

Diffusion-driven GAN Inversion for Multi-Modal Face Image Generation

Jihyun Kim^{1,2}, Changjae Oh³, Hoseok Do², Soohyun Kim², Kwanghoon Sohn^{1,4*}

¹ Yonsei University ² AI Lab, CTO Division, LG Electronics

³ Queen Mary University of London ⁴ Korea Institute of Science and Technology (KIST)

{hyunys21, khsohn}@yonsei.ac.kr c.oh@qmul.ac.uk {hoseok.do, soohyun1.kim}@lge.com

Abstract

We present a new multi-modal face image generation method that converts a text prompt and a visual input, such as a semantic mask or scribble map, into a photo-realistic face image. To do this, we combine the strengths of Generative Adversarial networks (GANs) and diffusion models (DMs) by employing the multi-modal features in the DM into the latent space of the pre-trained GANs. We present a simple mapping and a style modulation network to link two models and convert meaningful representations in feature maps and attention maps into latent codes. With GAN inversion, the estimated latent codes can be used to generate 2D or 3D-aware facial images. We further present a multi-step training strategy that reflects textual and structural representations into the generated image. Our proposed network produces realistic 2D, multi-view, and stylized face images, which align well with inputs. We validate our method by using pre-trained 2D and 3D GANs, and our results outperform existing methods. Our project page is available at https://github.com/1211sh/Diffusion-driven_GAN-Inversion/.

1. Introduction

In recent years, multi-modal image generation has achieved remarkable success, driven by the advancements in Generative Adversarial Networks (GANs) [15] and diffusion models (DMs) [11, 18, 48]. Facial image processing has become a popular application for a variety of tasks, including face image generation [21, 39], face editing [6, 12, 30, 36, 37, 46], and style transfer [7, 64]. Many tasks typically utilize the pre-trained StyleGAN [21, 22], which can generate realistic facial images and edit facial attributes by manipulating the latent space using GAN inversion [39, 42, 58]. In these tasks, using multiple modalities as conditions is becoming a popular approach, which improves the user’s controllability in generating realistic face images. However, existing GAN

*Corresponding author

This research was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIP) (NRF2021R1A2C2006703).

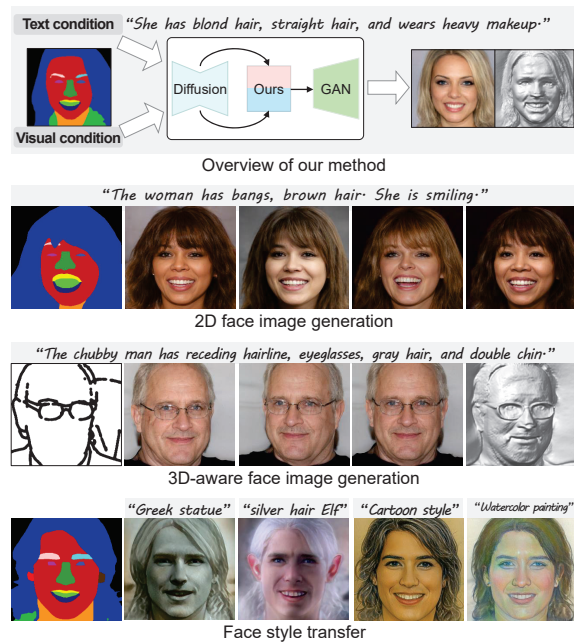


Figure 1. We present a method to map the diffusion features to the latent space of a pre-trained GAN, which enables diverse tasks in multi-modal face image generation and style transfer. Our method can be applied to 2D and 3D-aware face image generation.

inversion methods [51, 58] have poor alignment with inputs as they neglect the correlation between multi-modal inputs. They struggle to map the different modalities into the latent space of the pre-trained GAN, such as by mixing the latent codes or optimizing the latent code converted from a given image according to the input text.

Recently, DMs have increased attention in multi-modal image generation thanks to the stability of training and the flexibility of using multiple modalities as conditions. DMs [23, 53, 54] can control the multiple modalities and render diverse images by manipulating the latent or attention features across the time steps. However, existing text-to-image DMs rely on an autoencoder and text encoder, such as CLIP [41], trained on unstructured datasets collected from the web [40, 45] that may lead to unrealistic

image generation.

Moreover, some approaches address multi-modal face image generation in a 3D domain. In GAN inversion [14, 51], multi-view images can be easily acquired by manipulating the latent code with pre-trained 3D GANs. While DMs are inefficient in learning 3D representation, which has the challenge to generate multi-view images directly due to the lack of 3D ground-truth (GT) data for training [32, 47]. They can be used as a tool to acquire training datasets for 3D-aware image generation [24, 33].

In this paper, we present a versatile face generative model that uses text and visual inputs. We propose an approach that takes the strengths of DMs and GAN and generates photo-realistic images with flexible control over facial attributes, which can be adapted to 2D and 3D domains, as illustrated in Figure 1. Our method employs a latent mapping strategy that maps the diffusion features into the latent space of a pre-trained GAN using multi-denoising step learning, producing the latent code that encodes the details of text prompts and visual inputs.

In summary, our main contributions are:

- (i) We present a novel method to link a pre-trained GAN (StyleGAN [22], EG3D [4]) and DM (ControlNet [62]) for multi-modal face image generation.
- (ii) We propose a simple mapping network that links pre-trained GAN and DM’s latent spaces and an attention-based style modulation network that enables the use of meaningful features related to multi-modal inputs.
- (iii) We present a multi-denoising step training strategy that enhances the model’s ability to capture the textual and structural details of multi-modal inputs.
- (iv) Our model can be applied for both 2D- and 3D-aware face image generation without additional data or loss terms and outperforms existing DM- and GAN-based methods.

2. Related Work

2.1. GAN Inversion

GAN inversion approaches have gained significant popularity in the face image generation task [7, 31, 51, 59] using the pre-trained 2D GAN, such as StyleGAN [21, 22]. This method has been extended to 3D-aware image generation [27, 60, 61] by integrating 3D GANs, such as EG3D [4]. GAN inversion can be categorized into learning-based, optimization-based, and hybrid methods. Optimization-based methods [44, 67] estimate the latent code by minimizing the difference between an output and an input image. Learning-based methods [1, 52] train an encoder that maps an input image into the latent space of the pre-trained GAN. Hybrid methods [58, 66] combine these two methods, producing an initial latent code and then refining it with additional optimizations. Our work employs a learning-based GAN inversion, where a DM serves as the encoder. We produce latent codes by leveraging semantic

features in the denoising U-Net, which can generate images with controlled facial attributes.

