

Resolving Multi-Condition Confusion for Finetuning-Free Personalized Image Generation

Qihan Huang^{*1, 2}, Siming Fu^{*2}, Jinlong Liu², Hao Jiang², Yipeng Yu², Jie Song^{1†}

¹ Zhejiang University

² Alibaba Group

{qh.huang, sjie}@zju.edu.cn, fusiming.fsm@taobao.com,
LJLwykqh@126.com, aoshu.jh@alibaba-inc.com, yypzju@163.com

Abstract

Personalized text-to-image generation methods can generate customized images based on the reference images, which have garnered wide research interest. Recent methods propose a finetuning-free approach with a decoupled cross-attention mechanism to generate personalized images requiring no test-time finetuning. However, when multiple reference images are provided, the current decoupled cross-attention mechanism encounters the *object confusion* problem and fails to map each reference image to its corresponding object, thereby seriously limiting its scope of application. To address the object confusion problem, in this work we investigate the relevance of different positions of the latent image features to the target object in diffusion model, and accordingly propose a weighted-merge method to merge multiple reference image features into the corresponding objects. Next, we integrate this weighted-merge method into existing pre-trained models and continue to train the model on a multi-object dataset constructed from the open-sourced SA-1B dataset. To mitigate object confusion and reduce training costs, we propose an *object quality score* to estimate the image quality for the selection of high-quality training samples. Furthermore, our weighted-merge training framework can be employed on single-object generation when a single object has multiple reference images. The experiments verify that our method achieves superior performance to the state-of-the-arts on the Concept101 dataset and DreamBooth dataset of multi-object personalized image generation, and remarkably improves the performance on single-object personalized image generation.

Code — <https://github.com/hqhQAQ/MIP-Adapter>

Introduction

Personalized text-to-image generation methods generate images conditioned on the reference images that specify the details of the generated contents, sparking considerable research interest due to its diverse applications. The methodology in this domain is gradually shifting from a *finetuning-based* approach (e.g., DreamBooth (Ruiz et al. 2023),

^{*}These authors contributed equally.

[†]Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

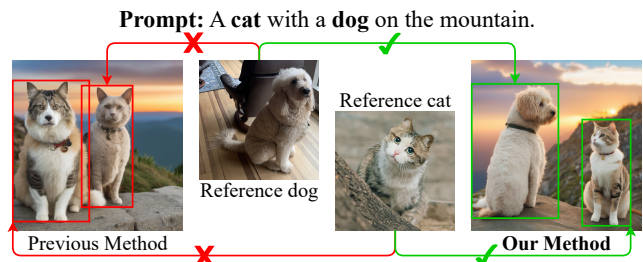


Figure 1: Left image demonstrates the object confusion problem in decoupled cross-attention mechanism, and right image presents the correct generation using our method.

Custom Diffusion (Kumari et al. 2023)) to a *finetuning-free* technique (e.g., IP-Adapter (Ye et al. 2023), Subject-Diffusion (Ma et al. 2024)), as finetuning-free methods eliminate the need for finetuning during test time and significantly reduce the usage cost.

Early finetuning-free methods, such as InstantBooth (Shi et al. 2024) and FastComposer (Xiao et al. 2023), simply integrate the features of the reference image into the text embeddings and feed them into the text encoder, without fully exploiting the information from the reference image. Recent finetuning-free methods, such as IP-Adapter (Ye et al. 2023), more comprehensively utilize the features of the reference image by training additional cross-attention layers to integrate reference image features into the intermediate layers of the diffusion model, and achieve comparable performance to the finetuning-based methods. However, the current decoupled cross-attention only considers one reference image for each generation. When multiple reference images are provided, the decoupled cross-attention suffers from the *object confusion* problem if applied straightforwardly, wherein object features in the reference images are assigned to the wrong objects in the generated images, as illustrated in Figure 1. Some previous image generation methods (Yang et al. 2024) attempt to mitigate the object confusion issue by incorporating the object features into the corresponding regions of latent image features in the diffusion model. Nevertheless, as the object information is distributed over the entire image feature space rather than confined to the corresponding local region owing to large receptive fields in deep networks (Luo et al. 2016; Araujo, Norris, and Sim 2019),

the generated images can be limited in faithfulness to the reference images (*i.e.*, the appearance differs between the generated and the reference images), as shown in Figure 2.

In this work, rather than splitting latent features into different regions, we propose a **weighted-merge method** to merge the reference image features into the whole latent image features with different weights on different positions. Specifically, this work estimates these weights as the relevance of different positions in latent image features to the target object, by ingeniously utilizing the cross-attention weights between the text features of the target object and the latent image features within the stable diffusion model. Besides, we design an experiment that adds different noise to the latent image features based on the predicted object relevance, verifying the effectiveness of this object relevance estimation method. We employ this method on the pre-trained finetuning-free personalized generation models (*e.g.*, IP-Adapter), enabling multi-object generation by simultaneously merging multiple conditions (reference images & text prompts) into the model. Experiment results indicate that our method can alleviate object confusion and significantly improve the performance of multi-object personalized image generation for these models without any training.

Although weighted-merge effectively alleviates object confusion, adding multiple reference images at once will interfere with the latent image features, causing them to deviate from their distribution in the original model and resulting in lower generation quality. To address this issue, this work trains the pre-trained finetuning-free model with the weighted-merge method on a multi-object dataset. Specifically, this dataset is constructed from the open-sourced SA-1B dataset (Kirillov et al. 2023) consisting of about 11 million images with multiple objects. Besides, this work proposes an object quality score to estimate the object quality of the image, according to the the degree of confusion between multiple objects, as well as the matching degree between object texts and images. Based on the object quality score, we can select high-quality images that alleviate the object confusion problem for higher performance while decreasing training costs.

