

Customizing Text-to-Image Models with a Single Image Pair

MAXWELL JONES, Carnegie Mellon University, United States of America
 SHENG-YU WANG, Carnegie Mellon University, United States of America
 NUPUR KUMARI, Carnegie Mellon University, United States of America
 DAVID BAU, Northeastern University, United States of America
 JUN-YAN ZHU, Carnegie Mellon University, United States of America

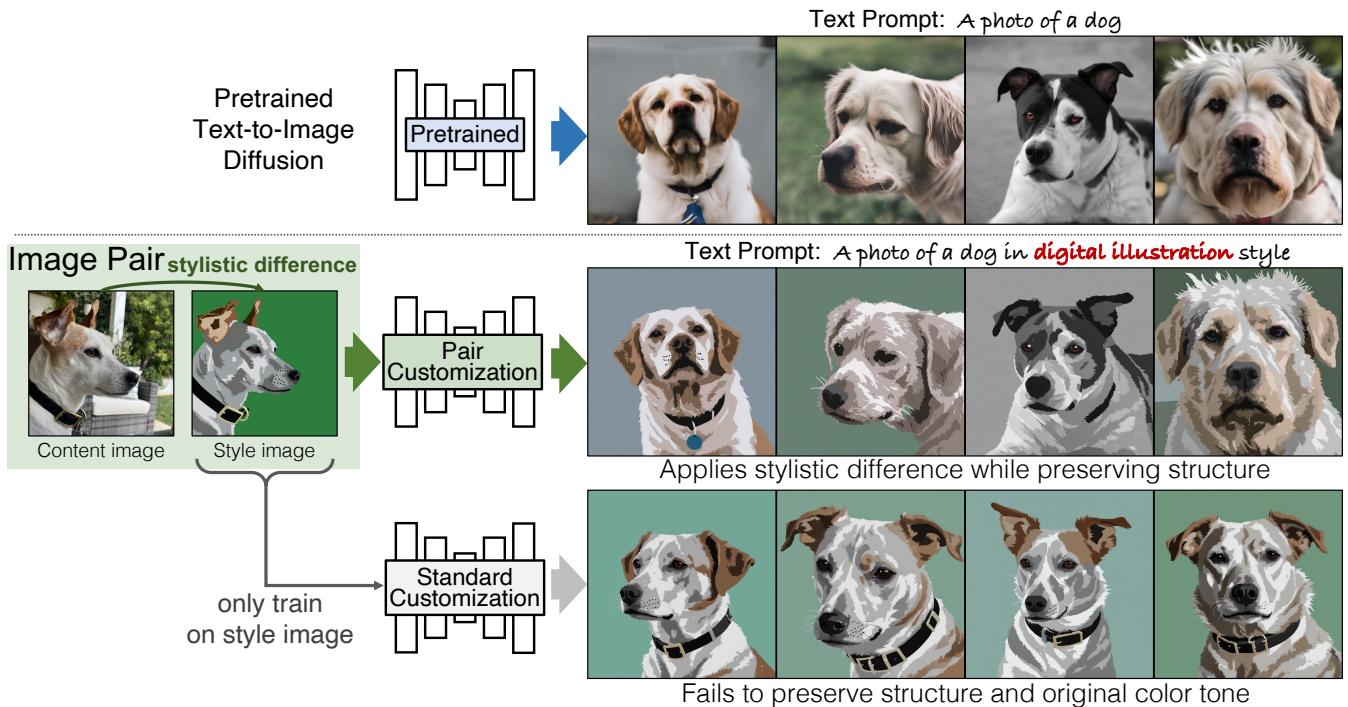


Fig. 1. Given a *single* image pair, we present *Pair Customization*, a method for customizing a pre-trained text-to-image model and learning a new style from the image pair's stylistic difference. Our method can apply the learned stylistic difference to new input images while preserving the input structure. Compared to Dreambooth LoRA[Hu et al. 2021; Ryu 2023a], a standard customization method that solely uses style images, our method effectively disentangles style and content, resulting in better structure, color preservation, and style application. Style image credit: Jack Parkhouse.

Art reinterpretation is the practice of creating a variation of a reference work, making a paired artwork that exhibits a distinct artistic style. We ask if such an image pair can be used to customize a generative model to capture the demonstrated stylistic difference. We propose *Pair Customization*, a new

Authors' Contact Information: Maxwell Jones, Carnegie Mellon University, United States of America, mjones2@andrew.cmu.edu; Sheng-Yu Wang, Carnegie Mellon University, United States of America, shengyu2@andrew.cmu.edu; Nupur Kumari, Carnegie Mellon University, United States of America, nkumari@andrew.cmu.edu; David Bau, Northeastern University, United States of America, david.bau@gmail.com; Jun-Yan Zhu, Carnegie Mellon University, United States of America, junyanz@andrew.cmu.edu.



This work is licensed under a Creative Commons Attribution International 4.0 License.

SA Conference Papers '24, December 03–06, 2024, Tokyo, Japan

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1131-2/24/12

<https://doi.org/10.1145/3680528.3687642>

customization method that learns stylistic difference from a *single* image pair and then applies the acquired style to the generation process. Unlike existing methods that learn to mimic a single concept from a collection of images, our method captures the stylistic difference between paired images. This allows us to apply a stylistic change without overfitting to the specific image content in the examples. To address this new task, we employ a joint optimization method that explicitly separates the style and content into distinct LoRA weight spaces. We optimize these style and content weights to reproduce the style and content images while encouraging their orthogonality. During inference, we modify the diffusion process via a new style guidance based on our learned weights. Both qualitative and quantitative experiments show that our method can effectively learn style while avoiding overfitting to image content, highlighting the potential of modeling such stylistic differences from a single image pair.

CCS Concepts: • Computing methodologies → Computer vision; Image manipulation.

Additional Key Words and Phrases: Image Generation, Diffusion Models, Deep Generative Models, Style Transfer

ACM Reference Format:

Maxwell Jones, Sheng-Yu Wang, Nupur Kumari, David Bau, and Jun-Yan Zhu. 2024. Customizing Text-to-Image Models with a Single Image Pair. In *SIGGRAPH Asia 2024 Conference Papers (SA Conference Papers '24), December 03–06, 2024, Tokyo, Japan*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3680528.3687642>

1 Introduction

Artistic works are often inspired by a reference image, a recurring scene, or even a previous piece of art [Markus 2019]. Such creations involve re-interpreting an original composition in the artist’s unique style. A notable example is Van Gogh’s *Repetitions* [Phillips Collection 2013], in which the artist created multiple versions of the same scenes with his distinctive expressiveness, including adaptations of other artists’ work. Such sets of variations allow close comparison of stylized art to a reference image, providing unique insights into an artist’s detailed techniques and choices.

In our work, we explore how such *content-style image pairs* can be used to customize a generative model to capture the demonstrated stylistic difference. Our goal is to customize a pre-trained generative model to synthesize stylized images, distilling the essence of the style from as few as a single pair without fixating on specific content. We wish to create a model capable of re-interpreting a variety of different content in the style demonstrated by the paired variation.

Prior works on model customization/personalization [Gal et al. 2022a; Kumari et al. 2023; Ruiz et al. 2023a] take one or a few images of a single concept to customize large-scale text-to-image models [Ramesh et al. 2022; Rombach et al. 2022]. While they aim to learn styles without using pairs, the generated samples from these customized models often resemble the training images’ content, such as specific objects, persons, and scene layouts. In Figure 1, we observe that standard single-image customization (3rd row) alters the subject, color tone, and pose of the original image (1st row). These issues arise because the artistic intent is difficult to discern from a single image: unlike image pairs that can demonstrate a style through contrasts, a singleton example will always intertwine choices of both style and content. Due to this ambiguity, the model fails to capture the artistic style accurately and, in some cases, overfits and generates the subject-specific details rather than the style, as shown in Figure 5.

On the other hand, our *Pair Customization* method exploits the contrast between image pairs to generate pairwise consistent images while better disentangling style and content. In Figure 1 (2nd row), our method accurately follows the given style, turning the background into a single color matching the original background and preserving the identity and pose for each dog. Our method achieves this by disentangling the intended style from the image pair.

Our new customization task is challenging since text-to-image models were not initially designed to generate *pairwise* content. Even when given specific text prompts like “*a portrait*” and “*a portrait with Picasso style*”, a text-to-image diffusion model often struggles to generate images with consistent structure from the same noise seed. Therefore, it remains unclear how a customized

model can generate stylized images while maintaining the original structure.

To address the challenges, we first propose a joint optimization method with separate sets of Low-Rank Adaptation [Hu et al. 2021] (LoRA) weights for style and content. The optimization encourages the content LoRA to reconstruct the content image and the style LoRA to apply the style to the content. We find that the resulting style LoRA can apply the same style to other unseen content. Furthermore, we enforce row-space orthogonality [Po et al. 2023] between style and content LoRA parameters to improve style and content disentanglement. Next, we extend the standard classifier-free guidance method [Ho and Salimans 2022] and propose style guidance. Style guidance integrates style LoRA predictions into the original denoising path, which aids in better content preservation and facilitates smoother control over the stylization strength. This method is more effective than the previous technique, where a customized model’s strength is controlled by the magnitude of LoRA weights [Ryu 2023b].

Our method is built upon Stable Diffusion XL [Podell et al. 2023]. We experiment with various image pairs, including different categories of content (e.g., portraits, animals, landscapes) and style (e.g., paintings, digital illustrations, filters). We evaluate our method on the above single image pairs and demonstrate the advantage of our method in preserving diverse structures while applying the stylization faithfully, compared to existing customization methods. Our code, models, and data are available on our webpage.

