

# MirrorGAN: Learning Text-to-image Generation by Redescription

Tingting Qiao<sup>1,3</sup>, Jing Zhang<sup>2,3,\*</sup>, Duanqing Xu<sup>1,\*</sup>, and Dacheng Tao<sup>3</sup>

<sup>1</sup>College of Computer Science and Technology, Zhejiang University, China

<sup>2</sup>School of Automation, Hangzhou Dianzi University, China

<sup>3</sup>UBTECH Sydney AI Centre, School of Computer Science, FEIT, The University of Sydney, Australia  
qiaott@zju.edu.cn, jing.zhang@uts.edu.au, xdq@zju.edu.cn, dacheng.tao@sydney.edu.au

## Abstract

Generating an image from a given text description has two goals: visual realism and semantic consistency. Although significant progress has been made in generating high-quality and visually realistic images using generative adversarial networks, guaranteeing semantic consistency between the text description and visual content remains very challenging. In this paper, we address this problem by proposing a novel global-local attentive and semantic-preserving text-to-image-to-text framework called MirrorGAN. MirrorGAN exploits the idea of learning text-to-image generation by redescription and consists of three modules: a semantic text embedding module (STEM), a global-local collaborative attentive module for cascaded image generation (GLAM), and a semantic text regeneration and alignment module (STREAM). STEM generates word- and sentence-level embeddings. GLAM has a cascaded architecture for generating target images from coarse to fine scales, leveraging both local word attention and global sentence attention to progressively enhance the diversity and semantic consistency of the generated images. STREAM seeks to regenerate the text description from the generated image, which semantically aligns with the given text description. Thorough experiments on two public benchmark datasets demonstrate the superiority of MirrorGAN over other representative state-of-the-art methods.

## 1. Introduction

Text-to-image (T2I) generation refers to generating a visually realistic image that matches a given text descrip-

1. The work was performed when Tingting Qiao was a visiting student at UBTECH Sydney AI Centre in the School of Computer Science, FEIT, in the University of Sydney

2.\*corresponding author

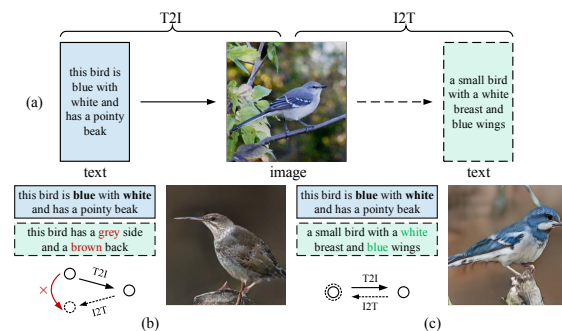


Figure 1: (a) Illustration of the mirror structure that embodies the idea of learning text-to-image generation by redescription. (b)-(c) Semantically inconsistent and consistent images/redescriptions generated by [35] and the proposed MirrorGAN, respectively.

tion. Due to its significant potential in a number of applications but its challenging nature, T2I generation has become an active research area in both natural language processing and computer vision communities. Although significant progress has been made in generating visually realistic images using generative adversarial networks (GANs) such as in [39, 42, 35, 13], guaranteeing semantic alignment of the generated image with the input text remains challenging.

In contrast to fundamental image generation problems, T2I generation is conditioned on text descriptions rather than starting with noise alone. Leveraging the power of GANs [10], different T2I methods have been proposed to generate visually realistic and text-relevant images. For instance, Reed *et al.* proposed to tackle text to image synthesis problem by finding a visually discriminative representation for the text descriptions and using this representation to generate realistic images [24]. Zhang *et al.* proposed StackGAN to generate images in two separate stages [39]. Hong *et al.* proposed extracting a semantic layout from the input text and then converting it into the image generator to guide

the generative process [13]. Zhang *et al.* proposed training a T2I generator with hierarchically nested adversarial objectives [42]. These methods all utilize a discriminator to distinguish between the generated image and corresponding text pair and the ground truth image and corresponding text pair. However, due to the domain gap between text and images, it is difficult and inefficient to model the underlying semantic consistency within each pair when relying on such a discriminator alone. Recently, the attention mechanism [35] has been exploited to address this problem, which guides the generator to focus on different words when generating different image regions. However, using word-level attention alone does not ensure global semantic consistency due to the diversity between text and image modalities. Figure 1 (b) shows an example generated by [35].

T2I generation can be regarded as the inverse problem of image captioning (or image-to-text generation, I2T) [34, 29, 16], which generates a text description given an image. Considering that tackling each task requires modeling and aligning the underlying semantics in both domains, it is natural and reasonable to model both tasks in a unified framework to leverage the underlying dual regulations. As shown in Figure 1 (a) and (c), if an image generated by T2I is semantically consistent with the given text description, its redescription by I2T should have exactly the same semantics with the given text description. In other words, the generated image should act like a mirror that precisely reflects the underlying text semantics. Motivated by this observation, we propose a novel text-to-image-to-text framework called MirrorGAN to improve T2I generation, which exploits the idea of learning T2I generation by redescription.

MirrorGAN has three modules: STEM, GLAM and STREAM. STEM generates word- and sentence-level embeddings, which are then used by the GLAM. GLAM is a cascaded architecture that generates target images from coarse to fine scales, leveraging both local word attention and global sentence attention to progressively enhance the diversity and semantic consistency of the generated images. STREAM tries to regenerate the text description from the generated image, which semantically aligns with the given text description.

