

DE-net: Dynamic Text-Guided Image Editing Adversarial Networks

Ming Tao,¹ Bing-Kun Bao,^{1*} Hao Tang,² Fei Wu,¹ Longhui Wei,³ Qi Tian³

¹Nanjing University of Posts and Telecommunications

²CVL, ETH Zürich

³Huawei Inc.

bingkunbao@njupt.edu.cn

Abstract

Text-guided image editing models have shown remarkable results. However, there remain two problems. First, they employ fixed manipulation modules for various editing requirements (e.g., color changing, texture changing, content adding and removing), which results in over-editing or insufficient editing. Second, they do not clearly distinguish between text-required and text-irrelevant parts, which leads to inaccurate editing. To solve these limitations, we propose: (i) a Dynamic Editing Block (DEBlock) that composes different editing modules dynamically for various editing requirements. (ii) a Composition Predictor (Comp-Pred), which predicts the composition weights for DEBlock according to the inference on target texts and source images. (iii) a Dynamic text-adaptive Convolution Block (DCBlock) that queries source image features to distinguish text-required parts and text-irrelevant parts. Extensive experiments demonstrate that our DE-Net achieves excellent performance and manipulates source images more correctly and accurately.

Introduction

The last few years have witnessed the great success of Generative Adversarial Networks (GANs) (Goodfellow et al. 2014) for a variety of applications (Tang et al. 2019; Zhang et al. 2019; Karras, Laine, and Aila 2019; Karras et al. 2020; Tang et al. 2020). Among them, image editing is one of the important applications of GANs. Recently, some works tried to bridge the gap between natural language and image editing (Dong et al. 2017; Liu et al. 2020b,a; Jiang et al. 2021; Li et al. 2020a,b; Nam, Kim, and Kim 2018; Xu et al. 2022). Different from other image editing tasks, text-guided image editing aims to translate source images to target images according to given text descriptions. Due to the advantages of natural language, the text guidance makes the image editing process more controllable and direct. Moreover, owing to its promising applications, this task has attracted many researchers and achieved many significant progresses (Dong et al. 2017; Nam, Kim, and Kim 2018; Li et al. 2020a,b).

However, current text-guided image editing models have two problems. First, they tackle various editing requirements equally with fixed manipulation modules. Different

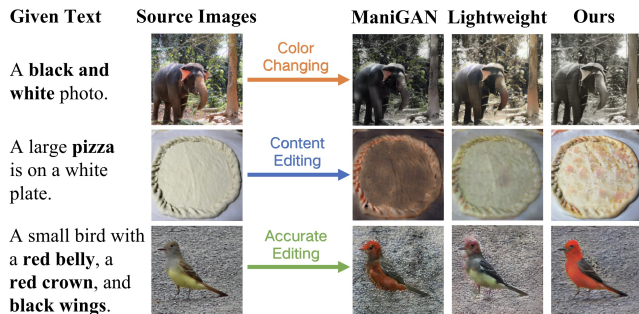


Figure 1: Benefiting from dynamic editing design, our DE-Net can deal with various editing tasks (e.g., color changing, and content editing). Furthermore, the text-adaptive convolution in DCBlock enables more accurate manipulations.

from other image-to-image tasks (e.g., image colorization, style transfer, and image inpainting), which design effective models for specific tasks, text-guided image editing is an open-target task. It contains various editing tasks (e.g., color changing, texture changing, content adding, and removing) prompted by various text guidance. But current models ignore the difference between varying editing requirements and employ fixed manipulation modules for all editing tasks. It leads the whole network tends to make a trade-off for different editing tasks. These compromises limit the manipulation ability and result in over-editing or insufficient editing. As the result shown in Figure 1, previous models (Li et al. 2020a,b) fail to deal with different editing requirements.

Second, the editing preciseness is deficient, as current manipulation modules cannot clearly distinguish between text-required and text-irrelevant parts. Text-visual feature concatenations (Dong et al. 2017; Nam, Kim, and Kim 2018) and ACM module (Li et al. 2020a,b) are widely used to manipulate image features. However, naive concatenation treats text-required and text-irrelevant parts equally, resulting in text information overwriting the source image features. And the ACM module predicts editing parameters from the encoded source images directly without introducing text information, which makes it unable to accurately edit text-required parts and preserve text-irrelevant parts. As the result shown in Figure 1, the previous model (Li et al. 2020a,b) diffuses the text information to the whole image.

*Corresponding author

To address these problems, we propose a novel Dynamic text-guided image Editing adversarial Network (i.e., DE-Net). Different from previous works (Dong et al. 2017; Nam, Kim, and Kim 2018; Li et al. 2020a,b), which employ fixed editing blocks, our DE-Net employs a Dynamic Editing Block (DEBlock). The DEBlock combines spatial- and channel-wise manipulations dynamically for different editing requirements. Furthermore, we propose the Composition Predictor (Comp-Pred) to predict the composition weights of DEBlock. The DEBlock and Comp-Pred enable our network can deal with various editing tasks through suitable combinations of spatial- and channel-wise manipulations. In addition, to improve the editing preciseness of the manipulation module, we propose a novel Dynamic text-adaptive Convolution Block (DCBlock). The DCBlock employs a dynamic text-adaptive convolutional layer that can automatically adjust the kernel weights according to the given text guidance. This enables our model to distinguish text-required and text-irrelevant image features by querying the visual features according to text guidance. Moreover, we add a lightweight semantic decoder that provides global-local visual features for DCBlock to obtain a more accurate prediction of editing parameters. Overall, our contributions can be summarized as follows:

- We propose a novel Dynamic text-guided Editing Block (DEBlock) to enable our model can deal with a variety of editing tasks adaptively through the dynamic composition of different editing modules.
- We propose a novel Composition Predictor (Comp-Pred), which predicts the composition weights for DEBlock according to the inference on the text and visual features.
- We propose a novel Dynamic text-adaptive Convolution Block (DCBlock), which can distinguish text-required parts and text-irrelevant parts of source images.
- Compared with the current state-of-the-art methods, our DE-Net achieves much better performance on commonly used public datasets.

