

Safety-Critical Scenario Generation Via Reinforcement Learning Based Editing

Haolan Liu¹, Liangjun Zhang², Siva Kumar Sastry Hari³ and Jishen Zhao⁴

Abstract—Generating safety-critical scenarios is essential for testing and verifying the safety of autonomous vehicles. Traditional optimization techniques suffer from the curse of dimensionality and limit the search space to fixed parameter spaces. To address these challenges, we propose a deep reinforcement learning approach that generates scenarios by sequential editing, such as adding new agents or modifying the trajectories of the existing agents. Our framework employs a reward function consisting of both risk and plausibility objectives. The plausibility objective leverages generative models, such as a variational autoencoder, to learn the likelihood of the generated parameters from the training datasets; It penalizes the generation of unlikely scenarios. Our approach overcomes the dimensionality challenge and explores a wide range of safety-critical scenarios. Our evaluation demonstrates that the proposed method generates safety-critical scenarios of higher quality compared with previous approaches.

I. INTRODUCTION

The long-tail problem remains a significant hurdle in the development of autonomous vehicles (AVs) [16], [14]. To ensure safe navigation in real-world scenarios, autonomous vehicles need to handle a diverse range of rare and safety-critical corner cases. Due to the rarity of these scenarios in the real world, conducting effective and scalable safety testing and verification is challenging. Furthermore, as modern autonomy stacks increasingly rely on data-driven approaches [8], [18], [13], they are becoming more vulnerable to corner cases that are poorly represented in training data.

Automatically generating safety-critical scenarios is crucial for addressing the long-tail problem. Previous studies treated the problem of generating adversarial scenarios as an optimization problem over the parameter space, using a specific cost function [19], [23], [22]. The cost function is typically constructed based on factors such as closest distance or time-to-collision [17], [1]. The scenario parameters typically include the initial positions, orientations, trajectories of the traffic agents, lane topologies, and even different weathers. By considering a wide range of parameters, the generated scenarios can cover a diverse range of situations that are difficult to encounter in the real world, allowing for a more comprehensive AV safety validation.

As illustrated in Figure 1, consider a scenario where the blue vehicle (AV) intends to change lanes to the left. In this




Fig. 1. (1) Adding exploration into scenario generation (newly added agents) helps generate high-quality safety-critical scenarios. (2) Jointly optimizing all the traffic agents (including agent B, C, D) in the scenario is not necessary, and can even hinder the optimization process, resulting in suboptimal results. Our framework edits the scenario efficiently in an agent-wise manner and search for desirable scenarios by sequential exploration.

case, the scenario parameter is the trajectory of traffic agent A, B, C, D . Previous studies have applied methods such as black-box optimization to iteratively search for perturbations that minimize the cost function, which is typically based on the closest distance metric. An optimal trajectory is one where the agent A decelerates to reduce the closest distance, challenging the AV to safely change lanes.

However, such an approach only considers a fixed parameter space, which is the trajectory of the five agents, and does not take into account the diverse range of parameters that can affect the scenario. For instance, introducing a new agent into the scenario as in Figure 1, may not necessarily minimize the closest distance. However, such an "exploration" step will increase the scenario's complexity, making it easier to optimize in the long run. For example, agent B can perform a lane change maneuver that intensifies the danger of the scenario. Moreover, this example also shows that previous manual cost function designs, e.g. closest distance are not general enough to accurately describe the risk of the scenario.

Furthermore, these approaches also face the challenge of high dimensionality. The widely-used black-box optimization techniques can only be applied to low-dimensional parameter spaces, typically with about 20 dimensions. To address this challenge, another branch of research has proposed optimizing scenarios in the latent space of a generative traffic model, using gradient-based methods [19]. However, such

¹Haolan Liu is a Ph.D. student with University of California San Diego, 9500 Gilman Dr, La Jolla, California hal022@ucsd.edu

²Liangjun Zhang is with the Baidu Research, 1195 Bordeaux Drive Sunnyvale, California liangjun.zhang@gmail.com

³Siva Kumar Sastry Hari is with the NVIDIA Corporation, 2788 San Tomas Expressway, Santa Clara, California shari@nvidia.com

⁴Jishen Zhao is an associate professor at the University of California San Diego, 9500 Gilman Dr, La Jolla, California jzhao@ucsd.edu

an approach requires performing joint optimization across all agents in the scenario, which may yield suboptimal results.

