

CaDRE: Controllable and Diverse Generation of Safety-Critical Driving Scenarios using Real-World Trajectories

Peide Huang¹ Wenhao Ding¹ Benjamin Stoler¹ Jonathan Francis^{1,2} Bingqing Chen² Ding Zhao¹

Abstract—Simulation is an indispensable tool in the development and testing of autonomous vehicles (AVs), offering an efficient and safe alternative to road testing. An outstanding challenge with simulation-based testing is the generation of safety-critical scenarios, which are essential to ensure that AVs can handle rare but potentially fatal situations. This paper addresses this challenge by introducing a novel framework, CaDRE, to generate realistic, diverse, and controllable safety-critical scenarios. Our approach optimizes for both the quality and diversity of scenarios by employing a unique formulation and algorithm that integrates real-world scenarios, domain knowledge, and black-box optimization. We validate the effectiveness of our framework through extensive testing in three representative types of traffic scenarios. The results demonstrate superior performance in generating diverse and high-quality scenarios with greater sample efficiency than existing reinforcement learning (RL) and sampling-based methods.

I. INTRODUCTION

Simulation plays a pivotal role in the domain of autonomous driving, serving crucial functions in both training and evaluation of autonomous vehicles (AVs) [1], [2], [3], [4]. Compared to the costly and time-consuming on-road testing, simulation offers efficient feedback to developers and avoids risky engagements in the real world [5].

It is widely recognized that traffic scenarios in the real world exhibit a long-tail distribution, where normal scenarios constitute the majority and safety-critical ones occur infrequently [6], [7]. AVs predominately trained on normal scenarios may not generalize to safety-critical ones, which can lead to fatal accidents when deployed at scale. Furthermore, evaluating AVs in such a long-tail distribution of naturalistic driving scenarios is neither sample-efficient nor comprehensive. Consequently, it is essential to generate safety-critical scenarios for both simulation-based training and evaluation.

There are three principal challenges in generating safety-critical scenarios. The first challenge is *realism*, that is, the scenarios generated should be sufficiently realistic to occur in the real world. This realism is often quantified by the similarity between real-world scenarios and generated ones [8], [9], [10]. To achieve this, existing methods either modify existing scenarios [11], [12] or employ generative models to approximate the distribution of real-world scenarios and then modify the generated samples [13], [14]. The second challenge is *diversity*, that is, the generation algorithm should cover a wide spectrum of scenarios rather

than focusing on a single type of failure. Existing methods either lack in diversity or achieve it by construction, e.g., by specifying a multi-modal generative model [6]. The final challenge is *controllability*, that is the scenarios are generated according to specified characteristics. These guidelines are often expressed through constraints [15], temporal logic [13], or language [14], [16], all of which require the framework to be specially designed to facilitate their integration.

To tackle these challenges, we reformulate safety-critical scenario generation within the Quality-Diversity(QD) framework and introduce our method CaDRE. QD is a branch of optimization that finds a collection of high-performing yet qualitatively different solutions [17], [18]. Through a novel design, CaDRE addresses the aforementioned challenges by integrating information from real-world data, domain knowledge, and black-box optimization. To maintain *realism*, we parameterize scenarios as bounded perturbations to real-world trajectories. To promote *diversity*, we improve on an existing QD algorithm and use it to efficiently explore and optimize the continuous parameter space. To achieve *controllability*, we retrieve archived scenarios according to the desired measure values.

The main contributions can be summarized below:

- We propose CaDRE, a novel QD formulation for safety-critical scenario generation, which achieves realism, diversity, and controllability.
- We propose an occupancy-aware restart (OAR) mechanism as a general extension to the QD algorithm family, which improves the exploration efficiency.
- We conduct experiments on three representative real-world traffic scenarios: unprotected cross-turn, high-speed lane-change, and U-turn, which demonstrate that CaDRE generates diverse and high-quality scenarios with better sample efficiency compared to both RL- and sampling-based methods.

II. RELATED WORK

Safety-critical Scenario Generation. One significant component of autonomous driving simulation is the traffic model, which governs the behavior of background vehicles. Many works, e.g., TrafficSim [19], TrafficGen [8], ScenarioNet [20], focus on generating realistic scenarios in the naturalistic distribution. In contrast, we focus on the long-tail of the distribution, that is the safety-critical scenarios, to provide efficient and comprehensive evaluations of the safety of AVs [21], [22], [23].

A majority of such works focuses on adversarial generation. L2C [24], MMG [25], and CausalAF [26] generate initial

*This work was partially performed during PH’s internship at the Bosch Center for Artificial Intelligence.

¹PH, WD, BS, JF, DZ are with Carnegie Mellon University, USA. ²JF, BC are with Bosch Center for Artificial Intelligence. Contact: peideh@andrew.cmu.edu

