

Hallucinating 360°: Panoramic Street-View Generation via Local Scenes Diffusion and Probabilistic Prompting

Fei Teng^{1,*}, Kai Luo^{1,*}, Sheng Wu^{1,*}, Siyu Li¹, Pujun Guo¹, Jiale Wei²,
Jiaming Zhang¹, Kunyu Peng^{2,3}, and Kailun Yang^{1,†}

Abstract—Panoramic perception holds significant potential for autonomous driving, enabling vehicles to acquire a comprehensive 360° surround view in a single shot. However, autonomous driving is a data-driven task. Complete panoramic data acquisition requires complex sampling systems and annotation pipelines, which are time-consuming and labor-intensive. Although existing street view generation models have demonstrated strong data regeneration capabilities, they can only learn from the fixed data distribution of existing datasets and cannot leverage noisy ground truth as a supervisory signal. In this paper, we propose the first panoramic generation method Percep360 for autonomous driving. Percep360 enables coherent generation of panoramic data with control signals based on the stitched panoramic data. Percep360 focuses on two key aspects: coherence and controllability. Specifically, to overcome the inherent information loss caused by the pinhole sampling process, we propose the Local Scenes Diffusion Method (LSDM). LSDM reformulates the panorama generation as a spatially continuous diffusion process, bridging the gaps between different data distributions. Additionally, to achieve the controllable generation of panoramic images, we propose a Probabilistic Prompting Method (PPM). PPM dynamically selects the most relevant control cues, enabling controllable panoramic image generation. We evaluate the effectiveness of the generated images from three perspectives: image quality assessment (*i.e.*, no-reference and with reference), controllability, and their utility in real-world Bird’s Eye View (BEV) segmentation. Notably, the generated data consistently outperforms the original stitched images in no-reference quality metrics and enhances downstream perception models, leading to an improvement of 2.5% in mIoU for panoramic BEV segmentation. The source code will be publicly available at <https://github.com/FeiT-FeiTeng/Percep360>.

I. INTRODUCTION

Autonomous vehicles require the integration of surrounding environmental information to achieve 360-degree perception. In this context, panoramic cameras have garnered increasing attention from the research community due to their large Field of View (FoV) [1], [2], [3]. Numerous

This work was supported in part by the National Natural Science Foundation of China (Grant No. 62473139), in part by the Hunan Provincial Research and Development Project (Grant No. 2025QK3019), in part by the Hunan Provincial Innovation Foundation For Postgraduate (Grant No. CX20250579) and in part by the State Key Laboratory of Autonomous Intelligent Unmanned Systems (the opening project No. ZZKF2025-2-10).

¹F. Teng, K. Luo, S. Wu, S. Li, P. Guo, J. Zhang, and K. Yang are with the School of Artificial Intelligence and Robotics and the National Engineering Research Center of Robot Visual Perception and Control Technology, Hunan University, Changsha 410082, China (email: kailun.yang@hnu.edu.cn).

²J. Wei and K. Peng are with the Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany.

³K. Peng is also with INSAIT, Sofia University “St. Kliment Ohridski”, Sofia 1784, Bulgaria.

[†]Corresponding author: Kailun Yang.

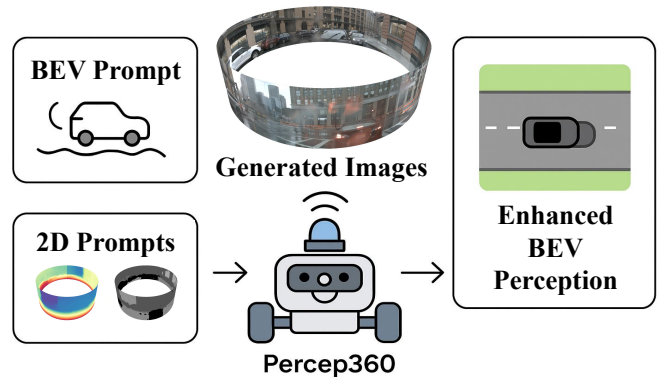


Fig. 1: Percep360 leverages diverse prompts to generate coherent panoramic images. The generated images exhibit improved quality and controllability. As a data generator, Percep360 synthesizes diverse data to avoid the expensive process of data processing, thereby enhancing the performance of semantic mapping.

studies have employed panoramic cameras for downstream perception tasks for street views, such as optical flow estimation [4], [5], object detection [6], [7], and semantic segmentation [1], [8], [9]. However, as a typical data-driven task, developing safe and effective autonomous driving systems heavily relies on the scale and diversity of the collected datasets. Due to the challenges in constructing panoramic data acquisition and annotation systems, the development of panoramic autonomous driving has been hindered by the scarcity of large-scale datasets. OneBEV [10] has attempted to use stitching-based methods to project pinhole images onto a spherical surface, thus simulating panoramic distortion, avoiding the need for data acquisition and offering a new feasible paradigm for panoramic autonomous driving. However, the stitching methods are unable to produce an infinite number of images from prompt input and fail to address the information misalignment caused by the sampling process.

The recent success of generative models for pinhole cameras has introduced an economic paradigm to avoid the high costs associated with data collection and annotation. Several models [11], [12], [13], [14] have proposed generation frameworks trained on prompt-image pairs, enabling fine-grained control over the diffusion process. In a similar vein, several works [15], [16], [17], [18] have emerged, offering cost-effective solutions for closed-loop autonomous driving data generation. However, existing methods that utilize pinhole samples to facilitate scene and style transfer in driving scenarios fail to ensure coherence in panoramic image synthesis. We modify the existing models for the panoramic domain. Those networks cannot leverage noisy

ground truth as a supervisory signal and learn the distribution of stitched images, including misalignment caused by the stitching process and inherent information loss introduced by the pinhole sampling process, as shown in Fig. 2. This can lead to perception errors and localization inaccuracies, further impairing the accuracy of environmental modeling and ultimately affecting the practical applications of robots.

Thus, effectively generating coherent panoramic images remains a crucial challenge for panoramic autonomous driving systems.