Related Work

Text-to-Image Synthesis bridges text information and image generation. Based on Generative Adversarial Network (GAN) (Goodfellow et al. 2014), some researchers tried to bridge the text and image modality. (Reed et al. 2016) first employed the conditional Generative Adversarial Network (cGAN) to synthesize plausible images from given text descriptions. (Zhang et al. 2017, 2018) proposed the stacked architecture, which stacks multiple generators and discriminators to synthesize high-resolution images. (Xu et al. 2018) introduced the attention mechanism to fuse the word information and image features. This text-image attention mechanism has been widely applied in many text-related generative works (Zhu et al. 2019; Cheng et al. 2020). (Tao et al. 2022) proposed the DF-GAN with deep fusion block and matching-aware gradient penalty, which enables one-stage high-quality generation. Recently, some large transformer-based text-to-image methods (Ramesh et al. 2021; Lin et al. 2021; Ding et al. 2021) show excellent performance on complex image synthesis. They tokenize the images and take

the image tokens and word tokens to make auto-regressive training by a unidirectional transformer (Radford et al. 2019; Brown et al. 2020).

Text-Guided Image Editing aims to manipulate source images according to given text descriptions. For example, (Dong et al. 2017) first proposed SISGAN based on an encoder-decoder architecture. It concatenates the encoded representations with text semantics to manipulate source images. The TAGAN (Nam, Kim, and Kim 2018) introduced a word-level attention mechanism in a discriminator to classify fine-grained attributes independently. However, both SISGAN (Dong et al. 2017) and TAGAN (Nam, Kim, and Kim 2018) only preserve coarse visual features from source images. (Li et al. 2020a) proposed the ManiGAN, which is a two-stage architecture with the ACM and DCM modules. The ACM module synthesizes text-matching images while preserving a rough shape from source images. And the DCM module rectifies mismatched image features and completes missing details. Then, (Li et al. 2020b) proposed a more efficient word-level discriminator to facilitate training a lightweight generator. Recently, the Mani-Trans (Wang et al. 2022) employ the extra knowledge of a pretrained autoregressive transformer for text-guided image editing. These methods have shown significant progress in text-guided image editing tasks. Recently, some text-guided image editing works focus on face editing (Xia et al. 2021a,b; Patashnik et al. 2021). They employ a pretrained StyleGAN (Karras, Laine, and Aila 2019) and manipulate its latent according to face descriptions.

In this work, we focus on the main research direction of the text-guided image editing task, which deals with open-target image editing (Dong et al. 2017; Nam, Kim, and Kim 2018; Li et al. 2020a,b; Xu et al. 2022). Unlike previous models, our DE-Net employs a Dynamic Editing Block (DEBlock) with a Composition Predictor (Comp-Pred) to combine the editing processes adaptively for various editing requirements. Moreover, we equip our DE-Net with a Dynamic text-adaptive Convolution Block (DCBlock) which can distinguish between text-required and text-irrelevant image features. Compared with previous works, our model edits the text-required image features more efficiently and preserves text-irrelevant source image features more accurately.

The Proposed DE-Net

In this paper, we propose a novel GAN model for text-guided image editing named Dynamic text-guided Image Editing adversarial Networks (DE-Net) (see Figure 2). To manipulate source images according to text guidance more effectively and accurately, we propose: (i) a novel manipulation module called Dynamic Editing Block (DEBlock), which manipulates the image features adaptively through the dynamic combination between spatial and channel-wise editing. (ii) a novel Composition Predictor (Comp-Pred) that can predict the combination weights according to the inference on text and visual features. (iii) a novel Dynamic text-adaptive Convolution Block (DCBlock) that improves the editing preciseness through querying the visual features according to text guidance. In the following of this section,