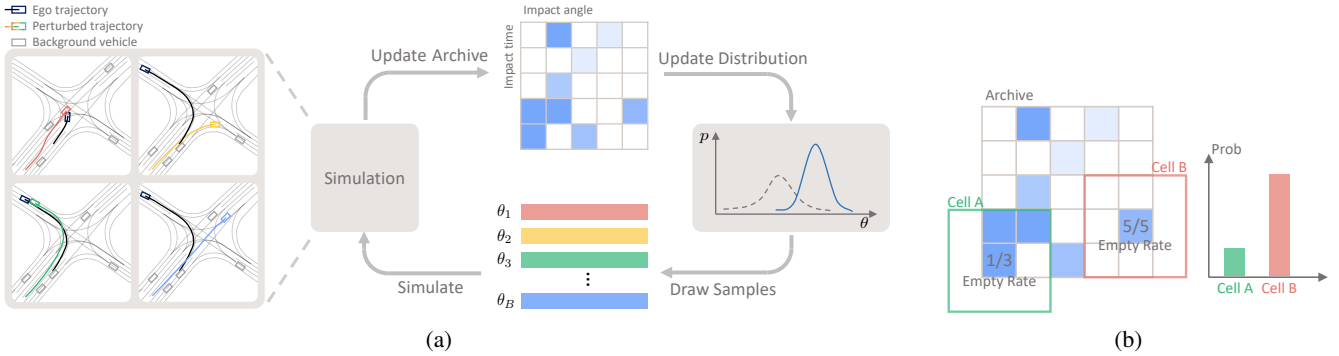


Fig. 1: (a) Overview of the CaDRE framework. CaDRE utilizes a black-box optimization algorithm to explicitly optimize for high-quality and diverse safety-critical scenarios. (b) Illustration of the Occupancy-Aware Restart (OAR) mechanism, a general extension to QD algorithms to improve exploration efficiency.

conditions for open-loop scenario generation using RL. The methods in [11], [27], [28], [29] optimize the trajectories of actors with black-box optimization to attack the ego vehicle. KING [12] and AdvDO [30] further assume access to differential dynamics models to improve the efficiency of finding safety-critical scenarios. However, adversarial attack-based methods do not explicitly optimize for the diversity and controllability of generated scenarios. To address this limitation, imitation learning [31], retrieval-augmented generation [32], causality [33], and evolutionary algorithms [34] have also been explored. To control the scenario generation with language, LCTGen [14] predefines an intermediate representation to bridge the large language model (LLM) [35] and the low-level trajectory generator, and CTG++ [36] uses LLMs to generate signal temporal logic to guide the sampling process of diffusion models. In this paper, we depart from the common practice of leveraging adversarial generation methods and instead focus on improving the diversity and controllability of generated samples through our novel application and extensions of Quality-Diversity algorithms.

Quality-Diversity Algorithms in Robotics. QD is a branch of optimization that finds a collection of high-performing, yet qualitatively different solutions [17], [18]. Specifically, QD optimizes an objective for each point in a measure space. Solving a QD problem in a continuous measure space requires infinite memory [18]. Thus, in practice, the measure space is discretized into a finite set, and an archive is maintained to keep track of the best-known solutions over the finite set. Some examples of popular algorithms include MAP-Elites [17], CMA-ME [37], and DQD [38].

Given QD’s ability to find a collection of high-performing solutions for different contexts, it is well-suited for many robotics applications. In [39], QD is used to learn a behavior-performance map, enabling the robot to quickly find a compensatory behavior and adapt after damage. QD has also been used on problems such as human-robot interaction [40], [41], robot manipulation [42], [43], locomotion [44], [45], and morphology design [46], [47]. To the best of our knowledge, we are the first to formulate QD for safety-critical scenario generation in autonomous driving; our approach enables us to generate an archive of diverse and high-quality scenes, and allows for fine-granular control by retrieving scenarios with

desired characteristics as defined by the measure space.

III. METHODOLOGY

Our method, **C**ontrollable and **D**iverse Generation of Safety-Critical Driving Scenarios using **RE**al-world trajectories (CaDRE), integrates real-world data, domain knowledge, and black-box optimization. As illustrated in Fig. 1, for each iteration, CaDRE maintains a grid archive of generated scenarios. First, it uses the QD algorithm to update the distribution from which the perturbations to the real-world trajectories are sampled. Then CaDRE simulates the perturbations to obtain diverse behavior measures and updates the archive according to the simulation results. Finally, we obtain an archive that contains thousands of critical scenarios, each with different behaviors according to the measure functions we defined using domain knowledge.

A. Representation of Traffic Scenarios

We utilize real-world trajectories and design the parameterization of traffic scenarios to retain *realism*. Let $\mathbf{x}_t^i \in \mathbb{R}^2, \psi_t^i \in [-\pi, \pi]$ and $v_t^i \in \mathbb{R}$ be the world-frame coordinate, orientation, and speed of the i -th vehicle agent at time t . The ego vehicle, with index $i = 0$, is the vehicle for which we want to generate critical scenarios. We denote the state of the vehicle as $\mathbf{s}_t = \{\mathbf{x}_t^i, \psi_t^i, v_t^i\}_{i=0}^N$, where N is the number of background agents. We define a specific traffic scenario as a sequence of these states $\mathcal{S} = \{\mathbf{s}_t\}_{t=0}^T$, where T is a fixed time horizon. We initialize a specific scenario from a real-world dataset that contains only naturalistic driving scenarios.

Safety-Critical Perturbation. In each generated scenario, we perturb the trajectory of a single background vehicle indexed by $i \in [1, \dots, N]$, which is selected according to criteria to be described in Section IV-A. We first recover the action sequence $\{\mathbf{a}_t^i\}_{t=0}^{T-1}$ from $\{\mathbf{s}_t^i\}_{t=0}^T$, assuming a kinematic bicycle model $[\dot{x}, \dot{y}, \dot{\psi}, \dot{v}] = [v \cos(\psi), v \sin(\psi), v \tan(\psi)/L, a]$, where L is the wheelbase. Each action consists of acceleration and steering input, i.e., $\mathbf{a}_t^i := [a_t^i, \delta_t^i]$. A new trajectory can be generated by 1) applying a sequence of bounded perturbations $\{\Delta \mathbf{a}_t^i\}_{t=0}^{T-1}$ to the recovered action sequence, and 2) unrolling the kinematics model from \mathbf{s}_0^i using the bicycle model. We then parameterize each safety-critical scenario with $\boldsymbol{\theta} = \{\Delta \mathbf{a}_0^i, \dots, \Delta \mathbf{a}_{T-1}^i\} \in \mathbb{R}^{T \times 2}$ and denote the allowable space of bounded perturbation as Θ .

