Research Article



Corner Case Generation and Analysis for Safety Assessment of Autonomous Vehicles

Transportation Research Record I-14 © National Academy of Sciences: Transportation Research Board 2021 Article reuse guidelines: sagepub.com/journals-permissions DOI: 10.1177/03611981211018697 journals.sagepub.com/home/trr



Haowei Sun¹, Shuo Feng¹, Xintao Yan¹, and Henry X. Liu¹

Abstract

Testing and evaluation is a crucial step in the development and deployment of connected and automated vehicles (CAVs). To comprehensively evaluate the performance of CAVs, it is necessary to test the CAVs in safety-critical scenarios, which rarely happen in a naturalistic driving environment. Therefore, how to purposely and systematically generate these corner cases becomes an important problem. Most existing studies focus on generating adversarial examples for perception systems of CAVs, whereas limited efforts have been put into decision-making systems, which is the highlight of this paper. As the CAVs need to interact with numerous background vehicles (BVs) for a long duration, variables that define the corner cases are usually high-dimensional, which makes the generation a challenging problem. In this paper, a unified framework is proposed to generate corner cases for decision-making systems. To address the challenge brought by high dimensionality, the driving environment is formulated based on the Markov decision process, and the deep reinforcement learning techniques are applied to learn the behavior policy of BVs. With the learned policy, BVs behave and interact with the CAVs more aggressively, resulting in more corner cases. To further analyze the generated corner cases, the techniques of feature extraction and clustering are utilized. By selecting representative cases of each cluster and outliers, the valuable corner cases can be identified from all generated corner cases. Simulation results of a highway driving environment show that the proposed methods can effectively generate and identify the valuable corner cases.

Testing and evaluation of connected and automated vehicles (CAVs) have been studied for years (1-11). To comprehensively evaluate the performance of CAVs, it is crucial to test the CAVs in different scenarios, especially the safety-critical ones. In a naturalistic driving environment (NDE), however, safety-critical scenarios rarely happen, so it is very time-consuming and inefficient to collect corner cases from either on-road test or simulation test of NDE (12, 13). Therefore, how to purposely and systematically generate corner cases becomes an important problem.

Most existing studies for corner case generation focus on the perception systems of CAVs. Utilizing the methods from the domain of computer vision, researchers aim to generate adversarial examples, which can fool the perception system of CAVs, such as misleading the object classification results (14–16) and hiding pedestrians from the perception results (17). Consequently, disturbed CAVs may miss safety-critical information and encounter dangerous situations.

The decision-making system is also essential in keeping the safety driving of CAVs. Even though the perception system is perfect in the sense that every object of interest can be correctly observed and recognized, the failure of the decision-making system can still cause severe accidents. However, how to generate corner cases for decision-making systems still lacks investigation. Toward solving this issue, several methods have been recently proposed to generate corner cases of simple scenarios (e.g., cross-walking) using adversarial machine learning techniques (18). For the real-world traffic environment (e.g., highway driving environment), however, CAVs need to interact with multiple background vehicles (BVs) for a long duration, so the variables that define the corner cases will be high-dimensional. To the best of our knowledge, no existing method can handle such a high dimensionality for corner case generation.

In this paper, a unified framework is proposed for the high-dimensional corner case generation problem. To

Corresponding Author: Shuo Feng, fshuo@umich.edu

¹University of Michigan, Ann Arbor, MI

address the challenge brought by high dimensionality, the Markov decision process (MDP) is introduced to formulate the traffic environment, which can simplify the temporal dependency between different snapshots of the simulation. Compared with existing driving models commonly used in prevailed simulation platforms, such as the intelligent driver model (IDM) (19), the "minimizing overall braking induced by lane change" (MOBIL) model (20), and the Krauss model (21), the MDP-based driving model can incorporate the behavior randomness of the real-world datasets, which makes the simulation more realistic and reliable. Moreover, with the reinforcement learning method, the simulation obtained can generate much more corner cases purposely, comparing with the NDE. To further simplify the spatial dependency, the BVs of the environment are assumed to make decisions simultaneously and independently during each time step, which is commonly accepted in existing studies (22, 23). Using naturalistic driving data (NDD), the empirical distributions of BV's actions can be obtained for every state. By sampling actions of BVs from the distributions at each time step, the NDE can be generated (12, 13). Based on this formulation, the corner case generation problem of the driving environment is equivalent to optimizing the behavior policy of BVs to improve the probability of corner cases. To achieve this optimization objective, deep reinforcement learning (DRL) (24) techniques are utilized to learn the optimal behavior policy of BVs. With the learned policy, the BVs behave more aggressively when interacting with CAVs, and therefore systematically generate corner cases for a complex driving environment, such as highway driving environment.

Owing to the diversity of corner cases, the generated cases can usually be divided into several clusters as well as outliers that do not belong to any specific cluster. Therefore, typical corner cases from each cluster and the outliers can represent the internal property of the generated corner cases, which are referred to as "valuable" corner cases hereafter. Given the generated corner cases, it is crucial to identify the valuable corner cases, which are usually more valuable for the evaluation and development of CAVs. To achieve this goal, feature extraction and clustering techniques are introduced. As the generated corner cases could be high-dimensional, principal component analysis (PCA) (25) is utilized to reduce the dimension of corner cases and extract the principal features. For the extracted features, clustering methods are applied to identify different clusters of corner cases as well as specific outliers. Consequently, valuable corner cases can be identified.

To validate the proposed method, experiments in a highway driving environment are simulated. To generate the corner cases, one of the commonly used DRL methods, the dueling deep Q network (DQN) (24), is applied to learn the optimal behavior policy of BVs. After the feature extraction by PCA, K-means (26) and the "density-based spatial clustering of applications with noise" (DBSCAN) algorithms (27) are utilized to analyze the corner cases. The experiment results show that the proposed method can effectively generate and analyze valuable corner cases.

The rest of the paper is organized as follows. First, the related works about corner case generation are introduced. Second, we propose a new unified framework for corner case generation using MDP and DRL techniques. Third, feature extraction and clustering techniques are applied for corner case analysis. After that, two case studies are provided to validate the proposed corner case generation and analysis methods. Finally, we conclude and discuss future works.

Related Works

In the CAV testing and evaluation domain, the ability to test the performance of CAVs under different scenarios is crucial. Therefore, researchers in this field have been working to find efficient and reasonable methods of generating corner cases for both CAV perception testing and decision-making testing.