To train the model end-to-end, we use two adversarial losses: visual realism adversarial loss and text-image paired semantic consistency adversarial loss. In addition, to leverage the dual regulation of T2I and I2T, we further employ a text-semantics reconstruction loss based on cross-entropy (CE). Thorough experiments on two public benchmark datasets demonstrate the superiority of MirrorGAN over other representative state-of-the-art methods with respect to both visual realism and semantic consistency.

The contributions of this work can be summarized as follows:

- We propose a novel unified framework called Mirror-

GAN for modeling T2I and I2T together, specifically targeting T2I generation by embodying the idea of learning T2I generation by redescription.

- We propose a global-local collaborative attention model that is seamlessly embedded in the cascaded generators to preserve cross-domain semantic consistency and to smoothen the generative process.
- Except commonly used GAN losses, we additionally propose a CE-based text-semantics reconstruction loss to supervise the generator to generate visually realistic and semantically consistent images. Consequently, we achieve new state-of-the-art performance on two benchmark datasets.

## 2. Related work

Similar ideas to our own have recently been used in CycleGAN and DualGAN, which handle the bi-directional translations within two domains together [43, 37, 1, 32], significantly advance image-to-image translation [14, 28, 15, 38, 23]. Our MirrorGAN is partly inspired by CycleGAN but has two main differences: 1) we specifically tackle the T2I problem rather than image-to-image translation. The cross-media domain gap between text and images is probably much larger than the one between images with different attributes, *e.g.*, styles. Moreover, the diverse semantics present in each domain make it much more challenging to maintain cross-domain semantic consistency. 2) MirrorGAN embodies a mirror structure rather than the cycle structure used in CycleGAN. MirrorGAN conducts supervised learning by using paired text-image data rather than training from unpaired image-image data. Moreover, to embody the idea of learning T2I generation by redescription, we use a CE-based reconstruction loss to regularize the semantic consistency of the redescribed text, which is different from the L1 cycle consistency loss in CycleGAN, which addresses visual similarities.

Attention models have been extensively exploited in computer vision and natural language processing, for instance in object detection [21, 6, 18, 41], image/video captioning [34, 9, 31], visual question answering [2, 33, 36, 22], and neural machine translation [19, 8]. Attention can be modeled spatially in images or temporally in language, or even both in video- or image-text-related tasks. Different attention models have been proposed for image captioning to enhance the embedded text feature representations during both encoding and decoding. Recently, Xu *et al.* proposed an attention model to guide the generator to focus on different words when generating different image subregions [35]. However, using only word-level attention does not ensure global semantic consistency due to the diverse nature of both the text and image modalities, *e.g.*, each image has 10 captions in CUB and 5 captions in COCO, however, they express the same underlying semantic

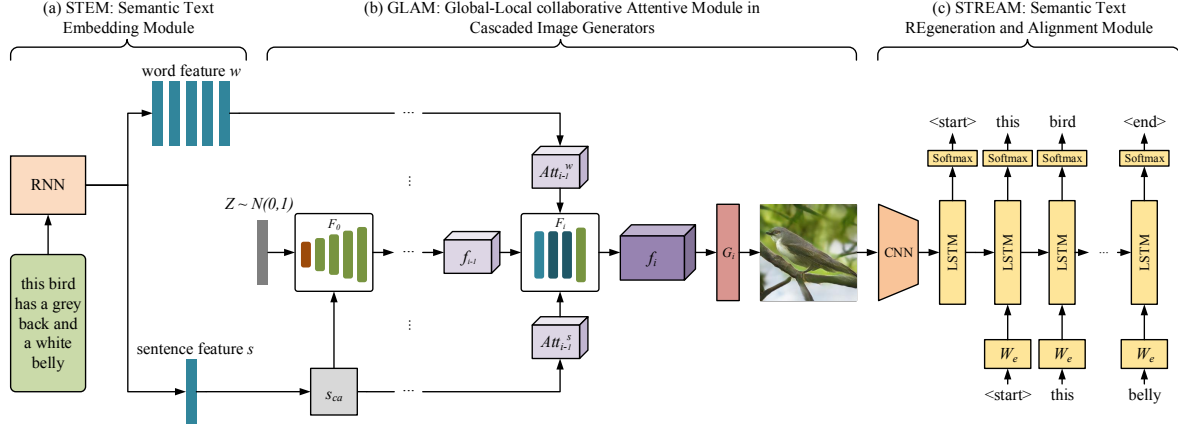


Figure 2: Schematic of the proposed MirrorGAN for text-to-image generation.

information. In particular, for multi-stage generators, it is crucial to make “semantically smooth” generations. Therefore, global sentence-level attention should also be considered in each stage such that it progressively and smoothly drives the generators towards semantically well-aligned targets. To this end, we propose a global-local collaborative attentive module to leverage both local word attention and global sentence attention and to enhance the diversity and semantic consistency of the generated images.

### 3. MirrorGAN for text-to-image generation

As shown in Figure 2, MirrorGAN embodies a mirror structure by integrating both T2I and I2T. It exploits the idea of learning T2I generation by redescription. After an image is generated, MirrorGAN regenerates its description, which aligns its underlying semantics with the given text description. Technically, MirrorGAN consists of three modules: STEM, GLAM and STREAM. Details of the model will be introduced below.