To address this problem, we propose the first solution: **Percep360**. Percep360 achieves coherent panoramic data generation by leveraging stitched pinhole images as a supervisory signal. As shown in Fig. 2(a), due to sampling errors from cameras, spatial misalignments occur when stitching multiple images into a panoramic view. Furthermore, we use BEVControl [14] as the baseline for controllable image generation. The generated images inherit the sampling errors and fail to produce consistent panoramic views. Meanwhile, since there is no supervisory guidance for coherence in the stitched region, the network inevitably alters the data distribution in that area when attempting to achieve coherence, which inevitably reduces the controllability of the image. From this perspective, ① to achieve coherence in panoramic images, we employ *Local Scenes Diffusion Method* (LSDM). LSDM circularly adjusts the correspondence between the network and the image-prompt pair data, leveraging locally coherent image features to bridge the information loss inherent to the pinhole sample. Furthermore, unlike panoramic image generation, street view synthesis must preserve controllability while maintaining image quality. As shown in Fig. 2(b), the spatial positions of roads and flower beds have changed. To address the decline in controllability resulting from compensation for sample errors, ② we introduce *Probabilistic Prompting Method* (PPM), which allows dynamic adjustment of prompt feature cues across different spatial regions, thereby achieving controllable generation. We evaluate the generated images on the nuScenes-360 dataset [10] from three perspectives: image quality assessment, no-reference image quality assessment, and controllability. Our method is the first to specifically target panoramic street-view image generation. It consistently improves both image quality and controllability, while boosting panoramic BEV segmentation performance by 2.5% in mIoU.

At a glance, we deliver the following contributions:

- This paper proposes Percep360, the first panoramic generation framework for autonomous driving. Percep360 enables the generation of coherent and controllable panoramic data using inconsistent stitched imagery.
- To fulfill the inherent information loss caused by the pinhole sampling process, we introduce the Local Scenes Diffusion Method (LSDM). LSDM leverages local information to realize a spatially continuous diffusion process, generating coherent panoramic scenes.
- To mitigate the controllability drop that arises when compensating for sample errors, we propose a

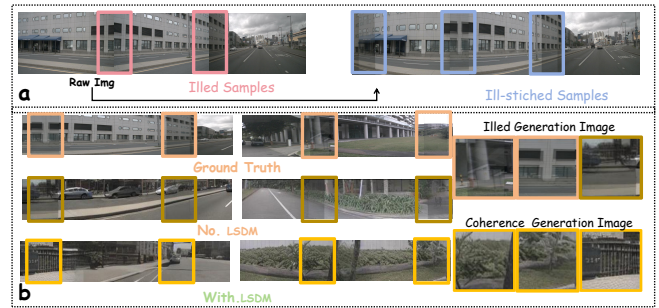


Fig. 2: Visualizations of the pinhole samples, stitched image, ground truth for training generation model, and the results of the existing approach are presented. Owing to the inherent information loss and misalignment in pinhole camera sampling, stitching-based methods struggle to produce coherent panoramic images. Moreover, when these stitched images are used as inputs, existing generation models often inherit the stitching errors. In contrast, the proposed LSDM approach enables the generation of images with improved coherence, particularly at stitching boundaries.

Probabilistic-Prompt Method (PPM). PPM adjusts the relationships among visual, textual, and other prompts, resulting in controllable panoramic generation.

- Our experiments demonstrate that Percep360 enables coherent and controllable panoramic scene generation. Moreover, the synthetic data it produces leads to an improvement of 2.5% in the performance of the downstream perception task of panoramic BEV segmentation.

II. RELATED WORK

To address the scarcity of data in panoramic autonomous driving, we explore a cost-effective data synthesis paradigm from the perspective of driving scenes generation. Furthermore, to achieve coherent data generation and compensate for stitched image errors, we investigate existing methods for panoramic image synthesis.

Driving Scenes Generation. Autonomous driving is a data-driven task, and data collection is labor-intensive. Leveraging generative paradigms to address information scarcity in the autonomous driving process has become an increasingly popular approach in recent research. BEVGen [13] utilizes BEV maps to generate multi-view images, while BEVControl [14] introduces a two-stage training paradigm to enhance the distinction between foreground and background, thereby improving the accuracy of generated images. Based on that, models such as DriveScape [19], DrivingDiffusion [20], HoloDrive [21], and MagicDrive [18] introduce bounding boxes to further enhance the controllability of the generated content. Meanwhile, Text2Street [22] leverages object and road topology prompts to achieve controllable single-image generation. Additionally, several works [23], [24], [25], [26], [27] propose the generation of multi-view videos using cross-frame methods. However, these works focus on generating controllable multi-view pinhole images under data-rich and clean conditions. In contrast, our goal is to generate coherent and controllable panoramic street-view images using imperfectly stitched images.

Panoramic Image Synthesis. The synthesis of panoramic images can be categorized into two aspects: text-based gen-

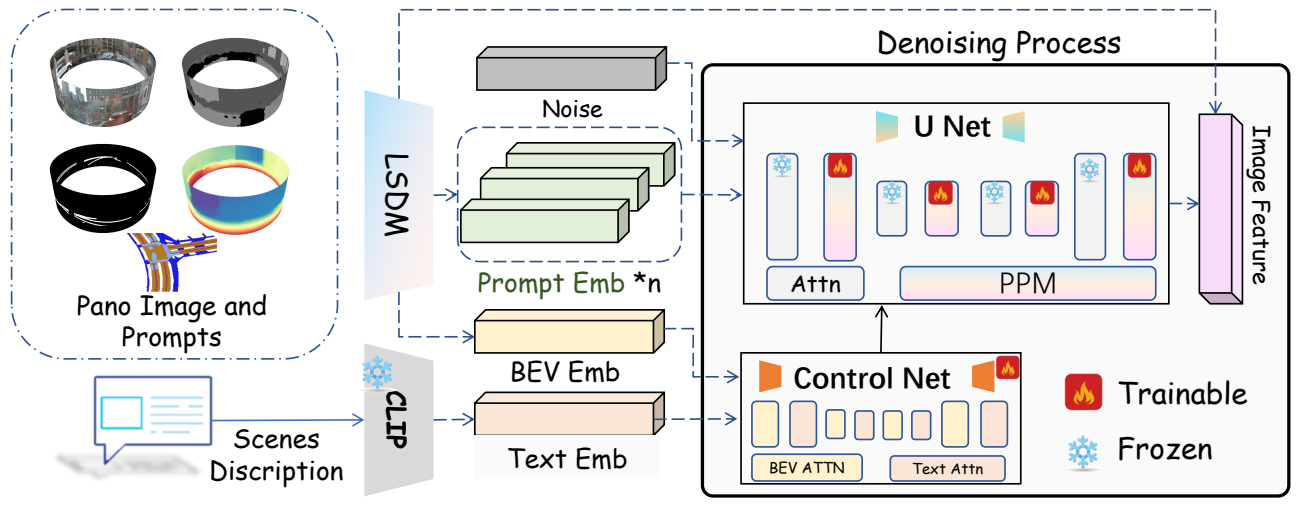


Fig. 3: The overall architecture of Percep360 leverages the BEV map, textual prompt, depth map, and mask map jointly as guidance signals. With the integration of the LSDM and PPM modules, the framework achieves coherent and controllable panoramic generation.

