FACEMUG: A Multimodal Generative and Fusion Framework for Local Facial Editing

Wanglong Lu, Jikai Wang, Xiaogang Jin, Xianta Jiang, and Hanli Zhao*

Abstract-Existing facial editing methods have achieved remarkable results, yet they often fall short in supporting multimodal conditional local facial editing. One of the significant evidences is that their output image quality degrades dramatically after several iterations of incremental editing, as they do not support local editing. In this paper, we present a novel multimodal generative and fusion framework for globally-consistent local facial editing (FACEMUG) that can handle a wide range of input modalities and enable fine-grained and semantic manipulation while remaining unedited parts unchanged. Different modalities, including sketches, semantic maps, color maps, exemplar images, text, and attribute labels, are adept at conveying diverse conditioning details, and their combined synergy can provide more explicit guidance for the editing process. We thus integrate all modalities into a unified generative latent space to enable multimodal local facial edits. Specifically, a novel multimodal feature fusion mechanism is proposed by utilizing multimodal aggregation and style fusion blocks to fuse facial priors and multimodalities in both latent and feature spaces. We further introduce a novel self-supervised latent warping algorithm to rectify misaligned facial features, efficiently transferring the pose of the edited image to the given latent codes. We evaluate our FACEMUG through extensive experiments and comparisons to state-of-the-art (SOTA) methods. The results demonstrate the superiority of FACEMUG in terms of editing quality, flexibility, and semantic control, making it a promising solution for a wide range of local facial editing tasks.

Index Terms—Generative adversarial networks, image-toimage translation, multimodal fusion, image editing, facial editing.

T HE rapid development of digital imaging and mobile computing has fueled the demand for personalized content in social media and various applications [1], [2], [3], making facial image editing an essential research area in computer graphics and computer vision. Faces are universally acknowledged as the most representative and expressive aspect of human beings, which makes facial editing a challenging task [4].

Many image editing tools provide convenient guidance information to allow users to edit facial features interactively [5]. In order to provide convenient user interfaces, many face image editing tools make use of various input modalities to guide the editing of facial features. In recent years, multimodal facial image editing has attracted considerable interest. Since multimodal models excel at conveying various types of conditioning information, their combined synergy can offer clearer descriptions to aid in facial editing. For example, semantics can define a facial image's coarse layout; sketches can detail its structure and texture; and text or attribute labels can adjust facial attributes, to name just a few. The support of local editing capability is also important in a variety of image editing applications. Local editing enables users to edit local image regions in an incremental manner while keeping the contents of unedited background regions unchanged.

There are some limitations in existing multimodal local facial editing methods. The first limitation is the difficulty in maintaining visual contents in unedited background regions. Existing multimodal facial editing techniques [6], [7], [8], [9] can only edit the facial image as a whole and are prone to introduce unwanted changes to unedited background regions. When users are not satisfied with some local effects, these techniques will fail to edit the local regions in an incremental manner. As shown in Fig. 1 (bottom two rows), while stateof-the-art (SOTA) methods may perform high-quality edits, they are likely to include unwanted changes of other facial features in incremental editing scenarios, where an already edited image is subject to further modifications using different modalities. Some existing methods [6], [7] have the limitation in manual annotations of paired data. These methods train their models with labeled paired data across different modalities but manual annotations of training datasets are label-intensive. Recent methods based on diffusion models [8], [9] instead train all uni-modal models first and perform multimodal facial editing by integrating these pre-trained uni-modal models. However, when the number of modalities grows, more unimodal models should be trained separately with these methods.

With this in mind, we explored ways to tackle these limitations. We investigated whether incorporating generative adversarial networks (GANs) [10] can improve global consistency for multimodal local facial editing. By learning the distribution of real facial images, adversarial training enforces the model to fill plausible contents for edited regions guided by multimodalities. To minimize the dependency on paired training data, we asked the question that if we can loosen the ties between the paired modalities data by aligning all modalities into a unified generative latent space to diminish the requirement for paired text, attribute label, and exemplar modalities. Instead of training a uni-modal model for each modality, we examined whether fusion and warping priors, along with multimodalities in both latent and feature space, could achieve seamless integration of multimodalities.

We thus introduce a **MU**ltimodal Generative and fusion framework for local **FAC**ial Editing (FACEMUG), which can solve the above problems. First, since the StyleGAN latent

W. Lu, J. Wang, and H. Zhao are with the College of Computer Science and Artificial Intelligence, Wenzhou University, Wenzhou 325035, China.

W. Lu and X. Jiang are with the Department of Computer Science, Memorial University of Newfoundland, St. John's, NL A1B 3X5, Canada.

X. Jin is with the State Key Laboratory of CAD&CG, Zhejiang University, Hangzhou 310058, China.

^{*} Corresponding author. E-mail: hanlizhao@wzu.edu.cn



Fig. 1. Examples demonstrating the superior performance of FACEMUG in high-quality globally consistent local facial editing, using subsets of the five modalities including semantic label, sketch, text, color, and exemplar image. Our method (top row) exhibits better visual quality and fidelity in incremental editing (the later editing taking the previous output image as input), compared to SOTA multimodal face editing methods: ColDiffusion [8] (middle row) and Unite&Conquer [9] (bottom row).

space [11] is disentangled well, we design our framework by bridging all modalities to the StyleGAN latent space. Second, for the seamless integration of multimodalities, we design our multimodal generator with fusion and warping in latent and feature space. To support the heterogeneity and sparsity of the pixel-wise conditional inputs, we aggregate multimodal conditional inputs into a homogeneous feature space. Since a fully trained GAN model excels at capturing rich textures and structural priors [12], we utilize a StyleGAN generator as a facial feature bank to provide candidate facial features and introduce style fusion blocks to fuse facial features for improving the generation quality. Moreover, to rectify the pose misalignment between the edited image and the given latent codes, we present a self-supervised latent warping method to efficiently transfer the pose of the edited image to that of the given latent codes in the latent space. To simulate the latent editing process during training and boost facial editing capabilities, a diversity-enhanced attribute loss is proposed.

To the best of our knowledge, our FACEMUG is the first method that generates realistic facial features in response to multimodal inputs on the edited regions while maintaining visual coherence with the unedited background to achieve global consistency. We have conducted extensive comparisons of FACEMUG against the SOTA methods and comprehensive experiments to demonstrate the superiority of FACEMUG in terms of editing quality, flexibility, and semantic control, illustrating its potential to significantly enhance various applications within facial editing.

In summary, our paper makes the following contributions:

- A novel globally-consistent local facial editing framework that enables diverse facial attribute manipulation.
- A novel multimodal feature fusion mechanism that utilizes multimodal aggregation and style fusion blocks to fuse facial features in both latent and feature spaces.

- A novel latent warping algorithm automatically aligns facial poses between edited and exemplar images in latent space, without relying on annotated labels or pose detection models.
- Our novel framework would benefit numerous practical applications, supporting incremental editing scenarios guided by multimodities (sketches, semantic maps, color maps, exemplar images, text, and attribute labels).

