

# FreeCustom: Tuning-Free Customized Image Generation for Multi-Concept Composition

Gangui Ding, Canyu Zhao, Wen Wang, Zhen Yang, Zide Liu, Hao Chen, Chunhua Shen  
Zhejiang University, China

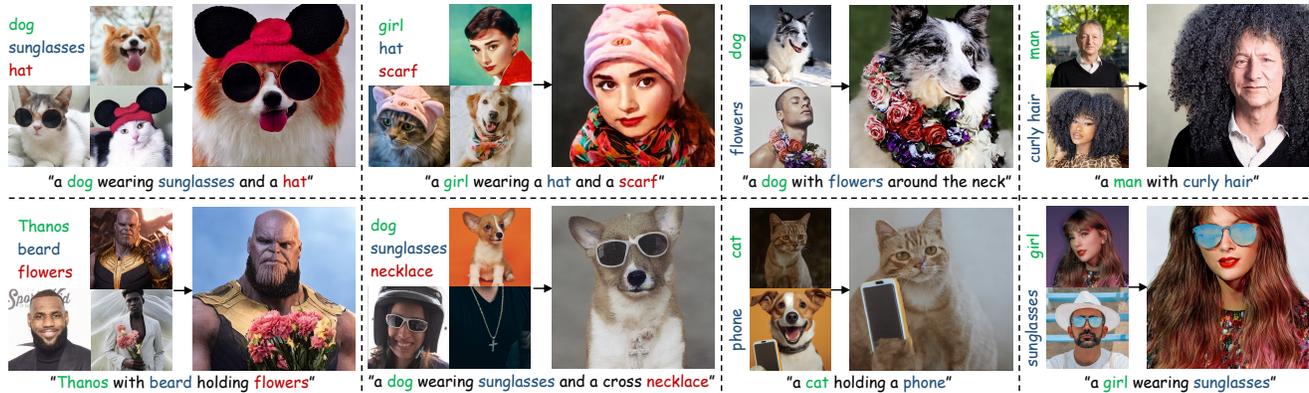


Figure 1. **Results of customized multi-concept composition.** Our method excels at *rapidly* generating high-quality images with multiple concept combinations, without any model parameter tuning. The identity of each concept is remarkably preserved. Furthermore, our method exhibits great versatility and robustness when dealing with different categories of concepts. This versatility allows users to generate customized images that involve diverse combinations of concepts, catering to their specific needs and preferences. Best viewed on screen.

## Abstract

Benefiting from large-scale pre-trained text-to-image (T2I) generative models, impressive progress has been achieved in customized image generation, which aims to generate user-specified concepts. Existing approaches have extensively focused on single-concept customization and still encounter challenges when it comes to complex scenarios that involve combining multiple concepts. These approaches often require retraining/fine-tuning using a few images, leading to time-consuming training processes and impeding their swift implementation. Furthermore, the reliance on multiple images to represent a singular concept increases the difficulty of customization.

To this end, we propose **FreeCustom**, a novel tuning-free method to generate customized images of multi-concept composition based on reference concepts, using only one image per concept as input. Specifically, we introduce a new multi-reference self-attention (MRSa) mechanism and a weighted mask strategy that enables the generated image to access and focus more on the reference concepts. In addition, MRSa leverages our key finding that input concepts

\*GD, CZ, and WW contributed equally. Correspondence should be addressed to HC and CS.

are better preserved when providing images with context interactions. Experiments show that our method’s produced images are consistent with the given concepts and better aligned with the input text. Our method outperforms or performs on par with other training-based methods in terms of multi-concept composition and single-concept customization, but is simpler. Codes can be found [here](#).

## 1. Introduction

With the joint scaling of data, computational resources, and model size, large-scale pre-trained diffusion models [9, 14, 26, 31, 34] have made unprecedented progress in text-to-image generation. Benefiting from the capacity of pre-trained models, customized generation, *i.e.*, generating user-specified objects, becomes possible and has achieved increasing attention due to its wide applications, such as advertisement production, virtual try-on, and art creation.

Existing methods for customization, such as DreamBooth [32], Textual Inversion [10], and BLIP Diffusion [19] have shown significant progress in single-concept customization [1], however, they encounter difficulties when dealing with more complex scenarios involving multiple concepts [18, 35]. In such cases, these methods are prone

to overfitting and demonstrate poor performance in maintaining image naturalness and preserving subject identities when combining multiple concepts. Additionally, as shown in Fig. 2, both approaches necessitate time-consuming training processes, either requiring fine-tuning using 3-5 images or retraining on large-scale datasets, further complicating their practical applicability.

To address these limitations, we propose FreeCustom, a tuning-free method for customized image generation that allows multi-concept composition using only one image per concept as input. Specifically, we employ a dual-path architecture to extract and combine features of input multiple concepts. Then, we introduce a novel multi-reference self-attention (MRSA) mechanism, which extends the original self-attention to access and query features of reference concepts. To highlight the input concepts and eliminate irrelevant information, we implement a weighted mask strategy that directs MRSA to focus more on the given concepts. In addition, we find that the context interaction of the input concept is of great importance for multi-concept composition, and the MRSA mechanism effectively captures the global context of the input images. Consequently, our proposed method can rapidly generate high-fidelity images that align precisely with the text and maintain consistency with reference concepts, without the need for training.

Extensive experiments have shown that our method achieves comparable results to other methods in terms of single-concept customization and exhibits noteworthy advantages when combining multiple concepts. It robustly and effectively generates high-quality images across diverse concepts, as shown in Figs. 1 and 3. Furthermore, it can be easily applied to other diffusion-based models.

To summarize, our contributions are listed as follows:

- We present FreeCustom, a novel tuning-free method that consistently delivers high-quality results for single-concept customization and multi-concept composition.
- We propose the MRSA mechanism and a weighted mask strategy, allowing the generated image to interact with and focus more on the input concepts.
- We pay attention to the significance of the context interaction and leverage it to generate high-fidelity customized images.