B. Quality-Diversity Formulation for Scenario Generation

Inspired by previous work [40], [18], we formulate the problem of generating a diverse set of safety-critical driving scenarios as a QD problem. First, we define an objective function $f : \mathbb{R}^{T \times 2} \rightarrow \mathbb{R}$ to quantify the safety-critical level. We further define K measure functions $m_k : \mathbb{R}^{T \times 2} \rightarrow \mathbb{R}$ and jointly represent them as $\mathbf{m} : \mathbb{R}^{T \times 2} \rightarrow \mathbb{R}^K$, which are a set of user-defined functions to quantify aspects of the scenario that we aim to diversify. We denote $\mathcal{M} = \mathbf{m}(\mathbb{R}^{T \times 2}) \subseteq \mathbb{R}^K$ as the measure space formed by the range of \mathbf{m} . Because f evaluates the quality of a scenario θ , the goal of the QD problem is to find, for each $m \in \mathcal{M}$, a scenario θ , such that $\mathbf{m}(\theta) = m$ and that $f(\theta)$ is maximized (Eqn. 1):

$$\begin{aligned} \max \quad & f(\theta) \\ \text{s.t.} \quad & \mathbf{m}(\theta) = m, \forall m \in \mathcal{M}. \end{aligned} \quad (1)$$

In practice, we discretize \mathcal{M} into a finite number of M cells, which we will refer to as the archive, and solve the simplified version of the problem:

$$\max_{\theta_1, \dots, \theta_M} \sum_{n=1}^M f(\theta_n). \quad (2)$$

With a slight overload of notation, we also use f to denote the objective value and \mathbf{m} to denote the values of the measure function. We denote the archive as M , and we can retrieve the scenarios from the archive by $M[\mathbf{m}]$. With properly defined objective and measure functions, we can optimize a diverse population of safety-critical scenarios and retrieve individual scenarios in a controllable manner by asking for specific measure values \mathbf{m} . We build a lightweight traffic scenario simulator $\text{Sim}(\mathcal{S}, \theta, i)$, which outputs the objective value f and the measure values \mathbf{m} , given the original scenario \mathcal{S} , perturbation θ , and the index of the perturbed vehicle i .

For this setting, the typical formulation of safety-critical scenario generation would have been $\max_{\theta \in \Theta} f(\theta)$. To promote diversity, existing methods either parameterize the scenarios as a multi-modal distribution or impose an additional distributional constraint/regularization. In comparison, the QD formulation incorporates diversity directly as part of the problem.

We specify the objective and measure functions here to demonstrate how concretely QD can be used to generate realistic, diverse, and controllable safety-critical scenarios. There could be many other alternative objective and measure function designs, and we hope other researchers will be inspired to adopt CaDRE for their problems.

Objective Function. The objective function f quantifies the safety-critical level, motivated by prior work on safety-critical scenario generation (where $d(\cdot, \cdot)$ is the l_2 distance):

$$f(\theta) := \begin{cases} 1, & \text{if vehicle } i \text{ collides with the ego vehicle} \\ 0, & \text{if vehicle } i \text{ collides with background vehicles} \\ \exp(-\min_t d(\mathbf{x}_t^0, \mathbf{x}_t^i)), & \text{otherwise} \end{cases} \quad (3)$$

Measure Functions. The measure functions are essential to capture different aspects of critical scenarios. We propose three measure functions to define the diverse behavior of perturbed vehicles. These measure functions collectively enable the definition and evaluation of a wide range of safety-critical

Algorithm 1: CaDRE: Controllable and Diverse Generation of Safety-Critical Driving Scenarios

Input: Real-world normal scenario \mathcal{S} , index of the perturbed background vehicle i , traffic simulator Sim , batch size B , empty grid archive M

Output: An grid archive M containing diverse safety-critical scenarios \mathcal{S}_c

Initialize emitter e

Recover $\{\mathbf{a}_t^i\}_{t=0}^{T-1}$ from $\{\mathbf{s}_t^i\}_{t=0}^T$ in \mathcal{S}

for $\text{iter} = 1, \dots, \text{total_iter}$ **do**

$\{\theta_b\}_{b=1}^B \sim \mathcal{N}(e.\mu, e.C)$

for $b = 1, \dots, B$ **do**

$\{f_b, \mathbf{m}_b\} \leftarrow \text{Sim}(\mathcal{S}, \theta_b, i)$

Unpack *parents*, sampling mean μ , covariance matrix C , and parameter set P from e .

for $b = 1, \dots, B$ **do**

if $M[\mathbf{m}_b]$ is empty **then**

$\Delta_b \leftarrow f_b$

Flag that θ discovered a new cell

Add θ_b to *parents*

else if $f_b > M[\mathbf{m}_b].f$ **then**

$\Delta_b \leftarrow f_b - M[\mathbf{m}_b].f$

Add θ_b to *parents*

if *parents* $\neq \emptyset$ **then**

Sort *parents* by (newCell, Δ_b)

Update μ, C, P according to *parents*

parents $\leftarrow \emptyset$

else

Occupancy-aware restart from an elite in M

scenarios, focusing on representative factors such as perturbation efforts (m_1), urgency of response (m_2), and collision behavior (m_3). Here, m_1 measures the mean magnitude of the steering perturbation. It partially reflects how much the generated trajectory would deviate from the original trajectory:

$$m_1 = \frac{1}{t_{\text{impact}}} \sum_{t=0}^{t_{\text{impact}}-1} |\delta_t^i|, \quad (4)$$

