. Tests with parameter values far away from the training set are observed to not perform so well. As depicted in Figure 3, this happens because varying test parameters can drastically change the scene.

Jaywalking pedestrian. We now test over the environment presented in Section 2. The environment consists of a straight road segment and a pedestrian. The pedestrian’s behavior is to cross the road at a specific walking speed when the autonomous vehicle is a specific distance away. The walking speed of the pedestrian and the distance of the autonomous vehicle when the pedestrian starts crossing the road are test parameters. The SUT is a CNN based on NVIDIA’s behavioral cloning framework [12]. It takes camera images as input, and produces the relevant steering angle or throttle control as output. The SUT is trained on 403 samples obtained by driving the vehicle manually and recording the camera and corresponding control data. The training environment has pedestrians crossing

the road at various time delays, but always at a fixed walking speed (1 m/s). In order to evaluate **RQ 2** completely, we evaluate the default coverage maximizing sampling approach, as well as explore two quantitative objectives: first, maximizing the collision speed, and second, finding new failing cases around samples that *almost* fail. For the default approach, the `CollisionMonitor` as presented in Section 2 is used. For the first quantitative objective, this `CollisionMonitor`'s code is prepended with the following calculation:

```
// Score is speed of car at time of collision
coll_speed = v.speed.CombineLatest(v.collider, (s,c) => s)
    .First()
```

The score `coll_speed` is used by the test generator for optimization. For the second quantitative objective, the `CollisionMonitor` is modified to give high scores to tests where the distance between the autonomous vehicle and pedestrian is very small:

```
CollisionMonitor(AutonomousVehicle v, Pedestrian p)
    extends Monitor {
    minDist = v.pos.Zip(p.pos).Map(1/abs(_-_)).Min()
    coll_score = v.collider.Map(0)
    // Score is either 0 (collision) or 1/minDist
    score = coll_score.DefaultIfEmpty(minDist)
    assert(v.collider.IsEmpty())
}
```

We evaluate the following test input generation strategies: (i) Random sampling (ii) Halton sampling, (iii) Random or Halton sampling with local search for the two quantitative objectives. We run 100 iterations of each strategy with a 15 second timeout. For random or Halton sampling, we sample 100 times. For the quantitative objectives, we first generate 85 random or Halton samples, then choose the top 5 scores, and finally run 3 simulated annealing iterations on each of these 5 configurations. Table 3 presents results from the various test input generation strategies. Clearly, Halton sampling offers the lowest dispersion (highest coverage) over the parameter space. This can also be visually confirmed from the plot of test parameters (Figure 4). There are no big gaps in the parameter space. Moreover, we find that test strategies optimizing for the first objective are successful in finding more collisions with higher speeds. As these techniques perform simulated annealing repetitions on top of already failing tests, they also find more failing tests overall. Finally, test strategies using the second objective are also successful in finding more (newer) failure cases than simple Random or Halton sampling.

Adaptive Cruise Control. We now create and test in an environment with our test vehicle following a car (lead car) on the same lane. The lead car's behavior is programmed to drive on the same lane as the test vehicle, with a certain maximum speed. This is a very typical driving scenario that engineers test their implementations on. We use 5 test parameters: the initial lead of the lead car to

Fig. 4: A comparison of the various test generation strategies for the jaywalking pedestrian case study. The X-axis is the walking speed of the pedestrian (2 to 10 m/s). The Y-axis is the distance from the car when the pedestrian starts crossing (30 to 60 m). Passing tests are labelled with a green dot. Failing tests (tests with a collision) are marked with a red cross.

the test vehicle ($[8m, 40m]$), the lead car’s maximum speed ($[3m/s, 8m/s]$), density of fog³ in the environment ($[0, 1]$), number of lanes on the road ($\{2, 4\}$), and color of the lead car ($\{Black, Red, Yello, Blue\}$). We use both, *CollisionMonitor*⁴ and *DistanceMonitor*, as presented in Section 2. A test *passes* if there is no collision and the autonomous vehicle moves atleast 5 m during the simulation duration (15 s).

