


# 2S-ODIS: Two-Stage Omni-Directional Image Synthesis by Geometric Distortion Correction

Atsuya Nakata<sup>✉</sup> and Takao Yamanaka<sup>✉</sup>

Sophia University, Tokyo, Japan

a-nakata-7r0@eagle.sophia.ac.jp, takao-y@sophia.ac.jp

**Abstract.** Omni-directional images have been increasingly used in various applications, including virtual reality and SNS (Social Networking Services). However, their availability is comparatively limited in contrast to normal field of view (NFoV) images, since specialized cameras are required to take omni-directional images. Consequently, several methods have been proposed based on generative adversarial networks (GAN) to synthesize omni-directional images, but these approaches have shown difficulties in training of the models, due to instability and/or significant time consumption in the training. To address these problems, this paper proposes a novel omni-directional image synthesis method, 2S-ODIS (Two-Stage Omni-Directional Image Synthesis), which generated high-quality omni-directional images but drastically reduced the training time. This was realized by utilizing the VQGAN (Vector Quantized GAN) model pre-trained on a large-scale NFoV image database such as ImageNet without fine-tuning. Since this pre-trained model does not represent distortions of omni-directional images in the equi-rectangular projection (ERP), it cannot be applied directly to the omni-directional image synthesis in ERP. Therefore, two-stage structure was adopted to first create a global coarse image in ERP and then refine the image by integrating multiple local NFoV images in the higher resolution to compensate the distortions in ERP, both of which are based on the pre-trained VQGAN model. As a result, the proposed method, 2S-ODIS, achieved the reduction of the training time from 14 days in OmniDreamer to four days in higher image quality.



**Fig. 1:** Overview of advantages of proposed method, 2S-ODIS. OmniDreamer [2] requires 14 days for training of the model, including 1-week training of the VQGAN model. In contrast, the proposed method only required 4 days for the training of the model since no training of VQGAN model was required.

## 1 Introduction

An omni-directional image is taken by an omni-directional camera, also known as a 360-degree camera, which captures lights from all directions falling onto the focal point to cover a full sphere, and is usually represented in the equi-rectangular projection (ERP) to represent it in a 2-dimensional plane. These images have been applied to various applications such as virtual reality, social networking services, and map tools such as Google Street View. However, the availability of the omni-directional images are still limited compared with Normal Field of View (NFoV) images captured by a standard camera, since the specialized camera is required to capture the omni-directional images.

Recently, several methods have been proposed to synthesize omni-directional images from a NFoV image [2,6,8,10,13], a text condition [5], or both [12]. These methods have been based on Generative Adversarial Networks (GAN) [6, 8, 10, 13], Vector Quantized GAN (VQGAN) [2,5], or auto-regressive outpainting using stable diffusion [12]. However, the GAN-based methods have faced challenges of instability in training, while the VQGAN-based methods and the auto-regression with stable diffusion require long training and inference time, respectively.

In this paper, we propose a novel method to synthesize omni-directional images from a NFoV image based on pre-trained VQGAN, trained on a large-scale NFoV-image dataset. The previous method with VQGAN [2] has required to train the VQGAN encoder and decoder to represent geometric distortions in omni-directional images in ERP, especially at poles (top and bottom regions in ERP). This training has taken long time, such as 1 week in the method [2]. In the proposed method, the pre-trained VQGAN encoder and decoder were applied without fine-tuning by synthesizing multiple NFoV images to integrate them into an omni-directional image based on geometric distortion correction. Since no training of VQGAN was required, the training of the model was shortened by removing the step of the fine-tuning of VQGAN, as shown in Fig. 1. Furthermore, a two-stage structure was adopted in the proposed method. At the first stage, a global coarse image in ERP is created using the pre-trained VQGAN encoder and decoder without the geometric distortion correction. Therefore, the generated omni-directional image at the first stage includes distortions. For example, a straight line in NFoV images at poles can not be reproduced at the first stage. At the second stage, this global coarse image is refined by synthesizing an omni-directional image from multiple NFoV images generated using the pre-trained VQGAN encoder and decoder. This second stage compensates the geometric distortions at the first stage, in addition to representing local detailed textures in a higher resolution. By using the two-stage structure, the model can produce globally plausible yet locally detailed omni-directional images without the geometric distortions.

The contributions of this paper include:

- A novel method to synthesize omni-directional images from a NFoV image was proposed using pre-trained VQGAN. Since no training of VQGAN was required, the training time was drastically reduced.

