Infinite-ID: Identity-preserved Personalization via ID-semantics Decoupling Paradigm

Yi Wu^{1*}, Ziqiang Li^{1,2*}, Heliang Zheng¹, Chaoyue Wang^{3**}, and Bin Li^{1**}

¹ University of Science and Technology of China, China
² Nanjing University of Information Science and Technology, China ³ The University of Sydney, Australia wuyi2021@mail.ustc.edu.cn, iceli@mail.ustc.edu.cn, zhenghl@mail.ustc.edu.cn, chaoyue.wang@outlook.com, binli@ustc.edu.cn

Abstract. Drawing on recent advancements in diffusion models for textto-image generation, identity-preserved personalization has made significant progress in accurately capturing specific identities with just a single reference image. However, existing methods primarily integrate reference images within the text embedding space, leading to a complex entanglement of image and text information, which poses challenges for preserving both identity fidelity and semantic consistency. To tackle this challenge, we propose Infinite-ID, an ID-semantics decoupling paradigm for identity-preserved personalization. Specifically, we introduce identity-enhanced training, incorporating an additional image cross-attention module to capture sufficient ID information while deactivating the original text cross-attention module of the diffusion model. This ensures that the image stream faithfully represents the identity provided by the reference image while mitigating interference from textual input. Additionally, we introduce a feature interaction mechanism that combines a mixed attention module with an AdaIN-mean operation to seamlessly merge the two streams. This mechanism not only enhances the fidelity of identity and semantic consistency but also enables convenient control over the styles of the generated images. Extensive experimental results on both raw photo generation and style image generation demonstrate the superior performance of our proposed method.

Keywords: Personalized Text-to-image Generation, Stable Diffusion, Identity-preserved Personalization

1 Introduction

Human photo synthesis [19,38] has experienced notable advancements, particularly with the introduction of large text-to-image diffusion models such as Stable

^{*} First two authors contributed equally to this work.

^{**} Corresponding Authors.



Fig. 1: With just a single reference image, our Infinite-ID framework excels in synthesizing high-quality images while maintaining superior identity fidelity and text semantic consistency in various styles.

Diffusion (SD) [27], Imagen [30], and DALL-E 3 [3]. Benefit from personalized text-to-image generation, recent researches focus on *Identity-preserved personalization*. This specialized area aims to produce highly customized photos that faithfully reflect a specific identity in novel scenes, actions, and styles, drawing inspiration from one or more reference images. This task has garnered considerable attention, leading to the development of numerous applications, including personalized AI portraits and virtual try-on scenarios. In the context of Identity-preserved personalization, the emphasis is placed on maintaining the invariance of human facial identity (ID), requiring a heightened level of detail and fidelity compared to more general styles or objects.

Recent tuning-free methods exhibit promise for large-scale deployment, yet they face a notable challenge in balancing the trade-off between the fidelity of identity representation (ID fidelity) and the consistency of semantic understanding conveyed by the text prompt. This challenge arises due to the inherent entanglement of image and text information. Typically, tuning-free methods extract ID information from reference images and integrate it into the semantic information in two distinct ways. The first type, exemplified by PhotoMaker [17], incorporates text information with ID details in the text embedding space of the text encoder. While this merging approach aids in achieving semantic consistency, it compresses image features into the text embedding space, thereby weakening the ID information of the image and compromising identity fidelity. The second type, demonstrated by IP-Adapter [40], directly injects ID information into the U-Net of the diffusion model through an additional trainable cross-attention module. Although this approach aims to enhance the strength of ID information for improved fidelity, it tends to favor the image branch during training, consequently weakening the text branch and compromising semantic consistency. In summary, existing methods entangle image and text information, resulting in a significant trade-off between ID fidelity and semantic consistency.

To address the entanglement between image and text information, we propose **Infinite-ID**, an innovative approach to personalized text-to-image generation. Our method tackles the trade-off between maintaining high fidelity of identity and ensuring semantic consistency of the text prompt by implementing the ID-semantics decoupling paradigm. Specifically, we adopt identity-enhanced training that introduces an additional image cross-attention module to capture sufficient ID information and deactivate the original text cross-attention module to avoid text interference during training stage. Accordingly, our method can faithfully capture identity information from reference image, significantly improving the ID fidelity. Additionally, we employ a novel feature interaction mechanism that leverages a mixed attention module and an AdaIN-mean operation to effectively merge text information and identity information. Notably, our feature interaction mechanism not only preserves both identity and semantic details effectively but also enables convenient control over the styles of the generated images (as depicted in Fig. 1). Our contributions are summarized as:

- We propose a novel ID-semantics decoupling paradigm to resolve the entanglement between image and text information, acquiring a remarkable balance between ID fidelity and semantic consistency in *Identity-preserved personalization*.
- We propose a novel feature interaction mechanism incorporating a mixed attention module and an AdaIN-mean operation to effectively merge ID information and text information and also conveniently control the styles of the generated image in diffusion models.
- Experimental results demonstrate the excellent performance of our proposed method as compared with current state-of-the-art methods on both raw photo generation and style image generation.

2 Related Works

2.1 Text-to-image Diffusion Models

Text-to-image diffusion models, such as those explored in [3, 13, 27, 30, 43, 45], have garnered significant attention due to their impressive image generation capabilities. Current research endeavors aim to further enhance these models along multiple fronts, including the utilization of high-quality and large-scale datasets [18, 20, 32, 33], refinements to foundational architectures [3, 27, 30], and advancements in controllability [12, 29, 45]. Present iterations of text-to-image diffusion models typically follow a two-step process: first, encoding the text prompt using pre-trained text encoders such as CLIP [24] or T5 [26], and then utilizing the resulting text embedding as a condition for generating corresponding images through the diffusion process. Notably, the widely adopted Stable Diffusion model [27] distinguishes itself by executing the diffusion process in latent space instead of the original pixel space, leading to significant reductions in computation and time costs. An important extension to this framework is the Stable Diffusion XL (SDXL) [23], which enhances performance by scaling

up the U-Net architecture and introducing an additional text encoder. Thus, our proposed method builds upon the SDXL. However, our method can also be extended to other text-to-image diffusion models.