eration and completion based on partial image information. **Generation:** Text-to-image synthesis [28], [29], [30], [31] has rapidly emerged as a prominent area of research. For instance, PanFusion [32] employs a dual-branch architecture to seamlessly integrate additional constraints, enabling more customizable and high-quality panoramic outputs. Several works [33], [34], [35] achieve panoramic image completion by extending the field of view, enabling seamless and realistic scene expansion. **Completion:** The image inpainting task aims to restore missing regions of an image while maintaining overall coherence. PICFormer [36] introduces the Codebook-based sharing paradigm and enhances restoration performance through plus-realism image completion [37]. RePaint [38] models the semantics and texture of missing regions by leveraging a denoising diffusion probabilistic model. BrushNet [39], [40] incorporates a mask branch to better distinguish masked regions. Several works [41], [42] introduce the rotation-based paradigm to address the continuity challenges in panoramic image completion, ensuring smoother and more coherent results. *However*, a fundamental limitation exists: These works are trained using consistent images. Such a paradigm is not suitable for generating coherent panoramic data from imperfect samples. Moreover, existing panoramic image synthesis methods utilize weak textual prompts, which fail to meet the requirements of autonomous driving for accurate road topology representation.

III. METHODOLOGY

In this section, we first introduce the basic process of diffusion and present the problem formulation of our work in Sec. III-A. Then, in Sec. III-B, we focus on describing the overall architecture of our network and the proposed method.

A. Preliminary and Problem Formulation

Preliminary: Latent Diffusion Models (LDM) [43] represent traditional diffusion models by conducting the diffusion process within a compact latent representation space, typically involving a two-step procedure: feature compression

and iterative denoising for modeling data distributions. In the feature compression, input images in the pixel space are first transformed into compact latent embeddings using a perceptual compression model E , significantly reducing the complexity associated with processing high-dimensional data: $z = E(x)$. This latent representation z passes through the diffusion process and can be decoded back into the original image domain using the decoder D , formally expressed as: $\hat{x} = D(z)$. In the iterative denoising for modeling data distributions, the diffusion process is carried out within this latent space. In the forward diffusion stage, gaussian noise is gradually added to the latent embeddings, progressively transforming the latent representation into a distribution close to a standard Gaussian: $z_t = \sqrt{\bar{\alpha}_t}z + \sqrt{1 - \bar{\alpha}_t}\epsilon$, $\epsilon \sim \mathcal{N}(0, I)$. Here, $\bar{\alpha}_t$ is a pre-defined noise scheduling parameter, and t denotes the diffusion timestep. In the reverse denoising stage, a U-Net is trained to estimate and remove the added noise from the noisy latent representation z_t . This network is parameterized as ϵ_θ , trained by minimizing the following mean-square error loss: $\min_\theta \mathbb{E}_{z, c, \epsilon, t} [\|\epsilon - \epsilon_\theta(z_t, t, c)\|^2]$, where c represents conditional information, guiding the conditional image generation. After training, the LDM can generate samples by iterative denoising a latent feature initialized from Gaussian noise. Formally, this iterative sampling process can be expressed as: $p_\theta(z_{t-1}|z_t, c) = \mathcal{N}(z_{t-1}; \mu_\theta(z_t, t, c), \Sigma_\theta(z_t, t, c))$, where μ_θ and Σ_θ are predicted by the denoising network.

Problem Formulation: Given a panoramic image defined over a pixel domain $\Omega \subset \mathbb{R}^2$, the entire data distribution of the image is denoted by $f(x)$, where $x \in \Omega$. The panoramic image is constructed by stitching multiple pinhole camera views, resulting in a composite pixel space partitioned into twelve regions: six coherent regions $\{\mathcal{R}_i\}_{i=1}^6$ and six aliased regions $\{\mathcal{A}_i\}_{i=1}^6$. Thus, the domain can be expressed as: $\Omega = \left(\bigcup_{i=1}^6 \mathcal{R}_i\right) \cup \left(\bigcup_{i=1}^6 \mathcal{A}_i\right)$, $\mathcal{R}_i \cap \mathcal{A}_j = \emptyset$, $\forall i, j$. In the initial state, the distribution of the image exhibits regional coherence. The local distribution over each coherent region

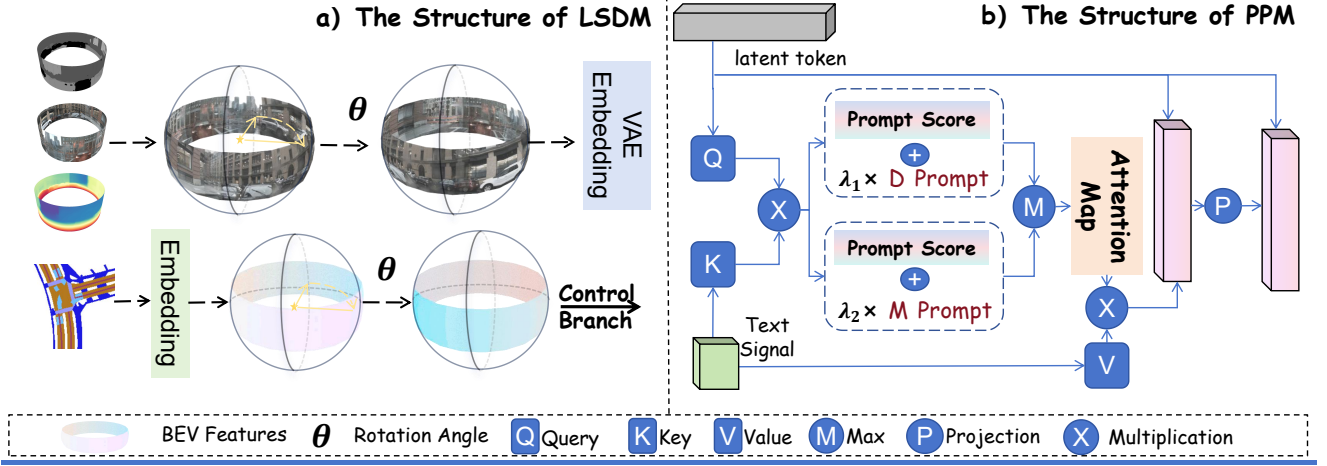


Fig. 4: The architecture of LSDM and PPM is illustrated in Figures (a) and (b), respectively. By applying circular rotations to diverse prompts and BEV features, LSDM reformulates panoramic image generation as a spatially continuous diffusion process. PPM further enhances controllability by dynamically selecting the most relevant control cues, leading to an accurate and consistent panoramic image.

