

# Conditional Flow-VAE for Safety-Critical Traffic Scenario Generation

Zimu Gong<sup>3,†,\*</sup>, Brian Zhaoning Zhang<sup>4,†,\*</sup>, Chris Zhang<sup>1,2,\*</sup>, Kelvin Wong<sup>1,2</sup>, and Raquel Urtasun<sup>1,2</sup>

**Abstract**—Safety-critical scenarios are essential for the development of autonomous vehicles (AVs) but are rare in real-world driving data. While simulation offers a way to generate such scenarios, manually designed test cases lack scalability, and adversarial optimization often produces unrealistic behaviors. In this work, we introduce a conditional latent flow matching approach for scalable and realistic safety-critical scenario generation. Our method uses distribution matching to transform nominal scenes into safety-critical rollouts. Furthermore, we demonstrate that incorporating both simulation and real-world data enables our framework to efficiently generate diverse, data-driven scenarios. Experimental results highlight that our approach is able to more consistently and realistically generate novel safety-critical scenarios, making it a valuable tool for training and benchmarking AV systems.

## I. INTRODUCTION

Safety-critical scenarios play a central role in the development of autonomous vehicles (AVs). Rare events such as sudden cut-ins, near-miss interactions, or unexpected braking are precisely the situations where an AV’s decision-making and planning policies are most challenged. Robust performance is essential, yet exposing AV systems to these conditions in the real world is costly and dangerous. Simulation is therefore critical: it enables evaluation under safety-critical conditions before deployment, reducing risk and accelerating development.

However, acquiring a sufficiently diverse and realistic set of safety-critical scenarios for simulation remains a major challenge. Traditionally, simulation-based approaches typically rely on heavy human curation. For example, [1] identifies potential hazardous events in their ODD and then recreates them in simulation, and [2] manually reconstructs safety-critical scenarios from police crash reports. However, this approach is far too tedious to scale efficiently and cost-effectively. Automated methods like adversarial optimization where agents are constructed to deliberately collide with the ego vehicle can be more easily scaled. However, real traffic participants are not inherently adversarial, and such methods neglect the prevalence of near-miss situations that pose genuine challenges for autonomy systems. In both cases, the resulting scenarios may not reflect the statistics of real world driving.

Alternatively, *distribution matching* approaches are a promising avenue towards traffic simulation that matches the real world. However, distribution matching approaches like imitation learning are notoriously data hungry, but real

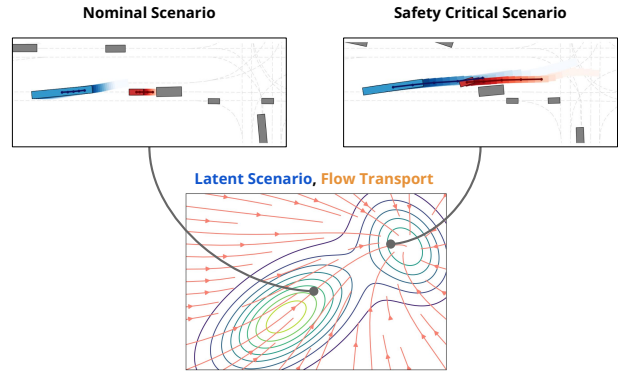


Fig. 1: Our method learns a latent space over scenarios and flow transport to map the nominal distribution to the safety-critical distribution.

safety critical scenarios are inherently rare and difficult to obtain. Standard traffic simulation models are trained predominantly on nominal data and thus are naturally biased toward reproducing nominal behaviors rather than generating safety-critical ones. Unfortunately, overly upsampling the limited number of safety-critical scenarios can easily lead to overfitting.

Towards a data-driven approach to generating safety-critical scenarios while preserving behavioral realism, we propose our conditional flow VAE. We supplement the limited number of high-fidelity real safety critical scenarios with lower fidelity but easily scalable synthetic data, and design a model which can maximally take advantage of this data mixture. Specifically, our model captures the semantics of a variety of driving scenarios with a conditional VAE encoder and uses a flow matching transformer to transform the VAE latents from the nominal distribution to the safety-critical counterpart. Then we utilize the VAE decoder to produce safety-critical rollouts from nominal driving scenarios, bridging the gap between rare real-world events and scalable synthetic generation. By incorporating both real-world and simulated data, our approach captures the realism of human driving while supplementing the long tail with synthetic diversity. Empirically, our flow approach outperforms alternative conditioning baselines. Furthermore, we enable controllable scenario difficulty by conditioning on automatically generated heuristic labels, allowing systematic evaluation across varying levels of criticality.

To summarize, our contributions are threefold:

- A generative framework for safety-critical scenario generation based on conditional latent flow matching.

The authors are with <sup>1</sup>Waabi Innovation Inc, <sup>2</sup>University of Toronto, <sup>3</sup>University of Michigan-Ann Arbor, <sup>4</sup>University of Waterloo.

<sup>†</sup>Work done during affiliation with Waabi.

\*Equal contribution.

- Realism through distribution alignment, enabling transitions from nominal to safety-critical outcomes without adversarial artifacts.
- Data efficiency and controllability, achieved through the integration of real and simulated data along with difficulty-conditioned generation.

Together, these advances provide a practical and principled way to generate realistic safety-critical scenarios.

## II. RELATED WORK

### A. Challenging Scenario Generation

Scenario generation includes a two-fold objective of actor placement and actor roll-out. The former objective generates a new scenario initialization from scratch or modifies existing scenarios. It often specifies the actor placement and initialization states so that the downstream simulation may yield safety-critical outcomes [3], [4]. In this paper, we focus on generating safety-critical scenarios from a nominal initialization with a roll-out model, where we train a model to directly control the actor maneuvers to be applied upon any nominal initialization.