#### 3.1. STEM: Semantic Text Embedding Module

First, we introduce the semantic text embedding module to embed the given text description into local word-level features and global sentence-level features. As shown in the leftmost part of Figure 2, a recurrent neural network (RNN) [4] is used to extract semantic embeddings from the given text description  $T$ , which include a word embedding  $w$  and a sentence embedding  $s$ .

$$w, s = RNN(T), \quad (1)$$

where  $T = \{T_l | l = 0, \dots, L-1\}$ ,  $L$  represents the sentence length,  $w = \{w^l | l = 0, \dots, L-1\} \in \mathbb{R}^{D \times L}$  is the concatenation of hidden state  $w^l$  of each word,  $s \in \mathbb{R}^D$  is the last hidden state, and  $D$  is the dimension of  $w^l$  and  $s$ . Due to the diversity of the text domain, text with few permutations may share similar semantics. Therefore, we follow

the common practice of using the conditioning augmentation method [39] to augment the text descriptions. This produces more image-text pairs and thus encourages robustness to small perturbations along the conditioning text manifold. Specifically, we use  $F_{ca}$  to represent the conditioning augmentation function and obtain the augmented sentence vector:

$$s_{ca} = F_{ca}(s), \quad (2)$$

where  $s_{ca} \in \mathbb{R}^{D'}$ ,  $D'$  is the dimension after augmentation.

#### 3.2. GLAM: Global-Local collaborative Attentive Module in Cascaded Image Generators

We next construct a multi-stage cascaded generator by stacking three image generation networks sequentially. We adopt the basic structure described in [35] due to its good performance in generating realistic images. Mathematically, we use  $\{F_0, F_1, \dots, F_{m-1}\}$  to denote the  $m$  visual feature transformers and  $\{G_0, G_1, \dots, G_{m-1}\}$  to denote the  $m$  image generators. The visual feature  $f_i$  and generated image  $I_i$  in each stage can be expressed as:

$$\begin{aligned} f_0 &= F_0(z, s_{ca}), \\ f_i &= F_i(f_{i-1}, F_{att_i}(f_{i-1}, w, s_{ca})), i \in \{1, 2, \dots, m-1\}, \\ I_i &= G_i(f_i), i \in \{0, 1, 2, \dots, m-1\}, \end{aligned} \quad (3)$$

where  $f_i \in \mathbb{R}^{M_i \times N_i}$  and  $I_i \in \mathbb{R}^{q_i \times q_i}$ ,  $z \sim N(0, 1)$  denotes random noises.  $F_{att_i}$  is the proposed global-local collaborative attention model which includes two components  $Att_{i-1}^w$  and  $Att_{i-1}^s$ , i.e.,  $F_{att_i}(f_{i-1}, w, s_{ca}) = \text{concat}(Att_{i-1}^w, Att_{i-1}^s)$ .

First, we use the word-level attention model proposed in [35] to generate an attentive word-context feature. It takes the word embedding  $w$  and the visual feature  $f$  as the input in each stage. The word embedding  $w$  is first converted into an underlying common semantic space of visual features by a perception layer  $U_{i-1}$  as  $U_{i-1}w$ . Then, it is multiplied

with the visual feature  $f_{i-1}$  to obtain the attention score. Finally, the attentive word-context feature is obtained by calculating the inner product between the attention score and  $U_{i-1}w$ :

$$Att_{i-1}^w = \sum_{l=0}^{L-1} (U_{i-1}w^l) (\text{softmax}(f_{i-1}^T (U_{i-1}w^l)))^T, \quad (4)$$

where  $U_{i-1} \in \mathbb{R}^{M_{i-1} \times D}$  and  $Att_{i-1}^w \in \mathbb{R}^{M_{i-1} \times N_{i-1}}$ . The attentive word-context feature  $Att_{i-1}^w$  has the exact same dimension as  $f_{i-1}$ , which is further used for generating the  $i^{\text{th}}$  visual features  $f_i$  by concatenation with  $f_{i-1}$ .

Then, we propose a sentence-level attention model to enforce a global constraint on the generators during generation. By analogy to the word-level attention model, the augmented sentence vector  $s_{ca}$  is first converted into an underlying common semantic space of visual features by a perception layer  $V_{i-1}$  as  $V_{i-1}s_{ca}$ . Then, it is element-wise multiplied with the visual feature  $f_{i-1}$  to obtain the attention score. Finally, the attentive sentence-context feature is obtained by calculating the element-wise multiplication of the attention score and  $V_{i-1}s_{ca}$ :

$$Att_{i-1}^s = (V_{i-1}s_{ca}) \circ (\text{softmax}(f_{i-1} \circ (V_{i-1}s_{ca}))), \quad (5)$$

where  $\circ$  denotes the element-wise multiplication,  $V_i \in \mathbb{R}^{M_i \times D'}$  and  $Att_{i-1}^s \in \mathbb{R}^{M_{i-1} \times N_{i-1}}$ . The attentive sentence-context feature  $Att_{i-1}^s$  is further concatenated with  $f_{i-1}$  and  $Att_{i-1}^w$  for generating the  $i^{\text{th}}$  visual features  $f_i$  as depicted in the second equality in Eq. (3).

### 3.3. STREAM: Semantic Text REgeneration and Alignment Module

As described above, MirrorGAN includes a semantic text regeneration and alignment module (STREAM) to regenerate the text description from the generated image, which semantically aligns with the given text description. Specifically, we employ a widely used encoder-decoder-based image caption framework [16, 29] as the basic STREAM architecture. Note that a more advanced image captioning model can also be used, which is likely to produce better results. However, in a first attempt to validate the proposed idea, we simply exploit the baseline in the current work.

