

# CamFreeDiff: Camera-free Image to Panorama Generation with Diffusion Model

Xiaoding Yuan<sup>1,2</sup>, Shitao Tang<sup>3</sup>, Kejie Li<sup>2</sup>, Alan Yuille<sup>1</sup>, and Peng Wang<sup>2</sup>

<sup>1</sup> The Johns Hopkins University

<sup>2</sup> ByteDance

<sup>3</sup> Simon Fraser University



**Fig. 1:** Camera-free image-to-panorama generation: our method generates panoramas with reasonable layout and canonical viewpoint with unknown input image camera parameters.

**Abstract.** This paper introduces Camera-free Diffusion (CamFreeDiff) model for 360-degree image outpainting from a single camera-free image and text description. This method distinguishes itself from existing strategies, such as MVDiffusion, by eliminating the requirement for pre-defined camera poses. Instead, our model incorporates a mechanism for predicting homography directly within the multi-view diffusion framework. The core of our approach is to formulate camera estimation by predicting the homography transformation from the input view to a pre-defined canonical view. The homography provides point-level correspondences between the input image and targeting panoramic images, allowing connections enforced by correspondence-aware attention in a fully differentiable manner. Qualitative and quantitative experimental results demonstrate our model’s strong robustness and generalization ability for

360-degree image outpainting in the challenging context of camera-free inputs.

## 1 Introduction

Unlike traditional movies, 360-degree movies create an immersive experience that allows viewers to feel as if they are part of the movies’s environment rather than merely observing it from a fixed perspective. This immersive aspect has been significantly enhanced with the advent and proliferation of AR/VR devices like VisionPro. Creating 360-degree movies typically requires specialized equipment, like 360-degree cameras, making it a highly professional endeavour. An alternative, yet underexplored, approach is outpainting existing movies into 360-degree formats.

Image-based panorama outpainting [2, 2, 4, 14, 27, 29, 32, 33, 35] is a necessary step towards video-based 360 movie creation. The advance of text-to-image diffusion models [9, 17–21] makes it possible to extrapolate an image into 360-degree view. PanoDiffusion [32] proposes a panorama outpainting methods by fine-tuning a pretrained latent diffusion model [20]. However, with a limited amount of training data available, this method disrupts the pre-trained model’s prior knowledge and diminishing its generalization capabilities. MVDiffusion [29], in contrast, maintains generalization by generating multi-view consistent panoramic images using a frozen pre-trained latent diffusion model. This method ensures geometric consistency through correspondence-aware attention, but it requires the input image to have known intrinsic and rotation matrices, limiting its application to panoramas from arbitrary images.

Extending panorama generation to camera-free inputs poses significant challenges. Our method estimates the input camera parameters by estimating the homography transformation from the input image to a predefined canonical view. The homography establishes a correspondence between input views and each panoramic views, allowing for the enforcement of multi-view consistency via correspondence-aware attention [29]. Furthermore, the homographic matrix estimation network is seamlessly integrated with the multi-view diffusion model in a fully differentiable manner. By doing so, the mechanism can effectively mitigate the errors introduced by the homography estimation process.

We fine-tuned our model on high-quality Matterport3D from the pre-trained stable diffusion inpainting model. We demonstrated the strong robustness of our model to random camera parameters. Our model also shows excellent generalization ability when tested on out-of-domain data from Structured3D without fine-tuning.

To summarize, we make the following contributions:

- We propose CamFreeDiff, a panorama generation model capable of handling camera-free input images. To the best of our knowledge, our model is the first of all designed to handle unknown input view and camera parameters.
- We formulate the camera parameter estimation by predicting the homography transformation from the input image to a predefined canonical view. We


- propose a three degrees of freedom(3-DoF) parameterization of homography instead of a standard 8-DoF way in the context of panorama outpainting.
- We propose a novel strategy for incorporating camera prediction into the panorama generation model to enforce multi-view consistency and high visual quality.

## 2 Related Work

**Multi-view image generation.** The development of large-scale image generation models [17, 20, 21] has paved the way for multi-view image generation [11–13, 23, 24, 28, 29, 31]. A notable advancement in this area is the Zero-1-to-3 model [11], which can deduce novel views of an object given a single image and specified camera pose. Building on this, SyncDreamer [12] introduces a method for generating images at fixed viewpoints through a combination of volumetric representation and epipolar attention, though it occasionally produces artifacts. In an effort to enhance consistency and image fidelity, models such as MVDream [24], ImageDream [31], and Wonder3D [13] have employed comprehensive self-attention mechanisms across all views. Despite their success, these models necessitate extensive fine-tuning of a latent diffusion model’s parameters, which can compromise pre-trained priors. While our method also adopts multi-view generation paradigms, our method is able to outpaint panorama for arbitrary images by freezing the pretrained stable diffusion model.