To address the aforementioned challenges, we propose a scenario editing framework based on reinforcement learning, to generate safety-critical scenarios. Specifically, we employ a reinforcement learning agent to sequentially construct the scenario, by taking actions such as adding new agents or modifying the trajectory of existing agents. Our framework naturally supports exploration in scenario generation and automatically selects the most relevant traffic agents for generating safety-critical scenarios.

Our framework is trained to optimize a reward that consists of the risk objective and plausibility objectives. The risk objective quantifies the number of feasible driving plans for the AV, offering a detailed description of the risk in the scenario. The plausibility objectives apply a generative model to learn the likelihood of the generated parameters from the training datasets and penalize the generation of unlikely scenarios. Besides, our approach can automatically identify the safety-critical agents.

Our approach leverages deep reinforcement learning (DRL) to overcome the curse of dimensionality that hampers traditional optimization techniques. Our method also allows us to generate scenarios with varying lengths of parameters, whereas previous approaches are restricted to fixed parameter spaces. The exploration-exploitation tradeoff in modern RL algorithms also facilitates the effective exploration of possible safety-critical scenarios.

We summarize our major contributions as follows:

- We propose a reinforcement learning-based scenario editing framework, which can generate a diverse range of safety-critical scenarios by supporting operations such as adding agents, perturbations, and resampling. Our framework sequentially edits the scenario and efficiently explores the feasible scenarios.
- We design a generative model to learn the plausibility constraints of the scenario and propose a novel anchor-based risky score to describe desirable scenarios in detail.
- We perform a comprehensive analysis on our framework. Experiments demonstrate that the framework generates a diverse range of high-quality safety-critical scenarios.

II. RELATED WORK

a) Adversarial Generation: Adversarial scenario generation aims to identify scenarios with high risk by exploring various parameters based on a specified cost function. The cost function defines the ideal attributes of the scenarios, considering factors like the distance to nearby agents and the presence of undesirable sideways collisions, like sideways contact, as shown in [1]. As this type of cost function can only be evaluated within non-differentiable simulators, previous research has used black-box optimization methods, such as Bayesian optimization (BO) [1] and evolutionary algorithms. However, these methods can become computationally demanding as the number of dimensions grows.

b) Generative Models: Our work is also related to the generative model. These models generate previously unseen scenarios by sampling from a model distribution that is similar to the data distribution. For example, DESIRE and HMSIM trains a conditional variational autoencoder to generate scenario-consistent trajectories for traffic agents [10], [12]. Similarly, TrafficGen trains an autoregressive model to sequentially insert new traffic agents [5].

However, since safety-critical scenarios are rare in real-world datasets, the scenarios generated by the generative models are often normal and safe. To address this issue, previous work trained an ensemble of learned policies with generative adversarial imitation learning (GAIL) and leveraged cross-entropy methods to upweight the sampling frequency of rare risky events [17]. Another approach, called STRIVE, generates new scenarios by encoding existing scenarios into latent space and optimizing the scenario based on a differentiable planning model [19].

c) Reinforcement Learning: The reinforcement learning approach trains an adversarial policy to control scenario agents, with the goal of facilitating collisions with AVs. For instance, [22] uses the Deep Deterministic Policy Gradient (DDPG) within a reinforcement learning framework to control surrounding agents and attack the AVs. The study designs an adversarial reward based on vehicle distance and adds a personal reward to discourage unrealistic agent behavior. The adversarial reward is allocated to each scenario agent based on their quantified contribution.

d) Structure Generation: Our work also relates to structured data generation (e.g., graph or set structured data), such as the architecture of the neural network [9], [15]. Such data exhibit conditional dependencies. For instance, the driving scenario can be defined by a large number of parameters, such as vehicle numbers, initial positions, and individual trajectories over time, which have complex interactions with each other.