\mathcal{R}_i , denoted $f_i(x)$, possesses strong spatial consistency. In contrast, the distribution over each aliased region \mathcal{A}_i , denoted $\tilde{f}_i(x)$, is complex and discontinuous due to stitching artifacts. Our objective is to enhance the model’s ability to capture the structure of the coherent regions by introducing an operation $R_\theta(\cdot)$. Meanwhile, by incorporating a control method, the network leverages prompt signals A and B to generate a controllable distribution: $f(x | A, B) \approx \sum_{i=1}^6 \mathbb{I}_{\mathcal{R}_i}(x) \cdot f_i(x | A, B)$, where $\mathbb{I}_{\mathcal{R}_i}(x)$ is the indicator function for region \mathcal{R}_i . Specifically, we seek to transform the initial discontinuous mixed distribution: $f(x) = \sum_{i=1}^6 \mathbb{I}_{\mathcal{R}_i}(x) \cdot f_i(x) + \sum_{i=1}^6 \mathbb{I}_{\mathcal{A}_i}(x) \cdot \tilde{f}_i(x)$, into a coherent, structure-dominant distribution: $f(x | A, B) \rightarrow \tilde{f}_{\text{coherent}}(x)$, $x \in \Omega$.

B. The Structure of Percep360

In this paper, we propose Percep360, the first paradigm for achieving high-quality controllable panoramic image generation. Specifically, Percep360 adopts a side-controlling architecture, following [11], [12]. The network introduces LSDM, enabling more coherent image generation under the panorama stitching setting. Furthermore, we introduce the PPM module, which injects prompt embeddings into the global features of panoramic images to enable controllable image generation.

Overall Architecture: As shown in Fig. 3, the control branch of U-Net processes the guidance signals, and the resulting features are injected into the main U-Net backbone to condition the generation. Specifically, the model takes as input three 2D data: the RGB image \mathbf{I} , the binary mask \mathbf{M} , and the depth map \mathbf{D} . In addition, the input includes a BEV map $\mathbf{B} \in \mathbb{R}^{B \times K \times 200 \times 200}$ and a text prompt \mathbf{T} . Each of the 2D inputs—RGB image, depth map, and mask—is individually encoded using a frozen variational autoencoder encoder \mathcal{E}_{VAE} with shared weights. This yields the following latent features:

$$\mathbf{F}_{\text{img,depth,mask}} = \mathcal{E}_{\text{VAE}}(\mathbf{I}, \mathbf{D}, \mathbf{M}). \quad (1)$$

Next, the BEV map is encoded using a dedicated BEV encoder \mathcal{E}_{BEV} , following BEVControl [14]. The resulting

feature map is resized to match the spatial resolution of the 2D features, as shown in the following equation:

$$\mathbf{F}_{\text{BEV}} = \text{Interp}(\mathcal{E}_{\text{BEV}}(\mathbf{B}), \text{size} = (h, w)). \quad (2)$$

Meanwhile, the text prompt \mathbf{T} is processed through a CLI text encoder $\mathcal{E}_{\text{CLIP}}$ [44], producing a semantic feature vector:

$$\mathbf{F}_{\text{TXT}} = \mathcal{E}_{\text{CLIP}}(\mathbf{T}) \in \mathbb{R}^{B \times d}. \quad (3)$$

Finally, the three encoded components (*i.e.*, \mathbf{F}_{2D} , \mathbf{F}_{BEV} , and \mathbf{F}_{TXT}) are fed into the control branch. This branch modulates the intermediate feature maps at multiple scales, enabling spatially and semantically controllable feature generation.

The design of LSDM: To achieve more coherent panoramic generation, we propose the Local Scenes Diffusion Method (**LSDM**), as shown in Fig. 4, which mitigates the discontinuities at image seams. Since a panoramic image spans a full 360° horizontal field of view, we treat its domain as a spherical manifold where the left and right boundaries are adjacent. For the 2D inputs, including the image \mathbf{I} , mask \mathbf{M} , and depth map \mathbf{D} , we perform a horizontal wrap-around rotation before feeding them into the network. A rotation angle $\theta \in [0, 2\pi)$ is randomly sampled and converted into a pixel displacement:

$$\Delta w = \left\lfloor \frac{\theta}{2\pi} \cdot W \right\rfloor, \quad (4)$$

where W is the width of the panoramic image. The rotated result is obtained via a circular shift:

$$\mathbf{X}_{\text{rot}} = \text{Roll}(\mathbf{X}, \Delta w), \quad \mathbf{X} \in \{\mathbf{I}, \mathbf{M}, \mathbf{D}\}, \quad (5)$$

where the Roll operator is defined as:

$$\text{Roll}(\mathbf{X}, \Delta w)[u, v] = \mathbf{X}[(u + \Delta w) \bmod W, v], \quad (6)$$

with \mathbf{X} denoting the input 2D tensor (image, mask, or depth map), W the horizontal width of \mathbf{X} , (u, v) the horizontal and vertical pixel coordinates, and mod the modulo operator ensuring wrap-around indexing.

The BEV input \mathbf{B} is encoded into a feature map:

$$\mathbf{F}_{\text{BEV}} = \mathcal{E}_{\text{BEV}}(\mathbf{B}). \quad (7)$$

Then we apply a similar circular shift in the feature space using the same angle θ :

$$\Delta w_{\text{BEV}} = \left\lfloor \frac{\theta}{2\pi} \cdot W_{\text{BEV}} \right\rfloor, \quad \mathbf{F}_{\text{BEV}}^{\text{rot}} = \text{Roll}(\mathbf{F}_{\text{BEV}}, \Delta w_{\text{BEV}}). \quad (8)$$

This feature-space rotation is consistent with the 2D image rotation, but applied directly to the encoded BEV representation. By synchronizing the angular shift across both 2D and BEV modalities, LSDM facilitates geometrically aligned and spatially consistent panoramic generation.