Next, m_2 measures the normalized impact time; it helps categorize scenarios based on the urgency of the response required. When there is no collision, we assume $t_{\text{impact}} = \arg \min_t d(\mathbf{x}_t^0, \mathbf{x}_t^i)$. This measure aids in the development of time-critical decision-making algorithms for AVs:

$$m_2 = t_{\text{impact}}/T, \quad (5)$$

Finally, m_3 measures the impact angle relative to the body frame of the ego vehicle. It allows for the evaluation of how well autonomous driving systems can recognize and react to threats from various directions, enhancing their ability to prevent accidents through appropriate maneuvering or braking:

$$m_3 = \text{atan2}(R_{\psi_t^0}(\mathbf{x}_t^i - \mathbf{x}_t^0)^T), \quad (6)$$

where $R_{\psi_t^0}$ is the rotation matrix, $t = t_{\text{impact}}$, and atan2 is the 2-argument arctangent function. The x -axis of the body frame points to the vehicle's front, and the y -axis points to the left.

C. CaDRE Algorithm

We summarize the algorithm of CaDRE in Algorithm 1. We modify a QD algorithm, namely Covariance Matrix Adaptation MAP-Elites (CMA-ME) [37], by incorporating the domain-specific formulation (as described in Section III-A, III-B) and introducing OAR mechanism as a general extension to QD algorithms with discretized grid archives.

Algorithm Overview. The algorithm begins with the initialization of an archive and a population of solutions, which are then evaluated using *Sim* and assigned to specific cells within the archive based on their measure values. The core of the algorithm involves iteratively selecting elite solutions according to 1) whether new cells are discovered (*exploration*) and 2) how much the new solutions' objectives improve upon the old solution (*exploitation*), which guides the adaptation of the CMA parameters, such as the mean μ and covariance C . These parameters shape the distribution from which new solutions are sampled. The algorithm proceeds by replacing less optimal solutions in the archive with better-performing ones, thereby refining the solution space exploration. This process continues until a predefined total iteration is reached, ultimately resulting in an archive that captures a diverse array of scenarios across the explored measure dimensions.

Occupancy-Aware Restart. Existing QD algorithms [37] restart from random regions in the search space when there is no improvement in the archive. This is not efficient due to the possibility of restarting in well-explored regions. To improve exploration efficiency, we propose Occupancy-Aware Restart (OAR), a restart mechanism to serve as a form of guidance for improved coverage and efficiency during exploration. .

Illustrated in Fig. 1(b), OAR considers the occupancy rate of neighboring cells and assigns a higher probability to elites with more empty neighboring cells. More specifically, given the neighbor empty rate of L elites r_1, \dots, r_L and the temperature τ , the softmax probability of restarting from elite i is computed by $p_i = \frac{e^{r_i/\tau}}{\sum_{j=1}^L e^{r_j/\tau}}$. OAR degenerates to the uniform sampling as $\tau \rightarrow \infty$. With a lower τ , OAR assigns a higher probability to those elites who have more empty neighbors. For an efficient implementation, we use a 3D convolution kernel to compute the number of empty cells around each elite.

IV. EXPERIMENTS

A. Experimental Setup

Real-world Trajectories. We pick three representative scenarios from nuPlan v1.1 [48]: *unprotected cross-turn*, *high-speed lane-change*, and *U-turn*. All scenarios have a total time horizon of 15 seconds.

Reactive Ego Policy. Following [12], [30], we assume that the ego vehicle is reactive to nearby vehicles. We implement a rule-based ego policy: The ego vehicle will try to follow the reference trajectory, and if there is a vehicle within $5m$ and $[-\pi/4, \pi/4]$ of the body frame of the ego vehicle, it will brake at $-7m/s^2$ and steer away from that vehicle with maximum steering angle $\pm\pi/8$. Since the problem is formulated as a black-box optimization, CaDRE does not require access to the reactive policy and would work with any other ego policies.

Selection of Perturbed Vehicles and Perturbation Range.

To ensure effective perturbation, we use a simple heuristic to select which vehicles to perturb: the top five background vehicles that have the smallest average distance to the ego vehicle. Acceleration perturbation is between ± 2 , and steering perturbation range is between $\pm\pi/8$.

Evaluation Metrics. We focus on three criteria that measure the quality and diversity of the generated scenarios, which are standard metrics in the QD literature [37], [18], [38], [49].

- **Coverage** $\in [0, 1]$: Proportion of cells in the archive that have an elite. This is a measure of *diversity* of the generated scenarios. Note that while coverage is a standard metric for QD, there lacks such a principled metric to evaluate *diversity* in safety-critical scenario generation literature.
- **Mean objective** $\in [0, 1]$: Mean objective value of elites in the archive.
- **QD score** $\in [0, 4000]$: Sum of the objective values of all elites in the archive. The theoretical maximum value of 4000 is due to our objective $f \in [0, 1]$, and we discretize the measure space into $10 \times 20 \times 20$ grid.

Baselines. We study a mixture of sampling- and RL-based methods that have been employed by existing literature.

- **Random Search (Random)**: uniformly-random sample from the solution space.
- **GOOSE** [50]: a state-of-the-art RL-based approach for safety-critical scenario generation. It parameterizes adversarial trajectories as non-uniform rational B-splines (NURBS) and iteratively modifies their control points for diverse scenario generation. Note that this is also an example of diversity by construction, where diversity is promoted through the representation of scenarios.
- **CMA-ES** [51]: CMA-ES is an evolutionary strategy that iteratively updates a population of solutions based on their fitness and adaptively adjusts the search distribution towards optimal regions of the solution space.
- **Multi-particle REINFORCE** [52]: policy gradient method employed by previous work [24], [25]. We set the number of particles to be the same as the batch size (36) of CMA-ME employed by our algorithm.
- **Stein Variational Policy Gradient (SVPG)** [53]: SVPG is an improved version of multi-particle REINFORCE by using a maximum entropy policy optimization framework that explicitly encourages diverse solutions and better exploration. The number of particles is the same as the batch size of CaDRE.

