

DAE-GAN: Dynamic Aspect-aware GAN for Text-to-Image Synthesis

Shulan Ruan^{1†}, Yong Zhang^{2*}, Kun Zhang³, Yanbo Fan², Fan Tang⁴, Qi Liu¹, Enhong Chen^{1*}

¹School of Computer Science and Technology, University of Science and Technology of China

²Tencent AI Lab, ³Hefei University of Technology, ⁴Jilin University

slruan@mail.ustc.edu.cn, {zhangyong201303, zhang1028kun, fanyanbo0124, tfan.108}@gmail.com, {qiliuql, cheneh}@ustc.edu.cn

Abstract

Text-to-image synthesis refers to generating an image from a given text description, the key goal of which lies in photo realism and semantic consistency. Previous methods usually generate an initial image with sentence embedding and then refine it with fine-grained word embedding. Despite the significant progress, the ‘aspect’ information (e.g., red eyes) contained in the text, referring to several words rather than a word that depicts ‘a particular part or feature of something’, is often ignored, which is highly helpful for synthesizing image details. How to make better utilization of aspect information in text-to-image synthesis still remains an unresolved challenge. To address this problem, in this paper, we propose a Dynamic Aspect-aware GAN (DAE-GAN) that represents text information comprehensively from multiple granularities, including sentence-level, word-level, and aspect-level. Moreover, inspired by human learning behaviors, we develop a novel Aspect-aware Dynamic Re-drawer (ADR) for image refinement, in which an Attended Global Refinement (AGR) module and an Aspect-aware Local Refinement (ALR) module are alternately employed. AGR utilizes word-level embedding to globally enhance the previously generated image, while ALR dynamically employs aspect-level embedding to refine image details from a local perspective. Finally, a corresponding matching loss function is designed to ensure the text-image semantic consistency at different levels. Extensive experiments on two well-studied and publicly available datasets (i.e., CUB-200 and COCO) demonstrate the superiority and rationality of our method.

1. Introduction

Text-to-image synthesis requires an agent to generate a photo-realistic image according to the given text description. Due to its significant potential in many applications

[†]Work done during an internship in Tencent AI Lab.

^{*}Corresponding Authors.



Figure 1. Comparisons between DM-GAN [45] and our DAE-GAN. DM-GAN first generates a low-resolution image with sentence-level information and then refines it with word-level features. DAE-GAN refines images from both global and local perspectives with word-level and aspect information contained.

such as art generation [41] and computer-aided design [2] but challenging nature, it is arousing extensive research attention in recent years.

In the past few years, Generative Adversarial Networks (GANs) [5] have been proved tremendously successful for this task [20]. Most existing methods make efforts on the two-stage framework by first generating initial low-resolution images and then refining them to high-resolution ones [38, 39, 36]. Among all these methods, the proposal of AttnGAN [36] plays an extremely important role. At the initial stage, sentence-level information is employed to generate a low-resolution image. Then, at the refinement stage, AttnGAN utilizes word-level features to refine the previously generated image by repeatedly adopting attention mechanism to select important words. Based on AttnGAN, text-to-image synthesis has been pushed by a large step forward [45, 4, 18]. A synthesis example by DM-GAN [45] is presented at the top of Figure 1.

Although remarkable performance has been accomplished with these efforts, there still exist several limitations to be unresolved. For example, most previous methods only employ sentence-level and word-level features, ignoring the ‘aspect-level’ features. ‘Aspect’ here refers to several words rather than a word that depicts ‘a particu-

lar part or feature of something’. There are often multiple aspect terms contained in a sentence to describe an object or a scene from different perspectives, *e.g.*, ‘the black bird’ and ‘red eyes’ in the text description in Figure 1. Semantic understanding of a sentence is highly dependent on both content and aspect[33]. Both industry and academia have realized the importance of the relationship between aspect term and sentence [3, 40, 17]. In fact, aspect information contained in the text could be helpful for image synthesis, especially for the refinement of local image details. Though the value of aspect information has been proved, how to make better utilization of aspect information in text-to-image synthesis still remains a big challenge.

Fortunately, some interesting studies of human learning behaviors could give us some inspirations. Researchers have already demonstrated that human eyes have central vision and peripheral vision [1, 34, 25]. Central vision concentrates on what a person needs at the current time, while peripheral vision uses observation of the surroundings to support central vision. Through the dynamic use of central vision and peripheral vision, we could make an in-depth semantic understanding of text and visual content.

To this end, in this paper, we propose a novel *Dynamic Aspect-aware GAN(DAE-GAN)* for text-to-image synthesis. To be specific, we firstly encode text information from multiple granularities comprehensively, including sentence-level, word-level, and aspect-level. Then, at the two-stage generation, we first generate a low-resolution image with sentence-level embedding at the initial stage. Next, at the refinement stage, by viewing aspect-level features as central vision and word-level features as peripheral vision, we develop an *Aspect-aware Dynamic Re-drawer (ADR)*, which alternately applies an *Attended Global Refinement (AGR)* module and an *Aspect-aware Local Refinement (ALR)* module for image refinement. *AGR* utilizes word-level embedding to globally enhance the previously generated images. *ALR* dynamically utilizes aspect-level embedding to refine image details from a local perspective. Finally, to provide supervision for intermediate synthesis procedures, a corresponding matching loss function is designed to ensure the text-image semantic consistency. The bottom of Figure 1 illustrates an example of our proposed method. When given the aspect ‘the black bird’, our method focuses on adjusting the color of the whole bird body based on the previously generated image. When dealing with the aspect ‘red eyes’, our model then correspondingly focuses on the refinement of bird eyes with remaining other parts.

Our main contributions are summarized as follows:

- We observe the great potential of aspect information and apply it to text-to-image synthesis.
- We propose a novel DAE-GAN, in which text information is comprehensively represented from multiple

granularities, and an *ADR* is developed to refine images from both local and global perspectives.