The design of PPM: To enhance scene-level controllability and geometric consistency in diffusion-based generation, we design the Probabilistic Prompting Method (PPM) module that incorporates structural features into image features. This module dynamically modulates attention based on two forms of masking: depth map and mask map for the entire image. Given the depth map $\mathbf{M}_d \in \mathbb{R}^{H_d \times W_d}$ and the mask map $\mathbf{M}_m \in \mathbb{R}^{H_m \times W_m}$, we first interpolate and reshape them to obtain spatially aligned maps $\tilde{\mathbf{M}}_d, \tilde{\mathbf{M}}_m \in \mathbb{R}^{H \times W}$, matching the spatial resolution of the visual token sequence $x \in \mathbb{R}^{N \times d}$, where $N = H \times W$:

$$\tilde{\mathbf{M}}_d = \text{Interpolate}(\mathbf{M}_d), \quad \tilde{\mathbf{M}}_m = \text{Interpolate}(\mathbf{M}_m). \quad (9)$$

The text signal $c \in \mathbb{R}^{L \times d}$ is linearly projected to the key and value spaces:

$$k = W_k c \in \mathbb{R}^{L \times d}, \quad v = W_v c \in \mathbb{R}^{L \times d}, \quad (10)$$

where $W_k, W_v \in \mathbb{R}^{d \times d}$ are learnable projection matrices. The latent token sequence x serves as the query after projection:

$$q = W_q x \in \mathbb{R}^{N \times d}. \quad (11)$$

The scaled dot-product attention logits $\mathbf{A} \in \mathbb{R}^{N \times L}$ are computed as:

$$\mathbf{A} = \frac{qk^\top}{\sqrt{d}}. \quad (12)$$

We then inject spatial priors from depth and mask prompts by reshaping $\tilde{\mathbf{M}}_d, \tilde{\mathbf{M}}_m$ into compatible matrices $\mathbf{M}'_d, \mathbf{M}'_m \in \mathbb{R}^{N \times L}$ and adding them with learnable scaling factors $\lambda_1, \lambda_2 \in \mathbb{R}$:

$$\mathbf{A}_d = \mathbf{A} + \lambda_1 \mathbf{M}'_d, \quad \mathbf{A}_m = \mathbf{A} + \lambda_2 \mathbf{M}'_m. \quad (13)$$

The final attention map is obtained by element-wise maximum:

$$\mathbf{A}_{\text{final}} = \max(\mathbf{A}_d, \mathbf{A}_m). \quad (14)$$

After normalization with softmax along the key dimension:

$$\text{Attn} = \text{softmax}(\mathbf{A}_{\text{final}}), \quad (15)$$

the output representation is computed via a weighted sum over the values, followed by a linear projection and residual connection, following [43]:

$$x_{\text{out}} = \text{Proj}(\text{Attn} \cdot v) + x. \quad (16)$$

This attention mechanism enhances the controllability of the network and mitigates the controllability degradation caused by the continuous diffusion process.

Benefiting from the proposed design, our method enhances the quality of generated images and has the potential to boost the performance of downstream panoramic BEV perception.

IV. EXPERIMENTS

In this section, we present the experimental configuration in Sec. IV-A, and then detail the quantitative and visualization results in Sec. IV-B.

A. Experiments Details

Dataset and Baselines. The experiments are conducted on the nuScenes-360 dataset [10], which consists of panoramic driving scenes stitched from the nuScenes dataset [48]. Following the official configuration, the dataset comprises 28,130 images for training and 6019 images for validation. We adopt BEVControl [14] as the baseline method. In our experiments, we replace the center view sample with a panoramic image.

Evaluation Metrics and Strategies. We employ multiple evaluation strategies to assess the effectiveness of our panorama generation method comprehensively. Specifically, we utilize Fréchet Inception Distance (FID) [47], Structural Similarity Index (SSIM) [46] for image quality assessment. Meanwhile, we also employ no-reference image quality assessment (*i.e.*, BRISQUE [45], PIQE [49]) to evaluate the quality by comparing the generated images with ground-truth images and those from other methods. Additionally, we evaluated the images, assessing their IoU score for road segmentation and mean intersection over union on the pre-trained BEV semantic segmentation network. Furthermore, to more intuitively compare the performance of different methods, we ranked the results for each score and obtained a ranking score.

Experiment Setting. As we are the first to propose the generation of panoramic images using stitched pinhole images, addressing the current challenges in image collection and annotation within the panoramic generation community, we modify the BEVControl [14] to serve as our baseline. Both the image and control signal are represented in 2D format (*i.e.*, BEV Map, Mask Map, Depth Map). Meanwhile, we introduce more comparisons with the 3D generation model, Magicdrive [18]. Our experiments were conducted on two A6000 GPUs over a training period of ten days. We set the diffusion step to 20 without using data augmentation strategies. To accelerate the diffusion process, the image is processed at a resolution of (60, 1200) following [50].

B. Main Results

As shown in Tab. I, we extend six-view image generation methods [14], [18] to the panoramic domain and establish a new benchmark. We conduct a comprehensive comparison between Percep360 and other state-of-the-art approaches. In the NRIQA evaluation, Percep360 achieves the best performance, particularly in the PIQE metric, where our method surpasses the original stitched images. Visualization results,

Method	NRIQA		IQA		Control		Rank.
	BRISQUE [45] (↓)	PIQE (↓)	SSIM [46] (↑)	FID [47] (↓)	Driv. (↑)	Mean. (↑)	(↓)
Oracle	19.39	12.06			0.59	0.47	N/A
MagicDrive* [18]	20.65 (2)	12.13 (2)	0.16 (1)	14.16 (1)	0.22 (2)	0.12 (2)	10 (2)
Baseline [14]	22.40 (3)	14.32 (3)	0.14 (3)	22.07 (3)	0.21 (3)	0.12 (2)	17 (3)
Percep360	20.24 (1)	11.44 (1)	0.16 (1)	14.43 (2)	0.25 (1)	0.13 (1)	7 (1)

TABLE I: Comparison with state-of-the-art methods on the nuScenes validation set in *Generation Quality*. ↑ / ↓ indicates that a higher/lower value is better. ‘Oracle’ presents the row results of corresponding models trained on the nuScenes training set. * indicates that the model has undergone structural modifications to accommodate the panoramic image generation task.