We adopt the QD algorithm library *pyribs* [18] to implement our framework. We aim to answer the following questions in our experimental study:

- How does CaDRE compare with baseline methods in terms of the evaluation metrics and sample efficiency?
- Is OAR effective in improving exploration?
- Can we retrieve diverse scenarios generated by CaDRE in a controllable manner?

TABLE I: The final performance of coverage, mean objective, QD score. We report the mean and variance over 5 perturbed vehicles for each scene. The QD score is shown in multiples of 1000.

Method	Unprotected cross-turn			High-speed lane-change			U-turn		
	Coverage (\uparrow)	Mean Obj (\uparrow)	QD Score (\uparrow)	Coverage (\uparrow)	Mean Obj (\uparrow)	QD Score (\uparrow)	Coverage (\uparrow)	Mean Obj (\uparrow)	QD Score (\uparrow)
Random	0.140±0.021	0.499±0.123	0.285±0.098	0.162±0.024	0.310±0.158	0.209±0.131	0.188±0.066	0.381±0.134	0.320±0.151
GOOSE	0.279±0.097	0.368±0.058	0.429±0.180	0.257±0.082	0.385±0.158	0.443±0.291	0.141±0.082	0.490±0.158	0.325±0.252
CMA-ES	0.182±0.033	0.672±0.076	0.489±0.090	0.163±0.018	0.540±0.143	0.347±0.086	0.228±0.118	0.447±0.228	0.502±0.253
REINFORCE	0.210±0.031	0.649±0.115	0.551±0.155	0.243±0.033	0.488±0.144	0.472±0.161	0.286±0.010	0.641±0.117	0.731±0.117
SVPG	0.219±0.032	0.607±0.130	0.553±0.155	0.265±0.026	0.410±0.129	0.437±0.166	0.290±0.020	0.577±0.120	0.665±0.110
CaDRE (ours)	0.565±0.054	0.829±0.062	1.884±0.309	0.541±0.079	0.627±0.142	1.375±0.436	0.542±0.094	0.793±0.073	1.699±0.219




Fig. 2: Coverage and QD score v.s. number of samples. Solid lines represent the mean, and the shaded area presents the standard deviation over 5 perturbed vehicles.

B. Sample-efficiency Compared with Baseline Methods

The coverage and QD score v.s. samples are shown in Figure 2. CaDRE outperforms all baselines with significant margins in three different scenarios, which demonstrates that CaDRE discovers both high-quality and diverse scenarios more sample-efficiently than Random Search, GOOSE, SVPG, and REINFORCE. Under the hood, CaDRE adapts the search distribution over generations to increase the likelihood of sampling in promising areas of the solution space, in contrast to Random Search, which samples uniformly across the solution space. CMA-ES maximizes the likelihood of increasing the objective and thus quickly converges to a single optimum. SVPG and REINFORCE, while exploring more than CMA-ES, still struggle with exploring complex problem spaces due to their focus on gradient-based optimization. Finally, GOOSE, a dedicated method for safety-critical scenario generation, claimed that the NURBS parameterization of traffic scenarios is sufficient for *diversity*. However, it is clear from the results that its diversity, as measured by coverage, lags far behind CaDRE, though being competitive with other baselines.

Table I shows the final performance of coverage, mean objective, and QD score. With the same number of samples, CaDRE achieves 158.0%, 122.6%, 89.5% more coverage, 36.6%, 28.5%, 23.7% higher mean objective, leading to a 240.7%, 191.3%, 132.4% improvement in QD score than the best-performing baselines in three representative scenarios, respectively. It again highlights the superior exploration and exploitation capability of CaDRE compared to the baselines.

TABLE II: QD score of occupancy-aware restart with different temperatures. The QD score is shown in multiples of 1000. The percentage improvements w.r.t $\tau \rightarrow \infty$ are in parentheses.

Index	1	2	3	4	5
$\tau \rightarrow \infty$	1.808	0.519	1.277	1.393	1.444
$\tau = 1/5$	1.855 (2.6%)	0.639 (22.9%)	1.133 (-11.3%)	1.515 (8.8%)	1.299 (-10.0%)
$\tau = 1/10$	2.026 (12.0%)	0.691 (33.0%)	1.229 (-3.7%)	1.367 (-1.9%)	1.563 (8.2%)

Table II shows the ablation of OAR in the high-speed lane-change scenario. $\tau \rightarrow \infty$ is equivalent to QD algorithms without OAR. OAR improves the QD score of individual vehicles by a maximum of 33.0%, which demonstrates the effectiveness of OAR. We also observe that the biggest improvements come from harder problems (low QD score). The intuition is that when it is harder to find a safety-critical scenario, then a more effective exploration strategy provides a more significant performance boost.

C. Analysis of the Generated Safety-Critical Scenarios

Visualization of Archives. We visualize the final archives in Fig. 4. It is observed that the proposed CaDRE leads to a much higher occupancy as well as mean objective in the final archives than the best-performing baseline SVPG. The main reason is that CaDRE explicitly encourages sustained exploration throughout optimization. Although the repulsive force in SVPG indeed introduces diversity among the particles to avoid premature convergence to local optima, the primary focus remains on optimizing a solution rather than explicitly seeking out diverse solutions across a range of measures.

