

# TARA: Token-Aware LoRA for Composable Personalization in Diffusion Models

Yuqi Peng<sup>1,2</sup>, Lingtao Zheng<sup>1</sup>, Yufeng Yang<sup>1</sup>, Yi Huang<sup>3,1</sup>, Mingfu Yan<sup>1</sup>,  
Jianzhuang Liu<sup>1</sup>, Shifeng Chen<sup>1,4\*</sup>

<sup>1</sup>Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences

<sup>2</sup>Northeastern University

<sup>3</sup>vivo AI Lab

<sup>4</sup>Shenzhen University of Advanced Technology

peng.yuq@northeastern.edu, {lt.zheng2, yf.yang2, yi.huang, mf.yan, jz.liu, shifeng.chen}@siat.ac.cn

## Abstract

Personalized text-to-image generation aims to synthesize novel images of a specific subject or style using only a few reference images. Recent methods based on Low-Rank Adaptation (LoRA) enable efficient single-concept customization by injecting lightweight, concept-specific adapters into pre-trained diffusion models. However, combining multiple LoRA modules for multi-concept generation often leads to identity missing and visual feature leakage. In this work, we identify two key issues behind these failures: (1) token-wise interference among different LoRA modules, and (2) spatial misalignment between the attention map of a rare token and its corresponding concept-specific region. To address these issues, we propose Token-Aware LoRA (TARA), which introduces a token mask to explicitly constrain each module to focus on its associated rare token to avoid interference, and a training objective that encourages the spatial attention of a rare token to align with its concept region. Our method enables training-free multi-concept composition by directly injecting multiple independently trained TARA modules at inference time. Experimental results demonstrate that TARA enables efficient multi-concept inference and effectively preserving the visual identity of each concept by avoiding mutual interference between LoRA modules.

## 1 Introduction

Diffusion models achieves remarkable progress in text-to-image generation, offering users high-quality and highly controllable image synthesis. Building on this success, a growing body of research explores how to leverage pre-trained diffusion models for personalized generation. The goal is to inject specific concepts (such as objects or styles) into the model using only 3–5 reference images, enabling the model to generate images of the personalized concept conditioned on user input (e.g., text prompt).

To enable personalized generation, various approaches have been proposed. Methods such as DreamBooth (Ruiz et al. 2023) inject concepts by fine-tuning the entire UNet and bind each concept to a rare token, such as “xlo”, while others like Textual Inversion (Gal et al. 2022) optimize a new token embedding to represent the personalized concept. Some works (Pang et al. 2024; Kumari et al.

\*Corresponding author: Shifeng Chen.

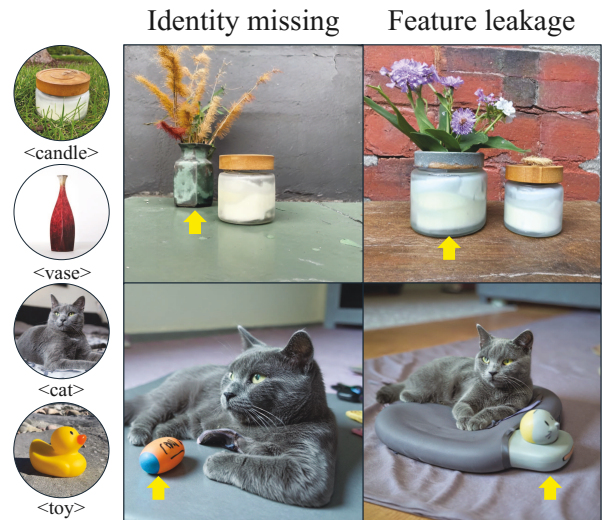


Figure 1: Illustration of two common issues when directly applying multiple LoRA modules for multi-concept generation. With prompts “A  $V_{\text{candle}}^*$  candle next to a  $V_{\text{vase}}^*$  vase” and “A  $V_{\text{cat}}^*$  cat sleeps beside a  $V_{\text{toy}}^*$  toy”, the left column shows *identity missing*, where objects (vase and toy) are present, but their visual identities (e.g., the distinctive shape or color of  $\langle \text{vase} \rangle$  and  $\langle \text{toy} \rangle$ ) are lost. The right column shows *feature leakage*, where visual features of one concept ( $\langle \text{candle} \rangle$  or  $\langle \text{cat} \rangle$ ) incorrectly appear on another object.

2023) further combine embedding learning with model fine-tuning. Among these approaches, Low-Rank Adaptation (LoRA) (Hu et al. 2022) has emerged as a popular and efficient solution for personalization tasks. By inserting trainable low-rank matrices, LoRA adapts pre-trained models without modifying their original parameters. Moreover, each personalized concept can be encapsulated in a compact set of LoRA weights, significantly reducing storage overhead and enabling efficient model sharing across users.

Although LoRA achieves promising results for single-concept customization, it still faces challenges in multi-concept composition, such as identity missing and visual feature leakage. As shown in Figure 1, when generating an image with two concepts, the model may generate a vase and