Previous works have used Bayesian Optimization (BO) methods on a predefined Euclidean space to generate safety-critical scenarios. However, black-box optimization methods are not practical for highly structured parameterization spaces, rendering the optimization process difficult.

III. PROPOSED METHOD

A. Problem Definition

In this section, we formulate the scenario generation problem and discuss major challenges, as well as opportunities for improvement. Given a scenario x , plausibility constraints \mathbb{C} , and a scoring function $s(x)$, our objective is to determine the optimal scenario, denoted by x^* :

$$x^* = \underset{x \in \mathbb{C}}{\operatorname{argmin}} s(x) \quad (1)$$

The scoring function $s(x)$ represents the risk associated with driving scenario x ; the scenario generation process must optimize this risk function while ensuring that the constraint \mathbb{C} is not violated. An example of the scoring function $s(x)$ is the closest distance, as used by Abeyesirigoonawardena et

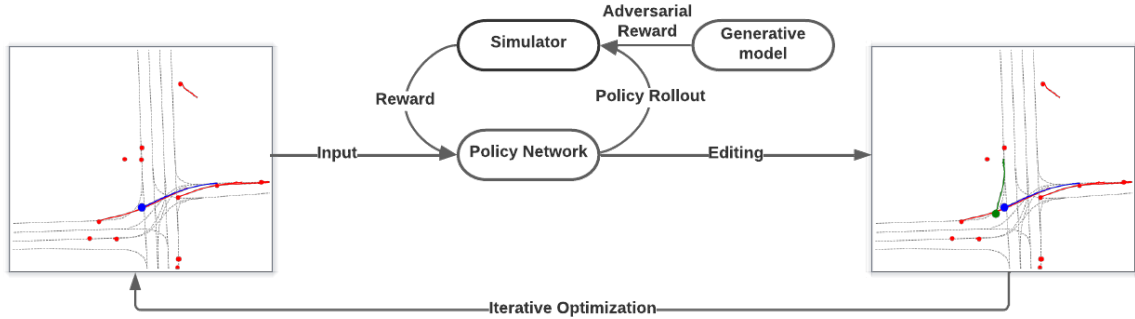


Fig. 2. Our scenario editor is based on reinforcement learning, with a policy network that optimizes the scenario iteratively. In each iteration, the network takes the current scenario as input and outputs a distribution of editing actions. To train the agent, we use the risk model and adversarial loss from the pretrained model as the reward. In this example, the policy network adds a new agent to the scenario. The circle indicates the endpoint for each agent’s trajectory. The blue, green, and red line indicate the trajectories of the ego vehicle, the newly added agent, and other agents in the scenario.

al. [1]. Minimizing the scoring function leads to closer agents in the scenario, increasing the possibilities of accidents.

Plausibility constraints \mathbb{C} are typically categorized into two categories. (1) Soft constraints, which aim to ensure that the generated traffic agents are reasonably realistic. However, these constraints may be violated in safety-critical scenarios, where irregularities are present. (2) Hard constraints, which impose structural requirements on the generated traffic agents. For instance, the trajectory of a dynamic agent, such as a vehicle should be physically feasible.

To specify \mathbb{C} , previous works add penalty terms to the optimization objectives, such as traffic rule violation [1], [23], [19]. Another approach is limiting the allowable range of parameters. e.g., Wang et al. applies local search from existing scenarios by perturbing the parameters, such as position and velocity profile, with limited magnitude [23]. But such constraints also make the objective function $s(x)$ complex and hard to optimize.

To solve the constrained optimization problem, most prior approaches rely on black-box optimization strategies. At each timestep i , the strategy M outputs d_i , which is iteratively improved to optimize the scoring function $s(x)$ [23], [1]. However, this approach works on a fixed and predefined parameter space and is subject to the curse of dimensionality. Moreover, there are no effective ways to explicitly control the plausibility of the output. To address these three challenges, we propose a scenario editing framework that utilizes deep reinforcement learning to edit the scenario, which is both flexible and scalable to high dimensions. Furthermore, we incorporate a pretrained generative model to explicitly define \mathbb{C} during generation.