2.2. Diffusion Model for Image Generation

Many studies have introduced text-to-image diffusion models [36, 43, 45] that generate images by encoding multi-modal inputs, such as text and image, into latent features via foundation models [41] and mapping them to the features of denoising U-Net via an attention mechanism. ControlNet [62] performs image generation by incorporating various visual conditions (*e.g.*, semantic mask, scribbles, edges) and text prompts. Image editing models using DMs [16, 20, 26, 28, 34] have exhibited excellent performance by controlling the latent features or the attention maps of a denoising U-Net. Moreover, DMs can generate and edit images by adjusting latent features over multiple denoising steps [2]. We focus on using latent features of DM, including intermediate features and cross-attention maps, across denoising steps to link them with the latent space of GAN and develop a multi-modal face image generation task.

2.3. Multi-Modal Face Image Generation

Face generative models have progressed by incorporating various modalities, such as text [25], semantic mask [38, 55], sketch [5, 9], and audio [65]. Several methods adopt StyleGAN, which can generate high-quality face images and edit facial attributes to control the style vectors. The transformer-based models [3, 13] are also utilized, which improves the performance of face image generation by handling the correlation between multi-modal conditions using image quantization. A primary challenge faced in face generative models is to modify the facial attributes based on given conditions while minimizing changes to other attributes. Some methods [39, 57] edit facial attributes by manipulating the latent codes in GAN models. TediGAN [58] controls multiple conditions by leveraging an encoder to convert an input image into latent codes and optimizing them with a pre-trained CLIP model. Recent works [19, 35] use DMs to exploit the flexibility of taking multiple modalities as conditions and generate facial images directly from DMs. Unlike existing methods, we use the pre-trained DM [62] as an encoder to further produce the latent codes for the pre-trained GAN models.

3. Method

3.1. Overview

Figure 2 illustrates the overall pipeline of our approach. During the reverse diffusion process, we use the middle and decoder blocks of a denoising U-Net in ControlNet [62] as an encoder \mathcal{E} . A text prompt c , along with a visual condition x , are taken as input to the denoising U-Net. Subsequently, \mathcal{E} produces the feature maps h from the middle block, and




Figure 2. Overview of our method. We use a diffusion-based encoder \mathcal{E} , the middle and decoder blocks of a denoising U-Net, that extracts the semantic features \mathbf{h}_t , intermediate features \mathbf{f}_t , and cross-attention maps \mathbf{a}_t at denoising step t . We present the mapping network \mathcal{M} (Sec. 3.2) and the attention-based style modulation network (AbSMNet) \mathcal{T} (Sec. 3.3) that are trained across t (Sec. 3.4). \mathcal{M} converts \mathbf{h}_t into the mapped latent code \mathbf{w}_t^m , and \mathcal{T} uses \mathbf{f}_t and \mathbf{a}_t to control the facial attributes from the text prompt c and visual input \mathbf{x} . The modulation codes \mathbf{w}_t^γ and \mathbf{w}_t^β are then used to scale and shift \mathbf{w}_t^m to produce the final latent code, \mathbf{w}_t , that is fed to the pre-trained GAN \mathcal{G} . We obtain the generation output I_t' from our model \mathcal{Y} and we use the image I_0^d from the U-Net after the entire denoising process for training \mathcal{T} (Sec. 3.4). Note that only the networks with the dashed line (---) are trainable, while others are frozen.

the intermediate features \mathbf{f} and the cross-attention maps \mathbf{a} from the decoder blocks. \mathbf{h} is then fed into the mapping network \mathcal{M} , which transforms the rich semantic feature into a latent code \mathbf{w}^m . The Attention-based Style Modulation Network (AbSMNet), \mathcal{T} , takes \mathbf{f} and \mathbf{a} as input to generate the modulation latent codes, \mathbf{w}^γ and \mathbf{w}^β , that determine facial attributes related to the inputs. The latent code \mathbf{w} is then forwarded to the pre-trained GAN \mathcal{G} that generates the output image I' . Our model is trained across multiple denoising steps, and we use the denoising step t to indicate the features and images obtained at each denoising step. With this pipeline, we aim to estimate the latent code, \mathbf{w}_t^* , that is used as input to \mathcal{G} to render a GT image, I^{gt} :

$$\mathbf{w}_t^* = \arg \min_{\mathbf{w}_t} \mathcal{L}(I^{gt}, \mathcal{G}(\mathbf{w}_t)), \quad (1)$$

where $\mathcal{L}(\cdot, \cdot)$ measures the distance between I^{gt} and the rendered image, $I' = \mathcal{G}(\mathbf{w}_t)$. We employ learning-based GAN inversion that estimates the latent code from an encoder to reconstruct an image according to given inputs.

3.2. Mapping Network

Our mapping network \mathcal{M} aims to build a bridge between the latent space of the diffusion-based encoder \mathcal{E} and that of the pre-trained GAN \mathcal{G} . \mathcal{E} uses a text prompt and a visual input, and these textual and image embeddings are aligned by the cross-attention layers [62]. The feature maps \mathbf{h} from the middle block of the denoising U-Net particularly contain rich semantics that resemble the latent space of the generator [28]. Here we establish the link between the latent

spaces of \mathcal{E} and \mathcal{G} by using \mathbf{h}_t across the denoising steps t .

Given \mathbf{h}_t , we design \mathcal{M} that produces a 512-dimensional latent code $\mathbf{w}_t^m \in \mathbb{R}^{L \times 512}$ that can be mapped to the latent space of \mathcal{G} :

$$\mathbf{w}_t^m = \mathcal{M}(\mathbf{h}_t). \quad (2)$$

\mathcal{M} is designed based on the structure of the *map2style* block in pSp [42], as seen in Figure 2. This network consists of convolutional layers downsampling feature maps and a fully connected layer producing the latent code \mathbf{w}_t^m .

3.3. Attention-based Style Modulation Network

By training \mathcal{M} with learning-based GAN inversion, we can obtain \mathbf{w}_t^m and use it as input to the pre-trained GAN for image generation. However, we observe that \mathbf{h}_t shows limitations in capturing fine details of the facial attributes due to its limited spatial resolution and data loss during the encoding. Conversely, the feature maps of the DM's decoder blocks show rich semantic representations [53], benefiting from aggregating features from DM's encoder blocks via skip connections. We hence propose a novel Attention-based Style Modulation Network (AbSMNet), \mathcal{T} , that produces style modulation latent codes, $\mathbf{w}_t^\gamma, \mathbf{w}_t^\beta \in \mathbb{R}^{L \times 512}$, by using \mathbf{f}_t and \mathbf{a}_t from \mathcal{E} . To improve reflecting the multimodal representations to the final latent code \mathbf{w}_t , we modulate \mathbf{w}_t^m from \mathcal{M} using \mathbf{w}_t^γ and \mathbf{w}_t^β , as shown in Figure 2.