2 Related Works

Text-to-image generative models. Deep generative models aim to model the data distribution of a given training set [Dinh et al. 2017; Goodfellow et al. 2020; Ho et al. 2020; Kingma and Welling 2014; Song et al. 2021b; Van den Oord et al. 2016]. Recently, large-scale text-to-image models [Balaji et al. 2022; Chang et al. 2023; Gokaslan et al. 2023; Kang et al. 2023; Luo et al. 2023; Peebles and Xie 2023; Podell et al. 2023; Ramesh et al. 2022; Rombach et al. 2022; Saharia et al. 2022b; Sauer et al. 2023; Yu et al. 2022] trained on internet-scale training data [Byeon et al. 2022; Schuhmann et al. 2021] have shown exceptional generalization. Notably, diffusion models [Ho et al. 2020; Song et al. 2021a] stand out as the most widely adopted model class. While existing models can generate a broad spectrum of objects and concepts, they often struggle with rare or unseen concepts. Our work focuses on teaching these models to understand and depict a new style concept. Conditional generative models [Brooks et al. 2023; Isola et al. 2017; Li et al. 2023; Mou et al. 2024; Park et al. 2019; Saharia et al. 2022a; Zhang and Agrawala 2023] learn to transform images across different domains, but the training often requires thousands to millions of image pairs. We focus on a more challenging case, where only a single image pair is available.

Customizing generative models. Model customization, or personalization, aims to adapt an existing generative model with additional data, with the goal of generating outputs tailored to specific user preferences. Earlier efforts mainly focus on customizing pre-trained GANs [Goodfellow et al. 2020; Karras et al. 2019, 2020b] for smaller datasets [Karras et al. 2020a; Nitzan et al. 2022; Zhao et al.

2020], incorporating user edits [Bau et al. 2020; Wang et al. 2021, 2022], or aligning with text prompts [Gal et al. 2022b; Nitzan et al. 2023]. Recently, the focus has pivoted towards adapting large-scale text-to-image models to generate user-provided concepts, typically presented as one or a few images. Simply fine-tuning on the concept leads to overfitting. To mitigate this and enable variations via free text, several works explored different regularizations, including prior preservation [Kumari et al. 2023; Ruiz et al. 2023a], human alignment [Sohn et al. 2023], patch-based learning [Zhang et al. 2023], as well as parameter update restriction, where we only update text tokens [Alaluf et al. 2023; Daras and Dimakis 2022; Gal et al. 2022a; Voynov et al. 2023], attention layers [Gal et al. 2023a; Han et al. 2023; Kumari et al. 2023], low-rank weights [Hu et al. 2021; Ryu 2023a; Tewel et al. 2023], or clusters of neurons [Liu et al. 2023]. More recent methods focus on encoder-based approaches for faster personalization [Arar et al. 2023; Chen et al. 2023a,b; Gal et al. 2023b; Li et al. 2024; Ma et al. 2023; Ruiz et al. 2023b; Shi et al. 2023; Valevski et al. 2023; Wei et al. 2023; Ye et al. 2023]. Instead of learning a single concept, several works further focus on learning multiple concepts [Avrahami et al. 2023; Gu et al. 2024; Kumari et al. 2023; Po et al. 2023; Shah et al. 2023]. Other methods [Materzynska et al. 2023; Ren et al. 2024] propose customizing text-to-video models to learn motion, while Guo et al. [2023] propose animating customized text-to-image models models by incorporating motion Low-Rank Adapter [Hu et al. 2021] modules. Our method takes inspiration from these techniques; however, we aim to address an inherently different task. Instead of learning concepts from an image collection, we customize the model to *learn stylistic differences* from an image pair.

Style and content separation. Various past works have explored learning a style while separating it from content [Chen and Schmidt 2016; Gatys et al. 2015; Huang and Belongie 2017; Li et al. 2017; Tenenbaum and Freeman 1996]. Our work is inspired by the seminal work Image Analogy [Hertzmann et al. 2001], a computational paradigm that takes an image pair and applies the same translation to unseen images. Common image analogy methods include patch-wise similarity matching [Hertzmann et al. 2001; Irony et al. 2005; Liao et al. 2017] and data-driven approaches [Bar et al. 2022; Park et al. 2020; Reed et al. 2015; Upchurch et al. 2016; Wang et al. 2023; Zhu et al. 2017]. Different from these, we aim to exploit the text-guided generation capabilities of large-scale models so that we can *directly* use the style concept with unseen context. Recently, StyleDrop [Sohn et al. 2023] has proposed to learn a custom style for masked generative transformer models. Unlike StyleDrop, we do not rely on human feedback in the process. Concurrent with our work, Hertz et al. [Hertz et al. 2023] introduced a method for generating images with style consistency, offering the option of using a style reference image. In contrast, we exploit an image pair to better discern the stylistic difference.

3 Method

Our method seeks to learn a new style from a single image pair. This task is challenging, as models tend to overfit when trained on a single image, especially when generating images in the same category as the training image (e.g., a model trained and tested

on dog photos). To reduce this overfitting, we introduce a new algorithm aimed at disentangling the structure of the subject from the style of the artwork. Specifically, we leverage the image pair to learn separate model weights for style and content. At inference time, we modify the standard classifier-free guidance formulation to help preserve the original image structure when applying the learned style. In this section, we give a brief overview of diffusion models, outline our design choices, and explain the final method in detail.

3.1 Preliminary: Model Customization

Diffusion models. Diffusion models [Ho et al. 2020; Sohl-Dickstein et al. 2015; Song et al. 2021b], map Gaussian noise to the image distribution through iterative denoising. Denoising is learned by reversing the forward diffusion process $\mathbf{x}_0, \dots, \mathbf{x}_T$, where image \mathbf{x}_0 is slowly *diffused* to random noise \mathbf{x}_T over T timesteps, defined by $\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$ for timestep $t \in [0, T]$. Noise $\epsilon \sim \mathcal{N}(0, I)$ is randomly sampled, and $\bar{\alpha}_t$ controls the noise strength. The training objective of diffusion models is to denoise any intermediate noisy image \mathbf{x}_t via noise prediction:

$$\mathbb{E}_{\epsilon, \mathbf{x}, \mathbf{c}, t} [\mathbf{w}_t \|\epsilon - \epsilon_\theta(\mathbf{x}_t, \mathbf{c}, t)\|^2], \quad (1)$$

where \mathbf{w}_t is a time-dependent weight, $\epsilon_\theta(\cdot)$ is the denoiser that learns to predict noise, and \mathbf{c} denotes extra conditioning input, such as text. At inference, the denoiser ϵ_θ will gradually denoise random Gaussian noise into images. The resulting distribution of generated images approximates the training data distribution [Ho et al. 2020].

In our work, we use Stable Diffusion XL [Podell et al. 2023], a large-scale text-to-image model built on Latent Diffusion Models [Rombach et al. 2022]. The model consists of a U-Net [Ronneberger et al. 2015] trained on the latent space of an auto-encoder, with text conditioning from two text encoders, CLIP [Radford et al. 2021] and OpenCLIP [Ilharco et al. 2021].

Model customization with Low-Rank Adapters. Low-Rank Adapters (LoRA) [Hu et al. 2021] is a parameter-efficient fine-tuning method [Houlsby et al. 2019] that applies low-rank weight changes $\Delta\theta_{\text{LoRA}}$ to pre-trained model weights θ_0 . For each layer with an initial weight $W_0 \in \mathbb{R}^{m \times n}$, the weight update is defined by $\Delta W_{\text{LoRA}} = BA$, a product of learnable matrices $B \in \mathbb{R}^{m \times r}$ and $A \in \mathbb{R}^{r \times n}$, where $r \ll \min(m, n)$ to enforce the low-rank constraint. The weight matrix of a particular layer with LoRA is:

$$W_{\text{LoRA}} = W_0 + \Delta W_{\text{LoRA}} = W_0 + BA. \quad (2)$$

At inference time, the LoRA strength is usually controlled by a scaling factor $\alpha \in [0, 1]$ applied to the weight update ΔW_{LoRA} [Ryu 2023b]:

$$W_{\text{LoRA}} = W_0 + \alpha \Delta W_{\text{LoRA}}. \quad (3)$$

LoRA has been applied for customizing text-to-image diffusion models to learn new concepts with as few as three to five images [Ryu 2023b].