B. Generation using RL

Overview: We represent the scenario as an unordered set (V, T) , where the traffic agent matrix $V \in \mathbb{R}^{n \times d_1}$ and the static context matrix $T \in \mathbb{R}^{m \times d_2}$ are feature matrices assuming each element has d_1 or d_2 features, respectively.

Our framework formulates the scenario generation as an agent making sequential modifications to the scenarios, which is well-suited for reinforcement learning. The scenario

editor iteratively edits the scenario representation by modifying the agent matrix V or adding new features. Figure 2 shows the process. At each time step, the policy network $\pi_{\theta} = p(a_t | s)$ takes the current scenario (V, T) and outputs an edit operation a . We then modify the scenario based on a and compute the reward R for the agent.

We iteratively edit the scenario representation (V, T) as a Markov Decision Process (MDP) $M = (S, A, P, R)$, where S is a set of states of possible scenarios, A is a set of actions that modify the scenario at each timestep. P is the transition dynamics that describe how the scenario changes given an action. R is the specified reward function to describe the desired properties of the generated scenario.

Our framework aims to optimize the expected cumulative reward $\sum_{t=0}^{\infty} \gamma^t r_t$, where t denotes each timestep and r_t denotes its reward. γ is the discount factor that determines whether to prioritize the immediate return over the long-term return. We set $\gamma = 1$ in our framework. In addition, the maximum episode length L_{max} bounds the number of edits to avoid overly complex scenarios.

a) State space: The state S represents a fully observable scenario (V, T) where T indicates the context, including details such as lanes, intersections, and traffic lights. V indicates editable factors in the scenario, such as traffic agents’ position and velocity.

b) Action space: The action describes how the scenario editor modifies the scenario. Figure 3 shows the action that is supported by our framework. Specifically, we support the following actions:

- **PerturbAgent:** perturb the trajectory of one of the existing agents based on its history and interaction with other agents. The continuous field includes the id of the selected agent and the perturbation vector to the original trajectory.
- **AddAgent:** regressing the position, orientation, and trajectory of the newly added agents (we implemented vehicle only at this stage).
- **Reconfigure:** resampling a trajectory (including history and future trajectories) of the existing agent
- **Undrivable:** add an undrivable region to the map,




Fig. 3. The editing action supported by our framework.

which can represent parked vehicles or a construction area. The action also includes its coordinates. We use fixed-size regions.

- **Termination:** terminate the generation.

For generality, our action space is a discrete-continuous hybrid action space.

$$\hat{A} = (f_{1:K}, x_{1:K}) | f_k \in \mathcal{F}, x_k \in \chi_k, k \in [K] \quad (2)$$

Here $\mathcal{F} \in \mathcal{R}$. f_k indicates the discrete meta-action, it can be determined using argmax or sampling based on softmax operations. Once the discrete action f is selected, e.g., adding a new agent, the rest of the action $x_{1:K}$ can be used to regress the initial position and trajectory of the generated agents. In this way, our framework can support a flexible set of scenario editing operations.

Consider a case when the action space only includes the perturbation, our framework degrades to the previous adversarial generation approach [1], [19], [23], which optimizes the representation (V, T) iteratively using black-box optimization. Therefore, our approach is a more generalized and expressive framework over prior work.

c) Transition Dynamics: With a given action a , the transition dynamics transform the scenario. We also perform necessary sanity checks for the generated action. For example, an action of adding a new agent, whose initial position collides with another vehicle, will be directly rejected.

d) Reward function: We apply two kinds of rewards to guide the search for safety-critical scenarios. (1) The final reward, which is only given after the termination actions, contains a risk model described in Section III-D. (2) Intermediate reward: the environment assigns an adversarial reward to penalize the unlikely behavior, by adding the data likelihood of each action given current scenarios. In addition, the environment also assigns a positive reward when the action passes the feasibility check, and assigns a negative reward otherwise.