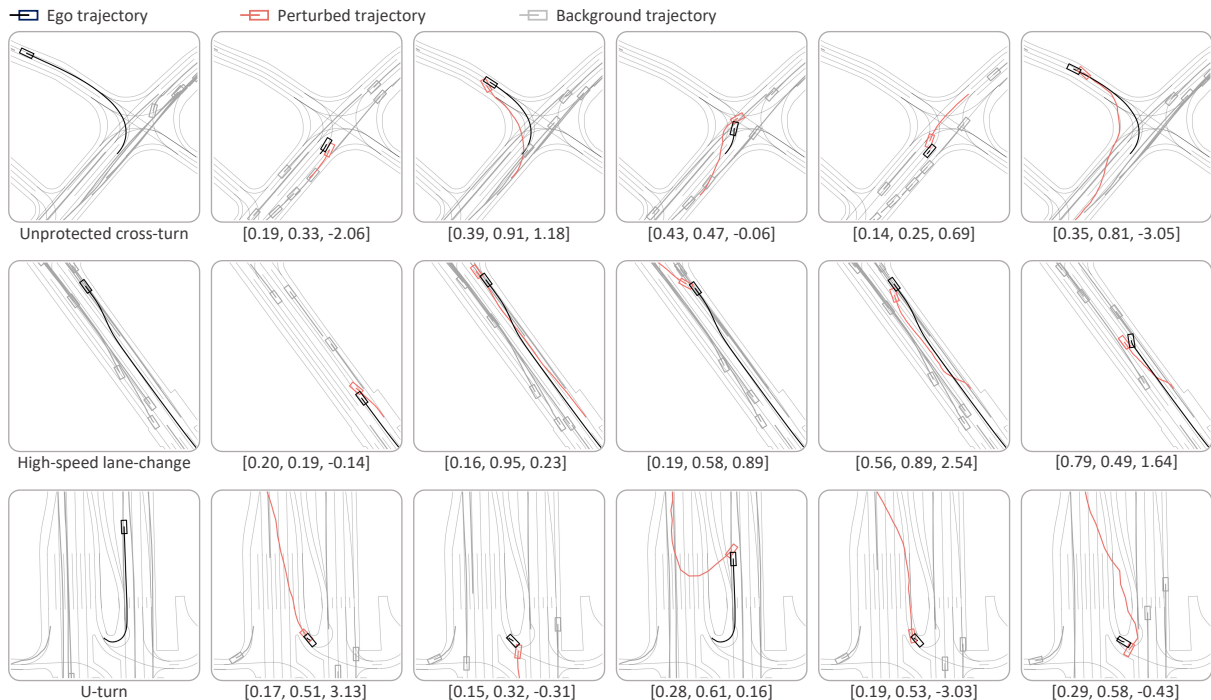


Fig. 3: Visualization of generated trajectories. The leftmost column shows the original unperturbed scenarios. The numbers below are the measure values $[m_1, m_2, m_3]$, representing the mean steering perturbation, impact time, and impact angle, respectively.




Fig. 4: Visualization of final archives. A darker color means a higher objective value. Transparent cells mean we cannot find scenarios. We visualize the perturbed vehicles that have the highest QD score for each scenario respectively.

Note that some cells are still unoccupied even for CaDRE. We hypothesize that it is due to the infeasibility of finding scenarios, which is induced by specific combinations of measure values, the vehicle states in the original scenarios, and the kinematics constraints. For example, it is extremely difficult to find a solution with a short impact time and a small impact angle (hitting from the front) in the unprotected cross-turn scenario, since there is no background vehicle starting

near the front of the ego vehicle.

Visualization of Generated Scenarios. In Fig. 3, we visualize five generated scenarios for the unprotected cross-turn, high-speed lane-change, and U-turn, respectively. The visualization shows that we are able to generate and retrieve diverse critical scenarios in a controllable manner. For example, in the unprotected cross-turn, we can control the perturbed vehicle hitting the right side of the ego vehicle by steering a little bit from the original trajectory or hitting the left side by overtaking from the left, simply by asking for different combinations of measure function values $[m_1, m_2, m_3]$ in the archive.

V. CONCLUSIONS AND DISCUSSION

In this work, we develop a framework CaDRE for generating safety-critical scenarios. CaDRE enhances the diversity and controllability of the scenario generation process while retaining realism. We conduct extensive experiments on three representative scenarios: unprotected cross-turn, high-speed lane-change, and U-turn. The experimental results show that CaDRE can generate and retrieve diverse and high-quality scenarios with better sample efficiency compared to existing methods.

We envision many extensions possible to CaDRE. At the moment, CaDRE only perturbs one background vehicles at a time. Alternative parameterization of traffic scenarios or objective and measure functions can make CaDRE more general or suited to specific settings. Second, CaDRE does not consider the lane information and road conditions such as barriers. It is possible to incorporate road information to CaDRE by augmenting the existing simulator or adding a post-hoc pruning step based on the drivable area.

