

MC-GAN: Multi-conditional Generative Adversarial Network for Image Synthesis

Hyojin Park¹
wolfrun@snu.ac.kr

Youngjoon Yoo²
youngjoon.yoo@navercorp.com

Nojun Kwak¹
nojunk@snu.ac.kr

¹ Department of Transdisciplinary Studies,
Seoul National University,
Republic of Korea.

² CLOVA AI Research, Naver Corp.,
Republic of Korea.

Abstract

In this paper, we introduce a new method for generating an object image from text attributes on a desired location, when the base image is given. One step further to the existing studies on text-to-image generation mainly focusing on the object's appearance, the proposed method aims to generate an object image preserving the given background information, which is the first attempt in this field. To tackle the problem, we propose a multi-conditional GAN (MC-GAN) which controls both the object and background information jointly. As a core component of MC-GAN, we propose a synthesis block which disentangles the object and background information in the training stage. This block enables MC-GAN to generate a realistic object image with the desired background by controlling the amount of the background information from the given base image using the foreground information from the text attributes. From the experiments with Caltech-200 bird and Oxford-102 flower datasets, we show that our model is able to generate photo-realistic images with a resolution of 128×128 .

1 Introduction

Recent studies on generative adversarial networks (GAN) [1, 2, 3, 4, 5] have achieved impressive success in generating images. However, since most of the generated images through GAN do not exactly satisfy the users' expectation, people use auxiliary information in various forms such as base images and texts as main cues for controlling the generated images. In this line of research, it is actively studied to create images under specific conditions such as transferring the style of an image [6, 7, 8, 9, 10] or generating an image based on text description [11, 12, 13, 14, 15].

Among these, text-to-image generation is meaningful in that it can fine-tune the generated image through the guide of text description. Although GAN can be applied to diverse text to image generation works, most of the applications are focused on controlling the shape and texture of the foreground and relatively less attention has been paid to the background. For example, Reed et al. [16] created images from the text containing information on the appearance of the object to generate. The method can generate a target image in a given location or with a specified pose. However, they also cannot control the background. There are some

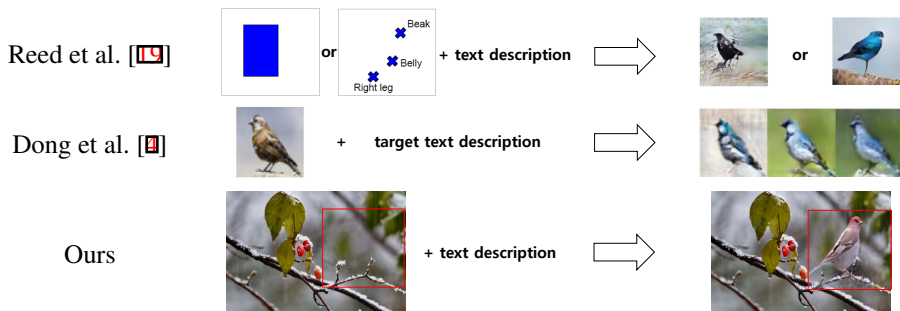


Figure 1: Comparison of different multi-modal conditional GAN problems. We try to synthesize an image based on a base background image with a specified location for an object and a text description on the object. The images on the right are from the respective algorithms.

works [9, 19] that have considered the background. Dong et al. [9] considered multi-modal conditions based both on image and text and can change the base image according to the text description. However, the method has a restriction that a similar object to the generated one should be in the base image. Thus, it can be considered as a style transfer problem. Ma et al. [19] also solved the multi-modal style transfer problem from a reference person image to a target pose. They kept the background and changed the reference person’s pose to a target pose.

In this paper, we define a novel problem of conditional GAN which generates a new image by synthesizing the background of an original base image and a new object described by the text description in a specific location. Different from the existing works [9, 19], we aim to draw a target object on a base image that does not contain similar objects. To the best of our knowledge, our research is the first attempt to synthesize a target image by combining the background of an original image and a text-described foreground object. As shown in Fig. 1, our approach is different from other studies that try to create a random image at a desired location [19] or to change the foreground style [9] in that we want to independently apply separate foreground and background conditions for image synthesis. This problem is not trivial because the generated foreground object and the background from the base image should be smoothly mixed with a plausible pose and a layout.

