

Photographic Text-to-Image Synthesis with a Hierarchically-nested Adversarial Network

Zizhao Zhang^{*}, Yuanpu Xie^{*}, Lin Yang University of Florida

*equal contribution

Abstract

This paper presents a novel method to deal with the challenging task of generating photographic images conditioned on semantic image descriptions. Our method introduces accompanying hierarchical-nested adversarial objectives inside the network hierarchies, which regularize mid-level representations and assist generator training to capture the complex image statistics. We present an extensile single-stream generator architecture to better adapt the jointed discriminators and push generated images up to high resolutions. We adopt a multi-purpose adversarial loss to encourage more effective image and text information usage in order to improve the semantic consistency and image fidelity simultaneously. Furthermore, we introduce a new visual-semantic similarity measure to evaluate the semantic consistency of generated images. With extensive experimental validation on three public datasets, our method significantly improves previous state of the arts on all datasets over different evaluation metrics.

1. Introduction

Photographic text-to-image synthesis is a significant problem in generative model research [33], which aims to learn a mapping from a semantic text space to a complex RGB image space. This task requires the generated images to be not only realistic but also *semantically consistent*, i.e., the generated images should preserve specific object sketches and semantic details described in text.

Recently, Generative adversarial networks (GANs) have become the main solution to this task. Reed *et al.* [33] address this task through a GAN based framework. But this method only generates 64^2 images and can barely generate vivid object details. Based on this method, StackGAN [46] proposes to stack another low-to-high resolution GAN to generate 256^2 images. But this method requires training two separate GANs. Later on, [7] proposes to bypass the difficulty of learning mappings from text to RGB images



This little bird has a white breast and belly, with a gray crown and black secondaries.



Figure 1: Top: Overview of our hierarchically-nested adversarial network, which produces side output images with growing resolutions. Each side output is associated with a discriminator. Bottom: Two test sample results where finegrained details are highlighted.

and treat it as a pixel-to-pixel translation problem [14]. It works by re-rendering an arbitrary-style 128² training image conditioned on a targeting description. However, its high-resolution synthesis capability is unclear. At present, training a generative model to map from a low-dimensional text space to a high-resolution image space in a fully endto-end manner still remains unsolved.

This paper pays attention to two major difficulties for text-to-image synthesis with GANs. The first is balancing the convergence between generators and discriminators [10, 36], which is a common problem in GANs. The second is stably modeling the huge pixel space in high-resolution images and guaranteeing semantic consistency [46]. An effective strategy to regularize generators is critical to stabilize the training and help capture the complex image statistics [12].

In this paper, we propose a novel end-to-end method that can directly model high-resolution image statistics and generate photographic images (see Figure 1 bottom). The contributions are described as follows.

Our generator resembles a simple vanilla GAN, without requiring multi-stage training and multiple internal text conditioning like [46] or additional class label supervision like [5]. To tackle the problem of the big leap from the text space to the image space, our insight is to leverage and regularize hierarchical representations with additional 'deep' adversarial constraints (see Figure 1 top). We introduce accompanying hierarchically-nested discriminators at multi-scale intermediate layers to play adversarial games and jointly encourage the generator to approach the real training data distribution. We also propose a new convolutional neural network (CNN) design for the generator to support accompanying discriminators more effectively. To guarantee the image diversity and semantic consistency, we enforce discriminators at multiple side outputs of the generator to simultaneously differentiate real-and-fake image-text pairs as well as real-and-fake local image patches.

We validate our proposed method on three datasets, CUB birds [40], Oxford-102 flowers [29], and large-scale MSCOCO [23]. In complement of existing evaluation metrics (e.g. Inception score [36]) for generative models, we also introduce a new visual-semantic similarity metric to evaluate the alignment between generated images and conditioned text. It alleviates the issue of the expensive human evaluation. Extensive experimental results and analysis demonstrate the effectiveness of our method and significantly improved performance compared against previous state of the arts on all three evaluation metrics. All source code will be released.

2. Related Work

Deep generative models have attracted wide interests recently, including GANs [10, 32], Variational Auto-encoders (VAE) [17], etc [31]. There are substantial existing methods investigating the better usage of GANs for different applications, such as image synthesis [32, 38], (unpaired) pixelto-pixel translation [14, 51], medical applications [4, 49], etc [20, 12, 52, 45].