- Extensive experiments including quantitative and qualitative evaluations show the superiority and rationality of our proposed method. Specially, the causality study demonstrates DAE-GAN as an interpretable model.

2. Related Work

Due to the great potential in broad applications, text-to-image synthesis though challenging yet is arousing extensive research attention. Earlier methods have achieved progress on this task due to the emergence of deep generative models [14, 6, 15, 30, 21].

Thanks to the advancement of GAN, recent approaches further improve the generation quality and have shown promising results on text-to-image synthesis. Reed et al. [20] first developed a simple and effective GAN architecture that enabled compelling text-to-image synthesis. Nevertheless, the size of the image was only as small as 64×64 . To this end, StackGAN [38] was proposed to generate higher resolution images with two stages. They initially sketched primary shape and colors, and then re-read the text to produce a photo-realistic image. With the aim of discarding stacking architectures, Tao et al. [29] proposed DF-GAN to directly synthesize images without extra networks. However, these works only took sentence-level features into consideration, which lacked fine-grained text understanding. As a consequence, fine-grained details are often missing in the generated images.

To address this issue, plenty of work has pushed text-to-image synthesis a large step forward by utilizing word-level features at the refinement stage to enhance image details. Among them, AttnGAN [36] played an important role. It utilized attention mechanism to repeatedly select important words at different steps for image refinement, which brought text-to-image synthesis research to a new height. Zhu et al. [45] proposed DM-GAN that substituted memory network for attention mechanism to dynamically pick important words at the refinement stage. To improve semantic consistency in text-to-image synthesis, Qiao et al. [18] proposed MirroGAN by semantically aligning the re-description of the generated image with the given text description. To explore the semantic correlation between different yet related sentences, RiFeGAN [4] exploited an attention-based caption matching model to select and refine the compatible candidate captions from prior knowledge. Yang et al. [37] proposed MA-GAN to reduce the variation between their generated images with similar captions and enhance the reliability of the generated results.

With demands for various applications as well as the emergence of new data sets, other compelling text-to-image research is also developed based on GAN. In [10, 8, 26, 11,

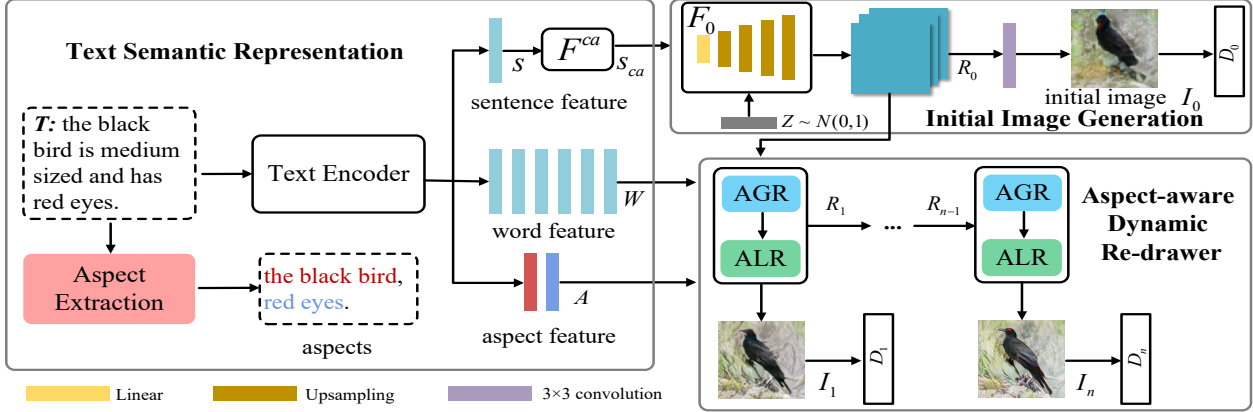


Figure 2. Overall framework of DAE-GAN.

9], image generation was studied for datasets with multiple objects. For example, Huang et al. [9] introduced an additional set of natural attentions between object-grid regions and word phrases. However, extra bounding box information of each object must be required as labels. To address the problem of food image synthesis with recipes, large efforts were made in [43, 16, 44]. Other recent work also has achieved impressive results in person image synthesis [42, 13]. Given the reference images and text descriptions, they could manipulate the visual appearance of a person. Aiming at text-guided multi-modal face generation and manipulation, TediGAN as well as a face image dataset were proposed by Xia et al. [35]. To model the text and image tokens as a single stream of data, Ramesh et al. [19] proposed DALL·E to train a transformer [31] autoregressively. Based on sufficient data and scale, DALL·E achieved comparable results with other domain-specific models.

However, most of the aforementioned methods only considered the sentence-level and word-level features for text utilization. They ignored the great potential of aspect information contained in the sentences, which is very helpful for image refinement (e.g., the example in Figure 1). To this end, in this paper, we argue that the aspect information should gain more attention and propose a novel text-to-image synthesis method to employ aspect-level features for local region refinement in a dynamic manner.

3. Dynamic Aspect-aware GAN (DAE-GAN)

As illustrated in Figure 2, our proposed DAE-GAN embodies three main components: 1) *Text Semantic Representation*: extracting text semantic representations from multiple granularities, i.e., sentence-level, word-level, as well as aspect-level; 2) *Initial Image Generation*: generating a low-resolution image with sentence-level text features and a random noise vector; and 3) *Aspect-aware Dynamic Re-drawer*: refining the initial image in a dynamic manner from both global and local perspectives, which is also the main focus in this paper.