3.2 Style Extraction from an image pair

We aim to customize a pre-trained model with an artistic style in order to stylize the original model outputs while preserving their

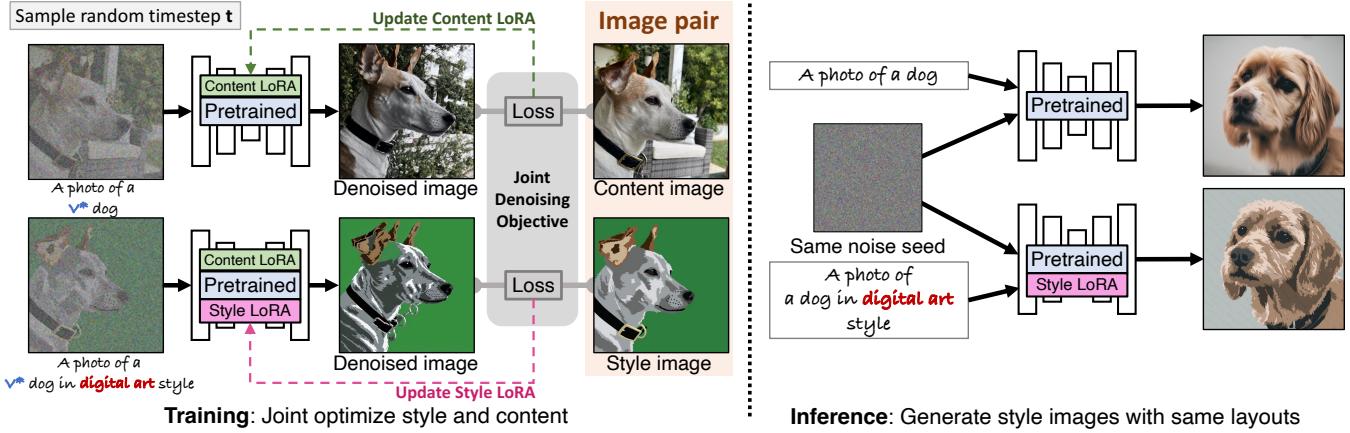


Fig. 2. **Method overview.** (Left) We disentangle style and content from an image pair by jointly training two low-rank adapters, StyleLoRA and ContentLoRA, representing style and content, respectively. Our training objective consists of two losses: The first loss fine-tunes ContentLoRA to reconstruct content image conditioned on a content prompt. The second loss encourages reconstructing the style image using *both* StyleLoRA and ContentLoRA conditioned on a style prompt, but we only optimize StyleLoRA for this loss. (Right) At inference time, we only apply StyleLoRA to customize the model. Given the same noise seed, the customized model generates a stylized counterpart of the original pre-trained model output. V^* is a fixed random rare token that is a prompt modifier for the content image. Style image credits: Jack Parkhouse

content, as shown in Figure 2 (right). To achieve this, we introduce style LoRA weight $\theta_{\text{style}} = \theta_0 + \Delta\theta_{\text{style}}$. While a pre-trained model generates content from a noise seed and text c , style LoRA’s goal is to generate a stylized counterpart of original content from the same noise seed and a style-specific text prompt c_{style} , where c_{style} is original text c appended by suffix “in <desc> style”. Here, <desc> is a placeholder for some worded description of the style (e.g., “digital art”), and style LoRA θ_{style} associates <desc> to the desired style.

Unfortunately, learning style LoRA θ_{style} from a single style image often leads to copying content (Figure 5). Hence, we explicitly learn disentanglement from a style and content image, denoted by x_{style} and x_{content} , respectively.

Disentangling style and content. We leverage the fact that the style image shares the same layout and structure as the content image. Our key idea is to learn a separate content LoRA $\theta_{\text{content}} = \theta_0 + \Delta\theta_{\text{content}}$ to reconstruct the content image. By explicitly modeling the content, we can train the style LoRA to “extract” the stylistic differences between the two images. We apply both style and content LoRA to reconstruct the style image, i.e., $\theta_{\text{combined}} = \theta_0 + \Delta\theta_{\text{content}} + \Delta\theta_{\text{style}}$. This approach prevents leaking the content image to style LoRA, resulting in a better stylization model.

During training, we feed the content LoRA θ_{content} with a content-specific text c_{content} , which contains a random rare token V^* , and feed the combined model θ_{combined} with c_{style} , where c_{style} is “ $\{c_{\text{content}}\}$ in <desc> style”. Figure 2 (Left) summarizes our training process.

Jointly learning style and content. We employ two different objectives during every training step. To learn the content of the image, we first employ the standard training objective for diffusion

models as described in Section 3.1 with the content image:

$$\mathcal{L}_{\text{content}} = \mathbb{E}_{\epsilon, x_{\text{content}}, t} [w_t \|\epsilon - \epsilon_{\theta_{\text{content}}}(x_{t, \text{content}}, c_{\text{content}}, t)\|^2], \quad (4)$$

where $\epsilon_{\theta_{\text{content}}}$ is the denoiser with content LoRA applied, $x_{t, \text{content}}$ is a noisy content image at timestep t , and c_{content} is text representing the content image, including some rare token V^* . Next, we optimize the combined style and content weights to reconstruct the style image. In particular, we only train the style LoRA weights during this step, while stopping the gradient flow to the content LoRA weights via stopgrad $\text{sg}[\cdot]$:

$$\theta_{\text{combined}} = \theta_0 + \text{sg}[\Delta\theta_{\text{content}}] + \Delta\theta_{\text{style}}. \quad (5)$$

We then apply diffusion objective to train θ_{combined} to denoise $x_{t, \text{style}}$, a noisy style image at timestep t :

$$\mathcal{L}_{\text{combined}} = \mathbb{E}_{\epsilon, x_{\text{style}}, t} [w_t \|\epsilon - \epsilon_{\theta_{\text{combined}}}(x_{t, \text{style}}, c_{\text{style}}, t)\|^2], \quad (6)$$

where $\epsilon_{\theta_{\text{combined}}}$ is the denoiser with both LoRAs applied as in Equation 5, c_{style} is “ $\{c_{\text{content}}\}$ in <desc> style”, and <desc> is a worded description of the style (e.g., “digital art”). Finally, we jointly optimize the LoRAs with the two losses:

$$\min_{\Delta\theta_{\text{content}}, \Delta\theta_{\text{style}}} \mathcal{L}_{\text{content}} + \mathcal{L}_{\text{combined}} \quad (7)$$

Figure 2 provides an overview of our method. Next, we discuss the regularization that promotes disentanglement of style from content.

Orthogonality between style and content LoRA. To further encourage style and content LoRAs to represent separate concepts, we enforce orthogonality upon the LoRA weights. We denote by W_0 the original weight matrix and W_{content} , W_{style} the LoRA modifications (layer index omitted for simplicity). Reiterating Equation 2, we decompose W_{content} , W_{style} into low-rank matrices:

$$W_{\text{content}} = W_0 + B_{\text{content}} A_{\text{content}}; \quad W_{\text{style}} = W_0 + B_{\text{style}} A_{\text{style}}. \quad (8)$$

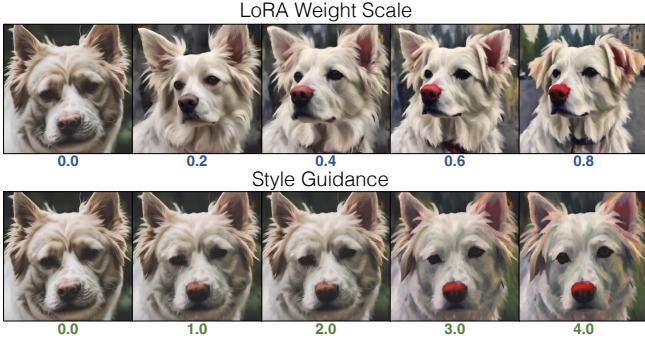


Fig. 3. **Style guidance.** We compare our style guidance and standard LoRA weight scaling [Ryu 2023b]. Style guidance better preserves content when the style is applied. **Blue** and **green** stand for the LoRA weight scale and style guidance scale, respectively. More details of style guidance formulation are in Section 3.3.

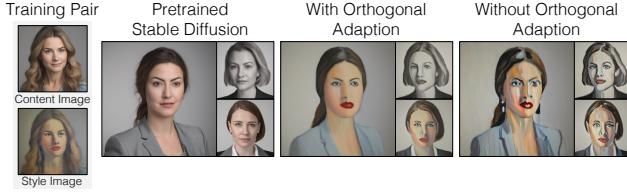


Fig. 4. **Orthogonal adaptation.** Enforcing row-space orthogonality between style and content LoRA improves image quality, where the images capture the style better and have fewer visual artifacts.

We initialize B_{content} , B_{style} with the zero matrix and choose the rows of A_{content} , A_{style} from an orthonormal basis. We then fix A_{content} , A_{style} and only update B_{content} , B_{style} in training. This forces the style and content LoRA updates to respond to orthogonal inputs, and empirically reduces visual artifacts, as shown in Figure 4. This technique is inspired by Po et al. [Po et al. 2023]. While their work focuses on merging multiple customized objects after each is trained separately, we apply the method for style-content separation during joint training.

3.3 Style Guidance

A common technique to improve text-to-image model's sample quality is via classifier-free guidance [Ho and Salimans 2022]:

$$\hat{\epsilon}_{\theta}(\mathbf{x}_t, \mathbf{c}) = \epsilon_{\theta}(\mathbf{x}_t, \emptyset) + \lambda_{\text{cfg}}(\epsilon_{\theta}(\mathbf{x}_t, \mathbf{c}) - \epsilon_{\theta}(\mathbf{x}_t, \emptyset)), \quad (9)$$

where $\hat{\epsilon}_{\theta}(\mathbf{x}_t, \mathbf{c}, t)$ is the new noise prediction, \emptyset denotes no conditioning, and λ_{cfg} controls the amplification of text guidance. For notation simplicity, we omit the timestep t in this equation and subsequent ones.

To improve pairwise consistency between original and stylized content, we propose an inference algorithm that preserves the original denoising path while adding controllable style guidance:

$$\begin{aligned} \hat{\epsilon}_{\theta_0, \theta_{\text{style}}}(\mathbf{x}_t, \mathbf{c}, \mathbf{c}_{\text{style}}) &= \epsilon_{\theta_0}(\mathbf{x}_t, \emptyset) \\ &+ \lambda_{\text{cfg}}(\epsilon_{\theta_0}(\mathbf{x}_t, \mathbf{c}) - \epsilon_{\theta_0}(\mathbf{x}_t, \emptyset)) \\ &+ \lambda_{\text{style}}(\epsilon_{\theta_{\text{style}}}(\mathbf{x}_t, \mathbf{c}_{\text{style}}) - \epsilon_{\theta_0}(\mathbf{x}_t, \mathbf{c})), \end{aligned} \quad (10)$$

where style guidance is the difference in noise prediction between style LoRA and the pre-trained model. Style guidance strength is controlled by λ_{style} and setting $\lambda_{\text{style}} = 0$ is equivalent to generating original content. In Figure 3, we compare our style guidance against scaling LoRA weights (Equation 3), and we find that our method better preserves the layout. More details and a derivation of our style guidance are in the supplement.