## 2. Related Work

The remarkable success of large-scale pre-trained text-to-image (T2I) diffusion models [3, 5, 9, 26, 27, 30, 31] inspires the domain of customized image generation [1, 4, 7, 11, 16, 35]. The goal is to generate new images of a subject of interest, specified by one or a few user-provided images, with varying poses and locations. These methods can be divided into two categories: training-based customization and tailored models for generalizable customization.

**Training-based customization.** In the training-based per-

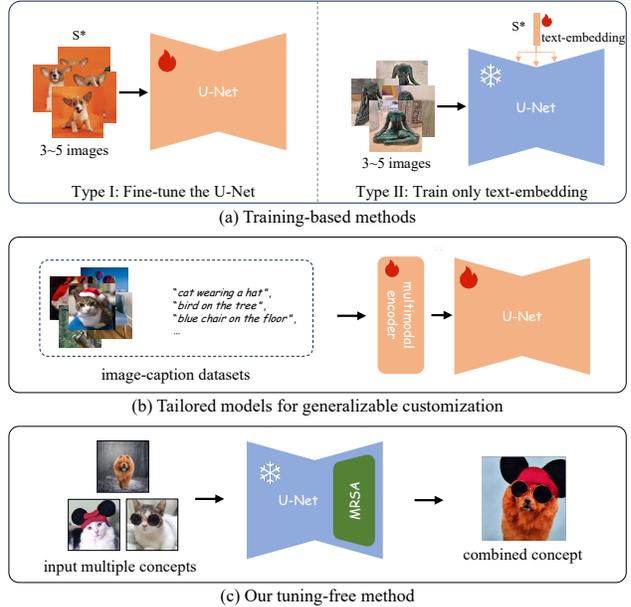


Figure 2. **Paradigm comparison.** Previous methods for customization can be categorized into two main categories: (a) training-based methods and (b) tailored models for generalizable customization. Training-based methods often involve fine-tuning an entire model (Type I) or learning a text embedding to represent a specific subject (Type II). Tailored models typically require re-training on large-scale image datasets to establish a versatile foundation. Unlike these two types of methods, our approach can directly generate customized images of multi-concept combinations without any additional training.

sonalization method, a pioneer work DreamBooth [32] fine-tunes the T2I diffusion model to bind a unique identifier with the subject of interest. It proposes a prior preservation loss to alleviate overfitting during few-shot tuning. A current work Textual Inversion [10] finds that a subject can be represented by a simple learnable text embedding. However, the capacity of a single learnable text embedding is limited, thus subsequent XTI [37] and NeTI [1] introduce layer-wise learnable embedding or implicit time-aware representation, to achieve better performance.

**Tailored models for generalizable customization.** Another line of works [6, 11, 19, 22, 33, 40] attempt to realize faster customization, they train a multimodal encoder and a text-to-image model on dataset-scale images, leading to tailored models for customized generation. In this way, they greatly reduce the number of fine-tuning steps required for customization. It only takes dozens to a hundred steps to achieve comparable outcomes as other methods, which require thousands of fine-tuning steps. For example, BLIP-Diffusion [19] is trained on a subset of OpenImage V6 [24, 25], and generates customized images from BLIP-2 [20] encoded subject representation. BLIP-Diffusion even



Figure 3. Results of single-concept customization.

supports customized images generation in a zero-shot manner, but the effect is diminished. Similarly, HyperDreamBooth [33] is trained on a massive dataset of the low-rank (LoRA) [15] weights for the customized subject and predicts the LoRA weights for the subject of interest during inference. These methods still cannot eliminate the need for fine-tuning, and the requirement for large-scale training further restricts their widespread application. For instance, our method supports direct application to various T2I basic models, However, BLIP-Diffusion requires large-scale re-training for each basic model.

**Multi-concept composition.** Different from these works that perform single object customization, some works [2, 18, 35] focus on multi-subject customization. Custom Diffusion [18] realizes it with a closed-form constrained optimization. Perfusion [35] learns text tokens and modifies the cross-attention of the T2I model, which supports multi-concept composition. Cones series [22, 23] find that a small cluster of neurons in the pre-trained diffusion model is connected to a subject, and further introduce layout guidance to achieve multi-subject customization. Similarly, Mix-of-show [12] introduces gradient fusion to merge several single concepts and alleviate concept conflicts among objects with regional sampling.

While progress has been made, these methods rely on a significant amount of time for few-shot or even dataset-scale training, before sampling for customized generation. What’s more, they struggle with identity preservation in the one-shot customization scenario. By contrast, we propose a completely tuning-free method for customized generation, while achieving superior performance with fast sample speed. The difference between these methods and our proposed method is shown in Fig. 2, more details can be seen in the Appendix.

### 3. Preliminaries

**Latent Diffusion Models.** Diffusion models [9, 14, 26, 34] employ the diffusion process within the image space, while Latent Diffusion Models (LDMs) [31] excel by performing it within the latent space of a pre-trained Variational Au-

toencoder (VAE).

LDMs include a forward and a reverse process. The forward process adds a degree of Gaussian noise to the sample  $\mathbf{x}_0$  to directly obtain  $\mathbf{x}_t$ :  $q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1 - \bar{\alpha}_t)\mathbf{I})$ , where  $\mathbf{x}_t$  is the noised sample at time step  $t$ ,  $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ ,  $\alpha_i$  represents the predefined noise schedule. While the reverse process is a denoising process, which iteratively removes the added noise to the  $\mathbf{x}_t$  until we obtain  $\mathbf{x}_0$ :  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \sigma_t^2)$ , where  $\mu_\theta(\mathbf{x}_t, t) = \frac{1}{\sqrt{\alpha_t}}(\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}}\epsilon_\theta(\mathbf{x}_t, t))$ ,  $\sigma_t = \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t}\beta_t$ ,  $\beta_t = 1 - \alpha_t$ ,  $\epsilon_\theta$  represents a neural network that is trained to predict the actual added noise  $\epsilon$ . The primary objective of the optimization process for training  $\epsilon_\theta$  is  $L = E_{\mathbf{x}_0, \epsilon \sim N(0, \mathbf{I})} \|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|_2^2$ . Moreover, it is possible to condition the text prompt  $P$  on the  $\epsilon_\theta$ . Thus,  $\epsilon_\theta$  can be formulated as  $\epsilon_\theta(\mathbf{x}_t, t, P)$ . By doing so, we can establish a text-to-image (T2I) diffusion model capable of generating images based on the provided prompt.