Early approaches rely on manually designed scenarios and heuristic rules [5], often embedded in simulation platforms such as CARLA [6]. These methods offer clear control over vehicle maneuvers by manually specifying the planning trajectory and kinematic constraints. However, they are often limited in scalability and diversity, as each scenario must be explicitly scripted by engineers. The scenario parameters also need to be decided carefully to ensure that there is no collision or other undesired behaviors.

The advent of data-driven approaches allows the machine learning model design that focuses on learning from mass-scale driving datasets like WOMD [7]. With generative machine learning model architectures like VAE, autoregressive models and diffusion models [8], these methods are often trained on behavior cloning objectives that allow high reconstruction L2 scores on the eval data [9], [10], [11]. They also utilize common knowledge like collision loss and traffic signals to enhance realism [12], [13]. However, the behavior cloning objective implies that the model is fitted to the distribution of the training dataset, where most of the data are nominal [14], which hinders the ability to diversify towards safety-critical rollouts.

Another approach is to apply an adversarial objective to the actors in the scenario. Many of those apply similar strategies where they manipulate the learned representation of the actors in the latent space. STRIVE [15] utilizes an optimization objective on the latent space of the CVAE. Other methods use reinforcement learning based editing [16], [17], [18]. The actors in the scenario are often set with an adversarial objective to cause collisions with the ego actor. While effective at exposing the weakness of planners, these methods often compromise realism: real traffic actors are not inherently adversarial, thus the resulting trajectories may be unnatural. Besides, the method also requires the presence of a planner module for the ego vehicle, leading to more complex training architecture.

### B. VAE and Flow Matching Model

Variational Autoencoder models are useful in representation learning, where it efficiently compresses the input features into a latent space, so that a decoder can utilize them for downstream tasks [19]. It is desired to manipulate in the latent space to achieve specific objectives. Research in computer vision has shown that the representation space can be decomposed into subcomponents and then used for domain translation [20]. In the context of autonomous driving, VAE encodes traffic scenarios efficiently conditioned on the past states and high-level scenario information, and a decoder head is usually applied to generate per-actor roll-outs. Previous work [21] have shown that the VAE embedding space contains useful information that is valuable for interpretable maneuver generation. In our work, we show that the nominal and safety-critical scenarios reside in different subsections of the latent space and a transfer map can be learned to map the nominal latents to the safety-critical latents.

A flow matching model is a type of generative model that learns to transform samples from a simple base distribution into samples from a complex target distribution following a continuous flow [22]. Instead of learning a discrete sequence of transformations (like normalizing flows) or stepping with random noises (like diffusion), flow matching learns a vector field that describes explicitly how data should move between source and target distributions over time. During inference, the flow model uses a sampler to step through the timesteps from 0 to 1 along the learned field to map the sample from the source distribution to the target distribution. Key developments, including Rectified Flow [23], allows flowing from an arbitrary source distribution instead of the standard Gaussian. The method is widely applied to various tasks like text-to-image generation [24], and robotics [25]. It is also easy to apply conditioning on the flow model to boost the performance [26]. Compared to diffusion, the computational efficiency and stable training objective make it preferable for learning the transfer from nominal scenarios into safety-critical ones within the latent space.

## III. PRELIMINARIES

### A. Traffic Modeling

We define the problem scope as generating a traffic scenario with  $N$  actors in a finite horizon of  $T$  time steps. We use  $Y^t = \{y_1^t, y_2^t, \dots, y_N^t\}$  to denote the actor states at time  $t$ . For each vehicle state we define  $y_i^t = (b_x, b_y, b_z, b_\theta, b_v, b_l, b_w, b_h)$ , which describes the 3D position, yaw, velocity, and the length, width, and height of a vehicle's bounding box. The model observes the high-definition map  $\mathbf{M}$  and past  $H$  states  $Y^{-H:0}$  and outputs the vehicle control action sequence  $A^t = \{a_1^t, a_2^t, \dots, a_n^t\}$ , where  $a_i^t = (a_{accel}, a_{steer})$ . The goal of traffic modeling is typically to learn to model the distribution over future actor states.