**Panorama generation.** Previous methods [1, 4, 25, 26, 29, 32, 34] for generating panoramas typically rely on generative adversarial networks [7], autoregressive [15] or diffusion models. Text2Light [4] emerged as a pioneering approach, utilizing a cascade autoregressive model for panorama creation, whereas LDM3D [26] and PanoDiffusion [32] have adopted latent diffusion techniques for generating or outpainting panoramas. Despite their innovations, these models often encounter a reduction in their ability to generalize when they are fine-tuned on limited datasets. On the other hand, the process of generating panoramic images inherently benefits from the direct correspondence between different views. Utilizing this, MVDiffusion [29] presents a novel approach by incorporating correspondence-aware attention to maintain multi-view consistency, without modifying the underlying pretrained model.

However, challenges persist with current outpainting techniques, notably their struggle to generate coherent panoramas due to their inability to infer camera parameters. PanoDiff [30] takes a step forward to estimate the longitude and latitude angles of a view, however still assumes camera roll angle and intrinsic. To address these shortcomings, we introduce a novel, fully differentiable pipeline designed to extrapolate from a single camera-free image to a complete 360-degree panorama. This method aims to bridge the gap between existing panoramic generation techniques and the need for accurate, consistent, and high-quality panoramic imagery, leveraging the strengths of direct view correspondence while overcoming the limitations of camera parameter estimation.



**Fig. 2:** In the context of panorama generation, we define the camera parameter for a canonical perspective view as  $fov = 90^\circ$ ,  $x$ -axis rotation  $\phi = 0$  and  $z$ -axis rotation  $\psi = 0$ . The  $y$ -axis rotation can be any  $\theta$  from  $0^\circ$  to  $360^\circ$ .

### 3 Preliminary

MVDiffusion is a multi-view diffusion model for coherent panoramic image generation. It splits the  $360^\circ$  scene into eight perspective views  $\mathbf{I} = \{I_0, I_1, \dots, I_7\}$  at fixed viewpoint, each with  $90^\circ$  horizontal and vertical field of view and a  $45^\circ$  horizontal overlapping. For image-based panorama outpainting, MVDiffusion employs a latent diffusion inpainting model and makes the first view (canonical view), conditional on the input images. Correspondence-aware attention (CAA) blocks are designed to enforce geometry consistency. Consider a pair of corresponding points located at  $p_s$  and  $p_t$  in the source view  $I_s$  and target view  $I_t$  respectively, the information is aggregated from a  $K \times K$  neighborhood  $\mathcal{N}(p_s)$  centered at  $p_s$  in the source view through a cross-attention mechanism:

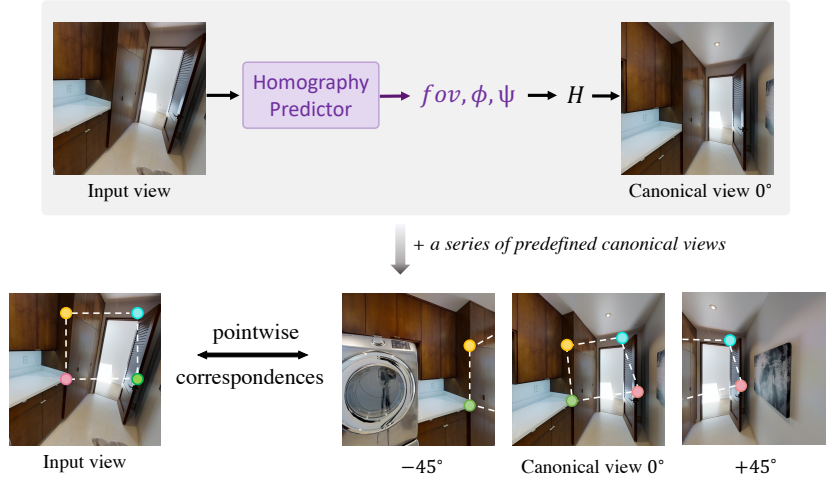
$$M = \sum_{p_s^* \in \mathcal{N}(p_s)} \text{Softmax}([W_Q f(p_t)] \cdot [W_K f(p_s^*)]) W_V f(p_s^*), \quad (1)$$

Where  $f(p_t)$  denotes image feature at location  $p_t$  from target view and  $f(p_s^*)$  represents image feature at location  $p_s^*$  in the neighborhood  $\mathcal{N}(p_s)$  from source view.  $W_Q, W_K, W_V$  are the query, key and value projections for cross-attention in CAA blocks. MVDiffusion assumes the camera parameters of the input are the same as the canonical view, and thus cannot handle camera-free image.

