

Multi-Concept Customization of Text-to-Image Diffusion

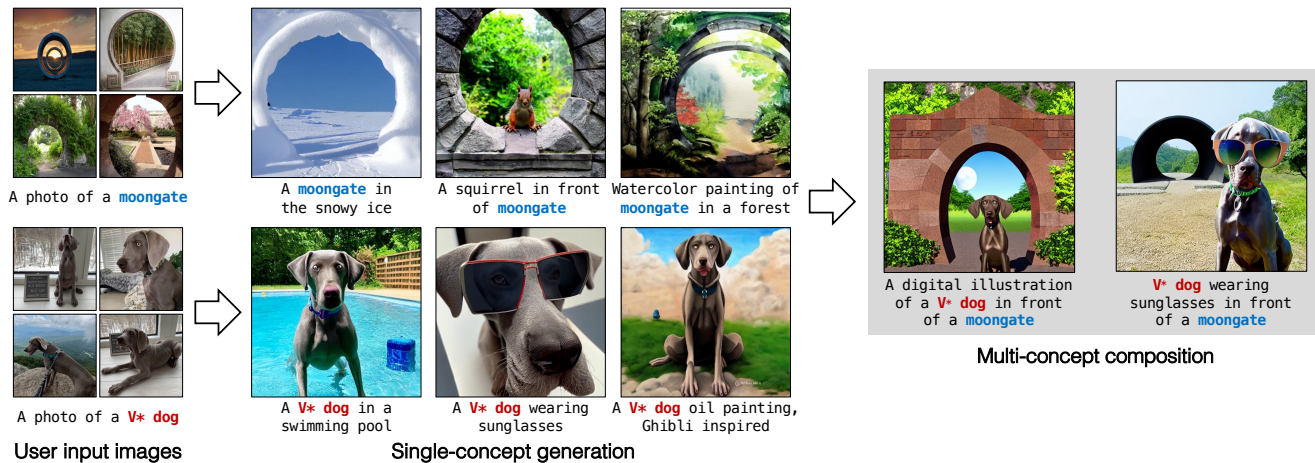
Nupur Kumari¹Bingliang Zhang²Richard Zhang³Eli Shechtman³Jun-Yan Zhu¹¹Carnegie Mellon University²Tsinghua University³Adobe Research

Figure 1. Given a few images of a new concept, our method augments a pre-trained text-to-image diffusion model, enabling new generations of the concept in unseen contexts. Example concepts include personal objects, animals, e.g., *a pet dog*, and classes not well generated by the model, e.g., *moongate* (a circular gate [76]). Furthermore, we propose a method for composing *multiple* new concepts together, for example, *V* dog wearing sunglasses in front of a moongate*. We denote personal categories with a new modifier token V^* .

Abstract

While generative models produce high-quality images of concepts learned from a large-scale database, a user often wishes to synthesize instantiations of their own concepts (for example, their family, pets, or items). Can we teach a model to quickly acquire a new concept, given a few examples? Furthermore, can we compose multiple new concepts together? We propose *Custom Diffusion*, an efficient method for augmenting existing text-to-image models. We find that only optimizing a few parameters in the text-to-image conditioning mechanism is sufficiently powerful to represent new concepts while enabling fast tuning (~ 6 minutes). Additionally, we can jointly train for multiple concepts or combine multiple fine-tuned models into one via closed-form constrained optimization. Our fine-tuned model generates variations of multiple new concepts and seamlessly composes them with existing concepts in novel settings. Our method outperforms or performs on par with several baselines and concurrent works in both qualitative and quantitative evaluations, while being memory and computationally efficient.

1. Introduction

Recently released text-to-image models [53, 57, 60, 79] have represented a watershed year in image generation. By simply querying a text prompt, users are able to generate images of unprecedented quality. Such systems can generate a wide variety of objects, styles, and scenes – seemingly “anything and everything”.

However, despite the diverse, *general* capability of such models, users often wish to synthesize *specific* concepts from their own personal lives. For example, loved ones such as family, friends, pets, or personal objects and places, such as a new sofa or a recently visited garden, make for intriguing concepts. As these concepts are by nature personal, they are unseen during large-scale model training. Describing these concepts after the fact, through text, is unwieldy and unable to produce the personal concept with sufficient fidelity.

This motivates a need for model *customization*. Given the few user-provided images, can we augment existing text-to-image diffusion models with the new concept (for example, their pet dog or a “moongate” as shown in Figure 1)? The fine-tuned model should be able to generalize and compose them with existing concepts to generate new variations. This

poses a few challenges – first, the model tends to forget [12, 35, 52] or change [34, 41] the meanings of existing concepts: e.g., the meaning of “moon” being lost when adding the “moongate” concept. Secondly, the model is prone to overfit the few training samples and reduce sampling variations.

Moreover, we study a more challenging problem, *compositional fine-tuning* – the ability to extend beyond tuning for a single, individual concept and compose multiple concepts together, e.g., pet dog in front of moongate (Figure 1). Improving compositional generation has been studied in recent works [40]. But composing multiple new concepts poses additional challenges, such as mixing unseen concepts.

In this work, we propose a fine-tuning technique, *Custom Diffusion* for text-to-image diffusion models. Our method is computationally and memory efficient. To overcome the above-mentioned challenges, we identify a small subset of model weights, namely the key and value mapping from text to latent features in the cross-attention layers [5, 70]. Fine-tuning these is sufficient to update the model with the new concept. To prevent model forgetting, we use a small set of real images with similar captions as the target images. We also introduce augmentation during fine-tuning, which leads to faster convergence and improved results. To inject multiple concepts, our method supports training on both simultaneously or training them separately and then merging.