I. RELATED WORK

A. Generative image synthesis

GANs have been at the forefront of image synthesis in recent years credited to their ability to efficiently sample high-resolution images with good perceptual quality [11]. The latent space learned by these networks enables an intuitive approach to controlling the image generation process, thereby facilitating semantic manipulation. Advancements in disentangled latent space learning, exemplified by developments in Fader Networks [13] and StyleGANs [14], [11], [15], have made the manipulation of latent codes in pre-trained GAN models an active area of research in image editing.

Another category of methods known as Diffusion Probabilistic Models [16], [17] has demonstrated remarkable potential in enhancing image synthesis quality [18]. For instance, Latent Diffusion Models (LDMs, Stable Diffusion) [19] were introduced to achieve better visual quality on image inpainting and class-conditional image synthesis, while significantly reducing computational requirements. Even though the sampling speed can be partially solved by sampling strategies [20] and hierarchical approaches [19], [21], multiple denoising steps are still needed for high-quality synthesis. Compared to diffusion probabilistic models [16], [17] with multiple denoising steps, FACEMUG ensures rapid inference without compromising the quality of the edited images.



Fig. 2. Overall pipeline of our FACEMUG globally-consistent local facial editing: the given attribute label (or text), random latent code z, and exemplar image I_{ex} are first processed through the exemplar style module, latent warping module, and the latent attribute editing module to get the edited latent codes. Simultaneously, the input pixel-wise multimodal inputs \mathcal{X} and a binary mask \mathbf{M} are fed into the multimodal aggregation module and the multimodal generator to get an edited realistic face image I_{out} , where the manipulation of the masked regions in \mathbf{M} is guided by multimodal inputs.

B. Facial attribute manipulation

Facial attribute manipulation is a crucial task in computer graphics and computer vision, focusing on the modification and control of various facial features, such as skin texture [22], structure [23], and expressions [24]. Existing image inpainting techniques [25], [26], [27] can directly remove unwanted or blemished facial features, replacing pixels in masked regions with plausible content. However, these techniques are limited in terms of controllable editing ability and do not support multimodal conditional editing. Low-level facial features such as texture and geometry can be manipulated utilizing auxiliary information, such as sketches [28], [22], colors [22], foreground contours [29], and structures [30]. However, these methods may overfit limited guidance information, requiring professional skills for semantic editing, such as facial expression and identity.

For semantic-level face manipulation, some methods [31], [32] can control a set of attributes using a domain label to index the mapped latent codes. However, they are limited to pre-defined attributes, thereby restricting editing freedom. Geometry-guided face manipulation methods mainly use semantic geometry, such as sketches [33], [34], [35], semantic maps [36], [37], [38], [39], [40], [41] to guide the generation of facial structure. However, due to information loss during the projection and reconstruction process between real photographs and corresponding latent representations, these methods might inadvertently alter fine facial details (i.e., unedited regions may be changed). Recent advancements allow the transfer of facial attributes from example images at the instance level [42], [43], [44], [45], [46], [47], manipulating attributes such as expression, identity, and decorative elements. However, these methods lack flexibility in selecting facial regions for local editing. Image-composition-based methods enable local editing by overlaying the foreground of a source image onto the background of a target image [48], [49], [50]. Although they can generate more diverse and realistic facial images, their method may fail when the poses in the edited and reference images differ.

Manipulating latent codes in pre-trained GAN models is another editing branch. Studies have explored semantically meaningful paths for editing latent codes through supervised [23], [51], [52], [53], [54] and unsupervised learning methods [55], [56]. With the advantages of StyleGAN's disentanglement in latent space, text descriptions employed by StyleCLIP [57], FFCLIP [58], and CLIPInverter [59] are gaining popularity for text-based image manipulation. These methods relied on image inversion techniques, might unintentionally alter unedited areas due to information loss in the inversion process, even though existing optimizationbased [60], [61], [62] and encoder-based [63], [33], [64], [65], [66] GAN inversion methods can alleviate this issue. In contrast, our method can manipulate local facial attributes with multimodal conditional information to produce realistic locally edited face images with fine-grained and semantic manipulation, preserving unaltered regions.

C. Multimodal image synthesis

Since different input modalities are well suited for manipulating different aspects of facial attributes, leading to recent studies exploring the multimodal image synthesis [67], [68], [69], [70]. Recently, TediGAN [6] leverages a pretrained unconditional generator for high-fidelity multimodal facial editing. However, its approach of fusing multimodalities solely in the latent space limits its responsiveness to specific input modalities, particularly for detailed facial features. PoE-GAN [7] improves the visual quality and control abilities using a product-of-experts inference framework to learn the image distribution conditioned on multimodalities. However, it requires a specific feature extractor for each modality, which may not make good use of feature correlations between multimodalities. Moreover, this approach increases training parameters and complexity as more modalities are added. Collaborative-Diffusion [8] and Unite&Conquer [9] demonstrate impressive multimodal generation and editing performances using pre-trained diffusion models. However, these methods necessitate multiple denoising steps for each unimodal model during inference, resulting in slow inference speeds and significant GPU memory usage, especially when more modalities are involved in the editing process. In this paper, our FACEMUG achieves seamless integration of multimodalities for globally consistent local facial editing and has the highest number of modalities.

II. METHOD

A. Overview

The overall editing pipeline of our FACEMUG framework is shown in Fig. 2. Given a collection of pixel-wise multimodal inputs $\mathcal{X} = \{\mathbf{I}_m, \mathbf{I}_1, \mathbf{I}_2, \dots, \mathbf{I}_n\}$, an attribute label $w^d \in \mathcal{W}+$ (or text t_{tar}), an exemplar image $\mathbf{I}_{ex} \in \mathbb{R}^{h \times w \times 3}$, a groundtruth face image $\mathbf{I}_{gt} \in \mathbb{R}^{h \times w \times 3}$ (with $h \times w$ pixels and three color channels), and a binary mask $\mathbf{M} \in \mathbb{R}^{h \times w \times 1}$ (with 1 for editing and 0 for unedited pixels), the masked image $\mathbf{I}_m \in \mathbb{R}^{h \times w \times 3}$ is obtained by $\mathbf{I}_m = \mathbf{I}_{gt} \odot (\mathbf{1} - \mathbf{M})$, where each input pixel-wise modality $\mathbf{I}_k \in \mathbb{R}^{h \times w \times c_k}$ contains c_k channels and \odot denotes the Hadamard product. Let $\mathcal{W}+$ denote the disentangled style latent space [11] and \mathcal{Z} denote the random latent space. The goal of FACEMUG is to generate an edited realistic face image \mathbf{I}_{out} , where the manipulation of the masked regions in \mathbf{M} is guided by multimodal inputs, while the unedited regions remain unchanged.

