

# RealGen: Retrieval Augmented Generation for Controllable Traffic Scenarios

Wenhao Ding<sup>\*1,2</sup>, Yulong Cao<sup>\*2</sup>, Ding Zhao<sup>1</sup>, Chaowei Xiao<sup>2,3</sup>, and Marco Pavone<sup>2,4</sup>

<sup>1</sup> Carnegie Mellon University

<sup>2</sup> NVIDIA Research

<sup>3</sup> UW-Madison

<sup>4</sup> Stanford University

{wenhaod, yulongc}@nvidia.com

**Abstract.** Simulation plays a crucial role in the development of autonomous vehicles (AVs) due to the potential risks associated with real-world testing. Although significant progress has been made in the visual aspects of simulators, generating complex behavior among agents remains a formidable challenge. It is not only imperative to ensure realism in the scenarios generated but also essential to incorporate preferences and conditions to facilitate controllable generation for AV training and evaluation. Traditional methods, which rely mainly on memorizing the distribution of training datasets, often fail to generate unseen scenarios. Inspired by the success of retrieval augmented generation in large language models, we present **RealGen**, a novel retrieval-based in-context learning framework for traffic scenario generation. RealGen synthesizes new scenarios by combining behaviors from multiple retrieved examples in a gradient-free way, which may originate from templates or tagged scenarios. This in-context learning framework endows versatile generative capabilities, including the ability to edit scenarios, compose various behaviors, and produce critical scenarios. Evaluations show that RealGen offers considerable flexibility and controllability, marking a new direction in the field of controllable traffic scenario generation. Check our project website for more information: <https://realgen.github.io>.

## 1 Introduction

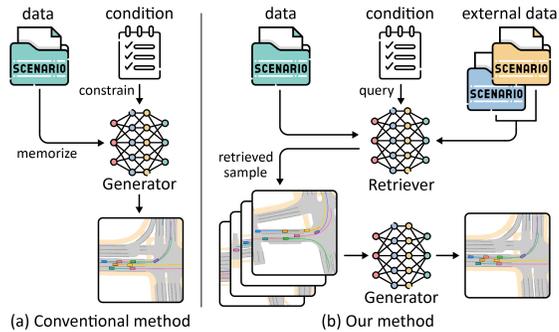
Simulation is indispensable in the development of autonomous vehicles (AVs), primarily due to the considerable risks associated with training and evaluating these systems in real-world conditions. The biggest challenge in simulations lies in achieving realistic driving scenarios, as this realism influences the discrepancy between AV performance in simulated and actual environments. Although advancements in high-quality graphical engines have significantly enhanced the perception quality of simulators, the realism of agent behavior remains constrained because of the complicated interactions among naturalistic agents. To counteract this issue, data-driven simulation has emerged as a promising approach in the realm of autonomous driving, which leverages real-world scenario datasets to accurately generate the behaviors of agents.

With the rapid achievement of deep generative models [29] and imitation learning algorithms [32], current data-driven simulations [23, 38] can generate scenarios that closely mimic human driver behavior. However, the effectiveness of these simulations in accelerating the development of AV is limited. This limitation stems from the need for scenarios that meet specific conditions tailored for targeted training and evaluation. Achieving such controllability in simulations is challenging due to the complex nature of driving scenarios, which involve intricate interactions, diverse road layouts, and varying traffic regulations.

In pursuit of controllability, existing work applies additional guidance, typically in the form of constraint functions [16, 64] or languages [49, 63], to pre-trained scenario generative models. Regularization of the generation process through these tools is straightforward and effective, yet it encounters two principal challenges. First, the training of scenario generative models typically utilizes naturalistic datasets, which might not encompass the specific scenarios desired as per the control signals. Even if such

scenarios exist within the dataset, they are often omitted because of the rarity of long-tail data. The second challenge is that the representation of the guidance to the generative model may not be sufficiently expressive to accurately depict complex scenarios, such as specifying intricate interactions among multiple vehicles, using language. These limitations underscore the need for a more sophisticated and nuanced framework for controllable scenario generation.

Retrieval Augmented Generation (RAG) [8, 25], which enhances the generative process by querying related information from external databases, represents great potential in the domain of large language models [37]. In contrast to conventional models that memorize all knowledge within their parameters, as shown in Figure 1(a), RAG models, shown in Figure 1(b), learn to generate comprehensive outputs by retrieving pertinent knowledge from a database, based on the input provided. A notable aspect of RAG is the ability of the database to undergo updates even after the model has been trained, allowing continuous improvement and adaptation. This flexible framework offers possibilities for controllable scenario generation by using appropriate template scenarios as input and facilitating the generation that is not only realistic but also aligned with specific training and evaluation requirements.



**Fig. 1:** (a) Conventional methods make the model memorize the data distribution for generating. (b) In contrast, our method employs a retriever to query datasets (including external data obtained after training) and uses a generative model to generate scenarios by integrating the information from the retrieved scenarios.

In this work, we present *RealGen*, a retrieval augmented generation framework for generating traffic scenarios. This framework, as shown in Figure 2, begins with the training of an encoder through contrastive self-supervised learning [42] to allow the retrieval process to query similar scenarios in a latent embedding space. Leveraging this latent representation, we subsequently train a generative model that combines retrieved scenarios to create novel scenarios. The key contributions of this paper are summarized below.

- We develop a novel contrastive autoencoder model to extract scenario embeddings as latent representations, which can be used for a wide range of downstream tasks.
- We propose the first retrieval augmented generation framework using the latent representation tailored for controllable driving scenario generation.
- We validate our framework through qualitative and quantitative metrics, demonstrating strong flexibility and controllability of generated scenarios.

## 2 Related Work

### 2.1 Traffic Modeling and Scenario Generation

Research on traffic modeling using generative models has received considerable attention in recent works. Notably, ScenarioNet [38] and Waymax [23] employ large-scale data to train imitation learning models, facilitating the generation of realistic multi-agent scenarios in simulations. SceneGen [50] utilizes the Long Short Term Memory [30] module to autoregressively generate the trajectories for vehicles and pedestrians based on provided maps. Additionally, TrafficSim [47] learns the multi-agent behaviors from real-world data, and TrafficGen [18] proposes a transformer-based autoencoder architecture to accurately model complex interactions involving multiple agents. MixSim [48] builds a reactive digital twin and finds safety-critical scenarios with black-box optimization. RTR [61] models reactive behaviors of vehicles using the combination of reinforcement learning and imitation learning.