We build our method on Stable Diffusion [1] and experiment on various datasets with as few as four training images. For adding single concepts, our method shows better text alignment and visual similarity to the target images than concurrent works and baselines. More importantly, our method can compose multiple new concepts efficiently, whereas concurrent methods struggle and often omit one. Finally, our method only requires storing a small subset of parameters (3% of the model weights) and reduces the fine-tuning time (6 minutes on 2 A100 GPUs, 2 – 4× faster compared to concurrent works). Full version of the paper is available at <https://arxiv.org/abs/2212.04488>.

2. Related Work

Deep generative models Generative models aim to synthesize samples from a data distribution, given a set of training examples. These include GANs [9, 20, 29], VAEs [31], autoregressive [15, 55, 69], flow-based [13, 14], and diffusion models [11, 25, 65]. To improve controllability, these models can be conditioned on a class [9, 62], image [27, 43, 73, 86], or text prompt [45, 67, 87]. Our work mainly relates to text-conditioned synthesis [42]. While earlier works [26, 56, 67, 78, 80, 87] were limited to a few classes [38, 71], recent text-to-image models [11, 45, 53, 54, 57, 60, 79], trained on extremely large-scale data, have demonstrated remarkable generalization ability. However, such models are by nature generalists and struggle to generate specific instances like personal toys or rare categories, e.g., “moongate”. We aim

to adapt these models to become specialists in new concepts.

Image and model editing. While generative models can sample random images, a user often wishes to edit a single, specific image. Several works aim at leveraging the capabilities of generative models, such as GANs [2–4, 49, 85] or diffusion models [10, 22, 30, 45] towards editing. A challenge is representing the specific image in the pre-trained model, and such methods employ per-image or per-edit optimization. A closely related line of work edits a generative model directly [6, 18, 72]. Whereas these methods aim to customize GANs, our focus is on text-to-image models.

Transfer learning. A method of efficiently producing a whole distribution of images is leveraging a pretrained model and then using transfer learning [18, 21, 36, 39, 44, 46–48, 74, 75, 82]. For example, one can adapt photorealistic faces into cartoons [18, 36, 44, 47, 48]. To adapt with just a few training images, efficient training techniques [28, 33, 61, 68, 83, 84] are often useful. Different from these works, which focus on tuning whole models to single domains, we wish to acquire multiple new concepts without catastrophic forgetting [16, 32, 35, 37, 52]. By preserving the millions of concepts captured in large-scale pretraining, we can synthesize the new concepts in composition with these existing concepts. Relatedly, several methods [19, 50, 64] propose to train adapter modules for large-scale models in the discriminative setting. In contrast, we adapt a small number of existing parameters and do not require additional parameters.

Adapting text-to-image models. Similar to our goals, two concurrent works, DreamBooth [59] and Textual Inversion [17], adopt transfer learning to text-to-image diffusion models [57, 60] via either fine-tuning all the parameters [59] or introducing and optimizing a word vector [17] for the new concept. Our work differs in several aspects. First, our work tackles a challenging setting: compositional fine-tuning of *multiple* concepts, where concurrent works struggle. Second, we only fine-tune a subset of cross-attention layer parameters, which significantly reduces the fine-tuning time. We find these design choices lead to better results, validated by automatic metrics and human preference studies.

3. Method

Our proposed method for model fine-tuning, as shown in Figure 2, only updates a small subset of weights in the cross-attention layers of the model. In addition, we use a regularization set of real images to prevent overfitting on the few training samples of the target concepts. In this section, we explain our design choices and final algorithm in detail.

3.1. Single-Concept Fine-tuning

Given a pretrained text-to-image diffusion model, we aim to embed a new concept in the model given as few as four images and the corresponding text description. The fine-

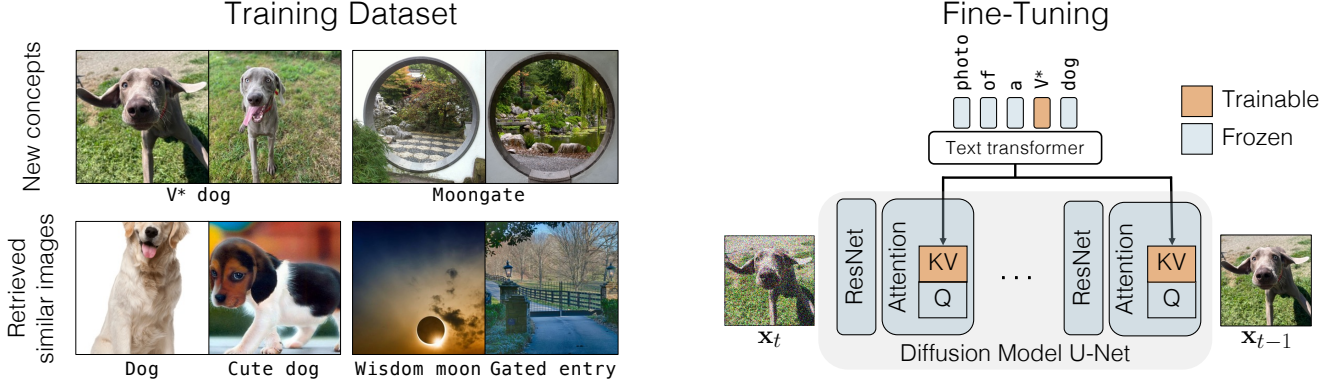


Figure 2. **Custom Diffusion.** Given images of new concepts, we retrieve real images with similar captions as the given concepts and create the training dataset for fine-tuning, as shown on the left. To represent personal concepts of a general category, we introduce a new modifier token V^* , used in front of the category name. During training, we optimize key and value projection matrices in the diffusion model cross-attention layers along with the modifier token. The retrieved real images are used as a regularization dataset during fine-tuning.

tuned model should retain its prior knowledge, allowing for novel generations with the new concept, based on the text prompt. This can be challenging as the updated text-to-image mapping might easily overfit the few available images.

In our experiments, we use Stable Diffusion [1] as our backbone model, which is built on the Latent Diffusion Model (LDM) [57]. LDM first encodes images into a latent representation, using hybrid objectives of VAE [31], PatchGAN [27], and LPIPS [81], such that running an encoder-decoder can recover an input image. They then train a diffusion model [25] on the latent representation with text condition injected in the model using cross-attention.