Corner Case Generation for Vehicle Perception

In the CAV perception field, one popular research area in generating corner cases is the adversarial examples generation method. Many autonomous vehicle manufacturers such as Tesla and Waymo have been leveraging neural networks for perception purposes. However, with the rapid development of deep learning and neural network training, many researchers have found that with a small perturbation on the training examples, machine learning systems can be fooled (28–30). These contaminated examples are defined as adversarial examples. For example, Szegedy et al. (31) proposed a method to generate gradient-based optimization adversary examples. By minimizing the difference between adversarial examples and normal training examples while modifying the predicted label from the machine learning system, the proposed algorithm can automatically generate adversarial examples for specific neural networks.

With regard to autonomous vehicle testing, there are also many approaches to interfering with vehicle perception, where most attention is focused on attacking the object detection system. Xie et al. (17) proposed an attack algorithm of object recognition system by removing the pedestrian segmentation. Kurakin et al. (32) showed that perturbation in the physical world instead of in the image perception area can lead to fatal errors of the CAV perception system. Furthermore, Evtimov et al. (15) and Eykholt et al. (16) designed a real-world stop sign with domain knowledge from adversarial examples and successfully let the object recognition system classify it as a speed limit sign. Several works such as the study by Lu et al. (33, 34) showed that the adversarial stop signs in the physical world will not fool the modern CAV perception system as the CAV is continuously moving and detecting the object at each time step. However, Chen et al. (14) showed that the perception system of CAV is still vulnerable to specific adversarial examples.

Corner Case Generation for Vehicle Decision-Making

Researchers have also been focusing on generating corner cases for CAV decision systems. Ma and Peng (35) proposed a worst-case evaluation method, which formulated the disturbance generator and the controller as two players in a game, and generated corner cases by finding the worst inputs from steering controllers and integrated chassis controllers. However, the proposed method focuses on a single vehicle, without considering the influence of other traffic participants, which is crucial in evaluating and testing CAVs.

To generate cases with multiple traffic participants, researchers introduced the risky index and the probabilistic model of the environment to help generate critical cases. For example, Zhao et al. (36, 37) introduced importance sampling techniques and generated testing cases for car-following and lane-changing maneuvers. To solve the overvalue problem of worst cases, Feng et al. (38-41) proposed a new definition of scenario criticality, which can be computed as a combination of maneuver challenge and exposure frequency, and generated critical cases using optimization methods and reinforcement learning techniques on various environment settings, including cut-in scenarios and car-following scenarios. To further address the challenge brought by the high dimensionality of complex environments (e.g., highway driving), Feng et al. (12) proposed a framework of generating a naturalistic and adversarial driving environment by adding sparse but adversarial adjustments to the NDE. Akagi et al. (42) use a self-defined risky index and NDD to sample critical cut-in scenarios. O'Kelly et al. (43) utilized neural network and imitation learning to calibrate the naturalistic driving model from the nextgeneration simulation data, and then chose a highway with six vehicles to generate testing cases. In summary, these critical case generation methods consider both the risky index and the naturalistic probability in the generating process. Although the critical cases are significant for the systematic evaluation of CAVs, the corner cases are also important especially for the vulnerability identification of CAVs, which is complementary to the critical cases. How to generate and identify corner cases with high coverage, variability, and representativeness remain open questions.

To deal with corner cases with long time duration, researchers introduced MDP and reinforcement learning techniques to reduce the temporal complexity. Ding et al. (18) modeled the environment as the combination of "blocks" and used the REINFORCE algorithm to generate corner cases. However, the modeling method can only be used in a simplified environment and cannot deal with multiple traffic participants. Koren et al. (4) proposed the adaptive stress testing method, which introduced Monte Carlo tree search and DRL to solve the pedestrian-crossing problem. In this study, however, it also only involves one autonomous vehicle and one or two pedestrians. Karunakaran et al. (44) utilized DQN to generate corner cases involving one pedestrian and one autonomous vehicle. However, the action chosen for the pedestrian is very simple, and the risk estimation only concentrates on the responsibility-sensitive safety metric (45), which may become misleading in predicting the crash probability.

To address the limitations of the existing decisionmaking corner case generation method, this paper proposes a corner case generation method over the decision domain. In most scenarios, the corner case is a sequence of snapshots of the environment with multiple traffic participants. Therefore, the corner cases in real-world traffic environments always have high dimensions. However, existing decision-making corner case generation methods are only validated under simplified scenarios and cannot process highly complex environments. In this paper, we propose a corner case generation method for high-dimensional and complex traffic simulations. By modeling the environment with MDP, the complexity of the temporal domain is simplified. Moreover, we model the scenario as interactions between multiple traffic participants, which can handle the challenge of dimensionality in space.

Corner Case Generation

In this section, we propose a unified framework for corner case generation based on MDP formulation and DRL techniques. To address the challenge brought by high dimensionality, we formulate the traffic simulation environment as an MDP. By utilizing the NDD, we build naturalistic driving models (e.g., car-following and lanechanging models). In this way, the NDE can be modeled as the interactions between multiple BVs with naturalistic driving models. To purposely generate the corner cases, we formulate the generation problem as an optimization problem of the driving models of BVs. The goal of the optimization problem is to improve the probability of crashes. By utilizing the DRL techniques, the optimization problem can be solved by learning the behavior policy of BVs. The learned behavior policy essentially leads to aggressive driving models of BVs, which can generate corner cases for the CAV under test.

Problem Formulation

In this paper, the problem formulation is consistent with (12, 38–41). Let θ describe the pre-determined parameters of the operational design domains (ODDs), such as the number of lanes, weather, and so forth. The definitions of scenario and scene are adopted from Ulbrich et al. (46). Under specific ODD parameters, a scene describes the snapshot of the traffic environment, which includes the states of static elements and dynamic traffic participants (e.g., pedestrians and vehicles). A scenario describes the temporal development among a sequence of scenes. Let X represent the decision variables of the scenario and s denote the state of the scene. In this paper, we only consider BVs, and each vehicle has three parameters to describe the overall state: position p, velocity v, and heading angle α . Therefore, the decision variable of one vehicle can be defined as $x = \{p, v, \alpha\}$. In each time step, the vehicle numbers in the CAV's neighborhood are different, so we define the number of observed vehicles in time step t as m_t . Then, the scene of time step t can be defined as $s_t = \{x_1^{(t)}, x_2^{(t)}, \cdots, x_{m_t}^{(t)}\}$, and the scenario can be defined as:

$$X = \{s_0, s_1, \cdots, s_n\}.$$
 (1)