2.2 Identity-preserved personalization

Identity-preserved personalization aims to generate highly customized photos that accurately reflect a specific identity across various scenes, actions, and styles, drawing inspiration from one or more reference images. Initially, tuningbased methods, exemplified by DreamBooth [29], Textual Inversion [8] and Celeb-Basis [41], employ images of the same identity (ID) to fine-tune the model. While these methods yield results with high fidelity in preserving facial identity (ID). a significant drawback emerges: the customization of each ID necessitates a time investment of 10-30 minutes [15], consuming substantial computing resources and time. This limitation poses a significant obstacle to large-scale deployment in commercial applications. Consequently, recent advancements in *tuning-free*. methods [6, 7, 9, 16, 17, 21, 35, 36, 39, 40] have been introduced to streamline the generation process. These methods specifically leverage the construction of a vast amount of domain-specific data and the training of an encoder or hyper-network to represent input ID images as embeddings or LoRA weights within the model. Post-training, users need only input an image of the ID for customization, enabling personalized generation within seconds during the inference phase. These tuning-free methods typically contain two distinct manners.

On one hand, methods [1, 17, 22, 39] incorporate text information alongside identity details within the text embedding space of the text encoder. For example, PhotoMaker [17] extracts identity embeddings from single or multiple reference images and merges them with corresponding class embeddings (e.g., "man" and "woman") in the text embedding space of the text encoder. While this stacking operation aids in achieving semantic consistency, it compresses image features into the text embedding space, leading to compromised identity fidelity. On the other hand, some studies [6, 36, 40] directly integrate identity information into the U-Net of the diffusion model. IP-Adapter [40] distinguishes itself by incorporating a additional cross-attention layer for each existing crossattention layer within the original UNet model. This approach merges identity information with semantic details directly within the U-net but leads to distortion of the semantic space of the U-Net model. Consequently, this compromises the semantic consistency of the text prompt.

In summary, existing methods entangle the identity and text information, leading to a significant trade-off between ID fidelity and semantic consistency. To mitigate these limitations, we propose an identity-enhanced training to capture ID and text information separately. Moreover, we design an effective feature interaction mechanism leveraging a mixed attention module and an AdaIN-mean operation to preserve both identity and semantic details while also enabling convenient control over the styles of generated images.

2.3 Attention Control in Diffusion model

Previous studies have investigated various attention control techniques within diffusion models. Hertz *et al.* [10] employed a shared attention mechanism, concatenating and applying an AdaIN module on the key and value between reference and synthesis images within the self-attention layer to ensure styleconsistent image generation using a reference style. Cao *et al.* [5] utilized a mutual self-attention approach to achieve consistent image generation and non-rigid image editing, wherein the key and value of the synthesis image were replaced with those of the reference image within the self-attention layers of the diffusion model. Similarly, Shi *et al.* [34] proposed a method termed reference attention, enabling consistent multi-view generation of target objects by concatenating the key and value features between the condition signal and the synthesis image in the self-attention layers. Wang *et al.* [39] and Avrahami *et al.* [2] exploited attention maps within the cross-attention layers to guide the optimization process towards disentangling learned concepts in personalized generation tasks.

3 Method

3.1 Preliminaries

Stable Diffusion XL. Our method builds upon Stable Diffusion XL [23], comprising three core components: a Variational AutoEncoder (VAE) denoted as $\xi(\cdot)$, a conditional U-Net [28] represented by $\epsilon_{\theta}(\cdot)$, and two pre-trained text encoders [25] denoted as $\Theta_1(\cdot)$ and $\Theta_2(\cdot)$. Specifically, given a training image x_0 and its corresponding text prompt T, the VAE encoder ξ transforms x_0 from its original space $\mathbb{R}^{H \times W \times 3}$ to a compressed latent representation $z_0 = \xi(x_0)$, where $z_0 \in \mathbb{R}^{h \times w \times c}$ and c denotes the latent dimension. Subsequently, the diffusion process operates within this compressed latent space to conserve computational resources and memory. Once the two text encoders process the text prompt Tinto a text embedding $c = \text{Concat}(\Theta_1(T), \Theta_2(T))$, the conditional U-Net ϵ_{θ} predicts the noise ϵ based on the current timestep t, the t-th latent representation z_t , and the text embedding c. The training objective is formulated as follows:

$$L_{\text{diffusion}} = E_{z_t, t, c, \epsilon \in N(0, I)}[||\epsilon - \epsilon_{\theta}(z_t, t, c)||_2^2].$$

$$\tag{1}$$

Attention mechanism in diffusion models. The fundamental unit of the stable diffusion model comprises a resblock, a self-attention layer, and a cross-attention layer. The attention mechanism is represented as follows:

$$\operatorname{Attn}(Q, K, V) = \operatorname{Softmax}(\frac{QK^T}{\sqrt{d}})V, \qquad (2)$$

where Q denotes the query feature projected from the spatial features generated by the preceding resolution, K and V represent the key and value features projected from the same spatial features as the query feature (in self-attention) or the text embedding extracted from the text prompt (in cross-attention).



Fig. 2: Framework of ID-semantics Decoupling Paradigm. In the training phase, we adopt Face embeddings extractor to extract rich identity information and identity-enhanced training for faithfully representing the identity provided by the reference image while mitigating interference from textual input. In the inference stage, a mixed attention module is introduced to replace the original self-attention mechanism within the denoising U-Net model, facilitating the fusion of both identity and text information.

3.2 Methodology

Overview. In this section, we introduce our ID-semantics decoupling paradigm, as illustrated in Fig. 2, which effectively addresses the severe trade-off between high-fidelity identity and semantic consistency within identity-preserved personalization. Subsequently, we present our mixed attention mechanism, depicted in Fig. 3, designed to seamlessly integrate ID information and semantic information within the diffusion model during the inference stage. Additionally, we utilize an adaptive mean normalization (AdaIN-mean) operation to precisely align the style of the synthesized image with the desired style prompts.

