

IDAdapter: Learning Mixed Features for Tuning-Free Personalization of Text-to-Image Models

Siying Cui^{1,3} Jia Guo^{2*} Xiang An^{2,3} Jiankang Deng² Yongle Zhao³ Xinyu Wei¹ Ziyong Feng³

¹Peking University ²InsightFace ³DeepGlint

¹hycsy@stu.pku.edu.cn

*Corresponding Author: guojia@gmail.com



Figure 1. Given a single facial photo of as the reference and a text prompt, our proposed method can generate images in a variety of styles, angles, and expressions without any test-time fine-tuning at the inference stage. The results exhibit dressing-up modifications, viewpoint control, recontextualization, art renditions, property alteration, as well as emotion integration, while preserving high fidelity to the face.

Abstract

Leveraging Stable Diffusion for the generation of personalized portraits has emerged as a powerful and noteworthy tool, enabling users to create high-fidelity, custom character avatars based on their specific prompts. However, existing personalization methods face challenges, including test-time fine-tuning, the requirement of multiple input images, low preservation of identity, and limited diversity in generated outcomes. To overcome these challenges, we introduce IDAdapter, a tuning-free approach that enhances the diversity and identity preservation in personalized image generation from a single face image. IDAdapter integrates a personalized concept into the generation process through a combination of textual and visual injections and a face identity loss. During the training phase, we incorporate mixed features from multiple reference images of a specific identity to enrich identity-related content details, guiding the model to generate images with more diverse styles, expressions, and angles. Extensive evaluations demonstrate the effectiveness of our method, achieving both diversity and identity fidelity.

1. Introduction

Recently, the field of text-to-image (T2I) synthesis has witnessed significant advancements, especially with the advent of diffusion models. Models such as Imagen proposed by [35], DALL-E2 by [31], and Stable Diffusion by [33] have gained attention for their ability to generate realistic images from natural language prompts. While these models excel in generating complex, high-fidelity images from extensive text-image datasets, the task of generating images of specific subjects from user-provided photos remains a significant challenge.

Personalization in text-to-image (T2I) synthesis has been primarily achieved through methodologies employing pre-trained models, as outlined in works such as [2, 13, 16, 23, 34, 38, 40]. These methods typically involve fine-tuning the models with a set of specific reference images (ranging from 3 to 20). While effective, this approach calls for specialized training of certain network layers, often requiring considerable computational resources and extended processing times on advanced GPUs, which may not be feasible for user-centric applications. An alternative strategy, discussed in studies like [7, 8, 44, 46], involves augmenting pre-trained

Work done when Siying Cui was an intern at DeepGlint.

diffusion models with additional parameters like adapters trained on large personalized image datasets. This approach enables tuning-free conditional generation but typically lacks the fidelity and diversity of fine-tuning methods. For example, as indicated in [6] and [37], this approach often restricts the generated images to the expressions present in the input image, thus limiting the expansive creative potential of diffusion models.

Drawing inspiration from test-time fine-tuning methods utilizing multiple reference images and the adapter series as described in works [26, 43, 47], we introduce IDAdapter. This innovative approach synthesizes features from various images of a person during training, effectively mitigating overfitting to non-identity attributes. IDAdapter operates by freezing the base diffusion model’s primary weights, with under 10 hours of training on a single GPU. During inference, IDAdapter requires only a single reference image and textual prompts to produce diverse, high-fidelity images that maintaining the person’s identity, as depicted in Figure 1. It broadens the range of what the base model can generate, making the results more diverse while preserving identity, which surpasses the limitations of previous models. Our contributions are threefold:

1. We present a method that incorporates mixed features from multiple reference images of the same person during training, yielding a T2I model that avoids the need for test-time fine-tuning.
2. This technique, without test-time fine-tuning, can generate varied angles and expressions in multiple styles guided by a single photo and text prompt, a capability not previously attainable.
3. Comprehensive experiments confirm that our model outperforms earlier models in producing images that closely resemble the input face, exhibit a variety of angles, and showcase a broader range of expressions.

2. Related Work

2.1. Text-to-Image Models

The field of computational image generation has witnessed remarkable advancements due to the evolution of deep generative models for text-to-image synthesis. Techniques like Generative Adversarial Networks (GANs) [22, 45], auto-regressive models [30], and diffusion models [17, 33] have played a crucial role. Initially, these models were limited to generating images under specific domains and textual conditions. However, the introduction of large-scale image-text datasets and advanced language model encoders has significantly improved text-to-image synthesis capabilities. The pioneering DALL-E [30] utilized autoregressive models for creating diverse images from text prompts. This was followed by GLIDE [27], which introduced more realistic and high-resolution images using diffusion mod-

els. Consequently, diffusion models have increasingly become the mainstream method for text-to-image synthesis. Recent developments like DALL-E 2 [31], Imagen [35], and LDM [33] have further enhanced these models, offering more realism, better language understanding and diverse outputs. The success of Stable Diffusion [33] in the open-source community has led to its widespread use and the development of various fine-tuned models. Our methodology, acknowledging this trend, is based on the Stable Diffusion model.