3.1. Text Semantic Representation

Comprehensive understanding of text semantics plays a vital role in text-to-image synthesis. Previous methods mainly extract text features from sentence-level and word-level. However, they overlook aspect information contained in the text, which refers to several words rather than a word that depicts a particular part or feature of something, e.g., ‘red eyes’ in ‘the black bird is medium sized and has red eyes’. The granularity of aspect-level information is between those of sentence-level and word-level information. It could be helpful for the refinement of image details and should gain more attention. As shown in Figure 2, we represent text features from multiple granularities, i.e., sentence-level, word-level, and aspect-level. We use a Long Short-Term Memory (LSTM) network to extract the semantic embedding of text description T , which is formulated as follows:

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

where $T = \{T_j | j = 0, 1, \dots, l-1\}$ consists of l words. $\mathbf{W} = \{\mathbf{W}_j | j = 0, 1, \dots, l-1\} \in \mathbb{R}^{l \times d_w}$ represents word-level features that are obtained from the hidden state of LSTM at each time step. Here, d_w means the dimension of text embedding. $\mathbf{s} \in \mathbb{R}^{d_w}$ stands for sentence-level semantic features from the last hidden state of LSTM.

We further employ the Conditioning Augmentation (CA) [38] to augment the training data and avoid overfitting by resampling the input sentence vector from an independent Gaussian distribution. Specifically, we enhance sentence features with CA, which is represented as follows:

$$\mathbf{s}_{ca} = F^{ca}(\mathbf{s}), \quad (2)$$

where $F^{ca}(\cdot)$ stands for the CA function, and \mathbf{s}_{ca} is the augmented sentence semantic representations with CA.

As mentioned before, aspect information is very crucial for the details of generated images. However, since the focus and description of different sentences are different, it

is hard to identify and extract the proper aspect information for each sentence. To this end, we employ the syntactic structures to tackle this problem. Specifically, we first adopt NLTK to make the POS Tagging for each sentence. Then, we manually design different rules to extract aspect information for different datasets. After that, we can obtain the aspect information $\{asp_i | i = 0, 1, \dots, n - 1\}$. Next, we use an LSTM to integrate this information and extract the aspect-level features, which is formulated as follows:

$$\mathbf{A} = \text{LSTM}(\{asp_i | i = 0, 1, \dots, n - 1\}), \quad (3)$$

where \mathbf{A} denotes the aspect-level feature representation for text description and n is the the number of extracted aspects.

3.2. Initial Image Generation

Following the common practice, we first generate a low-resolution image at the initial stage. As illustrated in Figure 2, we utilize the augmented sentence embeddings s_{ca} and a random noise vector z to generate an initial image I_0 . $z \sim N(0, 1)$ is sampled from a normal distribution. Mathematically, we use \mathbf{R}_0 to denote the corresponding image features at the initial stage:

$$\mathbf{R}_0 = F_0(s_{ca}, z), \quad (4)$$

where F_0 is the image generator at the initial generation stage. As depicted in Figure 2, it consists of one fully connected layer and four upsampling layers.

3.3. Aspect-aware Dynamic Re-drawer

To the best of our knowledge, we are the first to introduce aspect information contained in the given sentence into text-to-image synthesis. Therefore, how to integrate the aspect information into image refinement stage is the main challenge that we should tackle. Inspired by human learning behaviors, in this work, we develop a novel *Aspect-aware Dynamic Re-drawer (ADR)* to refine images with the consideration of aspect information in the sentence. Specifically, we design a novel *Attended Global Refinement (AGR)* module to employ fine-grained word-level features for global refinement, and a novel *Aspect-aware Local Refinement (ALR)* module to utilize aspect-level features for local enhancement. By alternately applying these two components in a dynamic way, we are able to refine image details from both global and local perspectives. In the following part, we will take the i^{th} refinement operation for generated image as an example to introduce the technical details of *AGR* and *ALR*.

3.3.1 Attended Global Refinement

To synthesize a photo-realistic and semantic consistent image, it is necessary to further globally refine the image with fine-grained features. Therefore, *AGR* is developed for global refinement based on the initial image.

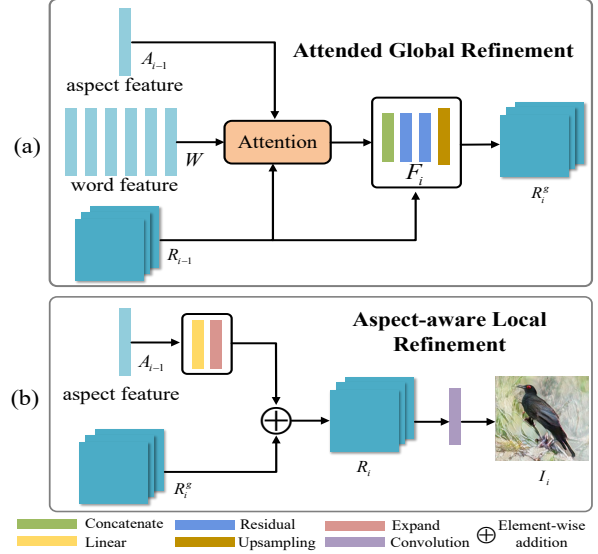


Figure 3. Architecture of *AGR* and *ALR*.