Moreover, this weighted-merge training framework can be applied to single-object generation, because a single object has multiple reference images in reality. Compared to previous approaches that only use a single reference image or simply average the features of multiple images, our weighted-merge method can extract diverse useful information from different reference images and adaptively merge them to achieve superior results.

We perform comprehensive experiments to validate the performance of our proposed framework. Experiment results demonstrate that with only 100,000 high-quality images (0.13% of the dataset from Subject Diffusion) selected from SA-1B, our model achieves state-of-the-art performance on the Concept101 dataset and DreamBooth dataset of multi-object personalized image generation. Besides, our weighted-merge training framework significantly improves the performance of pre-trained model on the DreamBooth dataset of single-object personalized image generation.

To sum up, the main contributions of this work can be




Figure 2: The reference image (IP) features, with the background masked, reduces generation quality in IP-Adapter.

summarized as follows:

- We extend the decoupled cross-attention mechanism of finetuning-free personalized image generation methods to merge multiple conditions, with a proposed weighted-merge method to tackle the object confusion problem.
- We construct a small but high-quality dataset from the open-sourced SA-1B dataset for model training, with a proposed object quality score for image selection.
- Experiment results demonstrate that our weighted-merge training framework outshines in merging multiple conditions, and our model achieves state-of-the-art performance on both the Concept101 dataset and DreamBooth dataset of multi-object personalized image generation.

Related Work

Finetuning-based personalized image generation. Early personalized image generation methods require finetuning the original diffusion model on the reference images. Specifically, DreamBooth finetunes the entire UNet network of diffusion model, Textual Inversion (Gal et al. 2023) finetunes only the special embedding vector of the target object, and Custom Diffusion finetunes only the K and V layers of the cross-attention in the UNet network. Cones (Liu et al. 2023b) detects the concept neurons in the K and V layers and updates them during training. Mix-of-Show (Gu et al. 2023) trains a separate LoRA model for each object and merges them with gradient fusion.

Finetuning-free personalized image generation. These methods train the model to directly incorporate the reference image features on a large dataset, without the need for additional finetuning during test time. Early finetuning-free methods (*e.g.*, InstantBooth, FastComposer, and Taming Encoder (Jia et al. 2023)) simply integrate the image features into the text embeddings, without fully utilizing the reference image information. Recent methods (*e.g.*, IP-Adapter, ELITE (Wei et al. 2023), PatchDPO (Huang et al. 2024a), and SSR-Encoder (Zhang et al. 2024)) make more extensive utilization of reference image information by integrating the image features into the middle layers of the diffusion model, using a decoupled cross-attention mechanism. These methods excel at merging a single reference image and achieve impressive performance. However, decoupled cross-attention encounters the object confusion prob-

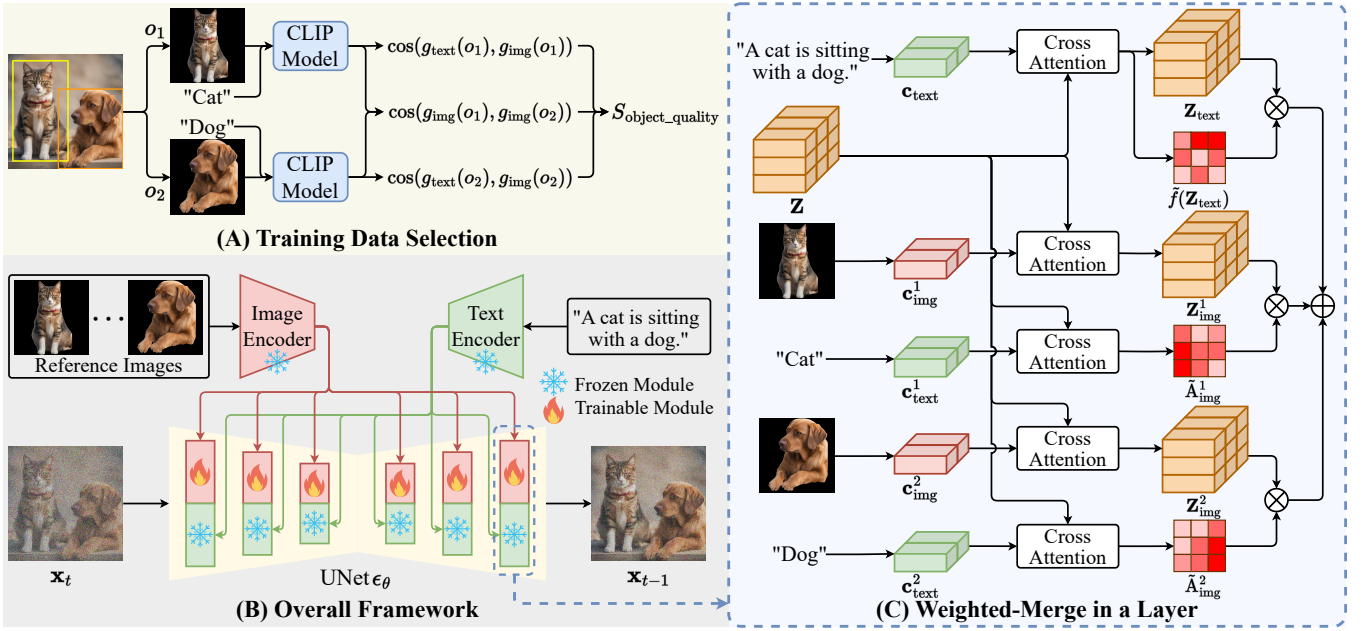


Figure 3: (A) demonstrates the calculation of $S_{\text{Object_relevance}}$, which is used for selecting training data. The overall framework in (B) consists of a UNet model for noise prediction conditioned on the text prompt and multiple reference images. (C) presents the proposed weighted-merge method in each cross-attention layer of UNet from (B). $\tilde{A}_{\text{img}}^i = \frac{A_{\text{img}}^i}{A_{\text{img}}^i}$, and $\tilde{f}(\mathbf{Z}_{\text{text}}) = \frac{f(\mathbf{Z}_{\text{text}})}{f(\mathbf{Z}_{\text{text}})}$.

lem when merging multiple reference images, a problem this study aims to address.