Numerous prior studies have employed adversarial generation to synthesize rare yet critical scenarios. L2C [17] utilizes a reinforcement learning framework, where the reward for scenario generation is based on the collision rate. To enhance the realism in purely adversarial generation methods, MMG [15] incorporates the data distribution as regularization. Further developments such as AdvSim [54], AdvDO [7], and KING [26] integrate vehicle dynamics to directly optimize the trajectory to find critical scenarios.

The proliferation of large language models (LLMs) has led to recent approaches in generating traffic scenarios with language as conditions to follow instructions from humans. CTG++ [63] replaces the gradient guidance process of CTG [64] with cost functions generated by an LLM. Leveraging the power of GPT-4 [59], this framework can generate diverse motion data with language conditions. Additionally, LCTGen [49] also harnesses the strengths of GPT-4 to generate a heuristic intermediate representation of scenarios from language

inputs. Subsequently, a generative model pre-trained on open-source datasets generates various scenarios with this representation as inputs.

## 2.2 Retrieval-augmented Generation

Information retrieval (IR) [33] is the procedure of representing and searching a collection of data with the goal of extracting knowledge to satisfy the query of users. Predominantly utilized in search engines and digital libraries [43], IR primarily deals with the information stored in textual formats. Recently, IR has extended across diverse sectors, notably enhancing the quality of outputs in areas such as language modeling [3, 39], question answering [25, 62], image creation [2, 11], and the generation of molecular [55].

One intuitive usage of the retrieved information is to enhance input data through several methods, such as merging the original data and retrieved data [37], employing attention mechanism [3], or extracting skeleton [6]. The rationale behind retrieval systems stems from the impracticality of encoding all knowledge within model parameters, especially considering the dynamic nature of knowledge that evolves with human activities. Consequently, the capability to access external knowledge databases can significantly improve the precision and quality of generated responses, as evidenced in applications involving LLMs [3, 37].

Another usage of retrieved data, particularly pertinent to this study, involves the controllable generation by integrating desired features retrieved from the dataset. In [35], the authors generate product reviews with controllable information about the user, product, and rating. Similarly, the process of story creation can be viewed as a blend of external story databases with selected text fragments [57]. Furthermore, [55] explores the controllable generation of molecules to meet various constraints, utilizing a database comprising simple molecules that typically meet only one of these constraints.

## 2.3 Self-supervised Feature Learning

The necessity for manual labeling in supervised learning often results in human biases, extraneous noises, and labor-intensive efforts. Therefore, the paradigm of self-supervised learning (SSL) [46] is garnering heightened interest, particularly for its promising applications in language modeling and image interpretation. SSL algorithms usually learn implicit representations from extensive pools of unlabeled data without relying on human annotations. Generally, this line of research can be bifurcated into two categories [60]: generative SSL and discriminative SSL.

In the domain of generative SSL, models employ an autoencoder to convert input data into a latent representation, followed by a reconstruction process. As an illustrative example, Denoising Autoencoder [53], an early example of this approach, reconstructs images from their noisy version to extract representative features. Another technique in this field is masked modeling, extensively utilized in language models such as GPT-3 [4] and BERT [14], which predicts the masked token to gain a semantic understanding of the text. More recently, the

Masked Autoencoder [27] has emerged as a powerful framework, demonstrating its efficacy across a variety of downstream tasks.

Discriminative SSL typically optimizes a discriminative loss to learn representations from sets of anchor, positive, and negative samples. Without ground-truth labels, these pairs are often constructed through solving jigsaw puzzles [41] or making geometry-based prediction [20]. A prominent example of discriminative SSL is contrastive learning, which brings samples from the same class closer while distancing those from different classes. Representative works such as MoCo [28] and SimCLR [10] learn embeddings by capturing invariant features between original data and its augmented variants. Additionally, InfoNCE [42], a method grounded in noise contrastive estimation [24], is widely used due to its effectiveness.

### 3 A Latent Representation of Scenario

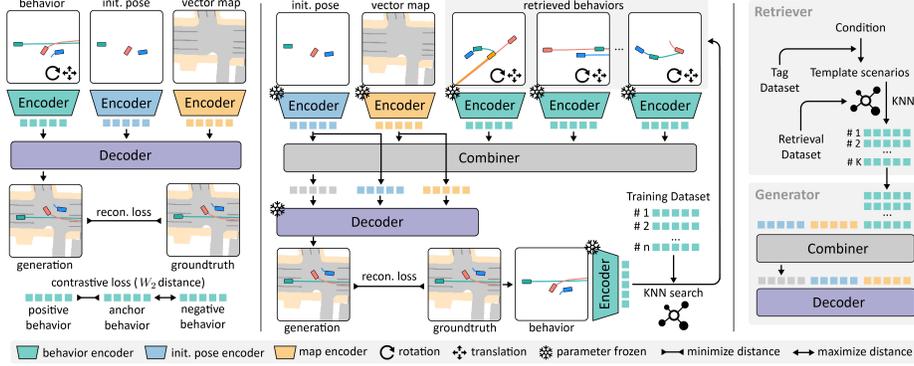
A crucial element within the retrieval system is the selection metric for data retrieval, which is typically implemented through a distance function between the query sample and the candidate samples in the database. Unlike text, which can be converted into word embeddings, traffic scenarios encompass sequential behaviors and intricate interactions among entities, complicating the establishment of a similarity metric for these scenarios. Consequently, in this section, we introduce a scenario autoencoder to extract latent representations that facilitate the assessment of similarity between various traffic scenarios, a vital component in our RAG framework.

#### 3.1 Scenario Definition

Each scenario is characterized by the trajectories  $\tau \in \mathbb{R}^{M \times T \times 5}$ , encompassing  $M$  agents over a maximum of  $T$  time steps. The trajectory of each agent is composed of parameters  $[x, y, v, c, s]$ , signifying position  $x$ , position  $y$ , velocity  $v$ , cosine of the heading  $c$ , and sine of the heading  $s$ . The initial state of these entities is denoted as  $\tau_0 \in \mathbb{R}^{M \times 5}$ . Furthermore, the map is encapsulated by  $m \in \mathbb{R}^{S \times 4}$ , composed of  $S$  lane segments. The attributes of these points  $[x_s, y_s, x_e, y_e]$  correspond to the starting and ending positions of each segment.

#### 3.2 Scenario Autoencoder

Autoencoders [44] are used to learn compressed latent representations of high-dimensional data, where an encoder projects the data into latent code and a decoder reconstructs the code in the data space. Given that traffic scenarios encompass spatial and temporal dynamics from multiple agents, we design a hierarchical encoder structure alternating between spatial and temporal layers based on transformer architecture [51]. As depicted on the left of Figure 2, our design incorporates three encoders for behavior, map, and initial position.