Specifically, we use word-level text features to help the refinement process by taking into account the contribution of each word. Current works mainly update the word-level features by employing image features from the previous step to select important words with attention mechanism [22]. Differently, we make a further step to integrate both image features and aspect-level features to update and enhance word-level features, as depicted in Figure 3 (a). This process can be mathematically formulated as follows:

$$\begin{aligned} \mathbf{R}_i^g &= F_i(\mathbf{R}_{i-1}, \mathbf{W}_i^g), i = 1, 2, \dots, n, \\ \mathbf{W}_i^g &= \sum_{j=0}^{l-1} (\mathbf{W}_j \mathbf{U}) \alpha_{i,j}, \\ \alpha_{i,j} &= \text{softmax}((\mathbf{W}_j \mathbf{U} + \mathbf{A}_{i-1} \mathbf{V}) \mathbf{R}_{i-1}), \end{aligned} \quad (5)$$

where $\mathbf{R}_i^g \in \mathbb{R}^{d_r \times N_i}$ represents the image features enriched globally with image features $\mathbf{R}_{i-1} \in \mathbb{R}^{d_r \times N_{i-1}}$ and attended word-level features. N_i is the size of \mathbf{R}_i^g at the i^{th} step. $F_i(\cdot, \cdot)$ denotes the image feature transformer. $\mathbf{W}_i^g \in \mathbb{R}^{d_r \times N_{i-1}}$ means the attended global features. $\alpha_{i,j}$ stands for the attention weights. $\mathbf{U} \in \mathbb{R}^{d_w \times d_r}$ and $\mathbf{V} \in \mathbb{R}^{d_w \times d_r}$ are perception layers to convert word embedding \mathbf{W} and aspect embedding \mathbf{A} into an underlying common semantic space of visual features.

3.3.2 Aspect-aware Local Refinement

In the previous part, we have introduced how to utilize word-level features to refine images from the global perspective. However, the enhancement of some specific image local details has not been completed yet. As mentioned above, aspects contained in the text description could be significant for synthesizing the corresponding local image

details. To this end, as shown in Figure 3 (b), *ALR* is developed to refine images from a local perspective with aspect-level features.

Technically, we combine the aspect features \mathbf{A}_{i-1} with globally refined image features \mathbf{R}_i^g by element-wise addition as follows:

$$\mathbf{R}_i = \mathbf{R}_i^g + [\mathbf{A}_{i-1} \mathbf{V}] \otimes N_i, i = 1, 2, \dots, n, \quad (6)$$

where operation $A_i \otimes N_i = [A_i; A_i; \dots; A_i]$ means repeatedly concatenating A_i for N_i times. To synthesize a photo-realistic image, we finally introduce a 3×3 convolution filter to transform the refined image feature \mathbf{R}_i into image I_i at the i^{th} refinement operation in *ADR*. In summary, *AGR* and *ALR* are alternately applied. Meanwhile, aspect-level features are dynamically added at each refinement step in *ADR*.

3.4. Objective Function

To generate a photo-realistic image and ensure the semantic consistency between text description and the corresponding image simultaneously, we carefully design the loss function. During each step, the generator G (e.g., *ADR*) and the discriminator D are trained in an alternative fashion. To start with common practice, the objective loss function of each generator at each step is defined as follows:

$$L_{G_i} = -\frac{1}{2} \left[\underbrace{\mathbb{E}_{I_i \sim p_{G_i}} \log D_i(I_i)}_{\text{unconditional loss}} + \underbrace{\mathbb{E}_{I_i \sim p_{G_i}} \log D_i(I_i, T)}_{\text{conditional loss}} \right], \quad (7)$$

where the first unconditional loss term is derived from the discriminator in distinguishing between real and fake images. The second term is a conditional loss to make the synthesized image match the input sentence.

Traditionally, the conditional loss term consists of sentence-image and word-image pairs. Different from previous works, we introduce aspect information through the generation process. To ensure the generated images truly contain local fine-grained details that match the corresponding aspects, we also include an aspect-image matching pair in the conditional loss as follows:

$$D(I, T) = D(I, \mathbf{s})^{\beta_1} \cdot D(I, \mathbf{W})^{\beta_2} \cdot D(I, \mathbf{A})^{\beta_3}, \quad (8)$$

where $D(I, \mathbf{s})$, $D(I, \mathbf{W})$, and $D(I, \mathbf{A})$ calculate the matching degrees between the image and the sentence, word, and aspect, respectively.

Following [45, 36], we further utilize the DAMSM loss [36] to compute the matching degree between images and text descriptions, mathematically denoted as L_{DAMSM} . And the CA loss is defined as the Kullback-Leibler divergence between the standard Gaussian distribution and the Gaussian distribution of training text, i.e.,

$$L_{CA} = D_{KL}(\mathcal{N}(\mu(\mathbf{s}), \sum(\mathbf{s})) || \mathcal{N}(0, I)). \quad (9)$$

The final objective function of the generator networks is composed of the aforementioned three terms:

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

For adversarial learning, each discriminator D_i is trained to precisely identify the input image as real or fake by minimizing the cross-entropy loss. The adversarial loss for each discriminator D_i is defined as:

$$L_{D_i} = -\frac{1}{2} \left[\underbrace{\mathbb{E}_{I_i^{GT} \sim p_{GT}} \log D_i(I_i^{GT}) + \mathbb{E}_{I_i \sim p_{G_i}} \log D_i(I_i)}_{\text{unconditional loss}} \right. \\ \left. + \underbrace{\mathbb{E}_{I_i^{GT} \sim p_{GT}} \log D_i(I_i^{GT}, T) + \mathbb{E}_{I_i \sim p_{G_i}} \log D_i(I_i, T)}_{\text{conditional loss}} \right], \quad (11)$$

where the unconditional loss is responsible for distinguishing synthesized images from real ones and the conditional term determines whether the image matches the input sentence. I_i^{GT} is sampled from the real image distribution p_{GT} at the i^{th} step. The final objective function of the discriminator networks is $L_D = \sum_i L_{D_i}$.

4. Experiment

In this section, we will first introduce the experiment setup. Next, we will evaluate DAE-GAN on two publicly available and well-studied datasets. Then, visualization study as well as causality analysis will be discussed to show the effectiveness and interpretability of DAE-GAN.

4.1. Experiment Setup

Datasets. To demonstrate the capability of our proposed method, we conduct extensive experiments on the CUB-200 [32] and COCO [12] datasets, following previous text-to-image synthesis works [36, 45, 18, 38]. The CUB-200 dataset contains 200 bird categories with 8,855 training images and 2,933 test images. Each image in CUB-200 has 10 text captions. For the COCO dataset, it consists of a training set with 80k images and a test set with 40k images. There are 5 captions for each image in COCO.