To tackle this problem of image synthesis, we introduce a new architecture of multi-conditional GAN (MC-GAN) using a synthesis block which acts like a pixel-wise gating function controlling the amount of information from the base background image by the help of the the text description for a foreground object. With the help of this synthesis block, MC-GAN can generate a natural image of an object described by the text in a specified location with the desired background. To show the effectiveness of our method, we trained MC-GAN using the Caltech-200 bird dataset [20] and the Oxford-102 flower dataset [13, 14] and compared the performances with those of a baseline model [9].

Our main contributions can be summarized as follows: (1) We define a novel multi-modal conditional synthesis problem using a base image, a text sentence and location information. (2) To handle complex multi-modal conditions in GAN, we suggest a new architecture of MC-GAN using synthesis blocks. (3) The proposed architecture was shown to generate a plausible natural scene as shown in Fig. 1 and 2 by training publicly available data regardless of whether the base image contains a similar object to create or not.




Figure 2: The examples of synthesized images using bird and flower datasets

2 Related Work

Among the diverse variants of GAN [3, 8, 17, 21, 24], we can mainly categorize the studies into three large groups: 1) the style transfer problem 2) the text-to-image problem 3) the multi-modal conditional problem.

Style Transfer: The style transfer problem uses an image as an input and converts the foreground to a different style. Choi et al. [8], Zhu et al. [24] and Kim et al. [8] transferred images to a different domain style, *e.g.* from a smile face to an angry one, or from a handbag to shoes. In addition to these applications, some works created a real image from a map of segmentation label [11, 21], or used a map of part location in combination with an original image of a person to generate images of a person with different poses [10].

Text-to-Image: The text-to-image problem uses text description as an input to generate an image. It has great advantages over other methods in that it can easily generate an image with the attributes that a user really wants, because text can express detailed high-level information on the appearance of an object with detailed attributions. The raw text is usually embedded according to the method in [16] which uses a hybrid of CNN (convolutional neural network) and RNN (recurrent neural network) structure. Reed et al. [17] proposed a novel text-to-image generation model, and Zhang et al. [22, 23] improved the image quality later by stacking multiple GANs.

Multi-modal Conditional Image Generation: A multi-modal conditional problem is to create images satisfying multiple input conditions in different modalities such as a pair of (image, location) or (image, text). Reed et al. [19] provided a desired object position by a bounding box or a set of object part locations by points in an empty image in combination with the text description to generate an object image (see Fig. 1). Dong et al. [9] used both an image and a text as inputs to GAN for image generation (see also Fig. 1). They intended to keep the image part irrelevant to the text and to change the style of the object contained in the base image based on the text description. Although [19] is similar to our work in terms of using the location information, our method generates an object with an appropriate pose by understanding the semantic information of the background image automatically. Compared to our method, the method of using parts' locations in [19], which requires a user to select parts' locations, is somewhat inconvenient and time-consuming. However, the bounding box condition in [19] is similar to our problem, thus it can be said that our study partially includes the problem defined in [19]. The method in [9] also uses both image and text conditions together. However, our study does not have a restriction that the same kind of object to generate must be in the base image. In other words, ours does not change the style of an already existing object but synthesize a new object with a slightly but properly changed background. (See the last row of Fig. 1 and Fig. 2.)




Figure 3: Overall structure of MC-GAN. The generator uses text description as well as the base image to create an image of the object as well as a segmentation mask. Using the image and the segmentation mask, the discriminator distinguishes whether the input is real or not.

3 Methodology

Fig. 3 represents the overall structure of the proposed MC-GAN. The generator of MC-GAN firstly encodes the input text sentence t into the text embedding $\phi(t)$ using the method in [16]. As in [20] and [23], $\phi(t)$ is concatenated with a noise vector z to which fully connected (FC) layers are applied to constitute a seed feature map. After then, we use a series of synthesis blocks whose inputs are the seed feature map, which in combination with the image features from the background image generates an output image and a segmentation mask. The synthesis blocks are used to prevent overlapping between the generated object and the background. In Section 3.1, we describe the characteristics of the proposed synthesis block in more detail and we introduce the detailed explanation of the model structure and the training strategy in Section 3.2,