Previous works have also used multiple guidance terms with diffusion models, including guidance from multiple text prompts using the same model [Liu et al. 2022a] and additional image conditions [Brooks et al. 2023]. Unlike these, we obtain additional guidance from a customized model and apply it to the original model. StyleDrop [Sohn et al. 2023] considers a similar formulation with two guidance terms but for masked generative transformers. SINE [Zhang et al. 2023] uses a customized content model to apply text-based image editing to a single image, like adding snow. However, we use a customized style model to generate any image with the desired style.

Blending multiple learned styles. With a collection of models customized by our method, we can blend the learned styles as follows. Given some set of styles \mathcal{S} and strengths $\lambda_{\text{style}_0}, \dots, \lambda_{\text{style}_n}$, we blend the style guidance from each model, and our new inference path is represented by

$$\begin{aligned} \hat{\epsilon}_{\theta_0, \theta_{\text{style}}}(\mathbf{x}_t, \mathbf{c}, \mathbf{c}_{\text{style}}) &= \epsilon_{\theta_0}(\mathbf{x}_t, \emptyset) \\ &+ \lambda_{\text{cfg}}(\epsilon_{\theta_0}(\mathbf{x}_t, \mathbf{c}) - \epsilon_{\theta_0}(\mathbf{x}_t, \emptyset)) \\ &+ \sum_{\text{style}_i \in \mathcal{S}} \lambda_{\text{style}_i}(\epsilon_{\theta_{\text{style}_i}}(\mathbf{x}_t, \mathbf{c}_{\text{style}_i}) - \epsilon_{\theta_0}(\mathbf{x}_t, \mathbf{c})), \end{aligned} \quad (11)$$

We can vary the strengths of any parameter λ_{style_i} to seamlessly increase or decrease style application while preserving content. Figure 10 gives a qualitative example of blending two different styles while preserving image content.

Implementation details. We train all models using an AdamW optimizer [Loshchilov and Hutter 2018] and learning rate 1×10^{-5} . For baselines, we train for 500 steps. For our method, we first train our content weights on the content image for 250 steps and then train jointly for 500 additional steps. All image generation is performed using 50 steps of a PNDMScheduler [Liu et al. 2022b]. For all methods using inference with LoRA adapters, we use SDEdit [Meng et al. 2022] to further preserve the structure. Specifically, normal classifier-free guidance on the original prompt without style is used for the first 10 steps. We then apply style guidance/LoRA scale for the rest of the timesteps.

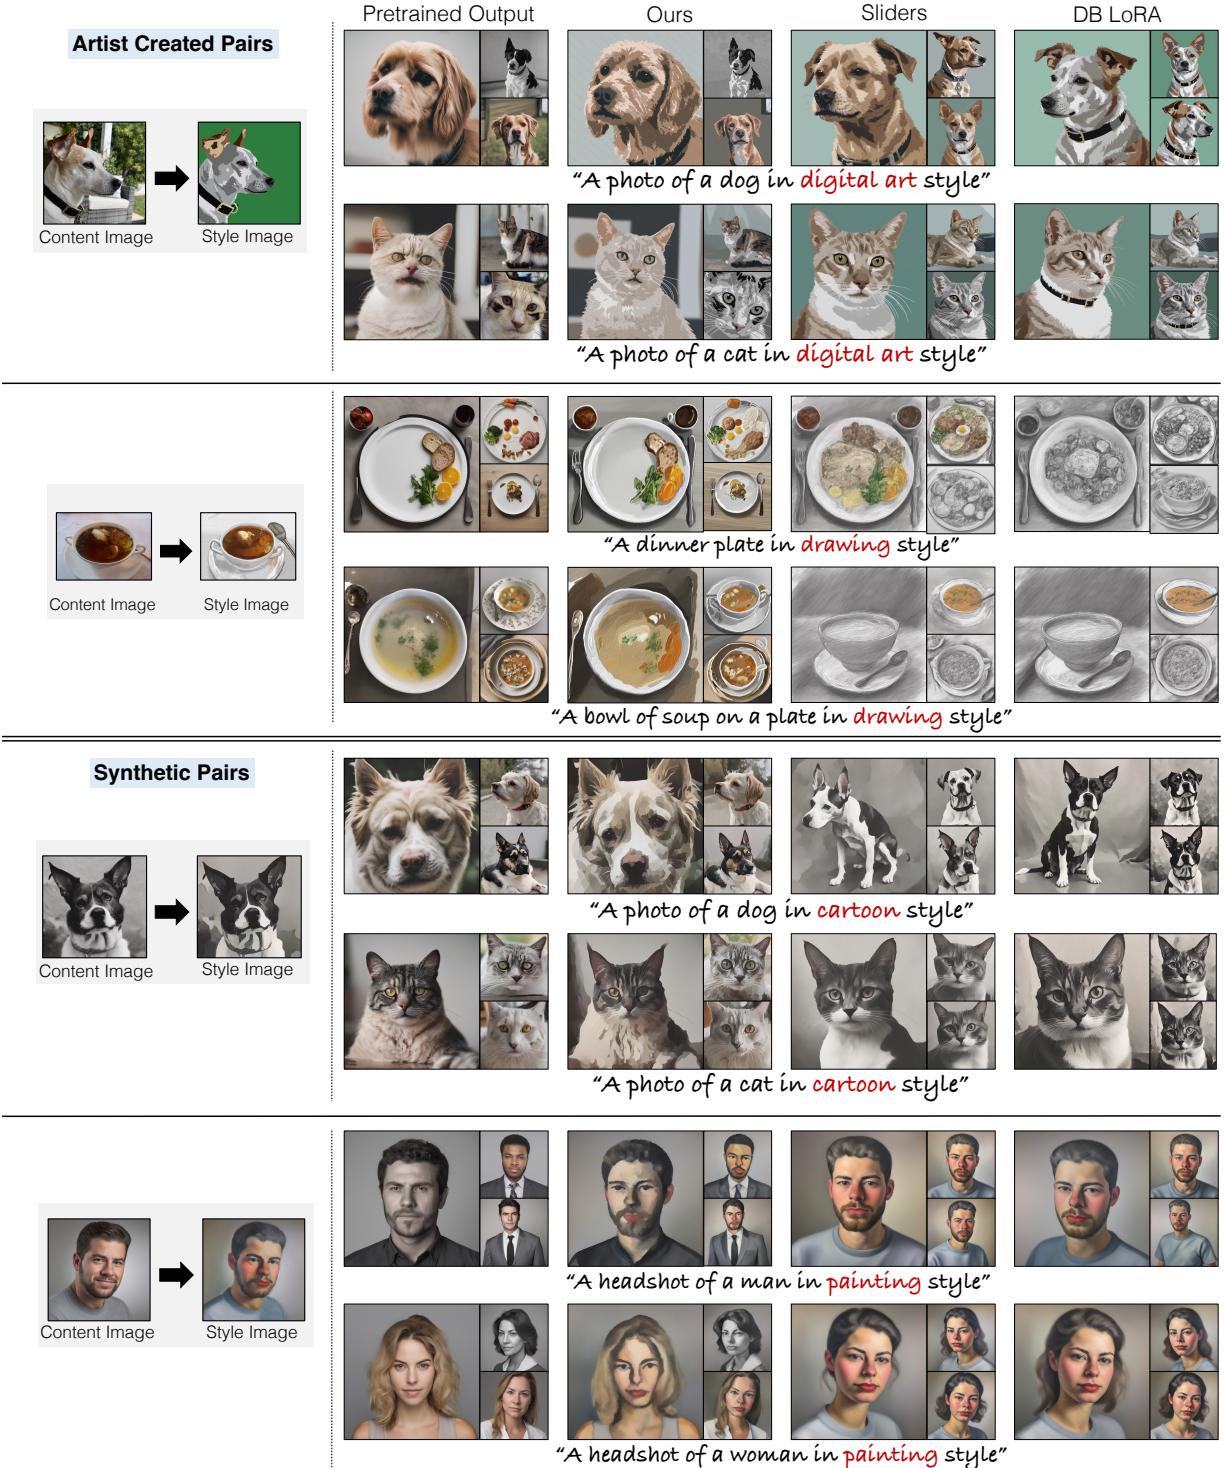


Fig. 5. Result of our method compared to the strongest baselines. When only training with the style image as in DB LoRA, the image structure is not preserved, and overfitting occurs. While Concept Slider's training scheme [Gandikota et al. 2023] uses both style and content images, it still exhibits overfitting and loss of structure in many cases. Our method preserves the structure of the input image while faithfully applying the desired style. We use style guidance strength 3 and classifier-free guidance strength 5. Style image credits: Jack Parkhouse (First row) and Aaron Hertzmann (Second row)