ID-semantics Decoupling Paradigm. To faithfully capture high-fidelity identity, we implement a novel identity-enhanced strategy during the training stage, as depicted in Fig. 2. Diverging from conventional methods [6,17,40] that utilize text-image pairs for training, we opt to exclude the text prompt input and deactivate cross-attention modules for text embeddings within the U-Net model. Instead, we establish a training pair consisting of an ID image, where the face is aligned to extract identity information, and denoising image (x_t in Fig. 2). Both the denoising image used for training and the ID image belong to the same individual, but they vary in factors such as viewpoints and facial expressions. This approach fosters a more comprehensive learning process [17]. We adopt Face embeddings extractor to accurately capture and leverage identity information from

7



Fig. 3: Mixed attention mechanism. On the left side, we employ mixed attention to fuse identity and text information. This involves concatenating their respective key and value features and subsequently applying mixed attention, where identity features are updated based on the concatenated key and value features. On the right side, for style merging, we introduce an additional AdaIN-mean operation (as depicted in Eq. 8) to the concatenated key and value features.

the input ID image. Additionally, to seamlessly integrate identity information into the denoising U-Net model, we introduce an extra trainable cross-attention mechanism (K'_{id} and V'_{id} in Fig. 2) for image embeddings.

Throughout the training phase, we exclusively optimize the parameters associated with the face mapper, CLIP mapper, and the image cross-attention module, while keeping the parameters of the pre-trained diffusion model fixed. The optimization loss closely resembles the original diffusion loss formulation (as delineated in Eq. 1), with the sole distinction being the shift from a text condition to an identity condition as the conditional input.

$$L_{\text{diffusion}} = E_{z_t, t, c_{id}, \epsilon \in N(0, I)} [||\epsilon - \epsilon_{\theta}(z_t, t, c_{id})||_2^2], \tag{3}$$

where the c_{id} is the identity embeddings of the input ID image.

During the inference phase, we align the desired face within the ID image to afford identity information. Following the approach adopted during the training phase, we utilize a face recognition backbone and a CLIP image encoder to extract identity features from the aligned face image. Leveraging the trained face mapper, clip mapper, and image cross-attention mechanisms, these identity features are seamlessly integrated into the denoising U-Net model. Subsequently, we compute the key and value features for self-attention in the original stable diffusion model, considering only the text prompt input. These text key and value features are instrumental in the mixed attention process (illustrated in Fig. 3), facilitating the fusion of text and identity information. Moreover, to further augment the text information during the denoising process, we incorporate the original text cross-attention mechanism, integrating the resulting text hidden states with the output image hidden states obtained from the image cross-attention module.

Face Embeddings Extractor. We adopt a multifaceted approach by incorporating pre-trained models to extract facial features. Firstly, following the methodologies of prior research, we utilize a pre-trained CLIP image encoder as one of our facial feature extractors. Specifically, we leverage the local embeddings, comprising the last hidden states obtained from the CLIP image encoder, forming a sequence of embeddings with a length of N (N=257 in our implementation). Subsequently, we employ a CLIP mapper to project these image embeddings (with a dimensionality of 1664) to the same dimension as the text features in the pre-trained diffusion model. As elucidated in [40], the features extracted by the CLIP image encoder are instrumental in capturing the structural information pertinent to the identity face within identity-preserved personalization tasks. Additionally, we leverage the backbone of a face recognition model as another facial feature extractor. As highlighted in [40], features extracted by the face recognition backbone are adept at capturing the characteristics associated with human facial features within identity-preserved personalization tasks. More specifically, we utilize the global image embedding derived from the extracted features and subsequently employ a face mapper to align the dimensionality (512 dimensions) of the extracted global image embedding with the dimensionality of the text features in the pre-trained diffusion model.

In summary, the identity embeddings c_{id} corresponding to the ID image x can be expressed as:

$$c_{id} = \text{Concat}\left(\text{M}_{\text{clip}}(\text{FA}(x))), \text{M}_{\text{face}}(\text{FA}(x))\right), \qquad (4)$$

where $Concat(\cdot, \cdot)$, $M_{clip}(\cdot)$, $E_{clip}(\cdot)$, $M_{face}(\cdot)$, $E_{face}(\cdot)$, and $FA(\cdot)$ denote the concatenation function, CLIP mapper, CLIP image encoder, face mapper, face recognition backbone, and face alignment module [44], respectively.

Mixed Attention Mechanism. As explored in previous studies [5, 14, 37], the features in the self attention layers play a crucial role in consistency image generation (across-frame in text-to-video works), which indicates these features provide a refined and detailed semantics information. In our study, we extract features from the self-attention layers of the original text-to-image diffusion model to capture rich semantic information, represented as K_t and V_t . We enhance the self-attention mechanism by incorporating it into a mixed attention framework, depicted in Fig. 3 (a). This fusion enables the integration of semantic features (K_t and V_t) with the identity-based features (K_{id} and V_{id}), thereby encapsulating identity information. Through this integration, the mixed attention mechanism seamlessly merges semantic details into the generated features

across different resolutions. The formulation of the mixed attention mechanism is as follows:

$$\begin{aligned} \operatorname{Attn}_{\operatorname{mix}}(Q, K, V) &\triangleq \operatorname{Attn}(Q, \hat{K}, \hat{V}) \\ \text{w.r.t} \quad \hat{K} = \operatorname{Concat}(K_{id}, K_t), \quad \hat{V} = \operatorname{Concat}(V_{id}, V_t), \\ K_{id} = W_k^{id} Z_{id}, \quad K_t = W_k^t Z_t, \\ V_{id} = W_v^{id} Z_{id}, \quad V_t = W_v^t Z_t, \end{aligned}$$
(5)

where Z_{id} and Z_t are the corresponding spatial features of generated features and semantic features, respectively. The parameters W_k^{id} , W_k^t , W_v^{id} , and W_v^t correspond to the weights of the corresponding fully connected layers.