We use PARACOSM’s default test generation strategy, i.e., Halton sampling for continuous parameters and Random sampling for discrete parameters (no optimization or fuzzing). The SUT is the same CNN as in the previous case study. It is trained on 1034 training samples, which are obtained by manually driving behind a red lead car on the same lane of a 2-lane road with the same maximum velocity (5.5 m/s) and no fog.

The results of this case study are presented in Table 4. Looking at the discrete parameters, the number of lanes does not seem to contribute towards a risk of collision. Surprisingly, though the training only involves a Red lead car, the results appear to be the best for a Blue lead car. Moving on to the continuous

³ 0 denotes no fog and 1 denotes very dense fog (exponential squared scale).

⁴ the monitor additionally calculates the mean distance of the test vehicle to the lead car during the test, which is used for later analysis.

Fig. 5: Continuous test parameters of the Adaptive Cruise Control study plotted against each other: the initial offset of the lead car (8 to 40 m), the lead car’s maximum speed (3 to 8 m/s) and the fog density (0 to 1). Green dots, red crosses, and blue triangles denote passing tests, collisions, and inactivity respectively.

Table 4: Parameterized test on Adaptive Cruise Control, separated for each value of discrete parameters, and low and high values of continuous parameters. A test *passes* if there are no collisions and no inactivity (the overall distance moved by the test vehicle is more than 5 m. The average offset (in m) maintained by the test vehicle to the lead car (for passing tests) is also presented.

	Discrete parameters						Continuous parameters					
	Num. lanes		Lead car color				Initial offset (m)		Speed (m/s)		Fog density	
	2	4	Black	Red	Yellow	Blue	< 24	≥ 24	< 5.5	≥ 5.5	< 0.5	≥ 0.5
Test iters	54	46	24	22	27	27	51	49	52	48	51	49
Collisions	7	7	3	3	6	2	6	8	8	6	12	0
Inactivity	12	4	4	4	6	2	9	7	9	7	1	15
Offset (m)	42.4	43.4	46.5	48.1	39.6	39.1	33.7	52.7	38.4	47.4	36.5	49.8

parameters, the fog density appears to have the most significant impact on test failures (collision or vehicle inactivity). In the presence of dense fog, the SUT behaves pessimistically and does not accelerate much (thereby causing a failure due to inactivity). These are all interesting and useful metrics about the performance of our SUT. Plots of the results projected on to continuous parameters are presented in Figure 5.

4.4 Results and Analysis

We now summarize the results of our evaluation with respect to our **RQs**:

RQ 1: All the three case studies involve varied, rich and dynamic environments. They are representative of tests engineers would typically want to do, and we parameterize many different aspects of the world and the dynamic behavior of its components. These designs are at most 70 lines of code. This provides confidence in PARACOSM’s ability of providing an easy interface for the design of realistic test environments.

RQ 2: Our default test generation strategies are found to be quite effective at exploring the parameter space systematically, eliminating large unexplored gaps,

and at the same time, successfully identifying problematic cases in all the three case studies. The jaywalking pedestrian study demonstrates that optimization and local search are possible on top of these strategies, and are quite effective in finding the relevant scenarios. The adaptive cruise control study tests over 5 parameters, which is more than most related works, and even guarantees good coverage of this parameter space. Therefore, it is amply clear that PARACOSM’s test input generation methods are useful and effective.

RQ 3: The road segmentation case study uses a well-performing neural network for object segmentation, and we are able to detect degraded performance for automatically generated test inputs. Whereas this study focuses on static image classification, the next two, i.e., the jaywalking pedestrian and the adaptive cruise control study uncover poor performance on simulated driving, using a popular neural network architecture for self driving cars. Therefore, we can safely conclude that PARACOSM can find bugs in various different kinds of systems related to autonomous driving.