- A two-stage structure was adopted to generate a global coarse omni-directional image at the first stage, and then generate a locally detailed image with geometric distortion correction at the second stage.
- Experimental results demonstrated that the proposed method synthesized higher quality omni-directional images in shortened training and inference time than the previous methods such as OmniDreamer [2].

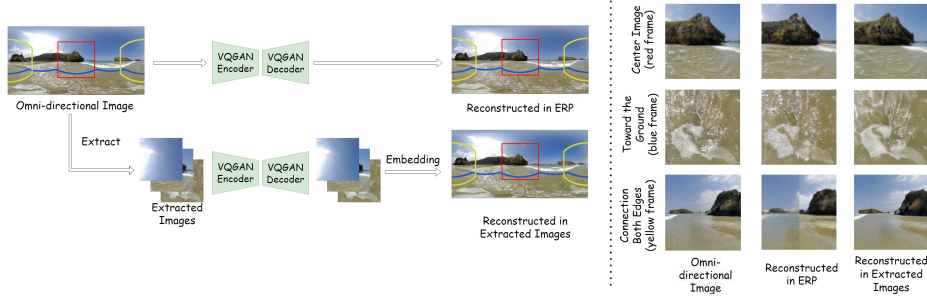
## 2 Related Works

### 2.1 Image Generation

VQVAE (Vector Quantized Variational AutoEncoder) [17] has been proposed to improve the generated image blurriness in VAE [11] by representing image patches with quantized latent vectors based on vector quantization. Furthermore, the adversarial loss has been introduced in VQVAE to make the generated images clearer, called VQGAN [7]. In this method, Transformer [18] has been used to sequentially predict image patches from neighbor patches based on auto-regressive prediction. The patches are represented with the quantized latent vectors called VQGAN codes to generate clear images with low computational cost. To improve the slow inference in VQGAN due to the sequential predictions of patches, MaskGIT [4] has been proposed by predicting multiple patches simultaneously. Although MaskGIT has succeeded in improving the inference speed, it has been difficult to generate high quality images in the high resolution. To solve this problem, Muse [3] has been proposed using a two-stage structure, where a low-resolution image is generated at the first stage, and then is refined to generate a higher-resolution image at the second stage. In our proposed method, this two-stage structure was adopted for the omni-directional image synthesis in the high resolution.

### 2.2 Omni-directional Image Generation

Several methods have been proposed for synthesizing omni-directional images from N FoV images. Okubo and Yamanaka [10] have proposed a method of generating omni-directional images based on conditional GAN from a single N FoV image with the class label. Hara et al. [8] have also proposed a method based on the symmetric property in the omni-directional images using GAN and VAE. Another work to synthesize omni-directional images is Guided ImmerseGAN [6], which generates omni-directional images from a N FoV image with the modulation guided by a given class label, which does not have to be the true class of the input N FoV image. In the work of OmniDreamer [2], VQGAN has been applied to the omni-directional image synthesis by using Transformer for auto-regressive prediction. In this method, VQGAN encoder and decoder have to be fine-tuned on an omni-directional image dataset since the geometric distortion in ERP has to be represented in the latent codes of VQGAN. Text2Light [5] also uses VQGAN with auto-regressive prediction for the generation, though only text



**Fig. 2:** Qualitative comparison of omni-directional image reconstruction using pre-trained VQGAN encoder and decoder. Omni-directional Image: original omni-directional image, Reconstructed in ERP: reconstructed in equirectangular projection, Reconstructed in Extracted Images: reconstructed by integrating multiple NFoV images in different directions. By extracting NFoV images, an omni-directional image can be correctly reconstructed without distortions.