Learning objective of diffusion models. Diffusion models [25, 65] are a class of generative models that aim to approximate the original data distribution $q(\mathbf{x}_0)$ with $p_\theta(\mathbf{x}_0)$:

$$p_\theta(\mathbf{x}_0) = \int \left[p_\theta(\mathbf{x}_T) \prod p_\theta^t(\mathbf{x}_{t-1} | \mathbf{x}_t) \right] d\mathbf{x}_{1:T}, \quad (1)$$

where \mathbf{x}_1 to \mathbf{x}_T are latent variables of a forward Markov chain s.t. $\mathbf{x}_t = \sqrt{\alpha_t}\mathbf{x}_0 + \sqrt{1 - \alpha_t}\epsilon$. The model is trained to learn the reverse process of a fixed-length (usually 1000) Markov chain. Given noisy image \mathbf{x}_t at timestep t , the model learns to denoise the input image to obtain \mathbf{x}_{t-1} . The training objective of the diffusion model can be simplified to:

$$\mathbb{E}_{\epsilon, \mathbf{x}, \mathbf{c}, t} [w_t \|\epsilon - \epsilon_\theta(\mathbf{x}_t, \mathbf{c}, t)\|], \quad (2)$$

where ϵ_θ is the model prediction and w_t is a time-dependent weight on the loss. The model is conditioned on timestep t and can be further conditioned on any other modality \mathbf{c} , e.g., text. During inference, a random Gaussian image (or latent) \mathbf{x}_T is denoised for fixed timesteps using the model [66].

A naïve baseline for the goal of fine-tuning is to update all layers to minimize the loss in Eqn. 2 for the given text-image pairs. This can be computationally inefficient for large-scale models and can easily lead to overfitting when training on a few images. Therefore, we aim to identify a minimal set of weights that is sufficient for the task of fine-tuning.




Figure 3. **Analysis of change in weights** on updating all network weights during fine-tuning. The mean change in the cross-attention layers is significantly higher than other layers, even though they only make up 5% of the total parameter count.

Rate of change of weights. Following Li *et al.* [36], we analyze the change in parameters for each layer in the fine-tuned model on the target dataset with the loss in Eqn. 2, i.e., $\Delta_l = \|\theta'_l - \theta_l\| / \|\theta_l\|$, where θ'_l and θ_l are the updated and pretrained model parameters of layer l . These parameters come from three types of layers – (1) cross-attention (between the text and image), (2) self-attention (within the image itself), and (3) the rest of the parameters, including convolutional blocks and normalization layers in the diffusion model U-Net. Figure 3 shows the mean Δ_l for the three categories when the model is fine-tuned on “moongate” images. We observe similar plots for other datasets. As we see, the cross-attention layer parameters have relatively higher Δ compared to the rest of the parameters. Moreover, cross-attention layers are only 5% of the total parameter count in the model. This suggests it plays a significant role during fine-tuning, and we leverage that in our method.

Model fine-tuning. Cross-attention block modifies the latent features of the network according to the condition features, i.e., text features in the case of text-to-image diffusion models. Given text features $\mathbf{c} \in \mathbb{R}^{s \times d}$ and latent image features $\mathbf{f} \in \mathbb{R}^{(h \times w) \times l}$, a single-head cross-attention [70] operation consists of $Q = W^q \mathbf{f}$, $K = W^k \mathbf{c}$, $V = W^v \mathbf{c}$, and a weighted sum over value features as:

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d'}}\right)V, \quad (3)$$

where W^q , W^k , and W^v map the inputs to a query, key, and

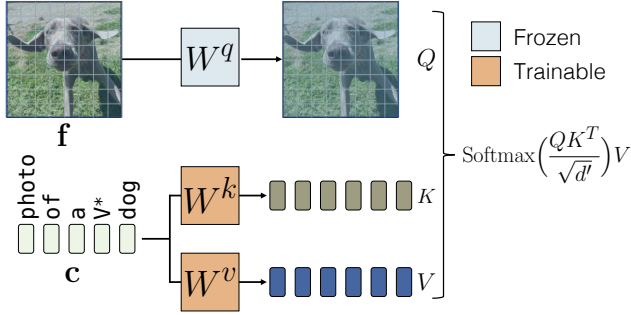


Figure 4. **Single-head Cross-Attention.** Latent image feature \mathbf{f} and text feature \mathbf{c} are projected into query Q , key K , and value V . Output is a weighted sum of values, weighted by the similarity between the query and key features. We highlight the updated parameters W^k and W^v in our method.

value feature, respectively, and d' is the output dimension of key and query features. The latent feature is then updated with the attention block output. The task of fine-tuning aims at updating the mapping from given text to image distribution, and the text features are only input to W^k and W^v projection matrix in the cross-attention block. Therefore, we propose to only update W^k and W^v parameters of the diffusion model during the fine-tuning process. As shown in our experiments, this is sufficient to update the model with a new text-image paired concept. Figure 4 shows an instance of the cross-attention layer and the trainable parameters.

Text encoding. Given target concept images, we require a text caption as well. If there exists a text description, e.g., *moongate*, we use that as a text caption. For personalization-related use-case where the target concept is a unique instance of a general category, e.g., *pet dog*, we introduce a new modifier token embedding, i.e., v^* *dog*. During training, v^* is initialized with a rare occurring token embedding and optimized along with cross-attention parameters. An example text caption used during training is, *photo of a v^* dog*.

Regularization dataset. Fine-tuning on the target concept and text caption pair can lead to the issue of language drift [34, 41]. For example, training on “moongate” will lead to the model forgetting the association of “moon” and “gate” with their previously trained visual concepts, as shown in Figure 5. Similarly, training on a personalized concept of v^* *tortoise plushy can leak*, causing all examples with *plushy* to produce the specific target images. To prevent this, we select a set of 200 regularization images from the LAION-400M [63] dataset with corresponding captions that have a high similarity with the target text prompt, above threshold 0.85 in CLIP [51] text encoder feature space.