Evaluation Metrics. Following [36, 45], for better comparison, we quantitatively measure the performance of DAE-GAN in terms of Inception Score (IS) [24], Fréchet Inception Distance (FID) [7], and R-precision [36].

We obtain IS by employing a pre-trained Inception-v3 network [27] to compute the KL-divergence between the conditional class distribution and the marginal class distribution. A large IS indicates generated images have a high diversity for all classes, and each of them could be clearly recognized as a specific class rather than an ambiguous one.

FID computes the Fréchet distance between the synthetic and real-world images based on the feature map output from

Table 1. Inception Score (higher is better) on different models.

Model	CUB-200	COCO
(1) GAN-INT-CLS [20]	2.88±0.04	7.88±0.07
(2) StackGAN [38]	3.70±0.04	8.45±0.03
(3) AttnGAN [36]	4.36±0.03	25.89±0.47
(4) MirroGAN [18]	4.54±0.17	26.47±0.41
(5) Huang et al. [9]	-	26.92±0.52
(6) DM-GAN [45]	4.75±0.07	30.49±0.57
(7) LostGAN [26]	-	13.8±0.4
(8) MA-GAN [37]	4.76±0.09	-
(9) KT-GAN [28]	4.85±0.04	31.67±0.36
(10) DF-GAN [29]	4.86±0.04	-
(11) RiFe-GAN [4]	5.23±0.09	-
(12) DAE-GAN	4.42±0.04	35.08±1.16

the pre-trained Inception_v3 network. Lower FID score means a closer distance between the generated image distribution and real image distribution and therefore implies the model is capable of synthesizing photo-realistic images.

R-precision is used to evaluate the semantic consistency between the synthetic image and the given text description. Similarly, we calculate the cosine distance between the global image vector and 100 candidate global sentence vectors to measure the image-text semantic similarity. The lower R-precision means better semantic consistency between synthesized images and given text descriptions.

Implementation Details. * For aspect rules, each (**adjective**, **noun**) pair is an aspect that describes an object or scene. For COCO that contains layout and location, if a **preposition** is before the pair that expresses a relative spacial relationship, we will also add it. Consistent with [36, 45], we adopt Inception-v3 [27] pre-trained on ImageNet [23] as image encoder, and use pre-trained LSTM [36] as text encoder. The size of the initial generated low-resolution image (*i.e.*, N_0) is set to 64×64 . The finally synthesized high-resolution image at the last step has the size (N_n) of 256×256 . During intermediate steps, all the image size (N_i) is fixed to 128×128 . Empirically, we set $d_w = 256$ and $d_r = 64$ to be the dimensions of text and image feature vectors, respectively. For other related hyperparameters, we set $(\beta_1, \beta_2, \beta_3) = (1, 1, 0.2)$. For CUB-200, we set $(\lambda_1, \lambda_2, n) = (1, 5, 2)$, and for COCO, $(\lambda_1, \lambda_2, n) = (1, 50, 3)$. During training, we use the Adam optimizer with the learning rate of 0.0002 to train the networks on 8 NVIDIA Tesla V100 GPUs in parallel with the batch size of 32 on each one. DAE-GAN is trained with 600 epochs and 120 epochs on CUB-200 and COCO respectively.

4.2. Quantitative Results

Performance on CUB-200. We compare our methods with state-of-the-art methods on CUB-200. The overall results are summarized in Table 1, 2, and 3.

It is clear that our proposed DAE-GAN achieves highly comparable performance, especially for FID and R-

Table 2. FID score (lower is better) on different models.

Model	CUB-200	COCO
(1) AttnGAN [36]	23.98	35.49
(2) Huang et al. [9]	-	34.52
(3) MA-GAN [37]	21.66	-
(4) DM-GAN [45]	<u>16.09</u>	32.64
(5) KT-GAN [28]	17.32	30.73
(6) LostGAN [26]	-	29.65
(7) DF-GAN [29]	19.24	<u>28.92</u>
(8) DAE-GAN	15.19	28.12

Table 3. R-precision (%) (higher is better) on different models.

Model	CUB-200	COCO
(1) AttnGAN [36]	67.82±4.43	72.31±0.91
(2) MirroGAN [18]	57.67	74.52
(3) DM-GAN [45]	<u>72.31±0.91</u>	88.56±0.28
(4) Huang et al. [9]	-	89.69±4.34
(5) RiFeGAN [4]	23.8±1.5	-
(6) DAE-GAN	85.45±0.57	92.61±0.50

precision scores that measure photo realism and semantic consistency respectively. Specifically, DAE-GAN first learns comprehensive text semantics from multiple granularities, *i.e.*, sentence-level, word-level, as well as aspect-level. This is one of the reasons that DAE-GAN could improve FID and R-precision scores against other baselines by a large margin. Moreover, *ADR*, the core component of DAE-GAN, is developed to refine images by alternately applying *AGR* and *ALR* in a dynamic manner, in which *AGR* enhances images from the global perspective with word-level features and *ALR* refines images from the local perspective with aspect-level features. This is another vital reason to make our synthesized image more photo-realistic and keep semantic consistency between text and image.

CUB-200 is a dataset full of description details. Therefore, models with comprehensive text semantic understanding achieve much better results than coarse-grained ones. For example, GAN-INT-CLS and StackGAN only leverage sentence-level features as input. Based on them, AttnGAN and DM-GAN employ word-level features to refine images and achieve higher performance. RiFeGAN is particularly designed for datasets with fine-grained visual details, *e.g.*, CUB-200. It synthesizes images from multiple captions, which on the one hand leads to very high IS due to more caption details, on the other hand causes low R-precision score. Differently, our DAE-GAN can synthesize images with high FID and R-precision scores only from one given caption. This is largely achieved by comprehensive representation and utilization of text information, including sentence-level, word-level as well as aspect-level.

