

Painting Outside as Inside: Edge Guided Image Outpainting via Bidirectional Rearrangement with Progressive Step Learning

Kyunghun Kim¹, Yeohun Yun¹, Keon-Woo Kang¹, Kyeongbo Kong², Siyeong Lee³, and Suk-Ju Kang¹

¹Department of Electronic Engineering, Sogang University, Seoul, South Korea,,

²Department of Electrical Engineering, POSTECH, Pohang, South Korea,

³NAVER LABS, South Korea

godgang@sogang.ac.kr, dugns159@gmail.com, gunkun00@gmail.com, kkb4723@postech.ac.kr,
siyeong.lee@naverlabs.com, sjkang@sogang.ac.kr

Abstract

Image outpainting is a very intriguing problem as the outside of a given image can be continuously filled by considering as the context of the image. This task has two main challenges. The first is to maintain the spatial consistency in contents of generated regions and the original input. The second is to generate a high-quality large image with a small amount of adjacent information. Conventional image outpainting methods generate inconsistent, blurry, and repeated pixels. To alleviate the difficulty of an outpainting problem, we propose a novel image outpainting method using bidirectional boundary region rearrangement. We rearrange the image to benefit from the image inpainting task by reflecting more directional information. The bidirectional boundary region rearrangement enables the generation of the missing region using bidirectional information similar to that of the image inpainting task, thereby generating the higher quality than the conventional methods using unidirectional information. Moreover, we use the edge map generator that considers images as original input with structural information and hallucinates the edges of unknown regions to generate the image. Our proposed method is compared with other state-of-the-art outpainting and inpainting methods both qualitatively and quantitatively. We further compared and evaluated them using BRISQUE, one of the No-Reference image quality assessment (IQA) metrics, to evaluate the naturalness of the output. The experimental results demonstrate that our method outperforms other methods and generates new images with 360°panoramic characteristics.

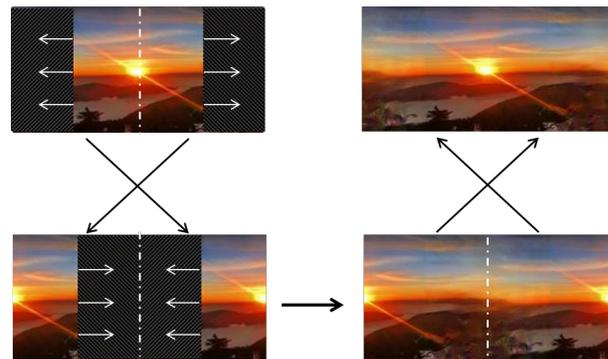


Figure 1. Main concept of the proposed method. We rearrange the outer area of the image into its inner area. This enables us to consider bidirectional information, and it can generate a high-quality natural image that is superior to images produced by conventional methods.

1. Introduction

With developments in deep learning, image completion has been actively researched and it has had the great performance for various applications. Typically, image completion includes image inpainting [22, 25, 42] and outpainting [14, 27, 35]. Image inpainting predicts relatively small missing or corrupted parts of the interiors in a given image based on the surrounding pixels. Image outpainting is used to extend images beyond their borders, thereby generating larger images with added peripheries. The task can be classified into two categories according to the amount of extrapolation desired: slightly extending images can be useful in such applications as panorama construction, when small additions may be needed to make the panorama rectangular after image stitching [37]. Classical image outpainting methods are patch-based methods [18, 31, 44] that fill in

the missing areas by copying similar information from the known areas of the image. These methods [18, 31, 44] perform well when filling small missing areas on a simple texture image, but the results are semantically unnatural and inconsistent with the given context. In recent years, generative adversarial network (GAN)-based approaches [10] have addressed the limitations of classical methods. In particular, the GAN network used in image outpainting that generates a large region is exceedingly difficult to train with stability.

Image inpainting and outpainting methods tend to perform a similar task with regard to filling the missing regions or unknown regions of a given image. Although image outpainting is closely related to image inpainting, which has been making considerable progress recently, image outpainting has attracted less attention because it is more difficult compared to inpainting. First, image inpainting fills the missing region using adjacent and bidirectional information; however, image outpainting fills the missing region using only unidirectional information. As image outpainting can utilize less information than image inpainting, image outpainting methods generate poor-quality output images as compared to those obtained by inpainting. Second, image outpainting generates a larger unknown region than that in image inpainting; hence, maintaining spatial consistency of contents between an input image and generated regions is difficult.

Existing image outpainting methods [35, 41] using GANs still use only one side of adjacent information, resulting in blurry textures. To solve image outpainting problems, how about restoring the image as if it were a panoramic image? We propose a bidirectional boundary region rearrangement method to increase the adjacent information. The overall of the rearrangement method shown in Fig. 1. We rearrange the outer area of the image into its inner area. This enables us to consider bidirectional information, and it can generate a high-quality natural image that is superior to images produced by conventional methods. In addition, we use the structural edge map generator to make the structural information clearer. Moreover, we present a progressive step learning method that divides the learning step into multiple steps to stabilize the GAN network. The method is based on the gains coarse-to-fine learning strategy [1] used in image inpainting. For more efficient outpainting task, we converted the progressive learning method [1] into a method that increases only the horizontal area of the mask to better connect the information at both ends of the horizontal area. We trained the model by step by step increasing the size of the mask per step.