3.2. Multiple-Concept Compositional Fine-tuning

Joint training on multiple concepts. For fine-tuning with multiple concepts, we combine the training datasets for each individual concept and train them jointly with our method.



Figure 5. **Role of regularization data in mitigating overfitting behavior during fine-tuning.** 1st row: samples from pre-trained models. In 2nd row, fine-tuning cross-attention key, value projection matrices without any regularization dataset leads to moongate-like images on the text prompt *photo of a moon*. We largely mitigate this issue with the use of regularization datasets as shown in 3rd row. More results can be found in our arxiv version.

To denote the target concepts, we use different modifier tokens, v_i^* , initialized with different rarely-occurring tokens and optimize them along with cross-attention key and value matrices for each layer. As shown in Figure 8, restricting the weight update to cross-attention key and value parameters leads to significantly better results for composing two concepts compared to methods like DreamBooth, which fine-tune all the weights.

Constrained optimization to merge concepts. As our method only updates the key and value projection matrices corresponding to the text features, we can subsequently merge them to allow generation with multiple fine-tuned concepts. Let set $\{W_{0,l}^k, W_{0,l}^v\}_{l=1}^L$ represent the key and value matrices for all L cross-attention layers in the pretrained model and $\{W_{n,l}^k, W_{n,l}^v\}_{l=1}^L$ represent the corresponding updated matrices for added concept $n \in \{1 \dots N\}$. As our subsequent optimization applies to all layers and key-value matrices, we will omit superscripts $\{k, v\}$ and layer l for notational clarity. We formulate the composition objective as the following constrained least squares problem:

$$\begin{aligned} \hat{W} &= \arg \min_W \|WC_{\text{reg}}^\top - W_0C_{\text{reg}}^\top\|_F \\ \text{s.t. } WC^\top &= V, \text{ where } C = [\mathbf{c}_1 \dots \mathbf{c}_N]^\top \\ &\text{and } V = [W_1\mathbf{c}_1^\top \dots W_N\mathbf{c}_N^\top]^\top. \end{aligned} \quad (4)$$

Here, $C \in \mathbb{R}^{s \times d}$ is the text features of dimension d . These are compiled of s target words across all N concepts, with all captions for each concept flattened out and concatenated. Similarly, $C_{\text{reg}} \in \mathbb{R}^{s_{\text{reg}} \times d}$ consists of text features of ~ 1000 randomly sampled captions for regularization. Intuitively, the above formulation aims to update the matrices in the original model, such that the words in target captions in C

are mapped consistently to the values obtained from fine-tuned concept matrices. The above objective can be solved in closed form, assuming C_{reg} is non-degenerate and the solution exists, by using the method of Lagrange multipliers [8]:

$$\hat{W} = W_0 + \mathbf{v}^\top \mathbf{d}, \text{ where } \mathbf{d} = C(C_{\text{reg}}^\top C_{\text{reg}})^{-1} \quad (5)$$

$$\text{and } \mathbf{v}^\top = (V - W_0 C^\top)(\mathbf{d} C^\top)^{-1}.$$

We show full derivation in the our arxiv version. Compared to joint training, our optimization-based method is faster (~ 2 seconds) if each individual fine-tuned model exists. Our proposed methods lead to the coherent generation of two new concepts in a single scene, as shown in Section 4.2.

Training details. We train the models with our method for 250 steps in single-concept and 500 steps in two-concept joint training, on a batch size of 8 and learning rate 8×10^{-5} . During training, we also randomly resize only the target images from $0.4-1.4\times$ and append the prompt “very small”, “far away” or “zoomed in”, “close up” accordingly to the text prompt based on resize ratio. We only backpropagate the loss on valid regions. This leads to faster convergence and improved results. We provide more training and architecture details in our arxiv version.

4. Experiments

In this section, we show the results of our method on multiple datasets in both single concept fine-tuning and composition of two concepts on the Stable Diffusion model [1].

Datasets. We perform experiments on ten target datasets spanning a variety of categories and varying training samples. It consists of two scene categories, two pets, and six objects, as shown in Figure 7. We show that the pretrained model cannot generate these concepts even with long text descriptions in our arxiv version.

Evaluation metrics. We evaluate our method on (1) *Image-alignment*, i.e., the visual similarity of generated images with the target concept, using similarity in CLIP image feature space [17], (2) *Text-alignment* of the generated images with given prompts, using text-image similarity in CLIP feature space [23], and (3) *KID* [7] on a validation set of 500 real images of a similar concept retrieved from LAION-400M to measure overfitting on the target concept (e.g., V^* dog) and forgetting of existing related concepts (e.g., dog). (4) We also perform a *human preference* study of our method with baselines. Unless mentioned otherwise, we use 200 steps of DDPM sampler with a scale 6. The prompts used for quantitative and human evaluation are shown on our [website](#).

Baselines. We compare our method with the two concurrent works, *DreamBooth* [59] (third-party implementation [77]) and *Textual Inversion* [17]. DreamBooth fine-tunes all the parameters in the diffusion model, keeping the text transformer frozen, and uses generated images as the regulariza-

tion dataset. Each target concept is represented by a unique identifier, followed by its category name, i.e., “[V] *category*”, where [V] is a rarely occurring token in the text token space and not optimized during fine-tuning. Textual Inversion optimizes a new v^* token for each new concept. We also compare with the competitive baseline of *Custom Diffusion (w/ fine-tune all)*, where we fine-tune all the parameters in the U-Net [58] diffusion model, along with the v^* token embedding in our method. We provide implementation details for all baselines in our arxiv version.