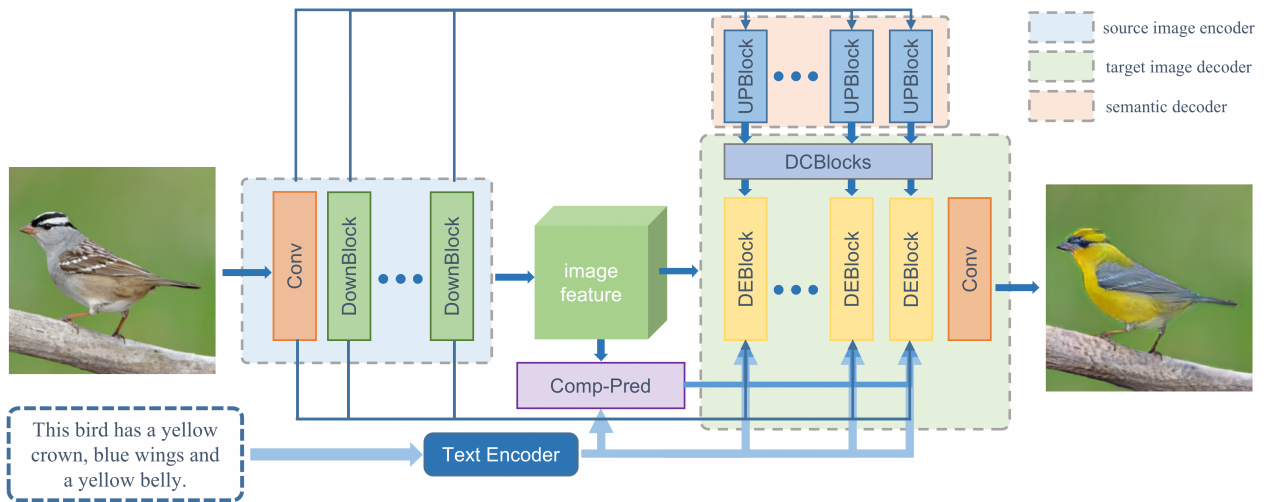


Figure 2: The architecture of the proposed DE-Net. The DE-Net comprises a source image encoder, a target image decoder, a semantic decoder, a Composition Predictor (Comp-Pred), and a pretrained text encoder (Xu et al. 2018). The DEBlock and DCBlock are introduced in the target image decoder to enable effective and accurate manipulations, respectively.

we first present the overall structure of our DE-Net and then introduce the Comp-Pred, DEBlock, and DCBlock in detail.

Dynamic Text-guided Image Editing Framework

Different from previous models (Dong et al. 2017; Nam, Kim, and Kim 2018; Li et al. 2020a,b; Wang et al. 2022), our DE-Net is based on a novel dynamic text-guided image editing framework. First, it can compose suitable editing processes based on input images and texts dynamically. Second, it can distinguish between text-required and text-irrelevant image features through dynamic text-adaptive convolution.

As shown in Figure 2, the whole editing framework is composed of a source image encoder, a target image decoder, a Composition Predictor (Comp-Pred), a semantic decoder, and a pre-trained text encoder (Li et al. 2020a,b). To introduce multi-scale source visual features for manipulation and preservation, we add full skip connections between the source image encoder and two decoders. In DE-Net, the source image is first encoded to 4×4 resolution by the source image encoder. Then, the encoded source image feature is fed into the semantic decoder and target image decoder, respectively. The semantic decoder comprises several upsampling and convolutional layers (UPBlock). It can capture long-range visual information and decode the global-local semantic features for the DCBlock through the growing size of semantic features. The target image decoder is composed of multiple DEBlocks and DCBlocks. The DEBlock manipulates source image features both from spatial and channel dimensions and dynamically combines these two manipulations at different image scales. And the Comp-Pred predicts the combination weights for DEBlock according to the inference on the text and visual features. The DCBlock adjusts its candidate convolution kernels and applies text-adaptive convolution on global-local semantic features provided by the semantic decoder. It enables our DE-Net to distinguish between text-required and text-irrelevant parts.

To stabilize the training process of adversarial learning, we further introduce the target-aware discriminator (Tao et al. 2022) in DE-Net. Finally, our DE-Net can be formulated as:

$$\begin{aligned}
 L_D &= -\mathbb{E}[\min(0, -1 + D(x, e))] \\
 &\quad - (1/2)\mathbb{E}[\min(0, -1 - D(G(x, \hat{e}), \hat{e}))] \\
 &\quad - (1/2)\mathbb{E}[\min(0, -1 - D(x, \hat{e}))] \\
 &\quad + k\mathbb{E}[(\|\nabla_x D(x, e)\| + \|\nabla_e D(x, e)\|)^p], \\
 L_G &= -\mathbb{E}[D(G(x, \hat{e}), \hat{e})] + \lambda_1 \mathbb{E}[\|G(x, \hat{e}) - x\|_2] \\
 &\quad - \lambda_2 \mathbb{E}[S(G(x, \hat{e}), \hat{e})],
 \end{aligned} \tag{1}$$

where x denotes the source images; e denotes the source sentence embeddings, \hat{e} denotes the target sentence embeddings, G denotes the generator network, D denotes the discriminator network, k and p are two hyperparameters of the discriminator to balance the effectiveness of gradient penalty. S represents the cosine similarity between the encoded visual and text features predicted by pre-trained DAMSM network (Xu et al. 2018; Li et al. 2020a,b). λ_1 and λ_2 are two balancing coefficients of the generator.

Composition Predictor

To enable our DE-Net to compose suitable editing processes based on current source images and text guidance, we propose the Composition Predictor (Comp-Pred). As shown in Figure 2 and Figure 3, for visual information, it takes the encoded 4×4 image feature and maps it to the visual vector V through two convolution layers. For text information, it takes the sentence embedding provided by the text encoder and maps it to a text vector T through one-hidden-layer MLP. Then we concatenate the V and T and get the V' and T' through an MLP to close the gap between the visual and text domain. The difference $V' - T'$ and the element-wise product $V' * T'$ are calculated and concatenated with the original vectors V and T to enhance the distance between

