## PreciseControl: Enhancing Text-To-Image Diffusion Models with Fine-Grained Attribute Control

Rishubh Parihar\*1, Sachidan<br/>and VS\*1, Sabariswaran Mani^\*\*, Tejan Karmali^1, and R. Venkatesh Babu<br/>1 $\,$ 

<sup>1</sup> Vision and AI Lab, IISc Bangalore <sup>2</sup> IIT Kharagpur Project Page

Abstract. Recently, we have seen a surge of personalization methods for text-to-image (T2I) diffusion models to learn a concept using a few images. Existing approaches, when used for face personalization, suffer to achieve convincing inversion with identity preservation and rely on semantic text-based editing of the generated face. However, a more finegrained control is desired for facial attribute editing, which is challenging to achieve solely with text prompts. In contrast, StyleGAN models learn a rich face prior and enable smooth control towards fine-grained attribute editing by latent manipulation. This work uses the disentangled W+ space of StyleGANs to condition the T2I model. This approach allows us to precisely manipulate facial attributes, such as smoothly introducing a smile, while preserving the existing coarse text-based control inherent in T2I models. To enable conditioning of the T2I model on the W+ space, we train a latent mapper to translate latent codes from  $\mathcal{W}$ + to the token embedding space of the T2I model. The proposed approach excels in the precise inversion of face images with attribute preservation and facilitates continuous control for fine-grained attribute editing. Furthermore, our approach can be readily extended to generate compositions involving multiple individuals. We perform extensive experiments to validate our method for face personalization and fine-grained attribute editing.

Keywords: Personalised Image Generation · Fine-grained editing

#### 1 Introduction

Recent personalization methods [10, 37] for large text-to-image (T2I) diffusion models [36, 39] aim to learn a new concept (e.g., your pet) given a few input images. The learned concept is then generated using text prompts in novel contexts (e.g. diverse backgrounds and poses) and styles, thus controlling coarse aspects of an image. Personalization of human portraits [51] is especially interesting due to the wide range of applications in entertainment and advertising.

<sup>\*</sup> Equal contribution

<sup>\*\*</sup> Work done during internship at VAL, IISc



Fig. 1: Given a single portrait image, we embed the subject into a text-to-image diffusion model for personalized image generation. The embedded subject can then be transformed or placed in a novel context using text conditioning. The proposed method can also compose multiple learned subjects with high fidelity and identity preservation. To obtain precise inversion of face, we condition the T2I model on the rich W+ latent space of StyleGAN2. This enables our method to additionally perform fine-grained control over the generated face with continuous control over facial attributes such as age and beard.

However, embedding faces into a generative model has its unique challenges, including faithful inversion of the subject's identity along with its fine facial features. More importantly, smooth control over facial attributes is crucial for precise editing of generated faces, which is challenging to achieve with only text (e.g., continuous increase in smile in Fig.1).

Advancements in StyleGAN models [17, 18] have enabled the generation of highly realistic face images by learning a rich prior over face images. Further, these models have semantically meaningful and disentangled W+ latent space [41] that enable fine-grained attribute control in the generated images [1,13,31]. However, as these models are domain-specific and trained only on faces, they are limited to editing and generating cropped portrait images.

This raises the following question - How can we combine the generalized knowledge from T2I models with the face-specific knowledge from StyleGAN models? Such a framework will enjoy benefits from both groups, enabling coarse control with text and fine-grained attribute control through latent manipulation in the generation process. In this work, we propose a novel approach to combine these two categories of models by conditioning the T2I model with W+ space from StyleGAN2. Conditioning on the W+ provides a natural way for embedding faces in T2I model by projecting them into W+ space using existing StyleGAN2 encoders [44]. The design of having W+ as the inversion bottleneck has two

major advantages: 1) excellent inversion of a face with precise reconstruction of attributes, and 2) explicit control over facial attributes for fine-grained attribute editing. To the best of our knowledge, this is the first work to demonstrate the combination of two powerful generative models StyleGANs and T2I diffusion models for controlled generation.

To condition the T2I model on the W+ space, we train a latent adaptor a lightweight MLP, conditioned on the diffusion process's denoising timestep. It takes a latent code  $w \in W^+$  of a face as input to predict a pair of time-dependent token embeddings that represent the input face and are used to condition the diffusion model cross-attention. We observe that having a different embedding for each timestep provides more expressivity to the inversion process. The latent adaptor is trained on a dataset of (image, w) latent pairs, guided by identity loss, a class regularization loss, and standard denoising loss. To further improve the inversion quality, we perform a few iterations of subject-specific U-Net tuning on the given input image using LoRA [15]. The embedded subject can then be edited in two ways: i) coarse semantic edits using text (for e.g., changing the layout and background) and ii) fine-grained attribute edits by latent manipulation in  $\mathcal{W}+$ (for e.g., smooth interpolation through a varying range of smiles, ages). Some example edits are provided in Fig. 1. Our method generalizes the fine-grained attribute edits from cropped faces (in StyleGANs) to in-the-wild and stylized face images generated by T2I diffusion model.