4.1. Single-Concept Fine-tuning Results

Qualitative evaluation. We test each fine-tuned model on a set of challenging prompts. This includes generating the target concept in a new scene, in a known art style, composing it with another known object, and changing certain properties of the target concept: e.g., color, shape, or expression. Figure 6 shows the sample generations with our method, DreamBooth, and Textual Inversion. Our method, Custom Diffusion, has higher text-image alignment while capturing the visual details of the target object. It performs better than Textual Inversion and is on par with DreamBooth while having a lower training time and model storage ($\sim 5\times$ faster and 75MB vs 3GB storage).

Quantitative evaluation. We evaluate on 20 text prompts and 50 samples per prompt for each dataset, resulting in a total of 1000 generated images. We use DDPM sampling with 50 steps and a classifier-free guidance scale of 6 across all methods. As shown in Figure 7, our method outperforms the concurrent methods [17, 59]. Also, Custom Diffusion works on par with our proposed baseline of fine-tuning all the weights in the diffusion model, while being more computationally and time efficient. Table 1 also shows the KID of generated images by each fine-tuned model on a reference dataset, with captions similar to the fine-tuned concept. As we observe, our method has lower KID than most baselines, which suggests less overfitting to the target concept. We show in our arxiv version that the updated matrices can be compressed to further reduce model storage.

Computational requirements Training time of our method is ~ 6 minutes (2 A100 GPUs), compared to 20 minutes for Ours (w/ fine-tune all) (4 A100s), 20 minutes for Textual Inversion (2 A100s), and ~ 1 hour for DreamBooth (4 A100s). Also, since we update only 75MB of weights, our method has low memory requirements for storing each concept model. We keep the batch size fixed at 8 across all.

4.2. Multiple-Concept Fine-tuning Results

We show the results of generating two new concepts in the same scene on the following five pairs: (1) Moongate + Dog, (2) Cat + Chair, (3) Wooden Pot + Cat, (4) Wooden Pot + flower, and (5) Table + Chair. We compare our method with DreamBooth training on the two datasets together, using

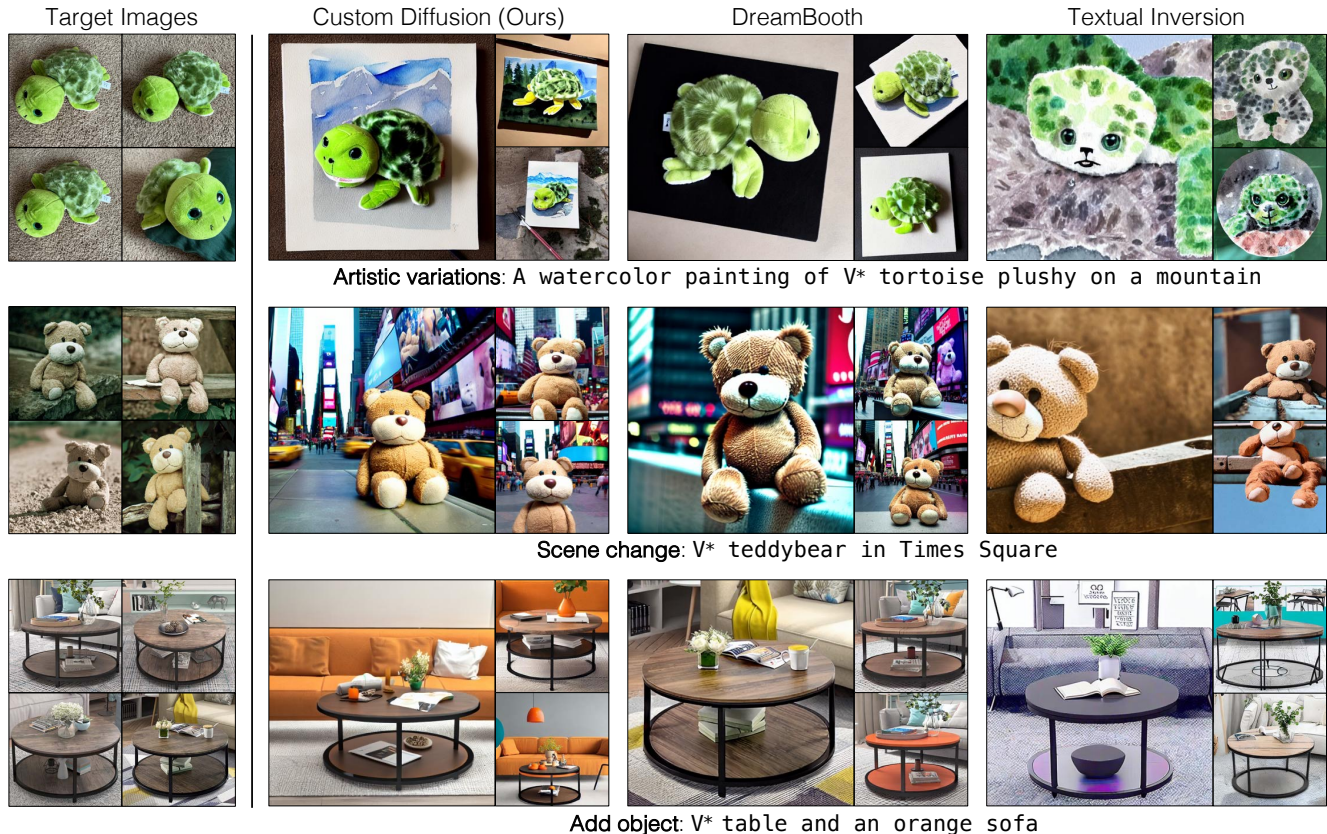


Figure 6. **Single-concept fine-tuning results.** Given images of a new concept (target images shown on the left), our method can generate images with the concept in unseen contexts and art styles. *First row*: representing the concept in an artistic style of watercolor paintings. Our method can also generate the mountains in the background, which DreamBooth and Textual Inversion omit. *Second row*: changing the background scene. Our method and DreamBooth perform similarly and better than Textual Inversion. *Third row*: adding another object, e.g., an orange sofa with the target table. Our method successfully adds the other object. We show more samples on our [website](#).