**Fig. 2:** **Left:** the training pipeline for the encoders and decoder aimed at learning latent embeddings of scenarios. A contrastive loss is applied to behavior embeddings to ensure invariance to absolute positions. **Middle:** the training pipeline for the combiner with frozen encoder and decoder parameters. We use K-Nearest Neighbors (KNN) to retrieve scenarios similar to a template scenario in the dataset and use the retrieved behaviors to reconstruct the template scenario. **Right:** the generation pipeline with a retriever and a generator.

First, we design a behavior encoder  $E_b$  with a spatial-temporal transformer structure that comprises  $L_e$  temporal transformer encoders ( $\text{Encoder}_t^i$ ) and  $L_e$  spatial transformer encoders ( $\text{Encoder}_s^i$ ), where  $i \in \{1, \dots, L_e\}$ . The trajectory  $\tau$  is first transformed into a latent embedding  $z_b$ , where  $H$  denotes the hidden dimension, through a Multiple Linear Perception (MLP) module. Subsequently, an alternating encoding procedure extracts spatial and temporal features from  $z_b$ . To preserve temporal information, we add a sinusoidal positional embedding (PE) to the embedding before employing the temporal transformer. To retain distinct agent information and further compress the behavior feature, we calculate the mean of the latent embedding across the temporal dimension, resulting in the final behavior embedding  $z_b \in \mathbb{R}^{M \times H}$ . Besides the behavior encoder, we also process the map information represented by lane vectors with the attention mechanism [51]. Similarly to the first step in  $E_b$ , we project  $m$  to the latent space with an MLP module. Then we apply a multi-head attention module (MHA) with layer normalization (LN) [1] to acquire the map embedding  $z_m \in \mathbb{R}^{S \times H}$  with a learnable query embedding  $q_m$ .

In the decoder, a spatial-temporal transformer architecture, similar to the behavior encoder, is designed to obtain the behavior embedding  $z_b$  through  $L_d$  times encoding. Given that  $z_b$  lacks the temporal dimension, we first replicate it  $T$  times and then add a PE before entering it into the temporal transformer encoder. Throughout the decoding process, the map embedding  $z_m$  is injected by a cross-attention mechanism:  $z_b \leftarrow z_b + \text{MHA}(z_b, z_m, z_m)$ .  $\text{MHA}(Q, K, V)$  is the multi-head attention with  $Q, K, V$  representing query, key, and value:

$$\begin{aligned} \text{MHA}(Q, K, V) &= \text{Concatenate}(h_1, \dots, h_i)W^O, \\ h_i &= \text{Attention}(QW_i^Q, KW_i^K, VW_i^V), \end{aligned} \quad (1)$$

**Algorithm 1:** Details of Encoder and Decoder

---

<pre> 1 <b>Behavior Encoder</b> <math>E_b(\tau)</math>: 2   <math>z_b \leftarrow \text{MLP}(\tau)</math> // projection 3   <b>for</b> <math>i</math> in <math>[1, \dots, L_e]</math> <b>do</b> 4     <math>z_b \leftarrow \text{Encoder}_t^i(\text{PE} + z_b)</math> 5     <math>z_b \leftarrow \text{Encoder}_s^i(z_b)</math> 6   <math>z_b \leftarrow \text{mean}(z_b)</math> 7   <b>return</b> behavior embedding <math>z_b</math> 8 <b>Map Encoder</b> <math>E_m(m)</math>: 9   Initialize a learnable query <math>q_m</math> 10  <math>z_m \leftarrow \text{MLP}(m)</math> // projection 11  <math>z_m \leftarrow \text{LN}(\text{MHA}(q_m, z_m, z_m))</math> 12  <math>z_m \leftarrow \text{LN}(z_m + \text{MLP}(z_m))</math> 13  <b>return</b> map embedding <math>z_m</math> </pre>	<pre> 14 <b>Initial Pose Encoder</b> <math>E_i(\tau_0)</math>: 15  <math>z_i \leftarrow \text{MLP}(\tau_0)</math> // projection 16  <b>return</b> initial pose embedding <math>z_i</math> 17 <b>Decoder</b> <math>D(z_i, z_b, z_m)</math>: 18  <math>z_r \leftarrow z_b + \text{MHA}(z_b, z_i, z_i)</math> 19  <b>for</b> <math>i</math> in <math>[1, \dots, L_d]</math> <b>do</b> 20    <math>z_r \leftarrow \text{Encoder}_t^i(\text{PE} + z_r)</math> 21    <math>z_r \leftarrow \text{Encoder}_s^i(z_r)</math> 22    <math>z_r \leftarrow z_r + \text{MHA}^i(z_r, z_m, z_m)</math> 23  <math>\hat{\tau} \leftarrow \text{MLP}(z_r)</math> // projection 24  <b>return</b> reconstructed trajectory <math>\hat{\tau}</math> </pre>
---	---

---

where  $W^O$ ,  $W_i^Q$ ,  $W_i^K$ , and  $W_i^V$  are learnable parameters. The last component of the decoder is an MLP that projects the hidden embedding back into the data domain, resulting in the reconstructed trajectory  $\hat{\tau}$ . To train the encoders and decoder, the mean square error is used as the reconstruction loss, expressed as  $\mathcal{L}_r = |\hat{\tau} - \tau|_2$ . Comprehensive descriptions of the encoding and decoding processes are provided in Algorithm 1.

It is important to note that, while the current design of the autoencoder is sufficient for acquiring compressed representations of scenarios, these representations are not invariant to the absolute coordinates and the order of the agents. This limitation can lead to substantial embedding distances between similar scenarios. To address these issues, two enhancements are proposed and detailed in the following section.

### 3.3 Invariant feature with contrastive loss

To enhance the representation of scenario similarity in the behavior embedding  $z_b$ , we employ contrastive learning to acquire invariant features. Specifically, we integrate InfoNCE [42] as an additional loss function  $\mathcal{L}_c$ , which optimizes categorical cross-entropy to distinguish a positive sample from a batch of negative samples. In practice, for a query embedding  $z_b$ , a positive sample  $z_b^+$  is generated by applying random rotation and translation to the original scenario  $\tau, m$ . Meanwhile, negative samples  $Z_b^-$  are selected from the remaining samples in the same batch. In conventional InfoNCE, to bring  $z_b$  and  $z_b^+$  closer together, the inner product is calculated between the embedding  $z_b$  and the set  $z_b^+, Z_b^-$ , employing cross-entropy loss to identify the index of the positive sample.

