

Fine-Grained Cross-Modal Fusion Based Refinement for Text-to-Image Synthesis

SUN Haoran², WANG Yang^{1,2}, LIU Haipeng², and QIAN Biao²

(1. *Key Laboratory of Knowledge Engineering with Big Data, Ministry of Education, Hefei University of Technology, Hefei 230000, China*)

(2. *Department of Computer Science and Information Engineering, Hefei University of Technology, Hefei 230000, China*)

Abstract — Text-to-image synthesis refers to generating visual-realistic and semantically consistent images from given textual descriptions. Previous approaches generate an initial low-resolution image and then refine it to be high-resolution. Despite the remarkable progress, these methods are limited in fully utilizing the given texts and could generate text-mismatched images, especially when the text description is complex. We propose a novel fine-grained text-image fusion based generative adversarial networks (FF-GAN), which consists of two modules: Fine-grained text-image fusion block (FF-Block) and global semantic refinement (GSR). The proposed FF-Block integrates an attention block and several convolution layers to effectively fuse the fine-grained word-context features into the corresponding visual features, in which the text information is fully used to refine the initial image with more details. And the GSR is proposed to improve the global semantic consistency between linguistic and visual features during the refinement process. Extensive experiments on CUB-200 and COCO datasets demonstrate the superiority of FF-GAN over other state-of-the-art approaches in generating images with semantic consistency to the given texts.

Key words — Text-to-image synthesis, Text-image fusion, Generative adversarial network.

I. Introduction

Text-to-image synthesis is one of the most significant tasks in the field of natural language processing [1]–[3] and computer vision [4]–[10], which aims to synthesize visual-realistic and text-matched images from the given linguistic descriptions. With the recent success of the generative adversarial networks (GANs) [9]–[14], text-to-image synthesis has drawn increasing attention

and a great number of advanced methods [15]–[19] have been proposed.

Most approaches adopt a fashion of multi-stage generation [4], [16], [19]–[22] to obtain high-quality images, which first generates a coarse image by utilizing sentence-level textual feature and improve it to be high-resolution. Although conventional approaches are impressive in generating high-quality images, most of these approaches often synthesize mismatch images and fail to fully utilize the text, in particular when the text is complicated. One major reason for this problem is the ineffective and inadequate fusion of text and image information during the refinement process of these methods. Early works [16], [18], [20], [23]–[25] attempt to simply concatenate the encoded text information with the visual feature or utilize the attention mechanism [7], [26], [27] to integrate the cross-modal features. However, the semantic gap between different modalities seriously impedes the fusion of texts and images. For example, Attn-GAN [20] employs an attention mechanism to fuse the fine-grained word-level linguistic and visual information, which first utilizes cross-modal attention to obtain word-context features for each image sub-region. And then concatenates the word-context features with the corresponding image features to refine the initial image. However, simply concatenating features of two different modalities is sub-optimal because it cannot explicitly distinguish which regions to be refined. Recently, SD-GAN [28] adopts conditional batch normalization [29]–[33] to inject text information into the image feature maps. However, conducting batch normalization on the visual feature maps and transforming it in-

to a normal distribution may reduce the conditional representation space of the generators.

By fusing cross-modal features, the fine-grained word-level features can be used to refine the visual features, which can make full use of text information while adding more vivid details to the initial images. However, previous methods cannot fully fuse the cross-modal features, and thus synthesize images inconsistent with the given text descriptions. As shown in Fig.1, images generated by Attn-GAN [20] and DM-GAN [22] failed to match the texts semantically. In order to synthesize more text-matched images, we need to effectively inject fine-grained linguistic features into visual features. Also, the global semantics is required to encourage the generated images semantically consistent with the texts.

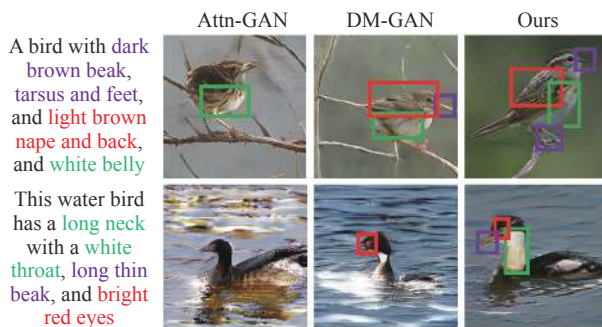


Fig. 1. Samples generated by Attn-GAN [20], DM-GAN [22] and our model on CUB-200 dataset. Images generated by our method are more realistic and semantically consistent with the given texts.

To this end, we develop a novel fine-grained text-image fusion based generative adversarial networks (FF-GAN) to improve the quality of images. Specifically, the first mechanism is the fine-grained text-image fusion block (FF-Block), which selects important word features that contains more fine-grained linguistic information and conducts affine transformation on the visual feature maps. In this way, fine-grained textual information can be added to its corresponding image sub-region in a sufficient and efficient manner. The second novel mechanism is global semantic refinement (GSR), which introduces a global semantic constraint during the refinement phase. Combing the GSR with the FF-Block in the refinement process makes it gradually and smoothly drives the generators toward the goal of fine-grained semantic alignment both globally and locally. We performed extensive experiments on two widely used and challenging benchmark datasets to verify the performance of the FF-GAN. Our FF-GAN shows remarkable superiority compared with the most advanced approaches in two evaluation metrics.

The main contributions of this paper are as follows:

- 1) We develop a novel end-to-end framework named FF-GAN, which makes full use of given textual descriptions to produce more visual-realistic and text-matched images; 2) An effective FF-Block is proposed to fuse cross-modal features more adequately and efficiently, and a GSR is developed to improve the global semantic alignment during the refinement phase.

II. Related Work

1. GAN for text-to-image synthesis

The success of GAN has greatly promoted the development of text-to-image generation, and many models based on GANs and its variants have been proposed. For the first time, Reeds *et al.* [15], [34] use conditional GAN [17], [35] to generate 64×64 fuzzy resolution images from texts. In a bid to improve the quality of synthesized pictures, Stack-GAN [16] is proposed, which synthesizes visual-authentic images with two stages by stacking a series of generators and discriminators. First, a coarse image is generated in the low-stage, and then the details of the image are modified to generate a high-resolution one. Attn-GAN [20] introduces an attention mechanism to align the semantics of texts and images. By combining semantically aligned image and text features, the generated images may maintain semantic consistency with the texts. DM-GAN [22] develops a dynamic memory based network to obtain the more significant linguistic feature at each image refinement process. DAE-GAN [36] introduces extra knowledge named aspect information to improve the details of synthesized images.