In this study, we adopt the Stable Diffusion (SD) model as the base model. The SD model belongs to LDMs and is a state-of-the-art (SOTA) T2I generation model based on U-Net architecture, which can generate high-quality images and show significant consistency with the given text prompt.

**Vanilla self-attention in Stable Diffusion.** The original SD U-Net  $\epsilon_\theta$  includes an encoder and a decoder, comprising a total of 7 basic blocks and 16 layers. Each layer contains a residual block, a self-attention module, and a cross-attention module. The self-attention (SA) in each layer is responsible for generating the layout and specific details of the content. At the same time, cross-attention utilizes text embedding to guide the model’s generation process. The output of self-attention is formulated as

$$\text{SA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right)\mathbf{V}, \quad (1)$$

where  $\mathbf{Q}$ ,  $\mathbf{K}$ , and  $\mathbf{V}$  are the query, key, and value features projected from self-attention layers with different projection matrices.

Previous work such as P2P [13], PnP [36], MasaCtrl [4], and Tune-A-Video [41] find the overall layout of the generated image is dominated by the query features, while the key and value features control the semantic contents, and the self-attention module in the U-Net architecture inherently supports a plug-and-play feature injection approach. Therefore, we employ a dual-path paradigm and modify the original self-attention into a meticulously crafted MRSA mechanism to obtain the input concepts’ features.

### 4. Method

We aim to generate customized images with multi-concept composition in a tuning-free manner. We employ the MRSA mechanism to integrate features of each reference concepts (Sec. 4.1), then implement a weighted mask

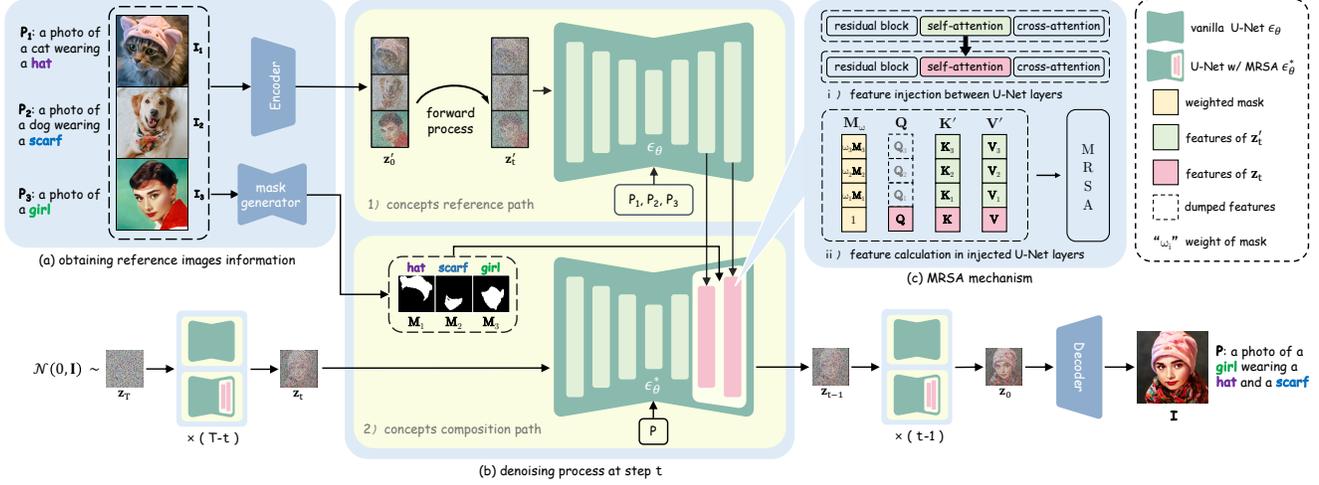


Figure 4. **Overview of the pipeline.** Given a set of reference images  $\mathcal{I} = \{I_1, I_2, I_3\}$  and their corresponding prompts  $\mathcal{P} = \{P_1, P_2, P_3\}$ , we generate a multi-concept customized composition image  $I$  aligned to the target prompt  $P$ . (a) We use a VAE encoder to convert reference images into the latent representation  $\mathbf{z}'_0$  and a segmentation network to extract masks of the concepts. (b) The denoising process involves two paths: 1) the concepts reference path and 2) the concepts composition path. In 1), we employ a diffusion forward process to transform  $\mathbf{z}'_0$  into  $\mathbf{z}'_t$ , subsequently passing  $\mathbf{z}'_t$  to the U-Net  $\epsilon_\theta$ . Notably, the output of  $\epsilon_\theta$  isn't used. In 2), we initially sample  $\mathbf{z}_T \sim \mathcal{N}(0, \mathbf{I})$  and iteratively denoise the latent until we obtain  $\mathbf{z}_0$ . At each time step  $t$ , we directly transmit the current latent  $\mathbf{z}_t$  to the modified U-Net  $\epsilon_\theta^*$  and employ the MRSA to integrate the features from the last two blocks of both the U-Net  $\epsilon_\theta$  and the U-Net  $\epsilon_\theta^*$ . Finally, we utilize a VAE decoder to convert  $\mathbf{z}_0$  into the final image  $I$ . (c) The MRSA mechanism. i) Feature injection happens in the self-attention module between U-Net layers, ii) we apply MRSA using Eq. (4).

strategy to emphasize the reference concept (Sec. 4.2), and selectively replace the original self-attention module with MRSA (Sec. 4.3), and finally find the importance of providing images incorporate context interactions. (Sec. 4.4).