2.2. Personalization via Subject-Driven Tuning

The goal of personalized generation is to create variations of a specific subject in diverse scenes and styles based on reference images. Originally, Generative Adversarial Networks (GANs) were employed for this purpose, as illustrated by [28], who achieved personalization by fine-tuning StyleGAN with around 100 facial images. Subsequently, pivotal tuning [32], which involved fine-tuning latent space codes in StyleGAN, enabled the creation of variant images. However, these GAN-based methods faced limitations in subject fidelity and style diversity. Recent advancements have been made with the Stable Diffusion Model, offering improvements in subject fidelity and outcome diversity. Textual Inversion [13] optimized input text embeddings with a small set of images for subject image generation. The study by [42] enhanced textual inversion to capture detailed subject information. DreamBooth [34] optimized the entire T2I network for higher fidelity. Following this, several methods like CustomDiffusion [23], SVDiff [15], LoRa [1, 19], StyleDrop [39], and the approach by [18] proposed partial optimizations. DreamArtist [12] demonstrated style personalization with a single image. Despite their effectiveness, these methods involve time-consuming multi-step fine-tuning for each new concept, limiting their practicality in real-world applications.

2.3. Tuning-Free Text-to-Image Personalization

A distinct research direction involves training models with extensive domain-specific data, thereby eliminating the need for additional fine-tuning at the inference stage. Facilitating object replacement and style variation, InstructPix2Pix [4] integrates latent features of reference images into the noise injection process. ELITE [44] introduced a training protocol combining global and local mappings, utilizing the OpenImages test set. UMM-Diffusion [25], leveraging LAION-400M dataset [36], proposed a multimodal latent diffusion approach that combines text and image inputs. Several studies, such as UMM [25], ELITE [44], and SuTI [7], have demonstrated subject image generation without fine-tuning. Similarly, Taming-Encoder [21] and InstantBooth [37] focus on human and animal subjects, employing a new conditional branch for diffusion models. FastComposer [46], Face0 [41]