4.5 Threats to Validity

The *internal validity* of our experiments depends on having implemented our system correctly and, more importantly, trained and used the neural networks considered in the case studies correctly. For training the networks, we followed the available documentation and inspected our examples to ensure that we use an appropriate training procedure. We watched some test runs and replays of tests we did not understand. Furthermore, our implementation logs events and we also capture images, which allow us to check a large number of tests.

In terms of threats to external validity, the biggest challenge in this project has been finding systems that we can easily train and test in complex driving scenarios. Publicly available systems have limited capabilities and tend to be brittle. Many networks trained on real world data do not work well in simulation. We therefore re-train these networks in simulation. An alternative is to run fewer tests, but use more expensive and visually realistic simulations. Our test generation strategy maximizes coverage, even when only a few test iterations can be performed due to high simulation cost.

5 Related Work

Traditionally, test-driven software development paradigms [9] have advocated testing and mocking frameworks to test software early and often. Mocking frameworks and mock objects [42, 47] allow programmers to test a piece of code against an API specification. Typically, mock objects are stubs providing outputs to explicitly provided lists of inputs of simple types, with little functionality of the actual code. Thus, they fall short of providing a rich environment for autonomous driving. PARACOSM can be seen as a mocking framework for reactive, physical systems embedded in the 3D world. Our notion of constraining streams is inspired by work on declarative mocking [59].

Testing Cyber-Physical Systems. There is a large body of work on automated test generation tools for cyber-physical systems through heuristic search of a high-dimensional continuous state space. While much of this work has focused on low-level controller interfaces [6, 17, 19, 20, 25, 60] rather than the system level, specification and test generation techniques arising from this work—for example, the use of metric and signal temporal logics or search heuristics—can be adapted to our setting. More recently, test generation tools have started targeting autonomous systems under a simulation-based semantic testing framework similar to ours. In most of these works, visual scenarios are either fixed by hand [1, 2, 10, 22, 27, 29, 66, 67], or are constrained due to the model or coverage criteria [3, 45, 50]. These analyses are shown to be preferable to the application of random noise on the input vector. Additionally, a simulation-based approach filters benign misclassifications from misclassifications that actually lead to bad or dangerous behavior. Our work extends this line of work and provides an expressive language to design parameterized environments and tests. ASFAULT [29] uses random search and mutation for procedural generation of road networks for testing. AC3R [28] reconstructs test cases from accident reports.

To address problems of high time and infrastructure cost of testing autonomous systems, several simulators have been developed. The most popular is Gazebo [26] for the ROS [54] robotics framework. It offers a modular and extensible architecture, however falls behind on visual realism and complexity of environments that can be generated with it. To counter this, game engines are used. Popular examples are TORCS [72], CARLA [21], and AirSim [61]. Modern game engines support creation of realistic urban environments. Though they enable visually realistic simulations, and enable detection of infractions such as collisions, the environments themselves are difficult to design. Designing a custom environment involves manual placement of road segments, buildings, and actors (as well as their properties). Performing many systematic tests is therefore time-consuming and difficult. While these systems and PARACOSM share the same aims and much of the same infrastructure, PARACOSM focuses on procedural design and systematic testing, backed by a relevant coverage criteria.

Adversarial Testing. Adversarial examples for neural networks [32, 64] introduce perturbations to inputs that cause a classifier to classify “perceptually identical” inputs differently. Much work has focused on finding adversarial examples in the context of autonomous driving as well as on training a network to be robust to perturbations [11, 30, 46, 51, 71]. Tools such as DEEPPXLORE [52], DEEPTTEST [65], DEEPGAUGE [41], and SADL [37] define a notion of coverage for neural networks based on the number of neurons activated during tests compared against the total number of neurons in the network and activation during training. However, these techniques focus mostly on individual classification tasks and apply 2D transformations on images. In comparison, we consider the closed-loop behavior of the system and our parameters directly change the world rather than apply transformations post facto. We can observe, over time, that certain vehicles are not detected, which is more useful to testers than a single misclassification [31]. Furthermore, it is already known that structural coverage criteria may not be an

effective strategy for finding errors in classification [39]. We use coverage metrics on the test space, rather than the structure of the neural network. Alternately, there are recent techniques to verify controllers implemented as neural networks through constraint solving or abstract interpretation [24, 30, 34, 58, 71]. While these tools do not focus on the problem of autonomous driving, their underlying techniques can be combined in the test generation phase for PARACOSM.