Text-to-image synthesis is an interesting application of GANs. Reed *et al.* [33] is the first to introduce a method that can generate 64^2 resolution images. This method also presents a new strategy for image-text matching aware adversarial training. Reed *et al.* [34] propose a generative

adversarial what-where network (GAWWN) to enable location and content instructions in text-to-image synthesis. Zhang *et al.* [46] propose a two-stage training strategy that is able to generate 256^2 compelling images. Recently, Dong *et al.* [7] propose to learn a joint embedding of images and text so as to re-render a prototype image conditioned on a targeting description. Cha *et al.* [2] explore the usage of the perceptional loss [15] with a CNN pretrained on ImageNet and Dash *et al.* [5] make use of auxiliary classifiers (similar with [30]) to assist GAN training for text-to-image synthesis. Xu *et al.* [42] shows an attention-driven method to improve fine-grained details.

Learning a continuous mapping from a low-dimensional manifold to a complex real image distribution is a longstanding problem. Although GANs have made significant progress, there are still many unsolved difficulties, e.g., training instability and high-resolution generation. Wide methods have been proposed to address the training instability, such as various training techniques [35, 1, 38, 30], regularization using extra knowledge (e.g. image labels, ImageNet CNNs) [8, 20, 5], or different generator and discriminator combinations [27, 9, 12]. While our method shows a new way to unite generators and discriminators and does not require any extra knowledge apart from training paired text and images. In addition, it is easy to see the training difficulty increases significantly as the targeting image resolution increases.

To synthesize high-resolution images, cascade networks are effective to decompose originally difficult tasks to multiple subtasks (Figure 2 A). Denton et al. [6] train a cascade of GANs in a Laplacian pyramid framework (LAP-GAN) and use each to synthesize and refine image details and push up the output resolution stage-by-stage. Stack-GAN also shares similar ideas with LAPGAN. Inspired by this strategy, Chen et al. [3] present a cascaded refinement network to synthesize high-resolution scenes from semantic maps. Recently, Karras *et al.* [16] propose a progressive training of GANs. The idea is to add symmetric generator and discriminator layers gradually for high-resolution image generation (Figure 2 C). Compared with these strategies that train low-to-high resolution GANs stage-by-stage or progressively, our method has the advantage of leveraging mid-level representations to encourage the integration of multiple subtasks, which makes end-to-end high-resolution image synthesis in a single vanilla-like GAN possible.

Leveraging hierarchical representations of CNNs is an effective way to enhance implicit multi-scaling and ensembling for tasks such as image recognition [21, 48] and pixel or object classification [41, 25, 22, 37, 44, 50]. Particularly, using deep supervision [21] at intermediate convolutional layers provides short error paths and increases the discriminativeness of feature representations. *Our hierarchically-nested adversarial objective is inspired by the family of*



Figure 2: Overviews of some typical GAN frameworks. A uses multi-stage GANs [46, 6]. **B** uses multiple discriminators with one generator [9, 28]. **C** progressively trains symmetric discriminators and generators [16, 12]. **A** and **C** can be viewed as decomposing high-resolution tasks to multi-stage low-to-high resolution tasks. **D** is our proposed framework that uses a single-stream generator with hierarchically-nested discriminators trained end-to-end.

deeply-supervised CNNs.

3. Method

3.1. Adversarial objective basics

In brief, a GAN [10] consists of a generator G and a discriminator D, which are alternatively trained to compete with each other. D is optimized to distinguish synthesized images from real images, meanwhile, G is trained to fool D by synthesizing fake images. Concretely, the optimal G and D can be obtained by playing the following two-player min-max game,

$$G^*, D^* = \arg\min_{G} \max_{D} \mathcal{V}(D, G, Y, \boldsymbol{z}), \qquad (1)$$

where Y and $z \sim \mathcal{N}(0,1)$ denote training images and random noises, respectively. \mathcal{V} is the overall GAN objective, which usually takes the form $\mathbb{E}_{Y \sim p_{data}} [\log D(Y)] + \mathbb{E}_{z \sim p_z} [\log(1 - D(G(z)))]$ (the cross-entropy loss) or other variations [26, 1].

3.2. Hierarchical-nested adversarial objectives

Numerous GAN methods have demonstrated ways to unite generators and discriminators for image synthesis. Figure 2 and Section 2 discuss some typical frameworks. Our method actually explores a new dimension of playing this adversarial game along the depth of a generator (Figure 2 D), which integrates additional hierarchical-nested discriminators at intermediate layers of the generator. The proposed objectives act as regularizers to the hidden space of \mathcal{G} , which also offer a short path for error signal flows and help reduce the training instability.

The proposed G is a CNN (defined in Section 3.4), which produces multiple side outputs:

$$X_1, \dots, X_s = \mathcal{G}(\boldsymbol{t}, \boldsymbol{z}), \tag{2}$$

where $t \sim p_{data}$ denotes a sentence embedding (generated by a pre-trained char-RNN text encoder [33]). $\{X_1, ..., X_{s-1}\}$ are images with gradually growing resolutions and X_s is the final output with the highest resolution.