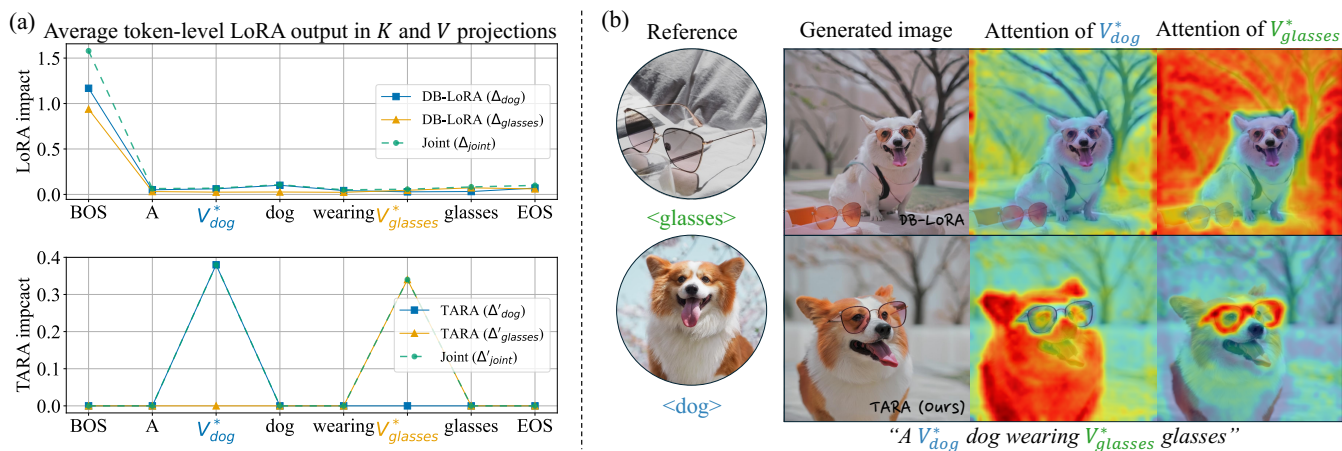


Figure 2: Illustration of the two key issues in multi-concept generation and how TARA addresses them. (a) Visualization of the average magnitude of LoRA output in Key and Value projections of each token in the cross-attention layers. The result on top shows DreamBooth LoRA modules, which all mainly affect the [BOS] token, resulting in overlapping influences and potential conflicts between concepts. In contrast, TARA (bottom) uses a mask to ensure that each module focuses only on its associated rare token, effectively avoiding the interference. (b) Visualization of the spatial attention of each rare token using DAAM (Tang et al. 2022). DreamBooth LoRA (top) causes dispersed and overlapping attention across concepts, while TARA (bottom) tightly aligns each rare token’s attention with the correct spatial region of its corresponding concept, reducing feature entanglement.

a toy, but it cannot preserve their personalized identities and shows visual feature leakage, leading to failed multi-concept generation. While some existing works explore various fusion strategies, they often require additional training (Gu et al. 2023; Shah et al. 2024), additional manual conditions or adapters to control the spatial placement of each concept (Po et al. 2024), or are limited to combining style and object LoRAs only (Frenkel et al. 2024; Ouyang, Li, and Hou 2025). In contrast to these approaches, our goal is to develop a lightweight and practical solution that enables multi-concept generation without any additional training or external conditions. To better understand these challenges, we investigate the behavior of LoRA modules in cross-attention layers during multi-concept generation.

Our approach is motivated by two key insights observed during multi-concept generation using LoRA, as shown in Figure 2. First, although each concept is associated with a rare token in the prompt, we observe that the LoRA module exerts limited influence on the rare token within the cross-attention layers. Instead, their primary impact is often concentrated on the [BOS] (beginning-of-sequence) token. When multiple LoRA modules are injected simultaneously, their shared focus raises the risk of interference, which can lead to concept identity distortion or even missing concepts. Second, we observe that rare tokens tend to distribute their attentions broadly across the image (especially on the background), rather than concentrating on the corresponding target objects. When multiple LoRA modules are injected together, the attention regions associated with different rare tokens significantly overlap with regions of other concepts, resulting in visual feature leakage (e.g., features of concept A incorrectly appear within concept B).

To address these issues, we propose Token-Aware LoRA (TARA), which incorporates *token focus masking* (TFM)

and a *token alignment loss* (TAL) to constrain the LoRA module in the cross-attention layers (specifically the Key and Value matrices). TFM filters out the impact of LoRA on non-target tokens during the forward pass using a binary mask, thereby focusing the module’s effect on its corresponding rare token. Meanwhile, TAL encourages the rare token to align its key representation with that of the class token, leading to better spatial attention localization by leveraging the model’s prior knowledge. Experimental results demonstrate that our method achieves efficient generation of multiple personalized concepts without relying on additional conditions or fusion-specific training, while effectively preserving each concept’s identity. This enables high-fidelity image synthesis based solely on user-provided text prompts.

## 2 Related Work

**Text-to-Image Diffusion Models.** Diffusion models are one of the foundations of modern text-to-image generation. Early methods such as DDPM (Ho, Jain, and Abbeel 2020) and DDIM (Song, Meng, and Ermon 2020) introduce the idea of generating images by reversing a gradual noising process. Stable Diffusion (Rombach et al. 2022; Podell et al. 2023) further advances this line of work by operating in a latent space, which reduces computational cost and accelerates inference. To enable text conditioning, these models are typically guided by language features extracted from pretrained encoders such as CLIP (Radford et al. 2021) or T5 (Raffel et al. 2020). Thanks to its open-source availability, modular design, and strong performance, Stable Diffusion is one of the most widely used models for text-to-image generation. Today, it serves as a core backbone for many applications, including personalization (Ruiz et al. 2023; Gal et al. 2022), image editing (Hertz et al. 2022; Huang

et al. 2025), and controllable generation (Zhang, Rao, and Agrawala 2023). In this work, we also build our method on Stable Diffusion due to its broad adoption and flexibility.

**Personalized Generation in Diffusion Models.** Personalization aims to adapt a text-to-image model to generate images of specific subjects (Shi et al. 2024; Ma et al. 2024), styles (Liu et al. 2024; Wang et al. 2025), or identities (Wu et al. 2024; Jang et al. 2024), based on user-provided examples. Early methods (Gal et al. 2022; Agarwal et al. 2025; Voynov et al. 2023; Zhang et al. 2023; Ruiz et al. 2023; Fan et al. 2024; Kumari et al. 2023) learn new token embeddings or fine-tune model weights to capture personalized concepts. However, these methods often require significant computational resources and are prone to overfitting or forgetting, especially when extended to multiple concepts.