The proposed method can be easily extended for multiple-person generation, which requires high fidelity identity of all the subjects (Fig. 1). We first predict separate token embeddings for each person and then perform subject-specific tuning to obtain personalized models. However, training a single personalized model for multiple subjects results in the problem of *attribute mixing* (Fig. 5) between faces, where attributes from one face are mixed with another. Instead, we learn separate subject-specific LoRA models, which are then jointly inferred with a chained diffusion process. The intermediate outputs for these processes are merged using an instance segmentation mask after each denoising step. This framework resolves *attribute-mixing* among subjects and preserves the identity from the fine-tuned models. The attributes of individual subjects can be edited in a fine-grained manner in W+ while preserving the other subjects' attributes as shown in Fig. 1-(Attribute Editing).

We perform extensive experiments for embedding single and multiple subjects in StableDiffusion [36] model. Compared to the existing personalization method, the proposed method is extremely efficient and achieves a good tradeoff between identity preservation and text alignment. Next, we present results for fine-grained attribute editing with continuous control in the W+ latent space and compare performance with existing editing methods. Finally, we compare the results for composing multiple subjects and attribute editing of individual subjects. In summary, our primary contributions are as follows:

– First approach to combine large text-to-image models with StlyeGAN2 by conditioning T2I on rich W+ latent space

- 4 R. Parihar *et al.* 
  - Effective personalization method using a single portrait image enabling finegrained attribute editing in W+ space and coarse editing with text prompts.
  - Novel approach to fuse multiple personalized models with chained diffusion processes for multi-person composition.

## 2 Related Work

**Text-based image generation.** Large text-to-image diffusion models [33, 36, 39] achieve excellent image generation performance when trained on internet scale captioned image datasets [40]. These models are scaled to high resolution by learning cascaded diffusion models [33, 34, 39] that generate low-resolution images followed by upsampling. Another promising approach is to train diffusion models in the compressed latent space of a pretrained autoencoder [36].

**Personalization** aims to embed a concept in T2I model, given a few input images. One group of methods optimize for object-specific token embeddings [2, 10,51] via optimization. These approaches preserve text editability, however, they struggle to preserve identity. Another direction is based on fine-tuning diffusion model with strong regularization to avoid overfitting [19, 22, 37]. The third set of methods [11, 38, 45, 48–50, 53, 54] learns a shared domain-specific encoder for faster inversion by leveraging the class-specific features.

**Embedding faces.** Recently embedding human faces in T2I models has received a lot of attention [8,11,49,51,53] as the generic personalization methods [10,37] often fail to faithfully embed human faces. Celeb-basis [51] learns a basis of the celebrity names in the token embedding space. The weights of these basis vectors are then predicted by an encoder model applied to the input image. Profusion [53] proposes a regularization-free encoder-based approach. Photoverse [8] applies a dual branch conditioning in text and image domains for faster and more accurate inversion of faces. Although these methods are able to achieve good inversion, they do not allow for fine-grained attribute control. A concurrent work [24] aims to map the W+ space to the T2I model, however, their method is limited in preserving identity.

**Image editing.** Trained T2I models serve as strong image priors and enable various image editing and restoration applications [5, 25, 29, 47]. For fine-grained image editing [14, 28, 29] localizes the object in the image space using attention masks and only allows editing in the specified region. However, these methods rely on text to change the localized object, which does not allow for fine-grained control. Promising approaches like [4, 12] provide a finer control by interpolation in the noise space or training special sliders per attribute. Another set of works explores the intermediate feature space of unconditional diffusion models to obtain finer attribute control in generation [23, 26]. However, they are limited to editing of generated images and not personalized subjects. We take inspiration from GAN models to attain fine-grained attribute control for real subjects by leveraging their disentangled and smooth latent spaces [16, 41]. This enables precise attribute editing through latent manipulation [1, 27, 31, 41] and we aim to embed these properties in pretrained T2I models.

### 3 Method

#### 3.1 Preliminaries

**Text-to-image Diffusion Models.** This work uses StableDiffusion-v2.1 [36] as a representative Text-to-image (T2I) diffusion model. Stable diffusion is based on the latent diffusion model, which applies the diffusion process in the latent space. Its training involves two stages: a) training a VAE or VQ-VAE autoencoder to map images to a compressed latent space, and b) training a diffusion model in its latent space conditioned on text for guiding the generation. This framework disentangles the learning of fine-grained details in the autoencoder and semantic features in the diffusion model, resulting in easier scaling.

Style-based GANs, [6,17,18] have been widely adapted to generate realistic object-specific images such as faces. Further, these models have disentangled latent space, which enables smooth interpolation between images and fine-grained attribute editing [31,41]. These properties are induced by mapping the Gaussian latent space to a learned latent space W/W+ with a mapper network. Further, GAN encoder models [35,44] can encode and edit real images that invert a given image into W+ space, allowing for fine-grained editing of real images.