To simplify the problem, we model the traffic environment as an MDP. Given a specific state s_t , the "agent" can represent the BVs in the environment and make a decision u_t (e.g., accelerations). As the u_t is only determined by state s_t , it can also be written as $u_t(s_t)$. Therefore, the scenario can be rewritten as follows:

$$X = \{s_0, u_0, s_1, u_1, s_2, \cdots, s_n\}.$$
 (2)

Then, we define A as the event of interest (e.g., crash event). For a given scenario X, we can clearly identify whether it is the event of interest. Therefore, $P(A|X, \theta)$ is known for given X. Under specific ODD, the probability of the event of interest can be written as follows:

$$P(A|\theta) = \sum_{X \in D} P(A|X, \theta) P(X|\theta),$$
(3)

where $P(X|\theta)$ denotes the probability of specific scenario X given ODD parameter θ , and D represents the set of available scenarios. By using the notation from Equation 2, we can further decompose $P(X|\theta)$ in a factorized way as:

$$P(X|\theta) = P(s_0|\theta) \prod_{k=0}^{m-1} P(s_{k+1}|u_k, s_k, \theta) P(u_k|s_k, \theta).$$
(4)

Here $P(s_{k+1}|u_k, s_k, \theta)$ is the state transition probability, which means the probability of the occurrence of s_{k+1} given state s_k and action u_k , and $P(u_k|s_k, \theta)$ denotes the probability of choosing action u_k at state s_k .

In Equation 4, $P(s_0|\theta)$ and $P(s_{k+1}|u_k, s_k, \theta)$ are determined by the environment, and P(A|X) is determined by the CAV under test, which cannot be modified. Therefore, to improve the exposure frequency of corner cases ($P(A|\theta)$ in Equation 3), we should modify the $P(u_k|s_k, \theta)$ such that the $P(X|\theta)$ can be increased for X where P(A|X) = 1. $P(u_k|s_k, \theta)$ denotes the probability of agent choosing action u_k under state s_k , so it can be viewed as a stochastic policy. If the policy can be optimized to improve $P(A|\theta)$, the corner cases can be generated more purposely. To achieve this goal, DRL techniques are utilized to train a new policy as the replacement of $P(u_k|s_k, \theta)$, which will be elaborated in the next section.

DRL-Based Method

By replacing $P(u_k|s_k, \theta)$ with the modified behavior policy $\pi_{\rho}(u_k|s_k, \theta)$, we can obtain higher probability of the event of interest. The new definition is shown in the following equation:

$$P(X|\theta,\rho) = P(s_0|\theta) \prod_{k=0}^{m-1} P(s_{k+1}|u_k, s_k, \theta) \pi_{\rho}(u_k|s_k, \theta).$$
(5)

The newly defined $\pi_{\rho}(u_k|s_k, \theta)$ represents a behavior policy to be learned in the DRL problem. Detailed optimization formulation can be seen in Equation 6:

$$\max_{\rho} \quad P(A|\theta,\rho) = \sum_{X \in D} P(A|X,\theta) P(X|\theta,\rho).$$
(6)

The optimization problem can be considered as training a specific behavior policy in the simulation environment, which lets BVs aggressively interact with the CAV. Therefore, given the complexity of the environment, DRL techniques can have a good performance in solving the optimization problem. As an unsupervised algorithm, DRL techniques can learn optimal policy from experience given specific reward settings. By implementing the reinforcement learning algorithms and using the deep neural network (DNN) as the function approximator, DRL can solve the problem with high complexity and derive well-behaved policy in aspects of video games (24), robotics (47), and so forth. Following the ideas of traditional reinforcement learning, DRL also has three different approaches: value-based DRL, policy-based DRL, and actor-critic DRL (24, 48-53). In this paper, we mainly use the value-based DRL to solve the optimization problem.

Value-based DRL, also commonly known as the DQN, aims at learning the optimal state-action value function $Q(s_t, u_t)$ by estimating the expected reward of action u_t given state s_t . Researchers introduce DNN here to represent the Q function, where the neural network will accept the state and action information as the input and return the estimated state-action value. Specifically, DQN uses a neural network with parameter ρ to approximate the optimal state-action value function:

$$Q_{\rho}^{*}(s,u) = \max_{\pi} E\left[\sum_{i=0}^{\infty} \gamma^{i} r_{t+1} | s_{t} = s, u_{t} = u, \pi_{\rho}\right], \quad (7)$$

where r_t denotes the immediate reward received at step t. Therefore, with the optimal state–action value function, we can easily derive the optimal policy that can increase the probability of corner cases. The optimal policy will be a deterministic policy as follows:

$$P(u_t|s_t) = \begin{cases} 1 & u_t = \arg\max_u Q(s_t, u) \\ 0 & \text{otherwise} \end{cases}.$$
 (8)

By only giving crash events positive reward r_A , with the training process of the DQN agent, the optimal value function and the optimal policy will bring about a higher probability of crash events.

To further improve the performance of DRL techniques, dueling network architecture (52) is introduced. By defining the new definition value function V(s) and advantage function A(s, u), the dueling network architecture can estimate the value function of the state-action more precisely. The connection between the dueling network and DQN can be seen as follows:

$$A(s, u) = Q(s, u) - V(s).$$
 (9)

Therefore, two neural networks are introduced to represent A(s, u) and V(s). These two networks often share some hidden layers, as shown in Figure 1.

Corner Case Analysis

With the learned behavior policy, many corner cases will be generated by simulating the interaction between the agent and the environment. With a detailed understanding of generated corner cases, the evaluation and development of CAVs can be more targeted and effective.

Owing to the diversity of corner cases, the generated cases can usually be divided into several clusters as well as outliers that do not belong to any specific cluster. Therefore, typical corner cases from each cluster and the outliers can represent the internal property of the generated corner cases, which are referred to as valuable



Figure 1. Dueling network architecture.

corner cases. To identify the valuable corner cases, we need two techniques: feature extraction and clustering. The long time period and multiple traffic participants lead to the high dimension of corner cases, which further bring about difficulty in the analysis. To address the difficulty, we introduce the feature extraction method to reduce the dimension and extract key information from the cases. In this paper, we utilize the PCA algorithm to reduce the dimension. Moreover, corner cases can usually be classified into different clusters owing to the diversity of the generated corner cases. Different types of cases may differ significantly with regard to frequency. Some types may occupy the majority of corner cases whereas other types may become the minority. To keep the diversity, the majority and minority need to be analyzed separately. To achieve this goal, the DBSCAN method is used to separate the majority and minority of corner cases as two groups. Then, for each group of the corner cases, the K-means method is utilized to cluster the corner cases. As unsupervised methods, the DBSCAN and K-means methods can be applied to cluster generic corner cases without knowing the predefined labels.