Method

In this section, we first give the preliminaries in **section 1**, then propose the object relevance estimation method in **section 2**. Next, **section 3** and **section 4** propose the weighted-merge method and directly apply it to the current pre-trained model. Finally, **section 5** proposes the training framework for further performance improvement.

1. Preliminaries

Diffusion model. Current personalized image generation methods adopt diffusion model (Ho, Jain, and Abbeel 2020; Rombach et al. 2022; Liu et al. 2024) as the base model. Diffusion model consists of two processes: a diffusion process which gradually adds noise into the original image with a Markov chain in T steps, and a denoising process which predicts the noise to generate the image using a deep neural network. Specifically, personalized image generation methods generate images simultaneously conditioned on the text prompt and the reference images. Typically, ϵ_θ denotes the deep neural network for noise prediction, and the training loss of personalized diffusion model is defined as below:

$$\mathcal{L} = \mathbb{E}_{\mathbf{x}_0, \epsilon \in \mathcal{N}(0, \mathbf{I}), \mathbf{c}_{\text{text}}, \mathbf{c}_{\text{img}}} \|\epsilon - \epsilon_\theta(\mathbf{x}_t, \mathbf{c}_{\text{text}}, \mathbf{c}_{\text{img}}, t)\|^2,$$

where \mathbf{x}_0 denotes the original real image, $t \in [0, T]$ denotes the time step in the diffusion process, $\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \epsilon$, and α_t, σ_t are predefined weights for step t in the diffusion process. \mathbf{c}_{text} denotes the text features, and \mathbf{c}_{img} denotes the reference image features. After training, the model

can generate images by progressively denoising Gaussian noise in multiple steps.

Decoupled cross-attention mechanism. Recent finetuning-free personalized image generation methods adopt decoupled cross-attention to merge the text features and reference image features into the middle layers of model ϵ_θ . Specifically, the latent image features $\mathbf{Z} \in \mathbb{R}^{(H \cdot W) \times D}$ in a middle layer are fed into a cross-attention module to interact with the text features $\mathbf{c}_{\text{text}} \in \mathbb{R}^{S_{\text{text}} \times D_{\text{text}}}$:

$$\mathbf{Z}_{\text{text}} = \text{Attn}(\mathbf{Q}, \mathbf{K}_{\text{text}}, \mathbf{V}_{\text{text}}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}_{\text{text}}^\top}{\sqrt{d}}\right)\mathbf{V}_{\text{text}}.$$

Here, $\mathbf{Q} = \mathbf{Z}\mathbf{W}^{\mathbf{Q}}$, $\mathbf{K}_{\text{text}} = \mathbf{c}_{\text{text}}\mathbf{W}_{\text{text}}^{\mathbf{K}}$, $\mathbf{V}_{\text{text}} = \mathbf{c}_{\text{text}}\mathbf{W}_{\text{text}}^{\mathbf{V}}$ are the query, key, and value matrices of the attention operation, respectively, and $\mathbf{W}^{\mathbf{Q}} \in \mathbb{R}^{D \times D}$, $\mathbf{W}_{\text{text}}^{\mathbf{K}} \in \mathbb{R}^{D_{\text{text}} \times D}$, $\mathbf{W}_{\text{text}}^{\mathbf{V}} \in \mathbb{R}^{D_{\text{text}} \times D}$ are the learnable weight matrices for feature projection. Besides, \mathbf{Z} is also fed into another cross-attention module to interact with the reference image features $\mathbf{c}_{\text{img}} \in \mathbb{R}^{S_{\text{img}} \times D_{\text{img}}}$:

$$\mathbf{Z}_{\text{img}} = \text{Attn}(\mathbf{Q}, \mathbf{K}_{\text{img}}, \mathbf{V}_{\text{img}}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}_{\text{img}}^\top}{\sqrt{d}}\right)\mathbf{V}_{\text{img}}.$$

Likewise, $\mathbf{K}_{\text{img}} = \mathbf{c}_{\text{img}}\mathbf{W}_{\text{img}}^{\mathbf{K}}$, $\mathbf{V}_{\text{img}} = \mathbf{c}_{\text{img}}\mathbf{W}_{\text{img}}^{\mathbf{V}}$, and $\mathbf{W}_{\text{img}}^{\mathbf{K}} \in \mathbb{R}^{D_{\text{img}} \times D}$, $\mathbf{W}_{\text{img}}^{\mathbf{V}} \in \mathbb{R}^{D_{\text{img}} \times D}$ are the learnable weight matrices for projecting the reference image features. Next, the final output of the decoupled cross-attention \mathbf{Z}_{new} is the addition of \mathbf{Z}_{text} and \mathbf{Z}_{img} :

$$\mathbf{Z}_{\text{new}} = \mathbf{Z}_{\text{text}} + \mathbf{Z}_{\text{img}}.$$

2. Object Relevance Estimation

Decoupled cross-attention mechanism excels at merging a single reference image into the model ϵ_θ than previous methods that only add the reference image features into the text embeddings. However, decoupled cross-attention simply merges the text-conditioned latent image features \mathbf{Z}_{text} and image-conditioned latent image features \mathbf{Z}_{img} with an addition operation, without constraining the reference image to the corresponding object in the text prompt. This results in an *object confusion problem* when merging multiple reference images, which incorrectly adds the reference image information to its unrelated objects, as shown in Figure 1. Therefore, given M reference images corresponding to M objects in the text prompt, this work strives to merge M image-conditioned latent image features $\{\mathbf{Z}_{\text{img}}^i\}_{i=1}^M$ into the text-conditioned latent image features \mathbf{Z}_{text} by resolving the object confusion problem.