Cross-attention Merging. To further refine semantic control, we incorporate text features into the identity feature within the cross-attention layers using the following formulation:

$$\operatorname{Attn}_{\operatorname{cross}}(Q, K, V) \triangleq \operatorname{Attn}(Q, K'_{id}, V'_{id}) + \operatorname{Attn}(Q, K'_t, V'_t), \tag{6}$$

where $K'_{id} = \hat{W}^{id}_k c_{id}$, $V'_{id} = \hat{W}^{id}_v c_{id}$, $K'_t = \hat{W}^t_k c_t$, and $V'_t = \hat{W}^t_v c_t$. c_{id} and c_t are the identity embedding and text embedding, respectively. \hat{W}^{id}_k , \hat{W}^{id}_v , \hat{W}^t_k , and \hat{W}^t_v correspond to the weights of the trainable fully connected layers within cross-attention module.

Style Information Merging. Inspired by [10], we propose an adaptive mean normalization (AdaIN-mean) operation to further align the style of the synthesis image with the style prompts. Concretely, we align the key and value features projected from identity features in both mixed attention and cross-attention with the key and value features projected from text features, formulated as follows:

$$K_{id} = \text{AdaIN-m}(K_{id}, K_t), \quad V_{id} = \text{AdaIN-m}(V_{id}, V_t), \quad \text{For Mixed Attention}$$
$$K'_{id} = \text{AdaIN-m}(K'_{id}, K'_t), \quad V'_{id} = \text{AdaIN-m}(V'_{id}, V'_t), \quad \text{For Cross Attention}$$
(7)

where the AdaIN-mean operation (AdaIN-m(\cdot)) is defined as:

AdaIN-m
$$(x, y) = x - \mu(x) + \mu(y),$$
(8)

where $\mu(x) \in \mathbb{R}^{d_k}$ is the mean of key and value features across different pixels. The mixed attention with AdaIN-mean has been illustrated in Fig. 3 (b).

4 Experiments

After outlining the experimental setup in Sec. 4.1, we conduct a comparative analysis of raw photo generation and style image generation in Sec. 4.2. Ablation studies, highlighting the significance of various components, are presented in Sec. 4.3. Additionally, experiments involving multiple input ID images including identity mixing and identity linear interpolation are detailed in Sec. A.7 of the supplementary materials. For further insights, Sec. A.8 and Sec. A.9 of the supplementary materials provide additional qualitative results for raw photo generation and style photo generation, respectively.



Fig. 4: Qualitative comparison on raw photo generation. The results demonstrate that our Infinite-ID consistently maintains identity fidelity and achieves highquality semantic consistency with just a single input image.

4.1 Experimental Setup

Implementation Details. Our experiments leverage a pre-trained Stable Diffusion XL (SDXL) model [23]. For image encoding, we utilize the OpenCLIP ViT-H/14 [42] and the backbone of ArcFace [11]. The SDXL model consists of 70 cross-attention layers, to each of which we append an additional image crossattention module. Training is conducted on 16 A100 GPUs for 1 million steps, with a batch size of 4 per GPU. We employ the AdamW optimizer with a fixed learning rate of 1e-4 and weight decay set to 0.01. During inference, we employ the DDIM Sampler with 30 steps and guidance scale is set to 5.0. Training data are sourced from multiple datasets, including the LAION-2B dataset [33], the LAION-Face dataset [46], and few high-quality images collected from the Websites⁴. We curate a dataset where each individual is represented by multiple photographs.

Evaluation. We assess the efficacy of our approach in preserving both identity fidelity and semantic consistency. Specifically, we measure identity fidelity, utilizing metrics such as M_{FaceNet} (measured by FaceNet [31]) and CLIP-I [8]. Identity fidelity is evaluated based on the similarity of detected faces between the

⁴ https://www.pexels.com, https://pixabay.com.



Fig. 5: Qualitative comparison on style image generation. The results demonstrate that our method maintains strong identity fidelity, high-quality semantic consistency, and precise stylization using only a single reference image.

reference image and generated images. Semantic consistency is quantified using CLIP text-image consistency (CLIP-T [24]), which compares the text prompt with the corresponding generated images. More details of evaluation are shown in Sec. A.1 and Sec. A.2 of the supplementary materials.

4.2 Comparison to Previous Methods

Raw Photo Generation. We benchmark our Infinite-ID against established identity-preserving personalization approaches. The qualitative outcomes are illustrated in Fig. 4, while quantitative results are provided in Table 1. Notably, all methods are tuning-free, necessitating no test-time adjustments. FastComposer [39] exhibits challenges in maintaining identity fidelity, often presenting undesired artifacts in synthesized images. While IP-Adapter [40] and IP-Adapter-Face [40] demonstrates relatively fewer artifacts, its semantic consistency fall short. This phenomenon arises from the direct fusion of identity information with semantic details within the U-net model, leading to a compromise in semantic consistency. In contrast, PhotoMaker [17] exhibits commendable semantic consistency but falls short in preserving identity fidelity. Leveraging our IDsemantics decoupling paradigm, our method excels in preserving identity fidelity. Furthermore, our mixed attention mechanism effectively integrate the semantic **Table 1: Quantitative comparison.** The evaluation metrics encompass CLIP-T, CLIP-I, and M_{FaceNet} . Our approach outperforms other methods in terms of identity fidelity while simultaneously achieving satisfactory semantic consistency. The best result is shown in **bold**, and the second best is <u>underlined</u>. Additionally, the quantitative comparison of ablation studies have been shown in the gray part.

	CLIP-T \uparrow	CLIP-I \uparrow	$M_{\rm FaceNet}$ \uparrow
FastComposer [39]	0.292	0.887	0.556
IP-Adapter [40]	0.274	0.905	0.474
IP-Adapter-Face [40]	0.313	0.919	0.513
PhotoMaker [17]	0.343	0.814	0.502
Ours	<u>0.340</u>	<u>0.913</u>	0.689
Ours (w/o identity-enhanced training)	0.329	0.891	0.593
Ours (w/o Mixed Attention)	0.331	0.905	0.700
Ours (Mixed Attention \Rightarrow Mutual Attention)	0.316	0.808	0.398

information into the denoising process, positioning our method favorably against existing techniques. More comparison results are shown in Sec. A.4, Sec. A.5 and Sec. A.6 of the supplementary materials.