C. Policy Network

The policy network takes an intermediate scenario and outputs a stochastic action distribution. At test time, we can




Fig. 4. The risk model evaluates the safety of the autonomous vehicle’s driving plans. In the given scenario, there are three predefined plans: one blue and two green. The blue anchor leads to a collision, while the two green anchors are safe. The risk model counts the number of feasible driving plans and uses this information to guide the scenario editing process.

sample or select the action with the maximum possibility. We leverage the graph neural network to compute the embedding of the scenario, as in VectorNet [6]. Specifically, the map and agent features (e.g., the lane geometry and the agent trajectory) are discretized into polylines, we use a graph neural network to extract polyline-level features and use a global self-attention layer to compute the embedding of the entire scenarios. We first pretrain the graph neural network, on the motion forecasting task, and later optimize each scenario. With the scenario embedding, we use a Multilayer Perceptron (MLP) network to compute the editing action, to generate the next scenario.

Our framework mainly employs Proximal Policy Optimization to train the policy network. Proximal Policy Optimization (PPO) is a state-of-the-art on-policy algorithm, that leverages surrogate objectives to avoid undesirable big updates in the policy [20]. We also pretrain the generative model for the adversarial reward, which will be further explained in III-E.

D. Risky Model

Previous studies have employed metrics, such as collision-or-not, time-to-collision, and closest distance, to quantify the risk associated with driving scenarios. However, these metrics are not always adequate in describing risk. For instance, in Figure 1, the addition of a new vehicle does not affect any of the aforementioned metrics, but it certainly makes the scenario more challenging to navigate correctly.

To overcome this limitation, we propose a risk model $r(x)$, which maps a scenario to a risk score based on the number of safe decisions available to the autonomous vehicle (AV). To achieve this, we discretize the feasible driving trajectory into M predefined anchors, which can be obtained through clustering the trajectories in the dataset or using predefined grid targets. This approach is also used in anchor-based motion forecasting methods [3]. Given a scenario x , we apply collision detection algorithms to the M trajectories and compute the number of safe anchors. Figure 4 illustrates an example where $M = 3$, and the predefined anchors include going straight, changing lanes to the right, and changing lanes to the left. We perform collision detection and determine that the two anchors shown in green are safe.

For the generated scenario to be meaningful for testing the AV, it is essential that the AV has a way to avoid the collision. To prevent the generation of scenarios that are not “solvable” [7], we have added a penalty term to the risk model in cases where no trajectory candidate is available.

E. Plausibility Control

We design an intermediate reward for each editing action, that penalizes unlikely behavior, such as generating a physically infeasible trajectory or placing a vehicle on the sidewalk. The previous approach designs heuristic rewards, such as penalty terms to the optimization objectives. Such heuristic design relies on human efforts and cannot scale to real-world complexities. We proposed an adversarial reward, that measures the deviation of the current scenario from the real-world scenarios. Such rewards are used to guide the learning of the scenario editor. In our paper, we implement the two kinds of adversarial reward, based on the data likelihood of the autoregressive model and the conditional variational autoencoder. The conditional variational autoencoder (CVAE) is trained to maximize the conditional log-likelihood, by practically optimizing the variational lower bound:

$$\hat{\mathcal{L}}_{CVAE}(\mathbf{x}, \mathbf{y}, \theta, \phi) = -KL(\mathbf{q}_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y})||\mathbf{p}_\theta(\mathbf{z}|\mathbf{x})) + \sum_{i=1}^L \log p_\theta(\mathbf{y}|\mathbf{x}, \mathbf{z}^{(i)}) \quad (3)$$

where L indicates the number of training samples, $z^{(i)}$ indicates each sample’s latent code. x and y respectively indicates the agent’s future trajectory and its context information (map and other agents’ history trajectories). The entire architecture includes a recognition network q_ϕ and a generation network p_θ . Our recognition network also uses graph neural networks as the feature extractor. We pretrain the CVAE model using real-world datasets and use the reconstruction error as the plausibility score to penalize unlikely editing. For certain actions such as `AddAgent`, we add a rule-based common sense loss that forces the newly added agents to stay on lanes, and not too close to other agents, which are proved useful by previous works [23], [19].