$$p(Y^{1:T} | \mathbf{M}, Y^{-H:0}) \quad (1)$$

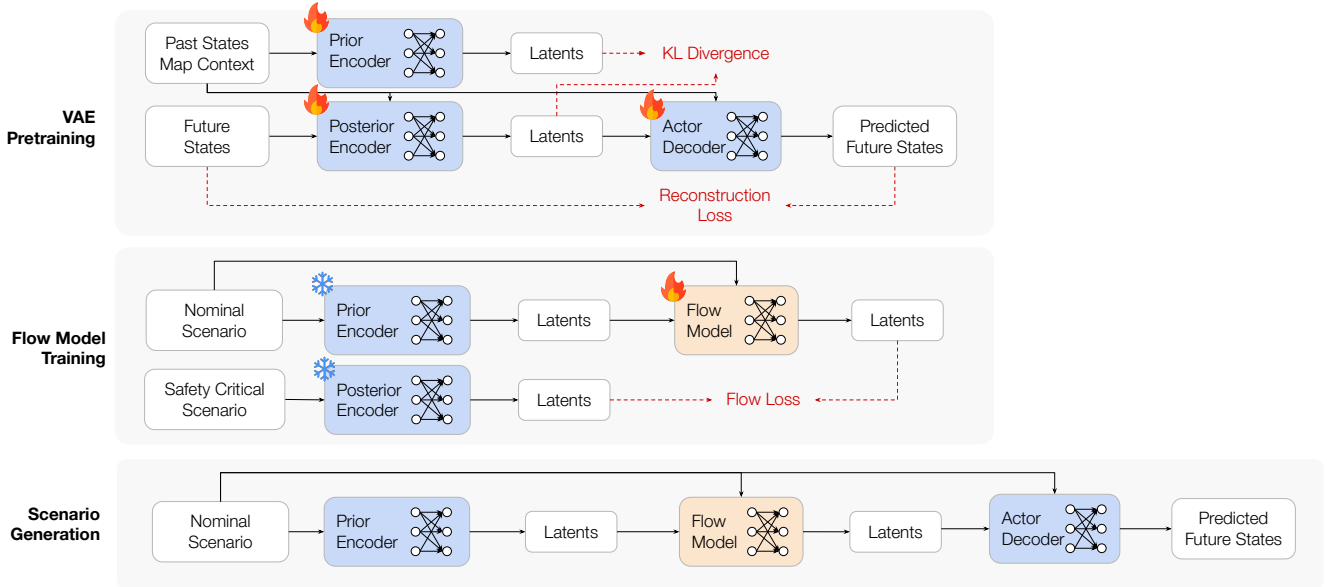


Fig. 2: **Conditional Flow VAE.** We first pretrain the VAE model on a mixture of data. Next, the VAE model is frozen and a flow model learns to map nominal latents to safety-critical latents. During inference, the flow transformer takes the prior latent and predicts the posterior latent. All the latents are passed through the CVAE decoder to generate the final actor states.

### B. Rectified Flow

Given observations of two distributions  $p_0, p_1$  on  $\mathbb{R}^d$  we wish to find a transport map  $T: \mathbb{R}^d \rightarrow \mathbb{R}^d$  such that  $X_1 := T(X_0) \sim p_1$  when  $X_0 \sim p_0$ . A rectified flow is an ODE on time  $t \in [0, 1]$

$$dZ_t = v_\theta(Z_t, t)dt \quad (2)$$

where  $v_\theta: \mathbb{R}^d \rightarrow \mathbb{R}^d$  is a velocity field that learns to drive the flow along the linear path from  $X_0$  to  $X_1$  by optimizing

$$\int_0^1 \mathbb{E} [ \| (X_1 - X_0) - v(X_t, t) \|^2 ] dt \quad (3)$$

where  $X_t = tX_1 + (1-t)X_0$ . After learning,  $v_\theta$  can be used to sample from  $p_1$  given samples from  $p_0$ .

### C. Variational Autoencoder

The variational autoencoder is a latent variable approach to generative modeling

$$p(x) = \int_z p(x|z)p(z)dz \quad (4)$$

where  $z$  is some latent variable meant to capture unobserved aspects of the generative process. An encoder  $q_\theta(z|x)$ , decoder  $p_\theta(x|z)$  and prior  $p_\theta(z)$  can be jointly learned to optimize the evidence lower bound:

$$\mathbb{E}_{z \sim q_\theta} [\log p_\theta(x|z)] + \mathcal{D}_{\text{KL}}(q_\theta(\cdot|x) \parallel p_\theta(\cdot)) \quad (5)$$

where  $\mathcal{D}_{\text{KL}}$  is the Kullback-Leibler divergence. After learning, sampling from  $p(x)$  amounts to sampling from the prior followed by the decoder. VAEs can be extended to support conditional generation [27] by extending the encoder, decoder and prior to be conditional distributions as well, e.g.  $q_\theta(z|x, c)$ ,  $p_\theta(x|z, c)$ ,  $p_\theta(z|c)$  respectively.

## IV. METHODOLOGY

The overall framework of the conditional flow VAE is depicted in Fig. 2.

### A. Latent flow matching

We take a distribution matching approach for safety-critical scenario generation. Let  $p_N$  be the nominal distribution of traffic scenarios for which we have many samples, and  $p_S$  be the safety-critical distribution, for which we have comparatively fewer samples. Our approach learns a flow from  $p_N$  to  $p_S$ . Doing so allows us to learn to sample new scenarios from  $p_S$  by transporting samples from  $p_N$ .

To begin, we follow [12] and model general traffic scenarios using a conditional VAE:

$$p(Y^{1:T} | \mathbf{M}, Y^{-H:0}) = \int_Z p_\theta(Y^{1:T} | \mathbf{M}, Y^{-H:0}, Z) p_\theta(Z | \mathbf{M}, Y^{-H:0}). \quad (6)$$

In this case, the encoder (posterior) is given as  $q_\theta(Z | \mathbf{M}, Y^{-H:0}, Y^{1:T})$ . Note that we learn the conditional VAE on the *mixture* of  $p_N$  and  $p_S$ , meaning the model is trained on samples from both distributions.

We now define the *safety-critical* aggregate posterior as

$$q_\theta^S(Z | \mathbf{M}, Y^{-H:0}) = \mathbb{E}_{p_S(Y^{1:T} | \mathbf{M}, Y^{-H:0})} [q_\theta(Z | \mathbf{M}, Y^{-H:0}, Y^{1:T})]. \quad (7)$$

Intuitively, this is the distribution over  $Z$  for safety-critical scenarios. We sample from  $q^S$  using the VAE posterior on safety-critical scenarios. The flow model then aims to learn a transport between the prior and this safety-critical latent

distribution by optimizing

$$\int_0^1 \mathbb{E} [ \| (Z_1 - Z_0) - v(Z_t, t) \|^2 ] dt \quad (8)$$

where  $Z_t = tZ_1 + (1 - t)Z_0$ , and  $Z_1 \sim q_\theta^S$  and  $Z_0 \sim p_\theta(Z|\mathbf{M}, Y^{-H:0})$ . This objective trains the model to map prior latents to safety-critical latents.