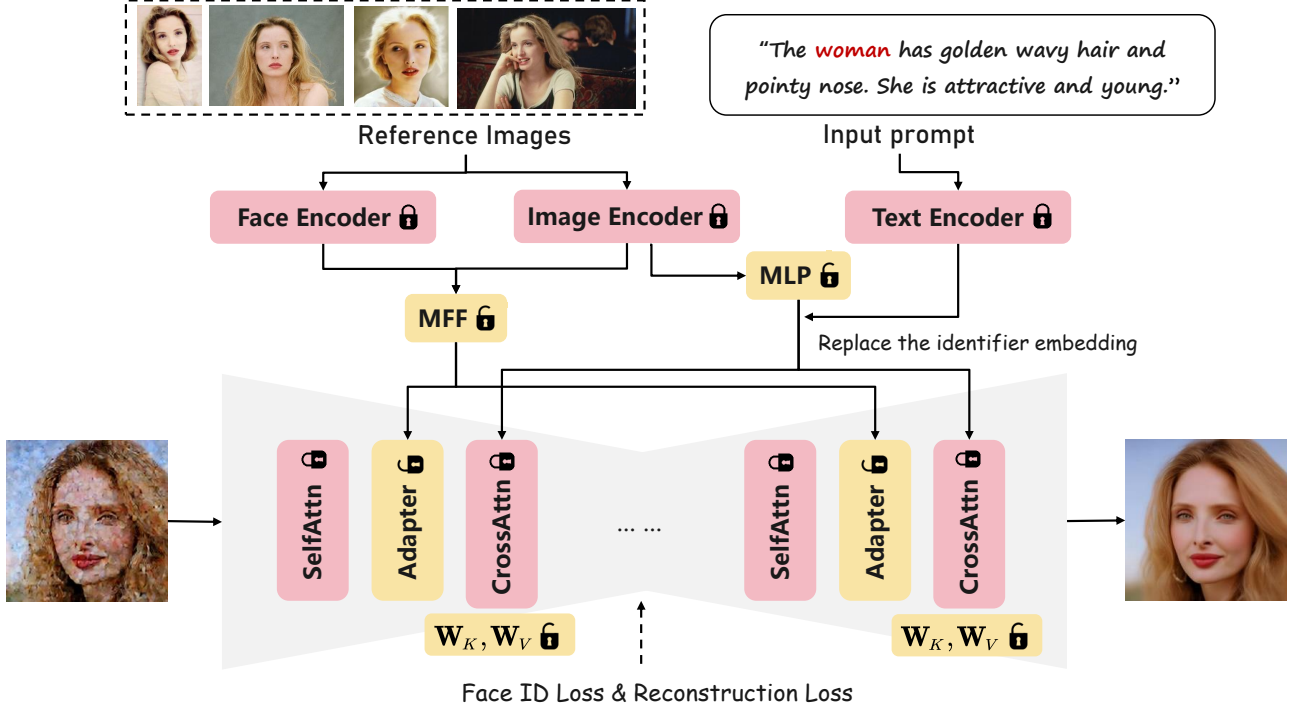


Figure 2. **The overview of IDAdapter training.** In each optimization step, we randomly select N different images of the same identity. We label the faces of all the reference images “[class noun]” (e.g. “woman”, “man”, etc.), and regard the text description and the reference images as a training pair. The features extracted from the reference images are then fused using a mixed facial features (MFF) module, which provides the model with rich detailed identity information and possibilities for variation. At the inference stage, only a single image is required, which is replicated to form a set of N reference images.

and PhotoVerse [6] have also contributed novel approaches in this domain. Despite these advancements, a key challenge remains in balancing ease of use with generation quality and diversity. Our proposed solution, IDAdapter, addresses this issue by coordinating model usability with output quality and diversity.

3. Method

Given only a single face image of a specific person, we intend to generate a range of vivid images of the person guided by text prompts with diversity. Example diversity includes not only adjusting dressing-up, properties, contexts, and other semantic modifications (these attributes are referred to as “styles” in this paper), but generating various facial expressions and poses. We next briefly review the necessary notations of Latent Diffusion Models (Sec. 3.1) as well as the method for simply extracting facial features from a single image (Sec. 3.2), then present our technique to extract mixed facial features from a few images (Sec. 3.3), and finally incorporate it as a new concept to the generator structure by the adapter layers (Sec. 3.4). Fig. 2 shows the overview of our approach based on the Stable Diffusion [33] structure.

3.1. Preliminaries

Text-to-Image (T2I) diffusion models ϵ_θ denoise a noise map $\epsilon \in \mathbb{R}^{h \times w}$ into an image \mathbf{x}_0 based on a textual prompt T . In this work, we utilize Stable Diffusion, a specific instance of Latent Diffusion Model (LDM), which comprises three key components: an image encoder, a decoder, and an iterative UNet denoising network for processing noisy latent representations.

The encoder \mathcal{E} maps an image \mathbf{x}_0 from the pixel space to a low-dimensional latent space $\mathbf{z} = \mathcal{E}(\mathbf{x}_0)$, while the decoder \mathcal{D} reconstructs the latent representation \mathbf{z} back into an image to achieve $\mathcal{D}(\mathcal{E}(\mathbf{x}_0)) \approx \mathbf{x}_0$. The diffusion model incorporates an input text embedding $\mathbf{C} = \Theta(T)$, which is generated using a text encoder Θ and a text prompt T and then employed in the intermediate layers of the UNet through a cross-attention mechanism:

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

where $\mathbf{Q} = \mathbf{W}_Q \cdot \varphi(\mathbf{z}_t)$, $\mathbf{K} = \mathbf{W}_K \cdot \mathbf{C}$, $\mathbf{V} = \mathbf{W}_V \cdot \mathbf{C}$, $\varphi(\mathbf{z}_t)$ is the hidden states through the UNet implementation, d is the scale factor utilized for attention mechanisms. The training goal for the latent diffusion model is to predict

the noise added to the image’s latent space, a formulation denoted as:

$$\mathcal{L}_{SD} = \mathbb{E}_{\mathbf{z} \sim \mathcal{E}(\mathbf{x}_0), \mathbf{C}, \epsilon \sim \mathcal{N}(0,1), t} \left[\|\epsilon - \epsilon_\theta(\mathbf{z}_t, t, \mathbf{C})\|_2^2 \right] \quad (2)$$

where ϵ is the ground-truth noise, \mathbf{z}_t is noisy latent representations at the diffusion process timestep t . See [33] for more details.

3.2. Facial Features

Our objective is to extract facial features from input images, inject them with the stylistic information denoted by text prompts, and generate a rich variety of images with fidelity to the identified facial characteristics. Intuitively, this diversity includes at least the following three aspects: A) Diversity of styles, where the generated images must conform to the styles indicated by the prompts; B) Diversity in facial angles, signifying the capability to produce images of the person from various facial poses; C) Diversity of expressions, which refers to the ability to generate images of the person displaying a range of different expressions or emotions.

An intuitive approach is learning the features of input facial images in the textual space and embedding these features into the generative guiding process of Stable Diffusion, so that we can control the image generation of the person via a specific identifier word. However, as noted by several studies [6, 20, 24, 37], the sole use of textual space embeddings constrains the ultimate quality of generated images. A potential cause for this pitfall could be the limitation of the textual space features in capturing identity (ID) characteristics. Consequently, it becomes imperative to supplement textual conditional guidance with guidance based on image features to augment the image generation capabilities.