For each side output X_i from the generator, a distinct discriminator D_i is used to compete with it. Therefore, our full min-max objective is defined as

$$\mathcal{G}^*, \mathcal{D}^* = \arg\min_{\mathcal{G}} \max_{\mathcal{D}} \mathcal{V}(\mathcal{G}, \mathcal{D}, \mathcal{Y}, \boldsymbol{t}, \boldsymbol{z}),$$
 (3)

where $\mathcal{D} = \{D_1, ..., D_s\}$ and $\mathcal{Y} = \{Y_1, ..., Y_s\}$ denotes training images at multi-scales, $\{1, ..., s\}$. Compared with Eq. (1), our generator competes with multiple discriminators $\{D_i\}$ at different hierarchies (Figure 2 D), which jointly learn discriminative features in different contextual scales.

In principle, the lower-resolution side output is used to learn semantic consistent image structures (e.g. object sketches, colors, and background), and the subsequent higher-resolution side outputs are used to render finegrained details. Since our method is trained in an endto-end fashion, the lower-resolution outputs can also fully utilize top-down knowledge from discriminators at higher resolutions. As a result, we can observe consistent image structures, color and styles in the outputs of both low and high resolution images. The experiment demonstrates this advantage compared with StackGAN.

3.3. Multi-purpose adversarial losses

Our generator produces resolution-growing side outputs composing an image pyramid. We leverage this hierarchy property and allow adversarial losses to capture hierarchical image statistics, with a goal to guarantee both semantic consistency and image fidelity.

In order to guarantee semantic consistency, we adopt the matching-aware pair loss proposed by [33]. The discriminator is designed to take image-text pairs as inputs and be trained to identify two types of errors: a real image with mismatched text and a fake image with conditioned text.

The pair loss is designed to guarantee the global semantic consistency. However, there is no explicit loss to guide the discriminator to differentiate real images from fake images. And combining both tasks (generating realistic images and matching image styles with text) into one network output complicates the already challenging learning tasks.



Figure 3: For each side out image in the pyramid from the generator, the corresponding discriminator D_i computes the matching-aware pair loss and the local image loss (outputting a $R_i \times R_i$ probability map to classify real or fake image patches).

Moreover, as the image resolution goes higher, it might be challenging for a global pair-loss discriminator to capture the local fine-grained details (results are validated in experiments). In addition, as pointed in [38], a single global discriminator may over-emphasize certain biased local features and lead to artifacts.

To alleviate these issues and guarantee image fidelity, our solution is to add local adversarial image losses. We expect the low-resolution discriminators to focus on global structures, while the high-resolution discriminators to focus on local image details. Each discriminator D_i consists of two branches (see Section 3.4), one computes a single scalar value for the pair loss and another branch computes a $R_i \times R_i$ 2D probability map O_i for the local image loss. For each D_i , we control R_i accordingly to tune the receptive field of each element in O_i , which differentiates whether a corresponding local image patch is real or fake. The local GAN loss is also used for pixel-to-pixel translation tasks [38, 51, 14]. Figure 3 illustrates how hierarchically-nested discriminators compute the two losses on the generated images in the pyramid.

Full Objective Overall, our full min-max adversarial objective can be defined as

$$\mathcal{V}(\mathcal{G}, \mathcal{D}, \mathcal{Y}, \boldsymbol{t}, \boldsymbol{z}) = \sum_{i=1}^{s} \left(L_2 \big(D_i(Y_i) \big) + L_2 \big(D_i(Y_i, \boldsymbol{t}_Y) \big) + \overline{L_2} \big(D_i(X_i) \big) + \overline{L_2} \big(D_i(X_i, \boldsymbol{t}_{X_i}) \big) + \overline{L_2} \big(D_i(Y_i, \boldsymbol{t}_{\overline{Y}}) \big) \Big),$$
(4)

where $L_2(x) = \mathbb{E}[(x-\mathbb{I})^2]$ is the mean-square loss (instead of the conventional cross-entropy loss) and $\overline{L_2}(x) = \mathbb{E}[x^2]$. This objective is minimized by \mathcal{D} . While in practice, \mathcal{G} minimizes $\sum_{i=1}^{s} (L_2(D_i(\mathcal{G}(t, z)_i)) + L_2(D_i(\mathcal{G}(t, z)_i, t_{X_i})))$ instead. For the local image loss, the shape of $x, \mathbb{I} \in \mathbb{R}^{R_i \times R_i}$ varies accordingly (see Figure 3). $R_i = 1$ refers to the (largest local) global range. $D_i(X_i)$ is the image loss branch and $D_i(X_i, t_{X_i})$ is the pair loss branch (conditioned on t_{X_i}). $\{Y_i, t_Y\}$ denotes a matched image-text pair and $\{Y_i, t_{\overline{Y}}\}$ denotes a mismatched image-text pair.