In summary, our contributions are as follows.

1. We propose a novel bidirectional boundary region rearrangement method to alleviate the difficulty of a problem by changing the problem domain from image outpainting to image inpainting. Using this approach, our proposed

method can generate a semantically more natural image better than conventional methods. Not only can we extend the image outside we also can extend the image inside the image.

2. We present a hinge loss specialized in an edge map generator. We use a structural edge map generator optimized for image outpainting. It generates an edge map to make the structural information clearer, thereby generating a photo-realistic image considering the surrounding regional characteristics.

3. We introduce a horizontally progressive step learning method that stabilizes the GAN generator and better connects the information in the horizontal region at both ends. This method is useful for horizontal image outpainting that generates a large unknown region. This can also achieve an augmentation effect for small datasets.

2. Related Work

Image Inpainting This method fills the unknown region in input image, and then make the image photo-realistic by extracting the image features of the image through the information of the known region. Conventional image inpainting methods [3, 7] rely on the similarity or diffusion of a patch to obtain information regarding unknown regions from known regions. These methods are effective when the damaged region is small. However, if the unknown region becomes large, they cannot perform semantic analysis and consequently generate image with low quality. The use of deep learning-based generative adversarial networks (GANs) [10] for image generation led to improved performance of image inpainting methods. Iizuka et al. [14] proposed a method using two discriminators—a global discriminator and a local discriminator—based on GANs. The global discriminator scans the entire image to assess its coherency, whereas the local discriminator scans only a small area centered at the completed region to ensure the local consistency of the generated patches. This allows them to generate a naturally unfragmented image. Liu et al. [19] proposed partial convolution that is masked and renormalized to be conditioned only on valid pixels. Typically, image inpainting methods use a standard convolutional network over the corrupted image, using convolutional filters conditioned on both valid pixels as well as pixels in the unknown regions. This often leads to artifacts such as color discrepancy and blurriness. However, as the partial convolution [19] renormalized convolution operation is conditioned only on valid pixels, it can robustly fill the missing regions of any shape, size, location, or distance from the image borders. Yu et al. [42] proposed a two-stage coarse-to-fine network architecture, in which the first network makes an initial coarse prediction, and the second network takes the coarse prediction as inputs and predicts the refined results. In addition, Yu et al. [42] proposed a contextual atten-

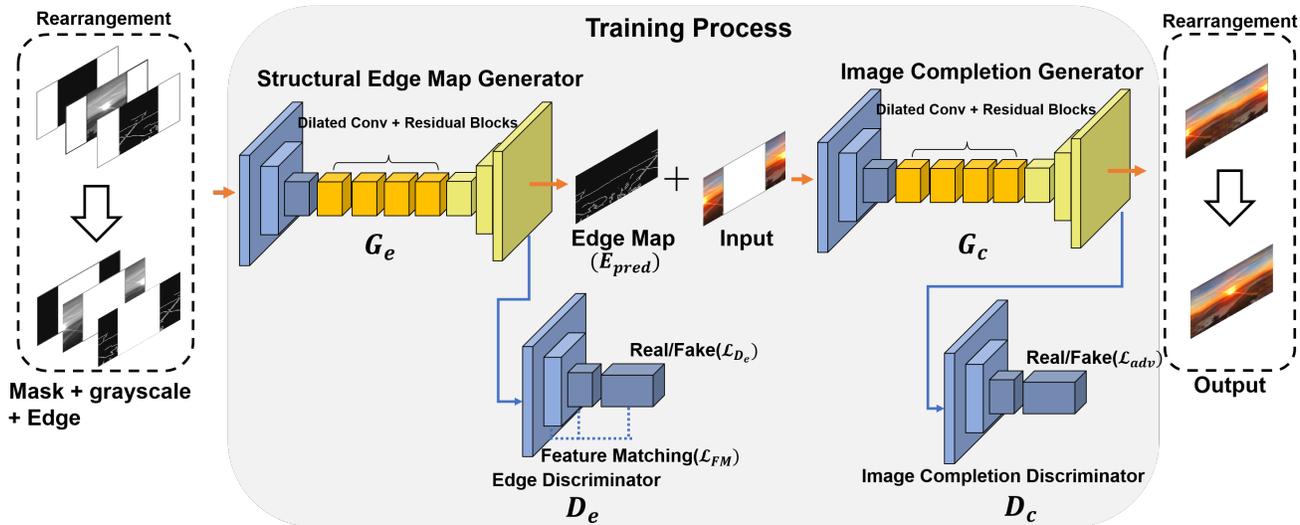


Figure 2. Overall architecture of our image outpainting network. Our model comprises two parts: an edge map generation network and an image completion network. Each part is composed of a generator-discriminator pair.

tion module that copied significant features from the known region to the unknown region that preserved the details of the filled region. Inpainting algorithms tend to have more predictable and higher-quality results compared to those of outpainting algorithms. However, we here in demonstrate that using inpainting algorithms with no modifications leads to poor results for image outpainting.