3.1 Synthesis Block

Fig. 4 (a) describes the framework of the proposed synthesis block. In the synthesis block, the background (BG) feature is extracted from the given image without non-linear function (*i.e.* only using convolution and batch normalization (BN)) [21] and the foreground (FG) feature is the feature map from the previous layer. As shown in the figure, the BG feature is controlled by multiplying it with an activated switch feature map, a fraction of the FG feature map. The sizes of the BG and FG feature maps are the same, and the depth of FG feature map is doubled as it passes through the convolution layers. A half of the doubled feature map denoted as switch in the figure is used as an input to the sigmoid function, while the other half is forwarded to generate a larger FG feature map for the next synthesis block. The switch determines what amount of BG information should be retained for the next synthesis block. After switch feature map is activated by the sigmoid function, it is multiplied element-wise with the BG feature map for the suppression of background information where the object is to be generated. Finally, the spatial dimension of feature map is doubled by the upsampling layer after element-wise addition of the suppressed BG feature map with the FG feature map. Because the MC-GAN has ℓ_1 loss comparing the background of the created image with that of the base image, the switch has an effect of suppressing the base image in the object area and mimicking the base image in the background. Therefore, a visualized switch map has

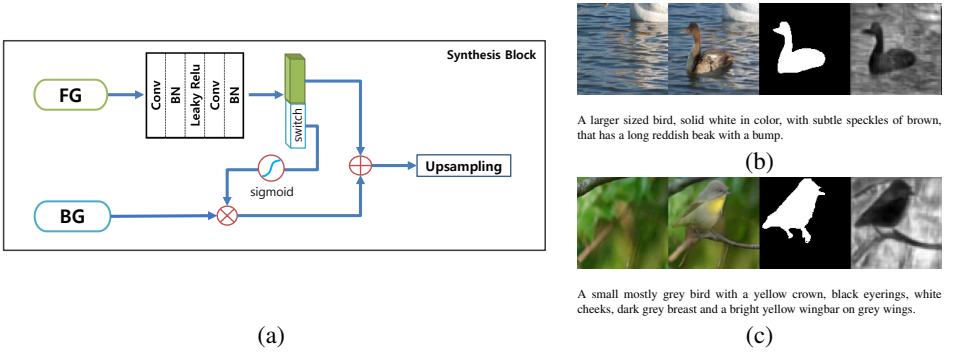


Figure 4: The structure of synthesis block (a) and the examples of generated image (b) (c)

the opposite concept to that of the segmentation mask.

Fig 4 (b) and (c) show a couple of output examples. From left to right they are 1) a cropped background image from a specific location, 2) the generated image, 3) the generated mask and 4) the switch feature map from the final synthesis block. A close look at the switch feature map in Fig 4 (b) shows that it does not change the original background much because the object naturally goes with the background. Thus, the background region is highly activated in the switch while the object region is suppressed (see the last column). On the other hand, in in Fig 4 (c), because a picture of bird standing on the air is unnatural, the generator adds a branch in the figure. At this time, since the original background should not overlap with the newly generated branch, the background information of the new branch area deactivates the corresponding area of the switch map.

3.2 Network Design and Loss Function

MC-GAN encodes the input text sentence t using the method in [16]. However, the vector $\varphi(t)$ generated by this method lies on a high dimensional manifold, while the number of available data is relatively small. Zhang et al. [22, 23] pointed out this problem and proposed the *conditioning augmentation* method and used fully connected layers to make an initial seed feature map from the text embedding $\varphi(t)$ and a noise vector z . Here, we follow this method of creating the initial seed feature map as in [22, 23]. A cropped region from the base image as well as the text sentence is inputted to MC-GAN. The spatial size of the input image is $W \times H$ and the number of synthesis blocks is N . The size of the seed feature map after fully connected layer is $1024 \times \frac{W}{2^N} \times \frac{H}{2^N}$ and the resolution is doubled with each pass through the upsampling layer. The upsampling layer uses the nearest neighborhood method to double the resolution and a 3×3 convolution is applied with BN and ReLU activation to improve the quality of the image. The number of channels is halved for each block of upstream. Conversely, the method of creating an image feature (downstream) does not use any non-linear function, and each step consists of a 3×3 convolution layer with BN. At each step of downstream, the spatial resolution is halved using stride 2 and the number of channels is doubled. By combining the upstream and downstream, the final $64 \times W \times H$ feature map is obtained which is converted into 4 channels (3 for RGB and 1 for segmentation masks).