1) Exemplar style module: As shown in Fig. 2 (b), our exemplar style module is designed to support randomized and exemplar-guided facial attribute editing. First, a multi-layer fully-connected neural mapping network $F_{\hat{\theta}_f}$ with the network parameters $\hat{\theta}_f$ linearly maps a random latent code $z \in \mathbb{R}^{512 \times 1}$ ($z \in \mathbb{Z}$) to style latent codes $w^z = F_{\hat{\theta}_f}(z) = \{w_i^z \in \mathbb{R}^{512 \times 1} | i \in T\} \in \mathcal{W}+$, where $T = \{1, 2, ..., t\}$ and t is the number of the style latent code. Simultaneously, a style encoder $E_{\hat{\theta}_e}$ with the network parameters $\hat{\theta}_e$ maps multimodalities to $\mathcal{W}+$. Given an exemplar image \mathbf{I}_{ex} , the style encoder extracts exemplar latent codes $w^e = \{w_i^e \in \mathbb{R}^{512 \times 1} | i \in T\} = E_{\hat{\theta}_e}(\mathbf{I}_{ex}) \in \mathcal{W}+$. Then we perform style interpolation between w^z and w^e to get the interpolated latent codes $\overline{w} \in \mathcal{W}+$.

2) Latent warping module: As shown in Fig. 2 (c), to alleviate visual artifacts due to the misalignment between the facial poses of the exemplar image and the edited image, we introduce a latent warping module H_{θ_h} that learns to transfer the pose of a target image to a source image in the style latent space with the learnable network parameters θ_h .

We first utilize the style encoder to project the pixel-wise multimodal inputs \mathcal{X} to the projected latent codes $w^p = E_{\hat{\theta}_e}(\mathcal{X}) \in \mathcal{W}+$. Then we obtain the warped latent codes as $\hat{w} = H_{\theta_h}(w^p - \overline{w}, \overline{w}) + \overline{w}, \hat{w} \in \mathcal{W}+$, under the guidance of w^p . As a consequence, \hat{w} aligns to the pose of projected latent codes w^p while preserving facial features of \overline{w} .

3) Latent attribute editing module: As shown in Fig. 2 (d), the latent attribute editing module is designed to edit the warped latent codes \hat{w} to support attribute-conditional facial attribute editing and text-driven facial attribute editing in the style latent space. In this module, we obtain the edited latent codes w^* .

4) Multimodal aggregation module: As shown in Fig. 2 (e), for better controllability and visual quality in the image space, a multimodal aggregation module A_{θ_a} with the learnable network parameters θ_a is proposed to deal with multiple heterogeneous and sparse conditional inputs by merging them into a homogeneous feature space containing c_a feature channels. Given the pixel-wise multimodal inputs \mathcal{X} , we obtain the aggregated feature tensor $\hat{\mathbf{F}}^a = A_{\theta_a}(\mathcal{X}) \in \mathbb{R}^{h \times w \times c_a}$. $\hat{\mathbf{F}}^a$ is then input into our multimodal generator.

5) Multimodal generator: As shown in Fig. 2 (f), our multimodal generator consists of a facial feature bank and a refinement auto-encoder G_{θ_g} with the trainable network parameters θ_g . We implement the facial feature bank using the StyleGAN generator $S_{\hat{\theta}_s}$. $S_{\hat{\theta}_s}$ produces multi-scale coarse facial feature maps from edited latent codes w^* while G_{θ_g} refines the editing results by utilizing the aggregated feature tensor, the latent codes, and the generated coarse features. A set of facial feature maps $\mathcal{F}^s = \{\mathbf{F}_i^s \in \mathbb{R}^{\hat{h}_i \times \hat{w}_i \times \hat{c}_i | i \in T\}$ and a reconstructed image \mathbf{I}_p are obtained as $(\mathcal{F}^s, \mathbf{I}_p) = S_{\hat{\theta}_s}(w^*)$. We define $\hat{h}_i \times \hat{w}_i \times \hat{c}_i$ as the size of feature maps at *i*-th layer. Then, the refinement auto-encoder G_{θ_g} leverages $\hat{\mathbf{F}}^a$, w^* , and \mathcal{F}^s to generate an edited image $\mathbf{I}_{out} \in \mathbb{R}^{h \times w \times 3}$:

$$\mathbf{I}_{out}(w^*) = \mathbf{I}_m \odot (\mathbf{1} - \mathbf{M}) + G_{\theta_g}(\hat{\mathbf{F}}^a, w^*, \mathcal{F}^s) \odot \mathbf{M}.$$
 (1)

The refinement auto-encoder can be further divided into an encoder G^{en} and a decoder G^{de} , i. e., $G_{\theta_a} = \{G^{en}, G^{de}\}$.

6) Discriminator: A discriminative network D_{θ_d} with the learnable network parameters θ_d learns to judge whether an image is a real or fake image. The discriminator maps an image (e.g., \mathbf{I}_{out} or \mathbf{I}_{gt}) to a scalar $D_{\theta_d}(\mathbf{I}) \in \mathbb{R}^{1 \times 1}$. Note that the discriminator is only applied during the training phase.

B. Latent warping module

To achieve realistic exemplar-guided or randomized facial attribute editing, the primary challenge lies in the potential differences in pose between the exemplar image and the edited image, which can readily result in noticeable misalignment between the two facial images. Several existing methods [71], [72] have been developed for facial pose alignment in the context of face reenactment. However, these methods often require additional pre-trained facial pose detection models or involve multi-stage training processes, which can be resourceintensive. We thus ask: Is it possible to directly conduct facial pose transfer in the latent space, thereby eliminating auxiliary steps such as preliminary GAN inversion, subsequent pose detection, and ultimately pose transfer?

To solve this, we propose a self-supervised latent warping method that eliminates the need for manual annotations and pre-trained facial pose detection models. Our method provides an intuitive and straightforward way of warping the pose of the exemplar image to match that of the edited image in the latent space while preserving the attributes of individual faces (e.g., identity or expressions). The key idea is that we use the pose of a target image to guide the pose of a source image by warping the latent codes in the style latent space.

To effectively predict the offsets in the latent space from given two latent codes for facial warping, we design our latent warping network H_{θ_h} by four stacked code-to-code modulation blocks. We construct our code-to-code modulation block by extending FFCLIP's semantic modulation block [58] to code-to-code embeddings. Moreover, the sigmoid activation is incorporated in the semantic injection for gate activation.

Let $w^{ta} \in \mathbb{R}^{t \times 512 \times 1}$ and $w^{so} \in \mathbb{R}^{t \times 512 \times 1}$ be the target and source latent codes, respectively. The warped latent codes $w^{wa} \in \mathbb{R}^{t \times 512 \times 1}$ is obtained by warping w^{so} guided by w^{ta} :

$$w^{wa} = H_{\theta_h}(w^r, w^{so}) + w^{so},$$
 (2)

where $w^r = w^{ta} - w^{so}$ is the residual latent codes between w^{ta} and w^{so} . By leveraging the code-to-code modulation mechanism, our latent warping module effectively aligns the pose of the warped latent codes with the target latent codes.