#### 3.2 Overview

While the T2I models trade off the diversity in generation with an attribute-rich latent space, our goal is to condition the T2I model with an attribute-rich - W+ space from StyleGAN2, that allows for disentangled and fine-grained control over the face attributes in the generated image. To condition the T2I model on  $\mathcal{W}+$ , we augment the T2I model with a learnable latent adaptor network  $\mathcal{M}$ that projects a latent code  $w \in \mathcal{W}+$ into the text embedding space. For embedding a new subject, we pass it through a pre-trained StyleGAN2 encoder  $\mathcal{E}_{GAN}$  [44] to obtain w latent code, which is passed through  $\mathcal{M}$  to obtain the corresponding textembedding as shown in Fig. 2a). Conditioning on  $\mathcal{W}+$  enables fine-grained attribute control in the generated image by latent manipulation. In the next sections, we discuss the details of the proposed latent adaptor, model training, and fine-grained attribute editing.



Fig. 2: Framework for personalization. Given a single portrait image, we extract its w latent representation from encoder  $\mathcal{E}_{GAN}$ . The latent w along with diffusion timestep t are passed through the latent adaptor  $\mathcal{M}$  to generate a pair of time-dependent token embeddings  $(v_t^1, v_t^2)$  representing the input subject. Finally, the token embeddings are combined with arbitrary prompts to generate customized images.

#### 3.3 Latent adaptor $\mathcal{M}$

We implement the latent adaptor  $\mathcal{M}$  as a shallow MLP network that maps the w latent code from StyleGAN to the token embedding space of the T2I model for any human face image as input. We learn two token embeddings  $(v^1, v^2)$  to represent a human subject as it is known to improve the embedding quality [51]. To extract the timestep-specific semantic information from the latent  $\mathbf{w}$ , we condition  $\mathcal{M}$  on the diffusion timestep t, as diffusion models represent the semantic hierarchy in a timestep-wise fashion [30]. The output of  $\mathcal{M}$  is a set of pair of embeddings allow for a richer representation space and improve identity preservation (shown in Fig. 12). The complete architecture of  $\mathcal{M}$  is shown in Fig.2b). The input t is first passed through positional encodings [43] and the flattened w latent code is passed through a self-attention layer to get relevant features. The encoded representations are then concatenated before passing through a set of linear layers. The obtained pair embedding  $(v_t^1, v_t^2)$  represents the person and is then passed at  $t^{th}$  denoising time-step in the U-Net for generation.

#### 3.4 Training

We perform a two-stage training, where we first pretrain latent adaptor  $\mathcal{M}$  on a face dataset, followed by an few iterations of subject-specific training of  $\mathcal{M}$  and diffusion U-Net with low-rank updates for improving identity as detailed below. **Pretraining.** The mapper  $\mathcal{M}$  is pretrained with a paired dataset  $\mathcal{D}_w$  consisting of (I, w) pairs, where I is a portrait face image and w is its corresponding latent code obtained as  $\mathcal{E}_{GAN}(I)$ . During training, we sample a pair (I, w) and a denoising timestep  $t \in (1, T)$  which are passed through  $\mathcal{M}$  to obtain the pair of token embeddings  $(v_t^1, v_t^2)$  corresponding to the input subject. We place the sampled tokens along with the neutral prompt - y = A photo of  $a \dots$  person' and pass through the text encoder to obtain the final text embeddings  $c(y(\mathcal{M}(t, w)))$ . We add the noise from the noise schedule at t to the image I and train diffusion loss along with additional regularization losses shown in Eq.1 to train  $\mathcal{M}$ . In this stage of training, all the modules -  $\mathcal{E}_{GAN}$ , text-encoder, and U-Net are frozen, except for  $\mathcal{M}$  which is a shallow MLP, making the training compute efficient.

Subject specific training. In the second stage, we finetune encoder  $\mathcal{M}$  and U-Net for a few iterations with the single input image. Specifically, we perform low-rank weight updates (LoRA [15])



Fig. 3: Delayed identity injection results in better text editability.

on the U-Net projection matrices and fine-tune  $\mathcal{M}$ , using the combined loss from Eq. 1. In LoRA training, the model weights are updated as  $\mathbf{W} = \mathbf{W} + \alpha \Delta W$ , where  $\Delta W$  is the learned low-rank residual weights. The hyper-parameter  $\alpha$ controls the extent of fine-tuning and allows for a trade-off between identity preservation and text editability. This second stage of low-rank tuning improves the subjects' identity without hurting the text editability (Fig. 12). Loss function. We train latent adaptor with a combination of denoising diffusion loss  $\mathcal{L}_{Diffusion}$  and regularization loss  $\mathcal{L}_{reg}$  following [11]. The diffusion loss enforces text-to-image consistency, and regularization loss ensures that the predicted token embedding is close to the token embedding of the superclass  $v_{cls}$ , such as *face*. Additionally, we add identity loss  $\mathcal{L}_{ID}$  defined as the MSE between the face recognition embeddings from [9] to preserve the identity during inversion. The final loss is computed as a linear combination of these losses:

$$Diffusion = E_{z,y,\epsilon,t}[||\epsilon - \epsilon_{\theta}(z, c(y(\mathcal{M}(t, w))))||_{2}^{2}]$$

$$\mathcal{L}_{reg} = ||\mathcal{M}(t, w) - v_{cls}||_{2}^{2}$$

$$\mathcal{L}_{ID} = ||\mathcal{E}_{ID}(x_{t}) - \mathcal{E}_{ID}(I)||_{2}^{2}$$

$$\mathcal{L} = \mathcal{L}_{Diffusion} + \lambda_{reg}\mathcal{L}_{reg} + \lambda_{ID}\mathcal{L}_{ID}$$
(1)

where  $\lambda_{reg}$  and  $\lambda_{ID}$  are hyper-parameters and  $\mathcal{E}_{ID}$  is pretrained face-recognition model [9]. To compute  $\mathcal{L}_{id}$  at the intermediate denoising step, we use DDIM [42] approximation of the clean image  $\hat{x}_0$  and pass it to the face detector.

#### 3.5 Inference

During inference, given a single image I, we obtain its token embedding as  $(v_t^1, v_t^2) = \mathcal{M}(\mathcal{E}_{GAN}(I), t)$  for all the time steps  $t \in (1,T)$ . These embeddings can be added with text prompts to generate a novel composition of the learned subject. The image generation process in diffusion models follows a hierarchical structure, where the layout is formed in the first few steps, followed by the formation of object shape and appearance [29]. As our primary aim is to embed a subject's identity, we inject the obtained token embedding only after a time threshold  $(t < \tau)$  to not hurt the layout generated during the initial timesteps.

 $\mathcal{L}$ 

For the initial denoising timesteps  $(t > \tau)$ , we use a celebrity name as a placeholder in the prompt, e.g. 'A



Fig. 4: Fine-grained attribute editing. We map the given input image into w latent code, which is shifted by a global linear attribute edit direction to obtain edited latent code w\*. The edited latent code w\* is then passed through the T2I model to obtain fine-grained attribute edits. The scalar edit strength parameter  $\beta$  can be changed to obtain continuous attribute control.

photo of <u>Brad Pitt</u> as a star wars character' as the model generates improved image layouts when prompted popular subjects. Empirically, we observe that the generations are not sensitive to the text celebrity name used, and it acts as a placeholder, so we fix a single celebrity name, and there is no overlap of the identity with the dataset used for evaluation. This delayed injection of the learned embedding improves text alignment, and passing the predicted token embeddings to all the timesteps results in poor compositions where the model output is a cropped face, as shown below in Fig. 3.

#### 3.6 Fine-grained control over face attributes

Once trained, the latent adaptor  $\mathcal{M}$  bridges between disentangled and smooth  $\mathcal{W}+$  latent space and the text-conditioning of the diffusion model. This enables the *transfer* of latent attribute editing methods that function in the  $\mathcal{W}+$  space of StyleGANs [13,31,41] to the diffusion model. Specifically, for a given source image  $I_s$ , we first obtain its corresponding w latent code with  $\mathcal{E}_{GAN}$ . Next, we edit the latent code w by adding a global linear attribute edit direction d with scalar weight  $\beta$  to obtain  $\hat{w} = w + \beta d$ . Note, the same global edit direction d generalizes for all the identities in the  $\mathcal{W}+$  space [41]. The edited latent code  $\hat{w}$ 

is then passed through  $\mathcal{M}$  to obtain the edited token embedding  $\hat{v}_t$ . Notably, one can precisely control the strength of the attribute edit by changing the scalar  $\beta$  as shown in Fig. 9. To preserve the scene layout during editing, we use the same starting noise and copy the self-attention maps obtained during the generation with unedited w similar to [29]. Further, one can easily combine multiple edit directions by taking a weighted combination of individual attribute edits (Fig. 10), thanks to the linearity of the  $\mathcal{W}+$ .

#### 3.7 Composing multiple persons

Our method can be extended to compose multiple subject identities in a single scene. Naively, embedding multiple token embeddings (one per subject) in the text prompt without subject-specific tuning results in identity distortion Fig. 5a). Jointly performing subject-specific tuning improves the identity but suffers from at*tribute mixing*, where facial attributes from one subject are transferred to another, such as age and hairs in Fig. 5b). This is a well-known issue in T2I generation, where the model struggles with multiple objects in a scene and binds incorrect at-



Fig. 5: Composing multiple persons without finetuning results in identity distortion. Finetuning a single model for both the identities results in *attribute mixing*, the age and facial hairs from v1 are transferred to v2. Combining outputs of individual finetuned models results in excellent identity preservation without *attribute mixing*. of the  $\mathcal{W}_{+}$ 



Fig. 6: Composing multiple subjects. We run multiple parallel diffusion processes, one per subject and one for the background, which are fused using instance masks at each denoising step. Importantly, the diffusion process for each subject is passed through its corresponding fine-tuned model, which results in excellent identity preservation.