The image encoder is a convolutional neural network (CNN) [11] pretrained on ImageNet [5], and the decoder is a RNN [12]. The image  $I_{m-1}$  generated by the final stage generator is fed into the CNN encoder and RNN decoder as follows:

$$\begin{aligned} x_{-1} &= CNN(I_{m-1}), \\ x_t &= W_e T_t, t \in \{0, \dots, L-1\}, \\ p_{t+1} &= RNN(x_t), t \in \{0, \dots, L-1\}, \end{aligned} \quad (6)$$

where  $x_{-1} \in \mathbb{R}^{M_{m-1}}$  is a visual feature used as the input at the beginning to inform the RNN about the image content.  $W_e \in \mathbb{R}^{M_{m-1} \times D}$  represents a word embedding matrix, which maps word features to the visual feature space.  $p_{t+1}$  is a predicted probability distribution over the words. We pre-trained STREAM as it helped MirrorGAN achieve a more stable training process and converge faster, while jointly optimizing STREAM with MirrorGAN is instable and very expensive in terms of time and space. The encoder-decoder structure in [29] and then their parameters keep fixed when training the other modules of MirrorGAN.

### 3.4. Objective functions

Following common practice, we first employ two adversarial losses: a visual realism adversarial loss and a text-image paired semantic consistency adversarial loss, which are defined as follows.

During each stage of training MirrorGAN, the generator  $G$  and discriminator  $D$  are trained alternately. Specially, the generator  $G_i$  in the  $i^{\text{th}}$  stage is trained by minimizing the loss as follows:

$$\mathcal{L}_{G_i} = -\frac{1}{2} E_{I_i \sim p_{I_i}} [\log(D_i(I_i))] - \frac{1}{2} E_{I_i \sim p_{I_i}} [\log(D_i(I_i, s))], \quad (7)$$

where  $I_i$  is a generated image sampled from the distribution  $p_{I_i}$  in the  $i^{\text{th}}$  stage. The first term is the visual realism adversarial loss, which is used to distinguish whether the image is visually real or fake, while the second term is the text-image paired semantic consistency adversarial loss, which is used to determine whether the underlying image and sentence semantics are consistent.

We further propose a CE-based text-semantic reconstruction loss to align the underlying semantics between the redescription of STREAM and the given text description. Mathematically, this loss can be expressed as:

$$\mathcal{L}_{stream} = - \sum_{t=0}^{L-1} \log p_t(T_t). \quad (8)$$

It is noteworthy that  $\mathcal{L}_{stream}$  is also used during STREAM pretraining. When training  $G_i$ , gradients from  $\mathcal{L}_{stream}$  are backpropagated to  $G_i$  through STREAM, whose network weights are kept fixed.

The final objective function of the generator is defined as:

$$\mathcal{L}_G = \sum_{i=0}^{m-1} \mathcal{L}_{G_i} + \lambda \mathcal{L}_{stream}, \quad (9)$$

where  $\lambda$  is a loss weight to handle the importance of adversarial loss and the text-semantic reconstruction loss.

The discriminator  $D_i$  is trained alternately to avoid being fooled by the generators by distinguishing the inputs as either real or fake. Similar to the generator, the objective

of the discriminators consists of a visual realism adversarial loss and a text-image paired semantic consistency adversarial loss. Mathematically, it can be defined as:

$$\begin{aligned} \mathcal{L}_{D_i} = & -\frac{1}{2}\mathbb{E}_{I_i^{GT} \sim p_{I_i^{GT}}} [\log (D_i (I_i^{GT}))] \\ & -\frac{1}{2}\mathbb{E}_{I_i \sim p_{I_i}} [\log (1 - D_i (I_i))] \\ & -\frac{1}{2}\mathbb{E}_{I_i^{GT} \sim p_{I_i^{GT}}} [\log (D_i (I_i^{GT}, s))] \\ & -\frac{1}{2}\mathbb{E}_{I_i \sim p_{I_i}} [\log (1 - D_i (I_i, s))], \end{aligned} \quad (10)$$

where  $I_i^{GT}$  is from the real image distribution  $p_{I_i^{GT}}$  in  $i^{th}$  stage. The final objective function of the discriminator is defined as:

$$\mathcal{L}_D = \sum_{i=0}^{m-1} \mathcal{L}_{D_i}. \quad (11)$$

## 4. Experiments

In this section, we present extensive experiments that evaluate the proposed model. We first compare MirrorGAN with the state-of-the-art T2I methods GAN-INT-CLS [24], GAWWN [25], StackGAN [39], StackGAN++ [40], PPGN [20] and AttnGAN [35]. Then, we present ablation studies on the key components of MirrorGAN including GLAM and STREAM.

### 4.1. Experiment setup

#### 4.1.1 Datasets

We evaluated our model on two commonly used datasets, CUB bird dataset [30] and MS COCO dataset [17]. The CUB bird dataset contains 8,855 training images and 2,933 test images belonging to 200 categories, each bird image has 10 text descriptions. The COCO dataset contains 82,783 training images and 40,504 validation images, each image has 5 text descriptions. Both datasets were pre-processed using the same pipeline as in [39, 35].

#### 4.1.2 Evaluation metric

Following common practice [39, 35], the Inception Score [26] was used to measure both the objectiveness and diversity of the generated images. Two fine-tuned inception models provided by [39] were used to calculate the score.

Then, the R-precision introduced in [35] was used to evaluate the visual-semantic similarity between the generated images and their corresponding text descriptions. For each generated image, its ground truth text description and 99 randomly selected mismatched descriptions from the test set were used to form a text description pool. We then calculated the cosine similarities between the image feature and the text feature of each description in the pool, before counting the average accuracy at three different settings: top-1, top-2, and top-3. The ground truth entry falling into the top-k candidates was treated as correct, otherwise, it was

wrong. A higher score represents a higher visual-semantic similarity between the generated images and input text.

The Inception Score and the R-precision were calculated accordingly as in [39, 35].

