This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# **Low-Biased General Annotated Dataset Generation**

Dengyang Jiang<sup>1\*</sup> Haoyu Wang<sup>1\*</sup> Lei Zhang<sup>1†</sup> Wei Wei<sup>1</sup> Guang Dai<sup>2</sup> Mengmeng Wang<sup>3</sup> Jingdong Wang<sup>4</sup> Yanning Zhang<sup>1</sup>

<sup>1</sup> Northwestern Polytechnical University <sup>2</sup> SGIT AI Lab, State Grid Corporation of China <sup>3</sup> Zhejiang University of Technology <sup>4</sup> Baidu Inc.



Figure 1. **Performance on eight transfer learning datasets.** Pre-training on our generated low-biased general annotated dataset can bring stable generalization capacity enhancement of different backbones.

#### Abstract

Pre-training backbone networks on a general annotated dataset (e.g., ImageNet) that comprises numerous manually collected images with category annotations has proven to be indispensable for enhancing the generalization capacity of downstream visual tasks. However, those manually collected images often exhibit bias, which is non-transferable across either categories or domains, thus causing the model's generalization capacity degeneration. To mitigate this problem, we present a low-biased general annotated dataset generation framework (lbGen). Instead of expensive manual collection, we aim at directly generating low-biased images with category annotations. To achieve this goal, we propose to leverage the advantage of a multimodal foundation model (e.g., CLIP), in terms of aligning images in a low-biased semantic space defined by language. Specifically, we develop a bi-level semantic alignment loss, which not only forces all generated images to be consistent with the semantic distribution of all categories belonging to the target dataset in an adversarial learning manner, but also requires each generated image to match the semantic description of its category name. In addition, we further cast an existing image quality scoring model into a quality assurance loss to preserve the quality of the generated image. By leveraging these two loss functions, we can obtain a low-biased image generation model by simply fine-tuning a pre-trained diffusion model using only all category names in the target dataset as input. Experimental results confirm that, compared with the manually labeled dataset or other synthetic datasets, the utilization of our generated low-biased dataset leads to stable generalization capacity enhancement of different backbone networks across various tasks, especially in tasks where the manually labeled samples are scarce. Code is available at: https://github.com/vvvvvjdy/lbGen

<sup>\*</sup>Equal contribution.

<sup>&</sup>lt;sup>†</sup>Corresponding author



Figure 2. Visualization of some randomly sampled images from 4 datasets. It is hard to tell from which dataset exhibits low bias through these images. However, models trained on these four datasets demonstrate a significant disparity in their generalization capabilities.

### **1. Introduction**

Deep neural networks have achieved great success in various computer vision tasks [31, 34, 58]. One indispensable premise of such success lies on pre-training the parameterextensive backbone network using a general annotated dataset (e.g., ImageNet [6]) that contains a large number of images with manually annotated categories. Profiting from the vast amount of annotated images and the diverse image categories, the pre-trained backbone networks often show pleasing generalization capacity and perform effectively in the target computer vision task through simple fine-tuning it together with a corresponding task head with a few parameters [15, 53]. Unfortunately, recent studies [27, 47, 57] uncover that these manually collected images often exhibit nontravail bias<sup>1</sup> (e.g., a certain background, image style, object position for a specific category etc.) which can be easily captured by backbone networks during pre-training, but hardly noticed by human collectors (see Figure 2). Such hidden bias is proven to be cast into a shortcut feature representation [12] to improve the in-domain performance but deteriorate the generalization capacity of pre-trained backbone networks on target tasks in the cross-category or cross-domain settings [4, 10] which shows an obvious image distribution gap from the utilized general annotated dataset. For example, when a specific category of images often shows a similar background, the pre-train backbone networks will consider the background as the discriminative feature of such a category while overlooking the cross-category or cross-domain

transferable semantic features (e.g., shape, structure, etc.). Therefore, it is crucial to obtain a low-biased general annotated dataset to enhance the cross-category or cross-domain generalization capacity of the pre-trained backbones.

To this end, a straight solution is to manually re-collect extensive low-biased images. However, it will not only produce expensive manual collection costs, but also inevitably incur some other undetectable bias. Recently, diffusion models [38, 40, 41] have shown powerful capacity in terms of generating high-quality synthetic images based on the text description of image contents, thus providing a feasible way to directly generate images with annotations without manual collection cost. Moreover, some studies [1, 42, 56] have demonstrated that those randomly generated images with annotations can be utilized for network training. Although most existing diffusion models can be directly utilized for general annotated dataset generation, they mainly focus on generating images with the distribution consistent with the conventional manually annotated general dataset (e.g., ImageNet) and scarcely attempt to generate low-biased images. Thus, pre-training the backbone networks on these generated general annotated datasets will not bring non-travail generalization capacity enhancement [44].

To mitigate this problem, we present a low-biased general annotated dataset **gen**eration framework (**lbGen**), which takes the first attempt to directly generating synthetic lowbiased images with category annotations. To achieve this goal, we first have to define a low-biased space where the feature representation of each image emphasizes transferable semantic characteristics. Considering that the observation that text information is closer to the ideal semantic information and the recent progress of multimodal foundation model

<sup>&</sup>lt;sup>1</sup>In this paper, 'dataset bias' refers to 'systematic bias introduced in data collection, selection, or processing that impair the generalization capacity of the model'.