The aforementioned methods proposed many impressive mechanisms to align the feature of different modalities, which have achieved remarkable success. However, simply concatenating aligned cross-modal representations is sub-optimal for the fusion of cross-modal features, which leads to insufficient use of text semantics to generate semantic-consistent images.

In order to fuse the text and visual features more adequate, SD-GAN [28] adopts conditional batch normalization to inject text information into the visual feature maps instead of concatenating the aligned cross-modal features directly, which is more sufficient for the fusion of cross-modal features. However, the normalization of visual features will reduce the conditional representation space [37] of the generators, which is very bad for our text-image generation task.

Different from the previous methods, we propose to use affine transformation to fuse the features of different modalities. In this way, we introduce a more adequate and efficient word-level text-image fusion that enables the model to generate fine-grained images with high quality.

2. Affine transformation

Affine transformation is widely used in the case of conditional batch normalization to introduce additional conditional information [29]–[31] and avoid the loss of information [32], [33] caused by normalization.

Mani-GAN [38] introduces affine transformation in semantic image manipulation. It uses an affine transformation to fuse cross-modal representations between texts and images for effectively editing the image. MS-GAN [21] introduces an attention-modulation block to fuse the representations of text and image. DF-GAN [37] also proposed a deep text-image fusion block that stacks multiple affine transformations to make a full text-image fusion.

Although the above methods have achieved remarkable results, only fusing sentence-level semantics and visual feature maps will lead to affine transformation work on the visual representation in a spatially uniform manner and result in inadequate fusion. Ideally, the linguistic information should be incorporated into the sub-region of images with corresponding semantics. Different from the above methods, our pro-

posed model introduces more fine-grained word-level affine transformation to fully integrate the cross-modal features, which reinforces the generative model to synthesize authentic and text-matched images with more vivid details.

III. Proposed Method

As Fig.2 shows, the framework of the proposed FF-GAN mainly consists of two modules: the initial image generation and fine-grained semantic fusion based refinement. In initial image generation, we first extract semantic representations from the given textual description into sentence-level and word-level features, then we synthesize a initial low-resolution image according to the sentence-level linguistic feature. In the refinement stage, we develop an effective fine-grained text-image fusion block (FF-Block) which fuses word-level features into visual features to refine the initial image with details, while a global semantic refinement to improve the globally semantic consistency, in a bid to synthesize high-quality images that match with the corresponding textual sentences.

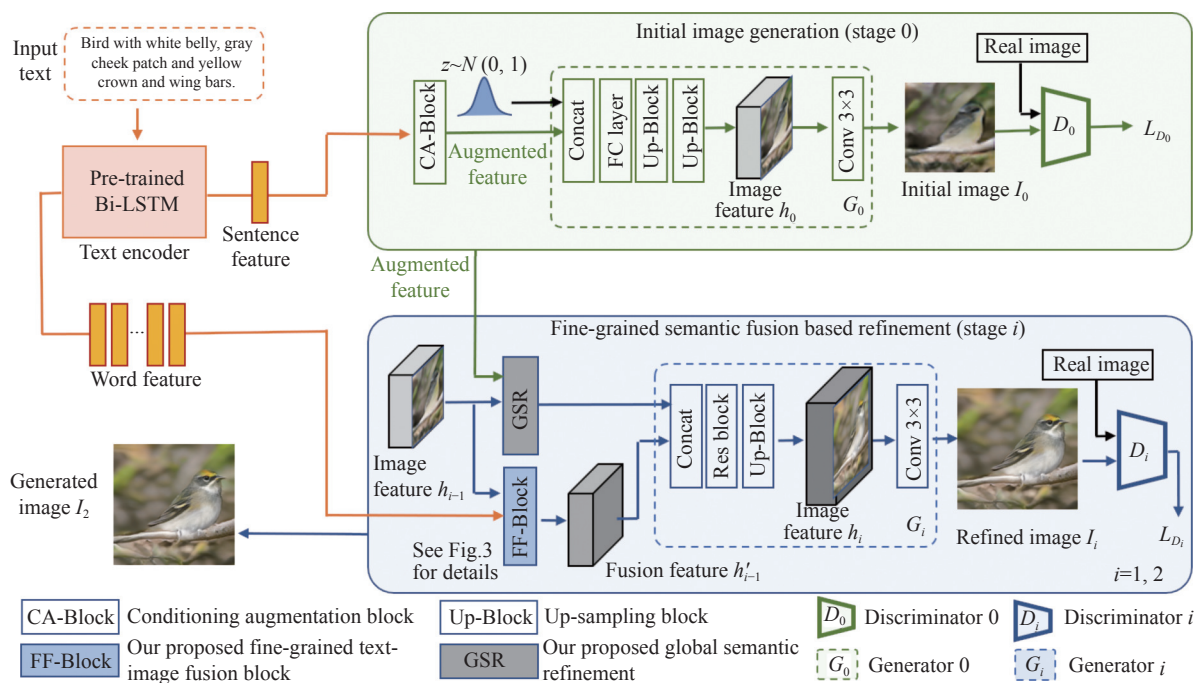


Fig. 2. Framework of our proposed FF-GAN. Our FF-GAN utilizes Bi-LSTM to extract linguistic features of two granularities, i.e., the sentence-level and word-level features, then it generates a low-resolution image by using the sentence-level features in the initial image generation (stage 0) and refines it to obtain a high-resolution image in fine-grained semantic fusion based refinement by using both the word-level features and sentence-level features (stage 1 and stage 2).

1. Initial image generation

We employ a bidirectional LSTM [1]–[3] to encode the input semantic representation into two granularities, namely sentence-level feature and the word-level feature. The sentence-level text feature S is leveraged to

synthesize the initial image in this phase and the word-level representation W is utilized to refine the initial image in the following refinement phase.

$$\{W, S\} = \text{LSTM}(T) \quad (1)$$

where $T = \{T_t | t = 0, \dots, L - 1\}$ is the text description which contains L words, $W = \{W_t | t = 0, 1, \dots, L - 1\} \in \mathbb{R}^{D_w \times L}$ is the word-level feature by concatenating the hidden states of LSTM, $S \in \mathbb{R}^{D_w}$ is the sentence-level feature from the last hidden layer of LSTM, and D_w is dimension of W_t and S .