The discriminator takes a tuple of image-mask-text (x - s - t) as an input. The convolution followed by BN and Leaky ReLU downsamples the image and the mask into an image code \tilde{x} and an image-mask code $[\hat{x}, \hat{s}]$ separately both of which have a resolution of $\frac{W}{2^N} \times \frac{H}{2^N}$. The image-mask code is concatenated with the replicated text code $\varphi(t)$, which is obtained by

the *conditioning augmentation technique* [22, 23] using the text embedding $\phi(t)$. We apply a convolution layer to the associated image-mask-text code, then perform BN and Leaky ReLU activation to reduce the dimension. The image code \tilde{x} , image-mask code $[\hat{x}, \hat{s}]$, and image-mask-text code $[\tilde{x}, \tilde{s}, \tilde{\phi}(t)]$ are trained by the method proposed in Mao et al. [22]. In our case, the discriminator learns the following four types of input tuples.

- | | |
|---|--|
| 1) image1(x_1), mask1(s_1), text1(t_1) | real image with matching mask and text |
| 2) image1(x_1), mask1(s_1), text2(t_2) | real image with matching mask but mismatching text |
| 3) image1(x_1), mask2(s_2), text1(t_1) | real image with mismatching mask but matching text |
| 4) fake image(x_g), fake mask(s_g), text(t) | generated image and mask with input text |

Here, the subscript indicates whether the tuple matches or not. (e.g. the tuple (x_1, s_2, t_1) means that the image matches with the text but the segmentation mask is mismatched.) Using the four types of tuples, the discriminator loss function for the output D_3 becomes

$$L_{D_3} = \mathbb{E}_{x_1, s_1, t_1 \sim p_d} [(D_3(x_1, s_1, t_1) - 1)^2] + \mathbb{E}_{x_1, s_1, t_2 \sim p_d} [D_3(x_1, s_1, t_2)^2] \\ + \mathbb{E}_{x_1, s_2, t_1 \sim p_d} [D_3(x_1, s_2, t_1)^2] + \mathbb{E}_{x_g, s_g \sim p_{G(b, t, z)}, t \sim p_d} [D_3(x_g, s_g, t)^2]. \quad (1)$$

Here, p_d and p_G denote the distributions of real and generated data respectively. $G(b, t, z)$ means the output of the generator using the base image b , text t and noise z . The first term enforces the discriminator to output 1 for the true input (type 1), the second term tries to distinguish mismatching texts (type 2), the third term for distinguishing false masks (type 3), and finally the last term is to distinguish the fake image and mask from the real one. Likewise, the loss functions for D_1 and D_2 become

$$L_{D_2} = \mathbb{E}_{x_1, s_1 \sim p_d} [(D_2(x_1, s_1) - 1)^2] + \mathbb{E}_{x_1, s_2 \sim p_d} [D_2(x_1, s_2)^2] + \mathbb{E}_{x_g, s_g \sim p_G} [D_2(x_g, s_g)^2] \\ L_{D_1} = \mathbb{E}_{x_1 \sim p_d} [(D_1(x_1) - 1)^2] + \mathbb{E}_{x_g \sim p_G} [D_1(x_g)^2] \quad (2)$$

In the training of the generator, to the general loss term of GAN, regularization terms for the *conditioning augmentation* and the background reconstruction are added as follows:

$$L_G = \mathbb{E}_{x, s \sim p_{G(b, t, z)}, t \sim p_d} [(D_1(x) - 1)^2 + (D_2(x, s) - 1)^2 + (D_3(x, s, t) - 1)^2] \\ + \lambda_1 D_{KL}(\mathcal{N}(\mu(\phi(t)), \Sigma(\phi(t))) || \mathcal{N}(0, I)) + \lambda_2 ||(x \odot (f \circ s)) - (b \odot (f \circ s))||_1. \quad (3)$$

Here, $\mu(\phi(t))$ and $\Sigma(\phi(t))$ are the mean and the diagonal covariance matrix from the text embedding by *conditioning augmentation* and the KL divergence loss term is used as in [22, 23]. The last term, the background reconstruction loss, affects the feature extraction of the base image and suggests the activation of the switch determining which part should be taken for synthesis. The operator $f \circ$ denotes the morphological erode operation to the mask s for smoothing and \odot is element-wise multiplication. The areas in the fake image x and the base image b excluding the object part are taken and trained by the ℓ_1 loss.