Recently, LoRA-based tuning methods have been widely adopted (Peng et al. 2025; Li, Han, and Ji 2024; Chavan et al. 2023). Instead of updating all model parameters, LoRA injects lightweight low-rank modules into the network for efficient fine-tuning. Among them, DB-LoRA (Ryu 2022) effectively personalizes large-scale diffusion models while preserving image quality. Although these methods support multi-concept generation by merging multiple LoRA modules, they often suffer from degraded identity preservation. Mix-of-Show (Gu et al. 2023) addresses this via additional training and external conditions, increasing complexity and limiting flexibility. Another method (Po et al. 2024) tackles the issue during single-module training by enforcing orthogonality between the basis vectors of different LoRA modules, mapping features into distinct subspaces and reducing interference. While this provides a useful perspective, it requires freezing the A matrix in each LoRA module to enforce orthogonality, reducing trainable parameters and harming generation performance.

Our method also addresses this issue from the perspective of single-concept training. Unlike prior approaches, we use Token Focus Masking (TFM) to distribute the influence of different LoRA modules across distinct tokens. The Token Alignment Loss (TAL) further concentrates each rare token’s attention on regions corresponding to its target concept, avoiding spatial conflicts. Our method enables multiple LoRA modules to be used jointly during inference without interference, effectively preserving identity consistency without additional training or input conditions.

## 3 Method

### 3.1 Preliminaries

**Latent Diffusion Models (LDMs).** We build our method upon the Stable Diffusion (SD V1.5) and Stable Diffusion XL (SDXL V1.0), which are latent diffusion models (LDMs) that generate high-resolution images from textual descriptions. LDMs operate by first encoding the input image into a latent space using a variational autoencoder (VAE) and then denoising a latent variable through a U-Net conditioned on text embeddings. The generation process is modeled as a Markov chain of denoising steps, aiming to learn the reverse process of a fixed forward noising schedule. The

models are trained with the following objective:

$$\mathcal{L}_{\text{denoise}} = \mathbf{E}_{\mathbf{z}_t, \mathbf{c}, \epsilon, t} \left[ \|\epsilon - \epsilon_{\theta}(\mathbf{z}_t, t, \mathbf{c})\|_2^2 \right], \quad (1)$$

where  $\mathbf{z}_t$  is the noisy latent at timestep  $t$ ,  $\mathbf{c}$  is the text condition from a pretrained language encoder,  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  and  $\epsilon_{\theta}$  is the denoising U-Net with parameters  $\theta$ .

**DreamBooth LoRA (DB-LoRA).** DreamBooth (Ruiz et al. 2023) enhances diffusion models with personalization capabilities by injecting a new concept into the model using a few (3–5) reference images. Specifically, it associates a user-defined rare token identifier  $V^*$  (e.g., “xlo”) with a visual concept by fine-tuning the entire U-Net of the diffusion model. Meanwhile, it leverages the prior knowledge of a class noun token  $V_{\text{class}}$  (e.g., “dog”) corresponding to the concept’s category to guide generation.

To improve training efficiency and modularity, DreamBooth can be adapted with LoRA (Hu et al. 2022; Ryu 2022), a parameter-efficient fine-tuning technique. LoRA inserts trainable low-rank matrices  $\Delta = \mathbf{B}\mathbf{A}$  into all linear layers of the U-Net. During training, the modified weights become  $\mathbf{W} + \Delta$ , with  $\mathbf{A} \in \mathbb{R}^{r \times d}$ ,  $\mathbf{B} \in \mathbb{R}^{d \times r}$ , and  $r \ll d$ . Only  $\mathbf{A}$  and  $\mathbf{B}$  are updated by the same objective function in Eq. (1), while the original parameters  $\mathbf{W}$  remain frozen. This design significantly reduces the number of trainable parameters and enables personalized LoRA weights to be stored and transferred independently.

### 3.2 Analysis

When multiple LoRA modules are simultaneously injected into a pre-trained diffusion model for personalized generation, the generated images often suffer from degraded quality, including *identity missing* (i.e., expected concept visual identities fail to appear) and *feature leakage* (i.e., visual features of different concepts undesirably blend together) as shown in Figure 1. We analyze the underlying causes of these problems and identify two critical issues.

**Identity Missing.** The core idea of DB-LoRA is to inject a LoRA module into the U-Net, binding a specific visual concept to a rare token identifier  $V^*$  and incorporating it into the model’s output domain. When a rare token appears in the prompt, the corresponding LoRA module is expected to influence the denoising trajectory, guiding the model to generate an image with the target concept. Intuitively, the injected LoRA module should focus primarily on the embedding of  $V^*$  within the cross-attention layers, so it can adjust the output of the linear layers when the rare token is present, thereby steering the denoising process. However, we observe that actually, the LoRA module does not focus on the rare token, but instead exerts considerable influence on the [BOS] token, as shown in Figure 2 (a).

The incorrect token focus poses no significant problem when a single LoRA module is used for generating a single concept. This is because the LoRA module in the cross-attention layers of the U-Net relies mainly on the [BOS] token to guide image generation. However, when multiple LoRA modules are combined, their effects on the [BOS]