Image Outpainting Image outpainting is a method that naturally fills the external region of the image to generate a natural image similar to that obtained by image inpainting. However, as mentioned in the previous section, the area to be generated is typically larger than that in image inpainting, and the adjacent information is insufficient. Conventional image outpainting methods [18, 31, 44] select candidates through the similarity of a patch, but their performance becomes less reliable as the size of the generated region increases. Recently, some methods [8, 17] using GANs have been proposed to enhance the output image quality. Sabini et al. [27] proposed a method based on GANs for the first time. This method comprises a simple encoder–decoder structure and uses only the mean squared error (MSE) and GAN loss [10]. It uses postprocessing to smooth the output image but shows repeated pixels in the generated image. Teterwak et al. [33] used semantic conditioning in the discriminator, which is a stabilization scheme for training, based on the semantic information from a pre-trained deep network to regulate the behavior of the discriminator. Yang et al. [41] proposed a recurrent content transfer (RCT) model based on long short-term memory (LSTM) [13]. RCT transfers the input region features to the prediction region, improving a natural connection between the input region and the prediction region. However, the generated region reflects the features of the adjacent input

region relatively largely; thus, the entire generated region tends to not match the color tone, which is an important feature of the input image. To solve this problem, we adapt two losses proposed in [8, 17], known as style loss and perceptual loss. Using these losses, the semantic contents are RAT consistent in styles such as color tones of the input image.

3. Methodology

We provide an overview of the overall architecture, which is shown in Fig. 2, then provide details on each component. Our model comprises two parts: an edge map generation network and an image completion network. Each part is composed of a generator-discriminator pair. We define the generator and discriminator of the structural edge map generator as G_e and D_e , respectively. G_e is used to predict missing structures, thereby generating the global edge map image, E_{pred} , which is used in the image completion network. The generator and discriminator of the image completion network are G_c and D_c , respectively. G_c draws details according to the edge map image and generates the final image, I_{pred} . Our generators follow an architecture similar to that proposed by Johnson et al. [17] and uses instance normalization [34] across all layers of the network.

3.1. Bidirectional Boundary Region Rearrangement

Image outpainting, which only utilizes unidirectional information, is less reliable than image inpainting, which generates missing areas through bidirectional information. Previous methods [27, 41, 33] use only unidirectional information and generate structurally and semantically insuffi-

cient images. To overcome this difficulty in image outpainting, we rearranged the image to benefit from the image inpainting by considering more directional information. We propose a novel bidirectional boundary region rearrangement to alleviate the difficulty of a problem by changing the problem domain from image outpainting to image inpainting. When testing, our model preprocesses an input image through bidirectional boundary region rearrangement and performs the same network process as training. The final image also rearranges the output image through the rearrangement module once again.

3.2. Structural Edge Map Generator

An edge map is usually used to enforce structural quality prior to image inpainting [22] and image super-resolution [23]. Nazeri et al. [22] proposed an edge generator to hallucinate edges in the edges of the missing regions, which can be regarded as an edge completion problem. Using edge images as structural guidance, high inpainting performance is achieved even for some highly structured scenes.

Let I_{gt} be the ground-truth image. E_{gt} and I_{gray} denote the edge map and grayscale image, respectively. In the edge generator, we use the masked grayscale image $\tilde{I}_{gray} = I_{gray} \odot (1 - M)$ and edge map $\tilde{E}_{gt} = E_{gt} \odot (1 - M)$ as the input. The image mask M is defined by the binary image (1 for the missing region and 0 for the background). Here, \odot denotes the Hadamard product. The generator predicts the edge map for the masked region.

$$E_{pred} = G_e(\tilde{I}_{gray}, \tilde{E}_{gt}, M). \quad (1)$$

Our edge discriminator D_e receives E_{gt} and E_{pred} conditioned on I_{gray} as inputs and predicts whether the edge map is real or fake. The edge generator is trained with an objective comprising of the hinge variant of GAN loss [21] and feature-matching loss [38]. The hinge loss is effectively used in binary classification [4]. We believed that the hinge loss we used in this task would be effective because we train the edge generator using binary edge maps.

$$\mathcal{L}_{G_e} = \lambda_{hinge} \mathcal{L}_{hinge} + \lambda_{FM} \mathcal{L}_{FM}, \quad (2)$$

where λ_{hinge} and λ_{FM} are regularization parameters. The generator and discriminator with hinge loss are defined as follows:

$$\mathcal{L}_{hinge} = -\mathbb{E}_{I_{gray}} [D_e(E_{pred}, I_{gray})], \quad (3)$$

$$\begin{aligned} \mathcal{L}_{D_e} = & \mathbb{E}_{(E_{gt}, I_{gray})} [\max(0, 1 - D_e(E_{gt}, I_{gray}))] \\ & + \mathbb{E}_{I_{gray}} [\max(0, 1 + D_e(E_{pred}, I_{gray}))]. \end{aligned} \quad (4)$$

The feature-matching loss, \mathcal{L}_{FM} , compares the activation maps in the intermediate layers of the discriminator. This stabilizes the training process by forcing the discriminator to produce an output that is similar to the real image. This

is similar to perceptual loss [8, 9] in which activation maps are compared with the feature maps of the pre-trained VGG network [30]. The feature-matching loss \mathcal{L}_{FM} is defined as

$$\mathcal{L}_{FM} = \mathbb{E} \left[\sum_i^L \frac{1}{N_i} \|D_e^{(i)}(E_{gt}) - D_e^{(i)}(E_{pred})\|_1 \right], \quad (5)$$

where N_i is the number of elements and D_e is the activation in the i th layer of the discriminator.

