

# Diffusion360: Seamless 360 Degree Panoramic Image Generation based on Diffusion Models

Mengyang Feng, Jinlin Liu, Miaomiao Cui, Xuansong Xie  
Alibaba Group

{mengyang.fmy, lj1191782, miaomiao.cmm, xingtong.xxs}@alibaba-inc.org




Figure 1. The proposed circular blending operation.

## Abstract

This is a technical report on the 360-degree panoramic image generation task based on diffusion models. Unlike ordinary 2D images, 360-degree panoramic images capture the entire  $360^\circ \times 180^\circ$  field of view. So the rightmost and the leftmost sides of the 360 panoramic image should be continued, which is the main challenge in this field. However, the current diffusion pipeline is not appropriate for generating such a seamless 360-degree panoramic image. To this end, we propose a circular blending strategy on both the denoising and VAE decoding stages to maintain the geometry continuity. Based on this, we present two models for **Text-to-360-panoramas** and **Single-Image-to-360-panoramas** tasks. The code has been released as an open-source project at <https://github.com/ArcherFMY/SD-T2I-360PanoImage> and [ModelScope](#).

## 1. Related Work

Recent studies like MVDiffusion [3], StitchDiffusion [4], and PanoDiff [5] have proved the feasibility of diffusion-based 360-degree panoramic images generation, but still have some drawbacks.

**MVDiffusion** needs 8 perspective views (user-provided or generated from Stable Diffusion [1]) as inputs. The resulting closed-loop panoramic image is more like a long-range image with a wide angle. So it has artifacts on the 'sky' and 'floor' when viewing in a 360 image viewer.

**StitchDiffusion** proposes a global cropping on the left and right side of the image to maintain the continuity. However, it still cracks on the junctions when zoom-in in the 360 image viewer.

**PanoDiff**, similar to the **StitchDiffusion**, proposes a circular padding scheme, which is the most related research to our work. The idea of our circular blending strategy is derived from the circular padding scheme. The differences are (1) we use an adaptive weighting policy for geometric continuity, (2) we do not need the *Rotating Schedule* at both training and inference time, which means that we can directly finetune a dreambooth [2] model using standard diffusion pipeline for this task, and just apply the circular blending at inference time, and (3) we can directly apply our technique into the ControlNet-Tile [8] model to produce high-resolution results.




Figure 2. The circular blending operation in different stages.




Figure 3. The pipeline of Text-to-360-Panoramas.

## 2. Method

### 2.1. Circular Blending

We propose a circular blending strategy at the inference time to generate seamless 360-degree panoramic images. Specifically, at each **denoising step**, the right part (of a such portion) of the latent feature and the left part (of the same portion as the right part) is blended with adaptive weights. This is illustrated in Fig. 1. Similarly, this strategy can be added to the `tiled_decode` function of the VAE decoder (see Fig. 2). We find that using the circular blending in the VAE decoder is more important than in the latent denoising stage for maintaining the geometric continuity.

### 2.2. Text-to-360-Panoramas

For the Text-to-360-Panoramas task, we propose a multi-stage framework to generate high resolution 360-degree panoramic images. As illustrated in Fig. 3, we first generate a low resolution image using a base model (finetuned on the SUN360 [7] dataset using the DreamBooth [2] training method), and then employ some super-resolution strategies (including diffusion-based and the GAN-based methods, like the ControlNet-Tile model and the RealESRGAN [6]) to up-scale the result to a high resolution one. For better results, we also finetune the ControlNet-Tile model on the SUN360 dataset by generate low-resolution and high-resolution image pairs.

### 2.3. Single-Image-to-360-Panoramas

For the Single-Image-to-360-Panoramas task, the framework is similar to the Text-to-360-Panoramas by replacing the base model to a controlnet-outpainting model. We

design a ControlNet-Outpainting model to generate a low resolution 360-degree panoramic image from a given single ordinary 2D image at perspective view. To generate the training pairs of perspective and panoramic images, we first convert the panoramic image to cube-maps and select the center-cube as its perspective image. The inputs of the ControlNet-Outpainting model consist of the converted center-cube map  $C$  with the other cubes filled by zeros and the mask  $M$ . At inference time, the perspective image can be generated from a certain generative model or captured by a camera (the image should be squared). The perspective image is converted to the center-cube map  $C$  as the input of the ControlNet-Outpainting model. For some reason, the trained models of this task can not be released. However, it should be easy to reproduce. See some results in Fig. 8.




## 3. Results

We show some testing results at different stages of the Text-to-360-Panoramas task in Fig. 4, Fig. 5, Fig. 6, and Fig. 7. The input prompts it fetch at the MVDiffusion project page (<https://mvdifffusion.github.io/>)

## 4. Limitations




The base model is trained using the DreamBooth [2] technique, so it can not be changed with the models from CIVITAI (<https://civitai.com/>) for stylizing purposes. Adding some style descriptions (such as 'cartoon style' and 'oil painting style') in the prompt does not work. One can generate an initial 360 image using our method, and then use ControlNets (like canny and depth) with different base models to change the style.

## Results from base model



*This kitchen is a charming blend of rustic and modern, featuring a large reclaimed wood island with marble countertop, a sink surrounded by cabinets. To the left of the island, a stainless-steel refrigerator stands tall. To the right of the sink, built-in wooden cabinets painted in a muted.*

*Majestically rising towards the heavens, the snow-capped mountain stood, its jagged peaks cloaked in a shroud of ethereal clouds, its rugged slopes a stark contrast against the serene azure sky, and its silent grandeur exuding an air of ancient wisdom and timeless solitude, commanding awe and reverence from all who beheld it.*



*Bathed in the soft, dappled light of the setting sun, the silent street lay undisturbed, revealing the grandeur of its cobblestone texture, the rusted lampposts bearing witness to forgotten stories, and the ancient, ivy-clad houses standing stoically, their shuttered windows and weather-beaten doors speaking volumes about their passage through time.*

*This kitchen is a charming blend of rustic and modern, featuring a large reclaimed wood island with marble countertop, a sink surrounded by cabinets. To the left of the island, a stainless-steel refrigerator stands tall. To the right of the sink, built-in wooden cabinets painted in a muted.*

Figure 4. Results from the Base Model.