Corner Case Feature Extraction

Corner cases are sequences of continuous snapshots of traffic environments, which always contain multiple traffic participants. Therefore, corner cases always have high dimensions, which bring about difficulties in the analysis. Feature extraction techniques can be utilized to reduce the dimension of the feature space while keeping essential information. PCA (25) is one of the most popular feature extraction methods. By maximizing the variance in every direction, the PCA method can project data points from a high-dimensional space to a low-dimensional space and

transform the data points to a new coordinate system. The first k directions given by PCA are also the directions on which the largest variance of data projection lies. Therefore, in most cases, only picking several coordinates given by PCA will bring about essential information of the original data. By introducing the PCA algorithm, we can extract several key features and simplify the analysis process in the next steps. Additionally, PCA can differentiate and visualize the distance and relatedness between different populations (i.e., data clusters). Therefore, the data points after the PCA projection will be more suitable for the next-step clustering analysis.

Corner Case Clustering

In the generated corner cases, some normal scenarios may occupy a large proportion. For example, some libraries may contain many homogeneous cases (e.g., many rear-end collisions). However, the value of normal scenarios is limited, as the increase in normal corner cases does not result in new information for CAV performance. Instead, we should pay more attention to valuable scenarios in the generated corner cases, such as typical corner cases of a specific corner case cluster, and the outliers that have significantly different properties compared with the majority of corner cases. In this paper, to detect the outliers, we utilize the DBSCAN method (27), which groups the data points that are closely packed together. Therefore, the outliers can be identified as the data points with low density, namely, the minority of corner cases.

After differentiating the majority and the minority, further analysis can be applied separately. Owing to the diversity of the generated corner cases, different cases have different internal patterns and are always distributed in different areas of the feature space. Therefore, by applying the clustering method, corner cases can be classified as different types. Additionally, some typical cases can represent many corner cases (in the same cluster). To automatically classify the type of different data points, the clustering method is commonly used. In this domain, the K-means method (26) is among the most popular algorithms. As a non-parametric unsupervised learning method, K-means provides the clustering result minimizing the in-group variance (squared Euclidean distances) for a given objective cluster number, which can help differentiate clusters with different internal properties.

Valuable Corner Case Identification

After corner case clustering, valuable corner case extraction becomes the next topic. The value of one specific corner case represents how much it can help in evaluating CAVs and improving the performance of CAVs.

From this perspective, two different types of cases can be defined as valuable: typical cases and rare cases. Typical cases usually represent multiple corner cases from the same cluster, and rare cases are usually the outliers, which have distinct properties compared with the majority and rarely happen in the NDE. Correspondingly, two techniques are utilized to extract valuable corner cases. First, after data clustering of corner cases, the typical corner cases can be selected by the distance to the center of a specific cluster. Second, by outlier detection methods powered by clustering, rare corner cases with various properties can be extracted. Then, the identification of valuable corner cases can be achieved. In this way, the identified valuable corner cases can well balance the consideration of coverage, variability, and representativeness of the cases.

Case Study

In this section, we validate our methodology using experiments in the highway simulation platform. This section consists of four parts. First, we introduce the NDD processing and NDD driving models. Second, we introduce the simulation platform used in the case study. Third, the proposed corner case generation method is validated in the simulation platform. Finally, by utilizing the corner case analysis method, we extract the valuable corner cases from different experiment settings.

NDD Processing

To build the naturalistic highway driving model, we implement a data-driven stochastic model from the integrated vehicle-based safety systems dataset (54) at the University of Michigan, Ann Arbor, Michigan, and the safety pilot model deployment dataset. By integrating the forward collision warning, lane departure warning, lane-change warning, and curve speed warning functions on passenger cars and heavy trucks, this project aims to prevent rear-end and other crashes. The project collected 650,000 mi of driving data on heavy trucks and 175,000 mi of driving data on regular vehicles.

To calibrate the data-driven naturalistic driving model for the highway environment, we collected data points in which the velocity of vehicles is between 20 and 40 m/s. As a result, we collected around 3×10^6 data points and built up the empirical distribution of BV's action for every state. By sampling actions from the empirical distribution, all BVs are essentially controlled by the naturalistic driving model, which formulate the NDE.

Simulation Platform

Our simulation platform is based on the open-source simulation platform HIGHWAY-ENV (55). This

Table I.	Experiments	Hyper-Parameters
----------	-------------	------------------

Hyper-parameter	Value 16
Mini-batch size	
Replay memory size	le6
Discount factor	I
Learning rate	le-6
Initial exploration	1
Final exploration	0.1
Replay start size	5,000
Target network update	1,000

simulation environment is fully compatible with the OpenAI gym environment (56), so it can be used to train the autonomous vehicle planning algorithm. Additionally, compared with other simulation platforms such as CARLA and SUMO, the HIGHWAY-ENV simulation platform is more lightweight and with higher computational efficiency. In the HIGHWAY-ENV environment, the IDM car-following model (19) and MOBIL model (20) for lane changing are used to provide continuous traffic flow and reasonable vehicle behaviors. However, default vehicle models are deterministic and cannot perform the naturalistic nature of vehicle behavior, which is not suitable for the corner case generation process. To overcome this limitation, we re-design the simulation environment and add control API to improve the controllability of BVs, so the BVs can be controlled by the naturalistic driving model.

Corner Case Generation

We implement the DQN (24) method using PyTorch and SGD as the optimizer. Detailed experiment hyperparameters are listed in Table 1. In the dueling neural network architecture (52), there are four fully connected layers, each with 128 units. As is shown in Figure 1, the dueling network splits into two streams of fully connected layers: the value stream and the advantage stream. Each stream has two fully-connected layers with 128 units. The final output layer of the state stream and the value stream are also fully connected. The value stream has one output and the advantage stream has 33 outputs (pre-defined discrete actions). The value stream output and the advantage stream output are combined using Equation 9. During the training process, +1reward is given when BVs successfully crash into the CAV, and -1 reward is given when BVs crash into each other. Otherwise, we give zero reward. The experiments are implemented on Ubuntu 18.04 LTS with i9-9900k CPU, RTX 2080 TI GPU, and 64 GB of RAM. The agent was trained for around 100,000 episodes in 4 days until it reaches the performance limit, and the training process of the agent can be seen in Figure 2. In the end,



Figure 2. Deep Q network training result.

the crash rate in the corner case environment converges to 0.3.