6 Future Work and Conclusion

Deploying autonomous systems like self-driving cars in urban environments raises several safety challenges. The complex software stack processes sensor data, builds a semantic model of the surrounding world, makes decisions, plans trajectories, and controls the car. The end-to-end testing of such systems requires the creation and simulation of whole worlds, with different tests representing different world and parameter configurations. PARACOSM tackles these problems by (i) enabling procedural construction of diverse scenarios, with precise control over elements like road layout, physical and visual properties of objects, and behaviors of actors in the system, and (ii) using quasi-random testing to obtain good coverage over large parameter spaces.

In our evaluation, we show that PARACOSM enables easy design of environments and automated testing of autonomous agents implemented using neural networks. While finding errors in sensing can be done with only a few static images, we show that PARACOSM also enables the creation of longer test scenarios which exercise the controller’s feedback on the environment. Our case studies focused on *qualitative* state space exploration. In future work, we shall perform *quantitative* statistical analysis to understand the sensitivity of autonomous vehicle behavior on individual parameters.

In the future, we plan to extend PARACOSM’s testing infrastructure to also aid in the training of deep neural networks that require large amounts of high quality training data. For instance, we show that small variations in the environment result in widely different results for road segmentation. Generating data is a time consuming and expensive task. PARACOSM can easily generate labelled data for static images. For driving scenarios, we can record a user manually driving in a parameterized PARACOSM environment and augment this data by varying parameters that should not impact the car’s behavior. For instance, we can vary the color of other cars, positions of pedestrians who are not crossing, or even the light conditions and sensor properties (within reasonable limits).

Acknowledgements This research was funded in part by the Deutsche Forschungsgemeinschaft project 389792660-TRR 248 and by the European Research Council under the Grant Agreement 610150 (ERC Synergy Grant ImPACT).

References

1. Abbas, H., O’Kelly, M., Rodionova, A., Mangharam, R.: Safe at any speed: A simulation-based test harness for autonomous vehicles. In: 7th Workshop on Design, Modeling and Evaluation of Cyber Physical Systems (CyPhy17) (October 2017)
2. Abdessalem, R.B., Nejati, S., Briand, L.C., Stifter, T.: Testing vision-based control systems using learnable evolutionary algorithms. In: Proceedings of the 40th International Conference on Software Engineering. p. 1016–1026. ICSE ’18, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3180155.3180160>, <https://doi.org/10.1145/3180155.3180160>
3. Alexander, R., Hawkins, H., Rae, A.: Situation coverage – a coverage criterion for testing autonomous robots (02 2015)
4. Alon, N., Spencer, J.H.: The Probabilistic Method. Wiley-Interscience series in discrete mathematics and optimization, Wiley (2004)
5. American Fuzzy Loop: Technical “whitepaper” for afl-fuzz, http://lcamtuf.coredump.cx/afl/technical_details.txt, accessed: 2019-08-23
6. Annpureddy, Y., Liu, C., Fainekos, G.E., Sankaranarayanan, S.: S-TaLiRo: A tool for temporal logic falsification for hybrid systems. In: TACAS 11. Lecture Notes in Computer Science, vol. 6605, pp. 254–257. Springer (2011)
7. Association for Advancement of international Standardization of Automation and Measuring Systems (ASAM): Opendrive (2018), <http://www.opendrive.org/index.html>, accessed: 2019-08-21
8. Association for Advancement of international Standardization of Automation and Measuring Systems (ASAM): Openscenario (2018), <http://www.opendrive.org/index.html>, accessed: 2019-08-21
9. Beck, K.L.: Test Driven Development: By Example. Addison-Wesley Professional (2002)
10. Ben Abdessalem, R., Nejati, S., Briand, L.C., Stifter, T.: Testing advanced driver assistance systems using multi-objective search and neural networks. In: Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering (ASE). pp. 63–74 (2016)
11. Bhagoji, A.N., He, W., Li, B., Song, D.: Exploring the space of black-box attacks on deep neural networks. CoRR **abs/1712.09491** (2017), <http://arxiv.org/abs/1712.09491>
12. Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L.D., Monfort, M., Muller, U., Zhang, J., et al.: End to end learning for self-driving cars. arXiv preprint arXiv:1604.07316 (2016)
13. Caspi, P., Pilaud, D., Halbwachs, N., Plaice, J.: Lustre: A declarative language for programming synchronous systems. In: Conference Record of the Fourteenth Annual ACM Symposium on Principles of Programming Languages, Munich, Germany, January 21–23, 1987. pp. 178–188 (1987)
14. Colbourn, C.J.: Combinatorial aspects of covering arrays. *Le Matematiche* **59**(1,2), 125–172 (2004), <https://lematematiche.dmi.unict.it/index.php/lematematiche/article/view/166>
15. comma.ai: openpilot: open source driving agent (2016), <https://github.com/commaai/openpilot>, accessed: 2018-11-13
16. Coporation, N.: Physx (2008), <https://developer.nvidia.com/gameworks-physx-overview>, accessed: 2018-11-13