#### 4.1. Multi-Reference Self-Attention

The overall pipeline is a dual-path paradigm as illustrated in Fig. 4. Given a set of  $N$  images containing reference concepts  $\mathcal{I} = \{I_1, I_2, \dots, I_N\}$  and their corresponding prompts  $\mathcal{P} = \{P_1, P_2, \dots, P_N\}$ , our goal is to generate a customized image  $I$  that combines multiple concepts and aligns with the target prompt  $P$ .

**Obtaining reference images information.** As depicted in Fig. 4 (a), we initially obtain information of  $\mathcal{I}$  for further processing. To begin, we employ the VAE encoder, denoted as  $\text{Enc}$ , to transform  $\mathcal{I}$  from the image space into the latent space, resulting in  $\mathbf{z}'_0 = \text{Enc}(\mathcal{I})$ . Subsequently, we utilize an existing segmentation model  $\Phi$  to extract masks  $\mathcal{M} = \{\mathbf{M}_1, \mathbf{M}_2, \dots, \mathbf{M}_N\}$  corresponding to the reference concepts from the input images.

**The concepts reference path and composition path.** Our method as a whole is a diffusion denoising process. We start by randomly sampling the latent  $\mathbf{z}_T$  from a Gaussian distribution  $\mathcal{N}(0, \mathbf{I})$  and gradually denoise  $\mathbf{z}_T$  until  $\mathbf{z}_0$ . The

process consists of two paths: the concepts reference path and the concepts composition path, as shown in Fig. 4 (b).

At each time step  $t$  in the denoising process of the reference path, we initially apply the conventional diffusion forward process to  $\mathbf{z}'_0$ , resulting in  $\mathbf{z}_t$ . Next, we input  $\mathbf{z}'_t$  and  $\mathcal{P}$  into the U-Net  $\epsilon_\theta$  to extract the query, key, and value features denoted as  $\mathcal{Q}_{nlt}$ ,  $\mathcal{K}_{nlt}$ , and  $\mathcal{V}_{nlt}$ , respectively, for the reference image  $n$ , attention layer  $l$ , and time step  $t$ . Note that we do not utilize the output of  $\epsilon_\theta$  at each time step.

In the composition path, we modify the original U-Net  $\epsilon_\theta$  by replacing the self-attention module with the MRSA module, referred to as  $\epsilon_\theta^*$ . At each time step  $t$ , we feed the denoised latent  $\mathbf{z}_t$  and the target prompt  $P$  into  $\epsilon_\theta^*$  and then calculate the query, key, and value features  $\mathcal{Q}_{lt}$ ,  $\mathcal{K}_{lt}$ , and  $\mathcal{V}_{lt}$  for the generated image at attention layer  $l$  and time step  $t$ . The final denoised  $\mathbf{z}_0$  is transformed back to the image space, yielding the ultimate customized image  $I = \text{Dec}(\mathbf{z}_0)$ , where  $\text{Dec}$  is the VAE decoder.

**MSRA mechanism.** As illustrated in Fig. 4 (c), features from the self-attention module in  $\epsilon_\theta^*$  are injected into the module in  $\epsilon_\theta$ , thus, MRSA perceives not only its own features but also the features obtained from the reference images. Specifically, at each attention layer  $l$  and time step  $t$  of the denoising process, the  $N$  injected key and value features from  $\mathbf{z}'_t$  are denoted as  $\mathcal{K}_{lt} = \{\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_N\}$  and  $\mathcal{V}_{lt} = \{\mathbf{V}_1, \mathbf{V}_2, \dots, \mathbf{V}_N\}$ . Meanwhile, the query,

key, and value features from  $\mathbf{z}_t$  are  $\mathbf{Q}_{lt}$ ,  $\mathbf{K}_{lt}$ , and  $\mathbf{V}_{lt}$ . We concatenate  $\mathbf{K}$  and  $\mathbf{V}$  with  $\mathcal{K}$  and  $\mathcal{V}$  to obtain  $\mathbf{K}' = [\mathbf{K}, \mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_N]$  and  $\mathbf{V}' = [\mathbf{V}, \mathbf{V}_1, \mathbf{V}_2, \dots, \mathbf{V}_N]$ . Here, we omit the subscript  $lt$  for brevity since we apply the same operation across selected layers and all time steps. Finally, the operation of MRSA is as follows:

$$\text{MRSA}(\mathbf{Q}, \mathbf{K}', \mathbf{V}') = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}'^T}{\sqrt{d}}\right)\mathbf{V}'. \quad (2)$$

This approach allows generated images to naturally incorporate features from the input images, however, we observe that it also retains unrelated features to a significant extent, resulting in concept confusion. To address this issue, we concatenate all masks  $\mathbf{M}_i \in \mathcal{M}$  with an all-ones matrix  $\mathbf{1}$ , yielding  $\mathbf{M} = [\mathbf{1}, \mathbf{M}_1, \mathbf{M}_2, \dots, \mathbf{M}_N]$  and then employ  $\mathbf{M}$  to constrain the regions attended by the MRSA, effectively masking out unrelated contents and guiding the model’s focus towards the target concept. Technically,

$$\begin{aligned} \text{MRSA}(\mathbf{Q}, \mathbf{K}', \mathbf{V}', \mathbf{M}) \\ = \text{Softmax}\left(\frac{\mathbf{M} \odot (\mathbf{Q}\mathbf{K}'^T)}{\sqrt{d}}\right)\mathbf{V}', \end{aligned} \quad (3)$$

$\mathbf{M}$  is downsampled here to match the resolution of the current feature, and  $\odot$  represents the Hadamard product. This masked MRSA mechanism can utilize the context interaction information within the input images and filter out irrelevant features during the image generation process.