In the case study, we test one commonly used CAV model, which is constructed by the IDM car-following model and MOBIL lane-change model. To better interpret the results, the CARLA simulator (57) is used to visualize corner cases. We implement the NDE (12, 13) as our baseline, which contains the CAV model and NDD vehicles that follow the naturalistic driving model. In the NDE, the crash case frequency is around 1×10^{-6} , which is rare. In the corner case generation environment, we control the nearest BV around the CAV using the trained DQN. Results show the crash case frequency is about 0.286 as shown in Figure 2. Therefore, for corner case generation frequency, we get about 3×10^{6} times more corner cases compared with the NDE.

To evaluate the corner case generation environment, we run about 50 mi in both NDE and the corner case generation environment. The distribution of the bumperto-bumper distance (BBD) and time-to-collision (TTC) metric are calculated to compare the difference between NDE and the corner case generation environment. The comparison of BBD can be seen in Figure 3a (from the CAV to the front vehicle) and Figure 3b (from CAV to the rear vehicle), and the comparison of TTC for front and rear vehicles can be seen in Figure 3, c and d, respectively. From these figures, we can see that the testing vehicle in the corner case generation environment has much smaller BBD and TTC with both the front car and the rear car. It suggests that the corner case generation environment is much riskier than NDE.

To further test whether the corner case environment could fit non-surrogate CAV models, a CAV model based on reinforcement learning is introduced as another



Figure 3. Comparison between naturalistic driving environment and corner case generation environment: (*a*) BBD front car, (*b*) BBD rear car, (*c*) TTC front car, and (*d*) TTC rear car.

Note: BBD = bumper-to-bumper distance; TTC = time to collision.

testing model. When the corner case environment is utilized for the surrogate IDM-based model, we get approximately 28.6% crash probability, whereas for the reinforcement learning-based testing model we get a 10% crash rate.

Additionally, a logic-based crash-type analysis is introduced to help illustrate the generated corner cases in detail. In this study, we adopted the crash-type diagram defined by the Fatality Analysis Reporting System, which is a nationwide census provided by the National Highway Traffic Safety Administration. Specifically, we categorize the generated corner cases into five types according to the positions and angles of the CAV and BV involved in each crash. Detailed illustrations of the categories can be seen in Figure 4, where the blue vehicles denote CAVs and the green vehicles denote BVs. The crash-type distribution of generated corner cases of IDM-based CAV and reinforcement learning-based CAV can be seen in Figure 5. From the figure, we can see that both experiments contain the type 1, 2, 4, 5 of the overall categories, and type 4 occupies the most proportion. For each crash type, a detailed corner case demonstration is attached to illustrate the case type in detail, which can be seen in Figure 6.

Corner Case Analysis

Corner Case Feature Extraction. In the generated corner cases (around 50,000 scenarios), different cases have different sizes. For the convenience of corner case analysis, we need one unified structure of selecting features from the corner cases. As discussed before, the scenario is composed of scenes, and the scene can be written as $\{x_1^{(t)}, x_2^{(t)}, \ldots, x_{m_t}^{(t)}\}$. Recall that the vehicle number m_t is determined by time t and is continuously changing at each time step. Therefore, to obtain features with the same shape from different time steps, we need to restrict the number of vehicles recorded. Based on domain



Figure 4. Illustration of crash types.



Figure 5. Distributions of crash types.

knowledge, we select the most critical BV in the crash corner cases (i.e., the nearest vehicle in the BVs). Then, we use the following equations as the extracted feature of one time step:

$$[r_{lon}, r_{lat}, rr, h_{CAV}, h_{BV}],$$
(10)

where r_{lon} refers to the longitudinal relative distance, r_{lat} refers to the lateral relative distance, rr refers to the velocity difference, and h_{CAV} , h_{BV} represent the heading angle of CAV and critical BV, respectively. The feature extracted can represent the key information in a snapshot of the traffic simulation. To characterize the long-term change of the traffic state, we continuously pick features from k time steps before crashes.

In the experiment, eight values of the time period are applied, starting from the one time step to 15 consecutive time steps. For each time step, five features are considered as shown in Equation 10. Therefore, for 15 time steps, 75 features are considered. To extract the crucial features from the high-dimensional data, we apply the PCA algorithm and pick the first two dimensions. After the feature extraction and projection, we implement the DBSCAN clustering method to differentiate the minority and the majority of the generated corner cases. The distribution of the data points projection on the twodimensional (2-D) plane can be seen in Figure 7. The data points located in the high-density area are classified



Figure 6. Demonstration of each crash type in CARLA.

as the majority (green), whereas the low-density ones are classified as the minority (red). From the clustering result, one obvious trend is that, with the increase of the time period, the data points seem to be less gathered. With regard to the features of 1 or 2 s, only several data points are identified as the minority (marked as red). However, when it turns to quiet a long time period (15 s), nearly all data points are randomly distributed over the feature space, which makes it harder to differentiate the majority and the minority. Therefore, in the following analysis, we use the 1 s data as an example to illustrate the result of the generated corner cases. It is also reasonable considering that most of the vehicle accidents in the real world involve only a small number of vehicles in a short period.

Corner Case Clustering. After applying the DBSCAN method on the generated corner cases, we further cluster the minority of generated corner cases. As shown in Figure 8, the majority (labelled area A) are distinguished from the minority. We can see that the majority of the corner cases form a rectangle in the projected PCA feature space. In this rectangle, most cases share some characteristics in common: the CAV is running straight on



Figure 7. Corner case extracted features.



Figure 8. Corner case clustering results.

the road, and the BV is directly crashing into the CAV from different angles and different positions as shown in Figure 9*a*. In this type of crash, the crash angle and crash relative position parameters are continuous in the parameter space, whereas other parameters are the same. Therefore, even though it seems that there are many different kinds of crashes in this crash type, they are closely connected in the PCA feature space.