To ensure that the ordering of the agents does not affect the outcomes,  $z_b \in \mathbb{R}^{M \times H}$  should exhibit permutation invariance across the dimension  $M$ . Otherwise, simply stacking them into a one-dimensional vector could erroneously represent distinct scenarios as significantly different. To establish permutation

invariance, we adopt the Wasserstein distance  $W_2$  [52] as a more suitable metric than cosine distance for assessing the similarity. Taking  $z_b$  as a distribution representing  $M$  individual behaviors, the intuition behind this choice is that the Wasserstein distance identifies the minimal adjustments necessary for behaviors in one scenario to reflect those in another. With such a powerful tool, we have the following contrastive loss:

$$\mathcal{L}_c = - \sum_{z_b} \log \frac{\exp[-W_2(z_b, z_b^+)]}{\sum_{z' \in \{z_b^+, z_b^-\}} \exp[-W_2(z_b, z')]}.$$
 (2)

The implementation of this loss results in the behavior embedding  $z_b$  lacking absolute coordinate data, making the decoder incapable of accurately reconstructing the precise trajectory. To handle this problem, we add the initial poses  $\tau_0$  of all agents as a supplementary input for the decoder. We encode these poses into  $z_i$  with an MLP encoder  $\text{Enc}_i$ , and then integrate it into the decoder with an MHA module  $z_r \leftarrow z_b + \text{MHA}(z_b, z_i, z_i)$ . Additional specifics regarding the initial pose encoder are detailed in Algorithm 1. In the training stage, the objective is to minimize a combined loss function  $\mathcal{L} = \mathcal{L}_r + \lambda \mathcal{L}_c$ , where  $\lambda$  serves as a weighting factor for contrastive loss and is uniformly set at 0.1 for all experiments.

## 4 Retrieval Augmented Scenario Generation

The autoencoder structure introduced in Section 3 builds a one-to-one generation framework, which still needs additional modules to support the retrieval augmented generation, which is a many-to-one framework. In this section, we introduce a module called Combiner, which takes input behavior embeddings from multiple scenarios and outputs the combined embedding. To train this module, we designed a KNN-based training pipeline that forces the model to learn to combine and edit existing scenarios.

### 4.1 The Training of Combiner

Unlike the trajectory prediction task [45] which has ground truth as the optimization target, the objective in training the combiner is to learn the alignment of behaviors in the retrieved scenarios with the initial pose and the specified map. A well-trained combiner should be able to compose behaviors from these retrieved scenarios, thereby generating new scenarios that have a resemblance to all the retrieved scenarios. This type of objective aligns with the concept of meta-learning, or ‘learning to learn’, as described in [31].

Within our training framework, for a given query scenario  $\tau, m$  in the dataset, we initially utilize the behavior encoder to obtain  $z_b$  and then use KNN to identify  $K$  similar behavior embeddings from the database, denoted  $z_{ret} = [z_{b,1}, \dots, z_{b,K}]$ . Building on this, we propose a model comprising two MHA modules as the combiner:

$$\begin{aligned} z_{rag} &\leftarrow z_i + \text{MHA}(z_i, z_{ret}, z_{ret}), \\ z_{rag} &\leftarrow z_{rag} + \text{MHA}(z_{rag}, z_m, z_m), \end{aligned}$$
 (3)

where  $z_i \leftarrow E_i(x_0)$  and  $z_m \leftarrow E_m(m)$  are the initial pose embedding and the map embedding of the query scenario. Gradients are stopped for  $z_i$  and  $z_m$ . Assuming that the  $K$  nearest scenarios adequately represent the query scenario, it should be feasible to reconstruct the behavior of the query scenario,  $z_b$ , using the retrieved scenario,  $z_{rag}$ . Consequently, we employ the following loss function as our training goal:

$$\mathcal{L}_{rag} = \|D(z_i, z_{rag}, z_m) - \tau\|_2, \quad (4)$$

wherein the parameters of the decoder  $D\phi$  remain fixed during training. Essentially, this training method is designed to learn an "inverse" operation of the KNN, aiming to reconstruct a query scenario that closely resembles all  $K$  retrieved scenarios in the database.

## 4.2 The generation pipeline

The pipeline for scenario generation, as delineated in Figure 2, is categorized into retriever and generator components. The generation process is expedited by pre-processing all scenarios in the database using a behavior encoder, which yields behavior embeddings that facilitate efficient similarity computation.

The retriever enhances the versatility of generation by dividing the process into two stages. In the initial stage, users can choose a set of template scenarios that depict specific conditions. These may include manually annotated scenarios with tags denoting actions like left or right turns, enabling the generation of additional scenarios under similar tags or even a combination thereof. In addition, templates can include critical and interesting scenarios collected from real-world data. Solely relying on these templates for generation could be limited, which is mitigated by incorporating a secondary phase, which employs a KNN approach to fetch high-quality scenarios from a vast and unlabeled database to augment adaptability.

Subsequently, the generator, comprising a combiner and a decoder, follows the same inference as combiner training with the initial pose and lane map specified by the user. We obtain the RAG embedding  $z_{rag}$  with Eq. (3), and then infer the generated scenario through  $\tau_{rag} \leftarrow D(z_i, z_{rag}, z_m)$ .

## 5 Experiment

In this section, we start with an overview of the experimental setup and details regarding the implementation of the RealGEN framework. Then we discuss the baselines, including different types of autoencoders and variants of RealGen, used for comparison. In the discussion of the findings, we respectively evaluate the quality of scenario embedding and the generation capability of RAG.

### 5.1 Settings and Implementation Details

We conducted all training and evaluation with the nuScenes [5] dataset using the trajdata [34] package for data loading and processing. Each scenario spans

a duration of 8 seconds, operating at a frequency of 2 Hz, and encompasses a maximum of 11 agents. We filter out agents that travel less than 3 meters in 8 seconds and select 11 agents closest to the ego vehicle. The map contains 100 lanes (each lane has 20 points) that are ordered according to the distance between the center of the lane and the center of the ego vehicle.

To train RealGen, we use Adam [36] as the optimizer to update the parameters of the autoencoder and the combiner. To make the calculation of the Wasserstein distance efficient, we use the Sinkhorn distance [13], an entropy regularized approximation of the Wasserstein distance [52], with implementation in the GeomLoss [19] package.