### 4.2. Weighted Mask

Although the introduction of masks helps to tackle the problem of unrelated features, the current model still struggles to accurately preserve the distinctive characteristics of the target concept in the generated results. It can only provide a rough representation of the appearance of the target concept, lacking precise details. To overcome this limitation, we introduce a scaling factor to each mask, denoted as  $\mathbf{w} = \{1, \omega_1, \omega_2, \dots, \omega_N\}$ , to enhance the model’s focus on the target concept in conjunction with the weighted



Figure 5. **Weighted mask strategy.**  $\mathbf{w}$  is the weight of the mask, where the first weight corresponds to the main edited subject, and the following three weights are for the input concepts.

masks, *i.e.*,  $\mathbf{M}_w = [\mathbf{1}, \omega_1\mathbf{M}_1, \omega_2\mathbf{M}_2, \dots, \omega_N\mathbf{M}_N]$ . Then the MRSA is formulated as

$$\begin{aligned} \text{MRSA}(\mathbf{Q}, \mathbf{K}', \mathbf{V}', \mathbf{M}_w) \\ = \text{Softmax}\left(\frac{\mathbf{M}_w \odot (\mathbf{Q}\mathbf{K}'^T)}{\sqrt{d}}\right)\mathbf{V}'. \end{aligned} \quad (4)$$

We have found that the model is not overly sensitive to the value of  $\mathbf{w}$ . Assigning values between 2 and 3 to  $\mathbf{w}$  consistently yields excellent performance across various scenarios. By incorporating the concept-specific weighted masks within MRSA, the model is encouraged to attend more selectively to the desired features and suppress the influence of irrelevant information. This refinement leads to improved generation results that align more closely with the intended concepts as shown in Fig. 5.

### 4.3. Selective MRSA Replacement

Experimental results indicate that a straightforward substitution of the self-attention modules in all 7 basic blocks with the MRSA module results in unnatural generation, loss of conceptual coherence, and textual inconsistencies. Previous research [4, 36] find that the query features in the deep layers of the U-Net possess the capacity for layout control and semantic information acquisition. Therefore, we adopt a similar strategy to MasaCtrl [4], where the self-attention module is replaced with our MRSA module only in a selected set of blocks, particularly in the deeper blocks of the U-Net, denoted as  $\Psi$ . Through empirical observation, we achieve superior results by setting  $\Psi = [5, 6]$ . These results demonstrate our ability to not only improve the natural and photorealistic appearance of the generated image but also effectively preserve the identity of the given concepts, as shown in Fig. 10.

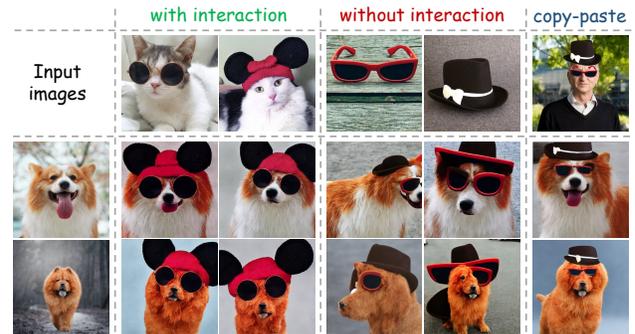


Figure 6. **Results of given concepts with and without context interaction.** The target prompt is “a dog wearing a hat and sunglasses”. The fidelity of the input concepts is effectively maintained when providing contextual images. On the contrary, unnatural results are generated. Manually constructing contexts by a simple copy-paste strategy also improves image quality.

#### 4.4. Preparing Images with Context Interaction

We find the context interaction of each concept in an image is crucial for achieving multi-concept composition, such as providing the context of “wearing a hat”, the hat concept is well preserved. However, the generated image suffers if the input image only shows a plain hat with no indication that it is being worn. Specifically, as can be seen in Fig. 6, when only images of isolated hats and glasses without context interactions are provided, the model struggles to maintain the desired concept features successfully (columns 4 and 5). However, when images of a cat wearing sunglasses and a hat are provided, the model can generate excellent customized results (columns 2 and 3). A simple copy-paste strategy can also provide enough context information to benefit the customization (column 6). Leveraging this strategy, we can manually produce reference images with contextual interactions, helping to generate images that possess high fidelity and text alignment.

### 5. Experiments

**Data.** We curate a diverse dataset comprising images sourced from the Internet and previous studies in customization [1, 10, 32] and editing [13, 36, 39, 42]. This dataset is used to provide reference concepts for customized image generation and encompasses a wide range of categories, including faces, anime characters, animals, accessories, and wearables, with each image representing one or more distinct concepts. We employ the Grounded-Segment-Anything [17, 21] to create masks for each concept, which enables us to extract masks by simply providing the name of the desired concept, streamlining the process of obtaining the concept-specific mask.

**Implementation setting.** We use Stable Diffusion V1.5<sup>1</sup> as our base model to generate high-quality images of 512×512 resolution using an NVIDIA 3090 GPU. The mask weight  $\omega_i$  for the reference concepts is set to 3, while the mask weight for the main edited subject is set to 1 to align with the original self-attention. We assign blocks  $\Psi$  as [5, 6] to replace the original self-attention module in the SD model with MRSA at each time step.

#### 5.1. Comparison with Existing Methods

**Evaluation metrics.** Similar to DreamBooth [32], we access image similarity using DINOv2 [8, 28] and CLIP-I [29], and image-text consistency using CLIP-T [29] and CLIP-T-L [29]. To assess the quality of the generated images, we employ CLIP-IQA [38], an evaluative criterion that considers the visual aesthetics and the abstract perception in an image. Furthermore, we evaluate the overall time required for implementing the method, including

<sup>1</sup><https://huggingface.co/runwayml/stable-diffusion-v1-5>

both the initial preprocessing time and the subsequent inference time. Finally, we conduct a user study for evaluation in terms of image quality, identity fidelity, and text alignment.

**Compared methods.** To evaluate the effectiveness of our approach in single concept customization, we compare it against several state-of-the-art methods, namely DreamBooth [32], NeTI [1], and BLIP Diffusion [19]. We compare our approach against Custom Diffusion [18] and Perfusion [35] for evaluating multi-concept composition. These competitors represent a range of approaches, including tuning-based models [1, 18, 32, 35] and models pre-trained on large-scale datasets [19], reflecting the diversity of established customization techniques. All comparisons are conducted using Stable Diffusion V1.5 as the common foundation, ensuring a consistent inference setup for fair evaluations among the different methods.