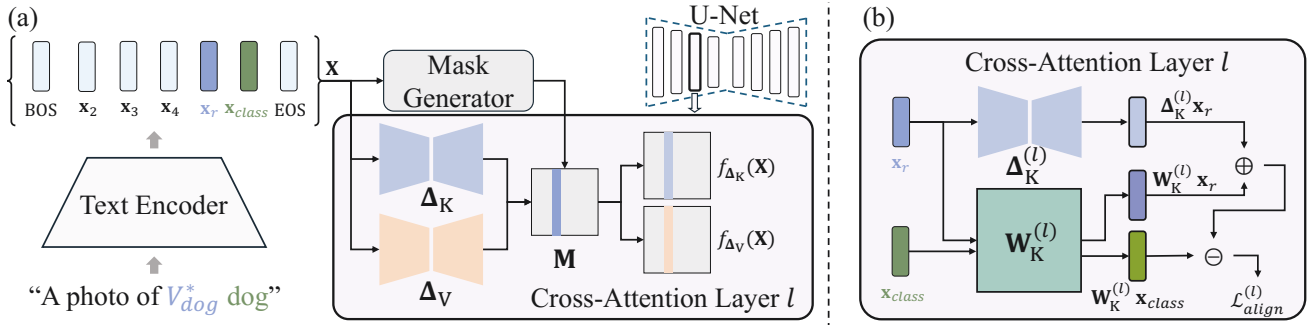


Figure 3: Overview of the proposed TARA framework. (a) Illustration of the Token Focus Masking (TFM). Given the prompt embeddings, a binary mask  $\mathbf{M}$  is generated to identify the rare token. The mask is applied to the outputs of the LoRA module inserted into the Key and Value paths of the cross-attention layer, producing masked outputs  $f_{\Delta_K}(\mathbf{X})$  and  $f_{\Delta_V}(\mathbf{X})$ . (b) Illustration of the Token Alignment Loss (TAL). The rare token embedding  $\mathbf{x}_r$  is passed through both the pre-trained attention weights  $\mathbf{W}_K^{(l)}$  and the LoRA module  $\Delta_K^{(l)}$ . The sum of these outputs is then compared to the class token embedding  $\mathbf{x}_{class}$  processed by  $\mathbf{W}_K^{(l)}$ . Their difference is used to compute the  $L_1$  loss  $\mathcal{L}_{align}^{(l)}$ , and the final alignment loss is the average over all layers.

token overlap and interfere, leading to diluted or even suppressed guidance from individual modules. As a result, the model fails to effectively steer the denoising process, ultimately causing identity missing in the output image.

Formally, let  $\mathbf{W}$  be the pre-trained Key or Value projection weight matrix of a cross-attention layer, and let  $\Delta_i$  and  $\Delta_j$  denote the LoRA modules trained for concepts  $i$  and  $j$ , respectively. Given a text condition embedding,  $\mathbf{X} = [\mathbf{x}_{BOS}, \mathbf{x}_2, \dots, \mathbf{x}_{r^i}, \dots, \mathbf{x}_{r^j}, \dots, \mathbf{x}_{n-1}, \mathbf{x}_{EOS}]$ , as the input to the Key or Value projection matrix in the cross-attention layer, where  $\mathbf{x}_{r^i}$  and  $\mathbf{x}_{r^j}$  are the rare tokens associated with concepts  $i$  and  $j$ , respectively. When only one LoRA module  $\Delta_i$  is injected, its main effect is concentrated on the  $[BOS]$  token. This effect, represented by  $\Delta_i \mathbf{x}_{BOS}$ , provides guidance for predicting noise. However, when another LoRA module  $\Delta_j$  is injected, the effect of the  $[BOS]$  token becomes  $\Delta_i \mathbf{x}_{BOS} + \Delta_j \mathbf{x}_{BOS}$ . Their effects interfere with each other, making it difficult for the model from accurately capturing the intended semantic information.

**Visual Feature Leakage.** In the DreamBooth setup, a prompt typically follows a format that combines a rare token and a class token (i.e.,  $[V^*] [V_{class}]$ ) to represent a concept. This design is intended to help the rare token inherit the prior knowledge of the class token from the pre-trained model. However, as shown in the upper row of Figure 2 (b), our attention analysis reveals that the spatial attention maps of rare tokens are often dispersed across the entire image. In contrast, class tokens tend to attend to more compact, concept-specific regions (see Figure 4). This misalignment suggests that rare tokens fail to effectively utilize the prior knowledge of their corresponding class tokens, and also fail to anchor the visual features of the concept to its correct spatial location in the generated image, potentially causing feature leakage into unrelated regions.

This issue becomes more pronounced in multi-concept generation, as the spatial attention of one rare token often overlaps with other concept-specific regions. Therefore, one concept’s features may be incorrectly associated with an-

other’s region. In the upper row of Figure 2 (b), the color associated with  $\langle dog \rangle$  is polluted by the background color of  $\langle glasses \rangle$ . We identify this phenomenon as a key contributor to visual feature leakage, where features of different subjects are mistakenly entangled.

To summarize, our analysis suggests that the primary reason why DB-LoRA modules struggle with identity preservation when injected simultaneously is twofold: different LoRA modules tend to focus on the same token  $[BOS]$ , and their spatial attentions often fail to localize to the correct concept regions. As a result, different modules interfere with each other by competing for influence over the same token representation and causing misalignment in their respective spatial attention regions.

### 3.3 Token Focus Masking

As discussed in the previous section, different LoRA modules tend to focus primarily on the same token (i.e.,  $[BOS]$ ) in the cross-attention layers, leading to interference. A solution is to constrain each injected LoRA module to operate solely on its corresponding rare token, preventing conflicts by localizing their effects to the distinct tokens.

To achieve this, we propose the *Token Focus Masking* (TFM) strategy, which is applied to Key and Value projections in cross-attention layers during both the training of a single LoRA and the inference with multiple LoRAs, constraining the LoRA module  $\Delta_i$  for concept  $i$  to operate solely on its associated rare token  $\mathbf{x}_{r^i}$ . Specifically, we introduce a binary mask  $\mathbf{M}_i \in \{0, 1\}^{d \times n}$  of the same size as the prompt embedding  $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_{r^i}, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$ , where  $n$  is the number of tokens and  $d$  is the embedding dimensionality. All entries of  $\mathbf{M}_i$  are set to 0 except for the column corresponding to the rare token  $\mathbf{x}_{r^i}$ , whose elements are all set to 1. The mask is applied to the output of a single LoRA module during the forward pass:

$$f_{\Delta_i}(\mathbf{X}) = \mathbf{M}_i \odot (\Delta_i \mathbf{X}) = [0, \dots, \Delta \mathbf{x}_{r^i}, \dots, 0], \quad (2)$$