We find that both commonly used general CLIP image encoders and feature vector encoders from face recognition networks exhibit a strong binding with non-identity (non-ID) information of the input images, such as facial poses and expressions. This binding results in the generated images lacking diversity at the person level. To address this issue, we propose the Mixed Facial Features module (MFF). This module is designed to control the decoupling of ID and non-ID features during the generation process of the diffusion model, thereby enabling the generation of images with enhanced diversity.

3.3. Mixed Facial Features (MFF)

The core idea behind MFF is to utilize rich detailed information from multiple reference images to help IDAdapter better extract identity features and achieve face fidelity, rather than simply learn from a single face paste. Specifically, we combine the features of N face images $\{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(N)}\}$ with the prompt T to guide the generation of Stable Diffusion, where $\mathbf{x}^{(i)} \in \mathbb{R}^{h \times w \times c}$ for




Figure 3. **Architecture of MFF:** Our MFF consists of a learnable transformer implemented with two attention blocks that translates identity feature \mathbf{f}_a and patch feature \mathbf{f}_v into a latent MFF vision embedding \mathbf{E}_r , which will be injected to the self-attention layers of the UNet through adapters.

$i = 1, \dots, N$, (h, w) is the resolution of $\mathbf{x}^{(i)}$ and c is the number of channels. We illustrate the idea in Figure 3.

Given a reference image set containing K images, we first enrich them to N images $\{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(K)}, \dots, \mathbf{x}^{(N)}\}$ if $K < N$, through various data augmentation operations such as random flipping, rotating, and color transformations. We first encode all of the reference images into visual patch features $\{\mathbf{f}_v^{(1)}, \mathbf{f}_v^{(2)}, \dots, \mathbf{f}_v^{(N)}\}$ and a learned class embedding $\{\mathbf{f}_c^{(1)}, \mathbf{f}_c^{(2)}, \dots, \mathbf{f}_c^{(N)}\}$ using the vision model of CLIP [29], where $\mathbf{f}_v^{(i)} \in \mathbb{R}^{p^2 \cdot c \times d_v}$ and $\mathbf{f}_c^{(i)} \in \mathbb{R}^{1 \times d_v}$. Here, (p, p) is the patch size, and d_v is the dimension of these embedded patches through a linear projection. Then we obtain an enriched patch feature \mathbf{f}_v by concatenating the patch features from all the reference images. We have:

$$\mathbf{f}_v = \text{Concat}(\{\mathbf{f}_v^{(i)}\}_1^N) \quad (3)$$

This enriched feature \mathbf{f}_v is derived from multiple images under the same identity, so their common characteristics (i.e., the identity information) will be greatly enhanced, while others (such as the face angle and expression of any specific image) will be somewhat weakened. Therefore, \mathbf{f}_v can greatly assist in diversifying the generation results as indicated in Sec. 4.3. We find that with $N = 4$, personalization results are strong and maintain identity fidelity, editability and facial variation.

To further guarantee the identity, we encode the faces from all the enriched reference images $\{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(N)}\}$ into identity features $\{\mathbf{f}_a^{(1)}, \mathbf{f}_a^{(2)}, \dots, \mathbf{f}_a^{(N)}\}$ using the face encoder of Arcface face recognition approach [9], where $\mathbf{f}_a^{(i)} \in \mathbb{R}^{1 \times d_a}$ for $i = 1, \dots, N$. Then we calculate the average

feature vector \mathbf{f}_a as an identity feature. We have:

$$\mathbf{f}_a = \sum_{i=1}^N \mathbf{f}_a^{(i)} / N \quad (4)$$

Then we append the identity feature \mathbf{f}_a to the patch feature \mathbf{f}_v as one and embed it into a learnable lightweight transformer $\mathcal{P}_{\text{visual}}$ implemented with two attention blocks as illustrated in Figure 3. We have:

$$\mathbf{E}_r = \mathcal{P}_{\text{visual}}([\mathbf{f}_v, \mathbf{f}_a]) \quad (5)$$

Finally we obtain a MFF vision embedding \mathbf{E}_r , which compresses the facial information and is adapted to the latent space of the diffusion model. The feature \mathbf{E}_r will be injected into the self-attention layers of the UNet through adapters.