We extract intermediate features, $\mathbf{f}_t = \{\mathbf{f}_t^n\}_{n=1}^N$, from N different blocks, and cross-attention maps, $\mathbf{a}_t = \{\mathbf{a}_t^k\}_{k=1}^K$, from K different cross-attention layers of the n -th block, in \mathcal{E} that is a decoder stage of denoising U-Net. The discrim-




Figure 3. Visualization of cross-attention maps and intermediate feature maps. (a) represents the semantic relation information between an input text and an input semantic mask in the spatial domain. The meaningful representations of inputs are shown across all denoising steps and N different blocks. (b) represents N different cross-attention maps, \mathbf{A}_t , at denoising steps $t = T$ and $t = 0$. (c) shows the example of refined intermediate feature map $\hat{\mathbf{F}}_T^1$ at 1st block and $t = T$ that is emphasized corresponding to input multi-modal conditions. The red and yellow regions of the map indicate higher attention scores. As the denoising step approaches T , the text-relevant features appear more clearly, and as the denoising step t approaches 0, the features of the visual input are more preserved.

inative representations are represented more faithfully because \mathbf{f}_t consists of N multi-scale feature maps that can capture different sizes of facial attributes, which allows for finer control over face attributes. For simplicity, we upsample each intermediate feature map of \mathbf{f}_t to same size intermediate feature maps $\mathbf{F}_t = \{\mathbf{F}_t^n\}_{n=1}^N$, where $\mathbf{F}_t^n \in \mathbb{R}^{H \times W \times C_n}$ has H , W , and C_n as height, width and depth.

Moreover, \mathbf{a}_t is used to amplify controlled facial attributes as it incorporates semantically related information in text and visual input. To match the dimension with \mathbf{F}_t , we convert \mathbf{a}_t to $\mathbf{A}_t = \{\mathbf{A}_t^n\}_{n=1}^N$, where $\mathbf{A}_t^n \in \mathbb{R}^{H \times W \times C_n}$, by max-pooling the output of the cross-attention layers in each decoder block and upsampling the max-pooling outputs. To capture the global representations, we additionally compute $\bar{\mathbf{A}}_t \in \mathbb{R}^{H \times W \times 1}$ by depth-wise averaging the max-pooling output of \mathbf{a}_t over each word in the text prompt and upsampling it. As illustrated in Figures 3 (a) and (b), \mathbf{A}_t and $\bar{\mathbf{A}}_t$ represent the specific regions aligned with input text prompt and visual input, such as semantic mask, across denoising steps t . By a pixel-wise multiplication between \mathbf{F}_t and \mathbf{A}_t , we can obtain the refined intermediate feature maps $\hat{\mathbf{F}}_t$ that emphasize the representations related to multi-




Figure 4. Style modulation network in \mathcal{T} . The refined intermediate feature maps $\hat{\mathbf{F}}_t$ and $\hat{\mathbf{F}}_t$ are used to capture local and global semantic representations, respectively. They are fed into the scale and shift network, respectively. The weighted summations of these outputs are used as input to the map2style network, which finally generates the scale and shift modulation latent codes, \mathbf{w}_t^γ , and \mathbf{w}_t^β .

modal inputs as shown in Figure 3 (c). The improved average feature map $\hat{\mathbf{F}}_t \in \mathbb{R}^{H \times W \times 1}$ is also obtained by multiplying $\bar{\mathbf{A}}_t$ with $\hat{\mathbf{F}}_t$, where $\hat{\mathbf{F}}_t \in \mathbb{R}^{H \times W \times 1}$ is obtained by first averaging the feature maps in $\mathbf{F}_t = \{\mathbf{F}_t^n\}_{n=1}^N$ and then depth-wise averaging the outputs. $\hat{\mathbf{F}}_t$ and $\hat{\mathbf{F}}_t$ distinguish text- and structural-relevant semantic features, which improves the alignment with the inputs. We use $\hat{\mathbf{F}}_t$ and $\hat{\mathbf{F}}_t$ as input to the style modulation network that produces the modulation codes \mathbf{w}_t^γ , and \mathbf{w}_t^β as shown in Figure 4. We capture both local and global features by using $\hat{\mathbf{F}}_t$, which consists of feature maps representing different local regions on the face, and $\hat{\mathbf{F}}_t$, which implies representations of the entire face. We concatenate N intermediate feature maps of $\hat{\mathbf{F}}_t$, $\text{concat}(\hat{\mathbf{F}}_t^1 \dots \hat{\mathbf{F}}_t^N)$, and it is forward to the scale and shift networks that consist of convolutional layers and Leaky ReLU, forming the local modulation feature maps, $\hat{\mathbf{F}}_t^{\gamma l}$ and $\hat{\mathbf{F}}_t^{\beta l}$. We also estimate global modulation feature maps, $\hat{\mathbf{F}}_t^{\gamma g}$ and $\hat{\mathbf{F}}_t^{\beta g}$, by feeding $\hat{\mathbf{F}}_t$ to the scale and shift network. The final scale, $\hat{\mathbf{F}}_t^\gamma$, and shift, $\hat{\mathbf{F}}_t^\beta$, feature maps are estimated by the weighted summation:

$$\begin{aligned} \hat{\mathbf{F}}_t^\gamma &= \alpha_t^\gamma \hat{\mathbf{F}}_t^{\gamma l} + (1 - \alpha_t^\gamma) \hat{\mathbf{F}}_t^{\gamma g}, \\ \hat{\mathbf{F}}_t^\beta &= \alpha_t^\beta \hat{\mathbf{F}}_t^{\beta g} + (1 - \alpha_t^\beta) \hat{\mathbf{F}}_t^{\beta l}, \end{aligned} \quad (3)$$

where α_t^γ and α_t^β are learnable weight parameters. Through the *map2style* module, we then convert $\hat{\mathbf{F}}_t^\gamma$ and $\hat{\mathbf{F}}_t^\beta$ into the final scale, $\mathbf{w}_t^\gamma \in \mathbb{R}^{L \times 512}$, and shift, $\mathbf{w}_t^\beta \in \mathbb{R}^{L \times 512}$, latent codes. With these modulation latent codes, we achieve more precise control over facial details while corresponding to the input multi-modal inputs at the pixel level.