From the 2-D projection of the feature vector, we can see that the minority can be assigned to several clusters, whereas the majority is closely connected. Therefore, in this case, we apply the clustering method (K-means) on the minority, resulting in four clusters (B, C, D, and E in Figure 8). The clustering process took 0.097 s on a laptop



Figure 9. Corner case examples.

with i7-8750h CPU and 24 GB RAM. For each cluster, cases of each cluster share similar internal properties. For example, cases of B cluster in Figures 8 and 9 demonstrate the rear-end cases, in which one BV cuts in the CAV's lane and forces the CAV to change its lane, making the CAV crash into another BV in the original lane behind the CAV. Cases of C cluster indicate another crash type: CAV and BV change their lanes simultaneously and crash into each other. Cases of D cluster show that the BV tries to change into the lane of the CAV and causes a crash when making the lane-change decision. The only case in the E cluster behaves similarly to the cases in the D cluster. However, there are some slight differences: the crash in the E cluster happens after the BV



Figure 10. Corner case clustering results in the naturalistic driving data (NDD)-bounded case study.

changes into the lane of the CAV and can be defined as a rear-end crash.

NDD-Bounded Case Study

We also provide the analysis results of another NDDbounded generation method. In Equation 6, the controlled BV can choose any actions in the pre-defined action space. However, in the NDE, the BVs generally have limited choices of actions. To solve this issue, we slightly modify the optimization problem in Equation 6 by adding action constraints. By applying the constraints, we can restrict the available actions of BVs, so the agent in the environment can only choose the actions that are possible in the NDE. The new optimization problem can be written as follows:

$$\max_{\rho} \quad P(A|\theta,\rho), \tag{11}$$

s.t.
$$\pi_{p}(u_{k}|s_{k},\theta) = 0$$
 if $P(u_{k}|s_{k},\theta) = 0.$ (12)

Therefore, the agent trained from the modified optimization problem can only reproduce the corner cases that are likely to happen in the NDE. In this way, the generated corner cases are much more realistic and valuable for the CAV evaluation. Using the same analysis method of the previous result, we apply PCA on the crash event data and reduce the data into two dimensions. After that, we apply the DBSCAN algorithm on the PCA projected features to get different clusters and outliers. The experiment is applied on 1,627 data samples and takes 0.0139 s on a laptop with i7-8750h CPU and 24 GB RAM. Detailed results can be seen in Figure 10.

From Figure 10, we can see that there are two main clusters identified and several outliers (D). To analyze



Figure 11. Naturalistic driving data (NDD)-bounded corner case examples.

the results, we split one large cluster into B and C, which are deeply connected but demonstrate different data layout. As shown in Figure 11, data points in area A share similar properties: the BV suddenly cuts in the CAV and causes the rear-end collision. We define it as the "aggressive cut-in" cluster. B, C, and D are in nearly the same situation: CAV and BV change their lanes at the same time and crash in the middle lane. We define it as the "lane conflict" cases. Even though the causes of crashes are similar, these three different clusters have different crash snapshots. In the B cluster, the CAV and BV are involved in a side-by-side collision. In the C cluster, CAV(BV) crashes into BV(CAV) at the rear of the car. In the D cluster, which is identified as the relatively rare events (minority) in the generated corner cases, we can see that the CAV and BV do not crash until they change to the same lane, after which the rear-end collision happens. Therefore, even though the last three clusters are corner cases caused by the same reason, the decision variables and conditions are different case by case. In the outlier part (D), we can get some cases with extreme decision variables.

The extracted valuable corner cases can provide insight into further improvement of the CAV model. Although some cases are inevitable, there are cases caused by the flaws of the CAV model, and thus can be avoided by more intelligent CAVs. For example, for the CAV model in the case study, one significant limitation is the lack of behavioral competency for lateral collision avoidance, especially during the lane-changing process, as shown in Figure 11 (B, C, and D clusters).

Conclusion

In this paper, we propose a decision-making corner case generation and analysis method for CAV testing and evaluation. By utilizing MDP formulation and DRL techniques, the corner cases of a highway driving environment are purposely generated with a higher probability, compared with NDE. After generating the corner cases, the valuable corner cases are further identified by the corner case analysis method, including feature extraction and clustering techniques. Two case studies are provided to validate the proposed methods. Results show that the valuable corner cases can be effectively generated and identified, which is helpful for CAV evaluation and development by revealing flaws of the given CAV model. Future studies may focus on introducing more different traffic participants (pedestrians, traffic lights, etc.). Furthermore, corner case generation for the urban driving environment deserves more investigation, including intersection and roundabout scenarios.

Author Contributions

The authors confirm contribution to the paper as follows: study concept and design: Haowei Sun, Shuo Feng, and Henry Liu; data collection: Haowei Sun, Xintao Yan; analysis and interpretation of results: Haowei Sun, Shuo Feng, and Henry Liu; draft manuscript preparation: Haowei Sun, Shuo Feng, and Henry Liu. All authors reviewed the results and approved the final version of the manuscript.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was funded by the U.S. Department of Transportation Region 5 University Transportation Center: Center for Connected and Automated Transportation of the University of Michigan.

ORCID iDs

Haowei Sun ib https://orcid.org/0000-0002-0232-117X Shuo Feng ib https://orcid.org/0000-0002-2117-4427 Xintao Yan ib https://orcid.org/0000-0002-0569-5628 Henry X. Liu ib https://orcid.org/0000-0002-3685-9920

References

- Kalra, N., and S. M. Paddock. Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability? *Transportation Research Part A: Policy and Practice*, Vol. 94, 2016, pp. 182–193.
- Li, L., W.-L. Huang, Y. Liu, N.-N. Zheng, and F.-Y. Wang. Intelligence Testing for Autonomous Vehicles: A New Approach. *IEEE Transactions on Intelligent Vehicles*, Vol. 1, No. 2, 2016, pp. 158–166.