Limited training data could lead to the sparsity in the textual conditioning manifold, so we follow StackGAN [16] by using conditioning augmentation (CA) to augment input text information. It yields more training data and thus improves the robustness of model against small perturbations. Explicitly, we enhance sentence-level linguistic feature S with conditioning augment function F_{ca} and obtain the augmented sentence semantic representation S_{ca} as follows:

$$S_{ca} = F_{ca}(S) \quad (2)$$

Then, we employ the augmented linguistic feature S_{ca} and a noise vector $\mathbf{z} \sim N(0, 1)$ that is stochastically sampled from a normal distribution to synthesize an initial image. Formally, we can obtain the initial image $I_0 \in \mathbb{R}^{D_m \times D_m}$ and its corresponding image feature $h_0 \in \mathbb{R}^{D_m \times D_m}$ as follows:

$$\{I_0, h_0\} = G_0(S_{ca}, \mathbf{z}) \quad (3)$$

where G_0 is the generator which is composed of a fully connected layer and multiple up-sampling layers.

2. Fine-grained semantic fusion based refinement

The image generated in the initial generation stage is coarse with only a rough shape and always ignores some important details of the text description. In this stage, we propose an FF-Block to fuse word-level features into visual feature maps, which contributes to refining the initial images with more fine-grain linguistic information. Also, we develop a GSR to improve the globally semantic consistency of cross-modal features.

1) Fine-grained semantic-aware fusion

In this section, we propose a novel FF-Block, as can be seen from Fig.3, which contains an attention block and several convolution layers, to fuse the fine-grained word-level feature into the image feature maps for refining the initial images with more vivid details.

Each word in the textual description possesses different importance, so it is sub-optimal to fuse the word-level linguistic feature and image feature directly. We employ an attention mechanism to obtain an attentive word-context feature, which contributes to determining the importance of every word for cross-modal fusion. Specifically, we firstly utilize a preception layer U_w to map the word-level linguistic feature W_i into the




Fig. 3. The architecture of our proposed FF-Block, which aligns the word feature and image feature by an attention block, and utilizes the affine transformation to fuse text and image features effectively.

identical latent space of visual feature. Then we compute an attention score between the word-level feature W_i and image feature h_0 by the softmax function. We obtain the attentive word-context linguistic feature f_{att}^w by conducting the inner product between $U_w W_i$ and the attention score as follows:

$$f_{att}^w = \sum_{i=0}^{L-1} (U_w W_i (\text{softmax}(h_0^T (U_w W_i))))^T \quad (4)$$

where $U_w \in \mathbb{R}^{D_m \times D_w}$ and $f_{att}^w \in \mathbb{R}^{D_m \times D_m}$, and the dimension of attentive word-context feature f_{att}^w is the same as visual feature h_0 .

To integrate linguistic and visual features efficiently, we utilize the attentive word-context feature f_{att}^w to conduct affine transformation on the visual feature h_0 . To be specific, we adopt several convolution layers to process the attentive word-context feature f_{att}^w and then predict the linguistic-conditional channel-wise scaling matrix $M(f_{att}^w)$ and shifting matrix $B(f_{att}^w)$ that both have the same size as h_0 . Finally, we obtain a cross-modal fusion representation h'_0 by fusing the word-context feature f_{att}^w and image feature h_0 as follows:

$$h'_0 = h_0 \odot M(f_{att}^w) + B(f_{att}^w) \quad (5)$$

where $M(f_{att}^w)$ and $B(f_{att}^w)$ are the learned weights and bias term based on the attentive word-context feature f_{att}^w , and \odot is the Hadamard element-wise product. Our FF-Block is able to effectively integrate word-level text features and image features to achieve a fine-grained cross-modal fusion. We could comprehend the effect of the affine transformation in cross-modal fusion from the following two aspects:

a) Scaling matrix $M(f_{att}^w)$: By multiplying with the scaling term $M(f_{att}^w)$, it helps establish the rela-

tionship between image feature h_0 and linguistic feature f_{att}^w , while re-weight the visual feature h_0 . Sub-regions of image feature h_0 that match with the word-context feature f_{att}^w will be precisely highlighted, otherwise it will be weakened, which plays a role in regional selection. In this way, FF-Block can accurately identify the attributes in the image that match the word-context semantics and establish the correlation between attributes and word-context semantics, so as to encourage the refined image of better image-text consistency.

b) Shifting matrix $\mathbf{B}(f_{att}^w)$: The bias term $\mathbf{B}(f_{att}^w)$ can effectively encode the text information and introduce the details in text descriptions which are ignored by initial images. Meanwhile, as mentioned in Feat-Trans [39], the bias term can achieve similar effects as the implementation of concatenation, which contributes to taking full advantage of linguistic information while retaining the invariant information in the visual features.

By conducting the above affine transformations, the word-context features can be efficiently fused into the visual feature maps to realize the fine-grained modification of the initial coarse images, and the refined images will contain more vivid details and be better consistent with the textual descriptions.

2) Global semantic refinement

We have introduced an FF-Block to effectively fuse cross-modal features to refine the initial images in the previous part. However, such a local fusion may lead to an inconsistency between the synthesized images and the linguistic semantics from the global semantic perspective. Some words with a higher importance in the descriptions may even affect the direction of the whole image refinement process, while those relatively less important word-context features may be ignored.

In order to maintain the global semantic consistency, we introduce the GSR inspired by Mirror-GAN [19] which exploits the sentence-level attention mechanism. The augmented linguistic feature S_{ca} is first mapped into the same latent space as the visual feature by a perceptron layer $U_s \in \mathbb{R}^{D_m \times D_w}$. Then it conducts the softmax function on the visual feature to attain the similarity score between the cross-modal features. Specifically, we could obtain a sentence-context feature f_{att}^S as follows:

$$f_{att}^S = (U_s S_{ca})(\text{softmax}(h_0^T (U_s S_{ca})))^T \quad (6)$$

To this end, the sentence-context feature f_{att}^S will be concatenated with cross-modal fusion representation h'_0 and then fed into the i th ($i = 1, 2$) generator G_i to synthesize a refined high-resolution image I_i and its corresponding visual feature h_i as follows:

$$\{I_i, h_i\} = G_i(h'_0, f_{att}^S) \quad (7)$$

where the generator G_i integrates a series of up-sampling layers and a residual block.

It is worthy to note that the combination of sentence-context features and word-level fusion representations can not only ensure the consistency of visual-linguistic semantics, but also help smoothen the refinement progress from both global and local perspectives.