In the spirit of variational auto-encoder [18] and the practice of StackGAN [46] (termed conditioning augmentation (CA)), instead of directly using the deterministic text embedding, we sample a stochastic vector from a Ganssian distribution $\mathcal{N}(\mu(t), \Sigma(t))$, where μ and Σ are parameterized functions of t. We add the Kullback-Leibler divergence regularization term, $D_{KL}(\mathcal{N}(\mu(t), \Sigma(t))||\mathcal{N}(0, I))$, to the GAN objective to prevent over-fitting and force smooth sampling over the text embedding distribution.

3.4. Architecture Design

Generator The generator is simply composed of three kinds of modules, termed K-repeat res-blocks, stretching layers, and linear compression layers. A single res-block in the K-repeat res-block is a modified¹ residual block [11], which contains two convolutional (conv) layers (with batch normalization (BN) [13] and ReLU). The stretching layer serves to change the feature map size and dimension. It simply contains a scale-2 nearest up-sampling layer followed by a conv layer with BN+ReLU. The linear compression layer is a single conv layer followed by a Tanh to directly compress feature maps to the RGB space. We prevent any non-linear function in the compression layer that could impede the gradient signals. Starting from a $1024 \times 4 \times 4$ embedding, which is computed by CA and a trained embedding matrix, the generator simply uses M K-repeat resblocks connected by M-1 in-between stretching layers until the feature maps reach to the targeting resolution. For example, for 256×256 resolution and K=1, there are M=61-repeat res-blocks and 5 stretching layers. At pre-defined side-output positions at scales $\{1, ..., s\}$, we apply the compression layer to generate side output images, for the inputs of discriminators.

Discriminator The discriminator simply contains consecutive stride-2 conv layers with BN+LeakyReLU. There are two branches are added to the upper layer of the discriminator. One branch is a direct fully convolutional layers to produce a $R_i \times R_i$ probability map (see Figure 3) and classify each location as real or fake. Another branch first concatenates a $512 \times 4 \times 4$ feature map and a $128 \times 4 \times 4$ text embedding (replicated by a reduced 128-d text embedding). Then we use an 1×1 conv to fuse text and image features and a 4×4 conv layer to classify an image-text pair to real or fake.

The optimization is similar to the standard alternative

¹We remove ReLU after the skip-addition of each residual block, with an intention to reduce sparse gradients.

training strategy in GANs. Please refer to the supplementary material for more training and network details.

4. Experiments

We denote our method as **HDGAN**, referring to High-Definition results and the idea of Hierarchically-nested Discriminators.

Dataset We evaluate our model on three widely used datasets. The CUB dataset [40] contains 11,788 bird images belonging to 200 categories. The Oxford-102 dataset [29] contains 8,189 flow images in 102 categories. Each image in both datasets is annotated with 10 descriptions provided by [33]. We pre-process and split the images of CUB and Oxford-102 following the same pipeline in [33, 46]. The COCO dataset [23] contains 82,783 training images and 40,504 validation images. Each image has 5 text annotations. We use the pre-trained char-RNN text encoder provided by [33] to encode each sentence into a 1024-d text embedding vector.

Evaluation metric We use three kinds of quantitative metrics to evaluate our method. 1) Inception score [36] is a measurement of both objectiveness and diversity of generated images. Evaluating this score needs a pre-trained Inception model [39] on ImageNet. For CUB and Oxford-102, we use the fine-tuned Inception models on the training sets of the two datasets, respectively, provided by Stack-GAN. 2) Multi-scale structural similarity (MS-SSIM) metric [36] is used for further validation. It tests pair-wise similarity of generated images and can identity mode collapses reliably [30]. Lower score indicates higher diversity of generated images (i.e. less model collapses).