### 4 Method

We aim to generate a text-guided panoramic image from a camera-free input, where we assume the camera parameters are unknown. To address this problem, we propose our Camera-free Diffusion(CamFreeDiff) model to predict the camera parameters of the input image by estimating the homography transformation from the input to a predefined canonical perspective view. The predicted homography provides correspondences between multiple target perspective views and the input view, where correspondence-aware attention could be used to enforce geometry consistency for the final panorama generation.





**Fig. 3:** We formulate the camera parameter estimation as estimating the homography matrix from the input view to a predefined canonical view of the scene. We define the canonical view as the perspective view with an absolute rotation angle of  $0^\circ$ . We use a 3-DoF parameterization of the homography matrix instead of standard 8-DoF (details in Sec. 4.1).

#### 4.1 Homography Matrix Estimation

To generate a panorama based on the input image with an unknown camera, we first estimate the homography  $H^{3 \times 3}$  matrix from the unknown input view to a predefined canonical view. Fig. 2 Shows how we define a canonical view, and we illustrated our method in Fig. 3.

**Parameterization of Homography Matrix.** Compared to the most straightforward way of parameterizing a homography into elements in the  $3 \times 3$  matrix, DeTone *et al.* [5] demonstrated a 4-point parameterization of homography matrix can well reduce the optimization difficulty for a deep neural network, since the 4-points are independent of each other while the  $3 \times 3$  matrix mixes both the rotational and the translational terms.

In our case, parameters can be further simplified. When expressing homography as  $H = K_2 R K_1^{-1}$ , where  $K_2$  and  $K_1$  are the camera intrinsics for the canonical view and the input image accordingly, some variables are constant by default for common cameras. Specifically, the intrinsic matrix of a canonical view  $K_2$  is

$$\begin{bmatrix} f_x & \gamma & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 256 & 0 & 255.5 \\ 0 & 256 & 255.5 \\ 0 & 0 & 1 \end{bmatrix} \quad (2)$$

We also assume the intrinsic of input image satisfies the pinhole camera and thus for  $K_1$ , in accord with our canonical view intrinsic, the axis skew coefficient

$\gamma$  for input view also defaults to 0 and principal point offsets  $(c_x, c_y)$  default to the centre of the image  $(w/2, h/2)$ .

Accordingly, for our task to generate 360-degree panorama, the homography matrix from input view to the canonical view has three degrees of freedom (3-DoF): 1) camera field of view ( $f$ ), 2) camera rotation around the  $x$ -axis ( $\phi$ ), 3) camera rotation around the  $z$ -axis ( $\psi$ ). Particularly, under the condition of a single input image, predicting the absolute rotation around the  $y$ -axis ( $\theta$ ) is considered meaningless since the input view can be mapped to any standard canonical view of a 360-degree panorama with  $0 \leq \theta \leq 360^\circ$ .

Therefore, we formulate the model that predicts homography matrix from an input image  $I \in \mathbb{R}^{H \times W \times 3}$  to a predefined canonical view as:

$$\mathcal{M}(I) \rightarrow (f, \phi, \psi) \quad (3)$$

The input image camera intrinsic  $K_1$  can be determined by predicted  $f$ . Along with known target perspective camera intrinsic and predicted rotation  $R$  from  $(\phi, \psi, \theta = 0)$ , the homography transformation  $H$  can be recovered from the predictions  $(f, \phi, \psi)$ .

**Homography Estimator** We build an MLP classifier with three hidden layers upon a general image encoder. The U-Net encoder pre-trained by stable diffusion model for image generation also serves as our homography estimation image encoder, but with weights frozen for efficiency. Only the MLP classifier is optimized to learn a homography estimator. Feature dimensions for each hidden layer in MLP are set to 5120, 2560 and 1280. *SiLU* [6] is used as activation functions in MLP block. Cross-entropy loss is applied as learning objectives to  $f, \phi$  and  $\psi$ , respectively.