(e.g., CLIP [37]) which aims at mapping images into a lowbiased semantic space defined by language, a straight idea is to constrain the image output of the existing diffusion models to follow the semantic distribution of the specific image category in such a low-biased semantic space. Following this idea, we develop a bi-level semantic alignment loss based on the CLIP model to fine-tune the pre-trained diffusion model. In a specific, on the one hand, such a loss forces all generated images to be consistent with the semantic distribution of all categories belonging to the target dataset in the CLIP feature space using an adversarial learning scheme. On the other hand, it also requires each generated image to match the semantic description of its category name in the CLIP feature space using a simple cosine similarity metric. By doing these, we can obtain a low-biased image generation model by simply fine-tuning a pre-trained diffusion model using only all category names in the target dataset as input. In addition, to sidestep image quality degradation caused by the low-biased image generation learning, we further cast an existing image quality scoring model into a quality assurance loss to assist the bi-level semantic alignment loss for diffusion model fine-tuning.

To testify the efficacy of the proposed framework, we pre-train two conventional backbone networks on our generated low-biased general annotated dataset and training specific heads on different downstream tasks. Compared with backbone networks pre-trained either on the manually collected generated dataset or that generated by existing diffusion models, our approach achieves obvious generalization performance improvement, especially when the manually annotated samples in the target task are scarce. Moreover, additional experiments prove that our pre-trained backbone networks capture lower specific bias (e.g. context, background, shape-texture), which further demonstrates the generality of our framework.

In summary, our main contributions are as follows:

- We propose the first low-biased general annotated dataset generation framework, which jumps out of the dilemma of traditional manual data collection in terms of mitigating dataset bias.
- We present a bi-level semantic alignment module assisted by a quality assurance loss to simply fine-tune the standard diffusion model using only all category names in the target dataset as input.
- With our generated general low-biased dataset, the pretrained backbone network shows state-of-the-art generalization capacity in different downstream tasks.

### 2. Related Work

#### 2.1. Datasets Bias

Since the deep learning revolution in 2012 [22], the largescale manually collected annotated dataset (e.g., ImageNet)

performs no longer a simple training dataset for its own tasks, but a general dataset utilized for backbone network pre-training which has become the indelible step for enhancing the generalization performance of various downstream tasks. However, recent studies [27, 47, 57] have consecutively revealed that these existing manually collected general dataset exhibit non-trial bias, which results in sup-optimal cross-categories and cross-domain generalization capacity, especially when the manually annotated samples in the target task are scarce. For example, Liu et al. [27] observe that deep neural networks can achieve excellent accuracy in classifying which dataset an image is from. In other words, the neural network discovers some dataset-specific patterns, a form of bias. In addition, studies in [17, 34, 44] attempt to implicitly measure the dataset bias by investigating the crosscategory or cross-domain generalization capacity as well as the robustness of the models pre-trained on the dataset.

However, these works mainly focus on raising the dataset bias problem or bias measurement. In contrast, in this study, we take the first attempt to solve this problem and aims at borrowing the advantage of diffusion model in image generation to directly generate a low-biased general dataset for better backbone pre-training.

### 2.2. Synthetic Dataset Generation

Different from manual collection, synthetic dataset generation aims at directly generating image using deep neural networks based on some text description. For example, in early days, Zhu et al. [54, 60] utilized adversarial generative networks to model the mapping relationship between the input text description and the output image. However, these methods require large-scale high-quality images from target categories for network training, and the generalization capacity to unknown text descriptions is limited. More recently, as diffusion model shows more powerful capacity in generalization to unknown text description [9, 35, 40], some works have attempted to utilize the diffusion model to generate ImageNet-like synthetic dataset for backbone pre-training. For example, Bansal et al. [1] directly fine-tune the diffusion model on ImagNet-1K [6] and use meticulously designed prompts to generate the dataset. Lei et al. [24] utilizes ViT-GPT2 [23] to get a unique prompt to generate each image. Yuan et al. [56] resort to learn the distribution of ImageNet and use Blip-captions [25] of ImageNet as prompts to synthesize the dataset.

Although these diffusion-based methods can be utilized for general dataset generation, they mainly focus on simulating the existing ImageNet without considering the dataset bias. In this study, we attempt to fine-tune the diffusion model to directly generate a low-biased general annotated dataset without using any image-text pairs but the category names of the target dataset as input.



Figure 3. **Overview of our training method.** The generator first generates an image according to the class name. Then the image is sent to bi-level semantic guidance module and quality assurance module respectively for loss calculation.

# 3. Approach

The overall training framework of lbGen is shown in Figure 3. In Section 3.1, we begin with elucidating the basic methodology employed in our training process. In Section 3.2, we then illustrate the bi-level semantic alignment module, which is the core part of our approach. Subsequently, in Section 3.3, we further introduce the quality assurance module for fidelity preserving the images, and finally we integrate the two components for joint learning.