## 5.2 Baselines

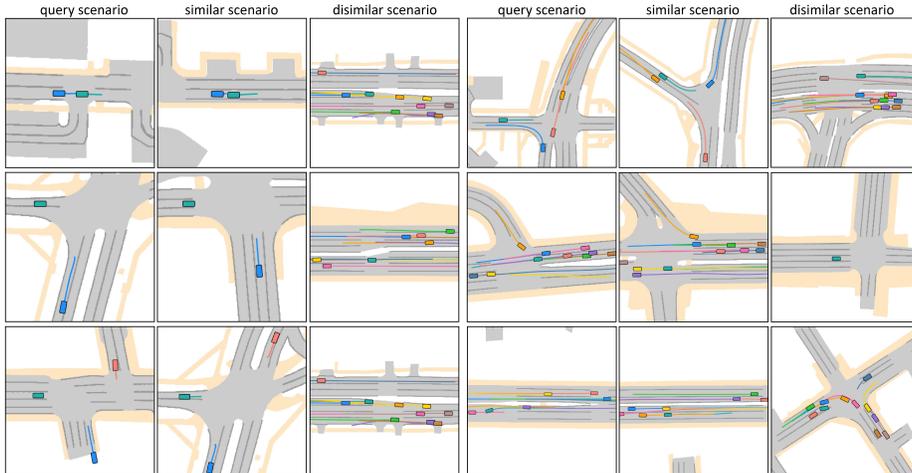
In the experimental section, we evaluate the following reconstruction-based generative models as baselines for comparison. Autoencoder (**AE**) shares the same behavior encoder, map encoder, and decoder structures as RealGen, serving as the most straightforward baseline for scenario reconstruction. **Contrastive AE** mirrors the structure of the RealGen autoencoder but omits the initial pose as absolute information. **Masked AE** is a self-supervised learning baseline, which has been investigated for trajectory data as described in [9, 12, 56, 58]. To evaluate the controllable generation capability, we select a state-of-the-art model **LCTGen** [49] as the baseline, which takes a high-level agent representation  $z$  and a vector map  $m$  as input. We also consider **LCTGen w/o  $z$** , a variant model proposed in the original paper to show the reconstruction results without using the agent information. Given that our method has two phases, we refer to the autoencoder component specifically as **RealGen-AE** and the complete model as **RealGen**.

## 5.3 Evaluation of Behavior Embedding

We first evaluate the quality of the learned representation of behavior, which is critical for the following retrieval and generation processes.

**Visualizing similar and dissimilar scenarios.** After training the autoencoder with contrastive loss, the distance between behavior embeddings can be used as an indicator of the similarity of the two scenarios. To validate this statement, we visualize qualitative examples of using a query to find the most similar (minimal  $W_2$  distance) and dissimilar (maximal  $W_2$  distance) scenarios in Figure 3. We observe that the most similar scenario contains the same behavior and number of vehicles as in the query scenario.

**Classifying Scene ID with behavior embedding.** To further provide quantitative results on how well the behavior embedding encodes behavior information, we leverage the Scene ID of the nuScenes dataset. Each scene typically lasts 20 seconds and our scenario segments last 8 seconds, so we get multiple segments belonging to the same Scene ID. We assume that the segments having the same Scene ID have similar behaviors so that we can calculate the accuracy by using the distance to find the closest segment. In Figure 4, we summarize



**Fig. 3:** Qualitative evaluation of similar and dissimilar scenarios calculated by our scenario embedding. Rectangles represent the initial poses of vehicles and the lines represent the future trajectories.

the results in the left part, where top- $k$  means we calculate the accuracy for the closest  $k$  segments. To validate the effectiveness of  $W_2$  distance, we also permute the order of agents in the behavior embedding, which is denoted as *cosine-permuted* and  $W_2$ -*permuted*. We find that the cosine distance cannot deal with the permuted setting, as the order is important when stacking the embedding of agents. Our method performs well in top-1 and top-5 settings but badly in others, which could be explained by the fact that segments in one scene have very different behaviors. To validate this, we plot the distance matrix between segments for 11 scenes in the right part of Figure 4. We find sub-blocks in the diagonal blocks, indicating that segments are similar in a small time interval but could be different when the time interval is large.

**Linear probing of behavior embedding.** Linear probing, a commonly employed method to assess representations in self-supervised learning (SSL), involves training a linear classifier using the derived embeddings [27].

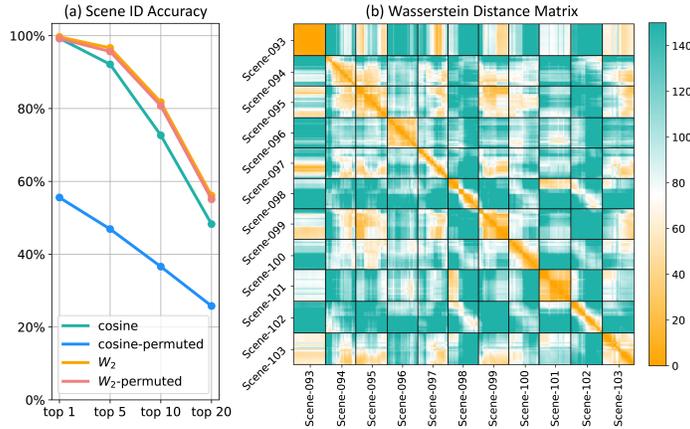
To train this classifier, we implemented heuristic rules to assign basic behavioral labels (acceleration, deceleration, stopping, keeping speed, left/right turn) to each agent’s embedding. The outcomes, along with comparisons with baseline models, are presented in Table 1. These findings demonstrate that the embeddings generated by RealGen surpass all baseline models in terms of accuracy.

**Table 1:** Accuracy of linear probing

AE	Contrastive AE	Masked AE	RealGen-AE
82.5%	67.1%	86.2%	<b>87.8%</b>

#### 5.4 Evaluation of Retrieval Augmented Generation

Evaluating scene-level controllability and quality of scenario generation is an open problem due to the lack of quantitative metrics. Following most previous



**Fig. 4:** (a) Scene ID accuracy using the behavior embedding with difference distance metrics. (b) A matrix shows the Wasserstein distance between scenario segments, where each block contains the segments that belong to the same Scene ID.

**Table 2:** Results of realism metrics. Recon.-based and Retrieval-based means using the target scenario and retrieved scenarios as input for generation, respectively.