C. Latent attribute editing module

Our method supports two types of latent attribute editing, including attribute-conditional facial attribute editing and textdriven facial attribute editing. It allows us to utilize conditional labels or text to manipulate latent codes for semantic-level editing with unedited portions unchanged, which is usually hard to achieve with GAN-inversion-based methods [65].

Attribute-conditional editing. Each attribute label corresponds to a semantic direction. For a user-specified target attribute label, we obtain the edited latent codes $w^* \in W+$ by moving the warped latent codes \hat{w} (w^{wa}) along the corresponding semantic direction w^d . The editing process can be expressed below [73]:

$$w^* = \hat{w} + \epsilon \cdot w^d, \tag{3}$$

where ϵ is a user-specified weight of the latent semantic direction (attribute label) $w^d \in W^+$ to control the degree of attribute adjustment.

Text-driven editing. For a user-specified target attribute text t_{tar} , we leverage CLIP [74] to find text-driven latent codes by solving the following latent codes optimization problem:

$$w^* = \operatorname*{arg\,min}_{w \in \mathcal{W}+} \left(\lambda_{clip} \cdot \mathcal{L}_{clip}(t_{tar}, t_{src}, w, \hat{w}) + \lambda_{reg} \cdot \|w - \hat{w}\|_2 \right)$$
(4)

where $t_{src} = "face"$; $\lambda_{clip} \in [0.1, 1.0]$ and λ_{reg} are used to balance the directional CLIP loss term and the regularization term. By default, we set $\lambda_{clip} = 0.05$ and $\lambda_{reg} = 0.08$. The directional CLIP loss \mathcal{L}_{clip} [75] is employed to align directions between the text-image pairs of the original and edited images in the CLIP space:

$$\mathcal{L}_{clip}(t_{tar}, t_{src}, w, \hat{w}) = 1 - \cos(\Delta T, \Delta I),$$

$$\Delta T = E_T (t_{tar}) - E_T (t_{src}),$$

$$\Delta I = E_I (\mathbf{I}_{out}(w)) - E_I (\mathbf{I}_{out}(\hat{w})),$$
(5)

where E_T and E_I are the text and image encoders of the CLIP model, $\mathbf{I}_{out}(w)$ and $\mathbf{I}_{out}(\hat{w})$ are obtained using Eq. 1.

D. Multimodal aggregation module

To integrate multimodalities within a unified framework, the varying density and value range among images, sketches, semantic maps, and color maps present challenges [76]. For pixel-wise multi-conditional image editing, the discrepancies in information content across modalities make it difficult to apply standard convolution layers to capture the diverse characteristics of each modality, leading to suboptimal generation quality. Moreover, the differing levels of details and densities in the inputs can impact the visual appearance and realism of the generated images. To mitigate this issue, we introduce a multimodal aggregation module that efficiently aggregates the multimodal inputs. We achieve this by using separate convolution layers for each modality and incorporating a normalized adaptive weighting mechanism to merge extracted features into a homogeneous feature space. As a result, the module provides more robust representations, making it wellsuited for handling multi-conditional image editing tasks.

For the pixel-wise multimodal inputs \mathcal{X} , we first employ a residual block to extract feature maps for each modality, resulting in a feature set. Using this set, a shared residual block is utilized to compute the contribution scores for each spatial point across all pixel-wise modalities, producing a contribution score map for each modality. Each score map adaptively weights the importance of each modality in a pixelwise fashion for the aggregation process. Thus, we can get the aggregated feature $\hat{\mathbf{F}}^a = A_{\theta_a}(\mathcal{X}) \in \mathbb{R}^{h \times w \times c_a}$. This adaptive weighting mechanism allows the model to assign higher importance to informative and detailed pixel-wise modalities while reducing the impact of less informative inputs.

E. Multimodal generator

To support multimodal conditional editing and generate high-quality editing results, we develop a multimodal generator that fully utilizes aggregated facial features and edited latent codes, while efficiently fusing feature maps from the refinement encoder, the facial feature bank, and the refinement decoder. To achieve this, we introduce a style fusion block to fuse features in both high-level and shallow-level feature spaces. These carefully designed modules enhance our approach's editing capabilities, enabling the generation of diverse 'and high-fidelity editing results.

Given the aggregated feature tensor $\hat{\mathbf{F}}^a$ and the edited latent codes w^* , the refinement encoder G_{en} extracts multiscale feature maps $\mathcal{F}^{en} = {\mathbf{F}_i^{en} \in \mathbb{R}^{\hat{h}_i \times \hat{w}_i \times \hat{c}_i} | i \in T}$ and outputs a global latent vector $c \in \mathbb{R}^{512 \times 2}$ from the aggregated multimodal feature $\hat{\mathbf{F}}^a$, i.e., $(\mathcal{F}^{en}, c) = G_{en}(\hat{\mathbf{F}}^a)$. At the same



Fig. 3. Illustration of our style fusion block. Conditioned by the modulated latent vector w_{i+1}^* , the block effectively integrates multi-scale facial features in both high-level and shallow-level feature spaces.

time, the facial priors $\mathcal{F}^s = \{\mathbf{F}_i^s \in \mathbb{R}^{\hat{h}_i \times \hat{w}_i \times \hat{c}_i} | i \in T\}$ are extracted from the edited latent codes w^* in the facial feature bank with the StyleGAN generator. Finally, the feature maps $\mathcal{F}^{de} = \{\mathbf{F}_i^{de} \in \mathbb{R}^{\hat{h}_i \times \hat{w}_i \times \hat{c}_i} | i \in T\}$ in the refinement decoder G_{de} are calculated as follows:

$$\mathbf{F}_{i}^{de} = \begin{cases} \operatorname{UP}(\operatorname{SC}(\mathbf{F}_{i-1}^{de}, [w_{i}^{*}, c])), & \text{if } i \mod 2 = 1; \\ \operatorname{SC}(\mathbf{F}_{i-1}^{g}, [w_{i}^{*}, c]), & \text{otherwise}, \end{cases}$$
(6)

where $[\cdot, \cdot]$ denotes concatenation; UP(\cdot) refers to upsampling; SC(\cdot, \cdot) indicates the style layer [11], $\mathbf{F}_0^{de} =$ Conv(\mathbf{F}_t^{en}), Conv(\cdot) is a Convolution layer; \mathbf{F}_j^g (j =1,3,5,...,2 $\lfloor \frac{t}{2} \rfloor + 1$) is the fused feature map by the proposed style fusion block. We set $\mathbf{I}_{out} = \text{Conv}(\mathbf{F}_t^{de}) \in \mathbb{R}^{h \times w \times 3}$ as the model output.