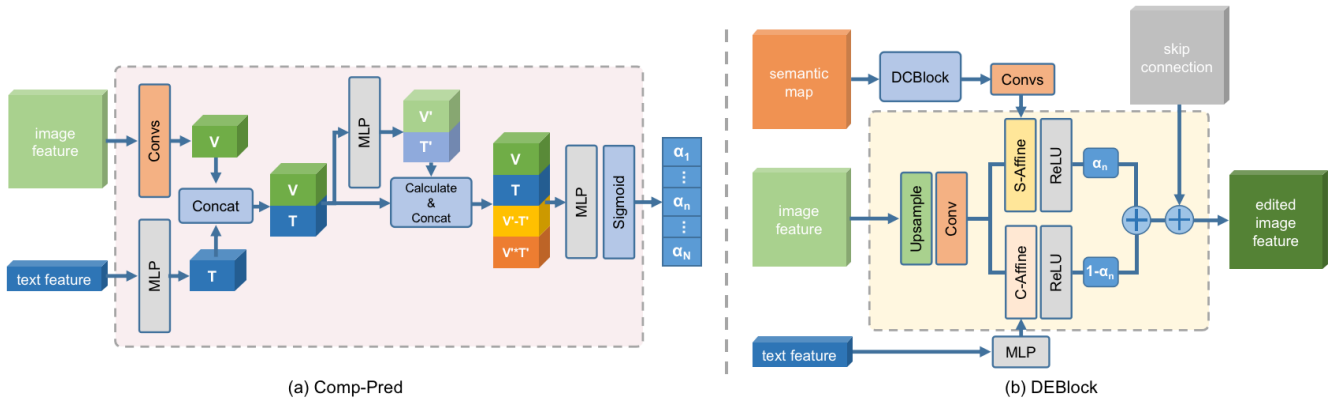


Figure 3: Illustration of the proposed Comp-Pred and DCBlock. (a) The Comp-Pred predicts the combination weights for DEBlock. (b) The DEBlock composes suitable editing processes for each input pair according to the combined weights of C-Affine and S-Affine.

text and visual information. Lastly, we adopt an MLP and a Sigmoid function to predict the combination weights α of spatial-wise and channel-wise editing for N DEBlocks in the target image decoder. The α_n is a vector, and its length is the channel size of the visual feature after the convolution layer in each DEBlock.

Dynamic Text-guided Image Editing Block

As mentioned before, the existing naive concatenation (Dong et al. 2017; Nam, Kim, and Kim 2018) or ACM module (Nam, Kim, and Kim 2018; Li et al. 2020b) cannot make an efficient interaction between text guidance and source image features. Moreover, these manipulation modules treat various editing requirements equally, ignoring the different model requirements. To address these problems, we first decompose the visual feature editing process into spatial and channel-wise editing. Then, we propose a novel DEBlock to re-compose suitable editing processes for each input pair through the dynamic combination of a spatial and channel-wise affine transformation.

As shown in Figure 3(a), the DEBlock is composed of upsampling layer, convolution layer, and two kinds of affine layers. One is the Channel-wise Affine layer (C-Affine), and the other is the Spatial Affine layer (S-Affine). To fully manipulate source image features, our DEBlock manipulates both spatial and channel dimensions and combines these two manipulations dynamically through combination weights α_n predicted by Comp-Pred. Specifically, in C-Affine, the text sentence vector t is sent to two different one-hidden-layer MLPs to predict the scaling parameters γ_c and shifting parameters θ_c for each channel:

$$\gamma_c = \text{MLP}_1(t), \quad \theta_c = \text{MLP}_2(t). \quad (2)$$

After obtaining γ_c and θ_c parameters from MLPs, the C-Affine transformation can be formally expressed as follows:

$$\text{C-Aff}(x_{c,h,w}|t) = \gamma_c \cdot x_{c,h,w} + \theta_c, \quad (3)$$

where x is the image feature; c , h , and w denote the channel size, feature height, and feature width, respectively.

Unlike the C-Affine, the S-Affine applies affine transformations on the spatial dimension of image features. The spatial scaling parameters $\gamma_{h,w}$ and shifting parameters $\theta_{h,w}$ are predicted by the DCBlock and two convolution layers for each pixel in the image feature. As shown in Figure 4, the DCBlock employs the text-adaptive convolution layer to get the text-attended visual features f . Then, the spatial scaling and shifting parameters are predicted by two convolution layers on the predicted text-attended visual features:

$$\gamma_{h,w} = \text{Conv}_1(f), \quad \theta_{h,w} = \text{Conv}_2(f). \quad (4)$$

Thus, the S-Affine can be formally expressed as follows:

$$\text{S-Aff}(x_{c,h,w}|t) = \gamma_{h,w} \cdot x_{c,h,w} + \theta_{h,w}. \quad (5)$$

To fully manipulate source image features, our model integrates the C-Affine and S-Affine in DEBlock. And the n th DEBlock in the target image decoder combines these two manipulations dynamically through the combination weights α_n predicted by Comp-Pred:

$$\text{DEBLK}(x|t) = \text{C-Aff}(x|t) \cdot \alpha_n + \text{S-Aff}(x|f) \cdot (1 - \alpha_n). \quad (6)$$

Compared with previous manipulation modules, our DEBlock fully manipulates source image features both from spatial and channel dimensions and considers different contributions of these two dimensions at different image scales. Moreover, our experiments and ablation studies also demonstrate the effectiveness of the DEBlock.

Dynamic Text-adaptive Convolution Block

To improve the editing preciseness of the manipulation module, we propose the DCBlock as an editing positioning module. As shown in Figure 4, the weights of convolution kernels in DCBlock are adjusted according to the given text guidance. However, directly generating all parameters of the convolution kernel not only requires a high computational cost but also results in overfitting to train text descriptions. To address this problem, we apply scaling and shifting operations for each candidate convolution kernel. First, we employ two MLPs to predict the scaling parameters $\varphi_{c,in}$ and

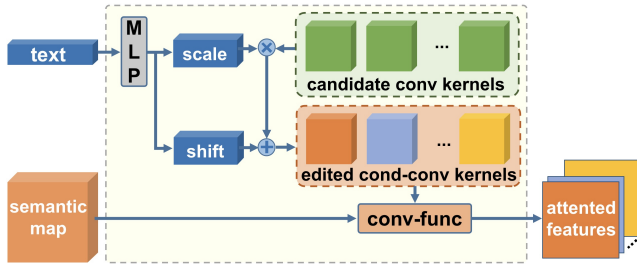


Figure 4: Illustration of the DCBlock. The text-attended features are predicted by the text-conditioned convolution layer.

shifting parameters $\omega_{c.in}$ for the convolution kernel. The modulation on one convolution kernel can be formally expressed as follows:

$$\text{Mod}(\mathbf{k}) = \varphi_{c.in} \cdot \mathbf{k}_{c.in,k_h,k_w} + \omega_{c.in}, \quad (7)$$

where $\mathbf{k}_{c.in,k_h,k_w}$ is the convolution kernel; $c.in$, k_h , and k_w denotes the input channel size, kernel height, and kernel width, respectively.