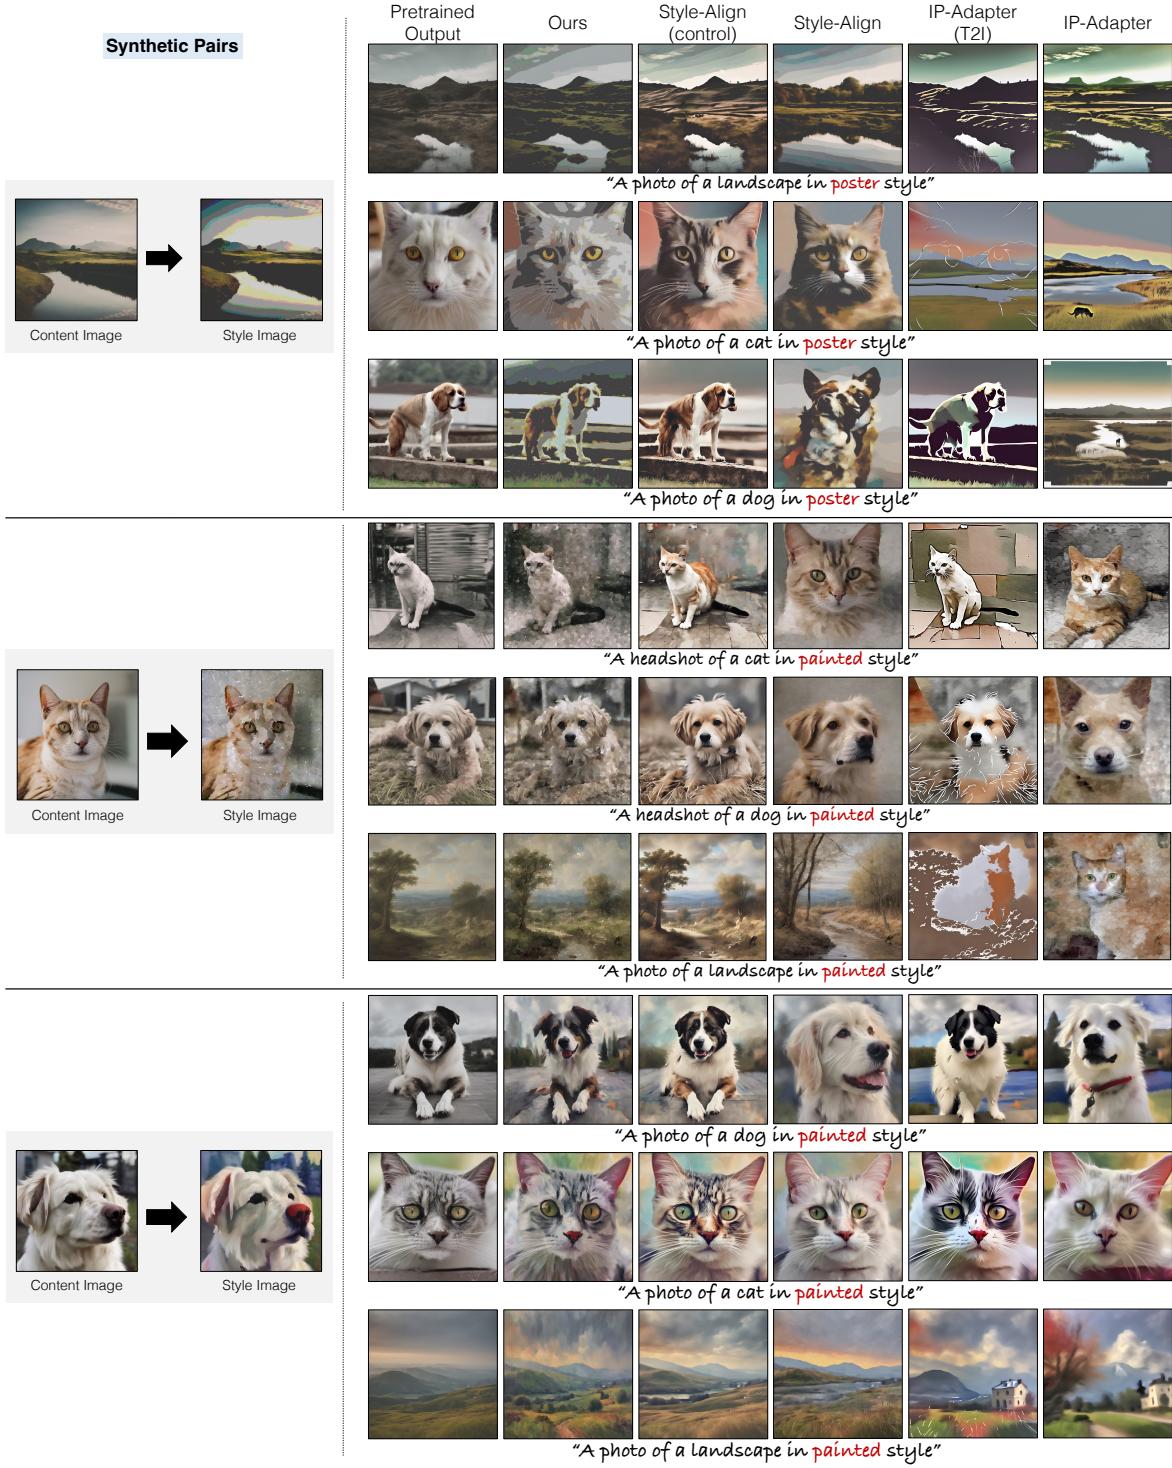


Fig. 6. Result of our method compared to the methods without finetuning (zoom in for best viewing). For the baseline methods, we also add the edge map from the pretrained output as an extra condition (3rd and 5th column). Without this edge map, other methods tend to lose the structure of the pretrained output. In some cases, however, an additional edge map can overly constrain the output of a model, like in the second image pair example. Our method preserves the structure of the Stable Diffusion image while faithfully applying the desired style. We use style guidance strength 3 and classifier-free guidance strength 5 for our method and set the IP-adapter scale and style-alignment scale to 0.5.

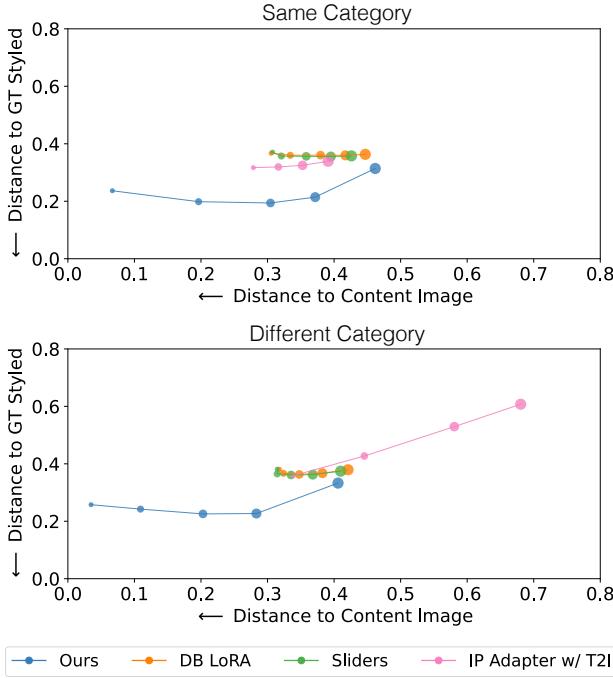


Fig. 7. **Quantitative comparison with baselines on learned style.** Given a fixed inference path, our method pareto dominates baselines for image generation both on the same category as training (left) and when evaluated on categories different from training, e.g., trained on human portraits but tested on dog images (right). We further evaluate the diversity of generated images in the supplement. We show that baselines often lose diversity, while our method leads to diverse generations while still achieving lower perceptual distance to the ground truth style. Increased marker size corresponds to an increase in the guidance scale.

4 Experiments

4.1 Dataset

In this section, we show our method’s results on various image pairs and compare them with several baselines. We explain our dataset, baselines, and metrics in detail, then we present quantitative and qualitative results.

Datasets. To enable large-scale quantitative evaluation, we construct a diverse set of paired style and content images as follows. First, we generate 40 content images for each class: headshots, animals, and landscapes. When generating images in the headshot class, we generate 20 images with the prompt “A professional headshot of a man” and 20 images with the prompt “A professional headshot of a woman”. Similarly, we split the animal class into photos of dogs and cats. To curate synthetic pairs, we then apply image editing or image-to-image translation methods to all the content images to obtain the stylized version. For each unique prompt, we choose a *single paired instance* as training data and hold out the other pairs with the same prompt as a test set (Same Category). For each prompt, we also choose 5 pairs from each of the other prompts as a secondary test set (Different Category). We

show all our synthetic training image pairs in the supplement. By leveraging synthetic pairs for evaluation, we can train on a single synthetic pair and test our results against held out synthetic style images. Secondly, we qualitatively compare against single artist pairs in Figure 5. Next, we describe the specific methods to create the paired dataset. First, we consider the diffusion-based image editing technique LEDITS++[Brack et al. 2023] to translate images into paintings. Next, we consider Cartoonization [Wang and Yu 2020], a GAN-based translation technique that applies a cartoon-like effect. We also consider Stylized Neural Painting [Zou et al. 2021], which turns photos into painting strokes using a rendering based approach. Finally, we consider the image filtering technique posterization. We provide a more detailed description of each method for creating synthetic pairs in the supplement.

4.2 Baselines and Evaluation Metrics

Baselines. We compare our method against – (1) DreamBooth LoRA [Hu et al. 2021; Ryu 2023b] (DB LoRA), (2) Concept Sliders [Gandikota et al. 2023] (3) IP-adapters [Ye et al. 2023], (4) IP-adapters w/ T2I, and (5) StyleDrop [Sohn et al. 2023]. DB LoRA uses only the style image and fine-tunes low-rank adapters in all the linear layers in the attention blocks of the diffusion model. We evaluate different amounts of style applications for DB LoRA using the standard LoRA scale [Ryu 2023b]. Concept sliders presents a paired image model customization method that trains a single low-rank adapter jointly on both images, with different reconstruction losses for the style and content images. We also evaluate using the standard LoRA scale. IP-adapters is an encoder-based method that does not require training for every style and takes a style image as an extra condition separate from the text prompt. Increasing or decreasing the guidance from the input style image is possible by scaling the weight of the image conditioning. We consider the SDXL [Podell et al. 2023] implementation of this method. For the IP-Adapter, we compare against the stronger baseline of providing extra conditioning of an edge map of the content image through T2I Adapters [Mou et al. 2024] to preserve the content image structure. The recently proposed Styledrop [Sohn et al. 2023] technique for learning new styles is based on MUSE [Chang et al. 2023], and uses human feedback in its method. Since MUSE is not publicly available, we follow Style-Aligned Image Generation’s [Hertz et al. 2023] setup, and implement a version of StyleDrop on SDXL. Specifically, we train low-rank linear layers following each Feed-Forward layer in the attention blocks of SDXL. For a fair comparison, we train Styledrop without human feedback.