There are several advantages to our approach. By training the VAE on the mixture of nominal and safety-critical data, we are able to learn better overall realistic driving by making use of all data, as opposed to learning only on safety-critical data. However, the explicit flow objective allows us to steer our sampling towards the safety-critical distribution. Compared to trajectory space, flowing in latent space also allows us to control the degree of safety criticality by doing a partial flow (e.g. until  $t = 0.5$ ), since the decoder still maps intermediate results in latent space to plausible futures.

### B. Training Recipe

We follow a two-stage training procedure inspired by latent space manipulation practices in computer vision [28]. In the first stage, we train a conditional VAE on both nominal and safety-critical scenarios to establish a stable and semantically meaningful latent representation for both types. In the second stage, we train the flow model exclusively on safety-critical scenarios to learn the distributional transport from nominal to safety-critical latents. We found that this staged design is essential for reliable convergence. This is because empirically we observed that for a single stage end-to-end training approach, the prior and posterior distributions of the VAE undergo large shifts in the early iterations, while the flow model simultaneously attempts to learn the mapping with high learning rates. As a result, the flow model is effectively trained on a non-stationary target, which often leads to instability [29]. By decoupling the stages, the VAE first provides a fixed latent space learned under a combination of imitation and traffic-compliance objectives. The subsequent flow model is then trained with a flow objective, which benefits from theoretical convergence guarantees under a fixed latent space [30].

### C. Architecture

We now describe the neural network architecture used for the different components of our approach.

*a) Backbone:* A transformer [31] based architecture is used as the backbone network for the VAE and flow transformer. We adopt several common techniques used in traffic modelling. We first extract state features from each actor using a simple MLP. Map features are extracted using an off-the-shelf map encoder [32]. Map and actor state features are then augmented with PairPose relative positional features [32] allowing for viewpoint-invariance. Our transformer comprises interleaved actor-to-map, actor-to-actor and actor-to-time attention layers [33]; relative positional encodings [32], [10], [34] between actors are used to preserve viewpoint-invariance. To save on computational cost, actor-to-actor and actor-map attention is limited to the

top- $k$  closest actors or lane graph nodes, essentially forming a local context for each actor.

*b) VAE:* Following prior work on multi-agent traffic simulation [12], [15], we employ a conditional variational autoencoder (CVAE) approach to learn latent embeddings that capture rich scene semantics and the multi-agent interactions. The prior  $p_\theta(Z|\mathbf{M}, Y^{-H:0})$  and posterior  $q_\theta(Z|\mathbf{M}, Y^{-H:0}, Y^{1:T})$  use the same backbone described above, differing only in the number of timesteps of actor states observed, sharing the map encoder. Note that, similar to [12], we predict a separate latent for each actor. The decoder also uses the same backbone; the latent is fused into the actor feature, and a steering and acceleration is predicted per actor.

*c) Flow model:* With the stable latent representation established from the CVAE, safety-critical rollout generation requires transforming nominal latents into their safety-critical counterparts. We frame this as a distribution matching problem: safety-critical behaviors (e.g., hard braking, aggressive cut-ins) occupy distinct subregions of the latent space, and our goal is to learn a mapping from nominal priors to these critical submanifolds.

We use the same transformer backbone for our flow model. Our transformer backbone is conditioned on the actor state feature to implement flow matching in latent space. We combine the flow matching context along with the actor-level features to form the following input features:

$$\begin{aligned} E_{actor} &= [X^t, E_T, E_x, E_c] \\ E_T &= \text{SinusoidalPE}(t_{denoise}) \\ E_x &= t_{denoise} Z_{posterior} + (1 - t_{denoise}) Z_{prior} \\ E_c &= \text{MLP}(\text{ManeuverLabel}), \end{aligned} \quad (9)$$

where  $t_{denoise}$  is the denoising time step drawn from the uniform distribution  $U(0, 1)$ .  $X$  is actor state and map features.  $E_T$  is the sinusoidal positional encoding of the denoising time step.  $E_x$  is an interpolation between the prior latent  $Z_{prior}$  and the posterior latent  $Z_{posterior}$ . Since the flow model only conditions on the scene initialization, an optional maneuver indicator is accepted to control the level of aggressiveness of maneuver. We classify each scenario in the training dataset into one of the three categories: nominal, safety-critical, and very safety-critical. The label is computed with heuristics on vehicle kinematics and time-to-collision.  $E_c$  is the maneuver label projected with an MLP.

During inference, we discard the posterior encoder and use the prior encoder only since ground-truth futures are unavailable. Given any initialization state, the prior encoder produces a latent  $Z_{prior}$ , which is then transformed by the Flow Transformer into a steered latent  $Z_{pred}$ . This latent is decoded through the CVAE decoder to acquire the final actor states.

### D. Dataset

Performing distribution matching between nominal and safety-critical rollouts requires access to safety-critical scenarios observed in real driving logs. However, due to their

TABLE I: **Realistic Safety-critical Scenario Generation.** We evaluate against baselines on our held out set of real safety critical scenarios, and obtain the highest near miss % (valid safety critical scenario), while having high distributional similarity.

	minSTTC ↓	Near Miss (%) ↑	SCR (%) ↓	Displ Error ↓	Distribution JSD		
					Velocity ↓	Accel ↓	Jerk ↓
VAE	2.927	22.9	0.7	5.71	0.245	0.101	0.017
VAE + Curation	3.776	31.2	0.8	7.32	0.278	0.136	0.019
STRIVE	1.885	30.2	7.0	7.22	0.277	0.151	0.013
FlowVAE	2.190	45.8	1.6	6.51	0.262	0.131	0.021