Category	Method	mADE	mFDE	Speed	Heading	SCR	ORR
Recon.-based	AE	<b>0.18±0.03</b>	0.41±0.06	<b>0.04±0.01</b>	0.10±0.01	<b>0.02±0.00</b>	<b>0.02±0.00</b>
	Masked AE	0.16±0.01	<b>0.39±0.01</b>	<b>0.04±0.01</b>	<b>0.09±0.01</b>	0.03±0.00	<b>0.02±0.00</b>
	Contrastive AE	0.92±0.02	1.47±0.04	0.12±0.00	0.36±0.02	0.04±0.00	0.04±0.00
	RealGen-AE	0.31±0.01	0.53±0.01	0.08±0.00	0.15±0.01	0.03±0.00	<b>0.02±0.00</b>
Retrieval-based	AE-KNN	14.3±0.03	16.4±0.05	0.57±0.01	0.59±0.02	0.15±0.01	0.15±0.01
	LCTGen	4.76±0.09	6.24±0.08	0.52±0.06	0.57±0.03	0.07±0.01	0.07±0.01
	LCTGen w/o $z$	14.2±0.07	16.7±0.09	2.04±0.04	1.42±0.00	0.16±0.02	0.13±0.04
	RealGen-AE-KNN	13.1±0.06	14.1±0.03	0.46±0.01	0.44±0.00	0.12±0.01	0.11±0.00
	RealGen	<b>1.54±0.04</b>	<b>1.21±0.03</b>	<b>0.21±0.03</b>	<b>0.21±0.01</b>	<b>0.05±0.00</b>	<b>0.04±0.00</b>

work [18, 49, 64], we provide the results of the realism metrics and show the qualitative results of using RealGen for various downstream tasks. More results can be found in the appendix.

**Realism of generated scenario.** To evaluate the realism of generated scenarios, we consider the following metrics and show the results in Table 2. We use the maximum mean discrepancy (MMD) [22] to measure the similarity in velocity and direction between the original scenarios and the generated scenarios. We also compare the mean average displacement error (mADE) and the mean final displacement error (mFDE) for the average reconstruction performance. To evaluate scene-level realism, we calculate the scene collision rate (SCR) and the off-road rate (ORR) following the metrics defined in [49]. The recon-based generation methods in Table 2 use the behavior of the target scenario as input. For this category, we compare RealGen-AE, only using the encoder and decoder modules, with three baseline methods. Due to the additional contrastive term, RealGen-AE achieves slightly worse performance than AE and Masked AE. However, RealGen is designed for retrieval-based generation, which uses retrieved scenarios rather than the target scenario as input. To fairly compare the generation performance, we design two baselines named AE-KNN and RealGen-AE-KNN,

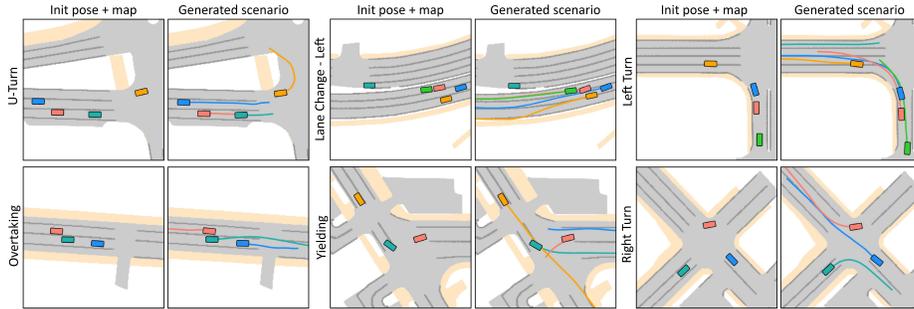


Fig. 5: Examples of tag-retrieved scenarios generated by RealGen for six different tags.

which use KNN to find the most similar behavior embedding to the target scenario and use this embedding as input to the decoder for generation. According to the results, we find that RealGen achieves comparable performance as recombined generation and outperforms baselines, which indicates the important role of the combiner in fusing the information of the retrieved scenarios.

**Generating tag-retrieved scenarios.** Now we explore the qualitative performance of using RealGen for tag-retrieved generation. Given a target behavior tag, we first obtain several template scenarios from a small dataset and use them to retrieve more scenarios from the training database, which will be used for generation in the combiner. Since there is no existing dataset with tags, we manually labeled six tags – U-Turn, Overtaking, Left Lane Change, Right Lane Change, Left Turn, Right Turn – for the nuScenes dataset to get template 1349 scenarios. We plot generated scenarios for each tag in Figure 5, where the left part of each example shows the given initial pose and map and the right part of each example shows the generated scenario from RealGen.

**Safety-critical scenario generation.** Beyond the tag mentioned above, RealGen demonstrates its capacity for in-context

learning by generating critical and unseen crash scenarios, divergent from those of the training datasets. This is initiated by manually creating several crash scenario templates, guided by the scenarios recorded in the NHTSA Crash Report [40]. Subsequently, we generate crash scenarios using existing initial poses and maps of the dataset. Figure 6 illustrates six instances in which the shadowed rectangles denote the initial positions of agents and the red box highlights

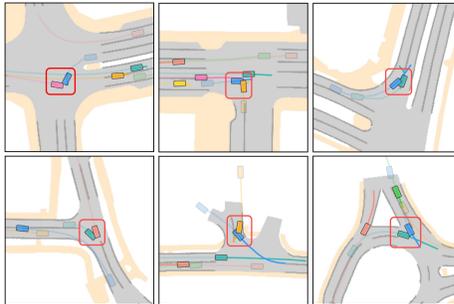


Fig. 6: Examples of generating crash scenarios from RealGen. The shadow rectangles represent the initial positions of agents.

Table 3: Safety-critical scenario generation.

Method	Realgen-AE-R	RealGen-R	RealGen
Collision Rate ( $\uparrow$ )	0.92	0.83	<b>0.59</b>

**Table 4:** Results of human evaluation of controllability. (Details in Appendix ??)

Category	Left-Turn	Right-Turn	Left-Lane-Change	Right-Lane-Change	Straight
RealGen Preferred (%)	81.8	91.7	97.8	93.3	100.0
RealGen Score (0-5)	<b>4.27±1.05</b>	<b>4.27±0.69</b>	<b>3.96±1.94</b>	<b>4.17±1.93</b>	<b>3.94±2.27</b>
LCTGen Score (0-5)	2.15±1.44	2.08±1.19	2.42±1.95	2.0±1.96	2.14±2.31

the point of collision. To quantitatively assess the performance of safety-critical scenarios, we compare RealGen with two baselines. *RealGen-AE-R* random samples behavior embeddings in the Realgen-AE model, and *RealGen-R* random retrieves behavior embeddings. According to the results in Table 3, we find that the scenarios generated by RealGen using crash scenarios as templates achieve the highest collision rate, which means that RealGen has more efficiency.