With this text-guided adjustment for candidate convolution kernels, the dynamic convolutional (Dy-Conv) operation acts as a querying process according to text descriptions. However, as the semantic information in one convolution kernel is limited, it is hard to incorporate the whole text semantic information into one convolution kernel. Inspired by the multi-head attention mechanism in Transformer (Vaswani et al. 2017), we employ multiple dynamic convolution kernels as different attention heads. In our work, we set 8 dynamic convolution kernels in DCBlock. It allows our model to jointly attend to different image features from different text representation subspaces. As shown in Figure 4, each dynamic convolution kernel predicts an attended feature map according to partial text semantics. Then, we stack these attention features along the channel direction as a whole attended feature map for predicting the spatial affine parameters for DEBlock.

However, directly applying Dy-Conv operations on image features from the source image encoder is inefficient. Each pixel of image features in the source image encoder only contains the semantic information of its local neighborhood through downsampling layers. It cannot capture long-range contextual information, resulting in low accuracy and efficiency of Dy-Conv. To cope with this problem, we employ a semantic decoder to decode global-local semantics, as shown in Figure 2. The source image is first downsampled through the image encoder and then upsampled by the semantic decoder with full skip connections from the source image encoder. Through this encoder-decoder structure with full skip connections, the semantic decoder can capture long-range visual information, and each pixel of decoded semantic features contains global-local semantic information. We apply Dy-Conv operations on the global-local semantic features to achieve more accurate queried results. After that, we employ two convolutional layers to predict the spatial scaling and shifting parameters for DEBlock.

The proposed DCBlock enables our DE-Net to distinguish between text-required and text-irrelevant parts. Furthermore,

the semantic decoder provides the global-local semantic features, and the multi-head dynamic convolutional operations query the global-local semantic features from different text representation subspaces. Armed with the DCBlock, the target image decoder can manipulate source image features more accurately.

Experiments

In this section, we introduce the datasets, training details, and evaluation metrics used in our experiments. Then we compare the text-guided image editing performance with previous models quantitatively and qualitatively.

Datasets. We conduct experiments on two challenging datasets: CUB bird (Wah et al. 2011) and COCO (Lin et al. 2014). For the CUB bird dataset, there are 11,788 images belonging to 200 bird species, with each image corresponding to ten language descriptions. The COCO dataset contains 80k images for training and 40k images for testing. Each image corresponds to 5 language descriptions.

Training and Evaluation Details. We employ the Adam optimizer (Kingma and Ba 2015) with $\beta_1=0.0$ and $\beta_2=0.9$ to train our model. According to the Two Timescale Update Rule (TTUR) (Heusel et al. 2017), the learning rate is set to 0.0001 for the generator and 0.0004 for the discriminator. The hyper-parameters of the discriminator k and p are set to 2 and 6 as (Tao et al. 2022). The hyper-parameters of the generator λ_1 and λ_2 are set to 40 and 4 for all the datasets.

Following previous works (Wang et al. 2022; Li et al. 2020a), we employ the Inception Score (IS) (Salimans et al. 2016), CLIP-sim, L2-error, and Manipulation Precision (MP) to evaluate the performance. Higher IS means higher fidelity of the edited images. CLIP-sim is the text-image cosine similarity calculated by a pretrained CLIP network (Xu et al. 2018) to measure the quality of text-guided manipulation. L2-error is the L_2 pixel difference between edited images and source images to measure the quality of source feature preservation. Lastly, MP simultaneously measures the quality of text-guided manipulation and source feature preservation. It can be formulated as:

$$\text{MP} = (1 - \text{L2_error}) \times (\text{CLIP_sim}). \quad (8)$$

It must be pointed out that the Fréchet Inception Distance (FID) (Heusel et al. 2017) widely adopted in text-to-image synthesis is not suitable for text-guided image editing. The FID calculates the Fréchet distance between synthetic images and ground truth in the feature space of a pre-trained Inception v3 network. But there is no ground truth (human-edited images) in this task, and calculating with source images makes the FID prefer lazy models which do not edit input images. So we adopt the IS to evaluate the image fidelity as previous works (Wang et al. 2022; Li et al. 2020a).

Quantitative Comparisons

We compare the proposed method with several state-of-the-art text-guided image editing models, including TAGAN (Nam, Kim, and Kim 2018), ManiGAN (Li et al. 2020a), ManiTrans (Wang et al. 2022). We also modify the state-of-the-art text-to-image model DF-GAN (Tao et al. 2022) and