4 Experiment

4.1 Dataset and training details

We validated the proposed algorithm using the publicly available Caltech-200 bird [24] and Oxford-102 flower [25] datasets. For comparison, in addition to several ablation methods, we tested recent Dong et al. [4]’s work which uses image and text based multi-modal conditions. Reed et al. [19] also used multi-modal conditions for generating images but they did not

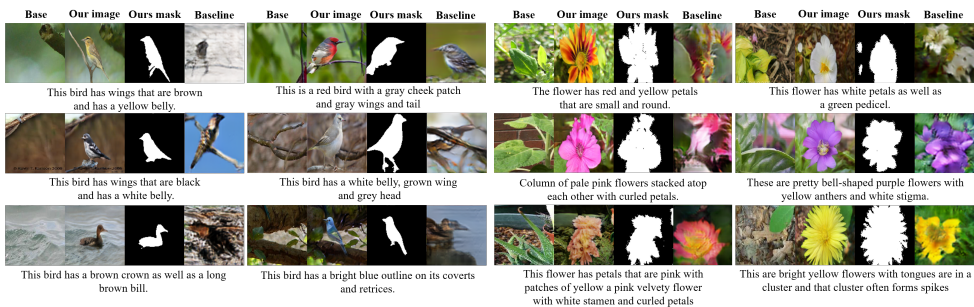


Figure 5: Examples of multi-modal image synthesis. From left to right: Base image, the synthesized image and mask from MC-GAN (128×128) and baseline method [9] (64×64)

consider image condition and both [9] and [19] are commonly based on the method in Reed et al. [17]. Thus, only the work of Dong et al. [9] was compared.

Caltech-200 bird dataset consists of 200 categories of bird images (150 categories for train and 50 categories for test), and gives ground-truth segmentation mask maps for all the 11,788 bird images. For the text attributes, we used the captions from Reed et al. [16] which contains 10 captions for each image. The captions describe the attributes of a bird such as appearance and colors. For the background image, we cropped the image patches from the Caltech-200 bird dataset excluding birds by using the segmentation mask. Separate sets of background images were used for training and test.

Oxford-102 flower dataset includes 102 categories of 8,189 flower images. The dataset is divide into a training set with 82 categories and a test set with 20 categories. To achieve the ground truth segmentation mask, we used a segmentation method of Nilsback and Zisserman [14]. For the background images, we crawled 1,352 images (1,217 for training and 135 for testing) from the web with keywords 'flower leaf' or 'flower foliage'. The captions of the flower images from [16] were used, which describes the shape and colors of the flowers.

In the training, an initial learning rate of 0.0002 and Adam optimization [9] with a momentum of 0.5 were used. To generate 128×128 bird images, we used a batchsize of 32 and trained the network for 1,200 \sim 1,500 epochs. For the flower dataset, we trained 900 \sim 1,200 epochs. We set $\lambda_1 = 2$ for experiments on both datasets while λ_2 was set to 15 for the bird and 30 for the flower. We also used an image augmentation technique including flipping, zooming and cropping randomly. The size of the seed feature map was 8×8 and 4 synthesis blocks were used to generate an 128×128 image.

4.2 Comparison with the Baseline Method

We compared our method with the baseline [9], which is also a multi-conditional GAN conditioned by an image and text. Originally, Dong et al. [9] reduced the learning rate by 0.5 in every 100 epochs and trained until 600 epochs. In our implementation, we decreased the learning rate in every 200 epochs, and trained the network with the same number of epochs, for fair comparison. Fig. 5 shows some examples of generated images from the proposed MC-GAN and the baseline work of [9]. From the figures, we confirmed that the results from [9] did not preserve the background information, or only generated background images without the target object.