Finally, the mapped latent code \mathbf{w}_t^m from \mathcal{M} is modulated by \mathbf{w}_t^γ and \mathbf{w}_t^β from \mathcal{T} to get the final latent code \mathbf{w}_t that is used to obtain the generated image I'_t as follows:

$$\mathbf{w}_t = \mathbf{w}_t^m \odot \mathbf{w}_t^\gamma \oplus \mathbf{w}_t^\beta, \quad (4)$$

$$I'_t = \mathcal{G}(\mathbf{w}_t). \quad (5)$$

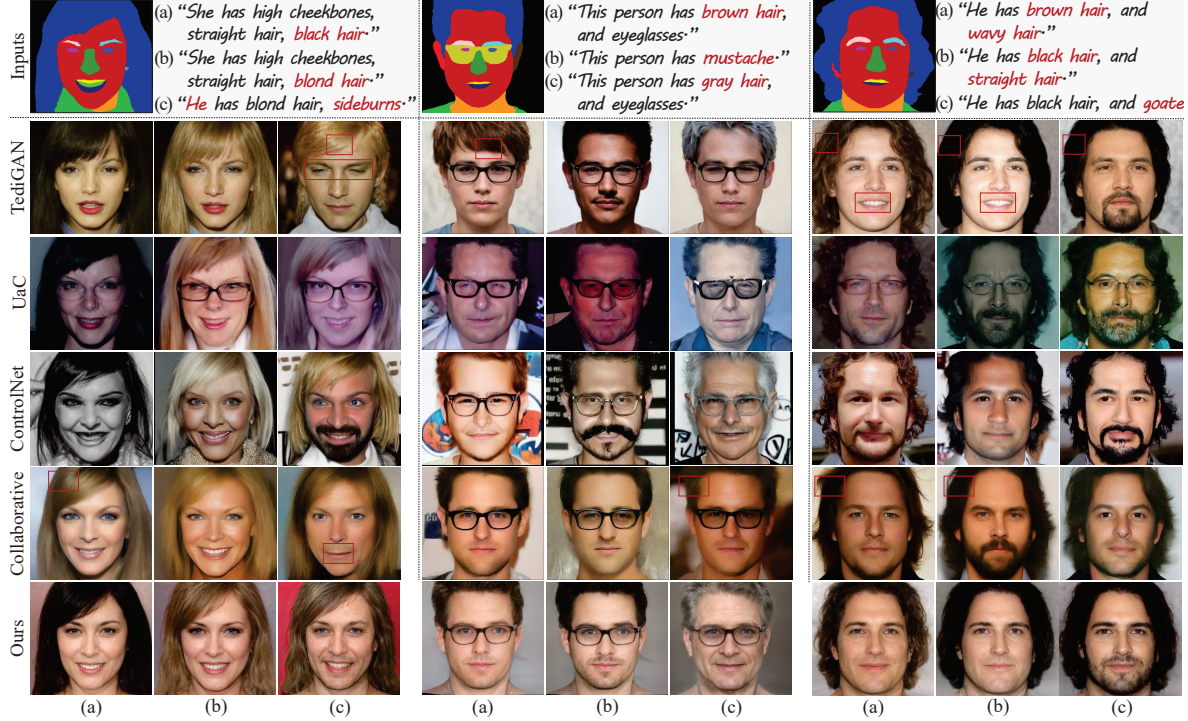


Figure 5. Visual examples of the 2D face image generation using a text prompt and a semantic mask. For each semantic mask, we use three different text prompts (a)-(c), resulting in different output images (a)-(c).

3.4. Loss Functions

To optimize \mathcal{M} and \mathcal{T} , we use reconstruction loss, perceptual loss, and identity loss for image generation, and regularization loss [42] that encourages the latent codes to be closer to the average latent code $\bar{\mathbf{w}}$.

For training \mathcal{M} , we use the GT image I^{gt} as reference to encourage the latent code \mathbf{w}_t^m to generate a photo-realistic image as follows:

$$\begin{aligned} \mathcal{L}_{\mathcal{M}} = & \lambda_0^m \|I^{gt} - \mathcal{G}(\mathbf{w}_t^m)\|_2 + \\ & \lambda_1^m \|\mathcal{F}(I^{gt}) - \mathcal{F}(\mathcal{G}(\mathbf{w}_t^m))\|_2 + \\ & \lambda_2^m (1 - \cos(\mathcal{R}(I^{gt}), \mathcal{R}(\mathcal{G}(\mathbf{w}_t^m)))) + \\ & \lambda_3^m \|\mathcal{E}(\mathbf{z}_t, t, \mathbf{x}, c) - \bar{\mathbf{w}}\|_2, \end{aligned} \quad (6)$$

where $\mathcal{R}(\cdot)$ is pre-trained ArcFace network [8], $\mathcal{F}(\cdot)$ is the feature extraction network [63], \mathbf{z}_t is noisy image, and the hyper-parameters $\lambda_{(\cdot)}^s$ guide the effect of losses. Note that we freeze \mathcal{T} while training \mathcal{M} .

For training \mathcal{T} , we use I_0^d produced by the encoder \mathcal{E} into the reconstruction and perceptual losses. With these losses, the loss $\mathcal{L}_{\mathcal{T}}$ encourages the network to control facial attributes while preserving the identity of I^{gt} :

$$\begin{aligned} \mathcal{L}_{\mathcal{T}} = & \lambda_0^s \|I_0^d - \mathcal{G}(\mathbf{w}_t)\|_2 + \\ & \lambda_1^s \|\mathcal{F}(I_0^d) - \mathcal{F}(\mathcal{G}(\mathbf{w}_t))\|_2 + \\ & \lambda_2^s (1 - \cos(\mathcal{R}(I^{gt}), \mathcal{R}(\mathcal{G}(\mathbf{w}_t)))) + \\ & \lambda_3^s \|\mathcal{E}(\mathbf{z}_t, t, \mathbf{x}, c) - \bar{\mathbf{w}}\|_2, \end{aligned} \quad (7)$$

where the hyper-parameters $\lambda_{(\cdot)}^s$ guide the effect of losses. Similar to Equation 6, we freeze \mathcal{M} while training \mathcal{T} .

We further introduce a multi-step training strategy that considers the evolution of the feature representation in \mathcal{E} over the denoising steps. We observe that \mathcal{E} tends to focus more on text-relevant features in an early step, $t = T$, and structure-relevant features in a later step, $t = 0$. Figure 3 (b) shows the attention maps $\bar{\mathbf{A}}$ showing variations across the denoising step. As the attention map, we can capture the textual and structural features by varying the denoising steps. To effectively capture the semantic details of multi-modal conditions, our model is trained across multiple denoising steps.

4. Experiments

4.1. Experimental Setup

We use ControlNet [62] as the diffusion-based encoder that receives multi-modal conditions, including text and visual conditions such as a semantic mask and scribble map. The StyleGAN [22] and EG3D [4] are exploited as pre-trained 2D and 3D GAN, respectively. See the Supplementary Material for the training details, the network architecture, and additional results.