Performance on COCO. We also evaluate our method on COCO that has multiple objects, complex layout and simple details. The relative results are reported in Table 1, 2 and 3. We also list the observations as follows:

DAE-GAN still achieves the best quantitative performance against baseline methods with regard to IS, FID and R-precision. The results demonstrate that DAE-GAN is

*<https://github.com/hisarsal/DAE-GAN>

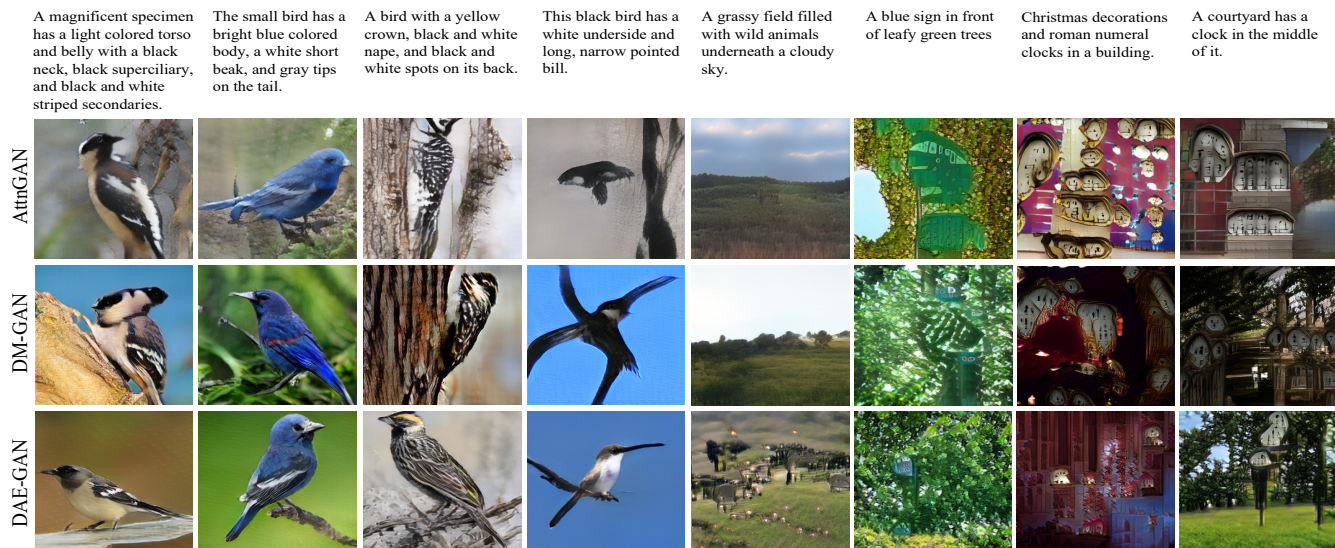


Figure 4. Example results for text-to-image synthesis by AttnGAN [36], DM-GAN [45] and our proposed DAE-GAN on CUB-200 (the left four ones) and COCO (the right four ones).

also capable of synthesizing well semantic consistent images with multi-object and complex layout. Comprehensive understanding of conditional text description and the newly proposed refinement paradigm *ADR* make it possible for DAE-GAN to refine different objects with dynamically provided aspect information. It is also the main reason that DAE-GAN can generalize well on different datasets.

4.3. Qualitative Results

To evaluate the visual quality of generated images, we show some subjective comparisons among AttnGAN [36], DM-GAN [45] and our proposed DAE-GAN in Figure 4.

In CUB-200, we can obtain that DAE-GAN generates better results. For example, when synthesizing a bird with the detail of long narrow bill (4th column), only DAE-GAN achieves this goal. Also in the 1st and 3rd columns, only DAE-GAN synthesizes photo-realistic and semantic consistency images. The reason is that DAE-GAN obtains comprehensive text representations, especially aspect-level features. Moreover, *ALR* is developed to dynamically enhance image details with aspect information.

In COCO dataset, we can also observe that images generated by DAE-GAN are more vivid and realistic. Taking examples in 6th and 7th columns of Figure 4, AttnGAN and DM-GAN often generate one object multiple times and the spatial distribution is also chaotic [36, 45], while DAE-GAN could address the problems well. By alternately applying *AGR* and *ALR*, DAE-GAN will not only enhance local details, but also refine images from a global perspective. This mechanism allows DAE-GAN to avoid getting stuck in a few most important words like other methods.

Table 4. Ablation performance on COCO about IS, FID and R-precision (%).

Model	IS	FID	R-precision
DAE-GAN (w/o AGR)	2.93±0.03	149.79	2.34±0.26
DAE-GAN (w/o ALR)	31.07±0.70	32.93	90.24±0.39
DAE-GAN (w/o asp in AGR)	34.70±0.64	28.60	92.28±0.46
DAE-GAN	35.08±1.16	28.12	92.61±0.50

4.4. Ablation Study

The overall experiments have proved the superiority of our proposed DAE-GAN. However, which component is really important for performance improvement is still unclear. Therefore, we perform an ablation study on COCO to verify the effectiveness of each part in *ADR*, including *AGR* and *ALR*. Corresponding results are illustrated in Table 4. According to the results, we can observe varying degrees of model performance decline when removing *AGR* and *ALR* separately from DAE-GAN. Recalling the aspect, since *ALR* is dependent on aspect information, we further remove aspect from *AGR*. The model performance also declines. The ablation study demonstrates comprehensive utilization of text information is helpful for image synthesis. *AGR* and *ALR* could employ the information well for image refinement.