Figure 6: Example results for the image synthesis problem defined in [2]. The baseline method [2] and our method (MC-GAN without mask) on Caltech-200 bird test dataset

4.3 Comparison using the Synthesis Problem in [2]

Here, we show that the synthesis block solves the multi-modal conditional problem more reliably than the baseline [2]. The semantic synthesis problem of Dong et al. [2] aims to keep the features of input images that are irrelevant to target text descriptions and to transfer the relevant part of the input image to the one that matches the target text description. We used the same text embedding method without a segmentation mask, and reduced the learning rate by 0.5 for every 100 epochs until 600 epochs under the same condition as [2]. The proposed structure of MC-GAN is applied to the generator and discriminator networks, but only the image-text pair loss is used for the discriminator as in [2].

Fig. 6 shows some examples of the baseline method and ours. Dong et al. [2] worked well if the background image is not complicated (column 3 and 6), but if the background is complex (column 5), it fails to generate a plausible object. Even though the object was generated well, the irrelevant part of image was also changed a lot. On the other hand, our generator using the proposed synthesis block stably maintained the shape of the object and the texture of the background even in a complex background, and the irrelevant background part to the text description rarely changes. Although segmentation masks were not used in the training, image features were provided to each synthesis block to keep the background and the shape of the object. Therefore, the background part which was irrelevant to the target texts was maintained and at least the shape was not changed strangely even if the color of the object changed as the text description.

4.4 Interpolation and Variety

To generate images appropriately for various sentences, the latent manifold of the text embedding should be continuously trained. We generated continuously changing images by linearly interpolating the two different text embeddings from the sentences shown at the bottom of Fig. 7. The figure shows some example images that changes its color smoothly (mainly from orange to blue / from gray to yellow) following interpolated text embedding




Figure 7: Generated images using the interpolation of two text embeddings corresponding to the text descriptions below. The leftmost is the base image.




Figure 8: Variety and stability of MC-GAN. The images are generated by linearly interpolating the noise vector z from all-zero to all-one vector.

under the same noise and image conditions.

As another experiment, we generated images using linearly interpolating the two noise vectors, which are all-zero (z_0) and all-one (z_1) vectors, under the same text and image conditions to demonstrate our model’s variety and stability. Fig. 8 shows some resultant images. Although it depends on the text and image conditions, but we usually got half of visually successful samples as can also be seen in Fig. 8.

4.5 The effect of switch

The switch in the synthesis block prevents the background and the foreground from overlapping each other. We compared the images generated by changing the switch value (the output value of the sigmoid function) under the fixed trained model of MC-GAN with the same base image, text and noise conditions. Fig. 9 shows some results. If we turn off all the switches (by zeroing the values) to prevent background information from being added, the original background disappears and only the generated object and the changed background are present. If all the switch values are set to 0.5 (half on), the background is reconstructed, but the object and the background slightly overlap and the image gets blurred compared to that of the original MC-GAN. Finally, when all the switches are turned on, the original background information are added without suppression. In this case, the object and the background overlap with each other and the object is not properly visualized. By this experiment, we can verify that the switch in the synthesis block analyzes the current image and text conditions and adjusts the image feature flexibly to assist the proper synthesis of an image.

4.6 High Resolution Image Generation by StackGAN

Based on the proposed MC-GAN, we additionally introduce a model to generate high resolution images by adding the StackGAN style two stage generator. The initial spatial size of feature map in MC-StackGAN is 4×4 and two multiple GANs were stacked for generating 128×128 images like [22, 23]. The structure of the first GAN is the same as our original method and the second GAN takes a $64 \times 64 \times 64$ final feature map from the first GAN and a text embedding code from *conditioning augmentation* without noise vector as mentioned in [22, 23]. The first GAN’s final feature map is concatenated with the replicated text embedding code. We applied a convolution to the associated feature map with batch normalization and ReLU. Finally, we used one more synthesis block and an upsampling layer to generate