3.3. Image Completion Network

After obtaining E_{pred} , G_c generates a complete colored image. The masked color image $\tilde{I}_{gt} = I_{gt} \odot (1 - M)$ and conditional composite edge map $E_{comp} = E_{gt} \odot (1 - M) + E_{pred} \odot M$ were used as inputs.

$$I_{pred} = G_c(\tilde{I}_{gt}, E_{comp}), \quad (6)$$

where I_{pred} denotes the final output result. This network is trained over a joint loss that comprises ℓ_1 loss, adversarial loss, perceptual loss, and style loss. To ensure proper scaling, the ℓ_1 loss is normalized by the mask size. We employ adversarial loss in our generator to generate realistic results.

$$\begin{aligned} \mathcal{L}_{adv} = & \mathbb{E}_{(I_{gt}, E_{comp})} [\log D_c(I_{gt}, E_{comp})] \\ & + \mathbb{E}_{E_{comp}} [\log [1 - D_c(I_{pred}, E_{comp})]]. \end{aligned} \quad (7)$$

Perceptual loss is defined as follows:

$$\mathcal{L}_{perc} = \mathbb{E} \left[\sum_i \frac{1}{N_i} \|\varphi_i(I_{gt}) - \varphi_i(I_{pred})\|_1 \right]. \quad (8)$$

\mathcal{L}_{perc} penalizes results that are not perceptually similar to features by defining a distance measure between activation maps of a pre-trained network. The style loss is defined as follows:

$$\mathcal{L}_{style} = \mathbb{E}_j \left[\|G_j^\varphi(\tilde{I}_{pred}) - G_j^\varphi(\tilde{I}_{gt})\|_1 \right]. \quad (9)$$

We choose to use style loss by Sajjadi et al. [28] to be an effective tool to combat ‘‘checkerboard’’ artifacts caused by the transpose convolution layers [24]. The final total loss is defined as

$$\mathcal{L}_{G_c} = \lambda_{\ell_1} \mathcal{L}_{\ell_1} + \lambda_{adv} \mathcal{L}_{adv} + \lambda_p \mathcal{L}_{perc} + \lambda_{style} \mathcal{L}_{style}. \quad (10)$$

To address color tone mismatch problems in the previous outpainting methods, we use a large proportion for style loss. For our experiments, we set λ_{ℓ_1} as 1, λ_{adv} as 0.2, λ_p as 0.1, and λ_{style} as 250.

| Method | IS | FID | PSNR | SSIM | BRISQUE |
|--------------------|-------------|--------------|--------------|-------------|--------------|
| Pix2Pix [15] | 2.82 | 19.73 | - | - | - |
| GLC[14] | 2.81 | 14.82 | - | - | - |
| CA[42] | 2.93 | 19.04 | 20.42 | 0.84 | 24.46 |
| StructureFlow[26] | 3.10 | 15.69 | 22.94 | 0.85 | 26.36 |
| NS-OUT[41] | 2.85 | 13.71 | 19.53 | 0.72 | 23.59 |
| Proposed(w/o BR) | 3.07 | 17.75 | 21.41 | 0.84 | 23.62 |
| Proposed (with BR) | 3.20 | 15.72 | 22.45 | 0.86 | 21.61 |

Table 1. Quantitative results for conventional and proposed models on SUN dataset [40]. Evaluation of Inception Score (IS) [29], PSNR and SSIM [39] (the higher, the better), Fréchet Inception Distance FID [12] and BRISQUE [20] (the lower, the better). Images from the validation set had an IS of 3.479 (The best result of each column is red, the second-best result is blue, and BR means Bidirectional Rearrangement).

3.4. Horizontally progressive step Learning Method

GANs are difficult to train because of problems such as mode collapse, non-convergence, and vanishing gradient [2]. These problems also apply to the image outpainting task where more areas than expected are generated without context. We propose a simple but effective training technique to stabilize GAN training from the perspective of image outpainting by horizontally increasing the mask size (Fig. 3.). Generators are usually successful in a small mask task but have difficulty in a wide mask task [1]. Therefore, we divide the learning step by mask size so that the model can learn more stably while increasing the mask size. The mask is divided into 32 steps and gradually grows at each step. The initial mask size is 3.125% of the input image and linearly increases, and the final mask size is set at 50% of the input image.

4. Training

When we training, we give the mask inside the images to proceed in the horizontally progressive step learning method. The training process corresponds to the gray box of Fig. 2. When testing, the masked outside of the images is given as input and generates the outpainting images using the rearrangement method described above.

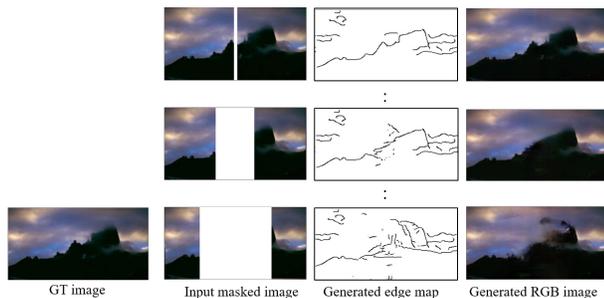


Figure 3. The process of image completion based on the edge map progressively generated by the edge map generator.