To this end, this work first investigates the information distribution of an object (as referenced in the text prompt) on \mathbf{Z}_{text} . Some methods (e.g., RPG (Yang et al. 2024)) assume that the position of the object in \mathbf{Z}_{text} is the same as that in the generated image, however, this assumption is not accurate. Actually, the deep neurons in the deep neural networks have large effective receptive fields (Luo et al. 2016; Araujo, Norris, and Sim 2019; Huang et al. 2023, 2024b), meaning that a wide range of latent image features can affect the target object in the generated image, rather than being limited to only the local latent image features with the same position as the target object. As shown in Figure 2, \mathbf{Z}_{img} , with the background masked, will decrease the generation quality of the foreground cat in the generated image. Therefore, simply adding the reference image information into some local regions of \mathbf{Z}_{text} will lead to information loss and degrade the performance.

To tackle this problem, we estimate the relevance of **all positions** in \mathbf{Z}_{text} to the target object, and merge $\mathbf{Z}_{\text{img}}^i$ into each position of \mathbf{Z}_{text} with different weights according to the estimated relevance. For estimating the object relevance, this work ingeniously utilizes the original cross-attention modules within model ϵ_θ and calculates the attention map between the text features of the object and the original latent image features \mathbf{Z} (note that \mathbf{Z}_{text} is calculated from \mathbf{Z}). Specifically, we first extract the text features $\mathbf{c}_{\text{text}}^i \in \mathbb{R}^{S_{\text{text}} \times D_{\text{text}}}$ of the i -th object (corresponding to the i -th reference image) by feeding the object text into the text encoder. Next, the object relevance $A_{\text{img}}^i \in \mathbb{R}^{(H \cdot W)}$ of the i -th object to $\mathbf{Z}_{\text{text}} \in \mathbb{R}^{(H \cdot W) \times D}$ is calculated by averaging the original cross-attention matrix:

$$A_{\text{img}}^i = \frac{1}{S_{\text{text}}} \sum_{j=1}^{S_{\text{text}}} \text{Softmax}\left(\frac{\mathbf{K}_{\text{text}}^i \mathbf{Q}^\top}{\sqrt{d}}\right)[j],$$

where $\mathbf{Q} = \mathbf{Z}\mathbf{W}^\mathbf{Q}$, $\mathbf{K}_{\text{text}}^i = \mathbf{c}_{\text{text}}^i \mathbf{W}_{\text{text}}^\mathbf{K}$ (note that $\mathbf{W}_{\text{text}}^\mathbf{K}$ is shared with the original text features \mathbf{c}_{text}), and $\text{Softmax}\left(\frac{\mathbf{K}_{\text{text}}^i \mathbf{Q}^\top}{\sqrt{d}}\right)[j] \in \mathbb{R}^{(H \cdot W)}$ is the j -th element of $\text{Softmax}\left(\frac{\mathbf{K}_{\text{text}}^i \mathbf{Q}^\top}{\sqrt{d}}\right) \in \mathbb{R}^{S_{\text{text}} \times (H \cdot W)}$.

Merging Method	$S_{\text{object.relevance}}$	CLIP-T	CLIP-I	DINO
Uniform-Merge	1.33	0.6343	0.6409	0.3481
Weighted-Merge	1.66	0.6427	0.6503	0.3624

Table 1: The performance of different merging methods for the pre-trained IP-Adapter (training-free) on *Concept101*.

3. Training-Free Personalized Image Generation

Based on the above object relevance estimation method, we propose a weighted-merge method to extend current pre-trained models (e.g., IP-Adapter) to multi-object personalized image generation, in a training-free manner. Specifically, this method first generates the text-conditioned latent image features \mathbf{Z}_{text} and M image-conditioned latent image features $\{\mathbf{Z}_{\text{img}}^i \in \mathbb{R}^{(H \cdot W) \times D}\}_{i=1}^M$ using the original model, then merges them using the estimated object relevance $\{A_{\text{img}}^i \in \mathbb{R}^{(H \cdot W)}\}_{i=1}^M$ as weights:

$$\mathbf{Z}_{\text{new}} = \mathbf{Z}_{\text{text}} + \sum_{i=1}^M \frac{A_{\text{img}}^i}{\bar{A}_{\text{img}}^i} \odot \mathbf{Z}_{\text{img}}^i,$$

where \odot is element-wise multiplication with $A_{\text{img}}^i[p, q] \in \mathbb{R}$ and $\mathbf{Z}_{\text{img}}^i[p, q] \in \mathbb{R}^D$ ($p \in \{1, 2, \dots, H\}, q \in \{1, 2, \dots, W\}$) as each element-pair. Here, $\bar{A}_{\text{img}}^i \in \mathbb{R}$ is the average of A_{img}^i , and the division operation is used for normalization (i.e., the average value of $\frac{A_{\text{img}}^i}{\bar{A}_{\text{img}}^i}$ equals 1). This method adds each $\mathbf{Z}_{\text{img}}^i$ more to the positions in \mathbf{Z}_{text} with higher relevance to the corresponding object, thus incorporating reference image information more accurately into the corresponding object and mitigating object confusion. Table 1 shows that this weighted-merge method can remarkably improve the performance of multi-object personalized image generation on the pre-trained IP-Adapter.