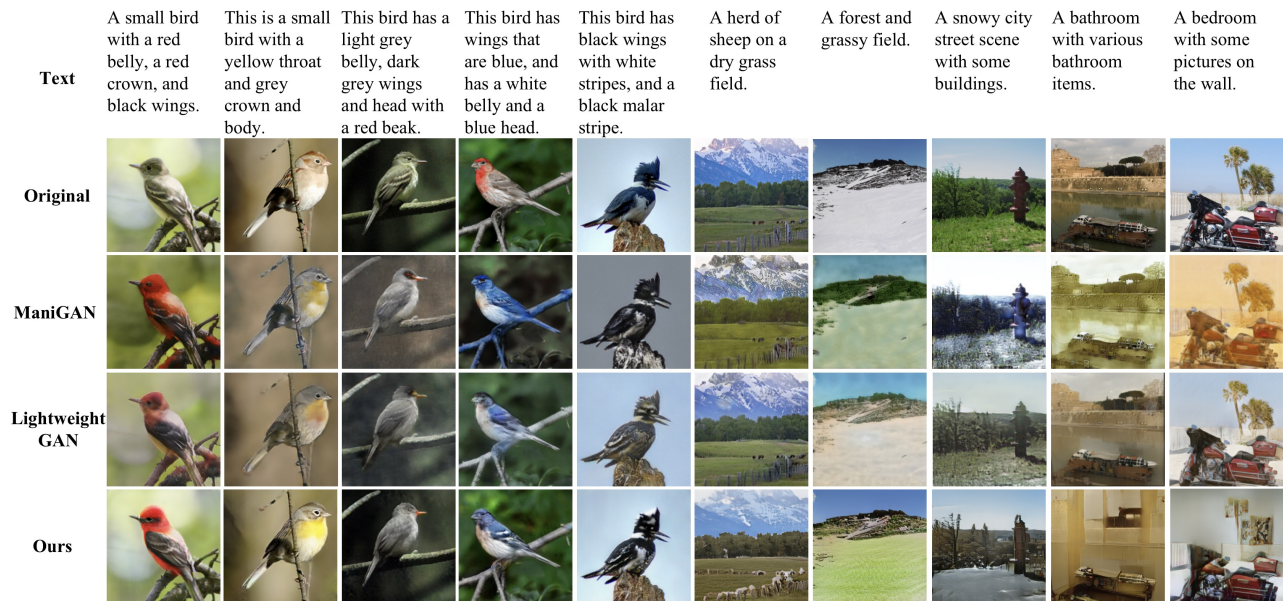


Figure 5: Qualitative comparison between different methods on the test set of CUB and COCO.

Model	CUB				COCO			
	IS \uparrow	CLIP-sim \uparrow	L2-error \downarrow	MP \uparrow	IS \uparrow	CLIP-sim \uparrow	L2-error \downarrow	MP \uparrow
TAGAN	3.72	0.202	0.185	0.164	9.33	0.125	0.102	0.112
ManiGAN	4.19	0.213	0.051	0.202	22.60	0.119	0.031	0.114
DF-GAN	4.54	0.221	0.046	0.210	20.10	0.129	0.027	0.125
Lightweight-GAN	4.66	0.188	0.132	0.163	24.80	0.136	0.025	0.132
ManiTrans	5.02	0.235	0.013	0.231	21.45	0.131	0.017	0.128
DE-Net (Ours)	5.08	0.240	0.010	0.237	25.81	0.192	0.015	0.189

Table 1: Quantitative comparison between different methods on the test set of CUB and COCO.

compare our model with it. As shown in Table 1, our proposed DE-Net achieves the highest IS, CLIP-sim, MP, and the lowest L2-error on the CUB and COCO test dataset. The highest IS shows that our method can produce more realistic manipulated results. And the highest CLIP-sim, MP, and lowest L2-error prove that our method can manipulate text-required image parts and preserve text-irrelevant parts more correctly and accurately. The extensive quantitative evaluation results on CUB and COCO demonstrate the superiority and effectiveness of our proposed DE-Net. Furthermore, the advantages of our model are more obvious in the complex COCO images. It demonstrates the significant improvements of our DE-Net when dealing with various editing requirements, e.g., textual changing, background changing, object adding and removing.

Qualitative Comparisons

In this subsection, we compare the images synthesized by ManiGAN (Li et al. 2020a), Lightweight GAN (Li et al. 2020b), and our DE-Net. We evaluate the quality of the edited images from two aspects, text-related source feature manipulation and text-irrelevant source feature preservation. First, we compare the quality of the edited images on the

CUB dataset, then compare the results on the more challenging COCO dataset.

As the results shown in 1st, 2nd, and 4th columns, both ManiGAN (Li et al. 2020a) and Lightweight GAN (Li et al. 2020b) cannot accurately synthesize the “black wings”, “yellow throat”, and “white belly”, respectively. Moreover, they tend to change the color of twigs and backgrounds (see 1st, 3rd, 4th, and 5th columns). While our DE-Net can manipulate these features successfully and preserve the text-irrelevant image features like twigs and backgrounds more accurately. The superiority is more obvious in the challenging COCO dataset, which contains various editing tasks. As the results are shown in the 6th, 7th, and 10th columns, both ManiGAN (Li et al. 2020a) and Lightweight GAN (Li et al. 2020b) cannot synthesize the “a herd of sheep”, “grassy field”, and “some pictures”, respectively. These results show that they only tend to edit the colors of source images (see 8th, 9th, and 10th columns). However, our DE-Net can both edit the color (see 6th, and 7th columns), texture (see 7nd and 8th columns) of source images, and synthesize new contents in target images (see 6th and 10th columns). Moreover, both ManiGAN (Li et al. 2020a) and Lightweight GAN (Li et al. 2020b) tend to edit the text-irrelevant regions. For ex-

Method	IS \uparrow	CLIP-sim \uparrow	L2 \downarrow	MP \uparrow
Baseline (B)	11.08	0.115	0.057	0.108
B w/ CAFF	12.16	0.124	0.051	0.117
B w/ SAFF	13.33	0.131	0.047	0.124
B w/ DE w/ CP	17.21	0.155	0.031	0.150
B w/ DE w/o CP	15.36	0.141	0.040	0.135
B w/ DE w/ CP*	14.13	0.145	0.039	0.139
B w/ DE w/ CP w/OC	18.47	0.160	0.030	0.155
B w/ DE w/ CP w/DC1	20.21	0.171	0.025	0.166
B w/ DE w/ CP w/DC4	22.18	0.176	0.021	0.172
B w/ DE w/ CP w/DC16	21.01	0.179	0.021	0.175
B w/ DE w/ CP w/DC	23.11	0.180	0.018	0.176
B w/ DE w/ CP w/DC w/SD	25.81	0.192	0.015	0.189

Table 2: Ablation Study of DE-Net on the test set of COCO.

ample, the color of the sky is also painted with the color of “grassy field” (see 7th column). While our DE-Net can preserve these text-irrelevant features or manipulate them to adapt to the text description (see 6th, and 7th columns).