two different $[V_1]$ and $[V_2]$ tokens for each concept. For Textual Inversion, we perform inference using the individually fine-tuned tokens for each concept in the same sentence. We compare our method with the baselines on a set of 400 images generated with 8 prompts for each composition setting in Figure 7 and Table 1. We also compare with our baseline of sequential training on the two concepts or fine-tuning all weights in our method. Our method performs better on all except the “Table+Chair” composition, where all methods perform comparably except Textual Inversion, which doesn’t perform well at composition, as also shown qualitatively in the appendix of our arxiv version. This shows the importance of fine-tuning only the cross-attention parameters for composition. In the case of sequential training, we observe forgetting of the first concept. Figure 8 shows sample images of our proposed two methods and DreamBooth. As we can see, our method is able to generate the two objects in the same scene in a coherent manner while having high alignment with the input text prompt. We show more samples on our [website](#).

	Method	Text-alignment	Image-alignment	KID (validation)
Single-Concept	Textual Inversion	0.670	0.827	22.27
	DreamBooth	0.781	0.776	32.53
	Ours (w/ fine-tune all)	0.795	0.748	19.27
	Ours	0.795	0.775	20.96
Multi-Concept	Textual Inversion	0.544	0.630	—
	DreamBooth	0.783	0.695	—
	Ours (w/ fine-tune all)	0.787	0.691	—
	Ours (Sequential)	0.797	0.700	—
	Ours (Optimization)	0.800	0.695	—
	Ours (Joint)	0.801	0.706	—

Table 1. **Quantitative comparisons.** *Top row*: single-concept fine-tuning evaluation averaged across datasets. The last column shows the KID ($\times 10^3$) between real validation set images and generated images with the same caption. Since our method uses a regularization set of real images, it achieves lower KID and even improves slightly over the pretrained model. Textual Inversion has the same KID as the pretrained model, as it does not update the model. *Bottom row*: evaluation on multi-concept averaged across the five composition pairs. We show individual scores for all in our arxiv version. We also evaluate single-concept fine-tuned models on FID [24] (MS-COCO [38]) and show the trend of image-, text-alignment with training steps in our arxiv version.

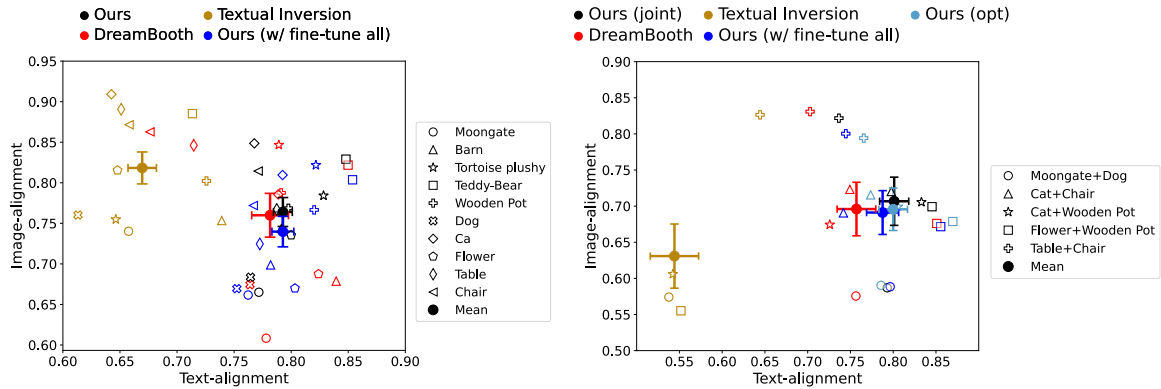


Figure 7. **Text- and image-alignment for single-concept (left) and multi-concept (right) fine-tuning.** Compared to others, our method lies further along the upper right corner (with less variance). Keeping in consideration the trade-off between image-alignment vs. text-alignment on new prompts, our method is on-par or better than the baselines.




Figure 8. **Multi-concept fine-tuning results.** *First row:* our method has higher visual similarity with the personal cat and chair while following the text condition. *Second row:* DreamBooth omits the cat sometimes, whereas our method generates both the cat and the wooden pot. *Third row:* our method better maintains the visual similarity to the target images. *Fourth row:* target table and chair together in a garden. For all settings, our methods perform better than DreamBooth, and joint training performs better than the optimization-based method.

Ours	Textual Inversion		DreamBooth		Ours (w/ fine-tune all)	
	Text Alignment	Image Alignment	Text Alignment	Image Alignment	Text Alignment	Image Alignment
Single-concept	72.62 ± 2.38%	51.62 ± 2.62%	53.50 ± 2.64%	56.62 ± 2.44%	55.17 ± 2.55%	53.99 ± 2.44%
Multi-concept (Joint)	86.65 ± 2.25%	81.89 ± 2.09%	56.39 ± 2.46%	61.80 ± 2.59%	59.00 ± 2.61%	59.12 ± 2.72%
Multi-concept (Optimization)	81.22 ± 2.72%	83.11 ± 2.18%	57.00 ± 2.62%	61.75 ± 2.68%	57.60 ± 2.43%	53.49 ± 2.71%

Table 2. **Human preference study.** For each paired comparison, our method is preferred (over $\geq 50\%$) over the baseline in both image- and text-alignment. Textual Inversion seems to overfit to target images and thus has a similar image-alignment as ours but performs worse on text-alignment in the single-concept setting.

Method	Text-alignment	Image-alignment	KID (validation)
Ours	0.795	0.775	20.96
Ours w/o Aug	0.800	0.736	20.67
Ours w/o Reg	0.799	0.756	32.64
Ours w/ Gen	0.791	0.768	34.70

Table 3. **Ablation Study.** No augmentation during training leads to lower image-alignment. No regularization dataset or using generated images as regularization produces much worse KID.