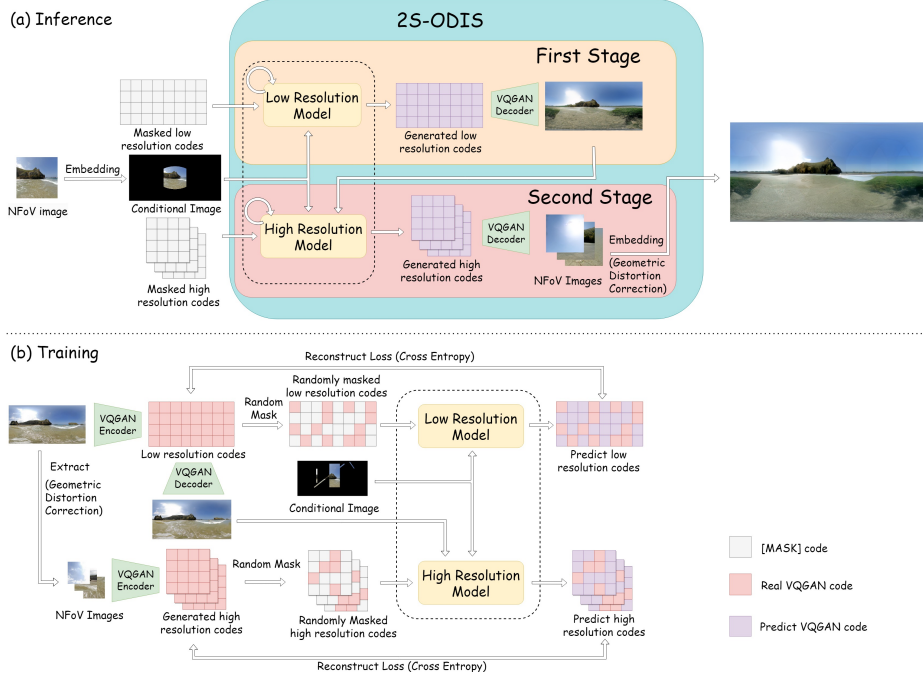
information is taken as the input instead of the NFoV image. Nakata et al. [13] have proposed a method to increase the diversity of generated omni-directional images based on MLP-Mixer by efficiently propagating the information of the NFoV image embedded at the center in ERP. AOGNet [12] has generated omni-directional images by out-painting an incomplete 360-degree image progressively with NFoV and text guidances jointly or individually. This has been realized using auto-regressive prediction based on the stable-diffusion backbone model. Due to the nature of sequential auto-regressive prediction, it takes long inference time.

In our proposed method, the pre-trained VQGAN model was used without fine-tuning on the omni-directional image dataset, since multiple NFoV images are synthesized based on the pre-trained VQGAN and then integrated into omni-directional images. By removing the step of VQGAN training, the overall training of the model was drastically shortened than the previous method with VQGAN such as OmniDreamer [2]. In addition, the proposed method is based on simultaneous synthesis of multiple NFoV images in different directions, whose inference was faster than auto-regressive prediction such as OmniDreamer [2] and AOGNet [12].

### 3 Proposed Method

#### 3.1 Two-Stage Structure

The proposed method consists of the two-stage structure, where a global coarse omni-directional image in ERP is synthesized in a low resolution ( $256 \times 512$  pixels) at the first stage without geometric distortion correction, and then is refined at the second stage by integrating the multiple synthesized NFoV images in different directions based on the geometric distortion correction, producing

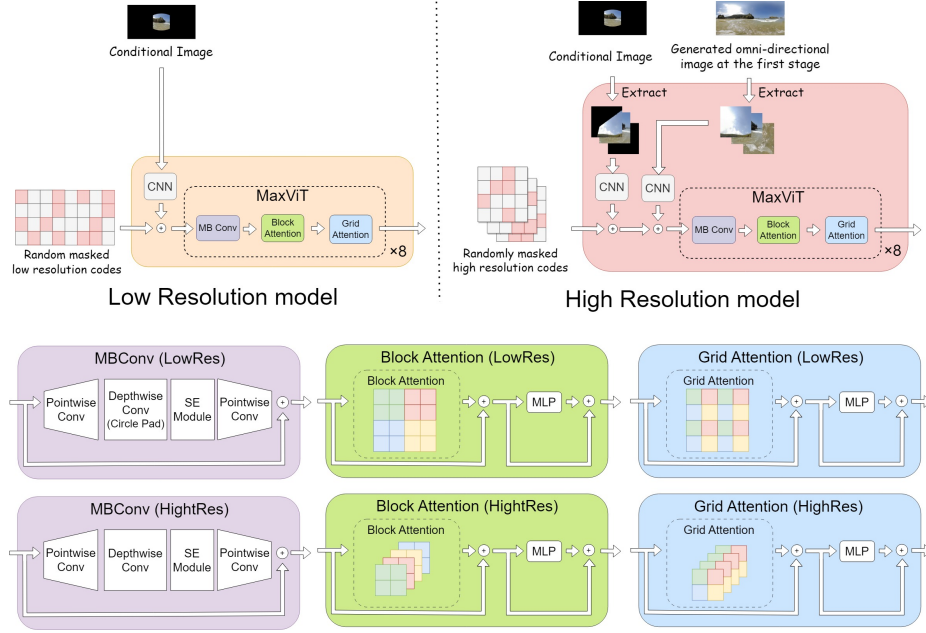


**Fig. 3:** Diagram of the proposed method. (a)Inference, (b)Training.

a high-quality omni-directional image in ERP in a high resolution ( $1024 \times 2048$  pixels). At both stages, the pre-trained VQGAN was utilized without fine-tuning on the omni-directional image dataset.