tributes [7]. We take an alternate approach inspired by MultiDiffusion [3], where we run multiple chained diffusion processes, one for each subject and one for the background. The outputs of these processes are combined at each denoising step



Fig. 7: Comparison for single subject personalization. Existing personalization methods designed for generic concepts either achieve good Identity similarity (Custom Diffusion) or Prompt similarity (Celeb Basis, Dreambooth, Dreambooth+Lora) but fail to achieve both simultaneously. Like ours, Celeb Basis is a single image, face-specific personalization method that achieves good prompt similarity. However, faces generated by Celeb Basis have a cartoonish look and lack realism. Our method strikes a perfect balance between Identity similarity and Prompt similarity, as shown in the plot, and generates highly photo realistic images following the text.

using an instance segmentation mask. We run the diffusion process for each subject through its corresponding subject-specific finetuned model. This preserves the subjects' details learned by each finetuned model and enables high-fidelity composition of multiple persons without *attribute mixing*. To obtain an instance segmentation mask, we run a single diffusion process with a prompt containing two persons and apply the off-the-shelf segmentation model SAM [20] on the generated image. Further, we can perform fine-grained attribute edits on a single subject with latent manipulation in W+ space while preserving other subjects, as shown in Fig. 1.

#### 4 Experiments

We perform all our experiments on StableDiffusion-v2.1 [36] as a representative T2I model. For inversion, we use pre-trained StyleGAN2 e4e encoder [44] trained on the face dataset to map images in W+. In the following sections, we first discuss the datasets and metrics (Sec.4.1), followed by results on single-subject and multi-subject personalization (Sec.4.2), fine-grained attribute editing (Sec.4.3), and ablation studies (Sec.4.4).

#### 4.1 Dataset and metrics

**Dataset.** The latent adaptor is trained with a combination of synthetic images generated from StyleGAN2 and real images from the FFHQ [17] dataset. The dataset contained 70K image and corresponding w latent codes obtained from e4e [44]. We collected a dataset of 30 subjects for evaluation, including scientists,

celebrities, sports persons, and tech executives. We also evaluate on 'non-famous' identities and synthetic faces in supplementary. We use a set of 25 diverse text prompts, including texts for stylization, background change, and doing certain actions. Further details about the setup is provided in the supplementary.

Metrics. We evaluate the personalization performance using two widely used metrics for subject personalization: **Prompt similarity** - to measure the alignment of the prompt with the generated image using CLIP [32], and **Identity similarity**(**CS**) - to measure the identity similarity between the input image and the generated image using cosine similarity between face embeddings from [46]. To evaluate fine-grained attribute editing, we compute the change in Prompt similarity ( $\Delta$  CLIP) with the attribute prompt (e.g., 'A v1 person smiling') before and after the edit. Additionally, we measure the change in the image during editing with LPIPS [52] and Identity similarity. For an ideal fine-grained attribute edit, a higher  $\Delta$  CLIP indicates a meaningful edit and a lower LPIPS and higher ID-sim denotes the preservation of source identity.

#### 4.2 Comparison with personalization methods.

Single-subject personalization. We perform single-image personalization on the evaluation set with diverse text prompts in Fig. 7, 13 & S8. We compare with following fine-tuningbased personalization methods: Custom Diffusion [22], Dreambooth [37], Dreambooth+LoRA, which is Dreambooth with low-rank updates to avoid overfitting, Textual Inversion [10] and Celeb Basis [51]. All the methods are trained with 5 images per subject except Celeb-basis and ours, which operate on a single input image. Details about hyper-parameters for competitor methods are provided in the supplementary. Custom Diffusion embeds a subject while preserving its identity; however, it mostly generates closeup faces and does not follow the text prompts to stylize the subject or to have it perform an action. Dreambooth cannot embed the subject's identity faithfully, whereas with LoRA training, the identity is improved along with text alignment,



Fig. 8: Comparison for multi-subject generation. Textual Inversion struggles to generate both the subjects and distorts their identities. Custom Diffusion and Celeb Basis can compose but suffer from *attribute mixing* - the same hairstyle and facial features are copied to both faces. This effect is more pronounced in Celeb Basis. Our method disentangles the attributes of the two subjects and generates identitypreserving composition.

which helps avoid overfitting. Textual Inversion and Celeb Basis have poor identity preservation as they fine-tune only the token embedding and not the U-Net. Celeb Basis achieves the highest text alignment due to the strong regularization imposed by basis spanning across celebrity names. Our method strikes a perfect balance between text alignment and identity preservation. Note that ours and the Celeb Basis use only 1 input image, which slightly affects the identity, against Custom-diffusion that requires 5 images. We have provided an additional comparison with encoder-based models and the recent IP-adaptor [50] method in supplementary material.