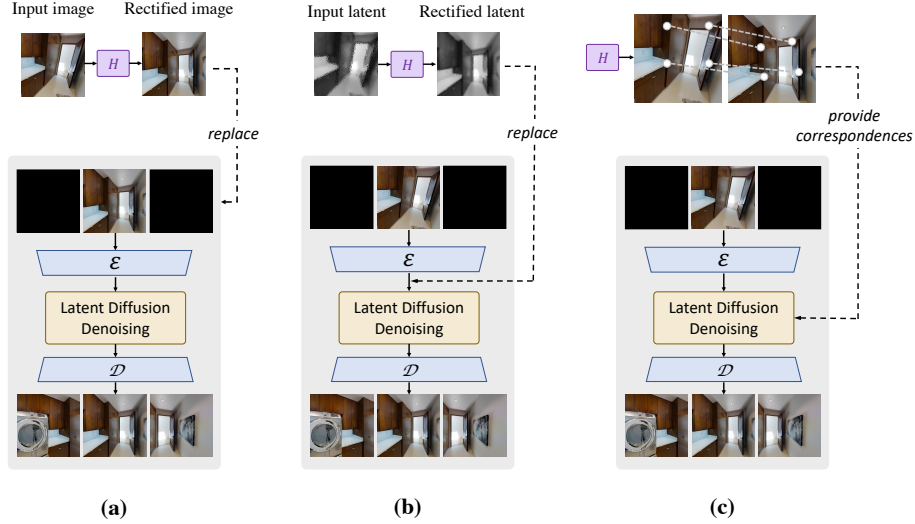
## 4.2 Camera-free 360-Degree Image Outpainting

Homography matrix from the camera-free input image to a predefined canonical image provides pixel-level correspondences between them. Consider the homography transformation from view  $I_a$  to view  $I_b$  as  $H_{a,b}$ . The projection from a point at location  $p_a = (u_a, v_a)$  in view  $I_a$  and its corresponding point at location  $p_b = (u_b, v_b)$  in view  $I_b$  is formulated as:

$$\begin{bmatrix} u_b \\ v_b \\ 1 \end{bmatrix} = H_{a,b} \begin{bmatrix} u_a \\ v_a \\ 1 \end{bmatrix} \quad (4)$$

Based on the estimated correspondences, we design three variants to generate 360-degree panoramic images, as shown in Fig. 4.

**Variation 1(unwarp image):** Initially, the input image is transformed to a canonical view utilizing estimated correspondences, as depicted in Fig. 4(a). Subsequently, we incorporate the latent diffusion inpainting model to generate eight panoramic views (eight diffusion branches with the same weight copy), following the design of MVDiffusion [29] image-based panorama outpainting model.