3) Visual-semantic similarity The aforementioned metrics are widely used for evaluating standard GANs. However, they can not measure the alignment between generated images and the conditioned text, i.e., semantic consistency. [46] resorts to human evaluation, but this procedure is expensive and difficult to conduct. To tackle this issue, we introduce a new measurement inspired by [19], namely visual-semantic similarity (VS similarity). The insight is to train a visual-semantic embedding model and use it to measure the distance between synthesized images and input text. Denote v as an image feature vector extracted by an Inception model f_{cnn} . We define a scoring function $c(x, y) = \frac{x \cdot y}{||x||_2 \cdot ||y||_2}$. Then, we train two mapping functions f_v and f_t , which map both real images and paired text embeddings into a common space in \mathbb{R}^{512} , by minimizing the following bi-directional ranking loss:

$$\sum_{\boldsymbol{v}} \sum_{\boldsymbol{t}_{\overline{v}}} \max(0, \delta - c(f_{v}(\boldsymbol{v}), f_{t}(\boldsymbol{t}_{v})) + c(f_{v}(\boldsymbol{v}), f_{t}(\boldsymbol{t}_{\overline{v}}))) + \sum_{\boldsymbol{t}} \sum_{\boldsymbol{v}_{\overline{t}}} \max(0, \delta - c(f_{t}(\boldsymbol{t}), f_{v}(\boldsymbol{v}_{v})) + c(f_{t}(\boldsymbol{t}), f_{v}(\boldsymbol{v}_{\overline{t}}))),$$
(5)

Method	Dataset			
Wiethou	CUB	Oxford-102	COCO	
GAN-INT-CLS	$2.88 \pm .04$	$2.66 {\pm} .03$	$7.88 {\pm}.07$	
GAWWN	$3.60 {\pm} .07$	-	-	
StackGAN	$3.70 {\pm}.04$	$3.20 {\pm}.01$	$8.45 {\pm} .03^{\star}$	
StackGAN++	$3.84 {\pm}.06$	-	-	
TAC-GAN	-	$3.45 {\pm} .05$	-	
HDGAN	$4.15{\pm}.05$	$3.45{\pm}.07$	$11.86{\pm}.18$	

*Recently, it updated to $10.62 \pm .19$ in its source code.

Table 1: The Inception-score comparison on the threedatasets. HDGAN outperforms others significantly.

Method	Dataset			
	CUB	Oxford-102	COCO	
Ground Truth	$.302 \pm .151$	$.336 {\pm} .138$	$.426 \pm .157$	
StackGAN	$.228 \pm .162$	$.278 \pm .134$	_	
HDGAN	$.246 {\pm} .157$	$.296 {\pm} .131$	$.199 {\pm} .183$	

Table 2: The VS similarity evaluation on the three datasets. The higher score represents higher semantic consistency between the generated images and conditioned text. The groundtruth score is shown in the first row.

where δ is the margin, which is set as 0.2. $\{v, t\}$ is a ground truth image-text pair, and $\{v, t_{\overline{v}}\}$ and $\{v_{\overline{t}}, t\}$ denote mismatched image-text pairs. In the testing stage, given an text embedding t, and the generated image x, the VS score can be calculated as $c(f_{cnn}(x), t)$. Higher score indicates better semantic consistency.

4.1. Comparative Results

To validate our proposed HDGAN, we compare our results with GAN-INT-CLS [33], GAWWN [34], TAC-GAN [5], Progressive GAN [16], StackGAN [46] and also its improved version StackGAN++ [47]². We especially compare with StackGAN in details (results are obtained from its provided models).

Table 1 compares the Inception score. We follow the experiment settings of StackGAN to sample $\sim 30,000\ 256^2$ images for computing the score. HDGAN achieves significant improvement compared against other methods. For example, it improves StackGAN by .45 and StackGAN++ by .31 on CUB. HDGAN achieves competitive results with TAC-GAN on Oxford-102. TAC-GAN uses image labels to increase the discriminability of generators, while we do not use any extra knowledge. Figure 4 and Figure 5 compare the qualitative results with StackGAN on CUB and Oxford-102, respectively, by demonstrating more, semantic details, natural color, and complex object structures. Moreover, we qualitatively compare the diversity of samples conditioned on the same text (with random input noises) in Figure 7 left.

²StackGAN++ and Prog.GAN are two very recently released preprints we noticed. We acknowledge them as they also target at generating high-resolution images.

This is a small bird with tis body covered in blue feathers, and some brown feathers on its wings A small bird with brown wings, vanilla break and a small black

The bird has a small black eyering and has a white belly

A small bird with brown and white feathers, and a red head



Figure 4: Generated images on CUB compared with StackGAN. Each sample shows the input text and generated 64^2 (left) and 256^2 (right) images. Our results have significantly higher quality and preserve more semantic details, for example, "the brown and white feathers and read head" in the last column is much better reflected in our images. Moreover, we observed our birds exhibit nicer poses (e.g. the frontal/back views in the second/forth columns). Zoom-in for better observation.



Figure 5: Generated images on Oxford-102 compared with StackGAN. Our generated images perform more natural satisfiability and higher contrast and can generate complex flower structures (e.g. spiked petals in the third example).

HDGAN can generate substantially more compelling samples.