As a preliminary experiment, an omni-directional image was reconstructed in ERP or in multiple NFOV images using pre-trained VQGAN encoder and decoder without the fine-tuning, as shown in Fig. 2. It can be seen from the figure that the reconstruction in ERP cannot correctly reproduce the texture in the region toward the ground (blue frame) and the continuity in the region at both edges (yellow frame), although it can reproduce the region at center in ERP (red frame). On the contrary, all the regions can be correctly reproduced in the extracted NFOV images. This indicates that the pre-trained VQGAN model can be applied without fine-tuning if it is applied to NFOV images.

Thus, the generated omni-directional image in ERP at the first stage in the proposed method includes distortions since the pre-trained VQGAN cannot represent the texture and continuities in the omni-directional images in ERP. However, these distortions are correctly compensated at the second stage by synthesizing the multiple NFOV images which can be correctly reproduced by the pre-trained VQGAN model. If only the second stage is used in the proposed method, it is difficult to synthesize multiple NFOV images simultaneously with global compatibility. Therefore, the two-stage structure was adopted in the pro-



**Fig. 4:** Structure of the proposed method. The range of attention differs between the high-resolution and low-resolution models.

posed method to produce globally plausible coarse omni-directional image at the first stage.

The structure of the proposed method is shown in Fig. 3. At the first stage, the low resolution model produces the low resolution codes, which are converted into patches of omni-directional images in ERP using the pre-trained VQGAN decoder. At the second stage, the high resolution model produces the high resolution codes, which are corresponding to the patches in the N FoV images in multiple directions (26 directions in our implementation) in an omni-directional image with overlapping. These 26 directions for the N FoV images were the same directions as the normal vectors in the faces of a rhombicuboctahedron. The field of view was set to 60 degrees for all directions. The generated N FoV images with the size of  $256 \times 256$  pixels were integrated into an omni-directional image with the size of  $1024 \times 2048$  pixels in ERP.

### 3.2 Inference

For synthesizing an omni-directional image in the inference, the low-resolution codes are first generated using the sampling strategy proposed in MaskGIT [4] at the first stage from the conditional image where an input N FoV image is embedded at the center in ERP, as shown in Fig. 3(a). In MaskGIT, the generation is started with ‘Masked low resolution codes’ which is filled with the [MASK] code,

the VQGAN code which indicates that it is masked, for all locations. Then, the low resolution model predicts the probabilities for all the [MASK] locations in parallel, and samples a VQGAN code based on its predicted probabilities over all possible VQGAN codes for each location. The location with the low probability is replaced with the [MASK] code again, and the VQGAN codes are resampled by predicting the probabilities using the low resolution model. This process is repeated in  $T$  steps. At each iteration, the model predicts all VQGAN code simultaneously but only keeps the most confident ones. The remaining VQGAN codes are replaced with the [MASK] code and re-predicted in the next iteration. The mask ratio during the iterations is determined by  $\cos\left(\frac{\pi}{2} \frac{t}{T}\right)$ , where  $t$  indicates the current iterations in the total steps  $T$ . This mask ratio is monotonically decreasing from 1 to 0 with respect to  $t$ , which ensures that most of the locations are masked during the early stage in the iterations to prevent producing the inconsistent codes.

At the second stage, the high-resolution codes are generated using the high resolution model, where the generation process is almost same as the low resolution model. The difference from the low-resolution model is that the model accepts the low-resolution image generated at the first stage as an additional conditional image, and generates NFoV images, as shown in Fig. 3(a). To integrate the generated NFoV images into an omni-directional image, the overlapped regions are merged with weights depending on the distance from the centers of the NFoV images. Specifically, let  $x_i$  and  $x_j$  be the two overlapping pixel values, and let  $d_i$  and  $d_j$  be the distances from the centers of the NFoV images at each position. The integrated pixel value  $y$  is given by


$$y = \frac{w_i}{(w_i + w_j)} x_i + \frac{w_j}{(w_i + w_j)} x_j \quad (1)$$

where  $w_i = 1 - \frac{d_i}{\max_k(d_k)}$ ,  $w_j = 1 - \frac{d_j}{\max_k(d_k)}$ .