- Li, L., Y.-L. Lin, N.-N. Zheng, F.-Y. Wang, Y. Liu, D. Cao, K. Wang, and W.-L. Huang. Artificial Intelligence Test: A Case Study of Intelligent Vehicles. *Artificial Intelligence Review*, Vol. 50, No. 3, 2018, pp. 441–465.
- Koren, M., S. Alsaif, R. Lee, and M. J. Kochenderfer. Adaptive Stress Testing for Autonomous Vehicles. *Proc.*, 2018 IEEE Intelligent Vehicles Symposium (IV), Changshu, China, IEEE, New York, 2018, pp. 1–7.
- Cui, L., J. Hu, B. B. Park, and P. Bujanovic. Development of a Simulation Platform for Safety Impact Analysis Considering Vehicle Dynamics, Sensor Errors, and Communication Latencies: Assessing Cooperative Adaptive Cruise Control under Cyber Attack. *Transportation Research Part C: Emerging Technologies*, Vol. 97, 2018, pp. 1–22.
- Thorn, E., S. C. Kimmel, M. Chaka, and B. A. Hamilton. *A Framework for Automated Driving System Testable Cases and Scenarios*. Technical Report. United States Department of Transportation, National Highway Traffic Safety Administration, Washington, D.C., 2018.
- Li, L., X. Wang, K. Wang, Y. Lin, J. Xin, L. Chen, L. Xu, B. Tian, Y. Ai, J. Wang, and D. Cao. Parallel Testing of Vehicle Intelligence via Virtual-Real Interaction. *Science Robotics*, Vol. 4, No. 28, 2019, Article eaaw4106.
- Wang, P., Y. Chen, C. Wang, F. Liu, J. Hu, and N. N. Van. Development and Verification of Cooperative Adaptive Cruise Control via LTE-V. *IET Intelligent Transport Systems*, Vol. 13, No. 6, 2019, pp. 991–1000.
- Liu, P., Z. Xu, and X. Zhao. Road Tests of Self-Driving Vehicles: Affective and Cognitive Pathways in Acceptance Formation. *Transportation Research Part A: Policy and Practice*, Vol. 124, 2019, pp. 354–369.
- Liu, P., R. Yang, and Z. Xu. How Safe Is Safe Enough for Selfdriving Vehicles? *Risk Analysis*, Vol. 39, No. 2, 2019, pp. 315–325.
- Li, L., N. Zheng, and F.-Y. Wang. A Theoretical Foundation of Intelligence Testing and Its Application for Intelligent Vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 2020, pp. 1–10.
- Feng, S., X. Yan, H. Sun, Y. Feng, and H. X. Liu. Intelligent Driving Intelligence Test for Autonomous Vehicles with Naturalistic and Adversarial Environment. *Nature Communications*, Vol. 12, 2021, p. 748. https://doi.org/10.1038/s41467-021-21007-8.
- Yan, X., S. Feng, H. Sun, and H. X. Liu. Distributionally Consistent Simulation of Naturalistic Driving Environment for Autonomous Vehicle Testing. *arXiv Preprint arXiv:2101.02828*, 2021.
- Chen, S.-T., C. Cornelius, J. Martin, and D. H. Chau. Robust Physical Adversarial Attack on Faster R-CNN Object Detector. *arXiv Preprint arXiv:1804.05810*, Vol. 2, No. 3, 2018, p. 4.
- Evtimov, I., K. Eykholt, E. Fernandes, T. Kohno, B. Li, A. Prakash, A. Rahmati, and D. Song. Robust Physical World Attacks on Machine Learning Models. *arXiv Preprint arXiv:1707.08945*, Vol. 2, No. 3, 2017, p. 4.
- Eykholt, K., I. Evtimov, E. Fernandes, B. Li, A. Rahmati, C. Xiao, A. Prakash, T. Kohno, and D. Song. Robust Physical-World Attacks on Deep Learning Visual Classification. *Proc.*, *IEEE Conference on Computer Vision and*

Pattern Recognition, Salt Lake City, UT, IEEE, New York, 2018, pp. 1625–1634.

- Xie, C., J. Wang, Z. Zhang, Y. Zhou, L. Xie, and A. Yuille. Adversarial Examples for Semantic Segmentation and Object Detection. *Proc., IEEE International Conference on Computer Vision*, Venice, Italy, IEEE, New York, 2017, pp. 1369–1378.
- Ding, W., M. Xu, and D. Zhao. Learning to Collide: An Adaptive Safety-Critical Scenarios Generating Method. *arXiv Preprint arXiv:2003.01197*, 2020.
- Treiber, M., A. Hennecke, and D. Helbing. Congested Traffic States in Empirical Observations and Microscopic Simulations. *Physical Review E*, Vol. 62, No. 2, 2000, p. 1805.
- 20. Treiber, M., and D. Helbing. MOBIL: General Lane-Changing Model for Car-Following Models. *Disponivel Acesso Dezembro*, 2016.
- 21. Krauß, S., P. Wagner, and C. Gawron. Metastable States in a Microscopic Model of Traffic Flow. *Physical Review E*, Vol. 55, No. 5, 1997, p. 5597.
- Bareiss, D., and J. van den Berg. Generalized Reciprocal Collision Avoidance. *The International Journal of Robotics Research*, Vol. 34, No. 12, 2015, pp. 1501–1514.
- Weng, B., S. J. Rao, E. Deosthale, S. Schnelle, and F. Barickman. Model Predictive Instantaneous Safety Metric for Evaluation of Automated Driving Systems. *Proc.*, 2020 *IEEE Intelligent Vehicles Symposium (IV)*, Las Vegas, NV, IEEE, New York, 2020, pp. 1899–1906.
- Mnih, V., K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, and S. Petersen. Human-Level Control through Deep Reinforcement Learning. *Nature*, Vol. 518, No. 7540, 2015, pp. 529–533.
- Wold, S., K. Esbensen, and P. Geladi. Principal Component Analysis. *Chemometrics and Intelligent Laboratory Systems*, Vol. 2, No. 1–3, 1987, pp. 37–52.
- Alsabti, K., S. Ranka, and V. Singh. An Efficient K-Means Clustering Algorithm. *Electrical Engineering and Computer Science*, Vol. 43, 1997.
- Ester, M., H.-P. Kriegel, J. Sander, and X. Xu. Density-Based Spatial Clustering of Applications with Noise. *Proc., International Conference on Knowledge Discovery and Data Mining*, Portland, OR, Vol. 240, 1996, p. 6.
- Goodfellow, I. J., J. Shlens, and C. Szegedy. Explaining and Harnessing Adversarial Examples. arXiv Preprint arXiv:1412.6572, 2014.
- Yuan, X., P. He, Q. Zhu, and X. Li. Adversarial Examples: Attacks and Defenses for Deep Learning. *IEEE Transactions on Neural Networks and Learning Systems*, Vol. 30, No. 9, 2019, pp. 2805–2824.
- Buckner, C. Understanding Adversarial Examples Requires a Theory of Artefacts for Deep Learning. *Nature Machine Intelligence*, Vol. 2, No. 12, 2020, pp. 731–736.
- Szegedy, C., W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing Properties of Neural Networks. *arXiv Preprint arXiv:1312.6199*, 2013.
- 32. Kurakin, A., I. Goodfellow, and S. Bengio. Adversarial Examples in the Physical World. *arXiv Preprint arXiv:1607.02533*, 2017.