4.1. Training Setup

Our model is implemented using the PyTorch framework. The network is trained with the SUN dataset [40] comprising 256×128 pixels. Considering our GPU memory, the batch sizes of G_e and G_c networks are 8 and 16, respectively. We use the AdamP optimizer [11] with $\beta_1 = 0$ and $\beta_2 = 0.99$. Generators G_e and G_c , are trained with the learning rate of 10^{-4} until the losses plateau separately. We lower the learning rate to 10^{-5} and continue to train G_e and G_c until convergence.

We use a Canny edge detector [5] to generate an edge map. The sensitivity of the Canny edge detector is controlled by the standard deviation of the Gaussian smoothing filter (γ). In the experimental results, it was the best when the γ is 2. Fig. 3 illustrates the process of image completion based on the edge map generated by the edge map generator in each step. The result of the image completion generator depends on the edge map.

5. Experiment

5.1. Quantitative Result

Evaluating the quality of generated images in image outpainting has the same difficulty as that in GAN, as there are few restrictions on the created images. Note that the well-outpainted image has only to be photo-realistic while sharing the context naturally with the input image. Hence, we include non reference image quality metrics to evaluate the generate images. We used structural similarity index (SSIM) [39], peak signal-to-noise ratio (PSNR), Inception Score (IS) [29], Fréchet Inception Distance (FID) [12], and Blind/Referenceless Image Spatial Quality Evaluator (BRISQUE) [20] to measure the resulting quality with the state-of-the-art outpainting and inpainting algorithm. FID uses a pre-trained Inception-V3 model to measure the Wasserstein-2 distance between the actual image and the representation of the shape space of the painted image [32]. We evaluated the task of generating 50% of the input image on the SUN dataset [40] with other methods. The results

| Method | SSIM | PSNR |
|-----------------------|--------------|---------------|
| Image-Outpainting[27] | 0.338 | 14.625 |
| Outpainting-srn[36] | 0.513 | 18.221 |
| SieNet[43] | 0.646 | 20.796 |
| Proposed | 0.810 | 18.957 |

Table 2. Performance on beach dataset [27] (The best result of each column is red and the second-best result is blue.).

for the SUN dataset are shown in Table. 1. Our model produced more natural images than other methods, as can be seen in the qualitative section. We further evaluated using the BRISQUE method [20] to evaluate how the naturalness of the image produced. A quantitative comparison shows that our model is superior to the other methods in all aspects except for the PSNR and the FID score. We believe that the combination of the listed metrics provide a better result in outpainting performance.

We additionally experimented with quantitative comparisons on the beach dataset [27] for a more objective comparison of our model with other conventional outpainting methods. We tested our model using the pre-trained weights with the sun dataset [40] without additional training on the beach dataset. We used the quantitative results from the SieNet [43] and compared them. The results for the beach dataset [27] are listed in Table. 2, and it can be seen that our model has better structural strength compared to other models.

5.2. Qualitative Result

Fig. 6 shows a qualitative result comparison image between the proposed method and the state-of-the-art methods NS-OUT [41], and the inpainting methods CA [42] and StructureFlow [26]. CA [42] produced images that show inconsistent objects near the missing image parts and show large color differences from the original image. The outpainting algorithm NS-OUT [41] and the inpainting algorithm StructureFlow [26] produced images that are more clearer, but image distortion and blurring in the generated image still exist. Our proposed method produced images that are clearly connected with the original image and are smoother and more consistent than those obtained by NS-OUT and StructureFlow.

5.3. Ablation

We conducted ablation studies to demonstrate the necessity of introducing bidirectional boundary rearrangement. A comparison of the quantitative results of our architecture and the model without bidirectional rearrangement (w/o BR) is shown in Table. 1. According to the results, the model with bidirectional rearrangement method (with BR) effectively improved the performance at all values.

As shown in Fig. 4, hinge loss is more effective than

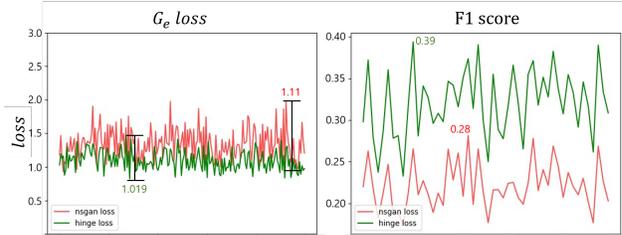


Figure 4. Loss value and F1 score of the edge map generator in the last training step. In the left image, the peak-to-peak amplitude of the hinge loss is 1.019 and the nsgan loss is 1.11. It can be seen that the hinge loss converges to a smaller amplitude. In the right image, the highest F1 score of hinge loss is 0.39 and nsgan loss is 0.28.

Non-Saturating GAN loss (nsgan loss) [10]. In the case of using the hinge loss function, the fluctuation of the loss vibrates to a smaller value during the learning process, and the learning process is more stable. we observed a higher F1 score than nsgan loss. The edge guided method useful not only for image outpainting, but also for image inpainting [22], super-resolution [23], and various other fields [6, 16]. Therefore, it is believed that the hinge loss in other fields as well as binary data such as edge maps will help improve the performance.