The qualitative comparison of both CUB and COCO datasets shows that our DE-Net can manipulate the source images according to the text guidance while preserving text-irrelevant image features more effectively and accurately.

Ablation Study

To verify the effectiveness of different components in DE-Net, we conduct ablation studies on the COCO test dataset. The components include Dynamic Text-guided Image Editing Block (DE), Composition Predictor (CP), Dynamic Text-adaptive Convolution Block (DC), and Semantic Decoder (SD). The results are provided in Table 2.

Baseline. Our baseline is based on an encoder-decoder architecture that employs the target-way discriminator (Tao et al. 2022). The baseline concatenates the text features with the encoded image features.

Effect of DEBlock (DE). The experimental results show that the DEBlock achieves better performances than only employing channel-wise (B w/CAFF) or spatial-wise affine transformation (B w/SAFF). The results demonstrate that our Dynamic Editing strategy, which composes suitable editing processes dynamically based on input text-image pairs, is more effective than fixed editing modules.

Effect of Comp-Pred (CP). Removing the Comp-Pred (B w/ DE w/o CP) reduces the editing performance. To further evaluate the effectiveness of the difference and element-wise product in Comp-Pred, we remove them from CP and name the new module CP*. Compared with CP*, the Comp-Pred (B w/ DE w/ CP) achieves better performances. The comparison results prove the effectiveness of Comp-Pred.

Effect of DCBlock (DC). To verify the effectiveness of DCBlock, we replace it with an Ordinary Convolution layer (OC). The results show that the DCBlock enables our model to manipulate more accurately (lower DIFF, higher MP). We further evaluate the multi-head mechanism of DCBlock by changing the number of text-adaptive convolution kernels to 1, 4, 16 (DC1, DC4, DC16). The results show that more convolution kernels do not mean better performance. The reason may be that more convolution kernels reduce the con-

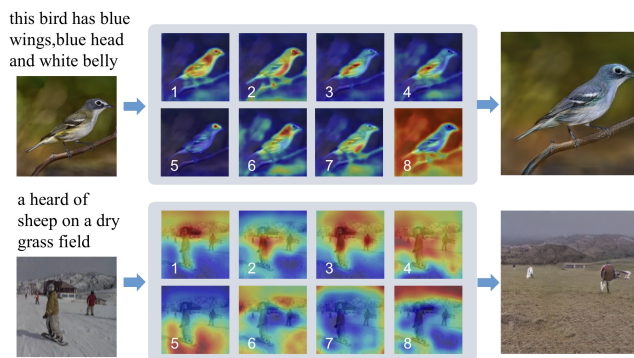


Figure 6: Visualization of the queried results by different text-adaptive convolution kernels.

vergence efficiency of the DE-Net. Therefore, we set 8 text-adaptive convolution kernels in DCBlock.

Effect of Semantic Decoder (SD). Without the Semantic Decoder (B w/ DE w/ CP w/DC), the model predicts the spatial editing parameters from the feature provided by the source image encoder. Armed with the Semantic Decoder, the model (B w/ DE w/ CP w/DC w/SD) improves the IS, SIM, MP, and decreases the DIFF. The results prove that the semantic decoder can promote the DCBlock to make more effective and accurate manipulations.

Visualisation Results of DCBlock

To further demonstrate the effectiveness of DCBlock, we visualize the text-conditioned queried results at the last DCBlock in the target image decoder. As shown in Figure 6, the kernel attends to different semantics in source image features. In addition, we find that the convolution kernel can be divided into two types, one focusing on text-related image parts and the other focusing on text-irrelevant image parts. For the CUB example, kernel 8 segments the whole background, which is text-irrelevant image parts. The other kernel attends to text-related image parts, e.g., “wings”, “head”, “belly”. For the COCO example, kernel 8 segments the sky, which is a text-irrelevant image part. The other kernel attends to text-related image parts, e.g., ground, human, and mountain, which need to be edited. It demonstrates that our DCBlock can attend to different semantic parts and distinguish between text-related and text-irrelevant parts.

Conclusion

In this paper, we propose a novel DE-Net for the text-guided image editing task. Via DE-Net, we propose a DEBlock to enable our model can deal with a variety of editing tasks adaptively through the dynamic composition of different editing modules. We also propose a novel Comp-Pred that compares the source image features and given text guidance and predict the combination weights for DEBlock. Furthermore, we propose a new DCBlock with a semantic decoder to help the target image decoder distinguish between text-required and text-irrelevant parts.

Acknowledgments

This work was supported by National Key Research and Development Project (No.2020AAA0106200), the National Nature Science Foundation of China under Grants (No.61936005, 61872424), and the Natural Science Foundation of Jiangsu Province (Grants No. BK20200037).