3. Objective functions

To generate visual-realistic images and maintain the cross-modal semantic consistency, we design the final objective function of generators which contains two practical adversarial losses: the unconditional loss and the conditional loss. To be specific, the loss function of the generator G_i ($i = 0, 1, 2$) can be expressed as

$$L_{G_i} = \underbrace{-\frac{1}{2}\mathbb{E}_{x \sim p_{G_i}}[\log(D_i(x))]}_{\text{Unconditional loss}} - \underbrace{\frac{1}{2}\mathbb{E}_{x \sim p_{G_i}}[\log(D_i(x, S_{ca}))]}_{\text{Conditional loss}} \quad (8)$$

where x is the generated image sampled from the distribution p_{G_i} in the i th stage, and S_{ca} is the augmented linguistic features. The first part is an unconditional loss that estimates the input image to be visually authentic or fake, while another one is the conditional loss, which is utilized to discriminate whether the image and input text is a correct match.

Following the common practice [19], [20], [22], we further adopt the practical DAMSM [20] loss to compute the alignment degree between linguistic and visual features, denoted as L_{DAMSM} . To smoothen the conditional manifold and avoid over-fitting [16], we also employ a regularization term L_{CA} during training, which is defined as the Kullback-Leibler divergence (KL divergence) between the standard Gaussian distribution and the conditioning Gaussian distribution of the texts. Mathematically, the L_{CA} is defined as follows:

$$L_{CA} = D_{KL}(\mathcal{N}(\mu(S_{ca}), \Sigma(S_{ca})) || \mathcal{N}(0, 1)) \quad (9)$$

where $\mu(S_{ca})$ is the mean and $\Sigma(S_{ca})$ is the diagonal covariance matrix of the linguistic feature, which are computed by a perception layer. Finally, we define the final loss function of the generative neural networks as follows:

$$L_G = \sum_i L_{G_i} + \lambda_1 L_{CA} + \lambda_2 L_{\text{DAMSM}} \quad (10)$$

Similarly, the adversarial loss of discriminator D_i also contains an unconditional loss and a conditional loss. Mathematically, the loss function of discriminator D_i can be defined as

$$L_{D_i} = -\frac{1}{2} \underbrace{[\mathbb{E}_{\hat{x} \sim p_{\text{data}}} [\log(D_i(\hat{x}))] + \mathbb{E}_{x \sim p_{G_i}} [\log(1 - D_i(x))]}_{\text{Unconditional loss}} + \underbrace{[\mathbb{E}_{\hat{x} \sim p_{\text{data}}} [\log(D_i(\hat{x}, S_{ca}))] + \mathbb{E}_{x \sim p_{G_i}} [\log(1 - D_i(x, S_{ca}))]}_{\text{Conditional loss}} \quad (11)$$

where the \hat{x} is the image sample from the realistic image distribution p_{data} . The final objective function of the discriminator is defined as follows:

$$L_D = \sum_i L_{D_i} \quad (12)$$

IV. Experiments

In this section, we performed extensive experiments on the CUB-200 [40] and COCO [41] benchmark datasets to evaluate the proposed FF-GAN. Section IV.1 details the datasets, evaluation metrics, and training details used in the experiments. Then we compare our FF-GAN quantitatively and qualitatively with the advanced GAN-based methods, e.g., Attn-GAN [20], Mirror-GAN [19], DM-GAN [22], etc. We also perform extensive ablation experiments on the key components of our proposed FF-GAN. Code is available at <https://github.com/haoranhfut/FF-GAN>.

1. Experiment setup

1) Datasets

We perform the experiments on two challenging open datasets, i.e., CUB-200 [40] and COCO [41]. The CUB-200 contains 200 species of bird with about 12k images. Each bird image is annotated with ten linguistic descriptions. The training set of CUB-200 includes 9k images of 150 bird categories while the test set has 50 categories with 3k images. The COCO dataset consists of 80k training images and 40k testing images. Each image is annotated with five textual sentences in the COCO dataset.

2) Evaluation metrics

Following the practice, we quantify the performance of our method and related competitors on two widely used quantitative metrics: Frechet inception distance (FID) [42] and R-precision [20]. FID evaluates the realism of the synthetic images by computing the Frechet distance between the visual feature distribution of the synthesized and authentic images, which are extracted by an inception-V3 network [5] pre-trained by Attn-GAN [20]. A lower FID score suggests a higher realism of the generated images. Following Attn-GAN [20], we use R-precision to evaluate whether the synthesized pictures are well-conditioned on the corresponding textual sentences. Specifically, we compute the cosine similarities between a query visual feature and 100 candidate linguistic features extracted from R matched text descriptions and $100 - R$ stochastically selected descriptions in the dataset. Then we rank the results and

retrieve the top- R matched sentences to get the R-precision score. In practice, we set $R = 1$ in our experiment. Higher R-precision values suggests that the synthetic image is much semantically consistent with the corresponding textual description.

To compute the FID score and R-precision score, we randomly select textual captions from the test set to generate 30,000 images from each model, with each image of 256×256 resolution.

3) Implementation details

Our proposed FF-GAN first generates a 64×64 image in the initial generation stage, then refines the initial image to 128×128 and 256×256 resolution. Note that we only repeat the refinement process twice for the GPU memory limitation. Following DM-GAN [22], we apply spectral normalization for all discriminators to enhance the stability of the training process and improve performance. We utilize a pre-trained bidirectional LSTM [1], [2] by Attn-GAN [20] to produce the sentence-level and word-level representations. Meanwhile, we set the dimension of word-level vector to 256, the dimension of the augmented sentence vector is 100 and the sentence length is 18 for the balance of performance and calculation. Following the Attn-GAN [20] and DM-GAN [22], we set the hyper-parameter λ_1 to 1, λ_2 is 5 on the CUB-200 and 50 on the COCO respectively. We adopt the Adam optimizer with a learning rate set to 0.0002 during training on one Geforce GTX 1080Ti GPU. Then, we train our proposed FF-GAN with 600 epochs on CUB-200 and 120 epochs on COCO.

2. Quantitative results

We compare our method quantitatively with several advanced methods on the test set of CUB-200 and COCO, the performance results on two evaluation metrics are reported in Tables 1 and 2 [43]–[46]. It is worth mentioning that recent approaches often employ extra supervisions, e.g., KT-GAN [44] uses the extra teacher network, and TIME [46] uses extra 2-D positional encoding.