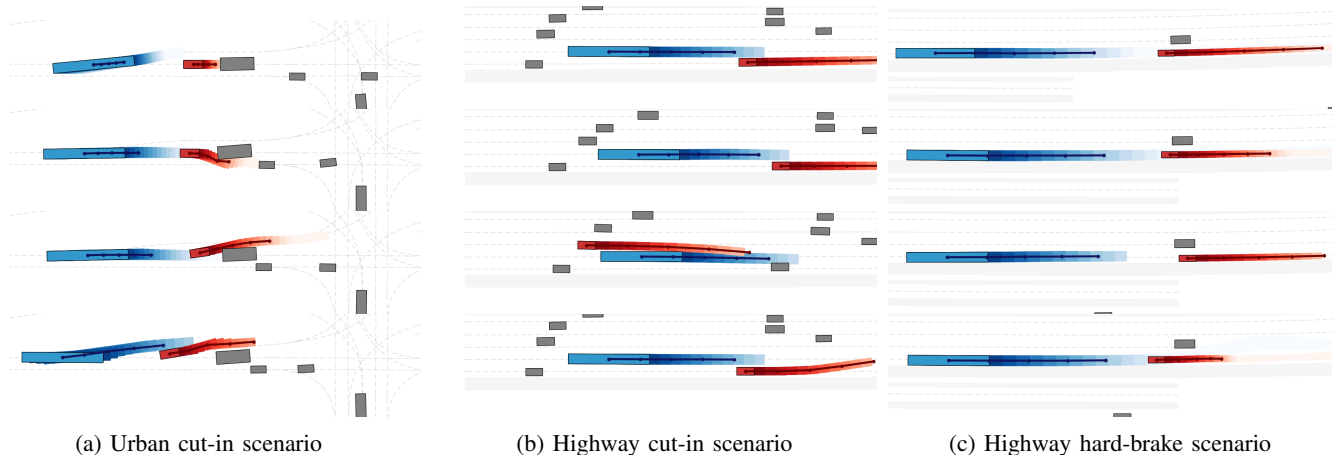


Fig. 3: **Qualitatives.** From top to bottom: original nominal scenario, VAE reconstruction, STRIVE, our model.

inherent sparsity in naturalistic datasets, none of the widely used open-source autonomous driving corpora explicitly curate such subsets. To address this limitation, we adopt a simulation–real data mixing strategy that balances scalability with realism. For real safety-critical scenarios, we conduct targeted data mining over a catalogue of real driving logs. We extract approximately 500 unique scenarios with challenging situations; e.g., actors perform abrupt cut-ins, aggressive braking, etc. The result of this data mining is a set of smaller scale but high-fidelity demonstrations of realistic human driving behaviors under challenging conditions.

To supplement the real safety critical scenarios, we leverage simulation. Our goal is to generate realistic scenarios which introduce diverse behaviors that can help transfer to real safety critical scenarios. We use existing deep-learning based traffic simulation models for nominal traffic, with the addition of a “hero actor” selected among existing actors with heuristics, or additionally inserted into the scene. It is parameterized by Intelligent Driver Model (IDM) heuristics and programmed to execute either a cut-in or a hard braking maneuver. Overall, this procedure provides reasonably fine-grained control over desired maneuvers and generates many new scenarios. We leverage rejection sampling to throw away simulations that fail a small set of basic checks due to a failure in the heuristics. This allows us to generate an order of magnitude more safety critical scenarios than we have mined from real logs. We also found that the diversity of the simulation generated scenarios helps improve the model performance. We created three versions of simulation data with different heuristics: one with the most safety-critical

maneuvers, one with kinematic constraints (deceleration, TTC, etc.) approximately tuned to the real data. It turns out that using both versions of simulation data achieves the best performance. While the resulting behaviors themselves remain limited in diversity, empirically we will show that they still provide benefits and partially transfer to the real evaluation set.

During training, we blend real and simulated scenarios using a hyperparameter  $\alpha_{real}$ , which determines the relative proportion of real and sim data. Each training batch has an  $\alpha_{real}\%$  chance to draw a sample from the real dataset, and a  $(1 - \alpha_{real})\%$  chance to draw from a sim sample. This mechanism allows us to smoothly adjust the balance between realism and scalability, and to study the effect of sim–real composition on downstream performance. As real safety-critical data alone is too scarce to provide sufficient coverage of the scenario space, while sim-only data introduces a domain gap,  $\alpha_{real}$  serves as a way to balance the two, which we empirically validate in Section V-D.

## V. EXPERIMENTS

We first evaluate end-to-end our approach’s ability to generate realistic safety-critical scenarios. Section V-B shows that compared to baselines, scenarios generated by our approach more closely match held-out real safety-critical scenarios both quantitatively and qualitatively. Next, we investigate the first key aspect of our approach: our conditional flow architecture. We ablate our design choices and show that both conditioning and flow are important to achieving good results (Section V-C). The other key aspect of our approach

TABLE II: **Conditional Flow Ablation.** We ablate the effect of our combined conditioning and flow approach. Our results show that both aspects are important, and together form the best results.

Flow	Conditioning	Distribution JSD						
		minSTTC ↓	Near Miss (%) ↑	SCR (%) ↓	Displ Error ↓	Velocity ↓	Accel ↓	Jerk ↓
		2.927	22.9	0.7	5.71	0.245	0.101	0.017
✓		2.190	31.9	1.7	8.33	0.279	0.151	0.025
	✓	2.863	16.7	1.7	6.39	0.268	0.120	0.016
✓	✓	2.190	45.8	1.6	6.51	0.262	0.131	0.021

TABLE III: **Simulation Data Transfers to Real.** We evaluate the effect of our simulation data and find that introducing a moderate amount of simulation data transfers to the real safety-critical distribution.