4. Verification with Object Relevance Score

To verify A_{img}^i accurately represents the object relevance of each position in \mathbf{Z}_{text} , we conduct an experiment in the original text-to-image diffusion model that evaluates the object relevance score $S_{\text{object.relevance}}$ by adding noise to \mathbf{Z}_{text} . Detailedly, we calculate $S_{\text{object.relevance}}$ in three steps: (1) Generate the bounding box bbox_x of the target object in the generated image x using the Grounding DINO (Liu et al. 2023a) detection model. (2) Let x_{noise} denote the generated image with noise added on \mathbf{Z}_{text} , $x_{\text{no.noise}}$ denote the generated image without adding noise, then calculate Δ_x^{bbox} as the averaged difference between the pixels of bbox_x in x_{noise} and $x_{\text{no.noise}}$. $\Delta_x^{\text{non.bbox}}$ is calculated likewise for the region outside the bounding box bbox_x . (3) Finally, $S_{\text{object.relevance}}$ is calculated as the ratio between the Δ_x^{bbox} and $\Delta_x^{\text{non.bbox}}$ averagely over all generated images \mathcal{X} ($\|\cdot\|$ denotes cardinality of a set):

$$S_{\text{object.relevance}} = \frac{1}{\|\mathcal{X}\|} \sum_{x \in \mathcal{X}} \frac{\Delta_x^{\text{bbox}}}{\Delta_x^{\text{non.bbox}}}.$$

Therefore, higher $S_{\text{object.relevance}}$ indicates that the added noise has a higher impact on the target object compared to other regions. We conduct this experiment on the total 1212 text prompts from Concept101 dataset (Kumari et al. 2023), and the seed for generating each pair of $\mathbf{x}_{\text{noise}} \in \mathbb{R}^{(H \cdot W) \times D}$ and $\mathbf{x}_{\text{no.noise}}$ is set to the same. Two strategies for adding noise are compared: **uniform-merge** and **weighted-merge**. Uniform-merge directly adds the noise ϵ_{object} equally into all positions of \mathbf{Z}_{text} , while weighted-merge adds the noise with different weights on different positions: $\epsilon_{\text{object}} \odot (A_{\text{img}}^i / \bar{A}_{\text{img}}^i)$. Note that the norm of the sampled noise is equal for these two strategies for fair comparison. As shown in Table 1, $S_{\text{object.relevance}} > 1$ for uniform-merge, indicating that the background of \mathbf{Z}_{text} also has a great influence on the object. Besides, weighted-merge achieves significantly higher $S_{\text{object.relevance}}$ than uniform-merge, implying that weighted-merge can effectively estimate the object relevance on \mathbf{Z}_{text} .

5. Training-Based Personalized Image Generation

However, this training-free weighted-merge method still lags behind other multi-object personalized image generation methods, because: (1) The pre-trained model are trained with only one reference image as input, and directly adding multiple $\mathbf{Z}_{\text{img}}^i$ will easily disrupt \mathbf{Z}_{new} from its original feature distribution and decrease the quality of the generated images. (2) Different $\mathbf{Z}_{\text{img}}^i$ may still conflict at the same position in \mathbf{Z}_{text} when the corresponding values of A_{img}^i are both high. To tackle these problems, we propose to continue to train the model with the weighted-merge method on a multi-object dataset, which is to align $\mathbf{Z}_{\text{text}} + \sum_{i=1}^M \frac{A_{\text{img}}^i}{\bar{A}_{\text{img}}^i} \odot$

$\mathbf{Z}_{\text{img}}^i$ with the original feature distribution for higher image quality and alleviate the conflict of different $\mathbf{Z}_{\text{img}}^i$.

To this end, we first construct the multi-object dataset from the open-sourced SA-1B dataset, following the data-construction paradigm of Subject-Diffusion. This data-construction paradigm adopts the pre-trained BLIP2 (Li et al. 2023), Grounding DINO (Liu et al. 2023a), and SAM (Kirillov et al. 2023) to generate the text prompts, bounding boxes, and segmentation maps of each image. Furthermore, we propose an object quality score $S_{\text{object.quality}}$ to estimate the object quality of each image and accordingly select the images with high $S_{\text{object.quality}}$. Detailedly, $S_{\text{object.quality}}$ is calculated based on two factors: (1) the quality of each individual object; (2) the quality of each pair of objects. The first factor is to ensure that the image of each object (cropped from the original image) is consistent with the object text. The second factor is to select the object pairs with **lower similarities**, which facilitates the model to resolve the conflict between multiple reference images and mitigate the object confusion problem, instead of continuing wrongly adding the information of **another similar** object into the current object. We utilize the CLIP model g to assess these two factors because of its excellent cross-modal ability. Let $\mathcal{O}_{\mathbf{x}}$ denote the objects in image \mathbf{x} , $g_{\text{text}}(o) \in \mathbb{R}^{D_{\text{clip}}}$ and $g_{\text{img}}(o) \in \mathbb{R}^{D_{\text{clip}}}$ denote the text and image features of object o , then $S_{\text{object.quality}}$ of image \mathbf{x} is calculated as be-

low ($\cos(\cdot, \cdot)$ denotes cosine similarity):

$$\begin{cases} S_{\text{object.quality}} = S_{\text{single.object}} + S_{\text{object.pair}} \\ S_{\text{single.object}} = \frac{1}{\mathcal{N}_1} \sum_{o \in \mathcal{O}_{\mathbf{x}}} \cos(g_{\text{text}}(o), g_{\text{img}}(o)) \\ S_{\text{object.pair}} = -\frac{1}{\mathcal{N}_2} \sum_{o', o'' \in \mathcal{O}_{\mathbf{x}}; o' \neq o''} \cos(g_{\text{img}}(o'), g_{\text{img}}(o'')) \end{cases}$$

Here, $\mathcal{N}_1 = \|\mathcal{O}_{\mathbf{x}}\|$ and $\mathcal{N}_2 = \|\mathcal{O}_{\mathbf{x}}\|(\|\mathcal{O}_{\mathbf{x}}\| - 1)$ are the normalization terms. Detailedly, for multi-object personalized image generation, we first filter 215,789 images with multiple annotated objects using the data construction paradigm of Subject-Diffusion, then utilize 100,000 images with the highest $S_{\text{object.quality}}$ for training.