Datasets. We employ the CelebAMask-HQ [29] dataset comprising 30,000 face RGB images and annotated semantic masks, including 19 facial-component categories such as *skin*, *eyes*, *mouth*, and *etc.* We also use textual de-

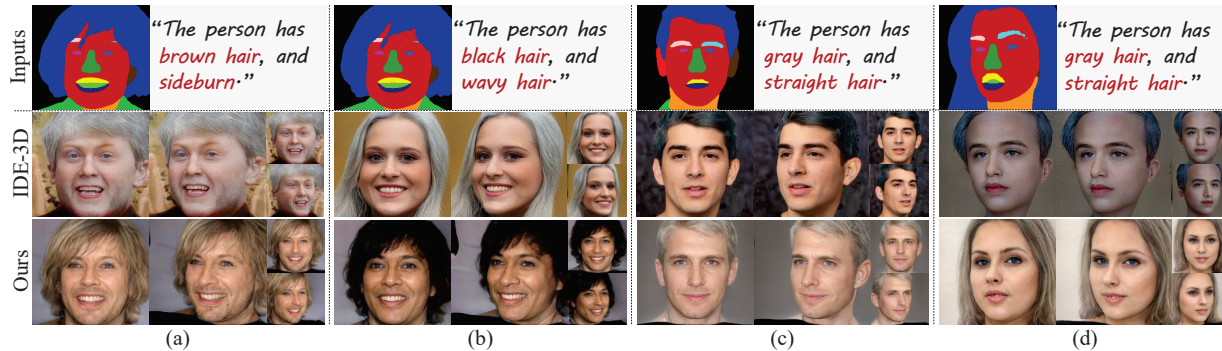


Figure 6. Visual examples of the 3D-aware face image generation using a text and a semantic mask. We show the images generated with inputs and arbitrary viewpoints.

| Input conditions | Method | Model | Domain | FID↓ | LPIPS↓ | SSIM↑ | ID↑ | ACC↑ | mIoU↑ |
|-------------------------|--------------------|-----------|--------|--------------|-------------|-------------|-------------|--------------|--------------|
| Text + semantic mask | TediGAN [58] | GAN | 2D | 54.83 | 0.31 | 0.62 | 0.63 | 81.68 | 40.01 |
| | IDE-3D [51] | GAN | 3D | 39.05 | 0.40 | 0.41 | 0.54 | 47.07 | 10.98 |
| | UaC [35] | Diffusion | 2D | 45.87 | 0.38 | 0.59 | 0.32 | 81.49 | 42.68 |
| | ControlNet [62] | Diffusion | 2D | 46.41 | 0.41 | 0.53 | 0.30 | 82.42 | 42.77 |
| | Collaborative [19] | Diffusion | 2D | 48.23 | 0.39 | 0.62 | 0.31 | 74.06 | 30.69 |
| | Ours | GAN | 2D | 46.68 | 0.30 | 0.63 | 0.76 | 83.41 | 43.82 |
| | Ours | GAN | 3D | 44.91 | 0.28 | 0.64 | 0.78 | 83.05 | 43.74 |
| Text + scribble map | ControlNet [62] | Diffusion | 2D | 93.26 | 0.52 | 0.25 | 0.21 | - | - |
| | Ours | GAN | 2D | 55.60 | 0.32 | 0.56 | 0.72 | - | - |
| | Ours | GAN | 3D | 48.76 | 0.34 | 0.49 | 0.62 | - | - |

Table 1. Quantitative results of multi-modal face image generation on CelebAMask-HQ [29] with annotated text prompts [58].

descriptions provided by [58] describing the facial attributes, such as *black hair*, *sideburns*, and *etc.*, corresponding to the CelebAMask-HQ dataset. For the face image generation task using a scribble map, we obtain the scribble maps by applying PiDiNet [49, 50] to the RGB images in CelebAMask-HQ. We additionally compute camera parameters based on [4, 10] for 3D-aware image generation.

Comparisons. We compare our method with GAN-based models, such as TediGAN [58] and IDE-3D [51], and DM-based models, such as Unite and Conquer (UaC) [35], ControlNet [62], and Collaborative diffusion (Collaborative) [19], for face generation task using a semantic mask and a text prompt. IDE-3D is trained by a CLIP loss term like TediGAN to apply a text prompt for 3D-aware face image generation. ControlNet is used for face image generation using a text prompt and a scribble map. We use the official codes provided by the authors, and we downsample the results into 256×256 for comparison.

Evaluation Metrics. For quantitative comparisons, we evaluate the image quality and semantic consistency using sampled 2k semantic mask- and scribble map-text prompt pairs. Frechet Inception Distance (FID) [17], LPIPS [63], and the Multiscale Structural Similarity (MS-SSIM) [56] are employed for the evaluation of visual quality and diversity, respectively. We also compute the ID similarity mean score (ID) [8, 57] before and after applying a text prompt.

Additionally, we assess the alignment accuracy between the input semantic masks and results using mean Intersection-over-Union (mIoU) and pixel accuracy (ACC) for the face generation task using a semantic mask.

4.2. Results

Qualitative Evaluations. Figure 5 shows the visual comparisons between ours and two existing methods for 2D face image generation using a text prompt and a semantic mask as input. We use the same semantic mask with different text prompts (a)-(c). TediGAN produces results consistent with the text prompt as the latent codes are optimized using the input text prompt. However, the results are inconsistent with the input semantic mask, as highlighted in the red boxes. UaC shows good facial alignment with the input semantic mask, but the results are generated with unexpected attributes, such as glasses, that are not indicated in the inputs. Collaborative and ControlNet produce inconsistent, blurry, and unrealistic images. Our model is capable of preserving semantic consistency with inputs and generating realistic facial images. As shown in Figure 5, our method preserves the structure of the semantic mask, such as the hairline, face position, and mouth shape, while changing the attributes through a text prompt.