Style Image Generation. We demonstrate the results of the stylization results in Fig. 5, which compare our method with state-of-the-art tuning-free identitypreserving personalization methods. The style of the synthesis images include anime style, comic book style, and line art style. According to the stylization results, the IP-Adapter [40] and IP-Adapter-Face fails to depict the desired style in the prompt and the style of generation results always obey the tone in the reference image. The training pipeline of the IP-Adapter and IP-Adapter-Face entangles the text embedding and image embedding which leads to the distortion of the text embeddings space. FastComposer [39] also fails in stylization generation and shows undesired artifacts. PhotoMaker [17] achieves a better semantic consistency and stylization, but the identity fidelity is still unsatisfactory. In contrast, our method achieves high identity fidelity, appealing semantic consistency, and precise stylization meantime.

4.3 Ablation Study

In this section, we begin by conducting ablation studies to assess the influence of our identity-enhanced training and mixed attention mechanism. Furthermore, we conduct style image generation to evaluate the effectiveness of our AdaINmean operation in regulating the style of the generated images. More ablation studies are demonstrated in Sec. A.3 of the supplementary materials.

Ablation of identity-enhanced training. In Fig. 6 and Table 1, we compare our method with "Ours (w/o identity-enhanced training)", which is implemented using an identity-semantics entangled training strategy. This strategy utilizes text-image pairs during training and activates cross-attention modules for text embeddings within the original U-Net model. It is noteworthy that both methods



Fig. 6: Ablation study of our identity-enhanced training and mixed attention (M-A). It is evident that identity-enhanced training significantly improves the identity fidelity, and mixed-attention mechanism enhances semantic consistency compared to mutual attention (MU-A) approach [5].



Fig. 7: Ablation study of our AdaIN-mean operation. The results show that AdaIN-mean play a crucial role in style image generation. Compare to AdaIN [10] module, our AdaIN-mean helps to achieve a higher identity fidelity.

share the same inference processing. The qualitative comparison demonstrates that our identity-enhanced training notably enhances identity fidelity.

Ablation of mixed attention mechanism. To assess the effectiveness of our proposed mixed attention (M-A) mechanism, we compare our method with "Ours (w/o M-A)" and "Ours (M-A \Rightarrow MU-A)" in Fig. 6 and Table 1. Mutual self-attention (MU-A) [5] converts the existing self-attention into 'cross-attention' for consistent image editing, where the crossing operation happens in the self-attention sof two related diffusion processes. The results show that our mixed attention mechanism demonstrates superior ability to improve semantic consistency while maintaining identity fidelity.

Ablation of AdaIN-mean operation. To assess the effectiveness of our proposed AdaIN-mean operation, we compare our Infinite-ID model with variations, namely Ours (w/o AdaIN-mean) and Ours (AdaIN-mean \Rightarrow AdaIN), as depicted

in Fig. 7. The results reveal that: i) Ours (w/o AdaIN-mean) exhibits superior ID fidelity compared to Infinite-ID but fails to achieve style consistency with the text prompt; ii) Both our AdaIN-mean and AdaIN modules successfully achieve style consistency with the text prompt, yet AdaIN-mean maintains better ID fidelity than AdaIN. In conclusion, our proposed AdaIN-mean operation facilitates precise stylization while concurrently preserving ID fidelity.

5 Conclusion and limitations.

In this paper, we introduce Infinite-ID, an innovative identity-preserved personalization method designed to meet the requirement of identity (ID) fidelity and semantic consistency of the text prompt, all achievable with just one reference image and completed within seconds. Infinite-ID comprises three key components: the identity-enhanced training, mixed attention mechanism, and adaptive mean normalization (AdaIN-mean). Through extensive experimentation, our results illustrate that Infinite-ID outperforms baseline methods, delivering strong ID fidelity, superior generation quality, and precise semantic consistency in both raw photo generation and style image generation tasks. However, it's important to note that our method lacks multi-object personalization capability. Moreover, artifacts may occur when the human face occupies only a small portion of the entire image, attributed to limitations inherent in the original diffusion model. Furthermore, the image encoders employed in the identity-preserved personalization remain to be further explored.

6 Copyright of Training Data and Ethical Considerations.

Copyright of Training Data. We construct our training dataset from the open-source public image dataset, including LAION-2B dataset, the LAION-Face dataset and few high-quality images collected from the Website⁵ which supports the CC0 (Creative Commons Zero) 1.0 Universal license⁶.

Ethical Considerations. We recognize the ethical considerations about improper use of such high-quality human photo synthesis techniques. Therefore, we emphasize the significance of establishing and complying to ethical guidelines [4] and using this technology responsibly. We hope that our work will drive further discussions and research on the security and ethical use of the human photo generation in computer vision. Moreover, we will employ the watermark and safety checker (NSFW classifiers) to attempt to avoid inappropriate content generation.

 $^{^5}$ https://www.pexels.com, https://pixabay.com.

⁶ The person who associated a work with this deed has dedicated the work to the public domain by waiving all of his or her rights to the work worldwide under copyright law, including all related and neighboring rights, to the extent allowed by law. You can copy, modify, distribute and perform the work, even for commercial purposes, all without asking permission.

Acknowledgements

This research was funded by the National Natural Science Foundation of China (U19B2044), supported by the Fundamental Research Funds for the Central Universities (WK2100000045).