Evaluation metrics. When evaluating the performance of each method, we consider two quantitative metrics: perceptual distance to ground truth style images and structure preservation from the original image. A better customization method will have a low perceptual distance to the ground truth style images while still preserving content of the original image before adding style. We measure these using – (1) *Distance to GT Styled*: given holdout ground truth style images, we measure the perceptual distance between our styled outputs and the ground truth style images using DreamSim [Fu et al. 2023], a recent method for measuring the perceptual distance

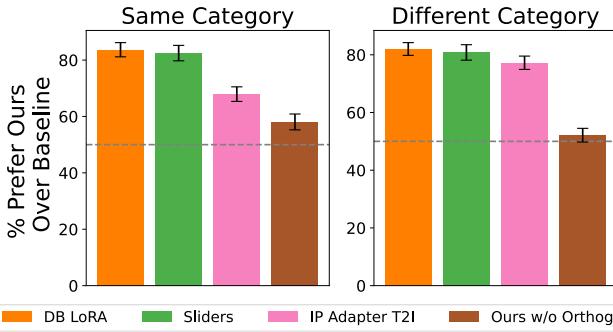


Fig. 8. **Human preference study.** Our method is preferred over the baselines ($\geq 60\%$). Further, our full method, including orthogonal weight matrices (Section 3.2), is preferred over the one w/o orthogonal weight matrices, specifically for the same category as training pair, e.g., trained on a headshot of a man and tested on other headshots of man. The Gray dashed line denotes 50% chance performance.

between images. DreamSim image embeddings are comprised of an ensemble of image embedding models, including CLIP [Radford et al. 2021] and DINO [Caron et al. 2021], which are then fine-tuned so the final embeddings respect human perception. We measure DreamSim distance as $(1 - \text{cosine similarity between DreamSim embeddings})$, where a lower value implies that the images are perceptually more similar. (2) *Distance to Content Image*: to measure content preservation after style application, we measure the perceptual distance of our generated style image to the original content image with no style guidance. We again use DreamSim, this time comparing styled and content images. Note here that a perceptual distance of zero to the content image is undesirable, as this would require no style to be applied. However, a better-performing method should obtain a better tradeoff between the two distances. (3) We also perform a *human preference study* of our method against baselines.

4.3 Results

Quantitative evaluation. We show quantitative results in Figure 7. Increased marker size (circles) indicates the higher application of style, and line color determines the method. When evaluating style similarity vs. structure preservation in Figure 7, we see that our training method’s Pareto dominates all baselines, yielding lower perceptual distance to style images while still being perceptually similar to the original content image. We find that DB LoRA and StyleDrop perform similarly, and report StyleDrop results in the supplement. Finally, we consider our method with LoRA scale during inference and other baselines with our style guidance during inference for ablation, and provide results in the supplement.

Qualitative evaluation. We compare our method with the highest performing baselines in Figure 5. The finetuning-based methods DB LoRA [Hu et al. 2021; Ryu 2023b] and Concept Sliders [Gandikota et al. 2023] outperform the encoder-based method [Ye et al. 2023] for our task. Hence, we compare against that in Figure 5. For both

baselines, we modulate style application with LoRA scale (Equation 3). We observe that DB LoRA often fails to generate the style-transformed version of the original image and overfits to the training pair image when generating similar concepts. There are two main reasons why this may occur. First, we are in a challenging case where there is only 1 training image instead of the usual 3 – 5 images that customization methods use. Second, we are prompting the model on the same or very similar text prompts to the training prompt, and the baseline method overfits to the training image for these prompts. Our method preserves the structure of the original image while applying the learned style. Moreover, applying our style guidance instead of the LoRA scale benefits the baseline method as well (Figure 5, last 2 columns), as it can better preserve the structure of the original image, though it still tends to overfit to the content of the training image. We observe a similar issue for other baselines as well.

We compare our method to non finetuning-based methods in Figure 6. We observe that these methods perform worse than finetuning-based methods, especially when generating images in a category different from the training style image. We also compare with baselines using our style guidance for style application at inference time in the supplement .

User preference study. We perform a user preference study using Amazon Mechanical Turk. We test our method against all baselines, as well as a version of our method trained without the orthogonality constraint. Specifically, we test on all datasets in Section 4.1. When evaluating against DB LoRA and Concept Sliders, we consider inference with both LoRA scale as in Equation 3 and style guidance as in Equation 10. For each method, we pick a single style strength that performs most optimally according to quantitative metrics as in Figure 7. Full details are available in the supplement . We collect 400 responses per paired test of ours vs the other method. The user is shown an image generated via our method and an image generated via the other method and asked to select the image that best applies the given style to the new content image. We provide a detailed setup of the user study as well as a user study on baseline methods using our style guidance in the supplement. As shown in Figure 8, our method is favored by users in comparison to baselines, whether evaluating images generated within the same category as the training image pair or across different categories. Secondly, users prefer our full method to ours without the orthogonality constraint, specifically when evaluating on the same category as training.

Real Image Editing. Our method can also stylize real images. We use DDIM inversion [Garibi et al. 2024; Song et al. 2021a] to invert images into their noisy latent codes at some intermediate step using a reference prompt c . From here, we use our style guidance (Equation 10) with reference prompt c and new prompt c_{style} to denoise the noisy latent code to a stylized image. In Figure 9, we show real image editing results. We provide more details in the supplement.

Blending learned styles. We show that we can blend the learned styles by applying a new inference path, defined in Equation 11. In Figure 10, we show the results of blending two models. We can

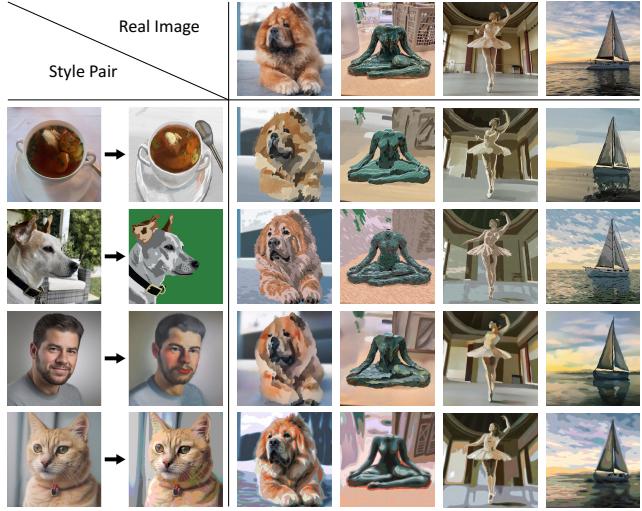


Fig. 9. **Real image editing.** We can edit real images by inverting a real image into a noisy latent code using a diffusion inversion pipeline [Garibi et al. 2024; Song et al. 2021a] and an image prompt c . From here, we apply style guidance (Equation 10) with the same prompt c and new style prompt c_{style} for the desired style application.

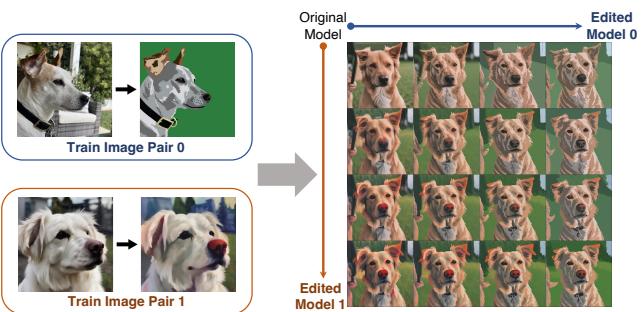
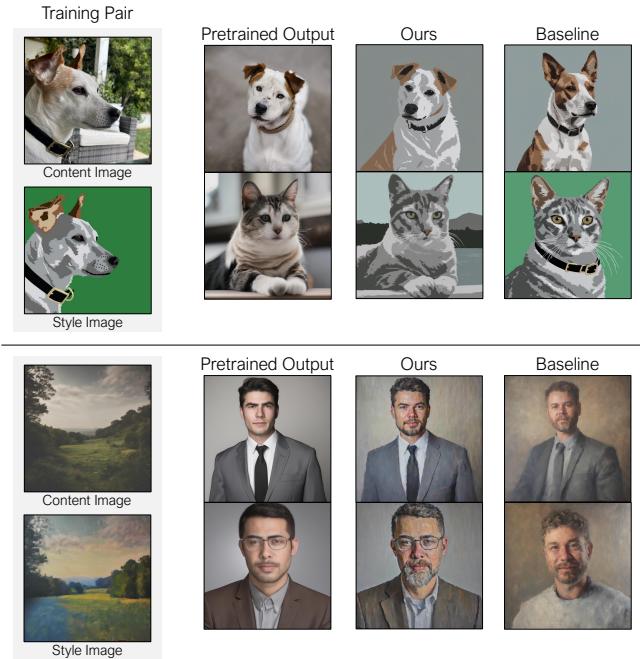


Fig. 10. **Blending multiple style guidances.** We can compose multiple customized models by directly blending each style guidance together. Adjusting the blending strength of each model allows us to acquire a smooth style transition. Each stylized image corresponds to different style guidance strengths. Train Image Pair 0 style image credits: Jack Parkhouse

seamlessly blend the two styles at varying strengths while still preserving the content.