### 4.1.3 Implementation details

MirrorGAN has three generators in total and GLAM is employed over the last two generators, as shown in Eq. (3).  $64 \times 64$ ,  $128 \times 128$ ,  $256 \times 256$  images are generated progressively. Followed [35], a pre-trained bi-directional LSTM [27] was used to calculate the semantic embedding from text descriptions. The dimension of the word embedding  $D$  was 256. The sentence length  $L$  was 18. The dimension  $M_i$  of the visual embedding was set to 32. The dimension of the visual feature was  $N_i = q_i \times q_i$ , where  $q_i$  was 64, 128, and 256 for the three stages. The dimension of augmented sentence embedding  $D'$  was set to 100. The loss weight  $\lambda$  of the text-semantic reconstruction loss was set to 20.

### 4.2. Main results

In this section, we present both qualitative and quantitative comparisons with other methods to verify the effectiveness of MirrorGAN. First, we compare MirrorGAN with state-of-the-art text-to-image methods [24, 25, 39, 40, 20, 35] using the Inception Score and R-precision score on both CUB and COCO datasets. Then, we present subjective visual comparisons between MirrorGAN and the state-of-the-art method AttnGAN [35]. We also present the results of a human study designed to test the authenticity and visual semantic similarity between input text and images generated by MirrorGAN and AttnGAN [35].

#### 4.2.1 Quantitative results

The Inception Scores of MirrorGAN and other methods are shown in Table 1. MirrorGAN achieved the highest Inception Score on both CUB and COCO datasets. Specifically, compared with the state-of-the-art method AttnGAN [35], MirrorGAN improved the Inception Score from 4.36 to 4.56 on CUB and from 25.89 to 26.47 on the more difficult COCO dataset. These results show that MirrorGAN can generate more diverse images of better quality.

The R-precision scores of AttnGAN [35] and MirrorGAN on CUB and COCO datasets are listed in Table 2. MirrorGAN consistently outperformed AttnGAN [35] at all settings by a large margin, demonstrating the superiority of the proposed text-to-image-to-text framework and the global-local collaborative attentive module, since MirrorGAN generated high-quality images with semantics consistent with the input text descriptions.



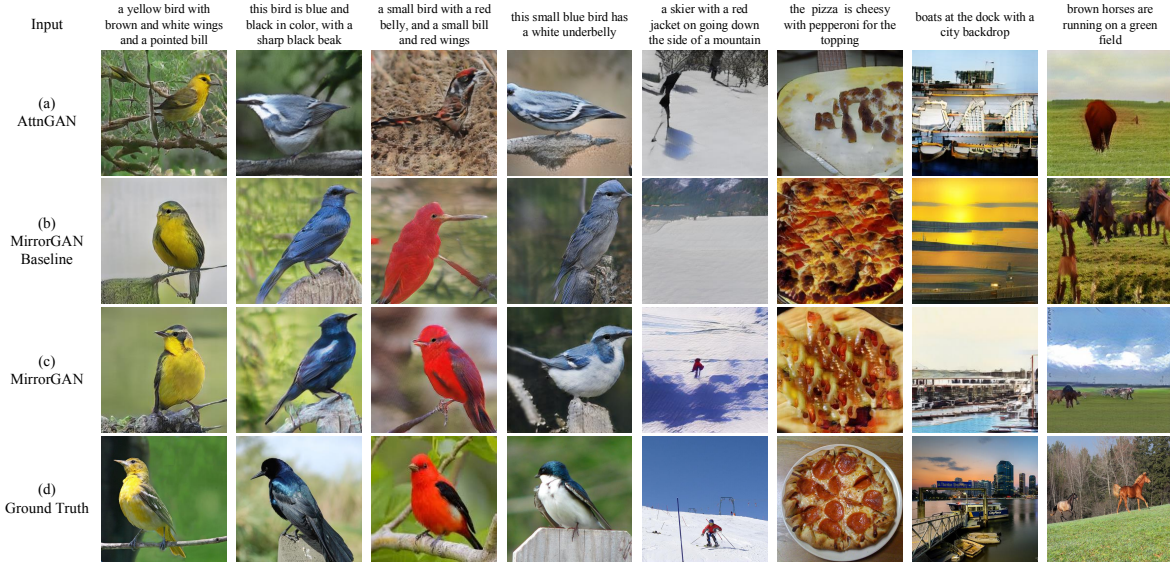


Figure 3: Examples of images generated by (a) AttnGAN [35], (b) MirrorGAN Baseline, and (c) MirrorGAN conditioned on text descriptions from CUB and COCO test sets and (d) the corresponding ground truth.

Table 1: Inception Scores of state-of-the-art methods and MirrorGAN on CUB and COCO datasets.

Model	CUB	COCO
GAN-INT-CLS [24]	$2.88 \pm 0.04$	$7.88 \pm 0.07$
GAWWN [25]	$3.62 \pm 0.07$	-
StackGAN [39]	$3.70 \pm 0.04$	$8.45 \pm 0.03$
StackGAN++ [40]	$3.82 \pm 0.06$	-
PPGN [20]	-	$9.58 \pm 0.21$
AttnGAN [35]	$4.36 \pm 0.03$	$25.89 \pm 0.47$
MirrorGAN	<b><math>4.56 \pm 0.05</math></b>	<b><math>26.47 \pm 0.41</math></b>

Table 2: R-precision [%] of the state-of-the-art AttnGAN [35] and MirrorGAN on CUB and COCO datasets.