5.4. User Study

We performed a user study to compare the image outpainting performance of our proposed method with other image inpainting/outpainting methods, such as CA [42], StructureFlow [26] and NS-OUT [41], using benchmark datasets. A total of 27 experts in image processing participated in anonymous voting to evaluate naturalness from randomly selected 30 images to ranks 1 to 4. Fig. 5 shows the summarized results. Each rank has a total of 810 votes, where our proposed method obtained 610 votes and 190 votes in rank 1 and rank 2, respectively. Consequently, the comparison result verified that our proposed method out-

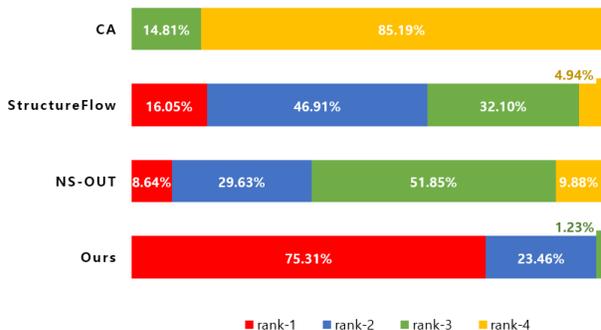


Figure 5. User study results between our proposed method and conventional methods (CA [42], StructureFlow [26], NS-OUT [41]). Our method obtained the most superior results.