The network architecture used in the low-resolution model and the high resolution model is shown in Fig. 4. The layer structure was adapted from MaxViT [16], although Transformer has been used in MaskGIT [4] and Muse [3]. The 8-layer MaxViT models were used in both the low-resolution model and the high-resolution model. In the low-resolution model, the padding in MBConv was replaced to the circular padding to encourage the continuity at the edges in ERP, whereas it was remained to the zero padding in the high-resolution model. The block attention was applied within each divided region at the low-resolution model and within each NFoV image in the high-resolution model, as shown in Fig. 4. The grid attention was also applied globally in sparse at the low-resolution model as in the original MaxViT model, whereas it was applied among same locations over the NFoV images at the high-resolution model.

### 3.3 Training

The low-resolution and high-resolution models are independently trained, as shown in Fig. 3(b). The objective of the training is to make the low-resolution



**Fig. 5:** Examples of conditional image in training. These images are generated from omnidirectional images in ERP by randomly masking.

and high-resolution models predict plausible VQGAN codes at [MASK]-code locations for each inference step. For the low-resolution model, the inputs are ‘randomly masked low-resolution codes’ and a conditional image which emulates the NFoV image embedded at the center in ERP (Fig. 3 a). The mask ratio in the randomly masked low-resolution codes is set to  $\cos(\frac{\pi}{2}r)$ , where  $r$  is sampled from a uniform distribution  $[0, 1)$ , since this emulates the single iteration in the inference. Examples of the conditional image in the training are shown in Fig. 5. They are prepared by randomly masking an original omnidirectional image in ERP to emulate the conditional image in the inference. Since they are not limited to the single NFoV image embedded in ERP, the trained model can be applied to various in-painting and out-painting tasks, as described in 5.2. The low-resolution model is trained to predict the original VQGAN codes in the real omnidirectional image at [MASK]-code locations, so that the cross entropy is used as the loss function to train the model.

For the high-resolution model, the inputs are ‘randomly masked high-resolution codes’ for multiple NFoV images (26 NFoV images in our implementation), conditional NFoV images converted from the conditional image in ERP, and the reconstructed low-resolution NFoV images converted from the low-resolution omnidirectional image reconstructed using the pre-trained VQGAN encoder and decoder. The ‘randomly masked high-resolution code’ and the conditional image in ERP are prepared in the same manner for the low-resolution model. The low-resolution omnidirectional image is required in the inputs of the high-resolution model to emulate the low-resolution image generated at the first stage. The high-resolution model is trained to predict the original VQGAN codes in the NFoV images converted from the real omnidirectional image at [MASK]-code locations based on the cross-entropy loss function.

Although the first and second stages are sequentially processed in the inference, they are independently trained in parallel, as shown in Fig. 3(b). This property is advantageous to shorten the required training time if multiple GPUs (Graphics Processing Units) can be used, although a single GPU was used in our implementation.

## 4 Experiments

The omnidirectional image dataset, SUN360 [19], was used in the experiments. The 5,000 outdoor images were used for test, while the remaining 47,938 outdoor images were used for training. The size of the images in the dataset is  $512 \times 1024$



pixels. Although the proposed method generates images in  $1024 \times 2048$  pixels, it was resized to  $512 \times 1024$  pixels for the evaluation. For comparison, several conventional methods were also evaluated, including GAN-based methods with CNN (Convolutional Neural Networks) [10] or MLP-Mixer [13], a VQGAN-based method (OmniDreamer) [2], and a GAN-based in-painting method (LAMA) [15]. The models were implemented in PyTorch, and were trained in a single GPU (NVIDIA RTX3090). The code for the network architecture in the proposed method is provided in the supplementary material.

The pre-trained VQGAN was obtained from [1], which is the model with 1024 codebooks trained on ImageNet. Eight MaxViT layers in Fig. 4 were used in both the low-resolution and high-resolution models, with 256 internal dimensions. The sizes of the VQGAN codes at the first and second stages were  $16 \times 32$  patches and  $16 \times 16$  patches  $\times 26$  NFoV images, respectively. These VQGAN codes were converted into trainable feature vectors in the 256 dimensions. The conditional image was also down-sampled into the same size as the VQGAN codes with 256 dimensions using CNN. At each iteration of MaskGIT in the first stage, these two feature vectors were added with trainable positional encoding vectors, and then were inputted into the low-resolution model composed of the 8 MaxViT layers. At each iteration in the second stage, the low-resolution image generated at the first stage (inference) or reconstructed using the pre-trained VQGAN (training) was converted into NFoV images. These NFoV images were down-sampled with CNN to be added with the feature vectors of the input VQGAN codes, the conditional NFoV images, and the positional encoding, and then were inputted into the high-resolution model. The total steps  $T$  in MaskGIT was set to 16 at both stages.