Style fusion block. The proposed style fusion block is as shown in Fig. 3. To fully leverage the guidance information extracted from the refinement auto-encoder and facial feature priors of the facial feature bank, we employ the gated fusion scheme to perform element-wise fusion between these features for enhancement. We first apply an adaptive gated fusion to activate features from the refinement encoder to obtain the intermediate generated feature $\hat{\mathbf{F}}_{i}^{g}$:

$$\hat{\mathbf{F}}_{i}^{g} = (\sigma(\mathrm{SC}(\mathbf{F}_{i}^{de}, w_{i+1}^{*})) + \mathbf{1}) \odot \mathbf{F}_{t-i}^{en} + \phi(\mathrm{SC}(\mathbf{F}_{i}^{de}, w_{i+1}^{*})),$$
(7)

where $\sigma(\cdot)$ denotes the sigmoid activation function and $\phi(\cdot)$ corresponds to the LeakyReLU activation function with the negative slope of 0.2. Then, the style layer computes the spatially-variant gate map \mathbf{F}_i^m from the facial priors \mathbf{F}_i^s and the modulated latent vector w_i^* for each *i*-th layer. The feature fusion is calculated as follows:

$$\mathbf{F}_{i}^{m} = \sigma(\mathrm{SC}(\mathbf{F}_{i}^{s}, w_{i+1}^{*})),
\mathbf{F}_{i}^{g} = \mathbf{F}_{i}^{m} \odot \mathbf{F}_{i}^{s} + (\mathbf{1} - \mathbf{F}_{i}^{m}) \odot \hat{\mathbf{F}}_{i}^{g},$$
(8)

where the spatially-variant gating map \mathbf{F}_{i}^{m} automatically selects the important features from generated feature maps $\hat{\mathbf{F}}_{i}^{g}$ and the facial priors \mathbf{F}_{i}^{s} .

F. Self-supervised module training

In our framework, we utilize the pre-trained StyleGAN generator $S_{\hat{\theta}_s}$ and mapping network $F_{\hat{\theta}_f}$ from StyleGAN-V2 [11]. We begin by the training of the style encoder $E_{\hat{\theta}_e}$. Next, we optimize the latent warping network H_{θ_h} with our proposed self-supervised warping learning. Finally, we detail the training process for the multimodal aggregation module



Fig. 4. Illustration of the self-supervised training of our latent warping module. We employ the style encoder to project the augmented image to obtain the target latent codes w^{ta} . The source latent codes w^{so} are sampled using interpolation between the initial latent codes w^{ini} and the flipped latent codes w^f . The identity loss, the LPIPS loss, and the attribute loss are utilized as constraints to disentangle the identity and pose. This module effectively transfers the pose of w^{ta} to the warped latent codes w^{wa} while remaining other facial features unchanged. The inversion process is utilized for the visualization purpose.

 A_{θ_a} , the refinement auto-encoder G_{θ_g} , and the discriminator D_{θ_d} . This training process can be in parallel with the optimization of the latent warping network.

1) Training of style encoder: The style encoder for exemplar images and that for multimodalities (sketch, color, semantic map, and mask) are the same encoder. The network of the style encoder $E_{\hat{\theta}_e}$ is borrowed from e4e [64]. In order to enable the encoder to project each modality individually into the latent space, we customized the first convolution layer of the e4e encoder to handle 26 channels from four modalities respectively: 3 channels for the exemplar image, 1 channel for sketches, 3 channels for colors, and 19 channels for semantic layouts. Then, we feed the concatenated randomly masked multimodal inputs into $E_{\hat{\theta}_e}$ for training. In addition, the loss functions defined in e4e [33], [64] are employed.

2) Training of latent warping module: As shown in Fig. 4, a triplet of initial, source, and target latent codes is utilized to learn the pose warping guided by the target latent codes. We begin by projecting an initial image \mathbf{I}_{ini} into the style latent space to get the initial latent codes $w^{ini} = E_{\hat{\theta}_e}(\mathbf{I}_{ini})$. Then, we obtain an augmented image \mathbf{I}_{ta} by applying bilinear scaling, color jittering, and region masking operations [77] to \mathbf{I}_{ini} . In addition, we obtain a flipped image \mathbf{I}_f with mirror flipping of \mathbf{I}_{ini} . Therefore, \mathbf{I}_f shares the same identity as \mathbf{I}_{ini} with a flipped pose. Then, the target latent codes w^{ta} and the flipped latent codes w^f can be obtained by $w^{ta} = E_{\hat{\theta}_e}(\mathbf{I}_{ta})$ and $w^f = E_{\hat{\theta}_e}(\mathbf{I}_f)$, respectively. Next, we obtain the source latent codes w^{so} with a linear interpolation:

$$w^{so} = \beta \cdot w^{ini} + (1 - \beta) \cdot w^f, \tag{9}$$

where $\beta \in [0,1]$ is a uniform random number. Since w^{ini} and w^f share the same facial features except for the pose, the interpolated source codes w^{so} maintain the same facial identity to w^{ini} but have a different pose. The warped latent

codes w^{wa} are obtained by transferring the pose of w^{ta} to w^{so} while maintaining the identity of w^{so} , using our latent warping network H_{θ_h} . Finally, we obtain the warped image as $(\mathcal{F}^{wa}, \mathbf{I}_{wa}) = S_{\hat{\theta}_s}(w^{wa})$.

Total loss. In order to disentangle the identity and pose during warping, we train H_{θ_h} by utilizing the identity loss, the LPIPS loss, and the attribute loss to constrain the identity and attribute similarities between w^{wa} and w^{ini} . The total training loss of H_{θ_h} is defined as:

$$O(\theta_h) = \lambda_{latent} \cdot \left(\mathcal{L}_{id}(\mathbf{I}_{wa}, \mathbf{I}_{ini}) + \mathcal{L}_{lpips}(\mathbf{I}_{wa}, \mathbf{I}_{ini}) + \mathcal{L}_{attr}(w^{wa}, w^{ini}) \right),$$
(10)

where λ_{latent} is empirically set to 0.1 in this work; \mathcal{L}_{id} , \mathcal{L}_{lpips} , and \mathcal{L}_{attr} are the identity loss, the LPIPS loss, and the attribute loss, respectively, and are defined below.

Identity loss. The identity loss [24], [58] is incorporated to constrain the identity similarity:

$$\mathcal{L}_{id}(\mathbf{I}_x, \mathbf{I}_y) = 1 - \cos(R(\mathbf{I}_x), R(\mathbf{I}_y)), \quad (11)$$

where $R(\cdot)$ is a pre-trained ArcFace network [78].

LPIPS loss. The Learned Perceptual Image Patch Similarity (LPIPS) [79] is applied to constrain the perceptual similarity:

$$\mathcal{L}_{lpips}(\mathbf{I}_x, \mathbf{I}_y) = \|P(\mathbf{I}_x) - P(\mathbf{I}_y)\|_2, \quad (12)$$

where $P(\cdot)$ is a pre-trained VGG feature extractor [80].

Attribute loss. The attribute loss [24], [81] is used to constrain the learning in the style latent space:

$$\mathcal{L}_{attr}(w^x, w^y) = \|w^x - w^y\|_2.$$
(13)

Consequently, we can obtain the optimized parameters θ_h^* via the minimization of $O(\theta_h)$. By taking advantage of the code-to-code modulation mechanism of H_{θ_h} with the loss constraints, w^{wa} effectively learns the pose from w^{ta} while remaining other facial features (e.g., identity) of w^{so} unchanged.