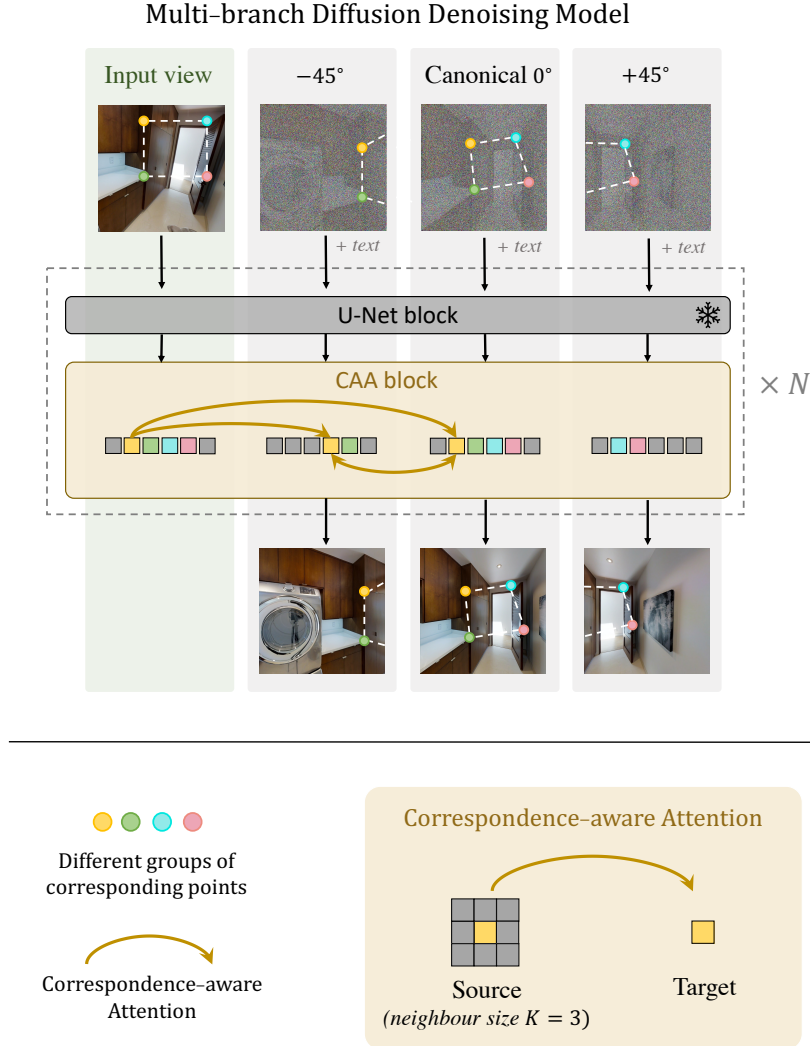


**Fig. 4:** Different strategies to generate panorama from a camera-free input. After estimating the homography matrix  $H$  from the input view to a predefined canonical view, alternatives are: (a) Rectify the input view by unwarping and replacing the original input image. (b) Rectify input view by unwarping and replacing the input latent after image encoding. (c) Provide point-level correspondences to the multi-view generation model to enforce consistency between corresponding points.

For the branch of canonical view  $0^\circ$ , input comprises a concatenation of noisy latent, the latent of the unwarped image, and a binary mask that identifies the areas requiring inpainting (*zero* for the visible region and *one* for the region to inpaint). The inputs for the remaining seven branches consist of the noisy latent, the latent of a purely white image, and a uniformly one-valued mask. It is important to note that the latent diffusion inpainting model is designed to preserve existing image content where the mask value is set to *zero*, and to generate new content in areas where the mask value is *one*.

**Variant 2(unwarp latent):** Unlike Variant 1, which initiates the process by unwarping the input image and subsequently encoding the unwarped image into a latent, in Variant 2, the image is first encoded into a latent space, followed by unwarping this latent into the canonical view  $0^\circ$ . Aside from this distinction, all other aspects of the design remain consistent with those outlined for Variant 1 as shown in Fig. 4(b).

**Variant 3(new view):** Variant 3 introduces a novel approach that diverges from the strategies that involve unwarping images or latent employed in variants 1 and 2. This design is structured into one conditional branch and eight generation branches. Unlike Variants 1 and 2, where the canonical view  $0^\circ$  branch depends on unwarped images or latent, the branch of canonical view is a generation branch in Variant 3, while the conditional branch depends on the input



**Fig. 5:** Our panorama generation pipeline based on multi-view diffusion denoising model. With the predicted homography matrix from the input view to a predefined canonical view, point-wise information can be aggregated from the input view to all target canonical views through correspondence-aware attention(CAA). Note that this figure only shows the cross-attention between one group of corresponding points for clear visualization. However, the same process is applied to all groups of corresponding points.

image. Input settings for other branches align with those specified in Variant 1. Particularly, Variant 3 enforces consistency between views by implementing CAA not only among the generation branches but also between the conditional



**Fig. 6:** Qualitative comparison between different panorama generation strategies. For hard samples, as shown above, The variants involving unwarping image and unwarping latent (Variant 1 and 2) could lead to different issues.

branch and the generation branches of canonical views. This strategy enables the effective reduction of the inaccuracies associated with homography estimation. We also illustrate our Variant 3 in Fig. 4(c).

## 5 Experiments

We provide experiment details in this section, including experiment setup (5.1, baselines (5.2), and experiment results (5.3).

### 5.1 Experiment setup

**Dataset.** We fine-tune our CamFreeView model on the real-world Matterport3D [3] dataset, which contains 90 building-scale indoor scenes with 10,912 high-resolution panoramic images. 9820 and 1092 images are split for training and evaluation, respectively, following MVDiffusion [29]. Each room in the dataset provides six distinct non-overlapping perspective images taken from identical



**Fig. 7:** Comparison between baseline MVDiffusion and proposed CamFreeDiff. Note that we provide camera parameter estimation results to MVDiffusion since it has not been trained to handle arbitrary camera views. We observe that CamFreeDiff still shows stronger robustness to camera-free input.

**Table 1:** We report the panorama generation results of our proposed CamFreeDiff on Matterport3D dataset. For baseline PanoDiffusion and MVDiffusion, we use our homography predictor to rectify the camera-free input image first, and then evaluate their panorama generations only. Therefore, PSNR is not reported for these baselines.

Model	FID ↓	IS ↑	CS↑	PSNR↑
PanoDiffusion	48.7	3.1	-	-
MVDiffusion	42.4	5.4	21.9	-
<b>CamFreeDiff</b>				
+ unwarped image	35.2	5.5	23.6	18.7
+ unwarped latent	34.3	5.6	22.4	15.6
+ <b>new view</b>	<b>27.0</b>	<b>5.6</b>	<b>24.4</b>	<b>19.3</b>

**Table 2:** We test our model on the Structured3D(S3D) dataset without applying fine-tuning or domain transfer techniques to demonstrate the generalization ability. Note that PanoDiffusion is trained on Structured3D without text guidance. Therefore, we provide no CLIP score (CS).