- Lu, J., H. Sibai, E. Fabry, and D. Forsyth. No Need to Worry about Adversarial Examples in Object Detection in Autonomous Vehicles. arXiv Preprint arXiv:1707.03501, 2017.
- Lu, J., H. Sibai, E. Fabry, and D. Forsyth. Standard Detectors Aren't (Currently) Fooled by Physical Adversarial Stop Signs. arXiv Preprint arXiv:1710.03337, 2017.
- Ma, W.-H., and H. Peng. A Worst-Case Evaluation Method for Dynamic Systems. *ASME Journal of Dynamic Systems, Measurement, and Control*, Vol. 121, No. 2, 1999, pp. 191–199.
- 36. Zhao, D., H. Lam, H. Peng, S. Bao, D. J. LeBlanc, K. Nobukawa, and C. S. Pan. Accelerated Evaluation of Automated Vehicles Safety in Lane-Change Scenarios Based on Importance Sampling Techniques. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 18, No. 3, 2016, pp. 595–607.
- Zhao, D., X. Huang, H. Peng, H. Lam, and D. J. LeBlanc. Accelerated Evaluation of Automated Vehicles in Car-Following Maneuvers. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 19, No. 3, 2017, pp. 733–744.
- Feng, S., Y. Feng, C. Yu, Y. Zhang, and H. X. Liu. Testing Scenario Library Generation for Connected and Automated Vehicles, Part I: Methodology. *IEEE Transactions* on *Intelligent Transportation Systems*, Vol. 22, No. 3, 2020, pp. 1573–1582.
- Feng, S., Y. Feng, H. Sun, S. Bao, Y. Zhang, and H. X. Liu. Testing Scenario Library Generation for Connected and Automated Vehicles, Part II: Case Studies. *IEEE Transactions on Intelligent Transportation Systems*, 2020, pp. 1–13.
- Feng, S., Y. Feng, H. Sun, Y. Zhang, and H. X. Liu. Testing Scenario Library Generation for Connected and Automated Vehicles: An Adaptive Framework. *IEEE Transactions on Intelligent Transportation Systems*, 2020, pp. 1–10. https://doi.org/10.1109/TITS.2020.3023668.
- Feng, S., Y. Feng, X. Yan, S. Shen, S. Xu, and H. X. Liu. Safety Assessment of Highly Automated Driving Systems in Test Tracks: A New Framework. *Accident Analysis & Prevention*, Vol. 144, 2020, Article 105664.
- 42. Akagi, Y., R. Kato, S. Kitajima, J. Antona-Makoshi, and N. Uchida. A Risk-Index Based Sampling Method to Generate Scenarios for the Evaluation of Automated Driving Vehicle Safety. *Proc., 2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, Auckland, New Zealand, IEEE, New York, 2019, pp. 667–672.
- O'Kelly, M., A. Sinha, H. Namkoong, R. Tedrake, and J. C. Duchi. Scalable End-to-End Autonomous Vehicle Testing via Rare-Event Simulation. *Advances in Neural Information Processing Systems*, Vol. 31, 2018, pp. 9827–9838.
- Karunakaran, D., S. Worrall, and E. Nebot. Efficient Statistical Validation with Edge Cases to Evaluate Highly Automated Vehicles. arXiv Preprint arXiv:2003.01886, 2020.
- 45. Shalev-Shwartz, S., S. Shammah, and A. Shashua. On a Formal Model of Safe and Scalable Self-Driving Cars. *arXiv Preprint arXiv:1708.06374*, 2017.
- 46. Ulbrich, S., T. Menzel, A. Reschka, F. Schuldt, and M. Maurer. Defining and Substantiating the Terms Scene,

Situation, and Scenario for Automated Driving. Proc., 2015 IEEE 18th International Conference on Intelligent Transportation Systems, Gran Canaria, Spain, IEEE, New York, 2015, pp. 982–988.

- 47. Tai, L., G. Paolo, and M. Liu. Virtual-to-Real Deep Reinforcement Learning: Continuous Control of Mobile Robots for Mapless Navigation. *Proc.*, 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Vancouver, BC, IEEE, New York, 2017, pp. 31–36.
- 48. Sutton, R. S., and A. G. Barto. *Reinforcement Learning: An Introduction.* MIT Press, Cambridge, MA, 2018.
- Mnih, V., K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing Atari with Deep Reinforcement Learning. arXiv Preprint arXiv:1312.5602, 2013.
- Van Hasselt, H., A. Guez, and D. Silver. Deep Reinforcement Learning with Double Q-Learning. arXiv Preprint arXiv:1509.06461, 2015.
- 51. Hessel, M., J. Modayil, H. Van Hasselt, T. Schaul, G. Ostrovski, W. Dabney, D. Horgan, B. Piot, M. Azar, and D. Silver. Rainbow: Combining Improvements in Deep Reinforcement Learning. *arXiv Preprint arXiv:1710.02298*, 2017.

- Wang, Z., T. Schaul, M. Hessel, H. Hasselt, M. Lanctot, and N. Freitas. Dueling Network Architectures for Deep Reinforcement Learning. *Proc., International Conference on Machine Learning, PMLR*, New York, NY, 2016, pp. 1995–2003.
- Lillicrap, T. P., J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous Control with Deep Reinforcement Learning. *arXiv Preprint arXiv:1509.02971*, 2015.
- Ference, J. J. The Integrated Vehicle-Based Safety Systems Initiative. National Highway Traffic Safety Administration, Washington, D.C., 2006.
- Leurent, E. An Environment for Autonomous Driving Decision-Making. 2018. https://github.com/eleurent/high way-env.
- Brockman, G., V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. OpenAI Gym. arXiv Preprint arXiv:1606.01540, 2016.
- Dosovitskiy, A., G. Ros, F. Codevilla, A. Lopez, and V. Koltun. CARLA: An Open Urban Driving Simulator. arXiv Preprint arXiv:1711.03938, 2017.

The views presented in this paper are those of the authors alone.