17. Deshmukh, J., Jin, X., Kapinski, J., Maler, O.: Stochastic local search for falsification of hybrid systems. In: ATVA. pp. 500–517. Springer (2015)
18. Deshmukh, J.V., Donzé, A., Ghosh, S., Jin, X., Juniwal, G., Seshia, S.A.: Robust online monitoring of signal temporal logic. *Formal Methods in System Design* **51**(1), 5–30 (2017). <https://doi.org/10.1007/s10703-017-0286-7>, <https://doi.org/10.1007/s10703-017-0286-7>
19. Deshmukh, J.V., Horvat, M., Jin, X., Majumdar, R., Prabhu, V.S.: Testing cyber-physical systems through bayesian optimization. *ACM Trans. Embedded Comput. Syst.* **16**(5), 170:1–170:18 (2017). <https://doi.org/10.1145/3126521>, <https://doi.org/10.1145/3126521>
20. Donzé, A.: *Breach, A Toolbox for Verification and Parameter Synthesis of Hybrid Systems*, pp. 167–170. Springer (2010)
21. Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., Koltun, V.: CARLA: An open urban driving simulator. In: *Proceedings of the 1st Annual Conference on Robot Learning*. pp. 1–16 (2017)
22. Dreossi, T., Donzé, A., Seshia, S.A.: Compositional falsification of cyber-physical systems with machine learning components. In: *NASA Formal Methods - 9th International Symposium, NFM 2017. Lecture Notes in Computer Science*, vol. 10227, pp. 357–372. Springer (2017)
23. Dreossi, T., Jha, S., Seshia, S.A.: Semantic adversarial deep learning **10981**, 3–26 (2018). https://doi.org/10.1007/978-3-319-96145-3_1, https://doi.org/10.1007/978-3-319-96145-3_1
24. Dutta, S., Chen, X., Jha, S., Sankaranarayanan, S., Tiwari, A.: Sherlock - A tool for verification of neural network feedback systems: demo abstract. In: Ozay, N., Prabhakar, P. (eds.) *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control, HSCC 2019, Montreal, QC, Canada, April 16-18, 2019*. pp. 262–263. ACM (2019). <https://doi.org/10.1145/3302504.3313351>, <https://doi.org/10.1145/3302504.3313351>
25. Fainekos, G.: Automotive control design bug-finding with the S-TaLiRo tool. In: *ACC 2015*. p. 4096 (2015)
26. Foundation, O.S.R.: Vehicle simulation in gazebo, <http://gazebo-sim.org/blog/vehicle%20simulation>, accessed: 2019-08-23
27. Fremont, D.J., Dreossi, T., Ghosh, S., Yue, X., Sangiovanni-Vincentelli, A.L., Seshia, S.A.: Scenic: A language for scenario specification and scene generation. In: *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*. pp. 63–78. PLDI 2019, ACM, New York, NY, USA (2019). <https://doi.org/10.1145/3314221.3314633>, <http://doi.acm.org/10.1145/3314221.3314633>
28. Gambi, A., Huynh, T., Fraser, G.: Generating effective test cases for self-driving cars from police reports. In: Dumas, M., Pfahl, D., Apel, S., Russo, A. (eds.) *Proceedings of the ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/SIGSOFT FSE 2019, Tallinn, Estonia, August 26-30, 2019*. pp. 257–267. ACM (2019). <https://doi.org/10.1145/3338906.3338942>, <https://doi.org/10.1145/3338906.3338942>
29. Gambi, A., Müller, M., Fraser, G.: Automatically testing self-driving cars with search-based procedural content generation. In: Zhang, D., Møller, A. (eds.) *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2019, Beijing, China, July 15-19, 2019*. pp. 318–328. ACM (2019). <https://doi.org/10.1145/3293882.3330566>, <https://doi.org/10.1145/3293882.3330566>