REFERENCES

- [1] C. Gulino, J. Fu, W. Luo, G. Tucker, E. Bronstein, Y. Lu, J. Harb, X. Pan, Y. Wang, X. Chen *et al.*, “Waymax: An accelerated, data-driven simulator for large-scale autonomous driving research,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [2] J. Herman, J. Francis, S. Ganju, B. Chen, A. Koul, A. Gupta, A. Skabelkin, I. Zhukov, M. Kumskey, and E. Nyberg, “Learn-to-race: A multimodal control environment for autonomous racing,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9793–9802.
- [3] S. H. Park, G. Lee, J. Seo, M. Bhat, M. Kang, J. Francis, A. Jadhav, P. P. Liang, and L.-P. Morency, “Diverse and admissible trajectory forecasting through multimodal context understanding,” in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*. Springer, 2020, pp. 282–298.
- [4] M. Xu, Z. Liu, P. Huang, W. Ding, Z. Cen, B. Li, and D. Zhao, “Trustworthy reinforcement learning against intrinsic vulnerabilities: Robustness, safety, and generalizability,” *arXiv preprint arXiv:2209.08025*, 2022.
- [5] P. Huang, X. Zhang, Z. Cao, S. Liu, M. Xu, W. Ding, J. Francis, B. Chen, and D. Zhao, “What went wrong? closing the sim-to-real gap via differentiable causal discovery,” in *Conference on Robot Learning*. PMLR, 2023, pp. 734–760.
- [6] W. Ding, C. Xu, M. Arief, H. Lin, B. Li, and D. Zhao, “A survey on safety-critical driving scenario generation—a methodological perspective,” *IEEE Transactions on Intelligent Transportation Systems*, 2023.
- [7] B. Stoler, I. Navarro, M. Jana, S. Hwang, J. Francis, and J. Oh, “Safeshift: Safety-informed distribution shifts for robust trajectory prediction in autonomous driving,” *arXiv preprint arXiv:2309.08889*, 2023.
- [8] L. Feng, Q. Li, Z. Peng, S. Tan, and B. Zhou, “Trafficgen: Learning to generate diverse and realistic traffic scenarios,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 3567–3575.
- [9] S. Tan, K. Wong, S. Wang, S. Manivasagam, M. Ren, and R. Urtasun, “Scenegen: Learning to generate realistic traffic scenes,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 892–901.
- [10] P. Huang, M. Xu, J. Zhu, L. Shi, F. Fang, and D. Zhao, “Curriculum reinforcement learning using optimal transport via gradual domain adaptation,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 10 656–10 670, 2022.
- [11] J. Wang, A. Pun, J. Tu, S. Manivasagam, A. Sadat, S. Casas, M. Ren, and R. Urtasun, “Advsim: Generating safety-critical scenarios for self-driving vehicles,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 9909–9918.
- [12] N. Hanselmann, K. Renz, K. Chitta, A. Bhattacharyya, and A. Geiger, “King: Generating safety-critical driving scenarios for robust imitation via kinematics gradients,” in *European Conference on Computer Vision*. Springer, 2022, pp. 335–352.
- [13] Z. Zhong, D. Rempe, D. Xu, Y. Chen, S. Veer, T. Che, B. Ray, and M. Pavone, “Guided conditional diffusion for controllable traffic simulation,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 3560–3566.
- [14] S. Tan, B. Ivanovic, X. Weng, M. Pavone, and P. Kraehenbuehl, “Language conditioned traffic generation,” *arXiv preprint arXiv:2307.07947*, 2023.
- [15] W. Ding, H. Lin, B. Li, K. J. Eun, and D. Zhao, “Semantically adversarial driving scenario generation with explicit knowledge integration,” *arXiv preprint arXiv:2106.04066*, vol. 1, 2021.
- [16] M. Xu, P. Huang, W. Yu, S. Liu, X. Zhang, Y. Niu, T. Zhang, F. Xia, J. Tan, and D. Zhao, “Creative robot tool use with large language models,” *arXiv preprint arXiv:2310.13065*, 2023.
- [17] J.-B. Mouret and J. Clune, “Illuminating search spaces by mapping elites,” *arXiv preprint arXiv:1504.04909*, 2015.
- [18] B. Tjanaka, M. C. Fontaine, D. H. Lee, Y. Zhang, N. R. Balam, N. Dennler, S. S. Garlanka, N. D. Klapsis, and S. Nikolaidis, “Pyribs: A bare-bones python library for quality diversity optimization,” in *Proceedings of the Genetic and Evolutionary Computation Conference*, ser. GECCO ’23. New York, NY, USA: Association for Computing Machinery, 2023, p. 220–229. [Online]. Available: <https://doi.org/10.1145/3583131.3590374>
- [19] S. Suo, S. Regalado, S. Casas, and R. Urtasun, “TrafficSim: Learning to simulate realistic multi-agent behaviors,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 10 400–10 409.
- [20] Q. Li, Z. Peng, L. Feng, C. Duan, W. Mo, B. Zhou *et al.*, “Scenarionet: Open-source platform for large-scale traffic scenario simulation and modeling,” *arXiv preprint arXiv:2306.12241*, 2023.
- [21] W. Ding, M. Xu, and D. Zhao, “Cmts: Conditional multiple trajectory synthesizer for generating safety-critical driving scenarios,” in *International Conference on Robotics and Automation (ICRA)*. IEEE, 2020.
- [22] W. Ding, W. Wang, and D. Zhao, “Multi-vehicle trajectories generation for vehicle-to-vehicle encounters,” in *2019 IEEE International Conference on Robotics and Automation (ICRA)*, 2019.
- [23] C. Xu, W. Ding, W. Lyu, Z. Liu, S. Wang, Y. He, H. Hu, D. Zhao, and B. Li, “Safebench: A benchmarking platform for safety evaluation of autonomous vehicles,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 25 667–25 682, 2022.
- [24] W. Ding, B. Chen, M. Xu, and D. Zhao, “Learning to collide: An adaptive safety-critical scenarios generating method,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 2243–2250.
- [25] W. Ding, B. Chen, B. Li, K. J. Eun, and D. Zhao, “Multimodal safety-critical scenarios generation for decision-making algorithms evaluation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1551–1558, 2021.
- [26] W. Ding, H. Lin, B. Li, and D. Zhao, “Causalaf: Causal autoregressive flow for goal-directed safety-critical scenes generation,” *arXiv preprint arXiv:2110.13939*, 2021.
- [27] M. Klischat and M. Althoff, “Generating critical test scenarios for automated vehicles with evolutionary algorithms,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 2352–2358.
- [28] M. Arief, Z. Huang, G. K. S. Kumar, Y. Bai, S. He, W. Ding, H. Lam, and D. Zhao, “Deep probabilistic accelerated evaluation: A robust certifiable rare-event simulation methodology for black-box safety-critical systems,” in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2021, pp. 595–603.
- [29] L. Zhang, Z. Peng, Q. Li, and B. Zhou, “Cat: Closed-loop adversarial training for safe end-to-end driving,” in *Conference on Robot Learning*. PMLR, 2023, pp. 2357–2372.
- [30] Y. Cao, C. Xiao, A. Anandkumar, D. Xu, and M. Pavone, “Advdo: Realistic adversarial attacks for trajectory prediction,” in *European Conference on Computer Vision*. Springer, 2022, pp. 36–52.
- [31] C. Zhang, J. Tu, L. Zhang, K. Wong, S. Suo, and R. Urtasun, “Learning realistic traffic agents in closed-loop,” in *7th Annual Conference on Robot Learning*, 2023.
- [32] W. Ding, Y. Cao, D. Zhao, C. Xiao, and M. Pavone, “Realgen: Retrieval augmented generation for controllable traffic scenarios,” *arXiv preprint arXiv:2312.13303*, 2023.
- [33] W. Ding, H. Lin, B. Li, and D. Zhao, “Generalizing goal-conditioned reinforcement learning with variational causal reasoning,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 26 532–26 548, 2022.
- [34] A. Li, S. Chen, L. Sun, N. Zheng, M. Tomizuka, and W. Zhan, “Scegen: Bio-inspired traffic scenario generation for autonomous driving testing,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 9, pp. 14 859–14 874, 2021.
- [35] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, “Gpt-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023.
- [36] Z. Zhong, D. Rempe, Y. Chen, B. Ivanovic, Y. Cao, D. Xu, M. Pavone, and B. Ray, “Language-guided traffic simulation via scene-level diffusion,” *arXiv preprint arXiv:2306.06344*, 2023.
- [37] M. C. Fontaine, J. Togelius, S. Nikolaidis, and A. K. Hoover, “Covariance matrix adaptation for the rapid illumination of behavior space,” in *Proceedings of the 2020 genetic and evolutionary computation conference*, 2020, pp. 94–102.
- [38] M. Fontaine and S. Nikolaidis, “Differentiable quality diversity,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 10 040–10 052, 2021.
- [39] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, “Robots that can adapt like animals,” *Nature*, vol. 521, no. 7553, pp. 503–507, 2015.
- [40] V. Bhatt, H. Nemlekar, M. Fontaine, B. Tjanaka, H. Zhang, Y.-C. Hsu, and S. Nikolaidis, “Surrogate assisted generation of human-robot interaction scenarios,” *arXiv preprint arXiv:2304.13787*, 2023.
- [41] Y.-S. Tung, M. B. Luebbbers, A. Roncone, and B. Hayes, “Workspace optimization techniques to improve prediction of human motion during human-robot collaboration,” *arXiv preprint arXiv:2401.12965*, 2024.
- [42] A. Morel, Y. Kunimoto, A. Coninx, and S. Doncieux, “Automatic acquisition of a repertoire of diverse grasping trajectories through behavior shaping and novelty search,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 755–761.