It is illustrated that our model achieves competitive performance compared with other methods in all evaluation metrics, especially in the R-precision score. Our method remarkably increases the R-precision score by a large margin from 72.31 to 80.49 on the CUB-200 as compared with DM-GAN [22]. As for the more challenging COCO dataset with multiple objects and complex backgrounds, our FF-GAN also outperforms other approaches and achieves a 91.28 R-precision score by a 1.59 numerical improvement compared with Huang *et*

Table 1. R-precision (higher is better) on the test set of CUB-200 and COCO

Method	CUB-200	COCO
Attn-GAN [20]	67.82 ± 4.43	72.31 ± 0.91
Mirror-GAN [19]	57.67	74.52
RiFe-GAN [11]	23.80 ± 1.5	–
Control-GAN [43]	69.33 ± 3.23	82.43 ± 2.43
DM-GAN [22]	72.31 ± 0.91	88.56 ± 0.28
KT-GAN* [44]	32.90	24.50
Huang <i>et al.</i> [45]	–	89.69 ± 4.34
TIME* [46]	71.57 ± 1.2	89.57 ± 0.9
Ours	80.49 ± 0.50	91.28 ± 0.51

Note: The two best scores are marked with red and blue color.

* means using extra supervisions.

Table 2. FID scores (lower is better) on the test set of CUB-200 and COCO

Method	CUB-200	COCO
HD-GAN [23]	18.23	75.34
Attn-GAN [20]	23.98	35.49
Mirror-GAN [19]	18.34	34.71
DM-GAN [22]	16.09	32.64
KT-GAN* [44]	17.32	30.73
Huang <i>et al.</i> [45]	–	34.52
TIME* [46]	14.30	31.14
Ours	15.13	29.44

Note: The two best scores are marked with red and blue color.

* means using extra supervisions.

al. [45], and a 2.72 numerical improvements compared with DM-GAN [22]. The superior performance of our method demonstrates that our FF-GAN is able to synthesize images that are more semantically consistent with the given texts.

Table 2 shows the performance on the CUB-200 and COCO datasets with respect to the FID score. As can be seen, the FID score of our FF-GAN is 15.13, which is only inferior to the result given by TIME [46] that uses extra 2-D positional encoding, but much better than other advanced models: 23.98 in Attn-GAN [20], 18.34 in Mirror-GAN [19] and 16.09 in DM-GAN [22]. As for the more challenging COCO dataset, our method shows remarkable superiority over all advanced approaches and decreases the FID score to 29.44, which indicates our FF-GAN is able to synthesize high-quality pictures with multiple complex sub-objects in highly realistic.

3. Qualitative results

1) Subjective visual comparisons

In order to qualitatively evaluate the visual quality of our proposed FF-GAN, we compare our qualitative results with the most advanced approaches including Attn-GAN [20] and DM-GAN [22]. As Fig.4(a) and Fig.4(b) shows, the images synthesized by FF-GAN contain more fine-grained details and are more se-

mantically alignment with the given textual descriptions.

Observing the visual results on the CUB-200 dataset in Fig.4(a), our method provides more authentic results than other models. Benefitting from the FF-Block, our model is able to take full advantage of the textual descriptions and mine more information. For instance, as can be seen in the 4th column of Fig.4(a), the given text “this small bird has a yellow breast, brown crown, and white superciliary”, the image generated by our method contains all attributes, especially the attribute “white superciliary” which is neglected by others. In the 5th column, the detailed attribute “short, roundish” in the text is perfectly represented by our generated image. However, Attn-GAN [20] and DM-GAN [22] failed to reflect the “roundish” and generate a slender bird. We could also observe some weakness of these methods in some cases. DM-GAN [22] generates an image in the 6th column that mismatched the attribution “speckled chest of brown and white” and the bird generated by Attn-GAN [20] in the 6th column is not photo-realistic. The qualitative results indicate that our FF-GAN fuses the visual and linguistic features in a more effective and sufficient manner, which helps synthesize high-quality images that are semantically matched with the conditioned texts.

The results on the COCO dataset are shown in Fig.4(b). In the 1st column, cows generated by our method are clearly recognized and separated, while the ones synthesized by Attn-GAN [20] and DM-GAN [22] are mixed together and difficult to distinguish. These qualitative results on COCO dataset demonstrate that our FF-GAN could synthesize visual-realistic pictures containing multiple complicated objects.

2) Low-to-high resolution synthesis

Fig.5 shows the images generated by our model in different stages. It can be seen that the pictures generated in the initial stage are very fuzzy and lose a lot of textual information. In the refinement progress, our model fully excavates the detailed linguistic information and fuses it into the visual features by FF-Block. So that the refined images can be matched with the input texts and generate more visual-realistic images. For example, as can be seen from the 1st column in Fig.5, the initial image with 64×64 resolution is a mass of a yellow body, which merely capture the attribute “yellow”. The refinement process helps to encode the attributes “black bill”, “white belly” and generate high-resolution images containing missed linguistic information.

3) Visual analysis on multi-stage refinement

In order to understand how our proposed FF-GAN utilizes word-level linguistic features to refine the initial images, we compute the attention weights between




Fig. 4. Qualitative comparison between our method and advanced AttnGAN, DM-GAN on (a) CUB-200 dataset and (b) COCO datasets.

word-level textual features and the image features, and visualize the intermediate attention maps of the word-level features in Fig.6. The 2nd and 3rd columns are the attention maps of some representative words at the first and second refinement stage respectively (stage 1 and stage 2 in Fig.6).

As can be seen from Fig.6, attention weights will

be more allocated to words related to the generated images (the bright area in Fig.6), such as “white”, “throat”, and other attributes that describe body parts and colors. By contrast, the irrelevant words will be assigned with less attention and displayed in black on the attention maps, such as the word “has”. In this way, by guiding the generator to focus on the most relevant




Fig. 5. Images of different stages generated by our FF-GAN on CUB and COCO datasets.

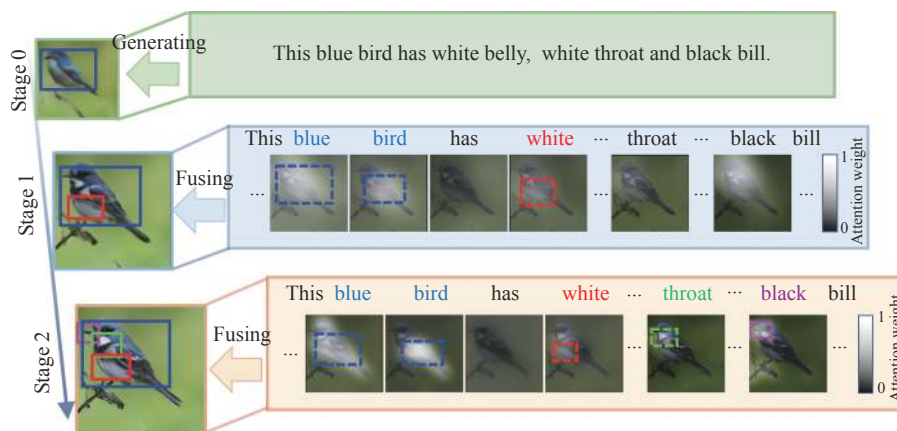


Fig. 6. Visualization of the multi-stage generation process on the CUB-200. Our FF-GAN first generates a low-resolution image by sentence-level features (stage 0), and exploits word-level features and additional sentence-level features to obtain a high-resolution image with more detailed information by two refinement stages (stage 1 and 2).

words in the given texts during the refinement process, it helps the generative model refine the generated images with more fine-grained details.