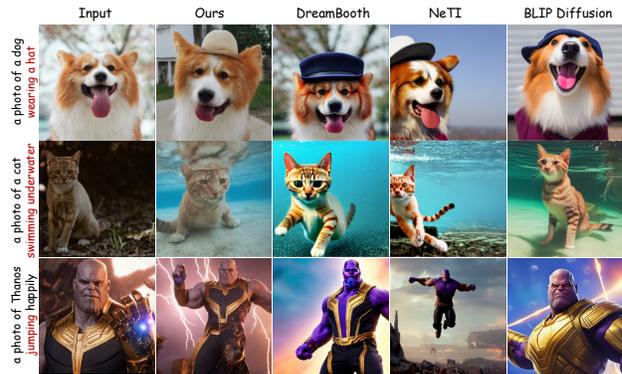


Figure 7. Comparisons of single-concept customization.

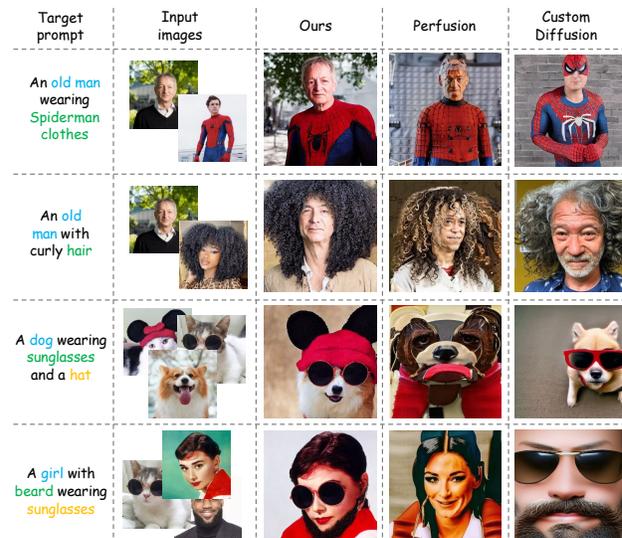


Figure 8. Comparisons of multi-concept composition.

Methods	DINOv2	CLIP-I	CLIP-T	CLIP-T-L	CLIP-IQA
<i>single-concept</i>					
DreamBooth [32]	<b>0.8948</b>	0.8906	27.3825	21.9413	0.7194
NeTI [1]	0.7839	0.8677	<u>29.9023</u>	<u>25.1220</u>	0.7239
BLIP Diffusion [19]	0.8734	<b>0.8975</b>	29.1278	24.0543	<b>0.7655</b>
FreeCustom	0.8376	0.8755	<b>32.0206</b>	<b>27.4440</b>	<u>0.7292</u>
<i>multi-concept</i>					
Custom Diffusion [18]	<u>0.6545</u>	<u>0.2393</u>	<u>29.0702</u>	<u>23.6657</u>	<u>0.8921</u>
Perfusion [35]	0.6399	0.2277	22.1371	16.1719	0.8624
FreeCustom	<b>0.7625</b>	<b>0.2871</b>	<b>33.7826</b>	<b>27.8758</b>	<b>0.9002</b>

Table 1. Image similarity (DINOv2, CLIP-I), image-text alignment (CLIP-T, CLIP-T-L), and image quality (CLIP-IQA).

Methods	Venue	Preprocessing	Inference	Total
<i>single-concept</i>				
DreamBooth [32]	CVPR’23	500s	3s	503s
NeTI [1]	SIGGRAPH’23	420s	30s	450s
BLIP Diffusion [19]	NeurIPS’23	6 days	3s	6 days
FreeCustom	this work	0	20s	<b>20s</b>
<i>multi-concept</i>				
Custom Diffusion [18]	CVPR’23	287s	13s	300s
Perfusion [35]	SIGGRAPH’23	821s	14s	835s
FreeCustom (2 concepts)	this work	0	36s	<b>36s</b>
FreeCustom (3 concepts)	this work	0	58s	<b>58s</b>

Table 2. Comparisons of preprocessing time and inference time. Preprocessing time includes fine-tuning or re-training time.

**Qualitative comparison.** In terms of multi-concept composition, as illustrated in Fig. 8, our method demonstrates superior performance compared to [18] and [35], despite being tuning-free. It efficiently preserves the identity of reference concepts in the generated customized images and performs well across a diverse range of concepts, including hairstyles, accessories, and clothing. Notably, our method exhibits significantly improved consistency between the generated images and the corresponding text prompts compared to the other methods evaluated. We also show the results compared with other methods in single concept customized image generation as depicted in Fig. 7. Our method achieves comparable or even more natural and realistic results, further emphasizing the high-quality image-generation capabilities of our approach.

**Image and text fidelity.** As shown in Tab. 1, our method ranks first in terms of 5 metrics when comparing multi-concept composition, highlighting our method’s superior advantage in multi-concept results. However, when it comes to single-concept customization, our image-text alignment is strong. Nonetheless, our performance in DINOv2 and CLIP-I is mediocre due to the trade-off between image similarity and text consistency [18].

**Time efficiency.** Thanks to our tuning-free paradigm, our method significantly outperforms other methods in terms of time efficiency, as reported in Tab. 2. Fine-tuning-based models require extensive preprocessing time, whereas our method can be efficiently applied to various models without

Methods	Venue	Alignment	Consistency	Quality
<i>single-concept</i>				
DreamBooth [32]	CVPR’23	1.77	1.48	1.53
NeTI [1]	SIGGRAPH’23	2.27	2.49	2.31
BLIP Diffusion [19]	NeurIPS’23	3.50	3.04	2.85
FreeCustom	this work	<b>3.62</b>	<b>3.11</b>	<b>2.89</b>
<i>multi-concept</i>				
Custom Diffusion [18]	CVPR’23	1.91	2.53	2.48
Perfusion [35]	SIGGRAPH’23	1.50	1.70	1.73
FreeCustom	this work	<b>4.40</b>	<b>4.65</b>	<b>4.17</b>

Table 3. **User study results.** The text-to-image correspondence (Alignment), identity preservation (Consistency) and image quality (Quality) are evaluated. Our approach demonstrates superior performance across all these aspects.

the need for additional training. On the other hand, BLIP Diffusion requires retraining when applied to different models, complicating its preprocessing.