Different from CUB and Oxford-102, COCO is a much more challenging dataset and contains largely diverse natural scenes. Our method significantly outperforms Stack-GAN as well (Table 1). Figure 6 also shows some generated samples in several different scenes. Please refer to the supplementary material for more results.

Furthermore, the right figure compares the multiresolution Inception score gon CUB. Our results are from the side outputs of a gsingle model. As can be gobserved, our 64^2 results outperform the 128^2 results of StackGAN and our 128^2 results also outperform 256^2



results of StackGAN substantially. It demonstrates that our HDGAN better preserves semantically consistent infor-

A kitchen with lots of clutter and an open drawer cloudy day

There is a tall tower connected to a building A man is standing on top of a snow covered hill



Figure 6: Samples on the COCO validation set, which contain descriptions across different scenes.

mation in all resolutions (as stated in Section 3.2). Figure 7 right validates this property qualitatively. On the other hand, we observe that, in StackGAN, the low-resolution images and high-resolution images sometimes are visually inconsistent (see examples in Figure 4 and 5).

Table 2 compares the proposed visual-semantic similarity (VS) results on three datasets. The scores of the groundtruth image-text pair are also shown for reference. HDGAN achieves consistently better performance on both CUB and Oxford-102. These results demonstrate that



Figure 7: Left: Multiple samples are shown given a single input text. The proposed HDGAN (top) show obviously more fine-grained details. Right: Side outputs of HDGAN with increasing resolutions. Different resolutions are semantically consistent and semantic details appear as the resolution increases.



Table 3: Left: Class-wise MS-SSIM evaluation. Lower score indicates higher intraclass dissimilarity. The points below the equality line represent classes our HDGAN wins. The inter-class std is shown in axis text. Right: Overall (not class-wised) MS-SSIM evaluation.

HDGAN can better capture the visual semantic information in generated images.

Table 3 compares the MS-SSIM score with StackGAN and Prog.GAN for bird image generation. StackGAN and our HDGAN use text as input so the generated images are separable in class. We randomly sample $\sim 20,000$ image pairs (400 per class) and compare the class-wise score in the left figure. HDGAN outperforms StackGAN in majority of classes and also has a lower standard deviation (.023 vs. .032). Note that Prog.GAN uses a noise input rather than text. We can compare with it for a general measure of the image diversity. Following the procedure of Prog.GAN, we randomly sample $\sim 10,000$ image pairs from all generated samples³ and show the results in Table 3 right. HDGAN outperforms both methods.

4.2. Style Transfer Using Sentence Interpolation

Ideally, a well-trained model can generalize to a smooth linear latent data manifold. To demonstrate this capability, we generate images using the linearly interpolated embeddings between two source sentences. As shown in Figure 8, our generated images show smooth style transformation and well reflect the semantic details in sentences. For example, in the second row, complicated sentences with detailed appearance descriptions (e.g. pointy peaks and black wings) are used, our model could still successfully capture these subtle features and tune the bird's appearance smoothly.

4.3. Ablation Study and Discussion

Hierarchically-nested adversarial training Our hierarchically-nested discriminators play a role of regularizing the layer representations (at scale $\{64, 128, 256\}$). In Table 4, we demonstrate their effectiveness and show the performance by removing parts of discriminators on both CUB and COCO datasets. As can be seen, increasing the usage of discriminators at different scales have positive effects. And using discriminators at 64^2 is critical (by comparing the 64-256 and 128-256 cases). For now, it is uncertain if adding more discriminators and even on lower resolutions would be helpful. Further validation will be conducted. StackGAN emphasizes the importance of using text embeddings not only at the input but also with intermediate features of the generator, by showing a large drop from 3.7 to 3.45 without doing so. While our method only uses text embeddings at the input. Our results strongly demonstrate the effectiveness of our hierarchically-nested adversarial training to maintain such semantic information and a high Inception score.

 $^{^3}We$ use 256^2 bird images provided by Prog.GAN at https://github.com/tkarras/progressive_growing_of_gans. Note that Prog.GAN is trained on the LSUN [43] bird set, which contains ${\sim}2$ million bird images.



Figure 8: Text embedding interpolation from the source to target sentence results in smooth image style changes to match the targeting sentence.

Discriminators		Inception score		
64	128	256	CUB	COCO
		\checkmark	$3.52 {\pm}.04$	-
	\checkmark	\checkmark	$3.99 {\pm} .04$	-
\checkmark		\checkmark	$4.14 {\pm} .03$	$11.29 \pm .18$
\checkmark	\checkmark	\checkmark	$4.15{\pm}.05$	$11.86 {\pm}.18$

Table 4: Ablation study of hierarchically-nested adversarial discriminators on CUB and COO. \checkmark indicates whether a discriminator at a certain scale is used. See text for detailed explanations.