Model architecture. Figure 3 demonstrates the whole pipeline of our method. We follow previous methods to freeze the original text-to-image diffusion model and only train the parameters ($\mathbf{W}_{\text{img}}^{\text{K}}$ and $\mathbf{W}_{\text{img}}^{\text{V}}$) for generating each $\mathbf{Z}_{\text{img}}^i$ in each layer. Note that $\mathbf{W}_{\text{img}}^{\text{K}}$ and $\mathbf{W}_{\text{img}}^{\text{V}}$ are shared for generating each $\mathbf{Z}_{\text{img}}^i$ to save training cost. Besides, we propose another weighted-merge method to predict the relevance of each position in \mathbf{Z}_{text} to object-unrelated texts, which is to resolve the conflict between \mathbf{Z}_{text} and $\{\mathbf{Z}_{\text{img}}^i\}_{i=1}^M$. However, it is difficult to directly extract the text features of these object-unrelated texts and calculate the corresponding cross-attention matrix like A_{img}^i . To address this problem, this work proposes to predict the weight for the text features with a trainable prediction layer. Specifically, let $f(\mathbf{Z}_{\text{text}}) \in \mathbb{R}^{(H \cdot W)}$ denote the predicted weight for \mathbf{Z}_{text} (f is the trainable linear layer followed with a Sigmoid activation function), then \mathbf{Z}_{new} is calculated as below:

$$\mathbf{Z}_{\text{new}} = \frac{f(\mathbf{Z}_{\text{text}})}{f(\mathbf{Z}_{\text{text}})} \odot \mathbf{Z}_{\text{text}} + \sum_{i=1}^M \frac{A_{\text{img}}^i}{\bar{A}_{\text{img}}^i} \odot \mathbf{Z}_{\text{img}}^i.$$

Single-object personalized generation. Our weighted-merge training framework can be extended to other scenarios of simultaneous merging multiple conditions, such as single-object personalized image generation with multiple reference images. In real applications, a single object may have multiple reference images (e.g., each object has 4 to 6 reference images in the DreamBooth dataset). However, previous decoupled cross-attention approaches can only use a single reference image or simply average the features of multiple images, without fully utilizing the information from different reference images. To tackle this problem, we continue to train the models using our weighted-merge training framework, enabling the model to extract diverse useful information from different reference images and adaptively merge them to achieve superior results.

Experiments

Implementation details. Our main experiments are conducted on the pre-trained IP-Adapter with sdxl model (Podell et al. 2023) and sdxl_plus (Jaegle et al. 2021) model as the text-to-image diffusion model and OpenCLIP ViT-bigG/14 as the image encoder. The parameters of sdxl &

Method	Type	CLIP-T	CLIP-I	DINO
DreamBooth •	FT	0.7383	0.6636	0.3849
Custom Diffusion (Opt) •	FT	0.7599	0.6595	0.3684
Custom Diffusion (Joint) •	FT	0.7534	0.6704	0.3799
Mix-of-Show §	FT	0.7280	0.6700	0.3940
MC ² §	FT	0.7670	0.6860	0.4060
FastComposer ★	no-FT	0.7456	0.6552	0.3574
λ-ECLIPSE ★	no-FT	0.7275	0.6902	0.3902
ELITE ★	no-FT	0.6814	0.6460	0.3347
IP-Adapter ★	no-FT	0.6343	0.6409	0.3481
SSR-Encoder ★	no-FT	0.7363	0.6895	0.3970
Ours (sdxl)	no-FT	0.7750	0.6943	0.4127
Ours (sdxl_plus)	no-FT	0.7765	0.6950	0.4397

Table 2: Performance comparison for multi-object personalized generation on *Concept101*. Here, “FT” denotes finetuning-based method, “no-FT” denotes finetuning-free method, and bold font denotes the best result. Each CLIP-T score is multiplied by 2.5 following Custom Diffusion.

Method	Type	CLIP-T	CLIP-I	DINO
DreamBooth †	FT	0.308	0.695	0.430
Custom Diffusion †	FT	0.300	0.698	0.464
Subject Diffusion †	no-FT	0.310	0.696	0.506
Ours (sdxl)	no-FT	0.311	0.726	0.482

Table 3: Performance comparison for multi-object personalized generation on *DreamBooth*.

sdxl_plus model and image encoder are frozen, and only the parameters for projecting image features and predicting text weights are trainable. During training, we adopt AdamW optimizer with a learning rate of 1e-4, and train the model on 8 PPUs for 30,000 steps with a batch size of 4 per PPU. To enable classifier-free guidance, we use a probability of 0.05 to drop text and image individually, and a probability of 0.05 to drop text and image simultaneously. We also conduct experiments on other pre-trained models based on decoupled cross-attention to verify the generalization ability of our method, in S2.2 of the appendix.

Test benchmark. For multi-object personalized image generation, we follow the *Concept101* (Kumari et al. 2023) benchmark that has evaluated many methods. Besides, we also evaluate our method on the *DreamBooth* benchmark for comparison with Subject-Diffusion.