3) Training of multimodal aggregation module, refinement auto-encoder and discriminator: The identity loss, the LPIPS loss, the diversity-enhanced attribute loss, and the adversarial loss are combined to optimize the multimodal aggregation module, the refinement auto-encoder, and the discriminator.

To learn the mapping between style latent codes and corresponding facial attributes, and to support attribute-conditional editing in the style latent space, a diversity-enhanced attribute loss $\mathcal{L}_{attr}(w^o, \overline{w})$ is employed to constrain the consistency between facial attributes of the edited image \mathbf{I}_{out} and the interpolated latent codes \overline{w} , where $w^o = E_{\hat{\theta}_x}(\mathbf{I}_{out})$.

Total loss. The total training loss is defined as:

$$O(\theta_a, \theta_g, \theta_d) = \lambda_{id} \mathcal{L}_{id}(\mathbf{I}_{out}, \mathbf{I}_{ex}) + \lambda_{attr} \mathcal{L}_{attr}(w^o, \overline{w}) + \lambda_{lpips} \mathcal{L}_{lpips}(\mathbf{I}_{out}, \mathbf{I}_{gt}) + \mathcal{L}_{adv}(\mathbf{I}_{out}, \mathbf{I}_{gt}),$$
(14)

where we empirically set $\lambda_{id} = 0.1$, $\lambda_{lpips} = 0.5$, and $\lambda_{attr} = 0.1$ in this work; \mathcal{L}_{adv} is the adversarial non-saturating logistic loss [10] with R_1 regularization [82].

The refinement network G_{θ_g} is trained to generate a realistic edited image \mathbf{I}_{out} while the discriminator D_{θ_g} tries to differentiate between \mathbf{I}_{gt} and \mathbf{I}_{out} . In an alternating fashion, A_{θ_a} and G_{θ_g} are trained in a phase while D_{θ_d} is trained in the other. For each iteration, we obtain the optimized parameters θ_a^* , θ_a^* and θ_d^* via the minimax game iteratively. 4) Implementation without manual annotation: Our framework was implemented using Python and PyTorch. We trained FACEMUG and our latent warping module independently. For more implementation details of the latent warping module, the multimodal aggregation module, and the training procedures of our networks, please refer to the supplementary document.

Manual annotations of labeled paired data across different modalities are required in existing multimodal editing methods [6], [7]. On the contrary, the training modalities (i.e., editing mask, exemplar, semantics, sketch, and color) in this paper were generated without any manual annotation using wellbuilt methods. The trained masks were generated randomly with the mask generation algorithm from CMOD [27]. The Face-parsing model [83] was used to extract semantic maps. Hand-drawn-like sketches were generated using a pencilsketch filter [33]. Color images were processed using a mean color of each semantic region. Exemplar images were sampled randomly from the ground-truth images. By aligning all modalities into a unified generative latent space, FACEMUG effectively loosens the ties between the paired modalities and enables model training without any human annotation.

The latent attribute editing module was implemented as follows. For the attribute-conditional editing, various attribute labels employed in InterfaceGAN [52], GANSpace [56], Style-CLIP [57], and CLIP2StyleGAN [54] were integrated in the module for semantic direction. For the text-driven editing, the edited latent codes were obtained through $100 \sim 300$ iterations of gradient descent [57] with the learning rate of 0.1.

Following the settings of StyleGANv2 [11], we employed the Adam optimizer with the first momentum coefficient of 0.5, the second momentum coefficient of 0.99, and the learning rate of 0.002. We trained the networks for 800,000 iterations with a batch size of 8.

III. EXPERIMENTAL RESULTS AND COMPARISONS

A. Settings

We conducted experimental evaluations on two publicly available and commonly used benchmark face image datasets: CelebA-HQ [84] and FFHQ [14]. Our FACEMUG was trained on the training set of FFHQ and evaluated on the testing set of CelebA-HQ and FFHQ, respectively. In the CelebA-HQ dataset, 2,000 images were randomly selected for testing. For the FFHQ dataset, 60,000 images were randomly chosen for training, and the remaining 10,000 images were used for testing. All images were resized to the resolution of 256×256 . To ensure a fair comparison, the same training and testing splits were used for all experiments.

All experiments were conducted on the NVIDIA Tesla V100 GPU. The training time of our FACEMUG was around one month. We also evaluated FACEMUG on a PC equipped with an NVIDIA GeForce RTX 4090 GPU. It takes 29 ms (34 FPS) for each inference.

We quantitatively evaluated the performance by using the Fréchet inception distance (FID) [85], the unpaired inception discriminative score (U-IDS) [27], and LPIPS [79] metrics which are robust assessment measures and correlate well with human perception for the image quality.



Fig. 5. Visual comparison to ColDiffusion [8] and Unite&Conquer [9] for text-driven multimodal facial editing. Our method produces visually appealing and globally consistent images with good responses to the corresponding multimodal inputs, and remains unmasked parts unchanged.



Fig. 6. Visual comparison to PoE-GAN [7]. We used the input modalities and results published in their paper. The consistency between generated images and multimodal inputs of FACEMUG is better than that of PoE-GAN.

We demonstrated the performance of FACEMUG by utilizing various multimodal inputs. For convenience, we used the following abbreviations for these multimodalities: "+Sk" for adding sketches, "+Se" for adding semantic maps, "+Co" for adding color information, "+Ex" for adding an exemplar image, "+Te" for adding text. We used the notation "+Ma" to represent the inclusion of a masked image for local editing. Different masks may affect the quantitative results because of the variation in position, size, and shape. We included various types of masks for comprehensive quantitative evaluation. Each mask was selected randomly from one of the following types of masks: hair, face, foreground subject, irregular region (50-60% mask ratio), and a fixed center (128×128) rectangle. The center mask was included because it effectively covers most of the facial region in a facial image. For quantitative comparisons, masks were obtained automatically. Semanticsbased masks (hair, face, etc.) were created with the Faceparsing model [83]. Irregular masks were obtained from the irregular mask templates [86]. For qualitative comparisons, masks were manually defined. A consistent set of inputs was used in each comparison to ensure fairness, both quantitatively and qualitatively. For more experimental results, please refer

to the supplementary document.

B. Comparison on multimodal facial editing

To evaluate the quality of generated images and the responsiveness of multimodal inputs, we performed comparisons against SOTA multimodal facial image editing techniques, including Unite&Conquer [9], Collaborative-Diffusion (ColDiffusion) [8], and PoE-GAN [7] using multimodal conditional inputs. Unless specified, we employed officially released pretrained models of compared methods.

Fig. 5 shows visual comparisons of our FACEMUG to Unite&Conquer and ColDiffusion. It shows that all the compared methods can use semantics or sketches to control the facial layout while adjusting appearance using given text. However, Unite&Conquer and ColDiffusion show low consistency between the output image and the text caption. In contrast, our method is capable of combining the three modalities to perform high-quality multimodal editing. Fig. 6 shows visual comparisons of our FACEMUG to PoE-GAN using text with sketches, semantics, and exemplar, respectively. The results of PoE-GAN show high-quality and good responses to the corresponding input modalities. However, for



Fig. 7. Visual comparison to Unite&Conquer [9]. We show more FACEMUG results by adding extra masks. Our method shows better visual quality and preserves background information when using masks.