**User study.** We collect a total of 23 questionnaires for single-concept and 42 questionnaires for multi-concept. Specifically, we generated results using different methods with identical prompts, base models, and image inputs. For single-concept, each method customizes 5 concepts, for a total of 20 questions. For multi-concept, each method generates 10 distinct sets of concept combinations, ranging from 2 to 4 concepts, resulting in a total of 30 questions. The results are randomly placed. Each participant is required to assess customized images based on three criteria: alignment between image and prompt, consistency of concepts between the given image and reference images, and overall image quality. Each question is rated on a 5-point scale, with 5 representing the best. The results unequivocally demonstrate the superiority of our method in both tasks as presented in Tab. 3.

## 5.2. Ablation Studies

**Weighted mask strategy.** To assess the significance of mask weights, we conduct ablation studies as depicted in Fig. 5. Insufficient weight assigned to a concept results in poor preservation of its identity. Additionally, we visual-

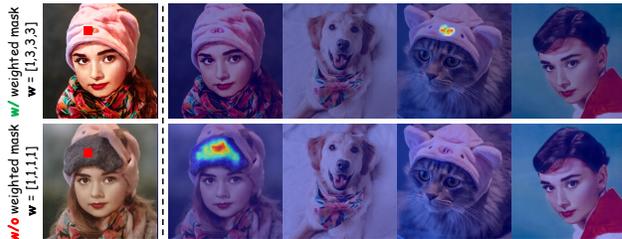


Figure 9. **Visualization of multi-attention map** in MRSA mechanism. We visualize the multi-attention map corresponding to the query feature indicated by the red box on the generated image.

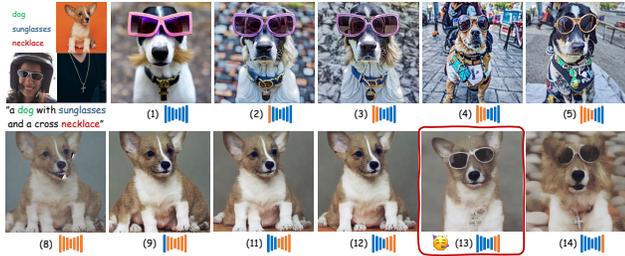


Figure 10. **Selective applying MRSA to basic blocks.** The blue color represents the original basic block and the yellow color indicates the basic block whose self-attention is replaced by MRSA.



Figure 11. **Appearance transfer.** Our method generates objects with similar appearance and materials as the input image.

ize the multi-attention map  $A = \text{Softmax}(\frac{M_w \odot (QK^T)}{\sqrt{d}})$ ,  $A \in \mathbb{R}^{H \times ((N+1)W)}$ ,  $H$  and  $W$  represents the resolution of features in current layer, in MRSA at the 16th layer and the 50th timestep. A weighted mask plays a critical role in enabling our model to effectively capture features from reference images, thereby facilitating accurate feature generation. In the absence of a weighted mask, attention is not adequately directed towards the input concepts. Consequently, the generated image fails to accurately capture the individual concepts portrayed in the reference image, such as the loss of the hat identity.

**Selective MRSA replacement.** Instead of directly replacing the self-attention of all 7 basic blocks of U-Net with MRSA, we conduct experiments using various strategies to selectively replace the self-attention modules in the basic blocks of the U-Net with MRSA, as depicted in Fig. 10. Results indicate that when  $\Psi = [5, 6]$  exhibits the highest efficacy in preserving the identity of intricate concepts, aligning with the provided prompt, and producing photo-realistic images.

### 5.3. More Applications

**Appearance transfer.** Appearance transfer aims to convert the appearance information of the input concept to a new object. As shown in Fig. 11, our method benefits from the ability to query the input image and can easily implement appearance transfer, the generated image effectively retains the appearance information of the input image.

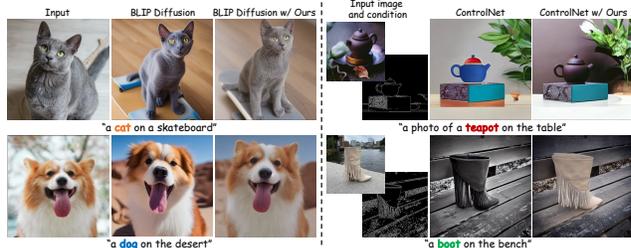


Figure 12. **Combined with existing methods.** (Left) Combined with our methods, BLIP Diffusion generates images more faithful to the input image. (Right) Using our approach, ControlNet yields results consistent with the input condition layout and the appearance of the input image.



Figure 13. **Applying our method to different base models** with the same images and prompt as input. The input concepts including cat, hat, sunglasses, and floral shirt are well preserved.

**Empower other methods.** Our method can enhance ControlNet [43] and BLIP Diffusion [19] in a plug-and-play manner. As illustrated in Fig. 12, by using our method, the output of BLIP diffusion becomes more faithful to the input image and better aligned with the input text. Furthermore, ControlNet can generate results that are consistent in layout and identity when combined with ours.

**Applied to different base models.** Our method exhibits strong flexibility and robustness, as it can directly work on different base models. As illustrated in Fig. 13, we apply our method to various versions of the SD model and different checkpoints on the Civitai platform.

## 6. Conclusion

In this paper, we propose a novel tuning-free method to address the challenges of multi-concept customized composition. Experimental results show that our proposed method enables flexible combinations of different objects from various categories. However, our method currently faces limitations in that it lacks an explicit module to perceive the structure of the input reference concepts. In our future work, we plan to integrate additional techniques such as image adapter, to overcome these limitations.