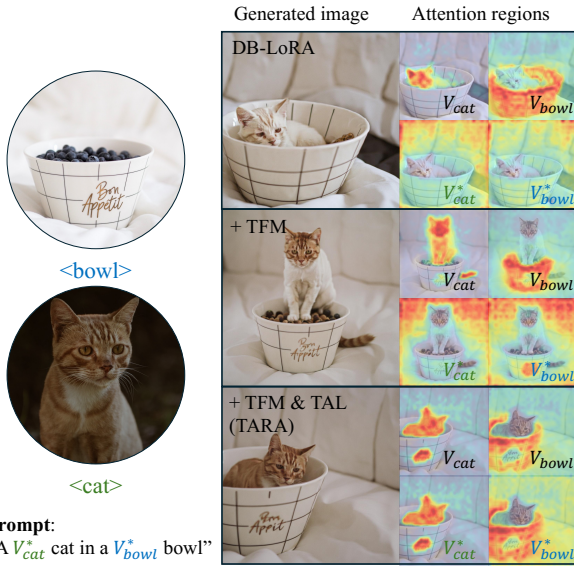


Figure 4: Visual examples illustrating the effect of each component in TARA. The left column shows generated images, and the right column visualizes the attention maps of class tokens ( $V_{cat}$ ,  $V_{bowl}$ ) and rare tokens ( $V_{cat}^*$ ,  $V_{bowl}^*$ ). DB-LoRA fails to reconstruct the concept identities. Applying TFM improves object identity preservation capability but feature leakage still exists. Combining TFM with TAL further enhances spatial alignment, enabling accurate reconstruction of both concept identity and visual features.

where  $\odot$  denotes element-wise multiplication. This ensures that only the rare token’s embedding  $\mathbf{x}_{r,i}$  contributes to the LoRA output. This mechanism is shown in Figure 3 (a).

To support multi-concept generation, we extend the token focus masking strategy to handle multiple LoRA modules simultaneously. Suppose we aim to inject  $p$  LoRA modules,  $\{\Delta_1, \dots, \Delta_p\}$ , each corresponding to a unique concept. For each concept  $i \in \{1, 2, \dots, p\}$ , we define a rare token  $\mathbf{x}_{r,i}$  and construct a binary mask  $\mathbf{M}_i \in \{0, 1\}^{d \times n}$ . Let the prompt embedding be  $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$ , which contains the set of rare tokens  $\{\mathbf{x}_{r,1}, \dots, \mathbf{x}_{r,p}\}$ . Each LoRA module’s output is masked by its corresponding token mask, as defined in Eq. (2).

During inference, the outputs of all LoRA modules are summed and added to the output of the frozen pre-trained projection weights  $\mathbf{W}$ , forming the final output:

$$F_{\mathbf{W}, \Delta_1, \dots, \Delta_p}(\mathbf{X}) = \mathbf{W}\mathbf{X} + \sum_{i=1}^p f_{\Delta_i}(\mathbf{X}). \quad (3)$$

In summary, TFM plays a critical role in enabling conflict-free multi-concept generation. By restricting each LoRA module to operate exclusively on its corresponding rare token, the masking ensures that the functional impacts of different LoRA modules do not overlap on the same token, thereby reducing the risk of identity missing. This effect is evident in the second row of Figure 4, where the identities of the  $\langle \text{cat} \rangle$  and  $\langle \text{bowl} \rangle$  concepts are preserved.

### 3.4 Token Alignment Loss

Although the proposed TFM strategy enables the injected LoRA modules in cross-attention layers to focus on their respective rare tokens, thus mitigating interference among modules and alleviating the issue of identity missing in multi-concept generation, the problem of visual feature leakage still persists. This is because the visual feature information encoded in these tokens often fails to align with the correct spatial regions of the corresponding concepts, as shown in the second row of Figure 4, where the body color of  $\langle \text{cat} \rangle$  is affected by  $\langle \text{bowl} \rangle$ .

In our experiments, we observe that the concept regions in the image often align well with the attention regions of the corresponding class tokens. This suggests that if the attention region of a rare token can be aligned with that of its class token, the visual features it encodes can be more effectively anchored to the correct concept region and can help prevent feature leakage.

In a cross-attention layer, each text token embedding  $\mathbf{x}_i$  is projected into a Key vector via a linear transformation:  $\mathbf{K}_i = \mathbf{W}_K \mathbf{x}_i$ , forming the Key matrix  $\mathbf{K} = [\mathbf{K}_1, \dots, \mathbf{K}_n]^T \in \mathbb{R}^{n \times d}$ . The latent image features are divided into  $m$  non-overlapping patch vectors  $\{\mathbf{f}_j\}_{j=1}^m$ , where each patch is projected into a Query vector  $\mathbf{Q}_j = \mathbf{W}_Q \mathbf{f}_j$ , forming the Query matrix  $\mathbf{Q} = [\mathbf{Q}_1, \dots, \mathbf{Q}_m]^T \in \mathbb{R}^{m \times d}$ . The spatial attention map is computed as:

$$\mathbf{A} = \text{softmax} \left( \frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}} \right), \quad (4)$$

where each column  $\mathbf{A}_i$  represents the attention distribution over spatial locations for token  $i$ . Since the attention scores come from the dot product between the Query and Key vectors, aligning the Key vector of the rare token with that of the class token helps align their spatial attention regions.

To this end, we propose the *Token Alignment Loss* (TAL) to guide the injected LoRA module in the Key projections. Specifically, let  $\mathbf{W}_K^{(l)}$  be the pre-trained Key projection matrix and  $\Delta_K^{(l)}$  be the corresponding injected LoRA module in the  $l$ -th cross-attention layer. Let  $\mathbf{x}_{\text{class}}$  and  $\mathbf{x}_r$  respectively denote the class token and the rare token associated with the personalized concept. TAL is defined as:

$$\mathcal{L}_{\text{align}} = \frac{1}{L} \sum_{l=1}^L \left\| \mathbf{W}_K^{(l)} \mathbf{x}_{\text{class}} - (\mathbf{W}_K^{(l)} + \Delta_K^{(l)}) \mathbf{x}_r \right\|, \quad (5)$$

where  $L$  is the total number of the cross-attention layers in the U-Net. The computation process at each layer is illustrated in Figure 3 (b). We combine  $\mathcal{L}_{\text{align}}$  with the diffusion model objective  $\mathcal{L}_{\text{denoise}}$  (defined in Equation (1)) to form the final single-LoRA training objective:

$$\mathcal{L} = \mathcal{L}_{\text{denoise}} + \lambda \mathcal{L}_{\text{align}}, \quad (6)$$

where  $\lambda$  is a weighting parameter that is set to 1.