4.5. Causality Interpretation

Visualization of Generation Process. To evaluate model rationality and interpretability, we study the synthesis process in Figure 5. In the left example, DAE-GAN initially generates a low-resolution image (*i.e.*, I_0) with the whole sentence. Then, based on I_0 , *ADR* further employs fine-grained information (*i.e.*, word-level and aspect-level features) to refine the image. Specifically, in image I_1 , *ADR* pours its attention to the refinement respecting the aspect information of ‘a metallic blue-black back’ and



Figure 5. Text-to-image synthesis visualization of different generation steps.

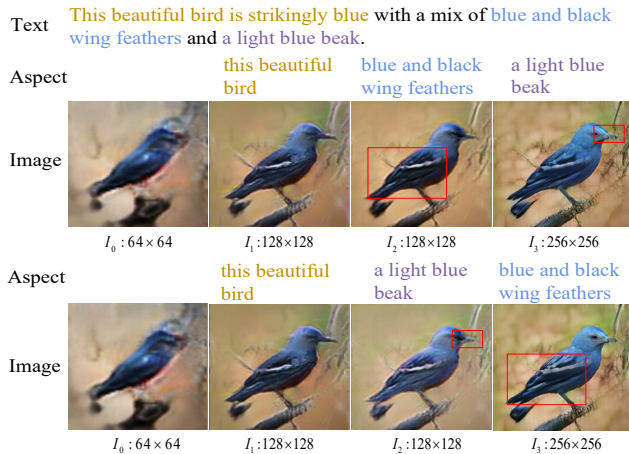


Figure 6. Cause-and-effect study of DAE-GAN. We exchange the input order of the last two aspect information to explore how the images generated in the related steps will correspondingly change with the variation of aspect input orders.

improves the image size to 128×128 . Finally in I_2 , *ADR* focuses on the refinement of aspect ‘an orange throat’ and further improves the resolution to 256×256 . Moreover, we can observe that in I_1 , almost the bird head is in orange, while in I_2 , only the throat is correctly drawn in orange and the image looks more vivid. In the right example, we present a visualization study in COCO with three aspects. We could observe that the images are also well generated under the corresponding guide of aspects ‘a grassy field’, ‘with wild animals’ and ‘underneath a cloudy sky’.

Influence of Aspect Order. As aspect information is placed in a significant position in our model, and it is dynamically utilized to refine local details, we are curious whether the order of aspect input will affect the generation results. Thus, an example case is studied in Figure 6. To be specific, the given text description has three aspect features as shown in Figure 6. We exchange the input order of the last two aspect features to explore how the images generated at the related steps will correspondingly change with the variation of aspect input orders. Meanwhile, we will not change the input of sentence-level features and word-level features. Taking the look at the examples of upper part, when refining the aspect of ‘blue and black wing

feathers’, we could find the bird feathers are vividly drawn in blue and black. Then in I_3 , it is obvious that the color of bird beak is changed from light black in I_2 to light blue. In the examples of bottom part, *ADR* focuses on the aspect ‘a light blue beak’ in I_2 . Visually compared with I_1 , the bird beak color turned light blue from light black. Then in I_3 , the feathers are refined more vividly in blue and black than the ones in I_2 .

The experimental results again fully confirmed the importance of aspect information to image refinement. Meanwhile, DAE-GAN can make full use of the aspect information to achieve image refinement in a dynamic manner. Furtherly, these examples also illustrate that our proposed DAE-GAN has good interpretation.

5. Conclusion

In this paper, we argued that aspect information contained in the text is quite helpful for image generation and should gain more attention. Then, we developed a novel DAE-GAN to make full use of aspect information for text-to-image synthesis. To be specific, we utilized text information from multiple granularities, including sentence-level, word-level and aspect-level. Moreover, a new generation paradigm *ADR* was developed to refine initial images, in which a novel *AGR* was proposed to refine images from the global perspective and a novel *ALR* was designed to enhance image details from the local perspective. By employing these two components dynamically, our proposed DAE-GAN had the ability to leverage aspect information to refine the details of the generated image, which is critical for image realism and semantic consistency. Extensive experiments demonstrated the superiority and rationality of our proposed method. In the future, we will pay a special focus on exploring a way to improve semantic consistency with self-supervised contrastive learning.

Acknowledgment

This research was partially supported by grants from the National Natural Science Foundation of China (Grants No. 61727809, 61922073, 62006066, and U20A20229), and the Open Project Program of the National Laboratory of Pattern Recognition (NLPR).