**Human evaluation of controllability.** As there is no automatic way to evaluate controllability, we follow the protocol in [49] to perform A/B testing using human evaluation. We report the ratio of our method preferred as well as the absolute score (0-5) of both our method and the baseline. In Table 4, we find that scenarios generated by RealGen are highly preferred in most categories. The absolute score of RealGen is also much higher than that of LCTGen.

**Downstream task evaluation.** A direct downstream task of our method is to use the generated data to augment the training dataset of trajectory

**Table 5:** Downstream task evaluation.

Method	Original	Random Aug.	RealGen Aug.
mADE (↓)	3.544	2.920	<b>2.309</b>
collision rate (↓)	0.049	0.037	<b>0.018</b>

prediction models. We use Autobots [21] as a predictor and report the results trained on different datasets in Table 5. *Original* means using the original data in nuScenes [5], *Random Aug.* means augmenting the original dataset with Gaussian noise, and *RealGen Aug.* means augmenting the original dataset with scenarios from RealGen. We observe that the model trained with the RealGen dataset achieves the lowest mADE and collision rate.

## 6 Conclusion and Limitation

This paper proposes RealGen, a novel framework for traffic scenario generation that utilizes retrieval-augmented generation. Unlike previous approaches, which primarily rely on models replicating training distributions, RealGen demonstrates in-context learning abilities that synthesize scenarios by combining and modifying provided examples, enabling controlled generation. These scenarios can be automatically obtained from a retrieval system, which only requires the users to provide a few template scenarios as examples. We comprehensively evaluated the similarity autoencoder model for retrieval and the combiner model for generation. The findings indicate that RealGen achieves low reconstruction error and high generation quality.

The primary limitation of our current method is that the behavior encoder focuses solely on agent trajectories, neglecting the intricate interactions between agents and lane maps in complex behaviors. Advancing the feature representation within the behavior encoder could significantly broaden RealGen’s capacity to generate diverse and controllable traffic scenarios.

## Acknowledgments

Wenhao Ding contributed to this paper while being an intern at NVIDIA Research. Ding Zhao was partially supported by the National Science Foundation under grant CNS-2047454.

## References

1. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
2. Blattmann, A., Rombach, R., Oktay, K., Müller, J., Ommer, B.: Retrieval-augmented diffusion models. *Advances in Neural Information Processing Systems* **35**, 15309–15324 (2022)
3. Borgeaud, S., Mensch, A., Hoffmann, J., Cai, T., Rutherford, E., Millican, K., Van Den Driessche, G.B., Lespiau, J.B., Damoc, B., Clark, A., et al.: Improving language models by retrieving from trillions of tokens. In: *International conference on machine learning*. pp. 2206–2240. PMLR (2022)
4. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in neural information processing systems* **33**, 1877–1901 (2020)
5. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11621–11631 (2020)
6. Cai, D., Wang, Y., Bi, W., Tu, Z., Liu, X., Shi, S.: Retrieval-guided dialogue response generation via a matching-to-generation framework. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. pp. 1866–1875 (2019)
7. Cao, Y., Xiao, C., Anandkumar, A., Xu, D., Pavone, M.: Advdo: Realistic adversarial attacks for trajectory prediction. In: *European Conference on Computer Vision*. pp. 36–52. Springer (2022)
8. Chen, D., Fisch, A., Weston, J., Bordes, A.: Reading wikipedia to answer open-domain questions. arXiv preprint arXiv:1704.00051 (2017)
9. Chen, H., Wang, J., Shao, K., Liu, F., Hao, J., Guan, C., Chen, G., Heng, P.A.: Traj-mae: Masked autoencoders for trajectory prediction. arXiv preprint arXiv:2303.06697 (2023)
10. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *International conference on machine learning*. pp. 1597–1607. PMLR (2020)
11. Chen, W., Hu, H., Saharia, C., Cohen, W.W.: Re-imagen: Retrieval-augmented text-to-image generator. arXiv preprint arXiv:2209.14491 (2022)
12. Cheng, J., Mei, X., Liu, M.: Forecast-mae: Self-supervised pre-training for motion forecasting with masked autoencoders. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 8679–8689 (2023)
13. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems* **26** (2013)
14. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)

15. Ding, W., Chen, B., Li, B., Eun, K.J., Zhao, D.: Multimodal safety-critical scenarios generation for decision-making algorithms evaluation. *IEEE Robotics and Automation Letters* **6**(2), 1551–1558 (2021)
16. Ding, W., Lin, H., Li, B., Eun, K.J., Zhao, D.: Semantically adversarial driving scenario generation with explicit knowledge integration. *arXiv preprint arXiv:2106.04066* (2021)
17. Ding, W., Xu, M., Zhao, D.: Learning to collide: An adaptive safety-critical scenarios generating method. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE (2020)
18. Feng, L., Li, Q., Peng, Z., Tan, S., Zhou, B.: Trafficgen: Learning to generate diverse and realistic traffic scenarios. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 3567–3575. IEEE (2023)
19. Feydy, J., Sejourne, T., Vialard, F.X., Amari, S.i., Troune, A., Peyre, G.: Interpolating between optimal transport and mmd using sinkhorn divergences. In: *The 22nd International Conference on Artificial Intelligence and Statistics*. pp. 2681–2690 (2019)
20. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728* (2018)
21. Girgis, R., Golemo, F., Codevilla, F., Weiss, M., D’Souza, J.A., Kahou, S.E., Heide, F., Pal, C.: Latent variable sequential set transformers for joint multi-agent motion prediction. *arXiv preprint arXiv:2104.00563* (2021)
22. Gretton, A., Borgwardt, K.M., Rasch, M.J., Scholkopf, B., Smola, A.: A kernel two-sample test. *The Journal of Machine Learning Research* **13**(1), 723–773 (2012)
23. Gulino, C., Fu, J., Luo, W., Tucker, G., Bronstein, E., Lu, Y., Harb, J., Pan, X., Wang, Y., Chen, X., et al.: Waymax: An accelerated, data-driven simulator for large-scale autonomous driving research. *arXiv preprint arXiv:2310.08710* (2023)
24. Gutmann, M., Hyvarinen, A.: Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In: *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. pp. 297–304. *JMLR Workshop and Conference Proceedings* (2010)
25. Guu, K., Lee, K., Tung, Z., Pasupat, P., Chang, M.: Retrieval augmented language model pre-training. In: *International conference on machine learning*. pp. 3929–3938. PMLR (2020)
26. Hanselmann, N., Renz, K., Chitta, K., Bhattacharyya, A., Geiger, A.: King: Generating safety-critical driving scenarios for robust imitation via kinematics gradients. In: *European Conference on Computer Vision*. pp. 335–352. Springer (2022)
27. He, K., Chen, X., Xie, S., Li, Y., Dollar, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 16000–16009 (2022)
28. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9729–9738 (2020)
29. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
30. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural computation* **9**(8), 1735–1780 (1997)
31. Hospedales, T., Antoniou, A., Micaelli, P., Storkey, A.: Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence* **44**(9), 5149–5169 (2021)