	Distribution JSD						
	minSTTC ↓	Near Miss (%) ↑	SCR (%) ↓	Displ Error ↓	Velocity ↓	Accel ↓	Jerk ↓
Sim-only	3.052	39.5	0.035	7.98	0.270	0.166	0.038
10% Real	2.944	39.6	0.030	7.97	0.276	0.151	0.029
20% Real	2.357	45.8	0.027	7.49	0.265	0.139	0.028
40% Real	1.945	58.3	0.029	7.19	0.251	0.146	0.032
60% Real	2.568	58.3	0.017	7.31	0.263	0.136	0.028
100% Real	5.078	35.4	0.024	8.28	0.284	0.152	0.029

is our data composition, and the use of simulation data to supplement real examples. In Section V-D, we evaluate our approach trained on various different data compositions and show that a simple balance of simulation and real safety-critical examples provides the best results. Finally, we study the controllability of our approach in Section V-E. We showcase how our model responds to the conditioning, and how intermediate flow timesteps can be an additional lever for controlling the specific degree of safety criticality.

#### A. Experimental Setup

1) *Dataset:* We conduct experiments with an in-house self-driving dataset. Our dataset spans both highway and urban driving, consisting of approximately 20,000 traffic scenarios. Each snippet contains 20s of driving data. Approximately 10,000 snippets are from real logs, 10,000 are simulated safety critical scenarios as described in Section IV-D. Additionally, as described in Section IV-D we have curated approximately 500 real safety critical snippets. The remaining training samples are also upsampled. We hold out 20% of the real safety critical data for evaluation.

2) *Metrics:* Evaluating a generative model for scenario generation is non-trivial and requires multiple metrics measuring realism and safety-criticality. We propose our experiments on the following suite of metrics.

- **minSTTC and Near Miss Rate** To evaluate if our method learns to construct near-miss cases from the training distribution, we propose a scenario-level minimum time to collision metric (minSTTC). For each scenario sample, we calculate the minimal time to collision between the ego actor and the closest leading actor throughout the entire rollout with an upper bound of 10 seconds. We report the median of the minSTTC because in the cases where no likely collision is going to happen, the minSTTC is likely to be large. We consider a more effective safety-critical scenario as inducing a

small TTC without causing any collision. We also report the percentage of scenarios where the minSTTC is less than 3 seconds, which we consider as a near miss that challenges the planner.

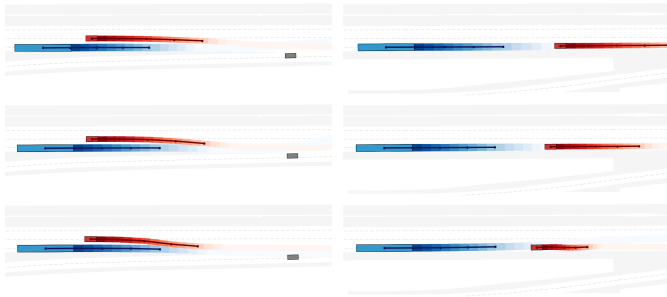
- **Distribution JSD** Following common practice [35], [36], we compute the distributional kinematics metrics of actors. Smaller divergence indicates that the model captures the essence of safety-critical maneuvers. We measure against common kinematic metrics that characterize actor maneuver: linear and angular speed and acceleration.
- **Collision Rate** We evaluate the average percentage of actors colliding in each scenario based on a small IOU threshold between the bounding boxes of the actors. Collision rate should generally be low even for safety-critical scenarios, since none of the ground truth data has any collisions. However, models can obtain low collision rate by generating nominal scenarios, so other metrics like minSTTC must also be considered.
- **Reconstruction** We also provide L2 reconstruction to the ground truth trajectory as another way to measure realism to supplement distribution JSD.

#### B. Generating Realistic Safety Critical Scenarios

We compare against 3 baselines

- 1) *VAE* is the base model, trained on the same data as our approach.
- 2) *VAE + Curation* is the base CVAE model, trained only on safety critical scenarios (both real and sim)
- 3) *Strive* [15] is a SOTA baseline which performs optimization in latent space. We use the same base CVAE as our approach.

Table I shows the results. Compared with the baseline CVAE models, our model achieves comparable kinematics metrics while generating more safety-critical scenarios. Our model generates smaller minSTTC with higher near-miss



(a) Cut-in scenario (b) Hard-brake scenario

Fig. 4: **Varying number of flow timesteps.** More flow steps corresponds to higher safety criticality.

rate. Because the baseline CVAE is trained on the base mixture distribution, it does not produce as many safety-critical scenarios, as we can see by its relatively worse minSTTC and Near Miss rate. Our VAE + Curation baseline obtains higher near miss rate but is less realistic overall. This is because the omitted nominal data still contains valuable learning signal, in particular for background traffic, etc. STRIVE also obtains a high near miss rate but suffers from realism. We believe that this is because the prior model does not provide strong enough regularization. Also, the adversarial optimization is not explicitly aware of the real world distribution of safety-critical scenarios. On the other hand our flow approach obtains the best of both worlds as it is able to generate a large percentage of near miss scenarios while maintaining good performance in the other metrics.

Qualitatively (Fig. 3), we can see that with an ordinary highway scenario, our model perturbs the maneuver of the leading actor by causing it to perform a hard brake, leading the ego vehicle to follow as well. On the second occasion, the model causes the actor in the neighbor lane to cut into ego vehicle’s lane aggressively, also causing a hard brake from the ego actor. One side benefit of the flow method is that the model automatically chooses the actor to interact with the ego vehicle as well as the maneuver to perform, so that there is no need for explicit interaction design.