Dataset	CUB			COCO		
	k=1	k=2	k=3	k=1	k=2	k=3
AttnGAN [35]	53.31	54.11	54.36	72.13	73.21	76.53
MirrorGAN	<b>57.67</b>	<b>58.52</b>	<b>60.42</b>	<b>74.52</b>	<b>76.87</b>	<b>80.21</b>

## 4.2.2 Qualitative results

**Subjective visual comparisons:** Subjective visual comparisons between AttnGAN [35], MirrorGAN Baseline, and MirrorGAN are presented in Figure 3. MirrorGAN Baseline refers to the model using only word-level attention for each generator in the MirrorGAN framework.

It can be seen that the image details generated by AttnGAN are lost, colors are inconsistent with the text descriptions (3<sup>rd</sup> and 4<sup>th</sup> column), and the shape looks strange (2<sup>nd</sup>, 3<sup>rd</sup>, 5<sup>th</sup> and 8<sup>th</sup> column) for some hard examples.

Furthermore, the skier is missing in the 5<sup>th</sup> column. MirrorGAN Baseline achieved better results with more details and consistent colors and shapes compared to AttnGAN. For example, the wings are vivid in the 1<sup>st</sup> and 2<sup>nd</sup> columns, demonstrating the superiority of MirrorGAN and that it takes advantage of the dual regularization by redescription, *i.e.*, a semantically consistent image should be generated if it can be redescriptioned correctly. By comparing MirrorGAN with MirrorGAN Baseline, we can see that GLAM contributes to producing fine-grained images with more details and better semantic consistency. For example, the color of the underbelly of the bird in the 4<sup>th</sup> column was corrected to white, and the skier with a red jacket was recovered. The boats and city backdrop in the 7<sup>th</sup> column and the horses on the green field in the 8<sup>th</sup> column look real at first glance. Generally, content in the CUB dataset is less diverse than in COCO dataset. Therefore, it is easier to generate visually realistic and semantically consistent results on CUB. These results confirm the impact of GLAM, which uses global and local attention collaboratively.

**Human perceptual test:** To compare the visual realism and semantic consistency of the images generated by AttnGAN and MirrorGAN, we next performed a human perceptual test on the CUB test dataset. We recruited 100 volunteers with different professional backgrounds to conduct two tests: the Image Authenticity Test and the Semantic Consistency Test. The Image Authenticity Test aimed to compare the authenticity of the images generated using different methods. Participants were presented with 100 groups of images consecutively. Each group had 2 images arranged in random order from AttnGAN and MirrorGAN

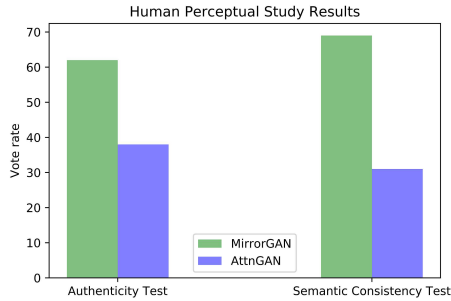


Figure 4: Results of Human perceptual test. A higher value of the Authenticity Test means more convincing images. A higher value of the Semantic Consistency Test means a closer semantics between input text and generated images.

Table 3: Inception Score and R-precision results of MirrorGAN with different weight settings.

Evaluation Metric	Inception Score		R-precision (top-1)	
	CUB	COCO	CUB	COCO
MirrorGAN w/o GA, $\lambda=0$	3.91 ± .09	19.01 ± .42	39.09	50.69
MirrorGAN w/o GA, $\lambda=20$	4.47 ± .07	25.99 ± .41	55.67	73.28
MirrorGAN, $\lambda=5$	4.01 ± .06	21.85 ± .43	32.07	52.55
MirrorGAN, $\lambda=10$	4.30 ± .07	24.11 ± .31	43.21	63.40
MirrorGAN, $\lambda=20$	<b>4.54 ± .17</b>	<b>26.47 ± .41</b>	<b>57.67</b>	<b>74.52</b>

given the same text description. Participants were given unlimited time to select the more convincing images. The Semantic Consistency Test aimed to compare the semantic consistency of the images generated using different methods. Each group had 3 images corresponding to the ground truth image and two images arranged at random from AttnGAN and MirrorGAN. The participants were asked to select the images that were more semantically consistent with the ground truth. Note that we used ground truth images instead of the text descriptions since it is easier to compare the semantics between images.

After the participants finished the experiment, we counted the votes for each method in the two scenarios. The results are shown in Figure 4. It can be seen that the images from MirrorGAN were preferred over ones from AttnGAN. MirrorGAN outperformed AttnGAN with respect to authenticity, MirrorGAN was even more effective in terms of semantic consistency. These results demonstrate the superiority of MirrorGAN for generating visually realistic and semantically consistent images.

### 4.3. Ablation studies

**Ablation studies on MirrorGAN components:** We next conducted ablation studies on the proposed model and its variants. To validate the effectiveness of STREAM and GLAM, we conducted several comparative experiments by

excluding/including these components in MirrorGAN. The results are listed in Table 3.

First, the hyper-parameter  $\lambda$  is important. A larger  $\lambda$  led to higher Inception Scores and R-precision on both datasets. On the CUB dataset, when  $\lambda$  increased from 5 to 20, the Inception Score increased from 4.01 to 4.54 and R-precision increased from 32.07% to 57.67%. On the COCO dataset, the Inception Score increased from 21.85 to 26.21 and R-precision increased from 52.55% to 74.52%. We set  $\lambda$  to 20 as default.