Fig. 8. Qualitative comparison of incremental editing (the later editing taking the previous output image as input), compared to SOTA methods (SemanticStyleGAN [39], ColDiffusion [8], DeepFaceEditing [34], SC-FEGAN [22], HFGI [65]) in row 2, using uni-modality. Our model shows better visual quality.

 TABLE I

 QUANTITATIVE COMPARISON OF OUR METHOD TO SOTA MULTIMODAL

 FACIAL EDITING METHODS WITH SKETCHES, SEMANTIC MAPS, AND TEXT

 ON THE CELEBA-HQ DATASET. BOLD: TOP-1 QUANTITY.

Method	FID↓	U-IDS [↑]	LPIPS↓
ColDiffusion [8] (+Se+Te) Unite&Conquer [9] (+Se+Te) Ours (+Se+Te) Ours (+Se+Te+Ma)	26.87 44.52 38.41 11.85	0 0 0.20%	0.5283 0.5809 0.4699 0.1645
Unite&Conquer [9] (+Sk+Se+Te) Ours (+Sk+Se+Te) Ours (+Sk+Se+Te+Ma)	45.29 32.65 11.24	0 0 0.28%	0.5493 0.4046 0.1448

the first case, our method shows a clear smile expression, and for the last case, the edited result exhibits more correlation with the exemplar. For FACEMUG, the semantic maps and sketches provide geometry information, while text controls the appearance of the generated content. Moreover, by leveraging the additional masked image, the generated content exhibits high consistency to the input modalities and is coherent to the editing boundary, preserving the unedited part unchanged. We also show the quantitative comparison of multimodal editing with ColDiffusion [8] and Unite&Conquer [9] on CelebA-HQ dataset, as shown in Table I. The semantic maps and text are from CelebAMask-HQ [36] and CelebA-Dialog [87] datasets, respectively. Guided by semantic maps and text, both ColDiffusion and Unite&Conquer produce high-quality images yet exhibit lower consistency with ground-truth images, as re-

TABLE II QUANTITATIVE COMPARISON TO SOTA MULTIMODAL FACIAL IMAGE EDITING METHODS WITH SKETCHES AND SEMANTIC MAPS ON CELEBA-HQ AND FFHQ DATASETS. **BOLD**: TOP-1 QUANTITY.

9

Method	CelebA-HQ			FFHQ		
Wiethou	FID↓	U-IDS↑	LPIPS↓	FID↓	U-IDS↑	LPIPS↓
Unite&Conquer [9] (+Sk+Se)	44.76	0	0.5350	52.39	0	0.5707
Ours (+Sk+Se) Ours (+Sk+Se+Ma)	29.96 10.36	0 0.63 %	0.3999 0.1371	23.16 2.26	0.03% 30.38%	0.4036 0.1162

flected by their LPIPS scores. In contrast, FACEMUG achieves lower LPIPS and competitive FID scores, indicating good fidelity of our results. When applying a mask to indicate the editing regions, FACEMUG not only demonstrates improved performance of FID and LPIPS scores but also illustrates superior editing quality. Our method achieves the lowest FID scores because it can modify specific facial attributes within the masked area while maintaining the unmasked regions unchanged and enhancing the overall image consistency.

We further conducted comparisons between FACEMUG and Unite&Conquer [9], focusing on facial editing using sketches and semantics. For a fair comparison, given that Unite&Conquer utilizes only the "skin" and "hair" semantic maps, we accordingly adjusted FACEMUG by removing other semantic labels. As shown in Fig. 7, Unite&Conquer unites multiple diffusion models trained on multiple sub-tasks to perform editing with the guidance of sketch and partial semantic labels. However, the presence of visual artifacts in



Fig. 9. Visual comparison between our FACEMUG and the SOTA sketch-guided editing methods (pSp [33], DeepFaceEditing [34], and ControlNet [35]). The sub-images in each group represent the guidance information for the editing process. FACEMUG produces images with superior quality and finer details by using more modalities, and shows global consistency when adding masks for local facial editing.

TABLE III The user study results on the CelebA-HQ dataset. We present the percentages (%) of cases where our results were preferred over those of the compared methods.

Configuration	Method-to-method comparison	Percentage
+Se+Te	Ours vs. ColDiffusion [8]	54.90%
+Se+Te	Ours vs. Unite&Conquer [9]	76.83%
+Se+Sk+Te	Ours vs. Unite&Conquer [9]	77.88%

the edited images indicates that employing disparate off-theshelf diffusion models, trained on different datasets, remains a challenging task. In comparison, FACEMUG adeptly leverages sketches and partial semantic labels to facilitate geometryguided facial generation (+Sk+Se) and editing (+Sk+Se+Ma). Moreover, the generated contents maintains consistency with the unedited portions. Table II displays the FID, U-IDS, and LPIPS scores of each method on CelebA-HQ and FFHQ datasets. Unite&Conquer introduces a novel reliability parameter to facilitate the multimodal mixing of contents generated from various uni-modal diffusion networks. Nevertheless, our method surpasses Unite&Conquer, exhibiting superior FID, U-IDS, and LPIPS scores.

Fig. 1 and Fig. 8 show examples of incremental multimodal local facial editing. Our FACEMUG incrementally edits the input facial images to achieve high-quality manipulation by taking advantage of multimodal inputs, including masks, exemplars, semantics, sketches, colors, and attribute labels. It is worth noting that incremental editing is achieved effectively with our unified FACEMUG model, while existing methods have to use multiple uni-modal models (see Fig. 8) and introduce unwanted modifications on unedited facial features (see Fig. 1). By leveraging provided multimodalities, our FACEMUG can edit various realistic facial attributes (e.g., facial geometry, hairstyle, and decorative goods) while preserving unedited parts unchanged.

User study: We conducted a user study for multimodal facial editing comparing ColDiffusion and Unite&Conquer with "Se+Te" and "+Sk+Se+Te" configurations. We sampled 100 images randomly from the CelebA-HQ test set and obtained corresponding edited images. 20 pairs of edited images were chosen randomly for each method-to-method comparison. For FACEMUG, ten edited images were produced with additional masks (+Ma), which were selected randomly

TABLE IV QUANTITATIVE COMPARISON TO SOTA SKETCH-GUIDED FACIAL IMAGE EDITING METHODS (+SK) ON CELEBA-HQ AND FFHQ DATASETS. **BOLD**: TOP-1 QUANTITY.