Model	Training	Cam-free	FID ↓	IS ↑	CS↑	PSNR↑
PanoDiffusion	✓	✗	35.3	3.2	-	-
<b>CamFreeDiff</b>	✗	✓	<b>31.1</b>	<b>6.2</b>	25.5	19.8



camera positions, with each offering a  $90^\circ$  field of view. To reach our goal of learning 360-degree image-to-panorama outpainting from camera-free input, we apply a random warp on each perspective image with a field of view from  $60^\circ$  to  $110^\circ$  and camera rotations of  $\pm 15^\circ$  to create camera-free images. After random warping, all input images are in  $512 \times 512$  resolution.

In addition to the primary dataset, we conducted further evaluation of our model on the Structured3D [36] dataset, a photo-realistic compilation of 3,500 indoor scenes encompassing 21,835 rooms, each rendered with panoramic images. Perspective images for each room are also generated from random camera positions and poses. This step aims to rigorously assess the model’s generalization abilities on out-of-domain data. We applied the same random warp approach in Structured3D as in Matterport3D.

BLIP-2 [10] captioning model is used to generate per-view text descriptions for both datasets mentioned above.

**Training Details.** Our CamFreeDiff model is fine-tuned from the stable diffusion inpainting model. We retain the weights of VAE image encoder/decoder and the latent denoising U-Net blocks frozen as pre-trained. We optimize the MLP block for homography prediction and the CAA blocks for multi-view consistency with a learning rate  $2 \times 10^{-4}$  for 30 epochs.

**Evaluation Metrics.** Our text-guided 360-degree image outpainting employs a series of standard image generation metrics to evaluate visual quality: Frechet Inception Distance (FID) [8], which quantifies the distance between real and generated images; Inception Score (IS) [22] offering insight into the diversity and quality of generated images; and CLIP score [16] to measure the alignment between text description and corresponding images. In addition, we use the Peak signal-to-noise ratio (PSNR) on the corresponding region between the generated and target canonical view  $0^\circ$  to evaluate our view estimation error. We also use Mean Absolute Error to assess the accuracy of homography estimation only.


## 5.2 Baselines

We consider the following baselines: MVDiffusion and PanoDiffusion.

*MVDiffusion* is a multi-view text-to-image diffusion model to generate view-consistent 360-degree scenes. For comparison, our model learns to generate from camera-free input with unknown camera parameters.

*PanoDiffusion* is designed for RGB-D panorama outpainting with different types of masks. A super-resolution model further enhances the outpainting results with higher resolution.

To the best of our knowledge, no existing panorama generation model was designed or learned to handle input images from different camera poses (rotation) and focal lengths (field of view). Therefore, by default, we provide all baselines with rectified input images using predicted camera parameters by our homography estimator for fair comparison. In addition, we also compare our three variants.



**Fig. 8:** More qualitative results of camera-free panorama generation from our Cam-FreeDiff model.

**Table 3:** Comparison of the architecture design for homography estimator.

Model	FID ↓	IS ↑	CS↑	PSNR↑
HomographyNet	29.2	5.6	24.2	19.2
SD encoder(frozen) + MLP	<b>27.0</b>	5.6	<b>24.4</b>	<b>19.3</b>

**Table 4:** Ablation study of different types of homography estimation objectives. Experiments are based on HomographyNet and evaluated on randomly warped images from Matterport3D. Predictions are evaluated in terms of Mean Absolute Error.

Objectives	fov↓	phi( $\phi$ )↓	psi( $\psi$ )↓
MSE (reg.)	10.6	2.5	2.4
CE (cls.)	<b>7.9</b>	<b>1.8</b>	<b>1.5</b>

### 5.3 Results

**Qualitative Results** In Fig. 7 we show qualitative comparison between baseline MVDiffusion and CamFreeDiff. Both MVDiffusion and CamFreeDiff are trained on Matterport3D dataset. CamFreeDiff is designed and trained to handle arbitrary camera parameters, while MVDiffusion is not. We report qualitative results for different variants of outpainting strategies as introduced in Sec. 4.2. Qualitative comparison for a hard sample is shown in Fig. 6. We observe that the homography prediction error can be propagated into later panorama generation for Variant 1(unwarping image), resulting in inconsistent scene layout. For Variant 2(unwarping latent), inaccurate homography prediction could lead to changed textures for the input view. Compared to the above variants, introducing a new conditional branch for the input view and leveraging correspondence-aware attention(Variant 3) enables the most consistent and robust multi-view generation. We provide more qualitative results for CamFreeDiff in Fig. 8.