#### **3.1. Preliminary**

We implement our method on the leading text-to-image diffusion model, Stable Diffusion [40], which belongs to the family of latent diffusion models (LDM) [39]. In the traditional training process, a normally distributed noise  $\epsilon$  is added to the original latent code  $z_0$  with a variable extent based on a timestep t sampling from  $\{1, ..., T\}$ . Then, a denoising function  $\epsilon_{\theta}$ , parameterized by a UNet backbone, is trained to predict the noise added to  $z_0$  with the text prompt  $\mathcal{P}_t$  and the current latent  $z_t$  as the input. Specifically, the text prompt is first encoded by a text encoder W, then incorporated into the denoising function  $\epsilon_{\theta}$  by the cross-attention mechanism. The denoising loss in diffusion models' training is formally expressed as:

$$\mathcal{L}_{\text{LDM}} = \mathbb{E}_{z_0, t, p, \epsilon \sim \mathcal{N}(0, I)} \left[ \left\| \epsilon - \epsilon_{\theta} \left( z_t, t, W(\mathcal{P}_t) \right) \right\|^2 \right].$$
(1)

For inference, the process can be formulated as a Markov decision process that iteratively estimates the noise and computes the next latent sample:

$$p_{\theta}(z_0|W(\mathcal{P}_t)) = p(z_t) \prod_{t=1}^T p_{\theta}(z_{t-1}|z_t, W(\mathcal{P}_t)).$$
 (2)

However, this method requires vast quantities of imagetext pairs for training and it takes extended times to converge, which is not feasible to fine-tune a model under our restrictive conditions. To overcome these difficulties, we follow some works [3, 19, 36] to fine-tune the diffusion model with reinforcement learning (RL). Different from the original loss function in Eq. (1), given a reward function R(.), the objective of RL is to maximize the expected reward:

$$J(\phi) = \mathbb{E}\left[R(Z_0, c]\right]. \tag{3}$$

With the denoising process showed in Eq. (2), the gradient when fine-tuning diffusion models with reward feedback Eq. (3) can be computed as:

$$\nabla_{\phi} J = \mathbb{E}\left[\sum_{t=0}^{T} \nabla_{\phi} \log p_{\theta}(z_{t-1}|z_t, W(\mathcal{P}_t)R(Z_0, c))\right].$$
 (4)

Noticing that only text prompt  $\mathcal{P}_t$  and conditional context c for reward models are required in this training paradigm instead of a dataset containing image-text pairs, which is consistent with the fact that we do not incorporate any external biased images in our training process.

#### 3.2. Bi-Level Semantic Alignment

As mentioned in Section 1, we assume that the semantic space defined by language can be a low-biased representation and our key insight is using this characteristic of language to fine-tune a pre-trained diffusion model as our lbGen generator. We achieve this by leveraging a simple Linear-ReLU-Linear based discriminator  $\mathcal{D}_{\phi}$  and utilizing CLIP to carry out a bi-level semantic alignment. We use only 1000 class names of ImageNet as our inputs, which ensures that

no other biased information is introduced except semantic information of the dataset.

Entire Dataset Alignment. Building on the advantages of CLIP which has unified the image and the text into one represention space, we can initially use CLIP text encoder to extract text features  $\{f_{c_1}, f_{c_2}, \dots, f_{c_{1000}}\}$  of classnames  $\{c_1, c_2, \ldots, c_{1000}\}$ . We consider these text features to be a low-biased semantic distribution of the entire ImageNet. Then, we generate an image using prompt  $c_i$  and send it to CLIP image encoder to get the image feature  $f_{im_i}$ . Next, we randomly choose a text feature  $f_{c_i}$  from  $\{f_{c_1}, f_{c_2}, \ldots, f_{c_{1000}}\}$ . It is worth emphasizing that we do not use the text feature which belongs to the same class as the image feature since we aim to align the whole synthetic dataset to its general semantic representation space regardless of the concrete class. Finally, features  $f_{im_i}$  and  $f_{c_i}$  are fed into  $\mathcal{D}_{\phi}$  for computing entire semantic alignment loss as follows:

$$\mathcal{L}_{en} = \log\left(\mathcal{D}_{\phi}\left(f_{c_{j}}\right)\right) + \log\left(1 - \mathcal{D}_{\phi}\left(f_{im_{i}}\right)\right). \quad (5)$$

Similar to the training object of Generative Adversarial Network(GAN) [13], we expect to fine-tune the diffusion model to minimize this adversarial loss, while concurrently training the discriminator to maximize it.

**Individual Image Alignment.** Expect for mapping the images to be consistent with the semantic distribution of all classes within the entire dataset, we need to precisely control each category of images to match their semantic description. To this end, we introduce the individual semantic alignment loss. In particular, given the generated image  $im_i$  using class name  $c_i$ , we use simple "photo of  $c_i$ " as the low-biased semantic description  $p_{c_i}$  and send them to CLIP. Different from using CLIP to align the dataset globally in the entire semantic space, we aim at forcing the semantic information of each image to be dovetailed with its class by maximizing the cosine similarity between the image and its corresponding semantic description through CLIP. Thus, we can obtain  $\mathcal{L}_{in}$  formulated as follows:

$$\mathcal{L}_{in} = 1 - \frac{f_{im_i} \cdot f_{p_{c_i}}}{\|f_{im_i}\| \cdot \|f_{p_{c_i}}\|},$$
(6)

where  $f_{im_i}$  and  $f_{p_{c_i}}$  represent the image and text feature vectors extracted by CLIP, the dot product of the vectors is denoted by  $\cdot$ , and  $\|\cdot\|$  denotes the norm of the vectors.

By considering these two distinct levels of losses, we can finally add them together and obtain bi-level semantic alignment loss  $\mathcal{L}_{bi}$  to refine the diffusion model to align more closely with the low-biased semantic reference.