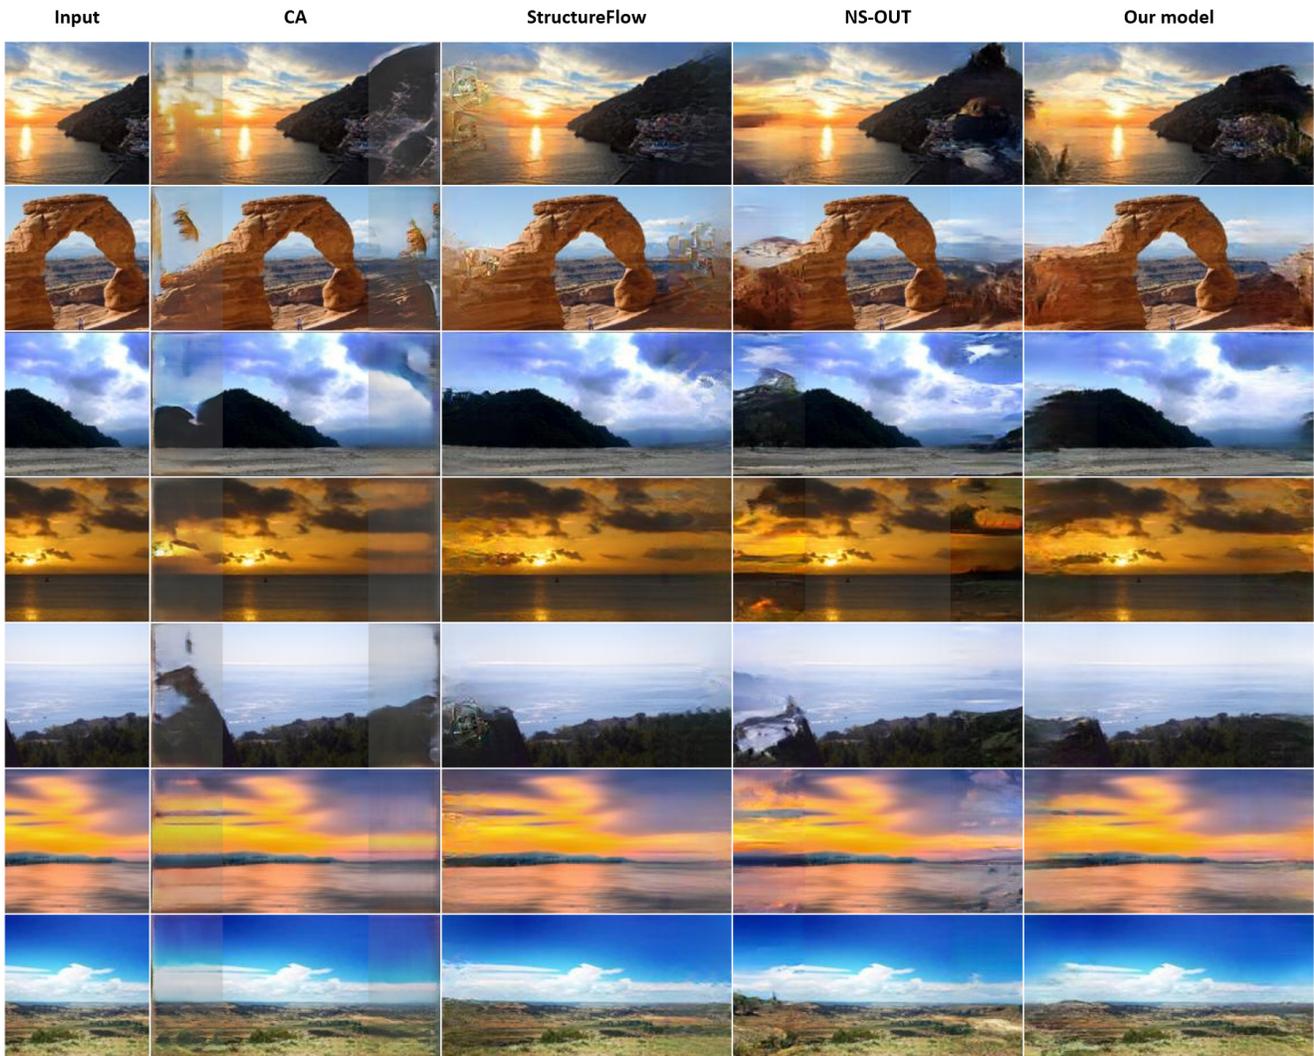


Figure 6. Qualitative results for conventional and proposed models on SUN dataset [40]: CA[42], StructureFlow[26], NS-OUT [41], and our model.

performs other methods in generating a more visually clear image to human raters.

6. Limitation

The proposed method produced a natural image well when the information at both ends of the image is similar. However, when the assumption, where information on both ends of an image is similar, is broken, our method generated an output image with different characteristics with the ground truth, but produced natural images. As can be seen from the quantitative and qualitative results above, at the sun and beach datasets showed overall good results. We expect results to be poor for datasets that are more complex and have little similar information at both ends. We will try to consider these factors in the future work.

7. Conclusion and Future Work

In this paper, we proposed a novel image outpainting method composed of three approaches. First, we rearranged the bidirectional boundary regions to address the lack of information when filling the image outward. The previous methods had difficulties in generating images because the adjacent information was not sufficient to generate large empty areas. However, the edges from our structural edge map generator worked as a guideline to reflect semantic information from the given regions to the unknown ones. In addition, with the rearrangement of the boundary regions, increased adjacent information prevented the large unknown areas from being filled with repetitive pixel values. Second, we present a hinge loss specialized in generating binary images such as edge maps. Third, the train-

ing step was divided according to different horizontal mask sizes so that the model was trained stably as well as naturally generated images outside and inside. Through multiple steps, we can obtain a wider and better-quality image. We evaluated our model compared to conventional image inpainting and outpainting methods in terms of qualitative and quantitative measurements. The experimental results show that our method outperformed the other methods.

In future work, we will explore how to generate outpainting images on horizontal and vertical directions with the same model simultaneously. Besides, we plan to design a novel model whether to selectively refer to different information on both ends so that a natural image can be robustly generated even in the case of images with different information on both ends.

8. Acknowledgments

This research was supported by the MSIT(Ministry of Science and ICT), Korea, under the ITRC(Information Technology Research Center) support program(IITP-2020-2018-0-01421) supervised by the IITP(Institute for Information & Communications Technology Planning & Evaluation), the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT)(No. 2020M3H4A1A02084899) and the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2018R1D1A1B07048421)

References

- [1] Abbas Hedjazi, Mohamed and Genc, Yakup. Learning to inpaint by progressively growing the mask regions. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [2] Arjovsky, Martin and Chintala, Soumith and Bottou, Léon. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- [3] Barnes, Connelly and Shechtman, Eli and Finkelstein, Adam and Goldman, Dan B. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3):24, 2009.
- [4] Peter L Bartlett and Marten H Wegkamp. Classification with a reject option using a hinge loss. *Journal of Machine Learning Research*, 9(Aug):1823–1840, 2008.
- [5] Canny, John. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence*, (6):679–698, 1986.
- [6] Yaqiong Chai, Botian Xu, Kangning Zhang, Natasha Lepore, and John C Wood. Mri restoration using edge-guided adversarial learning. *IEEE Access*, 8:83858–83870, 2020.
- [7] Efros, Alexei A and Freeman, William T. Image quilting for texture synthesis and transfer. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 341–346, 2001.
- [8] Gatys, Leon A and Ecker, Alexander S and Bethge, Matthias. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016.
- [9] Gatys, Leon and Ecker, Alexander S and Bethge, Matthias. Texture synthesis using convolutional neural networks. In *Advances in neural information processing systems*, pages 262–270, 2015.
- [10] Goodfellow, Ian and Pouget-Abadie, Jean and Mirza, Mehdi and Xu, Bing and Warde-Farley, David and Ozair, Sherjil and Courville, Aaron and Bengio, Yoshua. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [11] Byeongho Heo, Sanghyuk Chun, Seong Joon Oh, Dongyoon Han, Sangdoon Yun, Youngjung Uh, and Jung-Woo Ha. Slowing down the weight norm increase in momentum-based optimizers, 2020.
- [12] Heusel, Martin and Ramsauer, Hubert and Unterthiner, Thomas and Nessler, Bernhard and Hochreiter, Sepp. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, pages 6626–6637, 2017.
- [13] Hochreiter, Sepp and Schmidhuber, Jürgen. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [14] Iizuka, Satoshi and Simo-Serra, Edgar and Ishikawa, Hiroshi. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)*, 36(4):1–14, 2017.
- [15] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [16] Kui Jiang, Zhongyuan Wang, Peng Yi, Guangcheng Wang, Tao Lu, and Junjun Jiang. Edge-enhanced gan for remote sensing image superresolution. *IEEE Transactions on Geoscience and Remote Sensing*, 57(8):5799–5812, 2019.
- [17] Johnson, Justin and Alahi, Alexandre and Fei-Fei, Li. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [18] Kopf, Johannes and Kienzle, Wolf and Drucker, Steven and Kang, Sing Bing. Quality prediction for image completion. *ACM Transactions on Graphics (TOG)*, 31(6):1–8, 2012.
- [19] Liu, Guilin and Shih, Kevin J and Wang, Ting-Chun and Reda, Fitsum A and Sapra, Karan and Yu, Zhiding and Tao, Andrew and Catanzaro, Bryan. Partial convolution based padding. *arXiv preprint arXiv:1811.11718*, 2018.
- [20] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on image processing*, 21(12):4695–4708, 2012.
- [21] Miyato, Takeru and Kataoka, Toshiki and Koyama, Masanori and Yoshida, Yuichi. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- [22] Nazeri, Kamyar and Ng, Eric and Joseph, Tony and Qureshi, Faisal Z and Ebrahimi, Mehran. Edgeconnect: Genera-