References

- Achlioptas, P., Benetatos, A., Fostiropoulos, I., Skourtis, D.: Stellar: Systematic evaluation of human-centric personalized text-to-image methods. arXiv preprint arXiv:2312.06116 (2023)
- Avrahami, O., Aberman, K., Fried, O., Cohen-Or, D., Lischinski, D.: Break-a-scene: Extracting multiple concepts from a single image. arXiv preprint arXiv:2305.16311 (2023)
- Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., Ouyang, L., Zhuang, J., Lee, J., Guo, Y., et al.: Improving image generation with better captions. Computer Science. https://cdn. openai. com/papers/dall-e-3. pdf 2, 3 (2023)
- 4. Biden, J.R.: Executive order on the safe, secure, and trustworthy development and use of artificial intelligence (2023)
- Cao, M., Wang, X., Qi, Z., Shan, Y., Qie, X., Zheng, Y.: Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. arXiv preprint arXiv:2304.08465 (2023)
- Chen, L., Zhao, M., Liu, Y., Ding, M., Song, Y., Wang, S., Wang, X., Yang, H., Liu, J., Du, K., et al.: Photoverse: Tuning-free image customization with text-to-image diffusion models. arXiv preprint arXiv:2309.05793 (2023)
- Chen, X., Huang, L., Liu, Y., Shen, Y., Zhao, D., Zhao, H.: Anydoor: Zero-shot object-level image customization. arXiv preprint arXiv:2307.09481 (2023)
- Gal, R., Alaluf, Y., Atzmon, Y., Patashnik, O., Bermano, A.H., Chechik, G., Cohen-Or, D.: An image is worth one word: Personalizing text-to-image generation using textual inversion. arXiv preprint arXiv:2208.01618 (2022)
- Gal, R., Arar, M., Atzmon, Y., Bermano, A.H., Chechik, G., Cohen-Or, D.: Encoder-based domain tuning for fast personalization of text-to-image models. ACM Transactions on Graphics (TOG) 42(4), 1–13 (2023)
- Hertz, A., Voynov, A., Fruchter, S., Cohen-Or, D.: Style aligned image generation via shared attention. arXiv preprint arXiv:2312.02133 (2023)
- Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: 2017 IEEE International Conference on Computer Vision (ICCV) (Oct 2017). https://doi.org/10.1109/iccv.2017.167, http: //dx.doi.org/10.1109/iccv.2017.167
- Inoue, N., Kikuchi, K., Simo-Serra, E., Otani, M., Yamaguchi, K.: Layoutdm: Discrete diffusion model for controllable layout generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10167– 10176 (2023)
- Kawar, B., Zada, S., Lang, O., Tov, O., Chang, H., Dekel, T., Mosseri, I., Irani, M.: Imagic: Text-based real image editing with diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6007–6017 (2023)
- Khachatryan, L., Movsisyan, A., Tadevosyan, V., Henschel, R., Wang, Z., Navasardyan, S., Shi, H.: Text2video-zero: Text-to-image diffusion models are zeroshot video generators. arXiv preprint arXiv:2303.13439 (2023)

- Kumari, N., Zhang, B., Zhang, R., Shechtman, E., Zhu, J.Y.: Multi-concept customization of text-to-image diffusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1931–1941 (2023)
- Li, D., Li, J., Hoi, S.C.: Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. arXiv preprint arXiv:2305.14720 (2023)
- Li, Z., Cao, M., Wang, X., Qi, Z., Cheng, M.M., Shan, Y.: Photomaker: Customizing realistic human photos via stacked id embedding. arXiv preprint arXiv:2312.04461 (2023)
- Li, Z., Usman, M., Tao, R., Xia, P., Wang, C., Chen, H., Li, B.: A systematic survey of regularization and normalization in gans. ACM Computing Surveys 55(11), 1–37 (2023)
- Li, Z., Wang, C., Zheng, H., Zhang, J., Li, B.: Fakeclr: Exploring contrastive learning for solving latent discontinuity in data-efficient gans. In: European Conference on Computer Vision. pp. 598–615. Springer (2022)
- Li, Z., Xia, P., Tao, R., Niu, H., Li, B.: A new perspective on stabilizing gans training: Direct adversarial training. IEEE Transactions on Emerging Topics in Computational Intelligence 7(1), 178–189 (2022)
- Ma, J., Liang, J., Chen, C., Lu, H.: Subject-diffusion: Open domain personalized text-to-image generation without test-time fine-tuning. arXiv preprint arXiv:2307.11410 (2023)
- Peng, X., Zhu, J., Jiang, B., Tai, Y., Luo, D., Zhang, J., Lin, W., Jin, T., Wang, C., Ji, R.: Portraitbooth: A versatile portrait model for fast identity-preserved personalization. arXiv preprint arXiv:2312.06354 (2023)
- Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. arXiv preprint arXiv:2307.01952 (2023)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
- Radford, A., Kim, J., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Amanda, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. Cornell University arXiv,Cornell University - arXiv (Feb 2021)
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J.: Exploring the limits of transfer learning with a unified text-to-text transformer. The Journal of Machine Learning Research 21(1), 5485–5551 (2020)
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
- Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention-MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18. pp. 234–241. Springer (2015)
- Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: 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. pp. 22500–22510 (2023)

- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic textto-image diffusion models with deep language understanding. Advances in Neural Information Processing Systems 35, 36479–36494 (2022)
- Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 815–823 (2015)
- Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: Laion-5b: An open largescale dataset for training next generation image-text models. Advances in Neural Information Processing Systems 35, 25278–25294 (2022)
- Schuhmann, C., Vencu, R., Beaumont, R., Kaczmarczyk, R., Mullis, C., Katta, A., Coombes, T., Jitsev, J., Komatsuzaki, A.: Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. arXiv preprint arXiv:2111.02114 (2021)
- 34. Shi, R., Chen, H., Zhang, Z., Liu, M., Xu, C., Wei, X., Chen, L., Zeng, C., Su, H.: Zero123++: a single image to consistent multi-view diffusion base model. arXiv preprint arXiv:2310.15110 (2023)
- 35. Wang, Q., Bai, X., Wang, H., Qin, Z., Chen, A.: Instantid: Zero-shot identitypreserving generation in seconds. arXiv preprint arXiv:2401.07519 (2024)
- Wei, Y., Zhang, Y., Ji, Z., Bai, J., Zhang, L., Zuo, W.: Elite: Encoding visual concepts into textual embeddings for customized text-to-image generation. arXiv preprint arXiv:2302.13848 (2023)
- Wu, J.Z., Ge, Y., Wang, X., Lei, S.W., Gu, Y., Shi, Y., Hsu, W., Shan, Y., Qie, X., Shou, M.Z.: Tune-a-video: One-shot tuning of image diffusion models for textto-video generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7623–7633 (2023)
- Wu, Y., Li, Z., Wang, C., Zheng, H., Zhao, S., Li, B., Tao, D.: Domain remodulation for few-shot generative domain adaptation. Advances in Neural Information Processing Systems 36 (2024)
- Xiao, G., Yin, T., Freeman, W.T., Durand, F., Han, S.: Fastcomposer: Tuningfree multi-subject image generation with localized attention. arXiv preprint arXiv:2305.10431 (2023)
- Ye, H., Zhang, J., Liu, S., Han, X., Yang, W.: Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. arXiv preprint arXiv:2308.06721 (2023)
- 41. Yuan, G., Cun, X., Zhang, Y., Li, M., Qi, C., Wang, X., Shan, Y., Zheng, H.: Inserting anybody in diffusion models via celeb basis. In: Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S. (eds.) Advances in Neural Information Processing Systems. vol. 36, pp. 72958-72982. Curran Associates, Inc. (2023), https://proceedings.neurips.cc/paper_files/paper/2023/file/ e6d37cc5723e810b793c834bcb6647cf-Paper-Conference.pdf
- 42. Zhai, X., Puigcerver, J., Kolesnikov, A., Ruyssen, P., Riquelme, C., Lucic, M., Djolonga, J., Pinto, A.S., Neumann, M., Dosovitskiy, A., et al.: A large-scale study of representation learning with the visual task adaptation benchmark. arXiv preprint arXiv:1910.04867 (2019)
- Zhang, C., Zhang, C., Zhang, M., Kweon, I.S.: Text-to-image diffusion model in generative ai: A survey. arXiv preprint arXiv:2303.07909 (2023)
- Zhang, K., Zhang, Z., Li, Z., Qiao, Y.: Joint face detection and alignment using multitask cascaded convolutional networks. IEEE signal processing letters 23(10), 1499–1503 (2016)