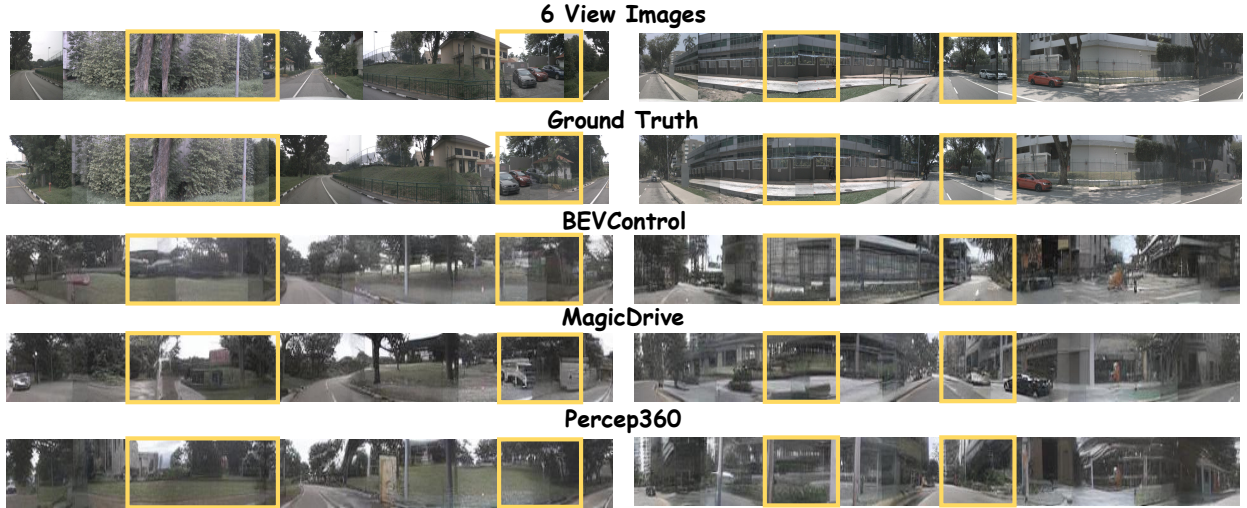


Fig. 5: Visualizations of the generation results for the ground truth, BEVControl [14], and MagicDrive [18] are provided. Regions of stitching misalignment are highlighted with yellow boxes. It can be observed that, due to the lack of coherent generation capability, existing methods produce discontinuous results. In contrast, our method achieves improved coherence without compromising controllability.

as shown in Fig. 5, further demonstrate that the stitched images suffer from obvious pixel inconsistencies, whereas our generated images effectively resolve this issue and produce more coherent results. Other methods also exhibit a tendency to learn the misalignment present in the ground truth during training, which leads to incoherent image generation in their visualization outputs. Regarding IQA and controllability, our network also attains strong performance. For MagicDrive, we modify its cross-view attention to self-attention, better suited for panoramic images (an advantage of panoramic imaging that eliminates the need for explicit viewpoint alignment). Although Percep360’s FID is slightly higher than that of MagicDrive, it achieves the best NRIQA, controllability, and overall score in the comprehensive evaluation.

V. ABLATION STUDIES

The effectiveness of our method is verified through ablation studies on its two main components. Coherent panorama generation strategies (LSDM) are compared in Sect. V-A, followed by an evaluation of the performance of prompt utilization in Sec V-B. We further validate that the generated data contributes to the improvement of downstream BEV segmentation performance in Sec. V-C.

A. Analysis of LSDM

To address the inherent incoherence in scene capture arising from multi-view camera systems during data collection,

we compared LSDM against three conventional approaches, as shown in Tab. II. It is worth noting that, to the best of our knowledge, there has been no prior discussion in the fields of panoramic and street-view image generation regarding the relationship between 3D projection into the BEV space and image coherence within the BEV view. First, the masking strategy (“Mask”) [51], [52], [38] serves as a training guidance mechanism within the diffusion paradigm. By intentionally omitting supervision in suboptimal or uncertain regions, the model is encouraged to leverage contextual information from surrounding areas to reconstruct or complete the masked content, thereby enhancing its reconstruction capabilities. Furthermore, we incorporate self-attention (“M.Crs”) over latent features within the U-Net architecture to enhance the diffusion of supervision signals from observed to unobserved regions, resulting in noticeable improvements in image quality at the feature level. However, those methods are not suitable for the present task. As street view generation requires road and geometric details, methods that rely solely on neighboring features, although effective in preserving semantic consistency, lead to a significant deterioration in the structural integrity of the masked regions. Furthermore, following the training paradigm of state-of-the-art panoramic generation methods [41], [35], we rotate the image pixels and apply the corresponding top-down rotation to the BEV image. However, this approach results in a significant degrada-

Settings	BRISQUE [45]↓	PIQE [49]↓	FID [47]↓	Driv.↑	Mean.↑	Rank.↓
Baseline	22.41 (5)	14.33 (5)	22.07 (5)	0.21 (1)	0.12 (1)	17 (3)
+ Mask	20.87 (4)	11.65 (4)	16.35 (3)	0.18 (3)	0.09 (3)	17 (4)
+ M.Crs	20.03 (1)	10.91 (2)	14.96 (1)	0.17 (4)	0.08 (4)	12 (1)
+ Rota	20.25 (2)	10.85 (1)	21.42 (4)	0.14 (5)	0.06 (5)	17 (5)
+ LSDM F	20.84 (3)	11.44 (3)	16.28 (2)	0.20 (2)	0.10 (2)	12 (1)

TABLE II: An in-depth analysis of LSDM. ↑/↓ indicates whether a higher or lower value is preferred for each metric. “Mask” refers to the masking strategy, while “M.Crs.” denotes the cross-attention method built upon the masking strategy. Additionally, “Rota” represents the rotation method. “Driv.” refers to the segmentation results of drivable area, and “Mean” denotes the average performance across all BEV segmentation classes.

Settings	BRISQUE [45]↓	PIQE [49]↓	FID [47]↓	Driv.↑	Mean.↑	Rank.↓
Baseline	22.41 (5)	14.33 (5)	22.07 (5)	0.21 (4)	0.12 (2)	21 (5)
LSDM	20.87 (3)	12.25 (4)	16.28 (4)	0.20 (5)	0.10 (5)	21 (5)
LSMD + D	20.41 (2)	11.54 (2)	12.76 (1)	0.23 (2)	0.12 (2)	9 (2)
Percep360	20.24 (1)	11.44 (1)	14.43 (3)	0.25 (1)	0.13 (1)	7 (1)

TABLE III: Ablation results under different prompt configurations. “D” refers to the Depth map. “Driv.” refers to the segmentation results of drivable area, and “Mean” denotes the average performance across all BEV segmentation classes.

tion in both image quality and layout controllability. LSDM decouples the correspondence between the image-BEV pair and dynamically reconfigures the relationship between the network and these pairs. This approach effectively mitigates the erroneous constraints introduced by sample incoherence from pinhole cameras, while leveraging the model’s learning capacity to alleviate the controllability degradation caused by the other methods. We further conduct a ranking of various image evaluation metrics, demonstrating that LSDM effectively achieves a balance between image quality and layout controllability.