The local image loss We analyze the effectiveness of the proposed local adversarial image loss. Table 4 compares the case without using it (denoted as 'w/o local image loss'). The local image loss helps improve the visualsemantic matching evidenced by a higher VS score. We hypothesize that it is because adding the separate local image loss can offer the pair loss more "focus" on learning the semantic consistency. Furthermore, the local image loss helps generate more vivid image details. As demonstrated in Figure 9, although both models can successfully capture the semantic details in the text, the 'w/ local' model generates complex object structures described in the conditioned text more precisely.

Design principles StackGAN claims the failure of directly training a vanilla 256² GAN to generate meaningful images. We test this extreme case using our method by removing all nested discriminators without the last one. Our method still generates fairly meaningful results (the first row of Table 4), which demonstrate the effectiveness of our proposed framework (see Section 3.4).

Initially, we tried to share the top layers of the hierarchical-nested discriminators of HDGAN inspired by [24]. The intuition is that all discriminators have a common goal to differentiate real and fake despite difficult scales, and such sharing would reduce their inter-variances. However, we did not observe benefits from this mechanism and our independent discriminators can be trained fairly stably.

HDGAN has a very succinct framework, compared most

	Inc. score	VS
w/o local image loss	$3.12 \pm .02$	$.263 \pm .130$
w/ local image loss	$3.45 {\pm}.07$	$.296 \pm .130$

Table 5: Ablation study of the local image loss on Oxford-102. See text for detailed explanations.

w/ local image loss w/o local image loss

This petals are purple and white, and are spiked outward like spines

This flower has plain white petals as well as some that have dark red stripes



Figure 9: Qualitative evaluation of the local image loss. The two images w/ the local image loss more precisely exhibit the complex flower petal structures described in the (colored) text.

existing methods, as they [42, 2] adds extra supervision on output images to 'inject' semantic information, which is shown helpful for improving the inception score. However, it is not clear that whether these strategies can substantially improve the visual quality, which is worth further study.

5. Conclusion

In this paper, we present a novel and effective method to tackle the problem of generating images conditioned on text descriptions. We explore a new dimension of playing adversarial games along the depth of the generator using the hierarchical-nested adversarial objectives. A multi-purpose adversarial loss is adopted to help render fine-grained image details. We also introduce a new evaluation metric to evaluate the semantic consistency between generated images and conditioned text. Extensive experiment results demonstrate that our method, namely HDGAN, can generate highresolution photographic images and performs significantly better than existing state of the arts on three public datasets.

References

- D. Berthelot, T. Schumm, and L. Metz. Began: Boundary equilibrium generative adversarial networks. *arXiv preprint arXiv:1703.10717*, 2017. 2, 3
- [2] M. Cha, Y. Gwon, and H. T. Kung. Adversarial nets with perceptual losses for text-to-image synthesis. *arXiv preprint* arXiv:1708.09321, 2017. 2, 8
- [3] Q. Chen and V. Koltun. Photographic image synthesis with cascaded refinement networks. *ICCV*, 2017. 2

- [4] P. Costa, A. Galdran, M. I. Meyer, M. D. Abràmoff, M. Niemeijer, A. M. Mendonça, and A. Campilho. Towards adversarial retinal image synthesis. *arXiv preprint* arXiv:1701.08974, 2017. 2
- [5] A. Dash, J. C. B. Gamboa, S. Ahmed, M. Z. Afzal, and M. Liwicki. Tac-gan-text conditioned auxiliary classifier generative adversarial network. *arXiv preprint* arXiv:1703.06412, 2017. 2, 5
- [6] E. L. Denton, S. Chintala, R. Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In *NIPS*, pages 1486–1494, 2015. 2, 3
- [7] H. Dong, S. Yu, C. Wu, and Y. Guo. Semantic image synthesis via adversarial learning. *ICCV*, 2017. 1, 2
- [8] A. Dosovitskiy and T. Brox. Generating images with perceptual similarity metrics based on deep networks. In *NIPS*, 2016. 2
- [9] I. Durugkar, I. Gemp, and S. Mahadevan. Generative multiadversarial networks. arXiv preprint arXiv:1611.01673, 2016. 2, 3
- [10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014. 2, 3
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In ECCV, 2016. 4
- [12] X. Huang, Y. Li, O. Poursaeed, J. Hopcroft, and S. Belongie. Stacked generative adversarial networks. *CVPR*, 2017. 2, 3
- [13] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 4
- [14] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *CVPR*, 2017. 1, 2, 4
- [15] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016.
 2
- [16] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *arXiv preprint arXiv:1710.10196*, 2016. 2, 3, 5
- [17] D. P. Kingma and M. Welling. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114, 2013. 2
- [18] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *ICLR*, 2013. 4
- [19] R. Kiros, R. Salakhutdinov, and R. S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models. arXiv preprint arXiv:1411.2539, 2014. 5
- [20] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. *CVPR*, 2017. 2
- [21] C.-Y. Lee, S. Xie, P. Gallagher, Z. Zhang, and Z. Tu. Deeplysupervised nets. In AIS, 2015. 2
- [22] K. Li, Z. Wu, K.-C. Peng, J. Ernst, and Y. Fu. Tell me where to look: Guided attention inference network. *arXiv preprint arXiv:1802.10171*, 2017. 2
- [23] 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 *ECCV*, 2014. 2, 5