Multi-subject personalization. We present results for embedding multipleperson composition in Fig. 8, 14, S7. Specifically, we combine the intermediate outputs of subject-specific tuned models during generation. We compare against multi-concept personalization methods, Textual Inversion, Custom Diffusion, and Celeb Basis. For Textual Inversion and Celeb Basis, we learned two separate token embeddings, one for each subject. For Custom Diffusion, we jointly fine-tune the projection matrices on both subjects. Textual Inversion fails to generate both the subjects in the scene. Celeb Basis and Custom Diffusion generate both the subjects but suffer from attribute mixing (eveglasses from v4 are transferred to v3). As noted earlier, Celeb Basis generates cartoonish faces in most cases. Our method resolves the attribute mixing by running multiple subjectspecific diffusion processes and results in highly realistic compositions.

# 4.3 Fine-grained control by latent manipulation

The proposed method matches the disentangled W+ latent space of StyleGANs to the token embedding space of T2I models, allowing for continuous control over the image at-



Fig. 9: Attribute Control. We perform continuous attribute editing by adding an attribute edit direction in W+ and increasing its edit strength  $\beta$ . Our method performs disentangled edits for various attributes while preserving identity and generalizing to in-the-wild faces, styles, and multiple persons. Identity Interpolation. We can perform smooth interpolation between identities by interpolating between the corresponding w codes.

tributes by latent space manipulation. We present two important image editing applications fueled by the disentangled latent space of StyleGANs: 1) fine-grained attribute editing and 2) smooth identity interpolation. Additionally, our model

					T		· · c	· · · ·				2	
		Imagic		InterfaceGA			N Ours + Prompt			$\mathrm{Ours} + \mathcal{W} +$			
Attribute	$ \Delta $	CLIP	$\uparrow$ LPIPS	$\downarrow CS \uparrow   \Delta$	$\operatorname{CLIP}$	$\uparrow \mathrm{LPIPS} \downarrow$	$CS\uparrow$	$\Delta$ CLIP	$\uparrow \text{LPIPS}$ .	$\downarrow CS \uparrow$	$\Delta$ CLIP	$\uparrow$ LPIPS $\downarrow$	. CS $\uparrow$
Beard	4	4.184	0.533	0.492	2.644	0.221	0.732	0.986	0.232	0.877	2.473	0.185	0.731
Age	4	1.432	0.571	0.474	1.359	0.220	0.744	0.429	0.202	0.929	1.777	0.209	0.698
Smile	1	1.945	0.499	0.666	1.800	0.188	0.804	0.784	0.317	0.813	1.104	0.190	0.779
Asian	5	5.908	0.585	0.478	4.198	0.139	0.708	1.008	0.330	0.804	4.917	0.165	0.625
Black	5	5.806	0.546	0.347	1.748	0.147	0.825	0.897	0.222	0.898	1.877	0.133	0.834

Table 1: Comparison of fine-grained attribute editing

can restore corrupted face images such as low resolution or inpainting masked facial features in supplementary.

Fine-grained attribute editing. We perform attribute editing by adding a global latent edit direction in W+ to w encoding of the input image. To have a unified method for all the attributes, we take a simplified approach to obtain edit directions, gathering a small set (< 20) of paired portrait images before

and after the attribute edit (generated using an off-the-shelf attribute editing method). Next, we take a difference between the corresponding paired wlatent and average them to obtain a global edit direction. We found global edit direction for smile, age, beard, gender, race, and eyeglasses. We also show edits with directions obtained using InterfaceGAN [41] in Fig. S6 in supplementary. The results for finegrained control editing are provided in Fig. 9 & 10, where we show disentangled continuous control for various attributes by changing  $\beta$  (ref. Sec. 3.6) while preserving the identity. Our method generalizes the edit directions in W+, originally defined for portrait faces, to in-the-wild and stylized face *images.* We evaluate attribute editing performance against 1) StyleGANbased global editing method Inter-



Fig. 10: Multi-attribute-control. We can perform continuous edits for two attributes simultaneously by taking a linear combination of attribute edit directions. Observe the smooth and disentangled edit transformations for age and beard attributes while preserving identity.

faceGAN [41], after encoding the image using e4e, 2) Prompt-based editing of the learned subject (by giving prompts like 'A photo of v1 smiling'), 3) Textbased editing method Imagic [19] build on single image personalization. The quantitative results are present in Tab. 1, and qualitative results are in Fig. 11. Our method achieves the lowest LPIPS scores with high  $\Delta$  CLIP, indicating highly disentangled attribute editing. Both text-based editing methods fail to preserve the image regions (higher LPIPS). We achieve high CS scores during edits with higher  $\Delta$  CLIP, indicating identity-preserving attribute edits. Ours prompt-based editing achieves a superior CS because the edit is not performed in many cases indicated by lower LPIPS [52]. Like ours, InterfaceGAN works in W+ latent space and performs similarly in preserving

PreciseControl

13

the image content, the identity of the subject, and editability. However, it is limited to the editing of portrait faces generated by StyleGANs and loses fine-facial features, whereas our method combines the best of both worlds, allowing for fine-grained latent editing with semantic editing in T2I models.