### 3.3. Quality Assurance

In practice, only under supervision in terms of text semantics, we observe that the quality of the generated images is sub-optimal after training. Thus, we introduce the quality assurance loss to assist the bi-level semantic alignment loss.

To be specific, we use the state-of-the-art image quality scoring model Q-ALIGN [49] as our quality assurance model. After fine-tuning a Leading open-source multimodal large language model (MLLM) mPLUG-Owl-2 [55] on carefully collected image quality assessment datasets, Q-ALIGN can achieve satisfactory image quality scoring performance. Feeding the generated image  $im_i$  with the system prompt "How would you rate the quality of this image?" into Q-ALIGN, we can obtain the quality score  $Q(im_i)$ , which ranges in [1, 5], of the image and using it to calculate quality assurance loss ( $\mathcal{L}_q$ ) for diffusion model to optimize. Details about how Q-ALIGN scores images can be found in Appendix.

$$\mathcal{L}_q = 1 - \frac{Q(im_i)}{5}.$$
(7)

Finally, we combine the losses in the bi-level semantic alignment module and quality assurance module to build up our final training object for lbGen generator as follows:

$$\mathcal{L} = \mathcal{L}_{bi} + \lambda_1 \mathcal{L}_q,\tag{8}$$

where  $\lambda_1$  is a scaling factor to balance the losses. The pseudocode of the integrated loss computation process can be found in Appendix.

### 4. Experiment

#### 4.1. Experimental Settings

**Datasets.** We choose two recent open-source synthetic ImageNet (GenRobust [1], RealFake [56]) and ImageNet-1K [6] for comparison. We test the generalization ability and robustness of the pre-trained model using eight transfer learning datasets (Aircraft [30], Cars196 [21], DTD[5], EuroSAT[16], Flowers[32], Pets[33], Food101[2], SUN397[51]), two visual perception datasets (COCO [26], ADE20K [59]), as well as three specific bias measurement datasets (FOCUS[20], Mixed-Rand & Mixed-Same [52], Cue Conflict [11])

**Implementation Details.** We implement our method on SD1.5 [40] and fine-tune it with LoRA [18]. We choose openai-CLIP-VIT-L [37] as our default CLIP model. For the visual backbones, we choose two representative models, ConvNets-based ResNet50 [14] and Transformer-based ViT-S [8]. During fine-tuning the generator, we follow Deep Reward [50] and CoMat [19] that only enable gradients in 5 steps out of those 50 steps to save GPU memory. During training visual backbones, we maintain the same training hyperparameters across all selected datasets to make a fair comparison.

More details about training hypersettings, data synthesis, datasets for evaluation, and computing resources are provided in Appendix.

Backbone	Pre-Trained Data	IN-val	Aircraft	Cars196	DTD	EuroSAT	Flowers	Pets	Food101	SUN397	Avg.
	IN-Real [6]	76.2	56.0	51.5	70.0	93.7	81.4	90.7	<u>67.2</u>	56.8	71.5
	IN-SD1.5 [40]	45.8	58.3	52.7	69.0	94.1	82.1	85.5	63.6	54.4	67.3
ResNet50	IN-GenRobust [1]	43.4	56.2	47.4	68.0	94.8	80.4	83.3	57.8	49.3	64.5
	IN-RealFake [56]	69.8	59.9	54.1	70.6	94.4	83.7	90.0	67.3	55.9	71.8
	IN-lbGen(ours)	46.1	62.1	58.5	72.8	95.0	86.4	87.2	65.3	64.3	73.2
	IN-Real [6]	78.7	59.4	56.4	69.5	<u>94.1</u>	83.0	<u>90.2</u>	<u>68.3</u>	<u>57.2</u>	72.3
ViT-S	IN-SD1.5 [40]	46.6	57.5	51.8	68.3	92.7	<u>84.0</u>	85.6	62.8	56.1	69.9
	IN-GenRobust [1]	44.9	52.3	54.0	63.5	94.4	78.3	82.1	54.7	49.7	66.3
	IN-RealFake [56]	72.3	57.3	53.1	67.2	93.3	82.1	91.7	65.6	55.9	70.8
	IN-lbGen(ours)	46.3	62.6	58.0	71.2	<u>94.1</u>	86.2	88.6	68.5	66.0	74.4

Table 1. **Top-1 accuracy on transfer learning datasets.** The average accuracy across eight transfer learning datasets is denoted as Avg. The best and <u>second-best</u> transfer learning performance of each backbone are highlighted in red and underlined. IN-SD1.5 denotes only using original SD1.5 to generate the data. We also present results on ImageNet validation set for reference.



Figure 4. Scaling down the number of training images of eight transfer learning datasets. The benefits of using pre-trained models on our lbGen images are even more pronounced when there is less data for training.

### 4.2. Generalization across Downstream Tasks

**Transfer Learning.** Transfer learning [34] is a widely known downstream visual task and can be significantly influenced by the generalization of the pre-trained model. In our work, we aim to indicate whether the utilization of our lbGen data can enable the backbones to learn better transferable patterns. To this end, we follow fakeit [42], which uses pre-trained visual backbones as feature extractors and train simple linear logistic regression classifiers from scratch.