18 Yi Wu et al.

- Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3836–3847 (2023)
- 46. Zheng, Y., Yang, H., Zhang, T., Bao, J., Chen, D., Huang, Y., Yuan, L., Chen, D., Zeng, M., Wen, F.: General facial representation learning in a visual-linguistic manner. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18697–18709 (2022)

A Supplementary Materials

A.1 Training and evaluation datasets.

Our training dataset comprises 200K images, categorized into 20K distinct identities via face recognition. For evaluation, similar to PhotoMaker, which constructs an evaluation dataset containing 25 celebrity identities, we curated a dataset of 24 diverse identities (as shown in Fig. 8): 12 celebrities and 12 ordinary individuals randomly sampled from FFHQ dataset. Evaluation scenarios feature individuals in different costumes and settings like Times Square, living rooms, and beaches, engaging in activities such as holding flowers or running, totally including 30 prompts.



Fig. 8: Evaluation dataset. The dataset contains 24 diverse identities: 12 identities collected independently and 12 identities randomly sampled from FFHQ dataset. These evaluation identities encompass various races, skin tones, appearances and genders, including celebrities and ordinary people.

A.2 Evaluation metrics

Evaluation of face similarity. To assess face similarity, we utilize the face alignment module $FA(\cdot)$, the face recognition backbone $E_{face}(\cdot)$, and the CLIP image encoder $E_{clip}(\cdot)$ to compute the metrics $M_{FaceNet}$ and CLIP-I. Specifically, for each generated image I_{gen} and its corresponding identity image I_{id}, we first employ the FA(\cdot) module to detect the face. Subsequently, we calculate the pairwise identity similarity using $E_{face}(\cdot)$ and $E_{clip}(\cdot)$, respectively:

$$M_{\text{FaceNet}} = \cos(\text{E}_{\text{face}}(\text{FA}(\text{I}_{\text{gen}})), \text{E}_{\text{face}}(\text{FA}(\text{I}_{\text{id}}))),$$

$$\text{CLIP-I} = \cos(\text{E}_{\text{clip}}(\text{FA}(\text{I}_{\text{gen}})), \text{E}_{\text{clip}}(\text{FA}(\text{I}_{\text{id}}))),$$
(9)

where $cos(\cdot, \cdot)$ is the cosine similarity function. Furthermore, in order to visualize (as shown in Fig. 9) the quantitative comparison in Table 1 of the main paper, we integrate both M_{FaceNet} and CLIP-I by utilizing z-score normalization:

$$mean(z-score(M_{FaceNet}), z-score(CLIP-I)),$$
 (10)

where z-score $(x) = (x - \mu)/\sigma$, μ and σ are the average and standard deviation of the x, respectively.

Definition of semantic consistency. We adopt the CLIP-T metric to assess semantic consistency. Specifically, for a generated image I_{gen} paired with its



Fig. 9: Visualization of the quantitative comparison in Table 1 of the main paper. The compared methods including IP-Adapter, IP-Adapter-Face, FastComposer, PhotoMaker, and ablation versions of our method including w/o identity-enhanced training (Ours-1), w/o mixed attention (Ours-2) and mixed attention \Rightarrow mutual attention (Ours-3).

Table 2: Quantitative ablation of cross-attention merge. The metrics includes CLIP-T (higher is better) measuring the semantic consistency, CLIP-I (higher is better) and M_{FaceNet} (higher is better) which are both reflect the identity fidelity. The best result is shown in **bold**.

	CLIP-T↑	CLIP-I↑	$M_{\rm FaceNet}$
$\hline Ours \ w/o \ Cross-attention \ merge$	0.335	0.910	0.681
Ours	0.340	0.913	0.689

corresponding prompt P, we compute the CLIP-T metric utilizing both the CLIP image encoder E_{clip} and the CLIP text encoder E_{text} :

$$CLIP-T = cos(E_{clip}(I_{gen}), E_{text}(P)),$$
(11)

where the $cos(\cdot, \cdot)$ is the cosine similarity function.

A.3 More Results on Ablation Study

Ablation Study of Cross-attention Merge. We conduct ablation experiments on the cross-attention merge to evaluate its effectiveness. As depicted in Fig. 10 and Table 2, the incorporation of cross-attention merge demonstrates improvement in semantic consistency.

Ablation Study of Input ID Images' Resolution. We perform an ablation study on the resolution of input ID images to assess the robustness of our method. Specifically, we utilize images with varying resolutions while maintaining the



Fig. 10: Quantitative ablation of cross-attention merge. It is obvious that the cross-attention merge helps to improve the semantic consistency.