4.3. Human Preference Study

We perform the human preference study using Amazon Mechanical Turk. We do a paired test of our method with DreamBooth, Textual Inversion, and Ours (w/ fine-tune all). As shown in Table 2, our method is preferred over baselines in both single-concept and multi-concept, even compared to Ours (w/ fine-tune all) method, which shows the importance of only updating cross-attention parameters. More details of the setup are provided in our arxiv version.

4.4. Ablation and Applications

In this section, we ablate various components of our method to show its contribution. We evaluate each experiment on the same setup as in Section 4.1. Sample generations for ablation experiments are shown on our [website](#). We also show more analysis and applications in our arxiv version.

Generated images as regularization (Ours w/ Gen). As detailed in Section 3.1, we retrieve similar category real images and captions to use as regularization during fine-tuning. Another way of creating the regularization dataset is to generate images from the pretrained model [59]. We compare our method with this setup, i.e., for the target concept of a “category” generate images using the prompt, `photo of a {category}`, and show results in Table 3. Using generated images results in a similar performance on the target concept. However, this shows signs of overfitting, as measured by KID on a validation set of similar category real images.

Without regularization dataset (Ours w/o Reg). We show results when no regularization dataset is used. We train the model for half the number of iterations (the same number of target images seen during training). Table 3 shows that the model has slightly lower image-alignment and tends to forget existing concepts, as evident from high KID on the



Figure 9. **Failure cases on multi-concept fine-tuning.** Our method fails at difficult compositions like a cat and dog together in a scene or similar category objects like teddybear and tortoise plushy as shown on the right. Though as shown on the left, the pretrained model also struggles with similar compositions.

validation set.

Without data augmentation (Ours w/o Aug). As mentioned in Section 3.2, we augment by randomly resizing the target images during training and append size-related prompts (e.g., “very small”) to the text. Here, we show the effect of not using these augmentations. The model is trained for the same number of steps. Table 3 shows that no augmentation leads to lower visual similarity with target images.

5. Discussion and Limitations

In conclusion, we have proposed a method for fine-tuning large-scale text-to-image diffusion models on new concepts, categories, personal objects, or artistic styles, using just a few image examples. Our computationally efficient method can generate novel variations of the fine-tuned concept in new contexts while preserving the visual similarity with the target images. Moreover, we only need to save a small subset of model weights. Furthermore, our method can coherently compose multiple new concepts in the same scene.

As shown in Figure 9, difficult compositions, e.g., a pet dog and a pet cat, remain challenging. In this case, the pretrained model also faces a similar difficulty, and our model inherits these limitations. Additionally, composing increasing three or more concepts together is also challenging. We show more analysis and visualization in our arxiv version.

Acknowledgment. We are grateful to Nick Kolkin, David Bau, Sheng-Yu Wang, Gaurav Parmar, John Nack, and Sylvain Paris for their helpful comments and discussion and to Allie Chang, Chen Wu, Sumith Kulal, Minguk Kang, and Taesung Park for proofreading the draft. We also thank Mia Tang and Aaron Hertzmann for sharing their artwork. This work was partly done by Nupur Kumari during the Adobe internship. The work is partly supported by Adobe Inc.

References

- [1] Stable diffusion. <https://huggingface.co/CompVis/stable-diffusion-v-1-4-original>, 2022. 2, 3, 5
- [2] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [3] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan++: How to edit the embedded images? In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [4] Rameen Abdal, Peihao Zhu, John Femiani, Niloy Mitra, and Peter Wonka. Clip2stylegan: Unsupervised extraction of stylegan edit directions. In *ACM SIGGRAPH*, 2022. 2
- [5] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*, 2015. 2
- [6] David Bau, Steven Liu, Tongzhou Wang, Jun-Yan Zhu, and Antonio Torralba. Rewriting a deep generative model. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [7] Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. In *International Conference on Learning Representations (ICLR)*, 2018. 5
- [8] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004. 5
- [9] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *International Conference on Learning Representations (ICLR)*, 2019. 2
- [10] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [12] Yuxuan Ding, Lingqiao Liu, Chunna Tian, Jingyuan Yang, and Haoxuan Ding. Don't stop learning: Towards continual learning for the clip model. *arXiv preprint arXiv:2207.09248*, 2022. 2
- [13] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. In *ICLR Workshop*, 2015. 2
- [14] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. In *International Conference on Learning Representations (ICLR)*, 2017. 2
- [15] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [16] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999. 2
- [17] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022. 2, 5
- [18] Rinon Gal, Or Patashnik, Haggai Maron, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022. 2
- [19] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*, 2021. 2
- [20] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 2
- [21] Zheng Gu, Wenbin Li, Jing Huo, Lei Wang, and Yang Gao. Lofgan: Fusing local representations for few-shot image generation. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [22] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 2
- [23] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. In *EMNLP*, 2021. 5
- [24] 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 *Conference on Neural Information Processing Systems (NeurIPS)*, 2017. 6
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020. 2, 3
- [26] Xun Huang, Arun Mallya, Ting-Chun Wang, and Ming-Yu Liu. Multimodal conditional image synthesis with product-of-experts gans. In *European Conference on Computer Vision*, pages 91–109. Springer, 2022. 2
- [27] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2, 3
- [28] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [29] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [30] Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [31] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations (ICLR)*, 2014. 2, 3