We also apply a piecewise smooth function to normalize the data likelihood \mathcal{L} , as it can become huge for some outliers. We split the reward value range into K bins, and apply an individual N_i ($i = 1 - K$) to normalize the reward. Such normalization can reduce the numerical instability for the training, and induce more smooth reward function.

IV. EXPERIMENTS

A. Experimental Setup

a) Dataset and Simulators: We evaluate the proposed algorithm in a highway simulator (highway-env), which supports highly customizable scenario setup and also basic physical simulation [11]. The evaluated scenarios include multi-lane, intersection, and lane merging. We use the Argoverse Motion Forecasting Dataset [4], which contains over 300,000 tracked scenarios that span five seconds each, with a specific vehicle identified for trajectory forecasting. The dataset includes high-definition maps covering 290 km of mapped lanes, with both geometric and semantic metadata. We build a simple simulator that can import the Argoverse

motion dataset, perform rollout of the AV planning policy, and also detect collision. The timestep frequency is 10 Hz.

b) Tested Systems: For our tested AV system, we use two kinds of planners, (a) Reinforcement learning agent in the highway-env, trained by Proximal Policy Optimization (PPO) [20] (b) an imitation learning agent, based on ChauffeurNet [2] that exhibits good driving performance comparable to humans. We evaluate both planners across various scenarios and compute the average metric results.

c) Metrics: To evaluate the quality of generated scenarios, we adopt a set of several metrics, which is a common practice used in previous works [23]. 1) **Collision rate**, which represents the percentage of scenarios in which the autonomous vehicle (AV) collides with another actor. This metric allows us to determine if the generated scenario can effectively identify potential safety issues with the AV planner. aggressive way. 2) **L2 distance** to the imitation learning planner’s trajectory as a measure of how closely the generated trajectories align with those of human experts. This metric helps us evaluate the deviation of the generated trajectories from human driving behavior. 3) The **Distribution JSD** metric. We compute the distribution of velocity, acceleration, and jerk (derivative of the acceleration) in the datasets, and discretize them into histograms. We compute the Jensen-Shannon Divergence between the generated scenarios and dataset scenarios [21]. This metric indicates the motion similarities between real-world scenarios and generated scenarios. 4) we also introduce the **Traffic Rules Violation** metric, quantifying the frequency of generating behaviors such as off-road driving.

B. Generation Quality Evaluation

We evaluate the generated scenario by running our tested AV planner in the scenario. To establish a baseline, we compared our work with Bayesian Optimization [1], AdvSim [23], and STRIVE [19]. STRIVE applies a preprocessing step to filter out scenarios that are unlikely to result in a collision [19]. To ensure a fair comparison, we did not apply this preprocessing step for any of the three algorithms. Using the same scenario from the dataset, we evaluated the performance of the four algorithms by optimizing until convergence. We also filter out some physically infeasible scenarios using the bicycle model and also rule-checking.

a) Risk Evaluation: We show the collision metric in Table I. We optimize the existing scenario from the dataset. Our experiments demonstrate that our approach effectively generates challenging scenarios for various systems, including our imitation learning and reinforcement learning planners. This results in an average collision rate of 7.8%, which is over 80% higher compared to the original state-of-the-art method, STRIVE.

b) Plausibility Evaluation: Table I also shows the plausibility metrics evaluated on the generated scenarios. We can see that the Bayesian Optimization (**BO** and **AdvSim**) approach exhibits unrealistic scenarios, compared with STRIVE and our approach, which can be mainly explained by the strong data prior in both models. Our approach

TABLE I
EVALUATION OF THE COLLISION RATE AND THE PLAUSIBILITY METRICS ACROSS PREVIOUS WORKS AND OUR WORK.