At the same time, by comparing the attention maps in the 2nd column and the 3rd column, we could find that our model can refine the images of the previous stage to be more consistent with the given texts. For example, the attention map of the attribute “black” incorrectly highlights the whole head of the bird at stage 1, which results in generating the bad shape of the “bill” in the generated image. In stage 2 (the 3rd column), the attention map of the word “black” highlights the area corresponding to the bill of the bird, which indicates that our model successfully learned this attribute at the 2nd refinement stage and refine the image finely, while the generated image also proves that our model learns the above detailed information correctly.

4. Ablation studies

1) Ablation studies on key components

We perform a series of ablation experiments on the CUB-200 dataset to evaluate the contribution of the key components of our FF-GAN. We define a baseline that removes FF-Block and GSR from FF-GAN, and some variants of baseline, such as “baseline+FF-Block” and “baseline+GSR”. The performance of the baseline and its variants is reported in Fig.7.

Comparing the yellow curve (baseline) and green curve (baseline+FF-Block) in Fig.7, we could draw a conclusion that FF-Block is able to ensure the textual-visual alignment and synthesize more visual-authentic pictures that more semantically match the given descriptions. FF-Block decreases the FID from 23.32 to 17.06 and increases R-precision by a large margin, from 67.82 to 78.52. The great promotion in R-precision indicates that FF-Block is able to fully integrate the fine-

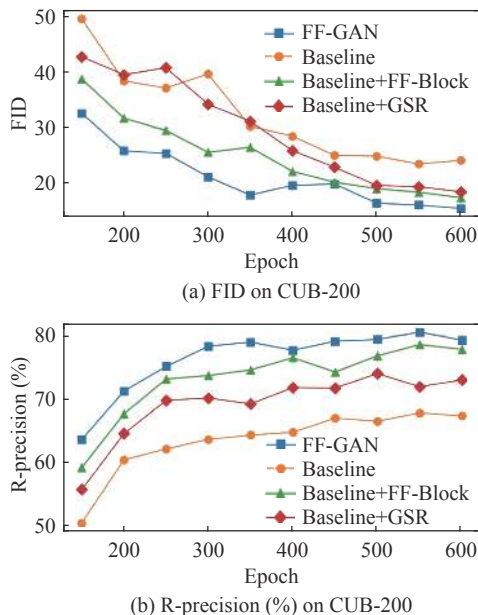


Fig. 7. FID and R-precision by our defined baseline and its variants at different epochs on CUB-200 test sets.

grained linguistic information into the visual features. Note that the R-precision of “baseline+FF-Block” is higher than R-precision of DM-GAN [22] and the FID exceeded most advanced methods, which indicate that the FF block is of great help in generating authentic and text-matched images.

The proposed GSR provides attentive sentence-level information to our FF-GAN in the refinement process. As can be seen from Fig.7, the baseline combined with GSR achieves 2.90 and 4.85 numerical improvements in FID and R-precision, respectively, which demonstrates that GSR is of great significance in improving the global semantic consistency between the synthesis images and the given texts. Moreover, the combination of GSR and FF-Block indicates that the two components can significantly reinforce the performance of the generative model.

2) Parameter studies

The hyperparameters λ_1 and λ_2 in equation (10) are used to balance each loss, and the parameter setting makes no difference on the specific methods, so we adopted the typical setting following Attn-GAN [20] and DM-GAN [22]. Specifically, the parameter λ_1 is used to balance the augmented data and real data, which is set to 1 by convention. The parameter λ_2 is used to control the text-image consistency and we set λ_2 to 5 on CUB-200 and 50 on COCO since λ_2 increases as the dataset becomes more complex.

To further verify the effectiveness of our parameter setting, we also conduct a series of verification experiments. Taking λ_2 on the CUB-200 dataset as an example, we set λ_1 to 1 by convention and set λ_2 from 1 to 10, with an interval of 1, and compare the results of

FID and R-precision. As can be seen from Fig.8, the FID score decreases first and then increases as the λ_2 becomes larger, while the R-precision is opposite. Setting λ_2 to 5 can achieve the minimal FID score and maximal R-precision, where our model achieves the best performance. To sum up, the above experiments prove the effectiveness of our parameter setting.

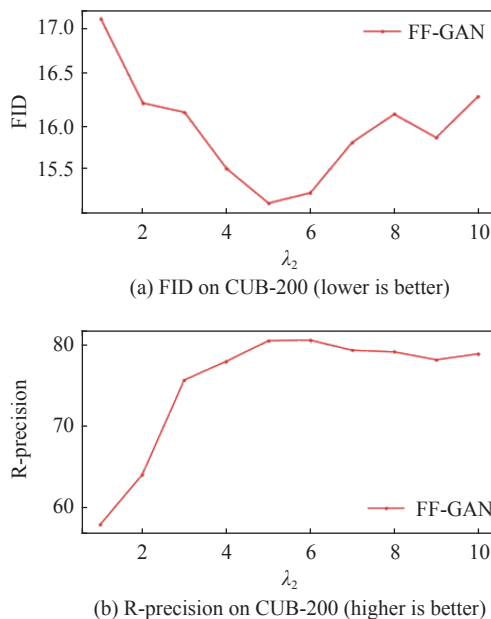


Fig. 8. FID and R-precision by different λ_2 settings (with an interval of 1) on CUB-200 test sets..

V. Conclusions

In this paper, we develop a novel fine-grained text-image fusion based generative adversarial networks (FF-GAN) to synthesize images conditioned on textual descriptions. We design an effective fine-grained text-image fusion block (FF-Block) to fully fuse visual and linguistic features, which helps synthesize more realistic and semantic consistent images based on the given linguistic descriptions. A global semantic refinement (GSR) is employed to strengthen the semantic consistency of texts and pictures from a global semantic perspective. Both qualitative and quantitative experiments on two real-world benchmark datasets indicate that our proposed FF-GAN is superior to other advanced approaches in synthesizing visual-authentic images that match the texts.