**Acknowledgement** This work was in part supported by National Key R&D Program of China (No. 2022ZD0118-700).

## References

- [1] Yuval Alaluf, Elad Richardson, Gal Metzer, and Daniel Cohen-Or. A neural space-time representation for text-to-image personalization. *ACM Trans. Graph.*, 42(6), 2023. 1, 2, 6, 7
- [2] Omri Avrahami, Kfir Aberman, Ohad Fried, Daniel Cohen-Or, and Dani Lischinski. Break-a-scene: Extracting multiple concepts from a single image. In *SIGGRAPH Asia 2023 Conference Papers*, New York, NY, USA, 2023. Association for Computing Machinery. 3
- [3] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. ediff-i: Text-to-image diffusion models with an ensemble of expert denoisers. corr, vol. abs/2211.01324 (2022), 2022. 2
- [4] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 22560–22570, 2023. 2, 3, 5
- [5] Kaixuan Chen, Jie Song, Shunyu Liu, Na Yu, Zunlei Feng, Gengshi Han, and Mingli Song. Distribution knowledge embedding for graph pooling. 2022. 2
- [6] Kaixuan Chen, Shunyu Liu, Tongtian Zhu, Ji Qiao, Yun Su, Yingjie Tian, et al. Improving expressivity of gnns with subgraph-specific factor embedded normalization. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2023. 2
- [7] Xi Chen, Lianghua Huang, Yu Liu, Yujun Shen, Deli Zhao, and Hengshuang Zhao. Anydoor: Zero-shot object-level image customization. *arXiv preprint*, 2023. 2
- [8] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers, 2023. 6
- [9] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 1, 2, 3
- [10] 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. 1, 2, 6
- [11] Rinon Gal, Moab Arar, Yuval Atzmon, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. Encoder-based domain tuning for fast personalization of text-to-image models. *ACM Transactions on Graphics (TOG)*, 42(4):1–13, 2023. 2
- [12] Yuchao Gu, Xintao Wang, Jay Zhangjie Wu, Yujun Shi, Yunpeng Chen, Zihan Fan, Wuyou Xiao, Rui Zhao, Shuning Chang, Weijia Wu, et al. Mix-of-show: Decentralized low-rank adaptation for multi-concept customization of diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [13] 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. 3, 6
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 1, 3
- [15] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. 3
- [16] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6007–6017, 2023. 2
- [17] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 6
- [18] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1931–1941, 2023. 1, 3, 6, 7
- [19] Dongxu Li, Junnan Li, and Steven Hoi. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 2, 6, 7, 8
- [20] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023. 2
- [21] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 6
- [22] Zhiheng Liu, Ruili Feng, Kai Zhu, Yifei Zhang, Kecheng Zheng, Yu Liu, Deli Zhao, Jingren Zhou, and Yang Cao. Cones: Concept neurons in diffusion models for customized generation. In *International Conference on Machine Learning*, 2023. 2, 3
- [23] Zhiheng Liu, Yifei Zhang, Yujun Shen, Kecheng Zheng, Kai Zhu, Ruili Feng, Yu Liu, Deli Zhao, Jingren Zhou, and Yang Cao. Cones 2: Customizable image synthesis with multiple subjects. *arXiv preprint arXiv:2305.19327*, 2023. 3
- [24] Shangbang Long, Siyang Qin, Dmitry Pantelev, Alessandro Bissacco, Yasuhisa Fujii, and Michalis Raptis. Towards end-to-end unified scene text detection and layout analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [25] Shangbang Long, Siyang Qin, Dmitry Pantelev, Alessandro Bissacco, Yasuhisa Fujii, and Michalis Raptis. Icdar 2023 competition on hierarchical text detection and recognition. *arXiv preprint arXiv:2305.09750*, 2023. 2
- [26] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International*

- Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. [1](#), [2](#), [3](#)
- [27] Alexander Quinn 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, 17-23 July 2022, Baltimore, Maryland, USA*, pages 16784–16804. PMLR, 2022. [2](#)
- [28] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2023. [6](#)
- [29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [6](#)
- [30] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022. [2](#)
- [31] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. [1](#), [2](#), [3](#)
- [32] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023. [1](#), [2](#), [6](#), [7](#)
- [33] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Wei Wei, Tingbo Hou, Yael Pritch, Neal Wadhwa, Michael Rubinstein, and Kfir Aberman. Hyperdreambooth: Hypernetworks for fast personalization of text-to-image models. *arXiv preprint arXiv:2307.06949*, 2023. [2](#), [3](#)
- [34] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. [1](#), [3](#)
- [35] Yoad Tewel, Rinon Gal, Gal Chechik, and Yuval Atzmon. Key-locked rank one editing for text-to-image personalization. In *ACM SIGGRAPH 2023 Conference Proceedings*, 2023. [1](#), [2](#), [3](#), [6](#), [7](#)
- [36] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1921–1930, 2023. [3](#), [5](#), [6](#)
- [37] Andrey Voynov, Qinghao Chu, Daniel Cohen-Or, and Kfir Aberman.  $p+$ : Extended textual conditioning in text-to-image generation. *arXiv preprint arXiv:2303.09522*, 2023. [2](#)
- [38] Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *AAAI*, 2023. [6](#)
- [39] Wen Wang, kangyang Xie, Zide Liu, Hao Chen, Yue Cao, Xinlong Wang, and Chunhua Shen. Zero-shot video editing using off-the-shelf image diffusion models. *arXiv preprint arXiv:2303.17599*, 2023. [6](#)
- [40] Yuxiang Wei, Yabo Zhang, Zhilong Ji, Jinfeng Bai, Lei Zhang, and Wangmeng Zuo. ELITE: encoding visual concepts into textual embeddings for customized text-to-image generation. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 15897–15907. IEEE, 2023. [2](#)
- [41] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7623–7633, 2023. [3](#)
- [42] Zhen Yang, Ganggui Ding, Wen Wang, Hao Chen, Bohan Zhuang, and Chunhua Shen. Object-aware inversion and re-assembly for image editing. 2023. [6](#)
- [43] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. [8](#)