As illustrated in Table 1, we observe that models pretrained on our dataset outperform all other candidates. Compared with the second-best synthetic datasets, we achieve +1.4% and +3.6% leading performance on ResNet50 and ViT-S respectively. More importantly, our method exhibits 1.7% and 2.1% average accuracy improvement and superior results in the vast majority of transfer learning datasets compared to real data. Furthermore, we investigate the transfer ability of models pre-trained on real images and our generated images when using less transfer learning training data. Such few-shot setting [48] requires an even higher general-

Pre-Trained Data	$COCO (AP^{box})$							
	1.0×	0.5  imes	$0.2 \times$	$0.1 \times$				
IN-Real [6]	39.32	34.97	29.14	25.51				
IN-SD1.5 [40]	38.89	34.68	28.60	24.05				
IN-GenRobust [1]	38.12	32.11	27.68	23.38				
IN-RealFake [56]	39.04	35.09	29.25	24.88				
IN-lbGen(ours)	<u>39.26</u>	35.24	30.68 25.64					
Pre-Trained Data		ADE20k	K (mIoU)	)				
Pre-Trained Data		ADE20 <b>k</b> 0.5×	K (mIoU) 0.2×	0.1×				
Pre-Trained Data IN-Real [6]		$\frac{\text{ADE20F}}{0.5\times}$ $\frac{38.05}{2}$	X (mIoU) 0.2× 32.10	0.1× 27.64				
Pre-Trained Data IN-Real [6] IN-SD1.5 [40]	$ \begin{vmatrix} 1.0 \times \\ 42.44 \\ 41.07 \end{vmatrix} $	ADE20 $k$ 0.5× $\frac{38.05}{37.62}$	X (mIoU) 0.2× 32.10 31.49	$0.1 \times$ 27.64 26.32				
Pre-Trained Data IN-Real [6] IN-SD1.5 [40] IN-GenRobust [1]	$ \begin{array}{ } \\ \hline 1.0 \times \\ \hline 42.44 \\ 41.07 \\ 40.77 \end{array} $	$     ADE20k     0.5 \times     \overline{)38.05}     37.62     37.13   $	X (mIoU) 0.2× 32.10 31.49 29.36	$     \begin{array}{r}       0.1 \times \\       \underline{27.64} \\       26.32 \\       24.70     \end{array} $				
Pre-Trained Data IN-Real [6] IN-SD1.5 [40] IN-GenRobust [1] IN-RealFake [56]	$\begin{array}{ c c c c c }\hline & & & \\ \hline & 1.0 \times \\ \hline & 1.0 \times \\ \hline & 42.44 \\ 41.07 \\ 40.77 \\ \hline & 40.77 \\ \hline & 41.89 \\ \hline \end{array}$	$     ADE20F     0.5 \times     \overline{)38.05}     37.62     37.13     37.76     $	$   \begin{array}{r}                                     $	$ \begin{array}{r} 0.1 \times \\ \underline{27.64} \\ 26.32 \\ 24.70 \\ 27.38 \end{array} $				

Table 2. Results on COCO object detection and ADE20K semantic segmentation of different number of training images. We gradually scaling down the number of downstream training images from original data size to 1/10 of it for testing the generalization ability of the backbones.

ization capacity of the model. As shown in Figure 4, it is important to note that the advantage of pre-trained models on our data is even greater when there are fewer downstream images for training. This phenomenon further underscores the lower bias of our data and the stronger generalization of the resulting model. Meanwhile, another striking finding is that achieving high accuracy on the ImageNet validation set does not necessarily correlate with enhanced cross-category generalization performance, thereby enabling us to draw more definitive conclusions regarding the dualistic impact of bias in the existing dataset on the model.

**Visual Perception.** Detection and segmentation are two of the most popular downstream visual perception tasks, at the same time, these two tasks can benefit a lot from a well pre-trained backbone. Hence, we want to find out if the utilization of our lbGen data can also enhance the performance on these tasks. To this end, we follow previous studies [28, 29, 43], which use Mask R-CNN [15] as detection head for COCO object detection and UperNet [53] as segmentation head for ADE20K semantic segmentation, to evaluate the performance of pre-trained ResNet50 backbone on different datasets. Moreover, to thoroughly test the generalization ability of the backbone, we progressively decrease the number of training samples and observe the outcomes.

The experimental results in Table 2 demonstrate the effectiveness of our lbGen datasets. Although pre-training on real data achieves marginally better performance with full training data, the performance of lbGen pre-trained model consistently outperforms all other pre-training data when downstream data is limited. Specifically, with only 20% of the original training data, the model achieves the highest performance across both tasks, showing gains of up to 1.54% AP<sup>box</sup> and 1.47% mIoU compared to that pre-trained on IN-Real. This result is particularly valuable for real-world applications where collecting and annotating task-specific training data is often costly and time-consuming.

#### 4.3. Robustness Against Specific Bias

In this section, we aim to figure out whether our lbGen data can help the backbones to learn good features instead of capturing specific bias as a shortcut. Hence, we follow one recent study [44] to test shape-texture bias, context bias, and background bias of the backbone networks. All these results are given in Table 3.

Shape-Texture Bias. Prior work shows humans primarily use shape for object recognition [7, 45], while neural networks often rely on texture cues [10, 11]. Hence, we evaluate whether our data can reduce texture bias using the Cue Conflict, where shape and texture cues intentionally conflict across 1200 images from 16 classes. We use TIwhich represents the texture inclination of the model to understand the decision-making of the model when facing a shape-texture conflicting image (e.g. a cat with the texture of an elephant). Our findings indicate that training on our lbGen images, models tend to be less texture-biased. Concretely speaking, the two types of models trained on our data show 4.8% and 9.8% texture inclination decline compared with those trained on real images.