References

- [1] M. Schuster and K. K. Paliwal, “Bidirectional recurrent neural networks,” *IEEE Transactions on Signal Processing*, vol.45, no.11, pp.2673–2681, 1997.
- [2] T. Mikolov, M. Karafiát, L. Burget, *et al.*, “Recurrent neural network based language model,” in *Proceedings of the*

- INTERSPEECH 2010, 11th Annual Conference of the International Speech Communication Association*, Makuhari, Japan, pp.1045–1048, 2010.
- [3] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol.9, no.8, pp.1735–1780, 1997.
- [4] Y. Wang, W. J. Zhang, L. Wu, *et al.*, “Iterative views agreement: An iterative low-rank based structured optimization method to multi-view spectral clustering,” in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, New York, NY, USA, pp.2153–2159, 2016.
- [5] C. Szegedy, V. Vanhoucke, S. Ioffe, *et al.*, “Rethinking the inception architecture for computer vision,” in *Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, NV, USA, pp.2818–2826, 2016.
- [6] Y. Wang, “Survey on deep multi-modal data analytics: Collaboration, rivalry, and fusion,” *ACM Transactions on Multimedia Computing, Communications, and Applications*, vol.17, no.1s, article no.10, 2021.
- [7] B. Qian, Y. Wang, R. C. Hong, *et al.*, “Diversifying inference path selection: Moving-mobile-network for landmark recognition,” *IEEE Transactions on Image Processing*, vol.30, pp.4894–4904, 2021.
- [8] B. Qian, Y. Wang, H. Z. Yin, *et al.*, “Switchable online knowledge distillation,” in *Proceedings of the 17th European Conference on Computer Vision*, Tel Aviv, Israel, pp.449–466, 2022.
- [9] L. Wu, Y. Wang, and L. Shao, “Cycle-consistent deep generative hashing for cross-modal retrieval,” *IEEE Transactions on Image Processing*, vol.28, no.4, pp.1602–1612, 2019.
- [10] H. P. Liu, Y. Wang, M. Wang, *et al.*, “Delving globally into texture and structure for image inpainting,” in *Proceedings of the 30th ACM International Conference on Multimedia*, Lisboa, Portugal, pp.1270–1278, 2022.
- [11] J. Cheng, F. X. Wu, Y. L. Tian, *et al.*, “RiFeGAN: Rich feature generation for text-to-image synthesis from prior knowledge,” in *Proceedings of 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Seattle, WA, USA, pp.10908–10917, 2020.
- [12] S. He, W. T. Liao, M. Y. Yang, *et al.*, “Context-aware layout to image generation with enhanced object appearance,” in *Proceedings of 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Nashville, TN, USA, pp.15044–15053, 2021.
- [13] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, *et al.*, “Generative adversarial nets,” in *Proceedings of the 27th International Conference on Neural Information Processing Systems*, Montreal, Canada, pp.2672–2680, 2014.
- [14] B. Qian, Y. Wang, R. C. Hong, *et al.*, “Rethinking data-free quantization as a zero-sum game,” *arXiv preprint*, arXiv:2302.09572, 2023.
- [15] S. E. Reed, Z. Akata, X. C. Yan, *et al.*, “Generative adversarial text to image synthesis,” in *Proceedings of the 33rd International Conference on Machine Learning*, New York City, NY, USA, pp.1060–1069, 2016.
- [16] H. Zhang, T. Xu, H. S. Li, *et al.*, “StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks,” in *Proceedings of 2017 IEEE International Conference on Computer Vision*, Venice, Italy, pp.5908–5916, 2017.
- [17] T. Miyato and M. Koyama, “cGANs with projection discriminator,” in *Proceedings of the 6th International Conference on Learning Representations*, Vancouver, BC, Canada, pp.2337–2346, 2018.
- [18] J. Y. Zhu, T. Park, P. Isola, *et al.*, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *Proceedings of 2017 IEEE International Conference on Computer Vision*, Venice, Italy, pp.2242–2251, 2017.
- [19] T. T. Qiao, J. Zhang, D. Q. Xu, *et al.*, “MirrorGAN: Learning text-to-image generation by redescription,” in *Proceedings of 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Long Beach, CA, USA, pp.1505–1514, 2019.
- [20] T. Xu, P. C. Zhang, Q. Y. Huang, *et al.*, “AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks,” in *Proceedings of 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Salt Lake City, UT, USA, pp.1316–1324, 2018.
- [21] F. L. Mao, B. P. Ma, H. Chang, *et al.*, “MS-GAN: Text to image synthesis with attention-modulated generators and similarity-aware discriminators,” in *Proceedings of the 30th British Machine Vision Conference*, Cardiff, UK, article no.150, 2019.
- [22] M. F. Zhu, P. B. Pan, W. Chen, *et al.*, “DM-GAN: Dynamic memory generative adversarial networks for text-to-image synthesis,” in *Proceedings of 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Long Beach, CA, USA, pp. 5795–5803, 2019.
- [23] Z. Z. Zhang, Y. P. Xie, and L. Yang, “Photographic text-to-image synthesis with a hierarchically-nested adversarial network,” in *Proceedings of 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Salt Lake City, UT, USA, pp.6199–6208, 2018.
- [24] Z. X. Zhang and L. Schomaker, “DTGAN: Dual attention generative adversarial networks for text-to-image generation,” in *Proceedings of 2021 International Joint Conference on Neural Networks*, Shenzhen, China, pp.1–8, 2021.
- [25] H. Zhang, T. Xu, H. S. Li, *et al.*, “StackGAN++: Realistic image synthesis with stacked generative adversarial networks,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol.41, no.8, pp.1947–1962, 2019.
- [26] A. Vaswani, N. Shazeer, N. Parmar, *et al.*, “Attention is all you need,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Long Beach, CA, USA, pp.6000–6010, 2017.
- [27] Y. Wang, J. J. Peng, H. B. Wang, *et al.*, “Progressive learning with multi-scale attention network for cross-domain vehicle re-identification,” *Science China Information Sciences*, vol.65, no.6, article no.160103, 2022.
- [28] G. J. Yin, B. Liu, L. Sheng, *et al.*, “Semantics disentangling for text-to-image generation,” in *Proceedings of 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Long Beach, CA, USA, pp.2322–2331, 2019.
- [29] E. Perez, F. Strub, H. De Vries, *et al.*, “FiLM: Visual reasoning with a general conditioning layer,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, New Orleans, LA, USA, pp.3942–3951, 2018.
- [30] A. El-Nouby, S. Sharma, H. Schulz, *et al.*, “Tell, draw, and repeat: Generating and modifying images based on continual linguistic instruction,” in *Proceedings of 2019 IEEE/CVF International Conference on Computer Vision*, Seoul, Korea (South), pp.10303–10311, 2019.
- [31] T. Park, M. Y. Liu, T. C. Wang, *et al.*, “Semantic image synthesis with spatially-adaptive normalization,” in *Proceedings of 2019 IEEE/CVF Conference on Computer Vis-*