- tive image inpainting with adversarial edge learning. *arXiv preprint arXiv:1901.00212*, 2019.
- [23] Nazeri, Kamyar and Thasarathan, Harrish and Ebrahimi, Mehran. Edge-informed single image super-resolution. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [24] Odena, Augustus and Dumoulin, Vincent and Olah, Chris. Deconvolution and checkerboard artifacts. *Distill*, 1(10):e3, 2016.
- [25] Pathak, Deepak and Krahenbuhl, Philipp and Donahue, Jeff and Darrell, Trevor and Efros, Alexei A. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016.
- [26] Yurui Ren, Xiaoming Yu, Ruonan Zhang, Thomas H Li, Shan Liu, and Ge Li. Structureflow: Image inpainting via structure-aware appearance flow. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 181–190, 2019.
- [27] Mark Sabini and Gili Rusak. Painting outside the box: Image outpainting with gans. *arXiv preprint arXiv:1808.08483*, 2018.
- [28] Sajjadi, Mehdi SM and Scholkopf, Bernhard and Hirsch, Michael. Enhancenet: Single image super-resolution through automated texture synthesis. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4491–4500, 2017.
- [29] Salimans, Tim and Goodfellow, Ian and Zaremba, Wojciech and Cheung, Vicki and Radford, Alec and Chen, Xi. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016.
- [30] Simonyan, Karen and Zisserman, Andrew. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [31] Sivic, Josef and Kaneva, Biliana and Torralba, Antonio and Avidan, Shai and Freeman, William T. Creating and exploring a large photorealistic virtual space. In *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 1–8. IEEE, 2008.
- [32] Szegedy, Christian and Vanhoucke, Vincent and Ioffe, Sergey and Shlens, Jon and Wojna, Zbigniew. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [33] Teterwak, Piotr and Sarna, Aaron and Krishnan, Dilip and Maschinot, Aaron and Belanger, David and Liu, Ce and Freeman, William T. Boundless: Generative adversarial networks for image extension. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10521–10530, 2019.
- [34] Ulyanov, Dmitry and Vedaldi, Andrea and Lempitsky, Victor. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6924–6932, 2017.
- [35] Van Hoorick, Basile. Image outpainting and harmonization using generative adversarial networks. *arXiv preprint arXiv:1912.10960*, 2019.
- [36] Yi Wang, Xin Tao, Xiaoyong Shen, and Jiaya Jia. Wide-context semantic image extrapolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1399–1408, 2019.
- [37] Wang, Miao and Lai, Yukun and Liang, Yuan and Martin, Ralph Robert and Hu, Shi-Min. Biggerpicture: data-driven image extrapolation using graph matching. *ACM Transactions on Graphics*, 33(6), 2014.
- [38] Wang, Ting-Chun and Liu, Ming-Yu and Zhu, Jun-Yan and Tao, Andrew and Kautz, Jan and Catanzaro, Bryan. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018.
- [39] Wang, Zhou and Bovik, Alan C and Sheikh, Hamid R and Simoncelli, Eero P. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [40] Xiao, Jianxiong and Hays, James and Ehinger, Krista A and Oliva, Aude and Torralba, Antonio. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3485–3492. IEEE, 2010.
- [41] Yang, Zongxin and Dong, Jian and Liu, Ping and Yang, Yi and Yan, Shuicheng. Very long natural scenery image prediction by outpainting. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10561–10570, 2019.
- [42] Yu, Jiahui and Lin, Zhe and Yang, Jimei and Shen, Xiaohui and Lu, Xin and Huang, Thomas S. Generative image inpainting with contextual attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5505–5514, 2018.
- [43] Xiaofeng Zhang, Feng Chen, Cailing Wang, Ming Tao, and Guo-Ping Jiang. Sienet: Siamese expansion network for image extrapolation. *IEEE Signal Processing Letters*, 2020.
- [44] Zhang, Yinda and Xiao, Jianxiong and Hays, James and Tan, Ping. Framebreak: Dramatic image extrapolation by guided shift-maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1171–1178, 2013.