References

- [1] Th Brandt, Johannes Dichgans, and Ellen Koenig. Differential effects of central versus peripheral vision on egocentric and exocentric motion perception. *Experimental brain research*, 16(5):476–491, 1973. 2
- [2] Kevin Chen, Christopher B Choy, Manolis Savva, Angel X Chang, Thomas Funkhouser, and Silvio Savarese. Text2shape: Generating shapes from natural language by learning joint embeddings. In *Asian Conference on Computer Vision*, pages 100–116. Springer, 2018. 1
- [3] Zhuang Chen and Tiejun Qian. Relation-aware collaborative learning for unified aspect-based sentiment analysis. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3685–3694, 2020. 2
- [4] Jun Cheng, Fuxiang Wu, Yanling Tian, Lei Wang, and Dapeng Tao. Rifegan: Rich feature generation for text-to-image synthesis from prior knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10911–10920, 2020. 1, 2, 6
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 1
- [6] Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Jimenez Rezende, and Daan Wierstra. Draw: A recurrent neural network for image generation. *arXiv preprint arXiv:1502.04623*, 2015. 2
- [7] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, pages 6626–6637, 2017. 5
- [8] Seunghoon Hong, Dingdong Yang, Jongwook Choi, and Honglak Lee. Inferring semantic layout for hierarchical text-to-image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7986–7994, 2018. 3
- [9] Wanming Huang, Richard Yi Da Xu, and Ian Oppermann. Realistic image generation using region-phrase attention. In *Asian Conference on Machine Learning*, pages 284–299. PMLR, 2019. 3, 6
- [10] Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image generation from scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1219–1228, 2018. 3
- [11] Wenbo Li, Pengchuan Zhang, Lei Zhang, Qiuyuan Huang, Xiaodong He, Siwei Lyu, and Jianfeng Gao. Object-driven text-to-image synthesis via adversarial training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12174–12182, 2019. 3
- [12] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 5
- [13] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *Advances in neural information processing systems*, pages 406–416, 2017. 3
- [14] Elman Mansimov, Emilio Parisotto, Jimmy Lei Ba, and Ruslan Salakhutdinov. Generating images from captions with attention. *arXiv preprint arXiv:1511.02793*, 2015. 2
- [15] Anh Nguyen, Jeff Clune, Yoshua Bengio, Alexey Dosovitskiy, and Jason Yosinski. Plug & play generative networks: Conditional iterative generation of images in latent space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4467–4477, 2017. 2
- [16] Dim P Papadopoulos, Youssef Tamaazousti, Ferda Ofli, Ingmar Weber, and Antonio Torralba. How to make a pizza: Learning a compositional layer-based gan model. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8002–8011, 2019. 3
- [17] Haiyun Peng, Lu Xu, Lidong Bing, Fei Huang, Wei Lu, and Luo Si. Knowing what, how and why: A near complete solution for aspect-based sentiment analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8600–8607, 2020. 2
- [18] Tingting Qiao, Jing Zhang, Duanqing Xu, and Dacheng Tao. Mirrorgan: Learning text-to-image generation by redescription. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1505–1514, 2019. 1, 2, 5, 6
- [19] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092*, 2021. 3
- [20] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *33rd International Conference on Machine Learning*, pages 1060–1069, 2016. 1, 2, 6
- [21] Scott E Reed, Aäron van den Oord, Nal Kalchbrenner, Sergio Gomez Colmenarejo, Ziyu Wang, Yutian Chen, Dan Belov, and Nando de Freitas. Parallel multiscale autoregressive density estimation. In *ICML*, 2017. 2
- [22] Shulan Ruan, Kun Zhang, Yijun Wang, Hanqing Tao, Weidong He, Guangyi Lv, and Enhong Chen. Context-aware generation-based net for multi-label visual emotion recognition. In *2020 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2020. 4
- [23] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 6
- [24] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29:2234–2242, 2016. 5
- [25] Hans Strasburger, Ingo Rentschler, and Martin Jüttner. Peripheral vision and pattern recognition: A review. *Journal of vision*, 11(5):13–13, 2011. 2

- [26] Wei Sun and Tianfu Wu. Image synthesis from reconfigurable layout and style. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10531–10540, 2019. 3, 6
- [27] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 5, 6
- [28] Hongchen Tan, Xiuping Liu, Meng Liu, Baocai Yin, and Xin Li. Kt-gan: Knowledge-transfer generative adversarial network for text-to-image synthesis. *IEEE Transactions on Image Processing*, 30:1275–1290, 2020. 6
- [29] Ming Tao, Hao Tang, Songsong Wu, Nicu Sebe, Fei Wu, and Xiao-Yuan Jing. Df-gan: Deep fusion generative adversarial networks for text-to-image synthesis. *arXiv preprint arXiv:2008.05865*, 2020. 2, 6
- [30] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. *Advances in neural information processing systems*, 29:4790–4798, 2016. 2
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017. 3
- [32] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 5
- [33] Yequan Wang, Minlie Huang, Xiaoyan Zhu, and Li Zhao. Attention-based lstm for aspect-level sentiment classification. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pages 606–615, 2016. 2
- [34] William H Warren and Kenneth J Kurtz. The role of central and peripheral vision in perceiving the direction of self-motion. *Perception & Psychophysics*, 51(5):443–454, 1992. 2
- [35] Weihao Xia, Yujiu Yang, Jing-Hao Xue, and Baoyuan Wu. Tedigan: Text-guided diverse image generation and manipulation. *arXiv preprint arXiv:2012.03308*, 2020. 3
- [36] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1316–1324, 2018. 1, 2, 5, 6, 7
- [37] Yanhua Yang, Lei Wang, De Xie, Cheng Deng, and Dacheng Tao. Multi-sentence auxiliary adversarial networks for fine-grained text-to-image synthesis. *IEEE Transactions on Image Processing*, 30:2798–2809, 2021. 2, 6
- [38] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 5907–5915, 2017. 1, 2, 3, 5, 6
- [39] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan++: Realistic image synthesis with stacked generative adversarial networks. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1947–1962, 2018. 1
- [40] Yaowei Zheng, Richong Zhang, Samuel Mensah, and Yongyi Mao. Replicate, walk, and stop on syntax: An effective neural network model for aspect-level sentiment classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9685–9692, 2020. 2
- [41] Jiale Zhi. Pixelbrush: Art generation from text with gans. In *Cl. Proj. Stanford CS231N Convolutional Neural Networks Vis. Recognition, Sprint 2017*, page 256, 2017. 1
- [42] Xingran Zhou, Siyu Huang, Bin Li, Yingming Li, Jiachen Li, and Zhongfei Zhang. Text guided person image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3663–3672, 2019. 3
- [43] Bin Zhu and Chong-Wah Ngo. Cookgan: Causality based text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5519–5527, 2020. 3
- [44] Bin Zhu, Chong-Wah Ngo, Jingjing Chen, and Yanbin Hao. R2gan: Cross-modal recipe retrieval with generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11477–11486, 2019. 3
- [45] Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. Dm-gan: Dynamic memory generative adversarial networks for text-to-image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5802–5810, 2019. 1, 2, 5, 6, 7