**Context Bias.** Context bias means that a model is biased towards using context cues to classify objects rather than learning real object appearance. In the Focus which we use to evaluate the context bias, each image is annotated with the object class, the time of day, location, and weather labels. These images are divided into common and uncommon sets. Uncommon samples are uncommon contexts like "airplane in the forest". Then we use mutually exclusive partitions of this dataset  $P_k$  where k is the number of uncommon attributes and report  $CB_{avg}$ . metrics, which is defined as the average relative accuracy between the accuracy on the

Backbone	Pre-Trained Data	$TI(\downarrow)$	$CB_{avg.}(\uparrow)$	$BG_{Gap}(\downarrow)$
ResNet50	IN-Real [6] IN-SD1.5 [40] IN-GenRobust [1] IN-RealFake [56] IN-IbGen(ours)	$ \begin{array}{r} 60.9\\ \underline{60.7}\\ 62.8\\ 69.2\\ \underline{56.1} \end{array} $	$ \begin{array}{r} 60.0 \\ 55.3 \\ 48.1 \\ \underline{60.1} \\ 64.7 \\ \end{array} $	6.8 8.0 7.5 8.2 6.4
ViT-S	IN-Real[6] IN-SD1.5 [40] IN-GenRobust [1] IN-RealFake [56] IN-lbGen(ours)	$ \begin{array}{r} 67.0\\ \underline{63.8}\\ 65.7\\ 70.6\\ 57.2 \end{array} $	$     \begin{array}{r}         \underline{61.8} \\         55.5 \\         47.3 \\         61.2 \\         \underline{66.0}     \end{array} $	$     \begin{array}{r}             \underline{6.7} \\             \overline{7.8} \\             \overline{7.9} \\             \overline{7.8} \\             \overline{6.1} \\             \hline         $

Table 3. Results on benchmarks of testing specific bias. TI (in %) denotes the texture inclination of the model.  $CB_{avg.}$  (in %) denotes the average relative accuracy when the number of uncommon attributes changes.  $BG_{Gap}$  (in %) metric reports the drop in performance by just changing the background to a different class than the foreground class.

partition with no uncommon attributes  $P_0$  and a partition with k uncommon attributes when k changes from 1 to 3:

$$CB_{avg.} = \frac{1}{3} \times \sum_{k=1}^{3} \frac{Acc_{P_k}}{Acc_{P_0}}.$$
 (9)

In our evaluation result, we find the models trained on our lbGen ImageNet demonstrate leading object recognition capabilities where achieve 64.7% and 66.0% average relative accuracy on each backbone in a constantly changing context. Background Bias. The background bias of models can be used to identify if the model is using the background of the image during training to improve the classification accuracy instead of using the object itself. For the two datasets that we utilize to evaluate the background bias, the Mixed-Rand segments the foreground object in an image and switches the original background with a random background from a different class label, while the Mixed-Same partition places the segmented foreground object on a random background from the same class label. Thus, we can use  $BG_{Gap}$  which measures the difference in performance between these two datasets to examine how decision-making processes can be influenced just by changing the background to a different class. As we report in Table 3, The ResNet50 and ViT-S trained on our data obtain 6.4% and 6.1% performance gaps, which shows lower gaps compared with training on other synthetic data and real data.

#### 4.4. Ablation Study

In this section, we investigate the design choices of lbGen training process. Due to computational cost, without loss of generality, we conducted the ablation study on the smaller ImageNet-100 datasets [46] for evaluation. Unless otherwise specified, we choose ResNet50 backbone and mainly report the results of transfer learning.

Pre-Trained Data	$\mid \mathcal{L}_{en}$	$\mathcal{L}_{in}$	$\mathcal{L}_q$	CLIP Size	IN100-val	Aircraft	Cars196	DTD	EuroSAT	Flowers	Pets	Food101	SUN397	Avg.
IN100-Real [46]	-	-	-	-	88.3	40.5	28.5	56.8	<u>92.4</u>	<u>68.8</u>	<u>72.1</u>	48.7	38.0	56.0
IN100-lbGen	✓	$\checkmark$	$\checkmark$	Large	66.0	44.8	34.3	59.1	92.7	71.6	72.0	50.8	42.3	58.5
	×	$\checkmark$	$\checkmark$	Large	64.6	39.3	27.0	54.8	91.5	65.7	70.3	45.4	35.3	53.6
	$\checkmark$	$\times$	$\checkmark$	Large	8.7	24.5	21.4	45.4	91.4	52.0	43.9	40.1	24.8	42.9
	<ul> <li>✓</li> </ul>	$\checkmark$	×	Large	51.3	41.7	25.1	56.2	92.1	64.4	68.3	46.2	36.4	53.7
IN100-SD1.5 [40]	×	$\times$	$\times$	-	65.5	39.7	26.9	53.2	91.1	64.2	69.3	45.6	35.9	53.2
	✓	$\checkmark$	$\checkmark$	Base	62.1	42.3	32.6	58.4	92.3	68.2	71.2	49.5	<u>39.2</u>	56.7

Table 4. Ablation studies on transfer learning datasets and IN100-val. Avg. is the average accuracy of eight transfer learning datasets. We also present the results of the model trained on real ImageNet100 for reference.



Figure 5. Impact of individual image alignmentloss. We observe that ambiguity problem between classes when discarding  $\mathcal{L}_{in}$ .