Data	BEV Segmentation	
	Driv. (↑)	Mean. (↑)
w/o synthetic data	0.593	0.470
w/ Baseline	0.589	0.469
w/ Percep360	0.610	0.495

TABLE IV: By leveraging BEV features along with additional conditioning signals as inputs, new panoramic street-view images are synthesized through different methods. The generated data is further used to boost the panoramic BEV segmentation task.

B. Analysis of Different Prompts

To achieve more controllable panoramic image generation and avoid the labor-intensive task of manual panoramic image annotation, we leverage large language models to pre-annotate the stitched panoramic images, thereby generating imprecise yet informative control prompts. Furthermore, we propose the PPM, which enables robust panoramic generation. By incorporating depth information, PPM dynamically selects relevant image tokens and depth cues, surpassing the baseline in terms of image quality while significantly improving controllability. With the joint contributions of PPM and

LSMD, Percep360 achieves comprehensive improvements in both image quality and layout controllability.

C. Training Support for BEV Segmentation.

As an important validation of the capabilities of the world model, we evaluate the effectiveness of newly generated data as augmentation for improving downstream network performance. Specifically, we adopt OneBEV [10] as the baseline for BEV segmentation. As shown in Table IV, incorporating synthetic data into the baseline model did not lead to performance improvement (mIoU dropped slightly from 0.470 to 0.469). In contrast, the synthetic images generated by Percep360 exhibit higher realism and better alignment with BEV representations, resulting in performance gains: the drivable region improved from 0.593 to 0.610, and the overall mIoU increased from 0.470 to 0.495. Notably, these benefits were achieved without the need for additional data cleaning or filtering, making Percep360 a cost-effective paradigm for enhancing agent perception.

VI. CONCLUSION AND FUTURE WORK

Contribution: In this paper, we propose the first panoramic generation method Percep360. By introducing the Local Scenes Diffusion Method, the panorama generation task is reformulated as a spatially continuous diffusion process, ensuring the coherence of panoramic images. Additionally, the Probabilistic Prompting Method dynamically selects the most relevant control cues, thereby improving the controllability of image generation. The generated data proves to be an effective form of data augmentation, leading to improved performance in downstream tasks.

Limitation and Future Work: In this paper, we focus on exploiting image coherence and the controllability of 2D features. The label mapping from six-view images to panoramic representations constitutes a critical component for efficient panoramic data generation. Furthermore, existing street-view

generation approaches have largely overlooked the effective utilization of prompt features, which we identify as another key direction for future investigation.