The optimizer for training the models was AdamW with the learning rate of 0.001, the weight decay of  $1e-5$ , Amsgrad, and the learning-rate scheduling of ExponentialLR reducing it by 0.95 every 5,000 iterations. OmniDreamer [2] was trained for 14 days (30 epochs in all training stages), while the proposed method was trained over 4 days (2 days with the batch size 16 at the first stage, and 2 days with the batch size 8 in the second stage, for 180,000 iterations at each stage). The other conventional methods were trained for 4 days, using their default batch sizes and hyper-parameters. FID (Frechet Inception Distance) [9], IS (Inception Score) [14], and LPIPS (Learned Perceptual Image Patch Similarity) [20] were used for evaluating the synthesized omni-directional images in ERP.


## 5 Results

### 5.1 Evaluation of Proposed Method

The proposed method, 2S-ODIS, was quantitatively evaluated, compared with the conventional methods, OmniDreamer [2], CNN-based cGAN [10], MLP-Mixer-based cGAN [13], and LAMA [15], as shown in Table 1. The models in the proposed method was trained for 4 days (2 days for low-resolution model and 2 days for high-resolution model). For comparison, the result with the models trained for 2 days (1 day + 1 day for low-resolution and high-resolution models) was also

**Table 1:** Quantitative comparison with conventional methods

Method	IS ( $\uparrow$ )	FID ( $\downarrow$ )	LPIPS ( $\uparrow$ )
2S-ODIS (Proposed method)	<b>5.969</b>	<b>18.263</b>	0.662
2S-ODIS (2days)	5.857	18.656	<b>0.668</b>
OmniDreamer [2]	4.458	23.101	0.655
CNN-based cGAN [10]	4.684	40.049	0.633
MLPMixer-based cGAN [13]	4.402	47.690	0.634
LAMA [15]	5.784	69.485	0.478

**Fig. 6:** Evaluation metrics during training of proposed method compared with conventional method, OmniDreamer [2]**Fig. 7:** Examples of synthesized omni-directional images compared with conventional method, OmniDreamer [2]




**Fig. 8:** Examples of N FoV images toward ground extracted from synthesized omni-directional images in proposed and conventional methods

shown in the table. The field of view for a N FoV image embedded in an input conditional image was set to 126.87 and 112.62 degrees for width and height in the experiments, respectively. It can be seen from the table that the proposed method achieved higher performance than the other conventional methods. Although the highest performance was achieved in the proposed method trained for 4 days, the performance already outperformed the other methods even with the 2-day training. To see this more clearly, the evaluation metrics during the training of the proposed method is shown in Fig. 6, where the performance exceeded OmniDreamer by 2 days and converged in around 4 days. Thus, the proposed method drastically shortened the training to 2-4 days from 14 days in OmniDreamer including the fine-tuning of the VQGAN model. In addition, the inference in the proposed method was much faster than in OmniDreamer: 1.54 seconds and 39.33 seconds for synthesizing each omni-directional image in the proposed method and OmniDreamer, respectively. This is because the proposed method is based on the simultaneous VQGAN-code prediction instead of the sequential auto-regressive prediction.

For the qualitative comparison, the examples of the synthesized omni-directional images are shown in Fig. 7. It is clear that the proposed method generated globally plausible and locally detailed omni-directional images, while OmniDreamer [2] sometimes failed in generating continuous images, especially along the edges in the input conditional images. To see if the proposed method can generate N FoV images without the distortion, the N FoV images toward the ground were extracted from the synthesized omni-directional images, since the geometric distortion is large at the poles in ERP. The examples of the N FoV images toward the ground are shown in Fig. 8. The left column shows the N FoV images at the first stage, while the middle column shows the N FoV images at the second stage. Since the pre-trained VQGAN code cannot represent the geometric distortion in ERP, the straight lines were not appropriately reproduced at the first stage.