Evaluation metrics. We follow previous methods to adopt three metrics (CLIP-T, CLIP-I, and DINO) for evaluation. Specifically, CLIP-T evaluates the similarity between the generated images and given text prompts; CLIP-I and DINO evaluate the similarity between the generated images and the reference images. 5 images are generated for each prompt to ensure the evaluation stability.

Baseline methods. We compare our method with both finetuning-based methods (e.g., Textual Inversion (Gal et al. 2023), DreamBooth (Ruiz et al. 2023), Custom Diffusion (Kumari et al. 2023)) and finetuning-free methods (e.g.,

Method	CLIP-T	CLIP-I	DINO
Uniformly Add	0.7702	0.6816	0.3937
Locally Add	0.7732	0.6851	0.3958
+ Image Weight	0.7734	0.6940	0.4079
+ Text Weight	0.7726	0.6924	0.4032
+ Image & Text Weights	0.7750	0.6943	0.4127

Table 4: Ablation experiments of weighted-merge methods for multi-object personalized generation on *Concept101*.

Method	CLIP-T	CLIP-I	DINO
100,000 images (lowest $S_{\text{object.pair}}$)	0.7708	0.6880	0.3963
100,000 images (highest $S_{\text{object.pair}}$)	0.7733	0.6923	0.4056
100,000 images (highest $S_{\text{object.quality}}$)	0.7750	0.6943	0.4127

Table 5: Ablation experiments of image selection strategies for multi-object personalized generation on *Concept101*.

SSR-Encoder (Zhang et al. 2024), Subject-Diffusion).

Multi-Object Personalized Generation

We conduct both quantitative and qualitative comparisons between our method and baseline methods.

Quantitative comparisons. Table 2 demonstrates the quantitative results of different methods on *Concept101*. Note that the results of methods marked with • are from the GitHub page of Custom Diffusion (Kumari et al. 2023), the results of methods marked with § are from the paper of MC² (Jiang et al. 2024), and the results of methods marked with ★ are re-implemented faithfully following their released code and weights (their original evaluation datasets have not been made public).

As shown in Table 2, recent methods enhance the utilization of image information with decoupled cross-attention to integrate image features into the middle layers of the model, but they have yet to achieve satisfactory results due to the object confusion problem. Differently, our method generalizes decoupled cross-attention to merging multiple reference images by resolving object confusion, which achieves significantly superior performance to existing methods.

Table 3 demonstrates the quantitative results of different methods on the DreamBooth dataset. The results of methods marked with † are from the paper of Subject-Diffusion. In this benchmark, our method outperforms Subject-Diffusion in 2 of 3 evaluation metrics, and surpasses it in the CLIP-I score by a large margin (0.726 vs. 0.696).

Qualitative comparisons. Figure 4 demonstrates the qualitative results of different methods on *Concept101*. The results of the original IP-Adapter indicate that it generates images with low image quality, due to the object confusion problem and the distortion of feature distribution when merging multiple images once. Next, after employing the weighted-merge training framework on the original IP-Adapter, our method can generate images with high image quality and mitigate object confusion, realizing the best qualitative results.

Besides, we provide more visualization results of our



Figure 4: Qualitative comparisons of different methods on multi-object personalized image generation.




Figure 5: Qualitative ablation experiment.

method in S3 of the appendix (e.g., simultaneously merging more than two objects).

Single-Object Personalized Generation

For single-object personalized image generation, we utilize the proposed $S_{\text{single-object}}$ ($S_{\text{object-pair}}$ is eliminated in the single-object scenario) to select 100,000 high-quality images for training. As shown in S2.1 of the appendix, our weighted-merge framework can improve all three scores of the original IP-Adapter and ELITE on the DreamBooth dataset. Besides, Figure 1 of the appendix shows the qualitative comparisons between our model and the original model, implying that our model can capture important image information from different images, instead of ignoring the unique details of some images by the original model.

Ablation Experiments

We conduct the ablation experiments mainly on the IP-Adapter based on sdxl model.

Weighted-merge training framework. We conduct ablation experiments on two proposed weight estimation methods (text weight $f(\mathbf{Z}_{\text{text}})$ & image weight $\{A_{\text{img}}^i\}_{i=1}^M$) of the weighted-merge training framework. Table 4 demonstrates that locally adding reference image features does not show obvious improvement compared to uniform adding. Besides, Table 4 indicates that these two weight estimation methods effectively enhance the performance of multi-object personalized generation. Moreover, the qualitative ablation ex-




Figure 6: The attention maps of reference image features on the latent image features \mathbf{Z} before/after training.

periment in Figure 5 also verifies the effectiveness of our weighted-merge method with the visualization results.

Image selection. Table 5 shows the performance of multi-object personalized generation with different image selection strategies, implying that the images selected by our proposed $S_{\text{object-quality}}$ lead to superior results.

Change of attention maps. We calculate the attention maps between reference image features of two objects (cat & dog from Figure 5) and the latent image features \mathbf{Z} in the middle cross-attention layer. As shown in Figure 6, the attention maps of the two objects become more distinct after training, thereby alleviating the object confusion problem.

Furthermore, we provide ablation experiments (e.g., the number of training images) in S2.3 of the appendix.

Conclusion

In this work, we generalize the finetuning-free methods with decoupled cross-attention for merging multiple reference images, by mitigating the object confusion problem. To this end, we propose a weighted-merge method to integrate reference image features with their corresponding objects. Next, we continue to train the pre-trained models on a multi-object dataset constructed with a proposed object quality score. Besides, our training framework can be applied to single-object generation when a single object has multiple reference images. Experiments demonstrate that our method achieves significantly superior performance to existing methods. We hope our method and dataset can contribute to the community of personalized image generation.

Acknowledgements

This work is supported by Alibaba-Zhejiang University Joint Research Institute of Frontier Technologies.