32. Hu, A., Corrado, G., Griffiths, N., Murez, Z., Gurau, C., Yeo, H., Kendall, A., Cipolla, R., Shotton, J.: Model-based imitation learning for urban driving. *Advances in Neural Information Processing Systems* **35**, 20703–20716 (2022)
33. Ibrihich, S., Oussous, A., Ibrihich, O., Esghir, M.: A review on recent research in information retrieval. *Procedia Computer Science* **201**, 777–782 (2022)
34. Ivanovic, B., Song, G., Gilitschenski, I., Pavone, M.: trajdata: A unified interface to multiple human trajectory datasets. *arXiv preprint arXiv:2307.13924* (2023)
35. Kim, J., Choi, S., Amplayo, R.K., Hwang, S.w.: Retrieval-augmented controllable review generation. In: *Proceedings of the 28th International Conference on Computational Linguistics*. pp. 2284–2295 (2020)
36. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
37. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.t., Rocktäschel, T., et al.: Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* **33**, 9459–9474 (2020)
38. Li, Q., Peng, Z., Feng, L., Duan, C., Mo, W., Zhou, B., et al.: Scenarionet: Open-source platform for large-scale traffic scenario simulation and modeling. *arXiv preprint arXiv:2306.12241* (2023)
39. Liu, Q., Yogatama, D., Blunsom, P.: Relational memory-augmented language models. *Transactions of the Association for Computational Linguistics* **10**, 555–572 (2022)
40. NHTSA: Nhtsa crash viewer (2023), <https://crashviewer.nhtsa.dot.gov/>
41. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: *European conference on computer vision*. pp. 69–84. Springer (2016)
42. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018)
43. Roshdi, A., Roohparvar, A.: Information retrieval techniques and applications. *International Journal of Computer Networks and Communications Security* **3**(9), 373–377 (2015)
44. Schmidhuber, J.: Deep learning in neural networks: An overview. *Neural networks* **61**, 85–117 (2015)
45. Shi, S., Jiang, L., Dai, D., Schiele, B.: Motion transformer with global intention localization and local movement refinement. *Advances in Neural Information Processing Systems* **35**, 6531–6543 (2022)
46. Shurrab, S., Duwairi, R.: Self-supervised learning methods and applications in medical imaging analysis: A survey. *PeerJ Computer Science* **8**, e1045 (2022)
47. Suo, S., Regalado, S., Casas, S., Urtasun, R.: Trafficsim: Learning to simulate realistic multi-agent behaviors. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10400–10409 (2021)
48. Suo, S., Wong, K., Xu, J., Tu, J., Cui, A., Casas, S., Urtasun, R.: Mixsim: A hierarchical framework for mixed reality traffic simulation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9622–9631 (2023)
49. Tan, S., Ivanovic, B., Weng, X., Pavone, M., Kraehenbuehl, P.: Language conditioned traffic generation. In: *Conference on Robot Learning*. pp. 2714–2752. PMLR (2023)
50. Tan, S., Wong, K., Wang, S., Manivasagam, S., Ren, M., Urtasun, R.: Scenegen: Learning to generate realistic traffic scenes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 892–901 (2021)

51. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
52. Villani, C., et al.: *Optimal transport: old and new*, vol. 338. Springer (2009)
53. Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.A.: Extracting and composing robust features with denoising autoencoders. In: *Proceedings of the 25th international conference on Machine learning*. pp. 1096–1103 (2008)
54. Wang, J., Pun, A., Tu, J., Manivasagam, S., Sadat, A., Casas, S., Ren, M., Urtasun, R.: Advsim: Generating safety-critical scenarios for self-driving vehicles. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9909–9918 (2021)
55. Wang, Z., Nie, W., Qiao, Z., Xiao, C., Baraniuk, R., Anandkumar, A.: Retrieval-based controllable molecule generation. *arXiv preprint arXiv:2208.11126* (2022)
56. Wu, P., Majumdar, A., Stone, K., Lin, Y., Mordatch, I., Abbeel, P., Rajeswaran, A.: Masked trajectory models for prediction, representation, and control. *arXiv preprint arXiv:2305.02968* (2023)
57. Xu, P., Patwary, M., Shoeybi, M., Puri, R., Fung, P., Anandkumar, A., Catanzaro, B.: Megatron-cntrl: Controllable story generation with external knowledge using large-scale language models. *arXiv preprint arXiv:2010.00840* (2020)
58. Yang, Y., Zhang, Q., Gilles, T., Batool, N., Folkesson, J.: Rmp: A random mask pretrain framework for motion prediction. *arXiv preprint arXiv:2309.08989* (2023)
59. Yang, Z., Li, L., Lin, K., Wang, J., Lin, C.C., Liu, Z., Wang, L.: The dawn of lmms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421* (2023)
60. Zhang, C., Zhang, C., Song, J., Yi, J.S.K., Zhang, K., Kweon, I.S.: A survey on masked autoencoder for self-supervised learning in vision and beyond. *arXiv preprint arXiv:2208.00173* (2022)
61. Zhang, C., Tu, J., Zhang, L., Wong, K., Suo, S., Urtasun, R.: Learning realistic traffic agents in closed-loop. In: *7th Annual Conference on Robot Learning* (2023)
62. Zhang, X., Bosselut, A., Yasunaga, M., Ren, H., Liang, P., Manning, C.D., Leskovec, J.: Greaselm: Graph reasoning enhanced language models. In: *International conference on learning representations* (2021)
63. Zhong, Z., Rempe, D., Chen, Y., Ivanovic, B., Cao, Y., Xu, D., Pavone, M., Ray, B.: Language-guided traffic simulation via scene-level diffusion. *arXiv preprint arXiv:2306.06344* (2023)
64. Zhong, Z., Rempe, D., Xu, D., Chen, Y., Veer, S., Che, T., Ray, B., Pavone, M.: Guided conditional diffusion for controllable traffic simulation. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 3560–3566. IEEE (2023)