- ion and Pattern Recognition*, Long Beach, CA, USA, pp.2332–2341, 2019.
- [32] H. de Vries, F. Strub, J. Mary, *et al.*, “Modulating early visual processing by language,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Long Beach, CA, USA, pp.6597–6607, 2017.
- [33] X. Huang and S. Belongie, “Arbitrary style transfer in real-time with adaptive instance normalization,” in *Proceedings of 2017 IEEE International Conference on Computer Vision*, Venice, Italy, pp.1510–1519, 2017.
- [34] S. E. Reed, Z. Akata, S. Mohan, *et al.*, “Learning what and where to draw,” in *Proceedings of the 30th International Conference on Neural Information Processing Systems*, Barcelona, Spain, pp.217–225, 2016.
- [35] A. Nguyen, J. Clune, Y. Bengio, *et al.*, “Plug & play generative networks: Conditional iterative generation of images in latent space,” in *Proceedings of 2017 IEEE Conference on Computer Vision and Pattern Recognition*, Honolulu, HI, USA, pp.3510–3520, 2017.
- [36] S. L. Ruan, Y. Zhang, K. Zhang, *et al.*, “DAE-GAN: Dynamic aspect-aware GAN for text-to-image synthesis,” in *Proceedings of 2021 IEEE/CVF International Conference on Computer Vision*, Montreal, QC, Canada, pp.13940–13949, 2021.
- [37] M. Tao, H. Tang, F. Wu, *et al.*, “DF-GAN: A simple and effective baseline for text-to-image synthesis,” in *Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, New Orleans, LA, USA, pp.16494–16504, 2022.
- [38] B. W. Li, X. J. Qi, T. Lukasiewicz, *et al.*, “ManiGAN: Text-guided image manipulation,” in *Proceedings of 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Seattle, WA, USA, pp.7877–7886, 2020.
- [39] V. Dumoulin, E. Perez, N. Schucher, *et al.*, “Feature-wise transformations: A simple and surprisingly effective family of conditioning mechanisms,” Available at: <https://distill.pub/2018/feature-wise-transformations/>, 2018-07-09.
- [40] C. Wah, S. Branson, P. Welinder, *et al.*, California Institute of Technology, “The Caltech-UCSD birds-200-2011 dataset,” *Computation & Neural Systems Technical Report*, 2010-001, <https://authors.library.caltech.edu/27452/>, 2011.
- [41] T. Y. Lin, M. Maire, S. J. Belongie, *et al.*, “Microsoft COCO: Common objects in context,” in *Proceedings of the 13th European Conference on Computer Vision*, Zurich, Switzerland, pp.740–755, 2014.
- [42] M. Heusel, H. Ramsauer, T. Unterthiner, *et al.*, “GANs trained by a two time-scale update rule converge to a local Nash equilibrium,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Long Beach, CA, USA, pp.6629–6640, 2017.
- [43] B. W. Li, X. J. Qi, T. Lukasiewicz, *et al.*, “Controllable text-to-image generation,” in *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Vancouver, BC, Canada, article no.185, 2019.
- [44] H. C. Tan, X. P. Liu, M. Liu, *et al.*, “KT-GAN: Knowledge-transfer generative adversarial network for text-to-image synthesis,” *IEEE Transactions on Image Processing*, vol.30, pp.1275–1290, 2021.
- [45] W. M. Huang, R. Y. D. Xu, and I. Oppermann, “Realistic

image generation using region-phrase attention,” in *Proceedings of the 11th Asian Conference on Machine Learning*, Nagoya, Japan, pp.284–299, 2019.

- [46] B. C. Liu, K. P. Song, Y. Z. Zhu, *et al.*, “TIME: Text and image mutual-translation adversarial networks,” in *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence*, virtual event, pp.2082–2090, 2021.



SUN Haoran received the B.E. degree in 2020. He is currently pursuing the M.S. degree at Hefei University of Technology. His research interests include artificial intelligence and computer vision. (Email: haoranhfut@gmail.com)



WANG Yang (corresponding author) received the Ph.D. degree from The University of New South Wales, Kensington, Australia, in 2015. He is currently a Huangshan Professor and Ph.D. Supervisor at Hefei University of Technology, China. He has published 90 research papers (featured 7 ESI highly cited papers, with all of them to be among top 1%) including two book chapters, most of which are (to be) appeared in the major venues, such as *Artificial Intelligence* (Elsevier), *International Journal of Computer Vision* (IJCV), *IEEE TIP*, *IEEE TNNLS*, *IEEE TMM*, *ACM TOIS*, *Machine Learning* (Springer), *IEEE TKDE*, *IEEE TCYB*, *VLDB Journal*, *KDD*, *ECCV*, *IJCAI*, *AAAI*, *ACM SIGIR*, *ACM Multimedia*, *IEEE ICDM*, *ACM CIKM*, *SCIENCE CHINA Information Sciences* etc. He currently serves as the Associate Editor of *ACM Trans. Information systems*. He was the winner of Best Research Paper Runner-up Award for PAKDD 2014, and was a program committee member for various leading conferences such as IJCAI, AAAI, CVPR, ECCV, EMNLP, ACM Multimedia, ACM Multimedia (Asia), ECM-LPKDD, etc. (Email: yangwang@hfut.edu.cn)



LIU Haipeng received the B.E. degree in 2018. He is a Ph.D. candidate at the Hefei University of Technology, Hefei, China. His current research interests include computer vision, deep learning and image inpainting. (Email: hpliu_hfut@hotmail.com)



QIAN Biao received the B.E. degree in 2017 and is a Ph.D. candidate at the Hefei University of Technology, Hefei, China. His current research interests include computer vision, deep learning and neural network compression and acceleration. (Email: hfutqian@gmail.com)