Figure 6 compares our method with IDE-3D [51] to validate the performance of 3D-aware face image generation

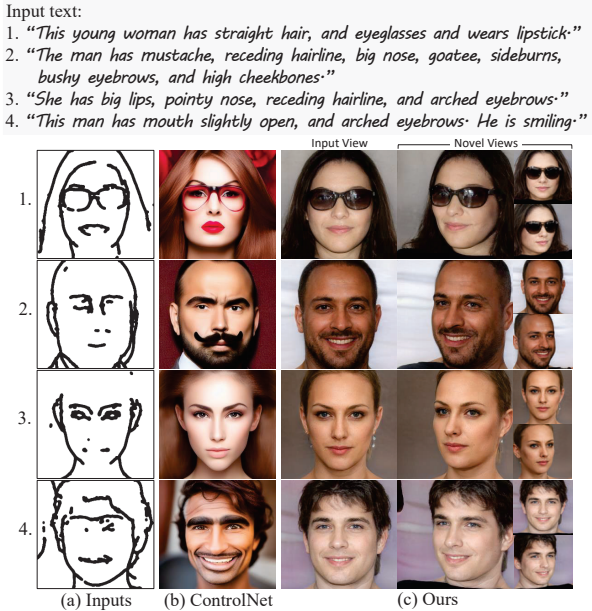


Figure 7. Visual examples of 3D-aware face image generation using text prompts and scribble maps. Using (1-4) the text prompts and their corresponding (a) scribble maps, we compare the results of (b) ControlNet with (c) multi-view images generated by ours.

using a semantic mask and a text prompt. We use the same semantic mask with different text prompts in Figures 6 (a) and (b), and use the same text prompt with different semantic masks in Figures 6 (c) and (d). The results of IDE-3D are well aligned with the semantic mask with the frontal face. However, IDE-3D fails to produce accurate results when the non-frontal face mask is used as input. Moreover, the results cannot reflect the text prompt. Our method can capture the details provided by input text prompts and semantic masks, even in a 3D domain.

Figure 7 shows visual comparisons with ControlNet on 2D face generation from a text prompt and a scribble map. The results from ControlNet and our method are consistent with both the text prompt and the scribble map. ControlNet, however, tends to over-emphasize the characteristic details related to input conditions. Our method can easily adapt to the pre-trained 3D GAN and produce photo-realistic multi-view images from various viewpoints.

Quantitative Evaluations. Table 1 reports the quantitative results on CelebAMask-HQ with text prompts [58]. Our method using text prompts and semantic masks shows performance increases in all metrics in 2D and 3D domains, compared with TediGAN and UaC. Our model using 2D GAN significantly improves LPIPS, ID, ACC, and mIoU scores, surpassing TediGAN, UaC, ControlNet, and Collaborative, respectively. It demonstrates our method’s strong ability to generate photo-realistic images while reflecting input multi-modal conditions better. For 3D-aware face image generation using a text prompt and a semantic mask, it

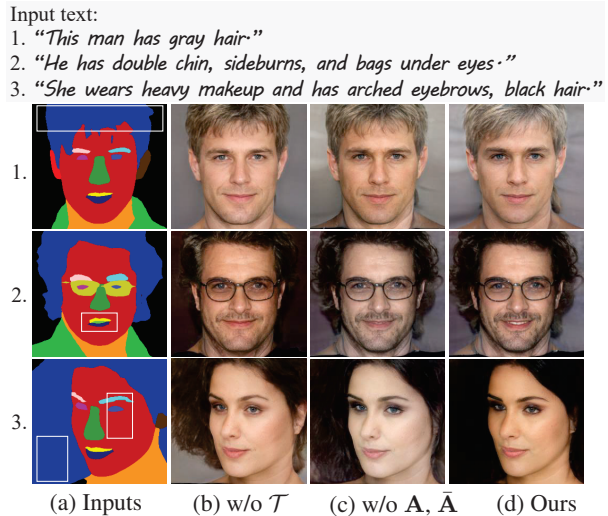


Figure 8. Effect of \mathcal{M} and \mathcal{T} . (b) shows the results using only \mathcal{M} , and (c) shows the effect of the cross-attention maps (\mathbf{A} and $\bar{\mathbf{A}}$) in \mathcal{T} . The major changes are highlighted with the white boxes.

| Method | \mathcal{M} | \mathcal{T} | \mathbf{A}_t | I^{gt} | I_0^d | FID↓ | LPIPS↓ | ID↑ | ACC↑ |
|--------|---------------|---------------|----------------|----------|---------|--------------|-------------|-------------|--------------|
| (a) | ✓ | | | ✓ | ✓ | 62.08 | 0.29 | 0.62 | 81.09 |
| (b) | ✓ | ✓ | | ✓ | ✓ | 48.68 | 0.28 | 0.66 | 82.86 |
| (c) | ✓ | ✓ | ✓ | | ✓ | 54.27 | 0.31 | 0.58 | 80.58 |
| (d) | ✓ | ✓ | ✓ | ✓ | | 61.60 | 0.29 | 0.62 | 80.04 |
| (e) | ✓ | ✓ | ✓ | ✓ | ✓ | 44.91 | 0.28 | 0.78 | 83.05 |

Table 2. Ablation analysis on 3D-aware face image generation using a text prompt and a semantic mask. We compare (a) and (b) with (e) to show the effect of our style modulation network and (c) and (d) with (e) to analyze the effect of I^{gt} and I^d in model training.

is reasonable that IDE-3D shows the highest FID score as the method additionally uses an RGB image as input to estimate the latent code for face generation. The LPIPS, SSIM, and ID scores are significantly higher than IDE-3D, with scores higher by 0.116, 0.23, and 0.24, respectively. Our method using 3D GAN exhibits superior ACC and mIoU scores for the 3D face generation task compared to IDE-3D, with the score difference of 35.98% and 32.76%, likely due to its ability to reflect textual representations into spatial information. In face image generation tasks using a text prompt and a scribble map, our method outperforms ControlNet in FID, LPIPS, SSIM, and ID scores in both 2D and 3D domains. Note that the ACC and mIoU scores are applicable for semantic mask-based methods.

4.3. Ablation Study

We conduct ablation studies to validate the effectiveness of our contributions, including the mapping network \mathcal{M} , the AbSM network \mathcal{T} , and the loss functions $\mathcal{L}_{\mathcal{M}}$ and $\mathcal{L}_{\mathcal{T}}$.

Effectiveness of \mathcal{M} and \mathcal{T} . We conduct experiments with different settings to assess the effectiveness of \mathcal{M} and \mathcal{T} .

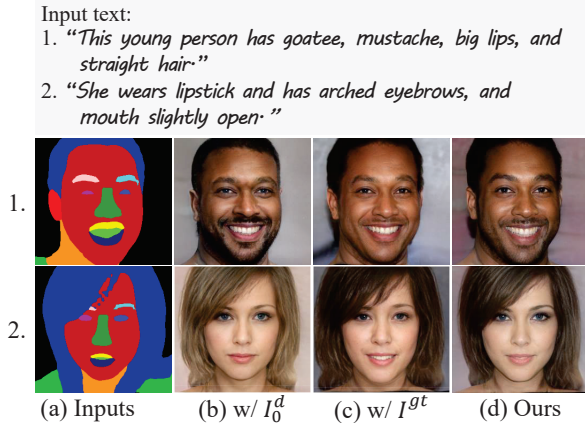


Figure 9. Effect of using I^d from the denoising U-Net and the GT image I^{gt} in model training. Using text prompts (1, 2) with (a) the semantic mask, we show face images using our model trained with (b) I_0^d , (c) I^{gt} , and (d) both.