3.4. Personalized Concept Integration

Textual Injection In addition to obtaining mixed facial features from the pixel space, we also aim at injecting a new personalized concept into Stable Diffusion’s “dictionary”. We label the faces of all the reference images “[*class noun*]” (e.g. “*woman*”, “*man*”, etc.), which can be specified by the user, and denote “*sks*” as an identifier word. In this paper, we assume that “[*class noun*] is *sks*” is appended to the end of each prompt by default, thereby linking the face features with the identifier word. As mentioned in the approach to generate patch features in Sec. 3.3 using the vision model of CLIP [29], we also obtain a learned class embedding $\{\mathbf{f}_c^{(1)}, \mathbf{f}_c^{(2)}, \dots, \mathbf{f}_c^{(N)}\}$ simultaneously. We adopt their average embedding to map all the reference images to a compact textual concept through a learnable multi-layer perceptron $\mathcal{P}_{\text{textual}}$:

$$\mathbf{E}_c = \mathcal{P}_{\text{textual}}\left(\sum_{i=1}^N \mathbf{f}_c^{(i)} / N\right) \quad (6)$$

where \mathbf{E}_c is the identity text embedding of the reference images, projected from the visual space to the textual space of Stable Diffusion in essence. At the first embedding layer of the text encoder, we replace the text embedding of the identifier word “*sks*” with the identity text embedding \mathbf{E}_c to inject textual personalized concept into the UNet. This final text embedding will be the condition in the cross-attention layers of Stable Diffusion.

Visual Injection We find that the model tends to generate overfitting results (e.g. fixed expressions, poses) if we fine-tune the entire Stable Diffusion since the prior is ruined. This motivates the need for key parameters to learn the personalized concept with the output of MFF. In this regard, some existing research [14, 23] have emphasized the significance of attention layers. Therefore, our approach is to extend the model with trainable adapter layers and optimize query and key matrices $\mathbf{W}_K, \mathbf{W}_V$ in the cross-attention modules.

Specifically, as for the injection of the MFF vision embedding \mathbf{E}_r , we employ a new learnable adapter layer between each frozen self-attention layer and cross-attention layer:

$$\mathbf{y} := \mathbf{y} + \beta \cdot \tanh(\gamma) \cdot S([\mathbf{y}, \mathbf{E}_r]) \quad (7)$$

where, \mathbf{y} is the output of the self-attention layer, S is the self-attention operator, γ is a learnable scalar initialized as 0, and β is a constant to balance the importance of the adapter layer.

Then, by updating the key and value projection matrices in each cross-attention block, the model is able to focus on the visual characteristics of the face and link them with the personalized concept in the textual space.

Face Identity Loss Our experiments will show the diversity of generation achieved by learning mixed face features, which loses the regularization of facial region. However, it gives rise to the problem of identity preservation. Accordingly, we introduce a face identity loss \mathcal{L}_{id} that supervises the model to preserve the identity of reference images. This allows the model to generate diverse appearances, as well as retain the identity feature. Specifically, we utilize a pre-trained face recognition model \mathcal{R} [9]:

$$\mathcal{L}_{\text{id}} = \mathbb{E}_{\hat{\mathbf{x}}_0} [1 - \cos(\mathcal{R}(\hat{\mathbf{x}}_0), \mathbf{f}_a)] \quad (8)$$

where \cos denotes the cosine similarity, $\hat{\mathbf{x}}_0$ is the predicted denoised image sample based on the model output \mathbf{z}_t at the diffusion timestep t , and \mathbf{f}_a is the average identity feature calculated by Equation 4. To prevent an unclear face of image $\hat{\mathbf{x}}_0$ misleading the model, we utilize a face detection model [10] \mathcal{F} . Face identity loss is applied only when a face is detected in $\hat{\mathbf{x}}_0$, i.e., when $\mathcal{F}(\hat{\mathbf{x}}_0) = 1$. It is often not possible to detect a face in $\hat{\mathbf{x}}_0$ with a large timestep t , at which $\mathcal{F}(\hat{\mathbf{x}}_0) = 0$. The loss becomes:

$$\mathcal{L} = \mathcal{L}_{\text{SD}} + \mathcal{F}(\hat{\mathbf{x}}_0) \cdot \lambda \mathcal{L}_{\text{id}} \quad (9)$$

where λ controls for the weight of the face identity loss. Sec. 4.3 shows that face identity loss is effective in preserving output identity. We find that ~ 50000 iterations, $\lambda = 0.1$ and learning rate 3×10^{-5} is enough to train a robustly performing model.

4. Experiments

4.1. Experimental settings

Datasets For our training process, we utilized the comprehensive collection of 30,000 image-text pairings from the Multi-Modal CelebA-HQ database, as detailed in [45]. This dataset includes 6,217 unique identities. To enhance the diversity of our dataset, we implemented various data augmentation techniques. These included random face swapping, utilizing the InsightFace [11] tool, alongside standard

methods such as image flipping, rotation, and color adjustments. For each identity, we ensured the presence of over N augmented images. During each iteration of training, N images per identity were randomly chosen to generate the MFF vision embedding E_r and the corresponding identity text embedding E_c .

For testing quantitative results, we methodically selected one image per individual for a total of 500 individuals from the VGGFace2 dataset [5] as reference for all methods. For the measurement of identity preservation, our prompts for generation were limited to a simple “[class noun]” word such as “woman” or “man”, and for the measurement of diversity, the prompts were a “[class noun]” word combined with an expression word (e.g. “happy”, “sad”, “angry”). It’s noteworthy that all facial imagery used for visualization purposes were acquired from SFHQ dataset [3] or publicly accessible channels.