**Table 2:** Quantitative comparison in various conditional images

Mask Setting	Method	IS ( $\uparrow$ )	FID ( $\downarrow$ )	LPIPS ( $\uparrow$ )
Inpainting	2S-ODIS (Proposed)	<b>5.582</b>	<b>15.044</b>	0.685
	OmniDreamer [2]	4.672	41.209	<b>0.708</b>
Inpainting of Ground Region	2S-ODIS (Proposed)	<b>6.084</b>	<b>13.038</b>	0.680
	OmniDreamer [2]	5.474	15.303	<b>0.699</b>
Outpainting from Two Images	2S-ODIS (Proposed)	<b>5.722</b>	<b>19.437</b>	0.663
	OmniDreamer [2]	3.952	33.403	<b>0.672</b>

**Fig. 9:** Examples of synthesized omni-directional images in various input conditions.

However, they were compensated at the second stage, which can be clearly seen in the sample images. On the other hand, OmniDreamer cannot appropriately reproduce the NFoV images toward the ground. Thus, it was confirmed that the proposed method synthesized omni-directional images without geometric distortion.

## 5.2 Evaluation in Various Input Conditions

As explained in 3.3, the models in the proposed method were trained with various conditional images in Fig. 5. They were not limited to the single NFoV image embedded in ERP, so that the proposed model can be applied various inpainting and out-painting tasks. For example, the model can be applied to the in-painting task to remove objects or people in the omni-directional image taken by a 360-degree camera, as shown in the top row in Fig. 9. Another example is the task to fill in the ground region of an omni-directional image as shown in the middle row in Fig. 9, since the omni-directional image often includes a hand or a camera stand at the bottom region in ERP. The last example in Fig. 9 is to synthesize an omni-directional image which includes two NFoV images such as

**Table 3:** Ablation study in propose method

	IS ( $\uparrow$ ) FID ( $\downarrow$ ) LPIPS ( $\uparrow$ )		
(1) Proposed	<b>5.969</b>	<b>18.263</b>	0.662
(2) 1 Stage: Low Resolution Model	5.798	28.329	<b>0.670</b>
(3) 1 Stage: High Resolution Model	4.821	52.453	0.638
(4) Direct use of low-resolution VQGAN codes	5.837	21.820	0.663

front and rear cameras of a smartphone. Although OmniDreamer failed in synthesizing the omni-directional images in these situations, the proposed method generated high-quality omni-directional images. The quantitative results shown in Table 2 also indicate that the proposed method achieved higher performance than OmniDreamer. Although the diversity of the synthesized images was higher in OmniDreamer than the proposed method, it may be due to generating random images as shown in Fig. 9.

### 5.3 Ablation Study

Ablation studies were conducted to investigate the effectiveness of each component in the proposed method: the low-resolution and high-resolution models in the 2-stage structure. The results are shown in Table 3. (1) is the proposed method with the 2-stage structure, while (2) and (3) are the results with 1-stage structure only using the low-resolution model and the high-resolution model, respectively. As can be seen from the table, the 2-stage structure was indispensable for the high-quality image synthesis.

Moreover, it was examined that the low-resolution VQGAN codes generated at the first stage were directly used at the second stage instead of the low-resolution image generated at the first stage, since the 2-stage structure in Muse [3] uses the low-resolution VQGAN codes directly. The result is shown in Table 3 (4). It was confirmed from the result that the low-resolution image was better to use at the second stage than the low-resolution VQGAN codes directly. This may be because the low-resolution image is compressed by CNN to properly extract the global information generated at the first stage.

## 6 Limitations and Future Prospects

The proposed method synthesizes an omni-directional image by merging multiple NFoV images with weights depending on the distance from the edges of the NFoV images. Although the generation of the NFoV images are conditioned by the global low-resolution image generated at the first stage, it may be possible to generate discontinuous NFoV images. One possible solution would be to add an additional network to refine the generated omni-directional images to improve the continuity between NFoV images. Another issue is that it takes 1-2 days to

convert omni-directional images in the dataset to VQGAN codes using the pre-trained VQGAN encoder. This may be alleviated by constructing a light-weight encoder using model distillation of the VQGAN encoder.