## Results from base+initsr



*This kitchen is a charming blend of rustic and modern, featuring a large reclaimed wood island with marble countertop, a sink surrounded by cabinets. To the left of the island, a stainless-steel refrigerator stands tall. To the right of the sink, built-in wooden cabinets painted in a muted.*



*Majestically rising towards the heavens, the snow-capped mountain stood, its jagged peaks cloaked in a shroud of ethereal clouds, its rugged slopes a stark contrast against the serene azure sky, and its silent grandeur exuding an air of ancient wisdom and timeless solitude, commanding awe and reverence from all who beheld it.*



*Bathed in the soft, dappled light of the setting sun, the silent street lay undisturbed, revealing the grandeur of its cobblestone texture, the rusted lampposts bearing witness to forgotten stories, and the ancient, ivy-clad houses standing stoically, their shuttered windows and weather-beaten doors speaking volumes about their passage through time.*





*This kitchen is a charming blend of rustic and modern, featuring a large reclaimed wood island with marble countertop, a sink surrounded by cabinets. To the left of the island, a stainless-steel refrigerator stands tall. To the right of the sink, built-in wooden cabinets painted in a muted.*

Figure 5. Results from Base+InitSR.




## Results from base+initsr+realesrgan



*This kitchen is a charming blend of rustic and modern, featuring a large reclaimed wood island with marble countertop, a sink surrounded by cabinets. To the left of the island, a stainless-steel refrigerator stands tall. To the right of the sink, built-in wooden cabinets painted in a muted.*

*Majestically rising towards the heavens, the snow-capped mountain stood, its jagged peaks cloaked in a shroud of ethereal clouds, its rugged slopes a stark contrast against the serene azure sky, and its silent grandeur exuding an air of ancient wisdom and timeless solitude, commanding awe and reverence from all who beheld it.*




*Bathed in the soft, dappled light of the setting sun, the silent street lay undisturbed, revealing the grandeur of its cobblestone texture, the rusted lampposts bearing witness to forgotten stories, and the ancient, ivy-clad houses standing stoically, their shuttered windows and weather-beaten doors speaking volumes about their passage through time.*

*This kitchen is a charming blend of rustic and modern, featuring a large reclaimed wood island with marble countertop, a sink surrounded by cabinets. To the left of the island, a stainless-steel refrigerator stands tall. To the right of the sink, built-in wooden cabinets painted in a muted.*


Figure 6. Results from Base+InitSR+ReslESRGAN. It can be observed that, the geometric continuity of the rightmost and the leftmost sides of our results are smooth and nearly no cracks. Some artifacts in the top two rows are cost by the RealESRGAN.

## Results from our full implementation



*This kitchen is a charming blend of rustic and modern, featuring a large reclaimed wood island with marble countertop, a sink surrounded by cabinets. To the left of the island, a stainless-steel refrigerator stands tall. To the right of the sink, built-in wooden cabinets painted in a muted.*

*Majestically rising towards the heavens, the snow-capped mountain stood, its jagged peaks cloaked in a shroud of ethereal clouds, its rugged slopes a stark contrast against the serene azure sky, and its silent grandeur exuding an air of ancient wisdom and timeless solitude, commanding awe and reverence from all who beheld it.*



*Bathed in the soft, dappled light of the setting sun, the silent street lay undisturbed, revealing the grandeur of its cobblestone texture, the rusted lampposts bearing witness to forgotten stories, and the ancient, ivy-clad houses standing stoically, their shuttered windows and weather-beaten doors speaking volumes about their passage through time.*

*This kitchen is a charming blend of rustic and modern, featuring a large reclaimed wood island with marble countertop, a sink surrounded by cabinets. To the left of the island, a stainless-steel refrigerator stands tall. To the right of the sink, built-in wooden cabinets painted in a muted.*


Figure 7. Results from the full implementation.



*A living  
room  
(generated)*



*The  
mountain  
road  
(generated)*



*The outer  
space  
(generated)*



*The Times  
Square  
(generated)*



*A office  
room  
(captured  
by a camera)*




Figure 8. Results of Single-Image-to-360-Panoramas.



## References

- [1] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. [1](#)
- [2] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. 2022. [1](#), [2](#)
- [3] Shitao Tang, Fuyang Zhang, Jiacheng Chen, Peng Wang, and Yasutaka Furukawa. Mvdiffusion: Enabling holistic multi-view image generation with correspondence-aware diffusion. *arXiv*, 2023. [1](#)
- [4] Hai Wang, Xiaoyu Xiang, Yuchen Fan, and Jing-Hao Xue. Customizing 360-degree panoramas through text-to-image diffusion models, 2023. [1](#)
- [5] Jionghao Wang, Ziyu Chen, Jun Ling, Rong Xie, and Li Song. 360-degree panorama generation from few unregistered nfov images. In *Proceedings of the 31th ACM International Conference on Multimedia*, 2023. [1](#)
- [6] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *International Conference on Computer Vision Workshops (ICCVW)*. [2](#)
- [7] Jianxiong Xiao, Krista A. Ehinger, Aude Oliva, and Antonio Torralba. Recognizing scene viewpoint using panoramic place representation. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2695–2702, 2012. [2](#)
- [8] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. [1](#)