30. Gehr, T., Mirman, M., Drachler-Cohen, D., Tsankov, P., Chaudhuri, S., Vechev, M.T.: AI2: safety and robustness certification of neural networks with abstract interpretation. In: 2018 IEEE Symposium on Security and Privacy, S&P 2018. pp. 3–18. IEEE (2018)
31. Gladisch, C., Heinz, T., Heinzemann, C., Oehlerking, J., von Vietinghoff, A., Pfitzer, T.: Experience paper: Search-based testing in automated driving control applications. In: Proceedings of the 34th IEEE/ACM International Conference on Automated Software Engineering (ASE). pp. 26–37 (2019)
32. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. CoRR **abs/1412.6572** (2014), <http://arxiv.org/abs/1412.6572>
33. Ho, H., Ouaknine, J., Worrell, J.: Online monitoring of metric temporal logic. In: Runtime Verification RV 2014. Lecture Notes in Computer Science, vol. 8734, pp. 178–192. Springer (2014)
34. Huang, X., Kwiatkowska, M., Wang, S., Wu, M.: Safety verification of deep neural networks. In: Majumdar, R., Kuncak, V. (eds.) Computer Aided Verification - 29th International Conference, CAV 2017, Heidelberg, Germany, July 24–28, 2017, Proceedings, Part I. Lecture Notes in Computer Science, vol. 10426, pp. 3–29. Springer (2017). https://doi.org/10.1007/978-3-319-63387-9_1, https://doi.org/10.1007/978-3-319-63387-9_1
35. Hudak, P., Courtney, A., Nilsson, H., Peterson, J.: Arrows, robots, and functional reactive programming. In: Advanced Functional Programming, 4th International School, AFP 2002, Oxford, UK, August 19–24, 2002, Revised Lectures. Lecture Notes in Computer Science, vol. 2638, pp. 159–187. Springer (2002). https://doi.org/10.1007/978-3-540-44833-4_6, https://doi.org/10.1007/978-3-540-44833-4_6
36. Kawai, Y.: Unixr: Reactive extensions for unity (2014), <https://github.com/neuecc/UniRx>, accessed: 2018-11-13
37. Kim, J., Feldt, R., Yoo, S.: Guiding deep learning system testing using surprise adequacy. In: Proceedings of the 41st International Conference on Software Engineering. pp. 1039–1049. ICSE '19, IEEE Press, Piscataway, NJ, USA (2019). <https://doi.org/10.1109/ICSE.2019.00108>, <https://doi.org/10.1109/ICSE.2019.00108>
38. Kuhn, D.R., Kacker, R.N., Lei, Y.: Combinatorial testing. In: Laplante, P.A. (ed.) Encyclopedia of Software Engineering, pp. 1–12. CRC Press (Nov 2010)
39. Li, Z., Ma, X., Xu, C., Cao, C.: Structural coverage criteria for neural networks could be misleading. In: Sarma, A., Murta, L. (eds.) Proceedings of the 41st International Conference on Software Engineering: New Ideas and Emerging Results, ICSE (NIER) 2019, Montreal, QC, Canada, May 29–31, 2019. pp. 89–92. IEEE / ACM (2019), <https://dl.acm.org/citation.cfm?id=3339171>
40. Liberty, J., Betts, P.: Programming Reactive Extensions and LINQ. Apress (2011)
41. Ma, L., Juefei-Xu, F., Zhang, F., Sun, J., Xue, M., Li, B., Chen, C., Su, T., Li, L., Liu, Y., Zhao, J., Wang, Y.: Deepgauge: Multi-granularity testing criteria for deep learning systems. In: Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering. pp. 120–131. ASE 2018, ACM, New York, NY, USA (2018). <https://doi.org/10.1145/3238147.3238202>, <http://doi.acm.org/10.1145/3238147.3238202>
42. Mackinnon, T., Freeman, S., Craig, P.: Endo-testing: Unit testing with mock objects. In: eXtreme Programming and Flexible Processes in Software Engineering - XP2000 (2000)