Implementation Details We utilize Stable Diffusion [33] V2.1 as the base model and fine-tune our IDAdapter at the training stage. We trained the model with Adam optimizer, learning rate of $3e - 5$ and batch size of 4 on a single A100 GPU. Our model was trained for 50,000 steps. At the testing and inference stage, we use only one image and simply duplicate it N times to serve as the input for the network.

Evaluation Metrics A critical aspect in our evaluation is the fidelity of facial identity in the generated images. To quantify this, we calculate the average identity preservation, which is the pairwise cosine similarity between facial features of generated images and their real counterparts (ID-Sim). This calculation is performed using a pretrained face recognition model, as outlined in [9]. Additionally, we have introduced two novel metrics to assess the diversity of the generated images: pose-diversity (Pos-Div) and expression-diversity (Expr-Div).

- **Pose-Diversity (Pose-Div)** This metric assesses the variance in facial angles between the generated image and the input image. To quantify this difference, we calculate the average deviation in facial angles across all test images. To better reflect real-world scenarios, we report the results specifically in terms of Pitch (Pose-Div pitch) and Yaw angles (Pose-Div yaw). This approach enables us to evaluate how well the model can generate images with a range of different facial orientations relative to the input image.
- **Expression-Diversity (Expr-Div)** This metric evaluates the variation in facial expressions between the generated images and the input image. Utilizing a pre-trained expression classification model, we measure the ratio of the generated images having different expression categories compared to the input across the entire test dataset. A higher value in this metric indicates a greater ability of the model to generate diverse facial expressions.

These metrics are crucial for determining the capability of our method to generate images that are not only personalized but also varied in terms of poses and expressions, reflecting a more comprehensive range of human facial appearance.

4.2. Comparisons

Qualitative Results Our methodology was benchmarked against several leading techniques, including Textual Inversion [13], Dreambooth [34], E4T [14], ProFusion [49], and Photoverse [6], as illustrated in Figure 4. The comparative results were sourced directly from the study of [6], where the “S*” in the prompts refers to the “[class noun]” we mentioned. We observe that our method surpasses both Textual Inversion and DreamBooth in terms of face fidelity. Unlike other methods, our approach effectively preserves identity without giving in to overfitting to expressions or poses as Figure 5 and Figure 6 shows, thereby facilitating the generation of images that are both more diverse and lifelike.

Quantitative Results In our quantitative experiments, the capability of IDAdapter was evaluated using three metrics: identity preservation (ID-Sim), pose-diversity (Pose-Div), and expression-diversity (Expr-Div). Moreover, these models demonstrate a lack of proficiency in generating varied facial expressions and poses. Consequently, we assessed Pos-Div and Expr-Div metrics exclusively on open-source models requiring fine-tuning [13, 23, 34, 48]. In this experiment, we have selected the parameter $N = 4$. As depicted in Table 1, our method achieved the highest scores across almost all metrics. It can be observed that IDAdapter effectively leverages the base model to generate more diverse results with identity preserved.

4.3. Ablation Studies

As illustrated in Table 2 and Figure 7, our analysis reveals the impact of different components of the IDAdapter method on the quality of generated images.

Impact of Identity Text Embedding When the identity text embedding component is removed from the process (No Text Embedding), there is a significant decrease in the identity preservation of the generated images. This drastic drop suggests that textual conditions play a crucial role in guiding Stable Diffusion to generate personalized images. Without the identity text embedding, the fundamental feature of personalized generation is almost lost.

Removal of MFF Vision Embedding Eliminating the vision embedding component output by MFF (No MFF) leads to a significant drop of both identity preservation and diversity. This indicates that the MFF module provides the model with rich identity-related content details. MFF is vital for counteracting overfitting and helps retain the ability of the base model to generate diverse images of the person.

Impact of Different N Values Changing the number of




Figure 4. **Comparisons with several baseline methods.** IDAdapter is stronger in the diversity of properties, poses, expressions and other non-ID appearance, achieving very strong editability while preserving identity.



Figure 5. In terms of diversity performance, we compare generated samples of our method, Subject-Diffusion and Photoverse. We observe that our method generally achieves very strong diversity while preserving identity without giving in to overfitting to expressions or poses.



Figure 6. **Visual comparison on expression with more state-of-the-art methods.** IDAdapter is better at generating expressions that match the prompt compared to other methods and retains ID features.