Effect of Bi-Level Semantic Alignment Loss. According to Table 4, it is notable that without the entire semantic alignment part, the average accuracy decreased by 4.9%. Furthermore, when we remove individual semantic alignment part, the accuracy on IN100-val and transfer learning data shows a significant decline. In our analysis, it appears that the absence of individual semantic alignment leads the model to solely learn the overall semantic distribution of the dataset, resulting in a lack of distinction or specific semantic meaning among different classes (see the left picture of each pair in Figure 5). This ultimately causes a collapse when training the backbone. In sum, the bi-level semantic alignment loss successfully help to align the generated images into low-biased semantic space from both the entire dataset distribution level and the specific object category level.

**Effect of Quality Assurance Loss.** As displayed in Figure 6, adding quality assurance loss sufficiently solves the quality deterioration when only using bi-level semantic alignment loss. Moreover, results in Table 4 indicate that the accuracy on both ImageNet100 and transfer learning tasks could decline in the absence of quality assurance loss. Thus, it can be drawn that the quality assurance loss successfully help to guarantee the low-level image quality during generation.

Effect of CLIP's Capacity. As depicted in Table 4, we aim to explore the effect of the knowledge of CLIP model as it plays a pivotal role in our method. Results indicate that the capacity of CLIP matters, the average accuracy decreased to 56.7% when changing to a smaller size which contains

Figure 6. Effectiveness of quality assurance loss. After adding  $\mathcal{L}_q$ , the image blur problem is solved.

lower image-text alignment capacity than a larger one.

# 5. Conclusion

In this study, we take the first attempt to directly generate a low-biased annotated dataset for more generalized backbone network pre-training. Specifically, we develop a novel bi-level semantic alignment loss, which not only forces all generated images to be consistent with the semantic distribution of all categories belonging to the target dataset, but also requires each generated image to match the semantic description of its category. Through fine-tuning the pre-trained diffusion model with the proposed loss together with a quality assurance loss which helps to guarantee the low-level image quality, we can obtain a low-biased annotated dataset generation model using only all category names in the target dataset as input. Experiments on various tasks demonstrate that pre-training backbone network on our generated dataset can lead to stable generalization capacity enhancement.

# 6. Acknowledgment

This work is supported in part by the National Natural Science Foundation of China under Grand 62372379, and Grant 62472359; in part by the Xi'an's Key Industrial Chain Core Technology Breakthrough Project: AI Core Technology Breakthrough under Grand 23ZDCYJSGG0003-2023.

# References

- Hritik Bansal and Aditya Grover. Leaving reality to imagination: Robust classification via generated datasets. *arXiv* preprint arXiv:2302.02503, 2023. 2, 3, 5, 6, 7
- [2] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101-mining discriminative components with random forests. In *Computer vision–ECCV 2014: 13th European conference, zurich, Switzerland, September 6-12, 2014, proceedings, part VI 13*, pages 446–461. Springer, 2014. 5
- [3] Chaofeng Chen, Annan Wang, Haoning Wu, Liang Liao, Wenxiu Sun, Qiong Yan, and Weisi Lin. Enhancing diffusion models with text-encoder reinforcement learning. arXiv preprint arXiv:2311.15657, 2023. 4
- [4] Hao Chen, Jindong Wang, Ankit Shah, Ran Tao, Hongxin Wei, Xing Xie, Masashi Sugiyama, and Bhiksha Raj. Understanding and mitigating the label noise in pre-training on downstream tasks. arXiv preprint arXiv:2309.17002, 2023. 2
- [5] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer* vision and pattern recognition, pages 3606–3613, 2014. 5
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009. 2, 3, 5, 6, 7
- [7] Gil Diesendruck and Paul Bloom. How specific is the shape bias? *Child development*, 74(1):168–178, 2003. 7
- [8] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929, 2020. 5
- [9] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. 3
- [10] Paul Gavrikov and Janis Keuper. Can biases in imagenet models explain generalization? In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 22184–22194, 2024. 2, 7
- [11] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. arXiv preprint arXiv:1811.12231, 2018. 5, 7
- [12] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020. 2
- [13] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 5
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceed-ings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5

- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In Proceedings of the IEEE international conference on computer vision, pages 2961–2969, 2017. 2, 7
- [16] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Introducing eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, pages 204–207. IEEE, 2018.
- [17] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8340–8349, 2021. 3
- [18] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. 5
- [19] Dongzhi Jiang, Guanglu Song, Xiaoshi Wu, Renrui Zhang, Dazhong Shen, Zhuofan Zong, Yu Liu, and Hongsheng Li. Comat: Aligning text-to-image diffusion model with imageto-text concept matching. *arXiv preprint arXiv:2404.03653*, 2024. 4, 5
- [20] Priyatham Kattakinda and Soheil Feizi. Focus: Familiar objects in common and uncommon settings. In *International Conference on Machine Learning*, pages 10825–10847. PMLR, 2022. 5
- [21] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops, 2013. 5
- [22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- [23] A Kumar. The illustrated image captioning using transformers. Ankur-NLP Enthusiast, 2022. 3
- [24] Shiye Lei, Hao Chen, Sen Zhang, Bo Zhao, and Dacheng Tao. Image captions are natural prompts for text-to-image models. arXiv preprint arXiv:2307.08526, 2023. 3
- [25] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip 2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023. 3
- [26] 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 Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13, pages 740–755. Springer, 2014. 5
- [27] Zhuang Liu and Kaiming He. A decade's battle on dataset bias: Are we there yet? *arXiv preprint arXiv:2403.08632*, 2024. 2, 3
- [28] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer:

Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 7

- [29] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 11976–11986, 2022. 7
- [30] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.
   5
- [31] Yujian Mo, Yan Wu, Xinneng Yang, Feilin Liu, and Yujun Liao. Review the state-of-the-art technologies of semantic segmentation based on deep learning. *Neurocomputing*, 493: 626–646, 2022. 2
- [32] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In 2008 Sixth Indian conference on computer vision, graphics & image processing, pages 722–729. IEEE, 2008. 5
- [33] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In 2012 IEEE conference on computer vision and pattern recognition, pages 3498–3505. IEEE, 2012. 5
- [34] Novi Patricia and Barbara Caputo. Learning to learn, from transfer learning to domain adaptation: A unifying perspective. In *Proceedings of the IEEE conference on computer* vision and pattern recognition, pages 1442–1449, 2014. 2, 3, 6
- [35] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. arXiv preprint arXiv:2307.01952, 2023. 3
- [36] Mihir Prabhudesai, Russell Mendonca, Zheyang Qin, Katerina Fragkiadaki, and Deepak Pathak. Video diffusion alignment via reward gradients. arXiv preprint arXiv:2407.08737, 2024. 4
- [37] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 3, 5
- [38] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125, 2022.
   2
- [39] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of* the IEEE/CVF conference on computer vision and pattern recognition, pages 10684–10695, 2022. 4
- [40] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of* the IEEE/CVF conference on computer vision and pattern recognition, pages 10684–10695, 2022. 2, 3, 4, 5, 6, 7, 8

- [41] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. Advances in neural information processing systems, 35:36479–36494, 2022. 2
- [42] Mert Bülent Sarıyıldız, Karteek Alahari, Diane Larlus, and Yannis Kalantidis. Fake it till you make it: Learning transferable representations from synthetic imagenet clones. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8011–8021, 2023. 2, 6
- [43] Dai Shi. Transnext: Robust foveal visual perception for vision transformers. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 17773– 17783, 2024. 7
- [44] Krishnakant Singh, Thanush Navaratnam, Jannik Holmer, Simone Schaub-Meyer, and Stefan Roth. Is synthetic data all we need? benchmarking the robustness of models trained with synthetic images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2505–2515, 2024. 2, 3, 7
- [45] Linda B Smith, Susan S Jones, Barbara Landau, Lisa Gershkoff-Stowe, and Larissa Samuelson. Object name learning provides on-the-job training for attention. *Psychological science*, 13(1):13–19, 2002. 7
- [46] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 776–794. Springer, 2020. 7, 8
- [47] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *CVPR 2011*, pages 1521–1528. IEEE, 2011.
   2, 3
- [48] Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. Generalizing from a few examples: A survey on few-shot learning. ACM computing surveys (csur), 53(3):1–34, 2020.
- [49] Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Yixuan Gao, Annan Wang, Erli Zhang, Wenxiu Sun, et al. Q-align: Teaching lmms for visual scoring via discrete text-defined levels. arXiv preprint arXiv:2312.17090, 2023. 5
- [50] Xiaoshi Wu, Yiming Hao, Manyuan Zhang, Keqiang Sun, Zhaoyang Huang, Guanglu Song, Yu Liu, and Hongsheng Li. Deep reward supervisions for tuning text-to-image diffusion models. arXiv preprint arXiv:2405.00760, 2024. 5
- [51] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In 2010 IEEE computer society conference on computer vision and pattern recognition, pages 3485–3492. IEEE, 2010. 5
- [52] Kai Xiao, Logan Engstrom, Andrew Ilyas, and Aleksander Madry. Noise or signal: The role of image backgrounds in object recognition. *ArXiv preprint arXiv:2006.09994*, 2020.
- [53] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In

Proceedings of the European conference on computer vision (ECCV), pages 418–434, 2018. 2, 7

- [54] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Finegrained 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. 3
- [55] Qinghao Ye, Haiyang Xu, Jiabo Ye, Ming Yan, Anwen Hu, Haowei Liu, Qi Qian, Ji Zhang, and Fei Huang. mplugowl2: Revolutionizing multi-modal large language model with modality collaboration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13040–13051, 2024. 5
- [56] Jianhao Yuan, Jie Zhang, Shuyang Sun, Philip Torr, and Bo Zhao. Real-fake: Effective training data synthesis through distribution matching. *arXiv preprint arXiv:2310.10402*, 2023. 2, 3, 5, 6, 7
- [57] Boya Zeng, Yida Yin, and Zhuang Liu. Understanding bias in large-scale visual datasets. Advances in Neural Information Processing Systems, 37:61839–61871, 2025. 2, 3
- [58] Zhong-Qiu Zhao, Peng Zheng, Shou-tao Xu, and Xindong Wu. Object detection with deep learning: A review. *IEEE transactions on neural networks and learning systems*, 30 (11):3212–3232, 2019. 2
- [59] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision*, 127:302–321, 2019. 5
- [60] Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. Dmgan: Dynamic memory generative adversarial networks for text-to-image synthesis. In *Proceedings of the IEEE/CVF* conference on computer vision and pattern recognition, pages 5802–5810, 2019. 3