References

- Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Cheng, J.; Wu, F.; Tian, Y.; Wang, L.; and Tao, D. 2020. RiFeGAN: Rich Feature Generation for Text-to-Image Synthesis From Prior Knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10911–10920.
- Ding, M.; Yang, Z.; Hong, W.; Zheng, W.; Zhou, C.; Yin, D.; Lin, J.; Zou, X.; Shao, Z.; Yang, H.; et al. 2021. Cogview: Mastering text-to-image generation via transformers. *Advances in Neural Information Processing Systems*, 34.
- Dong, H.; Yu, S.; Wu, C.; and Guo, Y. 2017. Semantic image synthesis via adversarial learning. In *Proceedings of the IEEE International Conference on Computer Vision*, 5706–5714.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2672–2680.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, 6626–6637.
- Jiang, W.; Xu, N.; Wang, J.; Gao, C.; Shi, J.; Lin, Z.; and Liu, S. 2021. Language-Guided Global Image Editing via Cross-Modal Cyclic Mechanism. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2115–2124.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4401–4410.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8110–8119.
- Kingma, D. P.; and Ba, J. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Li, B.; Qi, X.; Lukasiewicz, T.; and Torr, P. H. 2020a. Manigan: Text-guided image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7880–7889.
- Li, B.; Qi, X.; Torr, P. H.; and Lukasiewicz, T. 2020b. Lightweight generative adversarial networks for text-guided image manipulation. *arXiv preprint arXiv:2010.12136*.
- Lin, J.; Men, R.; Yang, A.; Zhou, C.; Ding, M.; Zhang, Y.; Wang, P.; Wang, A.; Jiang, L.; Jia, X.; et al. 2021. M6: A Chinese Multimodal Pretrainer. *arXiv preprint arXiv:2103.00823*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Liu, X.; Lin, Z.; Zhang, J.; Zhao, H.; Tran, Q.; Wang, X.; and Li, H. 2020a. Open-edit: Open-domain image manipulation with open-vocabulary instructions. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, 89–106. Springer.
- Liu, Y.; De Nadai, M.; Cai, D.; Li, H.; Alameda-Pineda, X.; Sebe, N.; and Lepri, B. 2020b. Describe what to change: A text-guided unsupervised image-to-image translation approach. In *Proceedings of the 28th ACM International Conference on Multimedia*, 1357–1365.
- Nam, S.; Kim, Y.; and Kim, S. J. 2018. Text-adaptive generative adversarial networks: Manipulating images with natural language. In *Advances in neural information processing systems*, 42–51.
- Patashnik, O.; Wu, Z.; Shechtman, E.; Cohen-Or, D.; and Lischinski, D. 2021. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2085–2094.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-Shot Text-to-Image Generation. *arXiv preprint arXiv:2102.12092*.
- Reed, S.; Akata, Z.; Yan, X.; Logeswaran, L.; Schiele, B.; and Lee, H. 2016. Generative adversarial text to image synthesis. In *Proceedings of the International Conference on Machine Learning*, 1060–1069.
- Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; and Chen, X. 2016. Improved techniques for training gans. In *Advances in neural information processing systems*, 2234–2242.
- Tang, H.; Xu, D.; Sebe, N.; Wang, Y.; Corso, J. J.; and Yan, Y. 2019. Multi-channel attention selection gan with cascaded semantic guidance for cross-view image translation. In *CVPR*.
- Tang, H.; Xu, D.; Yan, Y.; Torr, P. H.; and Sebe, N. 2020. Local class-specific and global image-level generative adversarial networks for semantic-guided scene generation. In *CVPR*.
- Tao, M.; Tang, H.; Wu, F.; Jing, X.-Y.; Bao, B.-K.; and Xu, C. 2022. DF-GAN: A Simple and Effective Baseline for Text-to-Image Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16515–16525.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.

Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology.

Wang, J.; Lu, G.; Xu, H.; Li, Z.; Xu, C.; and Fu, Y. 2022. ManiTrans: Entity-Level Text-Guided Image Manipulation via Token-wise Semantic Alignment and Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10707–10717.

Xia, W.; Yang, Y.; Xue, J.-H.; and Wu, B. 2021a. TediGAN: Text-Guided Diverse Face Image Generation and Manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2256–2265.

Xia, W.; Yang, Y.; Xue, J.-H.; and Wu, B. 2021b. Towards Open-World Text-Guided Face Image Generation and Manipulation. *arXiv preprint arXiv:2104.08910*.

Xu, T.; Zhang, P.; Huang, Q.; Zhang, H.; Gan, Z.; Huang, X.; and He, X. 2018. AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1316–1324.

Xu, Z.; Lin, T.; Tang, H.; Li, F.; He, D.; Sebe, N.; Timofte, R.; Van Gool, L.; and Ding, E. 2022. Predict, Prevent, and Evaluate: Disentangled Text-Driven Image Manipulation Empowered by Pre-Trained Vision-Language Model. In *CVPR*.

Zhang, H.; Goodfellow, I.; Metaxas, D.; and Odena, A. 2019. Self-attention generative adversarial networks. In *International Conference on Machine Learning*, 7354–7363.

Zhang, H.; Xu, T.; Li, H.; Zhang, S.; Wang, X.; Huang, X.; and Metaxas, D. N. 2017. StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, 5907–5915.

Zhang, H.; Xu, T.; Li, H.; Zhang, S.; Wang, X.; Huang, X.; and Metaxas, D. N. 2018. StackGAN++: Realistic image synthesis with stacked generative adversarial networks. *IEEE TPAMI*, 41(8): 1947–1962.

Zhu, M.; Pan, P.; Chen, W.; and Yang, Y. 2019. Dm-gan: Dynamic memory generative adversarial networks for text-to-image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5802–5810.