Method	Fine-tuning	Single Image	ID-Sim \uparrow	Expr-Div \uparrow	Pose-Div pitch \uparrow	Pose-Div yaw \uparrow
Ours	N	Y	0.603	65%	7.90	16.47
Profusion [14]	Y	Y	0.454	31%	1.95	2.31
Celeb Basis [48]	Y	Y	0.207	35%	4.92	12.04
DreamBooth [34]	Y	N	0.105	71%	6.93	12.23

Table 1. We compared our IDAdapter ($N = 4$) with several baseline methods in terms of identity preservation (ID-Sim) and diversity performance (Expr-Div, Pose-Div pitch and Pose-Div yaw).

Method	ID-Sim \uparrow	Expr-Div \uparrow	Pose-Div pitch \uparrow	Pose-Div yaw \uparrow
No Text Embedding	0.394	49%	6.08	13.49
No MFF	0.517	46%	5.31	13.26
IDAdapter ($N = 1$)	0.602	37%	5.02	12.90
IDAdapter ($N = 2$)	0.601	58%	6.97	15.39
IDAdapter ($N = 3$)	0.604	61%	7.03	15.44
IDAdapter ($N = 4$)	0.603	65%	7.90	16.47
IDAdapter ($N = 5$)	0.601	64%	7.88	16.42
No ID Loss	0.592	57%	7.64	16.38

Table 2. Ablation studies on identity preservation metric (ID-Sim) and diversity metrics (Expr-Div, Pose-Div pitch and Pose-Div yaw).

images N used in training process has varying impacts on diversity and identity preservation. After testing with different N values, we found that $N = 4$ offers the best balance. It achieves a superior compromise between maintaining the identity similarity and enhancing the diversity. This balance

is crucial for generating images that are both personalized and varied.

Impact of ID Loss We trained IDAdapter ($N = 4$) without face identity loss (No ID Loss). The model’s performance in learning facial features has declined, and the generated faces are not as similar to the input as when incorporating the ID loss.

5. Conclusion

We introduce a method named IDAdapter, which is the first to generate images of a person with a single input facial image in a variety of styles, angles, and expressions without fine-tuning during the inference stage, marking a significant breakthrough in personalized avatar generation.



Figure 7. **Visualization of generated results under different settings.** Fine-tuning without certain model structure can result in a decrease in the performance of identity preservation and diversity, overfitting to input image appearance.

References

- [1] Low-rank adaptation for fast text-to-image diffusion fine-tuning. <https://github.com/cloneofsimon/lora>, 2022. 2
- [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. 1
- [3] David Beniaguev. Synthetic faces high quality (sfhq) dataset. <https://github.com/SelfishGene/SFHQ-dataset>, 2022. 6
- [4] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, 2023. 2
- [5] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *FG*, 2018. 6
- [6] Li Chen, Mengyi Zhao, Yiheng Liu, Mingxu Ding, Yangyang Song, Shizun Wang, Xu Wang, Hao Yang, Jing Liu, Kang Du, et al. Photoverse: Tuning-free image customization with text-to-image diffusion models. *arXiv:2309.05793*, 2023. 2, 3, 4, 6
- [7] Wenhui Chen, Hexiang Hu, Yandong Li, Nataniel Ruiz, Xuhui Jia, Ming-Wei Chang, and William W. Cohen. Subject-driven text-to-image generation via apprenticeship learning. In *NIPS*, 2023. 1, 2
- [8] Wenhui Chen, Hexiang Hu, Chitwan Saharia, and William W. Cohen. Re-imagen: Retrieval-augmented text-to-image generator. In *ICLR*, 2023. 1
- [9] Jiankang Deng, Jia Guo, Xue Niannan, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. 4, 5, 6
- [10] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *CVPR*, 2020. 5
- [11] Jinakang Deng, Jia Guo, Xiang An, Jack Yu, and Baris Gecer. Insightface: 2d and 3d face analysis project. *Github*, 2022. 5
- [12] Ziyi Dong, Pengxu Wei, and Liang Lin. Dreamartist: Towards controllable one-shot text-to-image generation via contrastive prompt-tuning. *arXiv:2211.11337*, 2022. 2
- [13] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit Haim Bermano, Gal Chechik, and Daniel Cohen-or. An image is worth one word: Personalizing text-to-image generation using textual inversion. In *ICLR*, 2023. 1, 2, 6
- [14] 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 Trans Graph*, 2023. 5, 6, 8
- [15] Ligong Han, Yinxiao Li, Han Zhang, Peyman Milanfar, Dimitris Metaxas, and Feng Yang. Svdiff: Compact parameter space for diffusion fine-tuning. In *ICCV*, 2023. 2
- [16] Shaozhe Hao, Kai Han, Shihao Zhao, and Kwan-Yee K Wong. Vico: Detail-preserving visual condition for personalized text-to-image generation. *arXiv:2306.00971*, 2023. 1
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NIPS*, 2020. 2
- [18] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. *ICML*, 2019. 2
- [19] 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 *ICLR*, 2022. 2
- [20] Junha Hyung, Jaeyo Shin, and Jaegul Choo. Magicapture: High-resolution multi-concept portrait customization. *arXiv:2309.06895*, 2023. 4
- [21] Xuhui Jia, Yang Zhao, Kelvin CK Chan, Yandong Li, Han Zhang, Boqing Gong, Tingbo Hou, Huisheng Wang, and Yu-Chuan Su. Taming encoder for zero fine-tuning image customization with text-to-image diffusion models. *arXiv:2304.02642*, 2023. 2
- [22] Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. Scaling up gans for text-to-image synthesis. In *CVPR*, 2023. 2
- [23] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *CVPR*, 2023. 1, 2, 5, 6
- [24] Jian Ma, Junhao Liang, Chen Chen, and Haonan Lu. Subject-diffusion: Open domain personalized text-to-image generation without test-time fine-tuning. *arXiv:2307.11410*, 2023. 4
- [25] Yiyang Ma, Huan Yang, Wenjing Wang, Jianlong Fu, and Jiaying Liu. Unified multi-modal latent diffusion for joint subject and text conditional image generation. *arXiv:2303.09319*, 2023. 2
- [26] Chong Mou, Xintao Wang, Liangbin Xie, Jian Zhang, Zhongqiang Qi, Ying Shan, and Xiao Hu Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv:2302.08453*, 2023. 2
- [27] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, 2021. 2
- [28] Yotam Nitzan, Kfir Aberman, Qiurui He, Orly Liba, Michal Yarom, Yossi Gandelsman, Inbar Mosseri, Yael Pritch, and Daniel Cohen-Or. Mystyle: A personalized generative prior. *TOG*, 2022. 2
- [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 *ICML*, 2021. 4, 5
- [30] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, 2021. 2
- [31] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv:2204.06125*, 2022. 1, 2
- [32] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *TOG*, 2022. 2
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 1, 2, 3, 4, 6