Identity interpolation. W+ space also allows for smooth interpolation between two identities. Given two input images, we obtain their corresponding w latent codes and perform linear interpolation to obtain the intermediate latent codes. When used as conditioning through the latent



Fig. 11: Comparison for attribute editing. Using images for text-based attribute edits results in identity distortion and lacks realism after edit. As both InterfaceGAN and our method leverage the same disentangled latent space, they generate high-quality edits. However, Interface-GAN is limited to cropped faces while we can edit in-the-wild images too.

adaptor, these latents result in realistic face interpolations with smooth changes between the two faces, preserving background, as shown in Fig. 9-Bottom.

#### Ablations **4.4**

We ablate over the design choices made in the proposed approach for personalization in Fig. 12. The Identity loss and regularization loss have a similar effect in pushing the token embeddings close to an embedding region for faces. Time-dependent token embedding is crucial to preserving subjects' identity as it provides a more expressive space to represent the face. Finally, subject-specific tuning with combined loss improves the Identity similarity as well as the Prompt



similarity as the predicted token embeddings are pushed closer to the editable region with  $\mathcal{L}_{reg}$  and  $\mathcal{L}_{ID}$ .

#### 5 Discussion

**Conclusion.** We present a novel framework to condition T2I diffusion models on  $\mathcal{W}$ + space of StyleGAN2 models for fine-grained attribute control. Specifically, we learn a latent mapper that projects the latent codes from  $\mathcal{W}+$  to the input token embedding space of the T2I model learned with denoising, regularization, and identity preservation losses. This framework provides a natural way to embed

a real face image by obtaining its latent code using the GAN encoders model. The embedded face can then be edited in two manners - coarse text-based editing and fine-grained attribute editing by latent manipulation in W+.

**Limitations.** The primary limitation is that the encoder-based inversion in  $\mathcal{W}$ + is loose on some information hence we need to perform test time fine-tuning for a few iterations to recover identity similar to pivotal tuning. Additionally, the current approach utilizes multi-diffusion for composing multiple persons which requires multiple diffusion processes. Effectively composing more than two individuals with consistent identities proves challenging within the current method and is an interesting future direction to explore.



Fig. 13: Additional results for single subject personalization. Our method achieves excellent realism and text alignment with the prompts.



Fig. 14: Additional results for Multi-subject personalization. Our method preserves subject identity and follows the text prompt during generation.

**Ethics statement.** We propose a face personalization method that might be useful for fake news generation as it enables the generation of persons in novel contexts. However, this issue is not limited to this work as it exists in several personalization approaches [11,51] and generative models [36]. Nonetheless, recent works [21] have leveraged the power of the same personalization approaches for removing concepts or biases from T2I models.

Acknowledgements. We thank Aniket Dashpute, Ankit Dhiman and Abhijnya Bhat for reviewing the draft and providing helpful feedback. This work was partly supported by PMRF from Govt. of India (Rishubh Parihar) and Kotak IISc AI-ML Centre.

#### References

- Abdal, R., Zhu, P., Mitra, N.J., Wonka, P.: Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. ACM Transactions on Graphics (TOG) 40(3), 1–21 (2021)
- Alaluf, Y., Richardson, E., Metzer, G., Cohen-Or, D.: A neural space-time representation for text-to-image personalization. ACM Transactions on Graphics (TOG) 42(6), 1–10 (2023)
- 3. Bar-Tal, O., Yariv, L., Lipman, Y., Dekel, T.: Multidiffusion: Fusing diffusion paths for controlled image generation (2023)
- Brack, M., Friedrich, F., Hintersdorf, D., Struppek, L., Schramowski, P., Kersting, K.: Sega: Instructing text-to-image models using semantic guidance. In: Thirtyseventh Conference on Neural Information Processing Systems (2023)
- Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18392–18402 (2023)
- Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., et al.: Efficient geometry-aware 3d generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16123–16133 (2022)
- Chefer, H., Alaluf, Y., Vinker, Y., Wolf, L., Cohen-Or, D.: Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. ACM Transactions on Graphics (TOG) 42(4), 1–10 (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)
- 9. Esler, T.: Github face recognition using pytorch. https://github.com/timesler/facenet-pytorch (2021)
- 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)
- Gandikota, R., Materzynska, J., Zhou, T., Torralba, A., Bau, D.: Concept sliders: Lora adaptors for precise control in diffusion models. arXiv preprint arXiv:2311.12092 (2023)