**Quantitative Results** We show the numerical results for panorama generation between baselines and different variants of our CamFreeDiff model in Tab. 1. Among them, CamFreeDiff with Variant 3, treating the input as a new view, achieved the best results in terms of visual quality for panorama generation (FID, IS, CS) and reconstruction quality (PSNR).

To demonstrate the generalization ability of our model, we also tested our model on the Structured3D [36] dataset. Note that our model was never trained on or applied with domain transfer techniques from Structured3D. Results shown in Tab. 2 indicate the strong generalization ability of CamFreeDiff to out-of-domain data, even surpassing PanoDiffusion, which is trained directly on Structured3D but without learning from camera-free input.

**Table 5:** Ablation study on the neighbour size  $K$  in CAA blocks for information aggregation between two views.

Neighbour Size	FID ↓	IS ↑	CS↑	PSNR↑
$K = 1$	66.4	5.5	22.3	19.2
$K = 3$	51.1	5.8	23.9	<b>19.5</b>
$K = 5$	48.0	<b>6.2</b>	<b>24.1</b>	19.4
$K = 7$	<b>47.3</b>	6.1	<b>24.1</b>	19.4

## 5.4 Ablation Study

**Homography Estimator** We compared classification and regression as different types of homography estimators. Cross-entropy loss is used as the objective for the classifier, and mean squared error (MSE) loss is for the regression model. The classifier gives the best input view estimation results instead of the regression model, as shown in table Tab. 4. In addition, we ablate on the architecture design of the homography estimator. We compared our design of an MLP block built on a frozen stable diffusion image encoder with the HomographyNet, which is designed to predict homography matrix between views. Our designs achieve better generation results as shown in Tab. 3.

**Neighborhood Size in CAA blocks** Given correspondences between views, Correspondence-aware attention (CAA) aggregates information from source point  $p_s$  neighborhood to target point  $p_t$ , which is the key to yielding consistency between multiple views. The neighborhood in CAA refers to the  $K \times K$  neighboring points centered at  $p_s$ . We ablate the neighborhood size  $K$  for  $K = 1, 3, 5, 7$ . From results shown in Tab. 5, we find larger neighborhood size generally leads to better multi-view generation quality, but the improvement is limited. Note that larger  $K$  results in more computational and time complexity for CAA operation.

## 6 Conclusion

This paper introduces a novel method for generating panorama from a camera-free input image. We formulate the camera parameter estimation of the input as an estimation of the homography matrix from the input view to a predefined canonical view of the scene. Our method builds upon the MVDiffusion model for multi-view image generation and incorporates correspondences between the input and target canonical views for coherent and consistent panorama generation. We demonstrate the strong robustness of our model to camera-free inputs and generalization ability to out-of-domain data.

## References

1. Akimoto, N., Kasai, S., Hayashi, M., Aoki, Y.: 360-degree image completion by two-stage conditional gans. In: 2019 IEEE International Conference on Image Processing (ICIP). pp. 4704–4708. IEEE (2019) [3](#)
2. Akimoto, N., Matsuo, Y., Aoki, Y.: Diverse plausible 360-degree image outpainting for efficient 3dcg background creation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11441–11450 (2022) [2](#)
3. Chang, A., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., Zeng, A., Zhang, Y.: Matterport3d: Learning from rgb-d data in indoor environments. arXiv preprint arXiv:1709.06158 (2017) [9](#)
4. Chen, Z., Wang, G., Liu, Z.: Text2light: Zero-shot text-driven hdr panorama generation. ACM Transactions on Graphics (TOG) **41**(6), 1–16 (2022) [2](#), [3](#)
5. DeTone, D., Malisiewicz, T., Rabinovich, A.: Deep image homography estimation. arXiv preprint arXiv:1606.03798 (2016) [5](#)
6. Elfving, S., Uchibe, E., Doya, K.: Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. Neural networks **107**, 3–11 (2018) [6](#)
7. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. Communications of the ACM **63**(11), 139–144 (2020) [3](#)
8. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017) [11](#)
9. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020) [2](#)
10. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. arXiv preprint arXiv:2301.12597 (2023) [11](#)
11. Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9298–9309 (2023) [3](#)
12. Liu, Y., Lin, C., Zeng, Z., Long, X., Liu, L., Komura, T., Wang, W.: Syncdreamer: Generating multiview-consistent images from a single-view image. arXiv preprint arXiv:2309.03453 (2023) [3](#)
13. Long, X., Guo, Y.C., Lin, C., Liu, Y., Dou, Z., Liu, L., Ma, Y., Zhang, S.H., Habermann, M., Theobalt, C., et al.: Wonder3d: Single image to 3d using cross-domain diffusion. arXiv preprint arXiv:2310.15008 (2023) [3](#)
14. Oh, C., Cho, W., Chae, Y., Park, D., Wang, L., Yoon, K.J.: Bips: Bi-modal indoor panorama synthesis via residual depth-aided adversarial learning. In: European Conference on Computer Vision. pp. 352–371. Springer (2022) [2](#)
15. Van den Oord, A., Kalchbrenner, N., Espeholt, L., Vinyals, O., Graves, A., et al.: Conditional image generation with pixelcnn decoders. Advances in neural information processing systems **29** (2016) [3](#)
16. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021) [11](#)
17. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 **1**(2), 3 (2022) [2](#), [3](#)

18. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: International Conference on Machine Learning. pp. 8821–8831. PMLR (2021) [2](#)
19. Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H.: Generative adversarial text to image synthesis. In: International conference on machine learning. pp. 1060–1069. PMLR (2016) [2](#)
20. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022) [2](#), [3](#)
21. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems* **35**, 36479–36494 (2022) [2](#), [3](#)
22. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. *Advances in neural information processing systems* **29** (2016) [11](#)
23. Shi, R., Chen, H., Zhang, Z., Liu, M., Xu, C., Wei, X., Chen, L., Zeng, C., Su, H.: Zero123++: a single image to consistent multi-view diffusion base model. *arXiv preprint arXiv:2310.15110* (2023) [3](#)
24. Shi, Y., Wang, P., Ye, J., Long, M., Li, K., Yang, X.: Mvdream: Multi-view diffusion for 3d generation. *arXiv preprint arXiv:2308.16512* (2023) [3](#)
25. Somanath, G., Kurz, D.: Hdr environment map estimation for real-time augmented reality. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11298–11306 (2021) [3](#)
26. Stan, G.B.M., Wofk, D., Fox, S., Redden, A., Saxton, W., Yu, J., Aflalo, E., Tseng, S.Y., Nonato, F., Muller, M., et al.: Ldm3d: Latent diffusion model for 3d. *arXiv preprint arXiv:2305.10853* (2023) [3](#)
27. Sumantri, J.S., Park, I.K.: 360 panorama synthesis from a sparse set of images with unknown field of view. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 2386–2395 (2020) [2](#)
28. Tang, S., Chen, J., Wang, D., Tang, C., Zhang, F., Fan, Y., Chandra, V., Furukawa, Y., Ranjan, R.: Mvdifusion++: A dense high-resolution multi-view diffusion model for single or sparse-view 3d object reconstruction. *arXiv preprint arXiv:2402.12712* (2024) [3](#)
29. Tang, S., Zhang, F., Chen, J., Wang, P., Yasutaka, F.: Mvdifusion: Enabling holistic multi-view image generation with correspondence-aware diffusion. *arXiv preprint 2307.01097* (2023) [2](#), [3](#), [6](#), [9](#)
30. Wang, J., Chen, Z., Ling, J., Xie, R., Song, L.: 360-degree panorama generation from few unregistered nfov images. *arXiv preprint arXiv:2308.14686* (2023) [3](#)
31. Wang, P., Shi, Y.: Imagedream: Image-prompt multi-view diffusion for 3d generation. *arXiv preprint arXiv:2312.02201* (2023) [3](#)
32. Wu, T., Zheng, C., Cham, T.J.: Panodiffusion: 360-degree panorama outpainting via diffusion. In: The Twelfth International Conference on Learning Representations (2023) [2](#), [3](#)
33. Ying, Z., Bovik, A.: 180-degree outpainting from a single image. *arXiv preprint arXiv:2001.04568* (2020) [2](#)
34. Zheng, C., Cham, T.J., Cai, J.: Pluralistic image completion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1438–1447 (2019) [3](#)



35. Zheng, C., Cham, T.J., Cai, J., Phung, D.: Bridging global context interactions for high-fidelity image completion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11512–11522 (2022) [2](#)
36. Zheng, J., Zhang, J., Li, J., Tang, R., Gao, S., Zhou, Z.: Structured3d: A large photo-realistic dataset for structured 3d modeling. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX 16. pp. 519–535. Springer (2020) [11](#), [13](#)