TAL encourages the Key vector of the rare token to align with that of the class token, effectively guiding the spatial attention of the rare token to align with that of the class token. By doing so, the visual features carried by the rare token are more likely to be correctly associated to the concept region, thereby mitigating feature leakage during multi-concept generation, as shown in the third row of Figure 4.

Backbone	Method	Merge Time	CLIP-T $\uparrow$			CLIP-I $\uparrow$			DINO $\uparrow$		
			Single	Merged	$\delta$	Single	Merged	$\delta$	Single	Merged	$\delta$
SD V1.5	Custom Diffusion	<1 s	.213	.156	-.057	.722	.483	-.239	.433	.111	-.322
	DB-LoRA	<1 s	.314	.318	<u>+0.04</u>	.743	.678	-.065	.440	.333	-.107
	Mix-of-Show	~15 m	.278	.303	<b>+0.025</b>	.768	.760	<u>-.008</u>	.467	.458	<u>-.009</u>
	ROB	<1 s	.317	.315	-.002	.712	.702	-.010	.395	.380	-.015
	TARA (ours)	<1 s	.315	.314	-.001	.780	.776	<b>-0.004</b>	.489	.485	<b>-0.004</b>
SDXL V1.0	DB-LoRA	<1 s	.310	.333	<b>+0.023</b>	.802	.694	-.108	.535	.356	-.179
	ROB	<1 s	.321	.315	-.006	.777	.754	<u>-.023</u>	.495	.462	<u>-.033</u>
	TARA (ours)	<1 s	.313	.312	<u>-.001</u>	.799	.795	<b>-0.004</b>	.533	.528	<b>-0.005</b>

Table 1: Single-concept generation results. We provide quantitative comparisons for each method, evaluating single-concept generation performance both before and after the merging process. Our method achieves performance comparable to or better than others in the single-LoRA setting. More importantly, our method exhibits the smallest performance drop ( $\delta$ ) after merging another LoRA module, indicating strong robustness against interference. All results are obtained using the same prompts.

## 4 Experiments

In this section, we present the experimental results of TARA in various personalized image generation settings. Both qualitative and quantitative evaluations demonstrate that our method outperforms the DB-LoRA baseline and other popular approaches in terms of identity preservation. In addition, we conduct an ablation study to analyze the individual contributions of each component in our framework.

**Datasets.** All experiments are conducted on the DreamBooth dataset (Ruiz et al. 2023), which contains 30 diverse subjects, including 21 objects and 9 live subjects. Each subject is represented by 4–6 reference images. LoRA modules are individually fine-tuned using their corresponding image sets with a text prompt in the form of “A  $V^*V_{class}$ .” (e.g., “A  $V^*_{dog}$  dog.”) to learn personalized visual representations.

**Implementation Details.** Our method is implemented on both Stable Diffusion V1.5 and Stable Diffusion XL 1.0 backbones. For each concept, we fine-tune only the U-Net using a LoRA module with rank 8, a learning rate of 1e-5, and a batch size of 1 for 1000 epochs. Unlike some of previous methods, both the text encoder and the text embeddings in TARA are kept frozen during training. The token alignment loss weight is set to 1 by default. All experiments are conducted on one single NVIDIA RTX A6000 GPU.

**Baselines.** We compare our method with Custom Diffusion (Kumari et al. 2023), DB-LoRA (Ryu 2022), Mix-of-Show (Gu et al. 2023), and the Randomized Orthogonal Basis (ROB) method from prior works (Po et al. 2024; Zhang et al. 2025). Since the implementation of (Po et al. 2024) is unavailable and (Zhang et al. 2025) focuses on language models, we re-implement ROB by initializing the LoRA-A matrix with random orthogonal weights and keeping it frozen during training, following the core idea of both.

**Experimental Setup and Metrics.** We conduct both quantitative and qualitative comparisons with the recent methods. For quantitative evaluation, we assess both single- and multi-concept generation using CLIP (Radford et al. 2021) and DINO (Caron et al. 2021) features. Specifically,

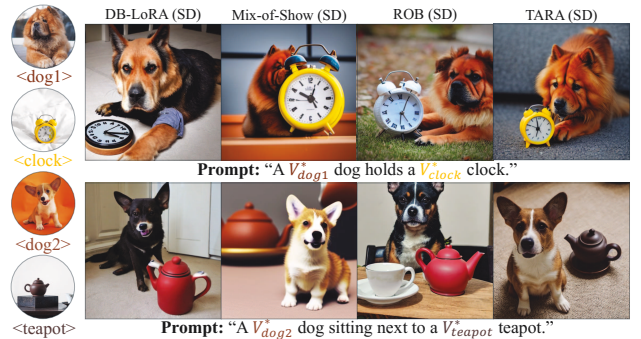


Figure 5: Qualitative comparisons between different LoRA-based methods. All images are generated using only text prompts, without additional training or external conditions. Compared to other methods, our approach better preserves the visual identity of each concept in the composed image.

CLIP-T refers to the similarity between the generated image and its corresponding text prompt in the CLIP text-image embedding space, evaluating prompt alignment. CLIP-I denotes the cosine similarity between generated and reference images in the CLIP image embedding space, measuring identity preservation. DINO score is computed as the cosine similarity between generated and reference images using features from the pretrained DINO-ViT model, offering an additional view of visual similarity. For multi-concept generation, we report the average similarity across all involved concepts. To ensure robust evaluation, we generate 10 images for each prompt across all subjects using the prompt templates from the DreamBooth dataset. We also perform an ablation study under the multi-concept generation setting to evaluate the effectiveness of Token Focus Masking and Token Alignment Loss.