References

- Araujo, A.; Norris, W.; and Sim, J. 2019. Computing Receptive Fields of Convolutional Neural Networks. *Distill*. <https://distill.pub/2019/computing-receptive-fields>.
- Gal, R.; Alaluf, Y.; Atzmon, Y.; Patashnik, O.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2023. An Image is Worth One Word: Personalizing Text-to-Image Generation using Textual Inversion. In *ICLR 2023*. OpenReview.net.
- Gu, Y.; Wang, X.; Wu, J. Z.; Shi, Y.; Chen, Y.; Fan, Z.; Xiao, W.; Zhao, R.; Chang, S.; Wu, W.; Ge, Y.; Shan, Y.; and Shou, M. Z. 2023. Mix-of-Show: Decentralized Low-Rank Adaptation for Multi-Concept Customization of Diffusion Models. In *NeurIPS 2023*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In *NeurIPS 2020*.
- Huang, Q.; Chan, L.; Liu, J.; He, W.; Jiang, H.; Song, M.; and Song, J. 2024a. PatchDPO: Patch-level DPO for Finetuning-free Personalized Image Generation. *arXiv preprint arXiv:2412.03177*.
- Huang, Q.; Song, J.; Hu, J.; Zhang, H.; Wang, Y.; and Song, M. 2024b. On the Concept Trustworthiness in Concept Bottleneck Models. In *AAAI 2024*, 21161–21168. AAAI Press.
- Huang, Q.; Xue, M.; Huang, W.; Zhang, H.; Song, J.; Jing, Y.; and Song, M. 2023. Evaluation and Improvement of Interpretability for Self-Explainable Part-Prototype Networks. In *ICCV 2023*, 2011–2020. IEEE.
- Jaegle, A.; Gimeno, F.; Brock, A.; Vinyals, O.; Zisserman, A.; and Carreira, J. 2021. Perceiver: General Perception with Iterative Attention. In *ICML 2021*, volume 139 of *Proceedings of Machine Learning Research*, 4651–4664. PMLR.
- Jia, X.; Zhao, Y.; Chan, K. C.; Li, Y.; Zhang, H.; Gong, B.; Hou, T.; Wang, H.; and Su, Y.-C. 2023. Taming encoder for zero fine-tuning image customization with text-to-image diffusion models. *arXiv preprint arXiv:2304.02642*.
- Jiang, J.; Zhang, Y.; Feng, K.; Wu, X.; and Zuo, W. 2024. MC²: Multi-concept Guidance for Customized Multi-concept Generation. *arXiv preprint arXiv:2404.05268*.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.; Dollár, P.; and Girshick, R. B. 2023. Segment Anything. In *ICCV 2023*, 3992–4003. IEEE.
- Kumari, N.; Zhang, B.; Zhang, R.; Shechtman, E.; and Zhu, J. 2023. Multi-Concept Customization of Text-to-Image Diffusion. In *CVPR 2023*, 1931–1941. IEEE.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. C. H. 2023. BLIP-2: Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models. In *ICML 2023*, volume 202 of *Proceedings of Machine Learning Research*, 19730–19742. PMLR.
- Liu, M.; Ma, Y.; Zhen, Y.; Dan, J.; Yu, Y.; Zhao, Z.; Hu, Z.; Liu, B.; and Fan, C. 2024. Llm4gen: Leveraging semantic representation of llms for text-to-image generation. *arXiv preprint arXiv:2407.00737*.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Li, C.; Yang, J.; Su, H.; Zhu, J.; et al. 2023a. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*.
- Liu, Z.; Feng, R.; Zhu, K.; Zhang, Y.; Zheng, K.; Liu, Y.; Zhao, D.; Zhou, J.; and Cao, Y. 2023b. Cones: Concept Neurons in Diffusion Models for Customized Generation. In *ICML 2023*, volume 202 of *Proceedings of Machine Learning Research*, 21548–21566. PMLR.
- Luo, W.; Li, Y.; Urtasun, R.; and Zemel, R. S. 2016. Understanding the Effective Receptive Field in Deep Convolutional Neural Networks. In *NeurIPS*, 4898–4906.
- 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 *SIGGRAPH 2024*, 25. ACM.
- 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*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *CVPR 2022*, 10674–10685. IEEE.
- 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 *CVPR 2023*, 22500–22510. IEEE.
- Shi, J.; Xiong, W.; Lin, Z.; and Jung, H. J. 2024. Instant-booth: Personalized text-to-image generation without test-time finetuning. In *CVPR 2024*, 8543–8552. IEEE.
- Wei, Y.; Zhang, Y.; Ji, Z.; Bai, J.; Zhang, L.; and Zuo, W. 2023. ELITE: Encoding Visual Concepts into Textual Embeddings for Customized Text-to-Image Generation. In *ICCV 2023*, 15897–15907. IEEE.
- Xiao, G.; Yin, T.; Freeman, W. T.; Durand, F.; and Han, S. 2023. Fastcomposer: Tuning-free multi-subject image generation with localized attention. *arXiv preprint arXiv:2305.10431*.
- Yang, L.; Yu, Z.; Meng, C.; Xu, M.; Ermon, S.; and Bin, C. 2024. Mastering text-to-image diffusion: Recaptioning, planning, and generating with multimodal llms. In *ICML 2024*, Proceedings of Machine Learning Research. PMLR.
- Ye, H.; Zhang, J.; Liu, S.; Han, X.; and Yang, W. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*.
- Zhang, Y.; Song, Y.; Liu, J.; Wang, R.; Yu, J.; Tang, H.; Li, H.; Tang, X.; Hu, Y.; Pan, H.; et al. 2024. Ssr-encoder: Encoding selective subject representation for subject-driven generation. In *CVPR 2024*, 8069–8078. IEEE.