We also show the advantages of using cross-attention maps in our model. The quantitative and qualitative results are presented in Table 2 and Figure 8, respectively. When using only \mathcal{M} , we can generate face images that roughly preserve the structures of a given semantic mask in Figure 8 (a), including the outline of the facial components (e.g. face, eye) in Figure 8 (b). On the other hand, \mathcal{T} enables the model to express face attribute details effectively, such as hair colors and mouth open, based on the multi-modal inputs in Figure 8 (c). The FID and ACC scores are higher than the model using only \mathcal{M} in Table 2 (b). We further present the impact of adopting cross-attention maps to \mathcal{T} for style modulation. Figure 8 (d) shows how the attention-based modulation approach enhances the quality of results, particularly in terms of the sharpness of desired face attributes and the overall consistency between the generated image and multi-modal conditions. Table 2 (e) demonstrates the effectiveness of our method by showing improvements in FID, LPIPS, ID, and ACC. Our method, including both \mathcal{M} and \mathcal{T} with cross-attention maps, significantly improves the FID showing our model’s ability to generate high-fidelity images. From the improvement of the ID score, the cross-attention maps enable relevantly applying the details of input conditions to facial components.

Model Training. We analyze the effect of loss terms $\mathcal{L}_{\mathcal{M}}$ and $\mathcal{L}_{\mathcal{T}}$ by comparing the performance with the model trained using either I_0^d from the denoising U-Net or GT image I^{gt} . The model trained using I_0^d produces the images in Figure 9 (b), which more closely reflected the multi-modal conditions (a), such as “goatee” and “hair contour”. In Table 2 (c), the ACC score of this model is higher than the model trained only using I^{gt} in Table 2 (d). The images generated by the model trained with I^{gt} in Figure 9 (c) are more perceptually realistic, as evidenced by the lower LPIPS score compared to the model trained with I_0^d in Ta-

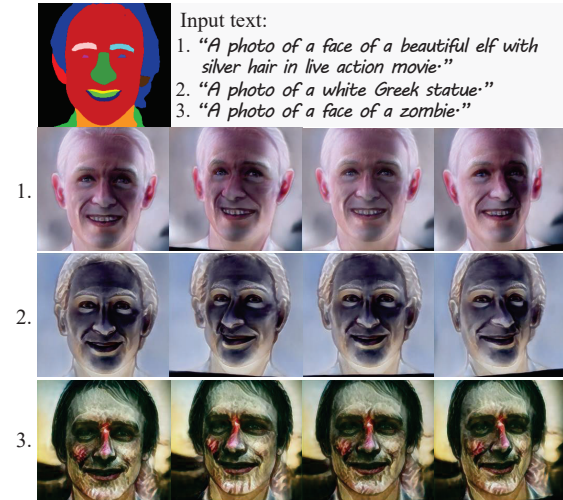


Figure 10. Visual examples of 3D face style transfer. Our method generates stylized multi-view images by mapping the latent features of DM and GAN.

ble 2 (c) and (d). Using I^{gt} also preserves more condition-irrelevant features inferred by the ID scores in Table 2 (c) and (d). In particular, our method combines the strengths of two models as shown in Figure 9 (d) and Table 2 (e).

4.4. Limitations and Future Works

Our method can be extended to multi-modal face style transfer (e.g. *face* \rightarrow *Greek statue*) by mapping the latent spaces of DM and GAN without CLIP losses and additional dataset, as shown in Figure 10. For the 3D-aware face style transfer task, we train our model using I_0^d that replaces GT image I^{gt} in our loss terms. This method, however, is limited as it cannot transfer extremely distinct style attributes from the artistic domain to the photo-realistic domain of GAN. To better transfer the facial style in the 3D domain, we will investigate methods to map the diffusion features related to the input pose into the latent space of GAN in future works.

5. Conclusion

We presented the diffusion-driven GAN inversion method that translates multi-modal inputs into photo-realistic face images in 2D and 3D domains. Our method interprets the pre-trained GAN’s latent space and maps the diffusion features into this latent space, which enables the model to easily adopt multi-modal inputs, such as a visual input and a text prompt, for face image generation. We also proposed to train our model across the multiple denoising steps, which further improves the output quality and consistency with the multiple inputs. We demonstrated the capability of our method by using text prompts with semantic masks or scribble maps as input for 2D or 3D-aware face image generation and style transfer.

References

- [1] Yuval Alaluf, Or Patashnik, and Daniel Cohen-Or. Restyle: A residual-based stylegan encoder via iterative refinement. In *ICCV*, 2021. 2
- [2] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *CVPR*, 2022. 2
- [3] Sam Bond-Taylor, Peter Hessey, Hiroshi Sasaki, Toby P Breckon, and Chris G Willcocks. Unleashing transformers: Parallel token prediction with discrete absorbing diffusion for fast high-resolution image generation from vector-quantized codes. In *ECCV*, 2022. 2
- [4] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In *CVPR*, 2022. 2, 5, 6
- [5] Shu-Yu Chen, Wanchao Su, Lin Gao, Shihong Xia, and Hongbo Fu. Deepfacedrawing: Deep generation of face images from sketches. *TOG*, 39(4):72–1, 2020. 2
- [6] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, 2018. 1
- [7] Min Jin Chong and David Forsyth. Jojogan: One shot face stylization. In *ECCV*, 2022. 1, 2
- [8] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. 5, 6
- [9] Kangle Deng, Gengshan Yang, Deva Ramanan, and Jun-Yan Zhu. 3d-aware conditional image synthesis. In *CVPR*, 2023. 2
- [10] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *CVPR Workshops*, 2019. 6
- [11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *NeurIPS*, 2021. 1
- [12] Hoseok Do, EunKyung Yoo, Taehyeong Kim, Chul Lee, and Jin Young Choi. Quantitative manipulation of custom attributes on 3d-aware image synthesis. In *CVPR*, 2023. 1
- [13] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, 2021. 2
- [14] Rinon Gal, Or Patashnik, Haggai Maron, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators. *TOG*, 41(4): 1–13, 2022. 2
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *NeurIPS*, 2014. 1
- [16] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 2
- [17] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*, 2017. 6
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 2020. 1
- [19] Ziqi Huang, Kelvin CK Chan, Yuming Jiang, and Ziwei Liu. Collaborative diffusion for multi-modal face generation and editing. In *CVPR*, 2023. 2, 6
- [20] Jaeseok Jeong, Mingi Kwon, and Youngjung Uh. Training-free style transfer emerges from h-space in diffusion models. In *WACV*, 2024. 2
- [21] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 1, 2
- [22] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020. 1, 2, 5
- [23] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. In *CVPR*, 2023. 1
- [24] Gwanghyun Kim and Se Young Chun. Datid-3d: Diversity-preserved domain adaptation using text-to-image diffusion for 3d generative model. In *CVPR*, 2023. 2
- [25] Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *CVPR*, 2022. 2
- [26] Yunji Kim, Jiyoung Lee, Jin-Hwa Kim, Jung-Woo Ha, and Jun-Yan Zhu. Dense text-to-image generation with attention modulation. In *ICCV*, 2023. 2
- [27] Jaehoon Ko, Kyusun Cho, Daewon Choi, Kwangrok Ryoo, and Seungryong Kim. 3d gan inversion with pose optimization. In *WACV*, 2023. 2
- [28] Gihyun Kwon and Jong Chul Ye. Diffusion-based image translation using disentangled style and content representation. *arXiv preprint arXiv:2209.15264*, 2022. 2, 3
- [29] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. In *CVPR*, 2020. 5, 6
- [30] Ming Liu, Yukang Ding, Min Xia, Xiao Liu, Errui Ding, Wangmeng Zuo, and Shilei Wen. Stgan: A unified selective transfer network for arbitrary image attribute editing. In *CVPR*, 2019. 1
- [31] Zhian Liu, Maomao Li, Yong Zhang, Cairong Wang, Qi Zhang, Jue Wang, and Yongwei Nie. Fine-grained face swapping via regional gan inversion. In *CVPR*, 2023. 2
- [32] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *CVPR*, 2021. 2
- [33] Tianxiang Ma, Kang Zhao, Jianxin Sun, Jing Dong, and Tieniu Tan. Free-style and fast 3d portrait synthesis. *arXiv preprint arXiv:2306.15419*, 2023. 2
- [34] Chenlin Meng, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. 2