### 4.1 Quantitative Comparison

**Single Concept Comparison.** Table 1 presents the quantitative results for single-concept generation before and after merging another LoRA module. The merging is used to

Method	4 concepts			3 concepts			2 concepts			Average		
	CLIP-T $\uparrow$	CLIP-I $\uparrow$	DINO $\uparrow$	CLIP-T $\uparrow$	CLIP-I $\uparrow$	DINO $\uparrow$	CLIP-T $\uparrow$	CLIP-I $\uparrow$	DINO $\uparrow$	CLIP-T $\uparrow$	CLIP-I $\uparrow$	DINO $\uparrow$
Custom Diffusion	.206	.554	.175	.184	.508	.153	.177	.506	.136	.189	.523	.155
DB-LoRA	<b>.301</b>	.585	.253	<b>.294</b>	.604	.284	.291	.621	.294	<b>.295</b>	.603	.277
Mix-of-Show	<u>.279</u>	<u>.630</u>	<u>.308</u>	.278	<u>.656</u>	<u>.347</u>	.278	<u>.696</u>	<u>.401</u>	.278	<u>.661</u>	<u>.352</u>
ROB	.274	.628	.302	.282	.653	<u>.347</u>	<b>.296</b>	.675	.378	<u>.284</u>	.652	.342
TARA (ours)	.273	<b>.633</b>	<b>.335</b>	<u>.286</u>	<b>.673</b>	<b>.386</b>	<u>.293</u>	<b>.711</b>	<b>.431</b>	<u>.284</u>	<b>.672</b>	<b>.384</b>

Table 2: Quantitative evaluation of customized multi-concept generation. TARA achieves the best performance across all multi-concept settings on CLIP-I and DINO, significantly outperforming the DB-LoRA baseline and other popular methods. The scores are computed per concept and averaged to assess identity preservation and prompt alignment quality.

	Component	CLIP-T	CLIP-I	DINO
(a)	LoRA	<b>.295</b>	.603	.277
(b)	+ TFM	<u>.284</u>	<u>.662</u>	<u>.374</u>
(c)	+ TAL	.266	.652	.342
(d)	TARA	<u>.284</u>	<b>.672</b>	<b>.384</b>

Table 3: Ablation study of TARA in multi-concept generation. Both Token Focus Masking (TFM) and Token Alignment Loss (TAL) effectively improve identity preservation (CLIP-I and DINO) during multi-concept generation. Combining both leads to the best overall performance.

test how the original concept changes after another concept (LoRA) is added. Our method achieves performance comparable to or better than other approaches in the single-LoRA setting, particularly in CLIP-I and DINO scores. More importantly, TARA exhibits the smallest performance drop after merging, with only a 0.004 drop in DINO on SD V1.5, while other methods often suffer significant degradation or require additional fusion time. These results highlight TARA’s robustness and efficiency in multi-LoRA composition, enabling near-instantaneous merging without sacrificing single-concept fidelity.

**Multi-Concept Comparison.** Table 2 presents the quantitative results for multi-concept generation using the SD V1.5 backbone. TARA consistently achieves the highest scores across the identity preservation metrics (CLIP-I and DINO). For example, it achieves a CLIP-I score of 0.707 and a DINO score of 0.426 in the 2-concept setting, outperforming DB-LoRA (0.621 / 0.294) by large margins. Even under the challenging 4-concept setting, TARA maintains strong identity preservation (CLIP-I: 0.633, DINO: 0.335), while DB-LoRA drops significantly. Meanwhile, TARA’s prompt alignment remains comparable to other methods, indicating that prompt following capability is not compromised.

Importantly, Mix-of-Show not only fine-tunes the text encoder and text embeddings, but also requires additional training for fusion mechanisms. In contrast, TARA achieves these improvements without any extra training. These results demonstrate that TARA enables robust and efficient multi-concept composition with superior identity preservation and competitive prompt alignment.

## 4.2 Qualitative Comparisons

We compare our method with other LoRA-based approaches on multi-concept personalized generation. As shown in Figure 5, our method achieves high-fidelity composition without requiring additional training or external conditions. Given only a user-provided prompt, our approach successfully integrates multiple concepts into a single coherent image, consistently outperforming other baselines in identity preservation and prompt adherence.

## 4.3 Ablation Study

To understand the contribution of each component in TARA, we conduct an ablation study under the multi-concept generation setting, as shown in Table 3. Starting from the standard LoRA baseline (a), we observe that adding Token Focus Masking (TFM) alone (b) significantly boosts identity preservation, with CLIP-I increasing from 0.603 to 0.662 and DINO from 0.277 to 0.374, while maintaining good prompt alignment (CLIP-T: 0.284). Applying Token Alignment Loss (TAL) alone (c) also improves CLIP-I (0.652) and DINO (0.342), though at the cost of a larger drop in CLIP-T. When both TFM and TAL are combined in the full TARA model (d), we achieve the best overall results: identity scores further improve (CLIP-I: 0.603  $\rightarrow$  0.672, DINO: 0.277  $\rightarrow$  0.384) while maintaining the same level of prompt alignment as TFM alone (CLIP-T: 0.284). These results confirm that TFM and TAL are complementary. TFM effectively identifies relevant tokens to focus on, and TAL strengthens spatial alignment. Together, they lead to more robust multi-concept generation.

## 5 Conclusions

We propose Token-Aware LoRA (TARA) for high-fidelity multi-concept personalization in diffusion models. Unlike DreamBooth LoRA, TARA applies token focus masking during the forward pass to constrain the effect of each injected LoRA module to its associated rare token, preventing interference caused by multiple LoRA modules acting on the same token. Additionally, we introduce a token alignment loss that guides the spatial attention of each rare token to align with its corresponding concept region, mitigating feature conflicts and improving generation quality. Together, these designs enable TARA to achieve strong identity preservation in multi-concept scenarios without requiring additional training or externally provided guidance.