43. Majumdar, R., Mathur, A.S., Pirron, M., Stegner, L., Zufferey, D.: Paracosm: A language and tool for testing autonomous driving systems. CoRR **abs/1902.01084** (2019), <http://arxiv.org/abs/1902.01084>
44. Majumdar, R., Niksic, F.: Why is random testing effective for partition tolerance bugs? PACMPL **2**(POPL), 46:1–46:24 (2018)
45. Majzik, I., Semeráth, O., Hajdu, C., Marussy, K., Szatmári, Z., Micskei, Z., Vörös, A., Babikian, A.A., Varró, D.: Towards system-level testing with coverage guarantees for autonomous vehicles. In: 2019 ACM/IEEE 22nd International Conference on Model Driven Engineering Languages and Systems (MODELS). pp. 89–94 (2019). <https://doi.org/10.1109/MODELS.2019.00-12>
46. Mirman, M., Gehr, T., Vechev, M.: Differentiable abstract interpretation for provably robust neural networks. In: International Conference on Machine Learning (ICML) (2018), <https://www.icml.cc/Conferences/2018/Schedule?showEvent=2477>
47. Mockito: Tasty mocking framework for unit tests in java, <http://site.mockito.org>, accessed: 2019-08-23
48. National Transportation Safety Board: Collision between vehicle controlled by developmental automated driving system and pedestrian, tempe, arizona, march 18, 2018. Highway Accident Report NTSB/HAR-19/03, National Transportation Safety Board (November 2019)
49. Niederreiter, H.: Random number generation and quasi-Monte Carlo methods. SIAM (1992)
50. O’Kelly, M., Sinha, A., Namkoong, H., Tedrake, R., Duchi, J.C.: Scalable end-to-end autonomous vehicle testing via rare-event simulation. Advances in Neural Information Processing Systems **31**, 9827–9838 (2018)
51. Papernot, N., McDaniel, P., Goodfellow, I., Jha, S., Celik, Z.B., Swami, A.: Practical black-box attacks against machine learning. In: Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security - ASIA CCS 17. ACM (2017). <https://doi.org/10.1145/3052973.3053009>, <https://doi.org/10.1145/3052973.3053009>
52. 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, Shanghai, China, October 28-31, 2017. pp. 1–18. ACM (2017). <https://doi.org/10.1145/3132747.3132785>, <https://doi.org/10.1145/3132747.3132785>
53. Pomerleau, D.: ALVINN: An autonomous land vehicle in a neural network. In: NIPS 88: Neural Information Processing Systems (1988)
54. Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., Wheeler, R., Ng, A.: Ros: an open-source robot operating system. In: ICRA workshop on open source software (2009)
55. Rawat, S., Jain, V., Kumar, A.J.S., Cojocar, L., Giuffrida, C., Bos, H.: Vuzzer: Application-aware evolutionary fuzzing. In: NDSS (2017)
56. ReactiveX: Reactivex, <http://reactivex.io/>, accessed: 2019-08-23
57. Rote, G., Tichy, R.: Quasi-Monte-Carlo methods and the dispersion of point sequences. Mathematical and Computer Modelling **23**, 9–23 (1996)
58. Ruan, W., Huang, X., Kwiatkowska, M.: Reachability analysis of deep neural networks with provable guarantees. In: Lang, J. (ed.) Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden. pp. 2651–2659. ijcai.org (2018). <https://doi.org/10.24963/ijcai.2018/368>, <https://doi.org/10.24963/ijcai.2018/368>