- [32] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. 2
- [33] Nupur Kumari, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Ensembling off-the-shelf models for gan training. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [34] Jason Lee, Kyunghyun Cho, and Douwe Kiela. Countering language drift via visual grounding. In *EMNLP*, 2019. 2, 4
- [35] Dingcheng Li, Zheng Chen, Eunah Cho, Jie Hao, Xiaohu Liu, Fan Xing, Chenlei Guo, and Yang Liu. Overcoming catastrophic forgetting during domain adaptation of seq2seq language generation. In *NAACL*, 2022. 2
- [36] Yijun Li, Richard Zhang, Jingwan Lu, and Eli Shechtman. Few-shot image generation with elastic weight consolidation. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020. 2, 3
- [37] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2017. 2
- [38] 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 (ECCV)*, 2014. 2, 6
- [39] Bingchen Liu, Yizhe Zhu, Kunpeng Song, and Ahmed Elgammal. Towards faster and stabilized gan training for high-fidelity few-shot image synthesis. In *International Conference on Learning Representations (ICLR)*, 2021. 2
- [40] Nan Liu, Shuang Li, Yilun Du, Antonio Torralba, and Joshua B Tenenbaum. Compositional visual generation with composable diffusion models. In *ECCV*, 2022. 2
- [41] Yuchen Lu, Soumye Singhal, Florian Strub, Aaron Courville, and Olivier Pietquin. Countering language drift with seeded iterated learning. In *International Conference on Machine Learning (ICML)*, 2020. 2, 4
- [42] Elman Mansimov, Emilio Parisotto, Jimmy Lei Ba, and Ruslan Salakhutdinov. Generating images from captions with attention. *arXiv preprint arXiv:1511.02793*, 2015. 2
- [43] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2022. 2
- [44] Sangwoo Mo, Minsu Cho, and Jinwoo Shin. Freeze the discriminator: a simple baseline for fine-tuning gans. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop*, 2020. 2
- [45] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *International Conference on Machine Learning (ICML)*, 2022. 2
- [46] Yotam Nitzan, Kfir Aberman, Qiurui He, Orly Liba, Michal Yarom, Yossi Gandelsman, Inbar Mosseri, Yael Pritch, and Daniel Cohen-Or. Mystyle: A personalized generative prior. In *SIGGRAPH ASIA*, 2022. 2
- [47] Atsuhiko Noguchi and Tatsuya Harada. Image generation from small datasets via batch statistics adaptation. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [48] Utkarsh Ojha, Yijun Li, Jingwan Lu, Alexei A Efros, Yong Jae Lee, Eli Shechtman, and Richard Zhang. Few-shot image generation via cross-domain correspondence. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [49] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [50] Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. Mad-x: An adapter-based framework for multi-task cross-lingual transfer. In *EMNLP*, 2020. 2
- [51] 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 (ICML)*, 2021. 4
- [52] Vinay Venkatesh Ramasesh, Aitor Lewkowycz, and Ethan Dyer. Effect of scale on catastrophic forgetting in neural networks. In *International Conference on Learning Representations (ICLR)*, 2021. 2
- [53] 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. 1, 2
- [54] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning (ICML)*, 2021. 2
- [55] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *NeurIPS*, 2019. 2
- [56] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *International Conference on Machine Learning (ICML)*, 2016. 2
- [57] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 3
- [58] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 5
- [59] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *arXiv preprint arXiv:2208.12242*, 2022. 2, 5, 8

- [60] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. In *NeurIPS*, 2022. 1, 2
- [61] Axel Sauer, Kashyap Chitta, Jens Müller, and Andreas Geiger. Projected gans converge faster. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [62] Axel Sauer, Katja Schwarz, and Andreas Geiger. Stylegan-xl: Scaling stylegan to large diverse datasets. In *ACM SIGGRAPH*, 2022. 2
- [63] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. 4
- [64] Gabriel Skantze and Bram Willemsen. Collie: Continual learning of language grounding from language-image embeddings. *Journal of Artificial Intelligence Research*, 74:1201–1223, 2022. 2
- [65] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning (ICML)*, 2015. 2, 3
- [66] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations (ICLR)*, 2021. 3
- [67] Ming Tao, Hao Tang, Songsong Wu, Nicu Sebe, Xiao-Yuan Jing, Fei Wu, and Bingkun Bao. Df-gan: Deep fusion generative adversarial networks for text-to-image synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [68] Ngoc-Trung Tran, Viet-Hung Tran, Ngoc-Bao Nguyen, Trung-Kien Nguyen, and Ngai-Man Cheung. Towards good practices for data augmentation in gan training. *arXiv preprint arXiv:2006.05338*, 2, 2020. 2
- [69] Aäron Van Den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International Conference on Machine Learning (ICML)*, 2016. 2
- [70] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 2, 3
- [71] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 2
- [72] Sheng-Yu Wang, David Bau, and Jun-Yan Zhu. Rewriting geometric rules of a gan. *ACM SIGGRAPH*, 2022. 2
- [73] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [74] Yaxing Wang, Abel Gonzalez-Garcia, David Berga, Luis Herranz, Fahad Shahbaz Khan, and Joost van de Weijer. Minegan: effective knowledge transfer from gans to target domains with few images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [75] Yaxing Wang, Chenshen Wu, Luis Herranz, Joost van de Weijer, Abel Gonzalez-Garcia, and Bogdan Raducanu. Transferring gans: generating images from limited data. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [76] Wikipedia. Moon gate. https://en.wikipedia.org/wiki/Moon_gate, 2022. 1
- [77] XavierXiao. Dreambooth on stable diffusion. <https://github.com/XavierXiao/Dreambooth-Stable-Diffusion>, 2022. 5
- [78] 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [79] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gungjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2022. 1, 2
- [80] 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 *IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [81] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [82] Miaoyun Zhao, Yulai Cong, and Lawrence Carin. On leveraging pretrained gans for generation with limited data. In *International Conference on Machine Learning (ICML)*, 2020. 2
- [83] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient gan training. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [84] Zhengli Zhao, Zizhao Zhang, Ting Chen, Sameer Singh, and Han Zhang. Image augmentations for gan training. *arXiv preprint arXiv:2006.02595*, 2020. 2
- [85] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros. Generative visual manipulation on the natural image manifold. In *European Conference on Computer Vision (ECCV)*, 2016. 2
- [86] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [87] Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. Dm-gan: Dynamic memory generative adversarial networks for text-to-image synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2