Currently, inputs in the proposed method were limited to the conditional images such as a single or several NFoV images embedded in ERP, or masked omni-directional images for in-painting. However, the proposed architecture can be applied to any conditional information such as text information, a class label, and guidance information for style modulation using an additional module such as cross attention similar to stable diffusion. In addition, the hyper-parameters have not been thoroughly explored in the evaluation of the proposed method. For example, the directions of NFoV images at the second stage were fixed to 26 directions corresponding to the faces in rhombicuboctahedron, in addition to the field of view fixed to 60 degrees, which may be optimized in the future work. Furthermore, the network structure such as MaxViT might be improved to more optimized architecture to the omni-directional image synthesis. Although this paper focused on the tasks of omni-directional image synthesis, the proposed architecture would be useful for other omni-directional image tasks, such as semantic segmentation and object detection.

## 7 Conclusion

A novel method for the omni-directional image synthesis is proposed in this paper. By using the pre-trained VQGAN encoder and decoder without fine-tuning, the training of the model was drastically shortened. To manage the distortion in an omni-directional image in ERP, a two-stage structure was adopted. At the first stage, an omni-directional image was generated in ERP without geometric distortion correction, so that it cannot reproduce straight lines at poles in a sphere. Therefore, it was corrected at the second stage by synthesizing an omni-directional image from multiple NFoV images based on geometric distortion correction. To realize fast inference, the sampling strategy in MaskGIT was adopted to predict VQGAN codes simultaneously. As a result, the proposed method achieved the high-quality omni-directional image synthesis with low computational costs both in training and inference.

## Acknowledgement

This work was supported by JSPS KAKENHI Grant Number JP21K11943

## References

1. Compvis/taming-transformers. <https://github.com/CompVis/taming-transformers>
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 (CVPR) (2022)

3. Chang, H., Zhang, H., Barber, J., Maschinot, A., Lezama, J., Jiang, L., Yang, M.H., Murphy, K.P., Freeman, W.T., Rubinstein, M., Li, Y., Krishnan, D.: Muse: Text-to-image generation via masked generative transformers. In: Proceedings of the 40th International Conference on Machine Learning (ICML) (2023)
4. Chang, H., Zhang, H., Jiang, L., Liu, C., Freeman, W.T.: Maskgit: Masked generative image transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2022)
5. Chen, Z., Wang, G., Liu, Z.: Text2light: Zero-shot text-driven hdr panorama generation. *ACM Transactions on Graphics (TOG)* **41**(6), 1–16 (2022)
6. Dastjerdi, M.R.K., Hold-Geoffroy, Y., Eisenmann, J., Khodadadeh, S., Lalonde, J.F.: Guided co-modulated gan for 360° field of view extrapolation. In: 2022 International Conference on 3D Vision (3DV). pp. 475–485 (2022)
7. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12873–12883 (2021)
8. Hara, T., Mukuta, Y., Harada, T.: Spherical image generation from a single image by considering scene symmetry. In: Thirty-Fifth AAAI Conference on Artificial Intelligence. pp. 1513–1521 (2021)
9. 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. In: Advances in Neural Information Processing Systems (NeurIPS) (2017)
10. Keisuke, O., Takao, Y.: Omni-directional image generation from single snapshot image. In: IEEE International Conference on Systems, Man, and Cybernetics (SMC) (2020)
11. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: International Conference on Learning Representations (ICLR) (2014)
12. Lu, Z., Hu, K., Wang, C., Bai, L., Wang, Z.: Autoregressive omni-aware outpainting for open-vocabulary 360-degree image generation. In: arXiv preprint arXiv:2309.03467 (2023)
13. Nakata, A., Miyazaki, R., Yamanaka, T.: Increasing diversity of omni-directional images generated from single image using cgan based on mlpmixer. In: Asian Conference on Pattern Recognition (ACPR)
14. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X., Chen, X.: Improved techniques for training gans. In: Advances in Neural Information Processing Systems (NeurIPS) (2016)
15. Suvorov, R., Logacheva, E., Mashikhin, A., Remizova, A., Ashukha, A., Silvestrov, A., Kong, N., Goka, H., Park, K., Lempitsky, V.: Resolution-robust large mask inpainting with fourier convolutions. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (2022)
16. Tu, Z., Talebi, H., Zhang, H., Yang, F., Milanfar, P., Bovik, A., Li, Y.: Maxvit: Multi-axis vision transformer. In: European Conference on Computer Vision (ECCV) (2022)
17. Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. In: Advances in Neural Information Processing Systems (NeurIPS). vol. 30 (2017)
18. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems (NeurIPS). vol. 30 (2017)
19. Xiao, J., Ehinger, K.A., Oliva, A., Torralba, A.: Recognizing scene viewpoint using panoramic place representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2012)

20. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)