- 16 R. Parihar *et al.*
- Härkönen, E., Hertzmann, A., Lehtinen, J., Paris, S.: Ganspace: Discovering interpretable gan controls. Advances in Neural Information Processing Systems 33, 9841–9850 (2020)
- Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-or, D.: Prompt-to-prompt image editing with cross-attention control. In: The Eleventh International Conference on Learning Representations (2022)
- Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. arXiv preprint arXiv:2106.09685 (2021)
- Karmali, T., Parihar, R., Agrawal, S., Rangwani, H., Jampani, V., Singh, M., Babu, R.V.: Hierarchical semantic regularization of latent spaces in stylegans. In: European Conference on Computer Vision. pp. 443–459. Springer (2022)
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4401–4410 (2019)
- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8110–8119 (2020)
- 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)
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. arXiv preprint arXiv:2304.02643 (2023)
- Kumari, N., Zhang, B., Wang, S.Y., Shechtman, E., Zhang, R., Zhu, J.Y.: Ablating concepts in text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 22691–22702 (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)
- Kwon, M., Jeong, J., Uh, Y.: Diffusion models already have a semantic latent space. arXiv preprint arXiv:2210.10960 (2022)
- 24. Li, X., Hou, X., Loy, C.C.: When stylegan meets stable diffusion: a w+ adapter for personalized image generation. arXiv preprint arXiv:2311.17461 (2023)
- Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: Sdedit: Guided image synthesis and editing with stochastic differential equations. arXiv preprint arXiv:2108.01073 (2021)
- Parihar, R., Bhat, A., Basu, A., Mallick, S., Kundu, J.N., Babu, R.V.: Balancing act: Distribution-guided debiasing in diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6668– 6678 (2024)
- Parihar, R., Dhiman, A., Karmali, T., Babu, R.V.: Everything is there in latent space: Attribute editing and attribute style manipulation by stylegan latent space exploration. In: Proceedings of the 30th ACM International Conference on Multimedia. pp. 1828–1836 (2022)
- Parmar, G., Kumar Singh, K., Zhang, R., Li, Y., Lu, J., Zhu, J.Y.: Zero-shot image-to-image translation. In: ACM SIGGRAPH 2023 Conference Proceedings. pp. 1–11 (2023)

- 29. Patashnik, O., Garibi, D., Azuri, I., Averbuch-Elor, H., Cohen-Or, D.: Localizing object-level shape variations with text-to-image diffusion models. arXiv preprint arXiv:2303.11306 (2023)
- Patashnik, O., Garibi, D., Azuri, I., Averbuch-Elor, H., Cohen-Or, D.: Localizing object-level shape variations with text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
- Patashnik, O., Wu, Z., Shechtman, E., Cohen-Or, D., Lischinski, D.: Styleclip: Text-driven manipulation of stylegan imagery. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2085–2094 (2021)
- 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)
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical textconditional image generation with clip latents. arXiv preprint arXiv:2204.06125 1(2), 3 (2022)
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: International Conference on Machine Learning. pp. 8821–8831. PMLR (2021)
- Richardson, E., Alaluf, Y., Patashnik, O., Nitzan, Y., Azar, Y., Shapiro, S., Cohen-Or, D.: Encoding in style: a stylegan encoder for image-to-image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2287–2296 (2021)
- 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)
- 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)
- Ruiz, N., Li, Y., Jampani, V., Wei, W., Hou, T., Pritch, Y., Wadhwa, N., Rubinstein, M., Aberman, K.: Hyperdreambooth: Hypernetworks for fast personalization of text-to-image models. arXiv preprint arXiv:2307.06949 (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)
- 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)
- Shen, Y., Gu, J., Tang, X., Zhou, B.: Interpreting the latent space of gans for semantic face editing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9243–9252 (2020)
- Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
- 43. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. Advances in Neural Information Processing Systems 33, 7537–7547 (2020)

- 18 R. Parihar *et al.*
- Tov, O., Alaluf, Y., Nitzan, Y., Patashnik, O., Cohen-Or, D.: Designing an encoder for stylegan image manipulation. ACM Transactions on Graphics (TOG) 40(4), 1–14 (2021)
- Valevski, D., Lumen, D., Matias, Y., Leviathan, Y.: Face0: Instantaneously conditioning a text-to-image model on a face. In: SIGGRAPH Asia 2023 Conference Papers. pp. 1–10 (2023)
- 46. Wang, H., Wang, Y., Zhou, Z., Ji, X., Gong, D., Zhou, J., Li, Z., Liu, W.: Cosface: Large margin cosine loss for deep face recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5265–5274 (2018)
- 47. Wang, J., Yue, Z., Zhou, S., Chan, K.C., Loy, C.C.: Exploiting diffusion prior for real-world image super-resolution. arXiv preprint arXiv:2305.07015 (2023)
- Wang, Q., Bai, X., Wang, H., Qin, Z., Chen, A.: Instantid: Zero-shot identitypreserving generation in seconds. arXiv preprint arXiv:2401.07519 (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)
- 50. 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)
- Yuan, G., Cun, X., Zhang, Y., Li, M., Qi, C., Wang, X., Shan, Y., Zheng, H.: Inserting anybody in diffusion models via celeb basis. arXiv preprint arXiv:2306.00926 (2023)
- 52. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)
- Zhou, Y., Zhang, R., Sun, T., Xu, J.: Enhancing detail preservation for customized text-to-image generation: A regularization-free approach. arXiv preprint arXiv:2305.13579 (2023)
- Zhou, Y., Zhang, R., Sun, T., Xu, J.: Enhancing detail preservation for customized text-to-image generation: A regularization-free approach. arXiv preprint arXiv:2305.13579 (2023)