	Risk		Plausibility			
	Safety	Dataset Reconstruction	Distribution JSD			Common Sense
	Collision Rate(Percent) \uparrow	L2 Distance \downarrow	Velocity \downarrow	Accel \downarrow	Jerk \downarrow	Traffic Rules Violation(Percent) \downarrow
BO[1]	1.1	9.81	0.83	0.61	0.62	8.4
AdvSim[23]	3.7	3.44	0.65	0.49	0.66	5.7
STRIVE[19]	4.3	1.93	0.21	0.17	0.35	2.1
Our work	7.8	1.72	0.17	0.18	0.23	1.6

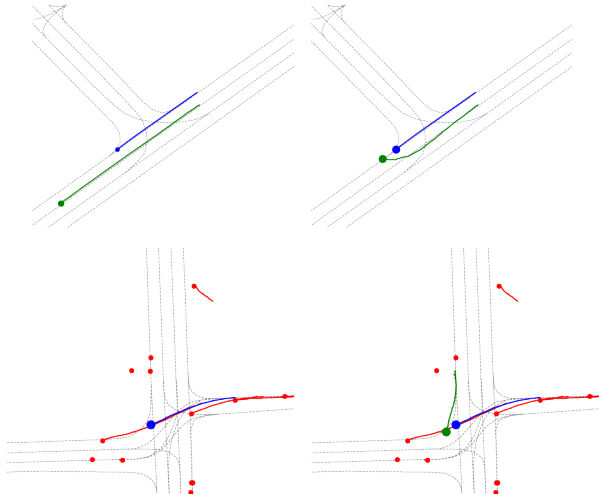


Fig. 5. The original scenario (left) and our edited version (right). The dotted line represents the lane centerline, and the circle denotes the final point of the trajectory. The green and red lines represent the trajectories of traffic agents, with the green line highlighting the safety-critical agents. For clarity, we only show the nearby agents in the figure. The blue line represents the AV’s trajectory. (1) lane changing. (2) intersection.

achieves better plausibility in every metric compared with previous works, while also leading to more challenging scenarios (higher collision rates).




Fig. 6. The blue vehicle indicates the ego vehicle (AV), the red vehicle indicates the original attacker vehicle and the grey vehicles indicate other agents in the scenario. STRIVE generates plain cutting-in behaviors, while our framework adds another agent (purple vehicle) and also an undrivable region (a purple parked bus).

C. Scenario Visualization

In Figure 5, we compare the original scenario with the edited version to demonstrate the diversity of safety-critical scenarios generated by our editing policy, which includes lane-changing and intersection scenarios. In the first example, our policy adds a new agent to the scenario and performs

TABLE II
ABLATION STUDY OF ACTION SPACE DESIGN AND THE RISK MODEL. M0 INDICATES A VANILLA RL BASELINE.

	AddAgent Action	Reconfigure Action	Risk Model	Collision Rate/Percent	Plausibility Score
M0				3.2	7.14
M1	✓			5.6	6.83
M2	✓	✓		7.1	7.52
M3	✓	✓	✓	7.8	7.44

a lane change in the intersection. In the second example, the agent of interest’s trajectory in an intersection is resampled and perturbed to a safety-critical lane merging scenario.

We also compare the scenario generated by our framework and STRIVE in Figure 6. Prior methods such as STRIVE can only generate risky behaviors from existing agents, leading to the cutting-in behaviors in the left figure. Our framework can add new agents or obstacles (the purple vehicle and the parked bus) and also reconfigure existing agents (red vehicle), to generate more natural but still challenging scenarios. These examples show the effectiveness of our framework in generating a wide range of realistic and complex safety-critical scenarios.

D. Ablation Study

We conduct an ablation study using various configurations (M0-3) to validate the effectiveness of our design choices. The results indicate that the AddAgent and Reconfigure actions not only increase the risk level of generated scenarios but also enhance their plausibility. Additionally, our risk model proves effective in guiding the search for riskier scenarios.

V. CONCLUSIONS

In this work, we proposed a scenario editing framework that leverages reinforcement learning to perform a set of operations on the driving scenario, including adding agents, perturbation, and resampling. We use generative models to learn the plausibility constraints of the scenarios and a novel anchor-based risky score to describe the desirable scenarios in a more fine-grained manner. Our evaluation shows that our framework can generate high-quality safety-critical scenarios, outperforming previous methods.

In future work, we plan to explore how our framework can be extended to handle more complex scenarios, such as those involving pedestrians or cyclists. We also plan to investigate how our framework can be integrated with existing scenario generation methods for autonomous vehicles.