- [35] Nithin Gopalakrishnan Nair, Wele Gedara Chaminda Bandara, and Vishal M Patel. Unite and conquer: Plug & play multi-modal synthesis using diffusion models. In *CVPR*, 2023. 2, 6
- [36] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 1, 2
- [37] Xingang Pan, Ayush Tewari, Thomas Leimkühler, Lingjie Liu, Abhimitra Meka, and Christian Theobalt. Drag your gun: Interactive point-based manipulation on the generative image manifold. In *SIGGRAPH 2023*, 2023. 1
- [38] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, 2019. 2
- [39] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *ICCV*, 2021. 1, 2
- [40] Malsha V Perera and Vishal M Patel. Analyzing bias in diffusion-based face generation models. *arXiv preprint arXiv:2305.06402*, 2023. 1
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*. PMLR, 2021. 1, 2
- [42] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *CVPR*, 2021. 1, 3, 5
- [43] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2
- [44] Rohit Saha, Brendan Duke, Florian Shkurti, Graham W Taylor, and Parham Aarabi. Loho: Latent optimization of hairstyles via orthogonalization. In *CVPR*, 2021. 2
- [45] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *NeurIPS*, 2022. 1, 2
- [46] Yujun Shen, Jinjin Gu, Xiaou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *CVPR*, 2020. 1
- [47] J Ryan Shue, Eric Ryan Chan, Ryan Po, Zachary Ankner, Jiajun Wu, and Gordon Wetzstein. 3d neural field generation using triplane diffusion. In *CVPR*, 2023. 2
- [48] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021. 1
- [49] Zhuo Su, Matti Pietikäinen, and Li Liu. Bird: Learning binary and illumination robust descriptor for face recognition. In *BMVC*, 2019. 6
- [50] Zhuo Su, Wenzhe Liu, Zitong Yu, Dewen Hu, Qing Liao, Qi Tian, Matti Pietikäinen, and Li Liu. Pixel difference networks for efficient edge detection. In *ICCV*, 2021. 6
- [51] Jingxiang Sun, Xuan Wang, Yichun Shi, Lizhen Wang, Jue Wang, and Yebin Liu. Ide-3d: Interactive disentangled editing for high-resolution 3d-aware portrait synthesis. *TOG*, 41: 1–10, 2022. 1, 2, 6
- [52] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *TOG*, 40(4):1–14, 2021. 2
- [53] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *CVPR*, 2023. 1, 3
- [54] Andrey Voynov, Kfir Aberman, and Daniel Cohen-Or. Sketch-guided text-to-image diffusion models. In *SIGGRAPH*, 2023. 1
- [55] Weilun Wang, Jianmin Bao, Wengang Zhou, Dongdong Chen, Dong Chen, Lu Yuan, and Houqiang Li. Semantic image synthesis via diffusion models. *arXiv preprint arXiv:2207.00050*, 2022. 2
- [56] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers*, 2003. 6
- [57] Tianyi Wei, Dongdong Chen, Wenbo Zhou, Jing Liao, Zhen-tao Tan, Lu Yuan, Weiming Zhang, and Nenghai Yu. Hairclip: Design your hair by text and reference image. In *CVPR*, 2022. 2, 6
- [58] Weihao Xia, Yujiu Yang, Jing-Hao Xue, and Baoyuan Wu. Tedigan: Text-guided diverse face image generation and manipulation. In *CVPR*, 2021. 1, 2, 6, 7
- [59] Xin Yang, Xiaogang Xu, and Yingcong Chen. Out-of-domain gan inversion via invertibility decomposition for photo-realistic human face manipulation. In *ICCV*, 2023. 2
- [60] Fei Yin, Yong Zhang, Xuan Wang, Tengfei Wang, Xiaoyu Li, Yuan Gong, Yanbo Fan, Xiaodong Cun, Ying Shan, Cengiz Oztireli, et al. 3d gan inversion with facial symmetry prior. In *CVPR*, 2023. 2
- [61] Yu Yin, Kamran Ghasedi, HsiangTao Wu, Jiaolong Yang, Xin Tong, and Yun Fu. Nerfinvator: High fidelity nerf-gan inversion for single-shot real image animation. In *CVPR*, 2023. 2
- [62] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023. 2, 3, 5, 6
- [63] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5, 6
- [64] Yuxin Zhang, Nisha Huang, Fan Tang, Haibin Huang, Chongyang Ma, Weiming Dong, and Changsheng Xu. Inversion-based style transfer with diffusion models. In *CVPR*, 2023. 1
- [65] Hang Zhou, Yu Liu, Ziwei Liu, Ping Luo, and Xiaogang Wang. Talking face generation by adversarially disentangled audio-visual representation. In *AAAI*, 2019. 2
- [66] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *ECCV*, 2020. 2
- [67] Peihao Zhu, Rameen Abdal, John Femiani, and Peter Wonka. Barbershop: Gan-based image compositing using segmentation masks. *arXiv preprint arXiv:2106.01505*, 2021. 2