### C. Conditional Flow Ablation

We now ablate the key architectural choices of our approach: the flow transformer, and the additional conditioning. The flow-only model simply removes the conditioning from the flow transformer. For the conditioning-only model, we add a similar conditioning encoder to the VAE prior. Table I shows the results. As expected, without any flow or conditioning to steer the sampling, the base model has trouble generating safety critical rollouts. We see that conditioning on its own is also ineffective. Our hypothesis is that adding conditioning during the VAE training potentially harms representation learning as it potentially provides too much of a shortcut. On its own, flow is already effective, but adding conditioning to flow results in the best overall model.




Fig. 5: **Flow timesteps vs. Reconstruction Error.** On average, reconstruction error smoothly drops with flow timesteps.

TABLE IV: **Varying conditioning.** We see that the model is controllable via the conditioning parameter.

	minSTTC ↓	Near Miss (%) ↑	Displ Error ↓
No Cond.	2.190	0.319	8.327
Cond. Nominal	3.263	0.281	9.727
Cond. Intermediate	2.293	0.343	9.271
Cond. Challenging	1.967	0.500	8.315

### D. Data Composition

We evaluate our model performance with different mixtures of real data in Table III. For these experiments, the same base VAE is used, but we adjust the mixed ratio of sim and real data when training the flow model. We see that using sim or real only is ineffective, due to lower fidelity data, and smaller scale data respectively. Combining them shows the best results, with a sweet spot at around 50%.

### E. Controllability Studies

Controllability is another desired property for our model because it allows us to generate different levels of safety-critical scenarios, which could be used to gradually test the performance upper bound of the autonomy system. In this part, we compare the difference of the rollouts from the same initialization but with different maneuver labels. In Table IV, we show the performance of the model with different maneuver labels. We see minSTTC decreases and near-miss rate increases as the maneuver label becomes more challenging, while the unconditional model lies somewhere in between. We also found that controllability could also be achieved by manipulating the flow time steps. We sample the latents along the inference time steps and decode them to visualize the reconstruction. Fig. 4 shows that the model rollout is nominal at  $t = 0$ , and becomes safety-critical at  $t = 1$ , providing more refined control of scenario generation when used in conjunction with the maneuver labels.

## VI. CONCLUSION

In this paper, we proposed Flow VAE, a method for data-driven safety-critical traffic scenario rollout generation. Flow VAE is a flow matching transformer that learns to transfer the latents from nominal initialization into safety-critical ones. We also show that using a mixture of sim and real data, we

are able to scalably generate safety-critical scenarios with a small dataset. Our ablation studies validate our architectural and data composition design choices, and we further show multiple methods to control our model at varying granularity.

Our real data curation and synthetic generation are proof-of-concept and can be improved. For instance higher fidelity vehicle models involving friction could better address sim-to-real and allow for more interesting scenarios involving slipping. Scaling up data collection and synthetic generation techniques could allow for increasingly diverse safety critical scenarios to be generated. While we showed controllability in the degree of safety criticality (through maneuver labels and number of flow timesteps), controlling the maneuver itself could be an interesting future direction.