- [43] J. Huber, F. H el enon, M. Kappel, E. Chelly, M. Khoramshahi, F. B. Amar, and S. Doncieux, "Speeding up 6-dof grasp sampling with quality-diversity," *arXiv preprint arXiv:2403.06173*, 2024.
- [44] S. Surana, B. Lim, and A. Cully, "Efficient learning of locomotion skills through the discovery of diverse environmental trajectory generator priors," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 12 134–12 141.
- [45] J. Nordmoen, F. Veenstra, K. O. Ellefsen, and K. Glette, "Map-elites enables powerful stepping stones and diversity for modular robotics," *Frontiers in Robotics and AI*, vol. 8, p. 639173, 2021.
- [46] E. Zardini, D. Zappetti, D. Zambrano, G. Iacca, and D. Floreano, "Seeking quality diversity in evolutionary co-design of morphology and control of soft tensegrity modular robots," in *Genetic and Evolutionary Computation Conference*, 2021.
- [47] P. Liu, Z. Guo, H. Yu, H. Linghu, Y. Li, Y. Hou, H. Ge, and Q. Zhang, "A preliminary study of multi-task MAP-elites with knowledge transfer for robotic arm design," in *2022 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, jul 2022. [Online]. Available: <https://doi.org/10.1109%2Fcec55065.2022.9870374>
- [48] H. Caesar, J. Kabzan, K. S. Tan, W. K. Fong, E. Wolff, A. Lang, L. Fletcher, O. Beijbom, and S. Omari, "nuplan: A closed-loop ml-based planning benchmark for autonomous vehicles," *arXiv preprint arXiv:2106.11810*, 2021.
- [49] M. Fontaine and S. Nikolaidis, "A quality diversity approach to automatically generating human-robot interaction scenarios in shared autonomy," *arXiv preprint arXiv:2012.04283*, 2020.
- [50] J. Ransiek, J. Plaum, J. Langner, and E. Sax, "Goose: Goal-conditioned reinforcement learning for safety-critical scenario generation," in *2024 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2024.
- [51] N. Hansen, "The cma evolution strategy: A tutorial," *arXiv preprint arXiv:1604.00772*, 2016.
- [52] R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," *Machine learning*, vol. 8, pp. 229–256, 1992.
- [53] Y. Liu, P. Ramachandran, Q. Liu, and J. Peng, "Stein variational policy gradient," *arXiv preprint arXiv:1704.02399*, 2017.