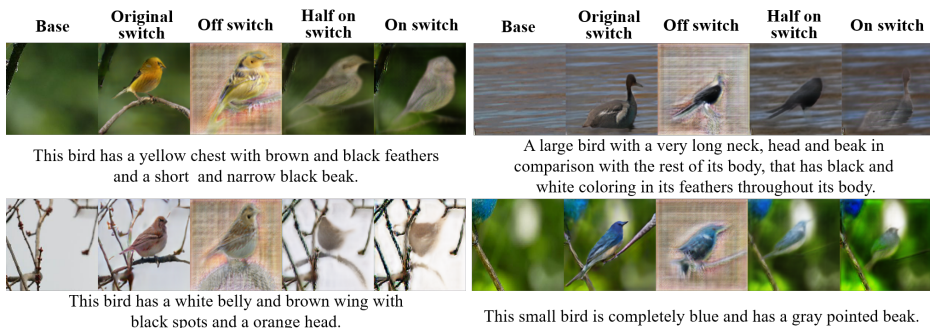


Figure 9: Experiment of generated images of switch transition in the Synthesis block

128×128 images. MC-StackGAN generates objects more stably than MC-GAN, but tends to transform the original base image compared to MC-GAN.



Figure 10: Examples of MC-GAN and MC-StackGAN. From left to right: Base image, the synthesized image from MC-GAN (128×128), the synthesized image from MC-StackGAN (128×128) on Caltech-200 bird test dataset

5 Conclusion

In this paper, we introduced a new method of GAN to generate an image given a text attribute and a base image. Different from the existing text-to-image synthesis algorithms only considering the foreground object, the proposed method aims to generate the proper target image, as well as preserving the semantics of the given background image. To solve the problem, we newly proposed an MC-GAN structure and a synthesis block which is a core component enabling a photo-realistic synthesis by smoothly mixing foreground and background information. Using the proposed method, we confirmed that our model can generate diverse forms of a target object according to the text attribute while preserving the information of the given background image. We also confirmed that our model can generate the object even when the background image does not include the same kind of object as the target, which is difficult for existing works.

6 Acknowledgement

This work was supported by Green Car Development project through the Korean MTIE (10063267) and Next-Generation Information Computing Development Program through the National Research Foundation of Korea (2017M3C4A7077582)

References

- [1] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *The IEEE International Conference on Computer Vision (ICCV)*, volume 1, 2017.
- [2] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2172–2180, 2016.
- [3] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [4] Hao Dong, Simiao Yu, Chao Wu, and Yike. Guo. Semantic image synthesis via adversarial learning. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [6] Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Jimenez Rezende, and Daan Wierstra. Draw: A recurrent neural network for image generation. 2015.
- [7] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. 2015.
- [8] Taeksoo Kim, Moonsu Cha, Hyunsoo Kim, Jungkwon Lee, and Jiwon Kim. Learning to discover cross-domain relations with generative adversarial networks. In *International Conference on Machine Learning (ICML) 2017*, 2017.
- [9] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. 2015.
- [10] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 405–415, 2017.
- [11] Elman Mansimov, Emilio Parisotto, Jimmy Ba, and Ruslan Salakhutdinov. Generating images from captions with attention. In *ICLR*, 2016.
- [12] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2813–2821. IEEE, 2017.
- [13] M-E. Nilsback and A. Zisserman. Automated flower classification over a large number of classes. In *Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing*, Dec 2008.

- [14] Maria-Elena Nilsback and Andrew Zisserman. Delving into the whorl of flower segmentation. In *BMVC*, pages 1–10, 2007.
- [15] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2016.
- [16] Scott Reed, Zeynep Akata, Honglak Lee, and Bernt Schiele. Learning deep representations of fine-grained visual descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 49–58, 2016.
- [17] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text-to-image synthesis. In *Proceedings of The 33rd International Conference on Machine Learning*, 2016.
- [18] Scott Reed, Aäron van den Oord, Nal Kalchbrenner, Victor Bapst, Matt Botvinick, and Nando de Freitas. Generating interpretable images with controllable structure. Technical report, 2016.
- [19] Scott E Reed, Zeynep Akata, Santosh Mohan, Samuel Tenka, Bernt Schiele, and Honglak Lee. Learning what and where to draw. In *Advances in Neural Information Processing Systems*, pages 217–225, 2016.
- [20] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- [21] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [22] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaolei Huang, Xiaogang Wang, and Dimitris Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 5907–5915, 2017.
- [23] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris Metaxas. Stackgan++: Realistic image synthesis with stacked generative adversarial networks. *arXiv: 1710.10916*, 2017.
- [24] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networkss. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017.