## REFERENCES

- [1] K. D. Kusano, K. Beatty, S. Schnelle, F. Favaro, C. Cray, and T. Victor, "Collision avoidance testing of the waymo automated driving system," 2022. [Online]. Available: <https://arxiv.org/abs/2212.08148>
- [2] J. M. Scanlon, K. D. Kusano, T. Daniel, C. J. Alderson, A. Ogle, and T. Victor, "Waymo simulated driving behavior in reconstructed fatal crashes within an autonomous vehicle operating domain." *Accident; analysis and prevention*, vol. 163, p. 106454, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:232285642>
- [3] L. Feng, Q. Li, Z. Peng, S. Tan, and B. Zhou, "Trafficgen: Learning to generate diverse and realistic traffic scenarios," 2023. [Online]. Available: <https://arxiv.org/abs/2210.06609>
- [4] S. Tan, B. Ivanovic, X. Weng, M. Pavone, and P. Kraehenbuehl, "Language conditioned traffic generation," 2023. [Online]. Available: <https://arxiv.org/abs/2307.07947>
- [5] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Phys. Rev. E*, vol. 62, pp. 1805–1824, Aug 2000. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevE.62.1805>
- [6] A. Dosovitskiy, G. Ros, F. Codevilla, A. M. López, and V. Koltun, "CARLA: an open urban driving simulator," *CoRR*, vol. abs/1711.03938, 2017. [Online]. Available: <http://arxiv.org/abs/1711.03938>
- [7] K. Chen, R. Ge, H. Qiu, R. Ai-Rfou, C. R. Qi, X. Zhou, Z. Yang, S. Etinger, P. Sun, Z. Leng, M. Mustafa, I. Bogun, W. Wang, M. Tan, and D. Anguelov, "Womd-lidar: Raw sensor dataset benchmark for motion forecasting," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, May 2024.
- [8] 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*, vol. 24, no. 7, pp. 6971–6988, 2023.
- [9] W. Wu, X. Feng, Z. Gao, and Y. Kan, "Smart: Scalable multi-agent real-time motion generation via next-token prediction," 2024. [Online]. Available: <https://arxiv.org/abs/2405.15677>
- [10] Z. Zhong, D. Rempe, D. Xu, Y. Chen, S. Veer, T. Che, B. Ray, and M. Pavone, "Guided conditional diffusion for controllable traffic simulation," 2022. [Online]. Available: <https://arxiv.org/abs/2210.17366>
- [11] D. Xu, Y. Chen, B. Ivanovic, and M. Pavone, "Bits: Bi-level imitation for traffic simulation," 2022. [Online]. Available: <https://arxiv.org/abs/2208.12403>
- [12] S. Suo, S. Regalado, S. Casas, and R. Urtasun, "TrafficSim: Learning to simulate realistic multi-agent behaviors," 2021. [Online]. Available: <https://arxiv.org/abs/2101.06557>
- [13] C. Zhang, J. Tu, L. Zhang, K. Wong, S. Suo, and R. Urtasun, "Learning realistic traffic agents in closed-loop," *arXiv preprint arXiv:2311.01394*, 2023.
- [14] J.-T. Zhai, Z. Feng, J. Du, Y. Mao, J.-J. Liu, Z. Tan, Y. Zhang, X. Ye, and J. Wang, "Rethinking the open-loop evaluation of end-to-end autonomous driving in nusences," 2023. [Online]. Available: <https://arxiv.org/abs/2305.10430>
- [15] D. Rempe, J. Pailion, L. J. Guibas, S. Fidler, and O. Litany, "Generating useful accident-prone driving scenarios via a learned traffic prior," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [16] H. Liu, L. Zhang, S. K. S. Hari, and J. Zhao, "Safety-critical scenario generation via reinforcement learning based editing," 2024. [Online]. Available: <https://arxiv.org/abs/2306.14131>
- [17] C. Zhang, S. Biswas, K. Wong, K. Fallah, L. Zhang, D. Chen, S. Casas, and R. Urtasun, "Learning to drive via asymmetric self-play," 2024. [Online]. Available: <https://arxiv.org/abs/2409.18218>
- [18] 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," 2023. [Online]. Available: <https://arxiv.org/abs/2101.06549>
- [19] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," 2022. [Online]. Available: <https://arxiv.org/abs/1312.6114>
- [20] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz, "Multimodal unsupervised image-to-image translation," 2018. [Online]. Available: <https://arxiv.org/abs/1804.04732>
- [21] A. Bairouk, M. Maras, S. Herlin, A. Amini, M. Blanchon, R. Hasani, P. Chareyre, and D. Rus, "Exploring latent pathways: Enhancing the interpretability of autonomous driving with a variational autoencoder," 2024. [Online]. Available: <https://arxiv.org/abs/2404.01750>
- [22] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le, "Flow matching for generative modeling," 2023. [Online]. Available: <https://arxiv.org/abs/2210.02747>
- [23] X. Liu, C. Gong, and Q. Liu, "Flow straight and fast: Learning to generate and transfer data with rectified flow," 2022. [Online]. Available: <https://arxiv.org/abs/2209.03003>
- [24] P. Esser, S. Kulal, A. Blattmann, R. Entezari, J. Müller, H. Saini, Y. Levi, D. Lorenz, A. Sauer, F. Boesel, D. Podell, T. Dockhorn, Z. English, K. Lacey, A. Goodwin, Y. Marek, and R. Rombach, "Scaling rectified flow transformers for high-resolution image synthesis," 2024. [Online]. Available: <https://arxiv.org/abs/2403.03206>
- [25] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, L. X. Shi, J. Tanner, Q. Vuong, A. Walling, H. Wang, and U. Zhilinsky, " $\pi_0$ : A vision-language-action flow model for general robot control," 2024. [Online]. Available: <https://arxiv.org/abs/2410.24164>
- [26] A. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen, "Glide: Towards photorealistic image generation and editing with text-guided diffusion models," 2022. [Online]. Available: <https://arxiv.org/abs/2112.10741>
- [27] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," *Advances in neural information processing systems*, vol. 28, 2015.
- [28] A. Razavi, A. van den Oord, and O. Vinyals, "Generating diverse high-fidelity images with vq-vae-2," 2019. [Online]. Available: <https://arxiv.org/abs/1906.00446>
- [29] J. He, D. Spokoyny, G. Neubig, and T. Berg-Kirkpatrick, "Lagging inference networks and posterior collapse in variational autoencoders," 2019. [Online]. Available: <https://arxiv.org/abs/1901.05534>
- [30] Y. Jiao, Y. Lai, Y. Wang, and B. Yan, "Convergence analysis of flow matching in latent space with transformers," 2024. [Online]. Available: <https://arxiv.org/abs/2404.02538>
- [31] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [32] A. Cui, S. Casas, K. Wong, S. Suo, and R. Urtasun, "Gorela: Go relative for viewpoint-invariant motion forecasting," *arXiv preprint arXiv:2211.02545*, 2022.
- [33] J. Ngiam, B. Caine, V. Vasudevan, Z. Zhang, H.-T. L. Chiang, J. Ling, R. Roelofs, A. Bewley, C. Liu, A. Venugopal *et al.*, "Scene transformer: A unified architecture for predicting multiple agent trajectories," *arXiv preprint arXiv:2106.08417*, 2021.
- [34] Z. Zhou, L. Ye, J. Wang, K. Wu, and K. Lu, "Hivt: Hierarchical vector transformer for multi-agent motion prediction," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 8823–8833.
- [35] M. Igl, D. Kim, A. Kuefler, P. Mougin, P. Shah, K. Shiari, D. Anguelov, M. Palatucci, B. White, and S. Whiteson, "Symphony: Learning realistic and diverse agents for autonomous driving simulation," 2022. [Online]. Available: <https://arxiv.org/abs/2205.03195>
- [36] S. Tan, J. Lambert, H. Jeon, S. Kulshrestha, Y. Bai, J. Luo, D. Anguelov, M. Tan, and C. M. Jiang, "Scenediffuser++: City-scale traffic simulation via a generative world model," 2025. [Online]. Available: <https://arxiv.org/abs/2506.21976>