59. Samimi, H., Hicks, R., Fogel, A., Millstein, T.: Declarative mocking. In: ISSTA 2013. pp. 246–256. ACM (2013)
60. Sankaranarayanan, S., Fainekos, G.: Falsification of temporal properties of hybrid systems using the cross-entropy method. In: HSCC 12. pp. 125–134. ACM (2012)
61. Shah, S., Dey, D., Lovett, C., Kapoor, A.: Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In: Field and Service Robotics (2017), <https://arxiv.org/abs/1705.05065>
62. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. CoRR **abs/1409.1556** (2014)
63. Stewart, L., Musa, M., Croce, N.: Look no hands: self-driving vehicles’ public trust problem (2019), <https://www.weforum.org/agenda/2019/08/self-driving-vehicles-public-trust/>, accessed: 2021-01-18
64. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I.J., Fergus, R.: Intriguing properties of neural networks. CoRR **abs/1312.6199** (2013)
65. Tian, Y., Pei, K., Jana, S., Ray, B.: Deeptest: Automated testing of deep-neural-network-driven autonomous cars. In: Proceedings of the 40th International Conference on Software Engineering. pp. 303–314. ACM (2018)
66. Tuncali, C.E., Fainekos, G., Prokhorov, D., Ito, H., Kapinski, J.: Requirements-driven test generation for autonomous vehicles with machine learning components. arXiv preprint arXiv:1908.01094 (2019)
67. Tuncali, C.E., Fainekos, G.E., Ito, H., Kapinski, J.: Sim-atax: Simulation-based adversarial testing framework for autonomous vehicles. In: Proceedings of the 21st International Conference on Hybrid Systems: Computation and Control (part of CPS Week), HSCC 2018, Porto, Portugal, April 11–13, 2018. pp. 283–284. ACM (2018). <https://doi.org/10.1145/3178126.3187004>, <http://doi.acm.org/10.1145/3178126.3187004>
68. Udacity: Self-driving car simulator, <https://github.com/udacity/self-driving-car-sim>, accessed: 2019-08-23
69. Unity3D: Unity game engine, <https://unity3d.com/>, accessed: 2019-08-23
70. Wan, Z., Hudak, P.: Functional reactive programming from first principles. In: Proceedings of the 2000 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Vancouver, British Columbia, Canada, June 18–21, 2000. pp. 242–252. ACM (2000). <https://doi.org/10.1145/349299.349331>, <https://doi.org/10.1145/349299.349331>
71. Wicker, M., Huang, X., Kwiatkowska, M.: Feature-guided black-box safety testing of deep neural networks. In: Beyer, D., Huisman, M. (eds.) Tools and Algorithms for the Construction and Analysis of Systems - 24th International Conference, TACAS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14–20, 2018, Proceedings, Part I. Lecture Notes in Computer Science, vol. 10805, pp. 408–426. Springer (2018). https://doi.org/10.1007/978-3-319-89960-2_22, https://doi.org/10.1007/978-3-319-89960-2_22
72. Wymann, B., Espié, E., Guionneau, C., Dimitrakakis, C., Coulom, R., Sumner, A.: TORCS, The Open Racing Car Simulator. <http://www.torcs.org> (2014)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.