REFERENCES

- [1] S. Gao, K. Yang, H. Shi, K. Wang, and J. Bai, "Review on panoramic imaging and its applications in scene understanding," *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–34, 2022.
- [2] Y. Xu, Z. Zhang, and S. Gao, "Spherical DNNs and their applications in 360° images and videos," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 10, pp. 7235–7252, 2022.
- [3] H. Ai, Z. Cao, and L. Wang, "A survey of representation learning, optimization strategies, and applications for omnidirectional vision," *International Journal of Computer Vision*, vol. 133, no. 8, pp. 4973–5012, 2025.
- [4] H. Shi *et al.*, "PanoFlow: Learning 360° optical flow for surrounding temporal understanding," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 5, pp. 5570–5585, 2023.
- [5] Y. Li, C. Barnes, K. Huang, and F.-L. Zhang, "Deep 360° optical flow estimation based on multi-projection fusion," in *ECCV*, 2022, pp. 336–352.
- [6] H. Xu *et al.*, "PANDORA: A panoramic detection dataset for object with orientation," in *ECCV*, 2022, pp. 237–252.
- [7] K. Luo *et al.*, "Omnidirectional multi-object tracking," in *CVPR*, 2025, pp. 21 959–21 969.
- [8] J. Zheng *et al.*, "Open panoramic segmentation," in *ECCV*, 2024, pp. 164–182.
- [9] X. Zheng, P. Zhou, A. V. Vasilakos, and L. Wang, "Semantics distortion and style matter: Towards source-free UDA for panoramic segmentation," in *CVPR*, 2024, pp. 27 885–27 895.
- [10] J. Wei, J. Zheng, R. Liu, J. Hu, J. Zhang, and R. Stiefelhagen, "OneBEV: Using one panoramic image for bird's-eye-view semantic mapping," in *ACCV*, 2024, pp. 583–596.
- [11] L. Zhang, A. Rao, and M. Agrawala, "Adding conditional control to text-to-image diffusion models," in *ICCV*, 2023, pp. 3813–3824.
- [12] B. Peng, J. Wang, Y. Zhang, W. Li, M.-C. Yang, and J. Jia, "Control-NeXt: Powerful and efficient control for image and video generation," *arXiv preprint arXiv:2408.06070*, 2024.
- [13] A. Swerdlow, R. Xu, and B. Zhou, "Street-view image generation from a bird's-eye view layout," *IEEE Robotics and Automation Letters*, vol. 9, no. 4, pp. 3578–3585, 2024.
- [14] K. Yang, E. Ma, J. Peng, Q. Guo, D. Lin, and K. Yu, "BEVControl: Accurately controlling street-view elements with multi-perspective consistency via BEV sketch layout," *arXiv preprint arXiv:2308.01661*, 2023.
- [15] Y. Wang, J. He, L. Fan, H. Li, Y. Chen, and Z. Zhang, "Driving into the future: Multiview visual forecasting and planning with world model for autonomous driving," in *CVPR*, 2024, pp. 14 749–14 759.
- [16] W. Zheng, W. Chen, Y. Huang, B. Zhang, Y. Duan, and J. Lu, "OccWorld: Learning a 3D occupancy world model for autonomous driving," in *ECCV*, 2025, pp. 55–72.
- [17] Y. Wei *et al.*, "Editable scene simulation for autonomous driving via collaborative LLM-agents," in *CVPR*, 2024, pp. 15 077–15 087.
- [18] R. Gao *et al.*, "MagicDrive: Street view generation with diverse 3D geometry control," in *ICLR*, 2024.
- [19] W. Wu *et al.*, "DriveScape: Towards high-resolution controllable multi-view driving video generation," *arXiv preprint arXiv:2409.05463*, 2024.
- [20] X. Li, Y. Zhang, and X. Ye, "DrivingDiffusion: Layout-guided multi-view driving scenarios video generation with latent diffusion model," in *ECCV*, 2024, pp. 469–485.
- [21] Z. Wu *et al.*, "HoloDrive: Holistic 2D-3D Multi-modal street scene generation for autonomous driving," *arXiv preprint arXiv:2412.01407*, 2024.
- [22] S. Gu, J. Su, Y. Duan, X. Chen, J. Luo, and H. Zhao, "Text2Street: Controllable text-to-image generation for street views," in *ICPR*, 2025, pp. 130–145.
- [23] Y. Wen *et al.*, "Panacea: Panoramic and controllable video generation for autonomous driving," in *CVPR*, 2024, pp. 6902–6912.
- [24] R. Gao, K. Chen, B. Xiao, L. Hong, Z. Li, and Q. Xu, "MagicDriveDiT: High-resolution long video generation for autonomous driving with adaptive control," *arXiv preprint arXiv:2411.13807*, 2024.
- [25] J. Zhang *et al.*, "PerDiff: Controllable street view synthesis using perspective-layout diffusion models," *arXiv preprint arXiv:2407.06109*, 2024.
- [26] R. Chen *et al.*, "UniMLVG: Unified framework for multi-view long video generation with comprehensive control capabilities for autonomous driving," *arXiv preprint arXiv:2412.04842*, 2024.
- [27] L. Wang *et al.*, "Stag-1: Towards realistic 4D driving simulation with video generation model," *arXiv preprint arXiv:2412.05280*, 2024.
- [28] J. Nam, H. Kim, D. Lee, S. Jin, S. Kim, and S. Chang, "DreamMatcher: Appearance matching self-attention for semantically-consistent text-to-image personalization," in *CVPR*, 2024, pp. 8100–8110.
- [29] M. Arar *et al.*, "PALP: Prompt aligned personalization of text-to-image models," in *SIGGRAPH Asia*, 2024, pp. 1–11.
- [30] Y. Hao, Z. Chi, L. Dong, and F. Wei, "Optimizing prompts for text-to-image generation," in *NeurIPS*, 2023, pp. 66 923–66 939.
- [31] Z. Xue *et al.*, "Raphael: Text-to-image generation via large mixture of diffusion paths," in *NeurIPS*, 2024, pp. 41 693–41 706.
- [32] C. Zhang *et al.*, "Taming stable diffusion for text to 360° panorama image generation," in *CVPR*, 2024, pp. 6347–6357.
- [33] A. Liu, Z. Li, Z. Chen, N. Li, Y. Xu, and B. A. Plummer, "PanoFree: Tuning-free holistic multi-view image generation with cross-view self-guidance," in *ECCV*, 2024, pp. 146–164.
- [34] J. Li and M. Bansal, "PanoGen: Text-conditioned panoramic environment generation for vision-and-language navigation," in *NeurIPS*, 2023, pp. 21 878–21 894.
- [35] Q. Wang, W. Li, C. Mou, X. Cheng, and J. Zhang, "360DVD: Controllable panorama video generation with 360-degree video diffusion model," in *CVPR*, 2024, pp. 6913–6923.
- [36] C. Zheng, G. Song, T.-J. Cham, J. Cai, L. Luo, and D. Phung, "Bridging global context interactions for high-fidelity pluralistic image completion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 12, pp. 8320–8333, 2024.
- [37] C. Zheng, T.-J. Cham, and J. Cai, "Pluralistic image completion," in *CVPR*, 2019, pp. 1438–1447.
- [38] A. Lugmayr, M. Danelljan, A. Romero, F. Yu, R. Timofte, and L. R. Van Gool, "RePaint: Inpainting using denoising diffusion probabilistic models," in *CVPR*, 2022, pp. 11 451–11 461.
- [39] X. Ju, X. Liu, X. Wang, Y. Bian, Y. Shan, and Q. Xu, "BrushNet: A plug-and-play image inpainting model with decomposed dual-branch diffusion," in *ECCV*, 2024, pp. 150–168.
- [40] Y. Li, Y. Bian, X. Ju, Z. Zhang, Y. Shan, and Q. Xu, "BrushEdit: All-in-one image inpainting and editing," *arXiv preprint arXiv:2412.10316*, 2024.
- [41] J. Wang, Z. Chen, J. Ling, R. Xie, and L. Song, "360-degree panorama generation from few unregistered NFOV images," in *MM*, 2023, pp. 6811–6821.
- [42] T. Wu, C. Zheng, and T.-J. Cham, "PanoDiffusion: 360-degree panorama outpainting via diffusion," in *ICLR*, 2023.
- [43] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *CVPR*, 2022, pp. 10 674–10 685.
- [44] A. Radford *et al.*, "Learning transferable visual models from natural language supervision," in *ICML*, 2021, pp. 8748–8763.
- [45] A. Mittal, A. K. Moorthy, and A. C. Bovik, "No-reference image quality assessment in the spatial domain," *IEEE Transactions on Image Processing*, vol. 21, no. 12, pp. 4695–4708, 2012.
- [46] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [47] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "GANs trained by a two time-scale update rule converge to a local nash equilibrium," in *NeurIPS*, 2017, pp. 6626–6637.
- [48] H. Caesar *et al.*, "nuScenes: A multimodal dataset for autonomous driving," in *CVPR*, 2020, pp. 11 618–11 628.
- [49] V. N., P. D., M. C. Bh., S. S. Channappayya, and S. S. Medasani, "Blind image quality evaluation using perception based features," in *NCC*, 2015, pp. 1–6.
- [50] E. Santana and G. Hotz, "Learning a driving simulator," *arXiv preprint arXiv:1608.01230*, 2016.
- [51] C. Liu, H. Wu, Y. Zhong, X. Zhang, Y. Wang, and W. Xie, "Intelligent grimm - Open-ended visual storytelling via latent diffusion models," in *CVPR*, 2024, pp. 6190–6200.
- [52] R. Suvorov *et al.*, "Resolution-robust large mask inpainting with fourier convolutions," in *WACV*, 2022, pp. 3172–3182.