- [34] 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 *CVPR*, 2023. 1, 2, 6, 8
- [35] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *NIPS*, 2022. 1, 2
- [36] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv:2111.02114*, 2021. 2
- [37] Jing Shi, Wei Xiong, Zhe Lin, and Hyun Joon Jung. Instantbooth: Personalized text-to-image generation without test-time finetuning. *arXiv:2304.03411*, 2023. 2, 4
- [38] James Seale Smith, Yen-Chang Hsu, Lingyu Zhang, Ting Hua, Zsolt Kira, Yilin Shen, and Hongxia Jin. Continual diffusion: Continual customization of text-to-image diffusion with c-lora. *arXiv:2304.06027*, 2023. 1
- [39] Kihyuk Sohn, Lu Jiang, Jarred Barber, Kimin Lee, Nataniel Ruiz, Dilip Krishnan, Huiwen Chang, Yuanzhen Li, Irfan Essa, Michael Rubinstein, Yuan Hao, Glenn Entis, Irina Blok, and Daniel Castro Chin. Styledrop: Text-to-image synthesis of any style. In *NIPS*, 2023. 2
- [40] Yoad Tewel, Rinon Gal, Gal Chechik, and Yuval Atzmon. Key-locked rank one editing for text-to-image personalization. *SIGGRAPH*, 2023. 1
- [41] Dani Valevski, Danny Wasserman, Yossi Matias, and Yaniv Leviathan. Face0: Instantaneously conditioning a text-to-image model on a face. *SIGGRAPH*, 2023. 2
- [42] Andrey Voynov, Qinghao Chu, Daniel Cohen-Or, and Kfir Aberman. $p+$: Extended textual conditioning in text-to-image generation. *arXiv:2303.09522*, 2023. 2
- [43] Zhouxia Wang, Xintao Wang, Liangbin Xie, Zhongang Qi, Ying Shan, Wenping Wang, and Ping Luo. Styleadapter: A single-pass lora-free model for stylized image generation. *arXiv:2309.01770*, 2023. 2
- [44] 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 *ICCV*, 2023. 1, 2
- [45] Weihao Xia, Yujiu Yang, Jing-Hao Xue, and Baoyuan Wu. Tedigan: Text-guided diverse face image generation and manipulation, 2021. 2, 5
- [46] Guangxuan Xiao, Tianwei Yin, William T Freeman, Frédo Durand, and Song Han. Fastcomposer: Tuning-free multi-subject image generation with localized attention. *arXiv:2305.10431*, 2023. 1, 2
- [47] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv:2308.06721*, 2023. 2
- [48] Ge Yuan, Xiaodong Cun, Yong Zhang, Maomao Li, Chenyang Qi, Xintao Wang, Ying Shan, and Huicheng Zheng. Inserting anybody in diffusion models via celeb basis. In *NIPS*, 2023. 6, 8
- [49] Yufan Zhou, Ruiyi Zhang, Tong Sun, and Jinhui Xu. Enhancing detail preservation for customized text-to-image generation: A regularization-free approach. *arXiv:2305.13579*, 2023. 6