MirrorGAN without STREAM ( $\lambda = 0$ ) and global attention (GA) achieved better results than StackGAN++ [40] and PPGN [20]. Integrating STREAM into MirrorGAN led to further significant performance gains. The Inception Score increased from 3.91 to 4.47 and from 19.01 to 25.99 on CUB and COCO, respectively, and R-precision showed the same trend. Note that MirrorGAN without GA already outperformed the state-of-the-art AttnGAN (Table 1) which also used the word-level attention. These results indicate that STREAM is more effective in helping the generators achieve better performance. This attributes to the introduction of a more strict semantic alignment between generated images and input text, which is provided by STREAM. Specifically, STREAM forces the generated images to be redescribed as the input text sequentially, which potentially prevents possible mismatched visual-text concept. Moreover, MirrorGAN integration with GLAM further improved the Inception Score and R-precision to achieve new state-of-the-art performance. These results show that the global and local attention in GLAM collaboratively help the generator to generate visually realistic and semantically consistent results by telling it where to focus on.

**Visual inspection on the cascaded generators:** To better understand the cascaded generation process of MirrorGAN, we visualized both the intermediate images and the attention maps in each stage (Figure 5). In the first stage, low-resolution images were generated with primitive shape and color but lacking details. With guidance from GLAM in the following stages, MirrorGAN generated images by focusing on the most relevant and important areas. Consequently, the quality of the generated images progressively improved, *e.g.*, the colors and details of the wings and crown. The top-5 global and local attention maps in each stage are shown below the images. It can be seen that: 1) the global attention concentrated more on the global context in the earlier stage and then the context around specific regions in later stages, 2) the local attention helped the generator synthesize images with fine-grained details by guiding it to focus on the most relevant words, and 3) the global attention is complementary to the local attention, they collaboratively contributed to the progressively improved generation.

In addition, we also present the images generated by MirrorGAN by modifying the text descriptions by a single

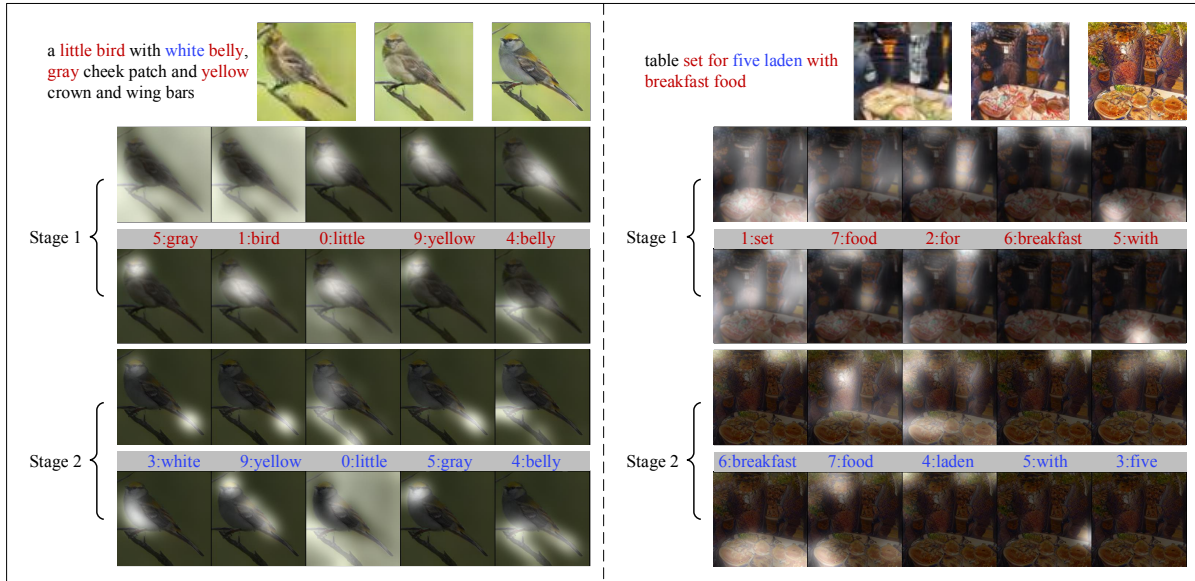


Figure 5: Attention visualization on the CUB and the COCO test sets. The first row shows the output  $64 \times 64$  images generated by  $G_0$ ,  $128 \times 128$  images generated by  $G_1$  and  $256 \times 256$  images generated by  $G_2$ . And the following rows show the Global-Local attention generated in stage 1 and 2. Please refer to the supplementary material for more examples.



Figure 6: Images generated by MirrorGAN by modifying the text descriptions by a single word and the corresponding top-2 attention maps in the last stage.

word (Figure 6). MirrorGAN captured subtle semantic differences in the text descriptions.

#### 4.4. Limitation and discussion

Although our proposed MirrorGAN shows superiority in generating visually realistic and semantically consistent images, some limitations must be taken into consideration in future studies. First, STREAM and other MirrorGAN modules are not jointly optimized with complete end-to-end training due to limited computational resources. Second, we only utilize a basic method for text embedding in STEM and image captioning in STREAM, which could be further improved, for example, by using the recently proposed BERT model [7] and state-of-the-art image captioning models [2, 3]. Third, although MirrorGAN is initially

designed for the T2I generation by aligning cross-media semantics, we believe that its complementarity to the state-of-the-art CycleGAN can be further exploited to enhance model capacity for jointly modeling cross-media content.

## 5. Conclusions

In this paper, we address the challenging T2I generation problem by proposing a novel global-local attentive and semantic-preserving text-to-image-to-text framework called MirrorGAN. MirrorGAN successfully exploits the idea of learning text-to-image generation by redescription. STEM generates word- and sentence-level embeddings. GLAM has a cascaded architecture for generating target images from coarse to fine scales, leveraging both local word attention and global sentence attention to progressively enhance the diversity and semantic consistency of the generated images. STREAM further supervises the generators by regenerating the text description from the generated image, which semantically aligns with the given text description. We show that MirrorGAN achieves new state-of-the-art performance on two benchmark datasets.