## Acknowledgments

This work was supported by the Shenzhen Science and Technology Program (JSGG20220831105002004) and the Shenzhen Key Laboratory of Computer Vision and Pattern Recognition.

## References

- Agarwal, A.; Karanam, S.; Shukla, T.; and Srinivasan, B. V. 2025. An image is worth multiple words: Multi-attribute inversion for constrained text-to-image synthesis. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 6053–6062. IEEE.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9650–9660.
- Chavan, A.; Liu, Z.; Gupta, D.; Xing, E.; and Shen, Z. 2023. One-for-all: Generalized lora for parameter-efficient fine-tuning. *arXiv preprint arXiv:2306.07967*.
- Fan, Z.; Yin, Z.; Li, G.; Zhan, Y.; and Zheng, H. 2024. Dreambooth++: Boosting subject-driven generation via region-level references packing. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 11013–11021.
- Frenkel, Y.; Vinker, Y.; Shamir, A.; and Cohen-Or, D. 2024. Implicit style-content separation using b-lora. In *European Conference on Computer Vision*, 181–198. Springer.
- Gal, R.; Alaluf, Y.; Atzmon, Y.; Patashnik, O.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2022. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*.
- Gu, Y.; Wang, X.; Wu, J. Z.; Shi, Y.; Chen, Y.; Fan, Z.; Xiao, W.; Zhao, R.; Chang, S.; Wu, W.; et al. 2023. Mix-of-show: Decentralized low-rank adaptation for multi-concept customization of diffusion models. *Advances in Neural Information Processing Systems*, 36: 15890–15902.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Huang, Y.; Huang, J.; Liu, Y.; Yan, M.; Lv, J.; Liu, J.; Xiong, W.; Zhang, H.; Cao, L.; and Chen, S. 2025. Diffusion Model-Based Image Editing: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(6): 4409–4437.
- Jang, S.; Jo, J.; Lee, K.; and Hwang, S. J. 2024. Identity decoupling for multi-subject personalization of text-to-image models. *Advances in Neural Information Processing Systems*, 37: 100895–100937.
- Kumari, N.; Zhang, B.; Zhang, R.; Shechtman, E.; and Zhu, J.-Y. 2023. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1931–1941.
- Li, Y.; Han, S.; and Ji, S. 2024. Vb-lora: Extreme parameter efficient fine-tuning with vector banks. *Advances in Neural Information Processing Systems*, 37: 16724–16751.
- Liu, J.; Li, C.; Sun, Q.; Ming, J.; Fang, C.; Wang, J.; Zeng, B.; and Liu, S. 2024. Ada-adapter: Fast few-shot style personalization of diffusion model with pre-trained image encoder. *arXiv preprint arXiv:2407.05552*.
- Ma, J.; Liang, J.; Chen, C.; and Lu, H. 2024. Subject-diffusion: Open domain personalized text-to-image generation without test-time fine-tuning. In *ACM SIGGRAPH 2024 Conference Papers*, 1–12.
- Ouyang, Z.; Li, Z.; and Hou, Q. 2025. K-lora: Unlocking training-free fusion of any subject and style loras. *arXiv preprint arXiv:2502.18461*.
- Pang, L.; Yin, J.; Zhao, B.; Wu, F.; Wang, F. L.; Li, Q.; and Mao, X. 2024. Atndreambooth: Towards text-aligned personalized text-to-image generation. *Advances in Neural Information Processing Systems*, 37: 39869–39900.
- Peng, Y.; Wang, P.; Liu, J.; and Chen, S. 2025. GLAD: Generalizable Tuning for Vision-Language Models. *arXiv preprint arXiv:2507.13089*.
- Po, R.; Yang, G.; Aberman, K.; and Wetzstein, G. 2024. Orthogonal adaptation for modular customization of diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7964–7973.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmlR.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140): 1–67.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2023. 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.
- Ryu, S. 2022. Low-rank adaptation for fast text-to-image diffusion fine-tuning. <https://github.com/cloneofsimo/lora>.

Shah, V.; Ruiz, N.; Cole, F.; Lu, E.; Lazebnik, S.; Li, Y.; and Jampani, V. 2024. Ziplora: Any subject in any style by effectively merging loras. In *European Conference on Computer Vision*, 422–438. Springer.

Shi, J.; Xiong, W.; Lin, Z.; and Jung, H. J. 2024. Instant-booth: Personalized text-to-image generation without test-time finetuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8543–8552.

Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.

Tang, R.; Liu, L.; Pandey, A.; Jiang, Z.; Yang, G.; Kumar, K.; Stenatorp, P.; Lin, J.; and Ture, F. 2022. What the daam: Interpreting stable diffusion using cross attention. *arXiv preprint arXiv:2210.04885*.

Voynov, A.; Chu, Q.; Cohen-Or, D.; and Aberman, K. 2023. p+: Extended textual conditioning in text-to-image generation. *arXiv preprint arXiv:2303.09522*.

Wang, Y.; Bai, T.; Xie, X.; Yi, Z.; Wang, Y.; and Ma, R. 2025. Sigstyle: Signature style transfer via personalized text-to-image models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 8051–8059.

Wu, Y.; Li, Z.; Zheng, H.; Wang, C.; and Li, B. 2024. Infinite-id: Identity-preserved personalization via id-semantics decoupling paradigm. In *European Conference on Computer Vision*, 279–296. Springer.

Zhang, J.; You, J.; Panda, A.; and Goldstein, T. 2025. Lori: Reducing cross-task interference in multi-task low-rank adaptation. *arXiv preprint arXiv:2504.07448*.

Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3836–3847.

Zhang, Y.; Dong, W.; Tang, F.; Huang, N.; Huang, H.; Ma, C.; Lee, T.-Y.; Deussen, O.; and Xu, C. 2023. Prospect: Prompt spectrum for attribute-aware personalization of diffusion models. *ACM Transactions on Graphics (TOG)*, 42(6): 1–14.