Fig. 11: Ablation study of input ID images' resolution. The identity fidelity slightly drops along with the lower image resolution and the semantic consistency is stable for all the resolution. Our method is robust to the resolution of input ID image.

same text prompt for personalization. As illustrated in Fig. 11, the identity fidelity exhibits only a marginal decrease with decreasing image resolution, while semantic consistency remains stable across all resolutions. In conclusion, our method demonstrates robustness to changes in input image resolution.

A.4 Comparison with DreamBooth and InstantID.

We compare our method with DreamBooth, representing optimization-based methods and tuning-free methods representing InstantID, as shown in Fig. 12 and Table 3. In our comparisons for raw and style photo generations, Dream-Booth excels in semantic consistency but struggles with identity fidelity and stylization. It often requires multiple training images to prevent training collapse, which can occur with a single ID image input. To address this, we use early-stop and train the LoRA weights for 300 iterations in DreamBooth. InstantID excels in identity fidelity but lacks semantic consistency due to overfitting to reference images from IdentityNet, and exhibits unrealistic style in raw photo generation.



Fig. 12: Comparison with DreamBooth and InstantID on raw photo generation and stlye photo generation. Our method consistently maintains identity fidelity and achieves high-quality semantic consistency with just a single input image in raw photo generation. Our method also achieves precise stylization.

Table 3: Quantitative comparison. We utilize CLIP-T metric to reflect the semantic consistency. The CLIP-I and M_{FaceNet} reflect the identity fidelity. FID score and LPIPS score are employed to reflect the quality and diversity of the generated faces. We also compare the speed of different methods.

	CLIP-T (\uparrow)	CLIP-I (\uparrow)	$M_{\rm FaceNet}$ (\uparrow)	FID (\downarrow)	LPIPS (\downarrow)	Speed (\downarrow)
PhotoMaker	0.34	0.81	0.50	266.1	0.664	$pprox 10 { m s}$
InstantID	0.30	0.85	0.70	269.1	0.725	$\underline{\approx}20s$
DreamBooth	0.38	0.80	0.48	268.5	0.677	$\approx 360 \mathrm{s}$
Ours	0.34	0.91	<u>0.68</u>	216.6	0.637	$\approx 20s$

While combining ControlNet and additional pose image inputs can improve semantic consistency for InstantID, this approach is not aligned with our study's use of only one reference image. In contrast, our method achieves high identity fidelity, semantic consistency, and precise stylization. Moreover, we also conduct the user study among different methods in the identity fidelity, image quality, text fidelity (semantic consistency) and stylization, as shown in Table 4.

A.5 Evaluation of the diversity and quality of the generated faces.

We utilize the FID score and LPIPS score to assess the quality and diversity of the generated faces, respectively. As shown in Table 3, our Infinite-ID surpasses other methods in both the diversity and quality of the generated faces. **Table 4: User study.** Users preferences on Identity fidelity (identity), image quality (quality), text fidelity (semantics) and stylization (style). The best result is shown in **bold**, and the second best is underlined.

	Identity	Quality	Semantics	Identity	Quality	Semantics	Style
PhotoMaker	12.5	<u>31.2</u>	25.0	27.5	25.0	<u>32.5</u>	37.5
InstantID	41.2	25.0	6.2	36.2	27.5	7.5	11.2
DreamBooth	7.5	0.0	22.5	3.8	2.5	12.5	2.5
Ours	38.8	43.8	46.3	32.5	45.0	47.5	48.8

(a) Raw photo generation

(b) Style photo generation

A.6 Inference time and storage space.

Compared to the original text-to-image pipeline, our method requires more storage space and inference time. Notably, Infinite-ID does not necessitate two U-Net models. Instead, we integrate an additional cross-attention module, consisting of two single-layer fully connected layers, into the original U-Net. This allows Infinite-ID to toggle between handling text and identity information by activating or deactivating the image cross-attention. Moreover, our method is comparable to other personalization methods in terms of storage space and inference time (as shown in Table 3).

A.7 Identity mixing

Upon receiving multiple images from distinct individuals, we stack all the identity embeddings to merge corresponding identities, as depicted in Fig. 13. The generated image can well retain the characteristics of different IDs, which releases possibilities for more applications. Additionally, by adjusting the interpolation of the identity embeddings, we can regulate the similarity between the generated identity and different input identities, as demonstrated in Fig. 14.

A.8 More Qualitative Results of Raw Photo Generation

Fig. 16 demonstrates the ability of our method to extract identity information from artworks while preserving identity for personalization purposes. Additionally, Fig. 15 illustrates the capability of our method to alter attributes of the extracted identities for raw photo generation. Additional visual samples for raw photo generation are provided in Fig. 17 and Fig. 18, showcasing identities of ordinary individuals sampled from the FFHQ dataset, spanning diverse races, skin tones, and genders.

A.9 More Qualitative Results of Style Photo Generation

Fig. 19 and Fig. 20 display the results of style photo generation. The identity samples consist of ordinary individuals randomly selected from the FFHQ dataset. A total of 12 stylization styles are employed, affirming the generalizability of our method.



Fig. 13: Identity mixing. When receiving multiple input ID images from different individuals, our method can mix these identities by stacking all the identity embeddings.



man facing and smiling to camera wearing a brown clothing in the snow

Fig. 14: Linear interpolation of different identities.



Fig. 15: Applications on attribute change.



Fig. 16: Applications on artworks to raw photo.



Reference



a photo of a woman wearing a brown clothing, in the snow



a photo of a woman wearing a yellow shirt, in the living room



Reference





Reference





Reference





Fig. 17: Raw photo generation. These identities are ordinary people sampled from FFHQ dataset, including various races, skin colors, male and female.



Fig. 18: Raw photo generation. These identities are ordinary people sampled from FFHQ dataset, including various races, skin colors, male and female.



Fig. 19: More visual examples for stylization. These identities are ordinary people sampled from FFHQ dataset, including various races, skin colors, male and female.



Fig. 20: More visual examples for stylization. These identities are ordinary people sampled from FFHQ dataset, including various races, skin colors, male and female.