5 Discussion and Limitations

In this work, we have introduced a new task: customizing a text-to-image model with a single image pair. To address this task, we have developed a customization method that explicitly disentangles style and content through both training objectives and a separated parameter space. Our method enables us to grasp the style concept without memorizing the content of input examples. While our approach outperforms existing customization methods, it still exhibits several limitations, as discussed below.

Fig. 11. **Limitations.** Top: our method can cause structure changes in some instances, like change of body position or background changes. Bottom: Our method can change the content in some cases from pretrained output, like the addition of facial hair. We display Baseline DB LoRA for comparison.

Limitations. First, our method may occasionally fail to completely maintain input structure, as demonstrated in Figure 11. This could occur as background/pose change (Top), or as additional features being added (Bottom).

Second, our current method relies on test-time optimization, which takes around 15 minutes on a single A5000 GPU. This can be computationally demanding if we need to process many image styles. Leveraging encoder-based approaches [Arar et al. 2023; Ruiz et al. 2023b] for predicting style and content weights in a feed-forward manner could potentially speed up the customization process.

Lastly, our current method relies on roughly aligned image pairs. Learning styles from unaligned images that depict similar content is an interesting problem for future work.

Acknowledgments

We would like to thank Ali Jahanian, Gaurav Parmar, Ruihan Gao, Sean Liu, and Or Patashnik for their insightful feedback and input that contributed to the finished work. We also thank Jack Parkhouse and Aaron Hertzmann for providing style images. Maxwell Jones is supported by the Siebel Scholar fellowship and the Rales fellowship. This project is partly supported by the Amazon Faculty Research Award, NSF IIS-2239076, Open Philanthropy, and the Packard Fellowship.

References

Yuval Alaluf, Elad Richardson, Gal Metzer, and Daniel Cohen-Or. 2023. A Neural Space-Time Representation for Text-to-Image Personalization. In *SIGGRAPH Asia*.

Moab Arar, Rinon Gal, Yuval Atzmon, Gal Chechik, Daniel Cohen-Or, Ariel Shamir, and Amit H. Bermano. 2023. Domain-agnostic tuning-encoder for fast personalization of text-to-image models. In *SIGGRAPH Asia 2023 Conference Papers*. 1–10.

Omri Avrahami, Kfir Aberman, Ohad Fried, Daniel Cohen-Or, and Dani Lischinski. 2023. Break-a-scene: Extracting multiple concepts from a single image. In *SIGGRAPH Asia 2023 Conference Papers*. 1–12.

Yogesh Balaji, Seungjoo Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. 2022. ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324* (2022).

Amir Bar, Yossi Gandalman, Trevor Darrell, Amir Globerson, and Alexei A. Efros. 2022. Visual Prompting via Image Inpainting. In *Conference on Neural Information Processing Systems (NeurIPS)*.

David Bau, Steven Liu, Tongzhou Wang, Jun-Yan Zhu, and Antonio Torralba. 2020. Rewriting a deep generative model. In *European Conference on Computer Vision (ECCV)*.

Manuel Brack, Felix Friedrich, Katharina Kormmeier, Linoy Tsaban, Patrick Schramowski, Kristian Kersting, and Apolinário Passos. 2023. EDIT5++: Lim-itless Image Editing using Text-to-Image Models. *arXiv preprint arXiv:2311.16711* (2023).

Tim Brooks, Aleksander Holynski, and Alexei A Efros. 2023. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 18392–18402.

Minwoo Byeon, Beomhee Park, Haecheon Kim, Sungjun Lee, Woonhyuk Baek, and Saehoon Kim. 2022. COYO-700M: Image-Text Pair Dataset. <https://github.com/kakaobrain/coyo-dataset>.

Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. 2021. Emerging properties in self-supervised vision transformers. In *IEEE International Conference on Computer Vision (ICCV)*.

Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, José Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T Freeman, Michael Rubinstein, et al. 2023. Muse: Text-to-image generation via masked generative transformers. In *Proceedings of the 40th International Conference on Machine Learning*. 4055–4075.

Tian Qi Chen and Mark Schmidt. 2016. Fast patch-based style transfer of arbitrary style. *arXiv preprint arXiv:1612.04337* (2016).

Wenhu Chen, Hexiang Hu, Yandong Li, Nataniel Rui, Xuhui Jia, Ming-Wei Chang, and William W Cohen. 2023a. Subject-driven text-to-image generation via apprenticeship learning. In *Conference on Neural Information Processing Systems (NeurIPS)*.

Xi Chen, Lianghua Huang, Yu Liu, Yujun Shen, Deli Zhao, and Hengshuang Zhao. 2023b. Anydoor: Zero-shot object-level image customization. *arXiv preprint arXiv:2307.09481* (2023).

Giannis Daras and Alexandros G Dimakis. 2022. Multiresolution textual inversion. *arXiv preprint arXiv:2211.17115* (2022).

Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. 2017. Density estimation using real nvp. In *International Conference on Learning Representations (ICLR)*.

Stephanie Fu, Netanel Tamir, Shobhit Sundaram, Lucy Chai, Richard Zhang, Tali Dekel, and Phillip Isola. 2023. DreamSim: Learning New Dimensions of Human Visual Similarity using Synthetic Data. *arXiv preprint arXiv:2306.09344* (2023).

Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit Haim Bermano, Gal Chechik, and Daniel Cohen-Or. 2022a. An Image is Worth One Word: Personalizing Text-to-Image Generation using Textual Inversion. In *The Eleventh International Conference on Learning Representations*.

Rinon Gal, Moab Arar, Yuval Atzmon, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. 2023a. Designing an encoder for fast personalization of text-to-image models. *ACM Transactions on Graphics (TOG)* (2023).

Rinon Gal, Moab Arar, Yuval Atzmon, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. 2023b. Encoder-based domain tuning for fast personalization of text-to-image models. *ACM Transactions on Graphics (TOG)* 42, 4 (2023), 1–13.

Rinon Gal, Or Patashnik, Haggai Maron, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. 2022b. Stylegan-nada: Clip-guided domain adaptation of image generators. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–13.

Rohit Gandikota, Joanna Materzynska, Tingrui Zhou, Antonio Torralba, and David Bau. 2023. Concept sliders: Lora adaptors for precise control in diffusion models. *arXiv preprint arXiv:2311.12092* (2023).

Daniel Garibi, Or Patashnik, Andrey Voynov, Hadar Averbuch-Elor, and Daniel Cohen-Or. 2024. ReNoise: Real Image Inversion Through Iterative Noising. *arXiv preprint arXiv:2403.14602* (2024).

Leon A Gatys, Alexander S Ecker, and Matthias Bethge. 2015. A neural algorithm of artistic style. *arXiv preprint arXiv:1508.06576* (2015).

Aaron Gokaslan, A Feder Cooper, Jasmine Collins, Landan Seguin, Austin Jacobson, Mihir Patel, Jonathan Frankle, Cory Stephenson, and Volodymyr Kuleshov. 2023. CommonCanvas: An Open Diffusion Model Trained with Creative-Commons Images. *arXiv preprint arXiv:2310.16825* (2023).

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.

Yuchao Gu, Xintao Wang, Jay Zhangjie Wu, Yujun Shi, Yunpeng Chen, Zihan Fan, Wuyou Xiao, Rui Zhao, Shuning Chang, Weijia Wu, et al. 2024. Mix-of-show: Decentralized low-rank adaptation for multi-concept customization of diffusion models. *Advances in Neural Information Processing Systems* 36 (2024).

Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. 2023. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv preprint arXiv:2307.04725* (2023).

Ligong Han, Yinxiao Li, Han Zhang, Peyman Milanfar, Dimitris Metaxas, and Feng Yang. 2023. Svdiff: Compact parameter space for diffusion fine-tuning. In *IEEE International Conference on Computer Vision (ICCV)*.

Amir Hertz, Andrey Voynov, Shlomi Fruchter, and Daniel Cohen-Or. 2023. Style aligned image generation via shared attention. *arXiv preprint arXiv:2312.02133* (2023).

Aaron Hertzmann, Charles E. Jacobs, Nuria Oliver, Brian Curless, and David H. Salesin. 2001. Image Analogies.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. In *Conference on Neural Information Processing Systems (NeurIPS)*.

Jonathan Ho and Tim Salimans. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598* (2022).

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for NLP. In *International Conference on Machine Learning*. PMLR, 2790–2799.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations (ICLR)*.

Xun Huang and Serge Belongie. 2017. Arbitrary style transfer in real-time with adaptive instance normalization. In *IEEE International Conference on Computer Vision (ICCV)*.

Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishala Shankar, Hongseok Namkoong, John Miller, Hannaneh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. 2021. OpenCLIP. <https://doi.org/10.5281/zenodo.5143773> If you use this software, please cite it as below..

Revital Irony, Daniel Cohen-Or, and Dani Lischinski. 2005. Colorization by Example. In *Eurographics Conference on Rendering Techniques*.

Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. 2023. Scaling up GANs for Text-to-Image Synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. 2020a. Training generative adversarial networks with limited data. In *Conference on Neural Information Processing Systems (NeurIPS)*.

Tero Karras, Samuli Laine, and Timo Aila. 2019. A style-based generator architecture for generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020b. Analyzing and improving the image quality of stylegan. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *International Conference on Learning Representations (ICLR)*.

Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. 2023. Multi-Concept Customization of Text-to-Image Diffusion. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, 1931–1941.

Dongxi Li, Junnan Li, and Steven Hoi. 2024. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. *Advances in Neural Information Processing Systems* 36 (2024).

Yuheng Li, Haotian Liu, Qinyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li, and Yong Jae Lee. 2023. Cligen: Open-set grounded text-to-image generation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Yanghao Li, Naixian Wang, Jiaying Liu, and Xiaodi Hou. 2017. Demystifying neural style transfer. *arXiv preprint arXiv:1701.01036* (2017).

Jing Liao, Yuan Yao, Lu Yuan, Gang Hua, and Sing Bing Kang. 2017. Visual Attribute Transfer through Deep Image Analogy. *ACM Trans. Graph.* 36, 4, Article 120 (Jul 2017), 15 pages.

Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. 2022b. Pseudo numerical methods for diffusion models on manifolds. In *International Conference on Learning Representations (ICLR)*.

Nan Liu, Shuang Li, Yilun Du, Antonio Torralba, and Joshua B Tenenbaum. 2022a. Compositional visual generation with composable diffusion models. In *European Conference on Computer Vision*. Springer, 423–439.

Zhiheng Liu, Ruili Feng, Kai Zhu, Yifei Zhang, Kecheng Zheng, Yu Liu, Deli Zhao, Jingren Zhou, and Yang Cao. 2023. Cones: Concept neurons in diffusion models for customized generation. In *International Conference on Machine Learning (ICML)*.

Ilya Loshchilov and Frank Hutter. 2018. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.

Simian Luo, Yiqin Tan, Longbo Huang, Jian Li, and Hang Zhao. 2023. Latent Consistency Models: Synthesizing High-Resolution Images with Few-Step Inference. *arXiv preprint arXiv:2310.04378* (2023).

Jian Ma, Junhao Liang, Chen Chen, and Haonan Lu. 2023. Subject-diffusion: Open domain personalized text-to-image generation without test-time fine-tuning. *arXiv preprint arXiv:2307.11410* (2023).

Csaba Marcus. 2019. How Six Different Artists Have Re-Interpreted Da Vinci’s ‘Mona Lisa’. <https://www.parkwestgallery.com/six-different-artists-da-vinci-mona-lisa/>.

Joanna Materzynska, Josef Sivic, Eli Shechtman, Antonio Torralba, Richard Zhang, and Bryan Russell. 2023. Customizing motion in text-to-video diffusion models. *arXiv preprint arXiv:2312.04966* (2023).

Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. 2022. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations (ICLR)*.

Chong Mou, Xiantao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, and Ying Shan. 2024. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 4296–4304.

Yotam Nitzan, Kfir Aberman, Qiurui He, Orly Liba, Michal Yarom, Yossi Gandalzman, Inbar Mosseri, Yael Pritch, and Daniel Cohen-Or. 2022. MyStyle: A Personalized Generative Prior. In *SIGGRAPH ASIA*.

Yotam Nitzan, Michaël Gharbi, Richard Zhang, Taesung Park, Jun-Yan Zhu, Daniel Cohen-Or, and Eli Shechtman. 2023. Domain Expansion of Image Generators. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. 2020. Contrastive learning for unpaired image-to-image translation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX*. Springer, 319–345.

Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. 2019. Semantic image synthesis with spatially-adaptive normalization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

William Peebles and Saining Xie. 2023. Scalable diffusion models with transformers. In *IEEE International Conference on Computer Vision (ICCV)*.

The Phillips Collection. 2013. Van Gogh Repetitions. <https://www.phillipscollection.org/event/2013-10-11-van-gogh-repetitions>.

Ryan Po, Guandao Yang, Kfir Aberman, and Gordon Wetzstein. 2023. Orthogonal Adaptation for Modular Customization of Diffusion Models. *arXiv preprint arXiv:2312.02432* (2023).

Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. 2023. Sdxt: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952* (2023).

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*.

Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* (2022).

Scott Reed, Yi Zhang, Yuting Zhang, and Honglak Lee. 2015. Deep Visual Analogy-Making. In *Conference on Neural Information Processing Systems (NeurIPS)*.

Yixuan Ren, Yang Zhou, Jimei Yang, Jing Shi, Difan Liu, Feng Liu, Mingi Kwon, and Abhinav Shrivastava. 2024. Customize-a-video: One-shot motion customization of text-to-video diffusion models. *arXiv preprint arXiv:2402.14780* (2024).

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*. Springer, 234–241.

Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2023a. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 22500–22510.

Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Wei Wei, Tingbo Hou, Yael Pritch, Neal Wadhwa, Michael Rubinstein, and Kfir Aberman. 2023b. HyperDreamBooth: HyperNetworks for Fast Personalization of Text-to-Image Models. *arXiv:2307.06949* [cs.CV].

Simo Ryu. 2023a. LoRA-Stable Diffusion. <https://github.com/cloneofsimo/lora>.

Simo Ryu. 2023b. Low-rank Adaptation for Fast Text-to-Image Diffusion Fine-tuning. <https://github.com/cloneofsimo/lora>.

Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. 2022a. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*. 1–10.

Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, Sara Mahdavi, Rapha Gontijo Lopes, et al. 2022b. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. In *NeurIPS*.

Axel Sauer, Tero Karras, Samuli Laine, Andreas Geiger, and Timo Aila. 2023. Stylegan-t: Unlocking the power of gans for fast large-scale text-to-image synthesis. In *International conference on machine learning*. PMLR, 30105–30118.

Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. 2021. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114* (2021).

Viraj Shah, Nataniel Ruiz, Forrester Cole, Erika Lu, Svetlana Lazebnik, Yuanzhen Li, and Varun Jampani. 2023. ZipLoRA: Any Subject in Any Style by Effectively Merging LoRAs. *arXiv preprint arXiv:2311.13600* (2023).

Jing Shi, Wei Xiong, Zhe Lin, and Hyun Joon Jung. 2023. Instantbooth: Personalized text-to-image generation without test-time finetuning. *arXiv preprint arXiv:2304.03411* (2023).

Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning (ICML)*.

Kihyuk Sohn, Nataniel Ruiz, Kimin Lee, Daniel Castro Chin, Irina Blok, Huiwen Chang, Jarred Barber, Lu Jiang, Glenn Ernis, Yuanzhen Li, et al. 2023. StyleDrop: Text-to-Image Generation in Any Style. *arXiv preprint arXiv:2306.00983* (2023).

Jiaming Song, Chenlin Meng, and Stefano Ermon. 2021a. Denoising diffusion implicit models. In *International Conference on Learning Representations (ICLR)*.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2021b. Score-based generative modeling through stochastic differential equations. In *ICLR*.

Joshua Tenenbaum and William Freeman. 1996. Separating style and content. *Advances in neural information processing systems* 9 (1996).

Yoad Tewel, Rinon Gal, Gal Chechik, and Yuval Atzmon. 2023. Key-Locked Rank One Editing for Text-to-Image Personalization. *ACM Transactions on Graphics (TOG)* (2023).

Paul Upchurch, Noah Snavely, and Kavita Bala. 2016. From a to z: supervised transfer of style and content using deep neural network generators. *arXiv preprint arXiv:1603.02003* (2016).

Dani Valevski, Danny Lumen, Yossi Matias, and Yaniv Leviathan. 2023. Face0: Instantaneously conditioning a text-to-image model on a face. In *SIGGRAPH Asia 2023 Conference Papers*. 1–10.

Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. 2016. Conditional image generation with pixelcnn decoders. In *Conference on Neural Information Processing Systems (NeurIPS)*.

Andrey Voynov, Qinghao Chu, Daniel Cohen-Or, and Kfir Aberman. 2023. P+: Extended Textual Conditioning in Text-to-Image Generation. *arXiv preprint arXiv:2303.09522* (2023).

Sheng-Yu Wang, David Bau, and Jun-Yan Zhu. 2021. Sketch your own gan. In *IEEE International Conference on Computer Vision (ICCV)*.

Sheng-Yu Wang, David Bau, and Jun-Yan Zhu. 2022. Rewriting Geometric Rules of a GAN. *ACM SIGGRAPH* (2022).

Xinlong Wang, Wen Wang, Yue Cao, Chunhua Shen, and Tiejun Huang. 2023. Images speak in images: A generalist painter for in-context visual learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Xinrui Wang and Jinze Yu. 2020. Learning to cartoonize using white-box cartoon representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 8090–8099.

Yuxiang Wei, Yubo Zhang, Zhilong Ji, Jinfeng Bai, Lei Zhang, and Wangmeng Zuo. 2023. Elite: Encoding visual concepts into textual embeddings for customized text-to-image generation. In *IEEE International Conference on Computer Vision (ICCV)*.

Hu Ye, Jun Zhang, Sibo Liu, Xiao Han, and Wei Yang. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721* (2023).

Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. 2022. Scaling Autoregressive Models for Content-Rich Text-to-Image Generation. *Transactions on Machine Learning Research* (2022).

Lvmin Zhang and Maneesh Agrawala. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. In *IEEE International Conference on Computer Vision (ICCV)*.

Zhixing Zhang, Ligong Han, Arnab Ghosh, Dimitris N Metaxas, and Jian Ren. 2023. Sine: Single image editing with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6027–6037.

Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. 2020. Differentiable Augmentation for Data-Efficient GAN Training. In *Conference on Neural Information Processing Systems (NeurIPS)*.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*.

Conference on Computer Vision (ICCV).

Zhengxia Zou, Tianyang Shi, Shuang Qiu, Yi Yuan, and Zhenwei Shi. 2021. Stylized neural painting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15689–15698.