**Acknowledgements:** This work is supported in part by Chinese National Double First-rate Project about digital protection of cultural relics in Grotto Temple and equipment upgrading of the Chinese National Cultural Heritage Administration scientific research institutes, the National Natural Science Foundation of China Project 61806062, and the Australian Research Council Projects FL-170100117, DP-180103424, and IH-180100002.



## References

- [1] A. Almahairi, S. Rajeswar, A. Sordoni, P. Bachman, and A. C. Courville. Augmented cyclegan: Learning many-to-many mappings from unpaired data. In *International Conference on Machine Learning (ICML)*, 2018.
- [2] P. Anderson, X. He, C. Buehler, D. Teney, M. Johnson, S. Gould, and L. Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [3] F. Chen, R. Ji, X. Sun, Y. Wu, and J. Su. Groupcap: Group-based image captioning with structured relevance and diversity constraints. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [4] K. Cho, B. van Merriënboer, C. Gulcehre, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *Proceedings of conference on Empirical Methods on Natural Language Processing (EMNLP)*, 2014.
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [6] H. Deubel and W. X. Schneider. Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Vision research*, 1996.
- [7] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [8] O. Firat, K. Cho, and Y. Bengio. Multi-way, multilingual neural machine translation with a shared attention mechanism. In *North American Association for Computational Linguistics (NAACL)*, 2016.
- [9] L. Gao, Z. Guo, H. Zhang, X. Xu, and H. T. Shen. Video captioning with attention-based lstm and semantic consistency. *IEEE Transactions on Multimedia*, 2017.
- [10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances In Neural Information Processing Systems (NIPS)*, 2014.
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [12] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 1997.
- [13] S. Hong, D. Yang, J. Choi, and H. Lee. Inferring semantic layout for hierarchical text-to-image synthesis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [14] P. Isola, J. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [15] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision (ECCV)*, 2016.
- [16] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [17] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision (ECCV)*, 2014.
- [18] J. Liu, C. Gao, D. Meng, and A. G. Hauptmann. Decidenet: Counting varying density crowds through attention guided detection and density estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [19] M.-T. Luong, H. Pham, and C. D. Manning. Effective approaches to attention-based neural machine translation. In *Proceedings of conference on Empirical Methods on Natural Language Processing (EMNLP)*, 2015.
- [20] A. Nguyen, J. Clune, Y. Bengio, A. Dosovitskiy, and J. Yosinski. Plug & play generative networks: Conditional iterative generation of images in latent space. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [21] A. Oliva, A. Torralba, M. S. Castelhano, and J. M. Henderson. Top-down control of visual attention in object detection. In *IEEE International Conference on Image Processing (ICIP)*, 2003.
- [22] T. Qiao, J. Dong, and D. Xu. Exploring human-like attention supervision in visual question answering. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [23] T. Qiao, W. Zhang, M. Zhang, Z. Ma, and D. Xu. Ancient painting to natural image: A new solution for painting processing. *IEEE Winter Conf. on Applications of Computer Vision (WACV)*, 2019.
- [24] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee. Generative adversarial text to image synthesis. In *International Conference on Machine Learning (ICML)*, 2016.
- [25] S. E. Reed, Z. Akata, S. Mohan, S. Tenka, B. Schiele, and H. Lee. Learning what and where to draw. In *Advances in Neural Information Processing Systems (NIPS)*, 2016.
- [26] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems (NIPS)*, 2016.
- [27] M. Schuster and K. K. Paliwal. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 1997.
- [28] Y. Taigman, A. Polyak, and L. Wolf. Unsupervised cross-domain image generation. *International Conference on Learning Representations (ICLR)*, 2017.
- [29] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [30] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011 dataset. *California Institute of Technology*, 2011.
- [31] J. Wang, W. Jiang, L. Ma, W. Liu, and Y. Xu. Bidirectional attentive fusion with context gating for dense video captioning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

- [32] J. Wen, R. Liu, N. Zheng, Q. Zheng, Z. Gong, and J. Yuan. Exploiting local feature patterns for unsupervised domain adaptation. In *Thirty-Third AAAI Conference on Artificial Intelligence*, 2019.
- [33] H. Xu and K. Saenko. Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. In *European Conference on Computer Vision (ECCV)*, 2016.
- [34] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International Conference on Machine Learning (ICML)*, 2015.
- [35] T. Xu, P. Zhang, Q. Huang, H. Zhang, Z. Gan, X. Huang, and X. He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [36] Z. Yang, X. He, J. Gao, L. Deng, and A. Smola. Stacked attention networks for image question answering. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [37] Z. Yi, H. Zhang, P. Tan, and M. Gong. Dualgan: Unsupervised dual learning for image-to-image translation. *The IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [38] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang. Generative image inpainting with contextual attention. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [39] H. Zhang, T. Xu, H. Li, S. Zhang, X. Huang, X. Wang, and D. Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [40] H. Zhang, T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, and D. Metaxas. Stackgan++: Realistic image synthesis with stacked generative adversarial networks. *arXiv preprint arXiv:1710.10916*, 2017.
- [41] X. Zhang, T. Wang, J. Qi, H. Lu, and G. Wang. Progressive attention guided recurrent network for salient object detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [42] Z. Zhang, Y. Xie, and L. Yang. Photographic text-to-image synthesis with a hierarchically-nested adversarial network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [43] J. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *The IEEE International Conference on Computer Vision (ICCV)*, 2017.