REFERENCES

- [1] Y. Abeysirigoonawardena, F. Shkurti, and G. Dudek. Generating adversarial driving scenarios in high-fidelity simulators. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8271–8277, 2019.
- [2] M. Bansal, A. Krizhevsky, and A. S. Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. *CoRR*, abs/1812.03079, 2018.
- [3] Y. Chai, B. Sapp, M. Bansal, and D. Anguelov. Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. In *Conference on Robot Learning*, 2019.
- [4] M.-F. Chang, J. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8748–8757, 2019.
- [5] L. Feng, Q. Li, Z. Peng, S. Tan, and B. Zhou. Trafficgen: Learning to generate diverse and realistic traffic scenarios. *arXiv preprint arXiv:2210.06609*, 2022.
- [6] J. Gao, C. Sun, H. Zhao, Y. Shen, D. Anguelov, C. Li, and C. Schmid. Vectornet: Encoding hd maps and agent dynamics from vectorized representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11525–11533, 2020.
- [7] Z. Ghodsi, S. K. S. Hari, I. Frosio, T. Tsai, A. Troccoli, S. W. Keckler, S. Garg, and A. Anandkumar. Generating and characterizing scenarios for safety testing of autonomous vehicles. In *2021 IEEE Intelligent Vehicles Symposium (IV)*, pages 157–164, 2021.
- [8] A. Jain, L. D. Pero, H. Grimmer, and P. Ondruska. *Autonomy 2.0: Why is self-driving always 5 years away?* *CoRR*, abs/2107.08142, 2021.
- [9] R. Jenatton, C. Archambeau, J. González, and M. Seeger. Bayesian optimization with tree-structured dependencies. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1655–1664. PMLR, 06–11 Aug 2017.
- [10] N. Lee, W. Choi, P. Vernaza, C. B. Choy, P. H. Torr, and M. Chandraker. Desire: Distant future prediction in dynamic scenes with interacting agents. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 336–345, 2017.
- [11] E. Leurent. An environment for autonomous driving decision-making, 2018.
- [12] H. Liu, J. Zhao, and L. Zhang. Hierarchical learning-based autonomy simulator.
- [13] H. Liu, J. Zhao, and L. Zhang. Interpretable and flexible target-conditioned neural planners for autonomous vehicles. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10076–10082. IEEE, 2023.
- [14] Z. Liu, Z. Miao, X. Zhan, J. Wang, B. Gong, and S. X. Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2537–2546, 2019.
- [15] X. Lu, J. I. González, Z. Dai, and N. D. Lawrence. Structured variationally auto-encoded optimization. In *International Conference on Machine Learning*, 2018.
- [16] O. Makansi, O. Çiçek, Y. Marrakchi, and T. Brox. On exposing the challenging long tail in future prediction of traffic actors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13147–13157, October 2021.
- [17] M. O’Kelly, A. Sinha, H. Namkoong, J. Duchi, and R. Tedrake. Scalable end-to-end autonomous vehicle testing via rare-event simulation. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, page 9849–9860, Red Hook, NY, USA, 2018. Curran Associates Inc.
- [18] L. Peng, Z. Chen, Z. Fu, P. Liang, and E. Cheng. Bevsegformer: Bird’s eye view semantic segmentation from arbitrary camera rigs, 2022.
- [19] D. Rempe, J. Philion, L. J. Guibas, S. Fidler, and O. Litany. Generating useful accident-prone driving scenarios via a learned traffic prior. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17284–17294, 2021.
- [20] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017.
- [21] S. Suo, K. Wong, J. Xu, J. Tu, A. Cui, S. Casas, and R. Urtasun. Mixsim: A hierarchical framework for mixed reality traffic simulation.

In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9622–9631, June 2023.

- [22] A. Wachi. Failure-scenario maker for rule-based agent using multi-agent adversarial reinforcement learning and its application to autonomous driving. In *International Joint Conference on Artificial Intelligence*, 2019.
- [23] 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 (CVPR)*, pages 9909–9918, June 2021.