- [24] M.-Y. Liu, T. Breuel, and J. Kautz. Unsupervised image-to-image translation networks. arXiv preprint arXiv:1703.00848, 2017. 8
- [25] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In CVPR, 2015. 2
- [26] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. P. Smolley. Least squares generative adversarial networks. arXiv preprint ArXiv:1611.04076, 2016. 3
- [27] L. Metz, B. Poole, D. Pfau, and J. Sohl-Dickstein. Unrolled generative adversarial networks. arXiv preprint arXiv:1611.02163, 2016. 2
- [28] T. D. Nguyen, T. Le, H. Vu, and D. Phung. Dual discriminator generative adversarial nets. In *NIPS*, 2017. 3
- [29] M.-E. Nilsback and A. Zisserman. Automated flower classification over a large number of classes. In *ICCVGIP*, 2008. 2, 5
- [30] A. Odena, C. Olah, and J. Shlens. Conditional image synthesis with auxiliary classifier gans. arXiv preprint arXiv:1610.09585, 2016. 2, 5
- [31] A. v. d. Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. *ICML*, 2016. 2
- [32] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434, 2015. 2
- [33] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee. Generative adversarial text to image synthesis. *ICML*, 2016. 1, 2, 3, 5
- [34] S. E. Reed, Z. Akata, S. Mohan, S. Tenka, B. Schiele, and H. Lee. Learning what and where to draw. In *NIPS*, 2016. 2, 5
- [35] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training gans. In *NIPS*, 2016. 2
- [36] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen, and X. Chen. Improved techniques for training gans. In *NIPS*. 2016. 2, 5
- [37] Z. Shen, Z. Liu, J. Li, Y.-G. Jiang, Y. Chen, and X. Xue. Dsod: Learning deeply supervised object detectors from scratch. In *ICCV*, 2017. 2
- [38] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb. Learning from simulated and unsupervised images through adversarial training. *CVPR*, 2017. 2, 4
- [39] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, June 2015. 5
- [40] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-ucsd birds 200. 2010. 2, 5
- [41] S. Xie and Z. Tu. Holistically-nested edge detection. In ICCV, 2015. 2
- [42] 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. *arXiv preprint arXiv:1711.10485*, 2017. 2, 8
- [43] F. Yu, A. Seff, Y. Zhang, S. Song, T. Funkhouser, and J. Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 7

- [44] H. Zhang, K. Dana, J. Shi, Z. Zhang, X. Wang, A. Tyagi, and A. Agrawal. Context encoding for semantic segmentation. In *CVPR*, 2018. 2
- [45] H. Zhang, V. Sindagi, and V. M. Patel. Image de-raining using a conditional generative adversarial network. arXiv preprint arXiv:1701.05957, 2017. 2
- [46] H. Zhang, T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, and D. Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *ICCV*, 2017. 1, 2, 3, 4, 5
- [47] H. Zhang, T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, and D. Metaxas. Stackgan++: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *arXiv* preprint arXiv:1710.10916, 2017. 5
- [48] Z. Zhang, Y. Xie, F. Xing, M. Mcgough, and L. Yang. Mdnet: A semantically and visually interpretable medical image diagnosis network. In *CVPR*, 2017. 2
- [49] Z. Zhang, L. Yang, and Y. Zheng. Translating and segmenting multimodal medical volumes with cycle- and shapeconsistency generative adversarial network. In *CVPR*, 2018.
- [50] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid scene parsing network. In CVPR, 2017. 2
- [51] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired imageto-image translation using cycle-consistent adversarial networks. *ICCV*, 2017. 2, 4
- [52] Y. Zhu, M. Elhoseiny, B. Liu, and A. Elgammal. Imagine it for me: Generative adversarial approach for zero-shot learning from noisy texts. *arXiv preprint arXiv:1712.01381*, 2017. 2