Method	CelebA-HQ			FFHQ		
	FID↓	U-IDS ⁺	LPIPS↓	FID↓	U-IDS ⁺	LPIPS↓
TediGAN [6]	42.83	0	0.4909	85.17	0	0.6031
pSp [33]	42.70	0	0.4911	85.16	0	0.6031
DeepFaceEditing [34]	19.78	0	0.2796	11.59	5.13%	0.2850
ControlNet [35]	64.59	0	0.5530	62.31	0	0.5475
Ours (+Sk)	36.58	0	0.4071	24.29	0	0.4096
Ours (+Sk+Co)	17.43	0	0.2936	9.12	9.59%	0.2815
Ours (+Sk+Ma)	11.43	0.52%	0.1444	2.49	29.51%	0.1259

from hair, face, and foreground subject masks; the other ten images were produced without masks. The participants were asked to perform two-alternative forced choices (2AFCs) based on visual realism and consistency with input modalities. Finally, we recruited 52 participants, resulting in 1040 votes per comparison.

Table III shows the user study results. With the "+Se+Te" configuration, our FACEMUG received 54.90% and 76.83% of the preference votes compared to ColDiffusion (45.10%) and Unite&Conquer (23.17%), respectively. Our FACEMUG also surpassed Unite&Conquer with 77.88% of majority votes with the "+Sk+Se+Te" configuration. The user study validated that our FACEMUG effectively produces realistic facial images by taking advantage of multimodal local facial editing.

C. Comparison on sketch-guided facial editing

We compared FACEMUG to the SOTA sketch-guided facial editing methods, including pSp [33], DeepFaceEditing [34] and ControlNet [35]. The pre-trained models of pSp and DeepFaceEditing provided in the official online repository were used in this experiment. For ControlNet, we fine-tuned the officially provided pre-trained weights on the FFHQ training set with corresponding sketches. To ensure optimal performance, we used default forms of sketches that were used during training for each method.

As depicted in Fig. 9, all the compared methods are capable of generating high-quality images guided by sketches. However, pSp may neglect some facial attributes, like glasses from input sketches. Even conditioned on facial appearance and sketches, DeepFaceEditing exhibits some visual artifacts in edited outputs. While ControlNet can generate high-quality



Fig. 10. Visual comparison between our FACEMUG and the SOTA semantic-guided editing methods (pSp [33], SEAN [38], SofGAN [37], SDM [40], and ColDiffusion [8]). The sub-images in each group represent the guidance information for the editing process. FACEMUG produces more visually appealing results using more conditional modalities, and achieves high-quality local facial editing by incorporating masks.

TABLE V QUANTITATIVE COMPARISON TO SOTA SEMANTIC-GUIDED FACIAL IMAGE EDITING METHODS (+SE) ON CELEBA-HQ AND FFHQ DATASETS. BOLD: TOP-1 QUANTITY.

Method	CelebA-HQ			FFHQ		
wiethou	FID↓	U-IDS↑	LPIPS↓	FID↓	U-IDS↑	LPIPS↓
TediGAN [6]	45.19	0	0.5208	85.50	0	0.6109
pSp [33]	45.19	0	0.5208	85.50	0	0.6109
SEAN [38]	28.74	0	0.2595	26.55	1.5%	0.3801
SofGAN [37]	38.88	0	0.6369	15.43	8.14%	0.6297
SDM [40]	27.61	0	0.4995	-	-	-
ColDiffusion [8]	30.22	0	0.5280	76.73	0	0.5937
Ours (+Se)	41.18	0	0.4462	43.78	0	0.4536
Ours (+Se+Co)	31.02	0	0.3691	14.08	8.09%	0.3018
Ours (+Se+Ma)	10.65	0.53%	0.1570	2.41	28.99%	0.1356

results, it needs several denoising steps and depends on highquality prompts. Without colors and unmasked background information (third last column), the color richness generated by FACEMUG may be affected. By incorporating both colors and sketches as guidance (second last column), FACEMUG can produce facial images with high fidelity. Our method (last column) demonstrates sketch-guided local facial editing, showcasing our method's ability to achieve high-quality facial image editing by leveraging unmasked pixels and input sketches while preserving the unedited regions unchanged.

Table IV presents a quantitative comparison between FACE-MUG and the existing SOTA approaches. All methods demonstrate good FID scores on both datasets. DeepFaceEditing achieves competitive FID and U-IDS scores by utilizing sketch and appearance information extracted from groundtruth images. Since sketches lack color and appearance information, the performance of FACEMUG (+Sk) may be limited. FACEMUG (+Sk+Co), which incorporates color as guidance information, further enhances the visual quality. FACEMUG (+Sk+Ma) surpasses all the compared methods by leveraging background information, demonstrating superior performance.

D. Comparison on semantic-guided facial editing

We compared FACEMUG to the SOTA semantic-guided facial editing methods, including pSp [33], SEAN [38], Sof-GAN [37], SDM [40], and ColDiffusion [8]. The pre-trained models of the compared methods provided in the official online repository were used in this experiment.

As shown in Fig. 10, all methods can generate highquality images guided by semantic maps. However, pSp and



Fig. 11. Visual comparison of the semantic-guided local facial editing with zoom-in details between IDE-3D [41] and FACEMUG. Sub-images are the guidance information. The result of IDE-3D was obtained from the paper [41].

ColDiffusion may not accurately capture certain attributes (e.g., the shape of the glasses). SofGAN may produce facial images with less defined details. SDM and ColDiffusion are capable of generating superior images but at the cost of longer computation time. Additionally, above compared methods fail to preserve the unedited regions surrounding the editing area. In contrast, FACEMUG (third last column) conditioned by semantic maps exhibits clear manifestations of semantic maps. When guided by color information as well (second last column), our method demonstrates superior visual performance. Furthermore, FACEMUG (last column) allows for local facial editing guided by semantic maps, preserving known regions while providing more flexibility for interactive facial attribute manipulation. FACEMUG consistently produces high-quality results with facial attributes guided by semantic maps.

Table V presents the quantitative performance of the compared methods. SDM and ColDiffusion achieve good FID scores when conditioned by semantic maps. Without colors and unmasked pixels, the performance of FACEMUG (+Se) is limited. Incorporating colors significantly improves the performance of FACEMUG (+Se+Co). With the ability to edit masked regions while preserving unmasked pixels, FACEMUG (+Se+Ma) achieves the best FID, U-IDS, and LPIPS scores for both datasets. FACEMUG outperforms the compared models in terms of the authenticity of locally edited facial images guided by semantic maps.

Fig. 11 shows the visual comparison between FACEMUG and the SOTA 3D-aware facial editing method, IDE-3D [41]. Both methods achieve high-quality editing. IDE-3D supports 3D face synthesis but may result in some losses in facial details during GAN inversion, such as moles and eyelids. In comparison, FACEMUG excels at preserving unedited facial attributes



Fig. 12. Visual performance of our FACEMUG with various modal inputs. The masked image was utilized for all outputs. There are a total of 32 combinations (subsets) of five modalities. Our FACEMUG generates visually appealing results and shows high global consistency to the unedited regions.



Fig. 13. Visual comparison of the ablation study on the latent warping module (from left to right in each group): the masked image, the exemplar image, inversion of latent codes of the exemplar image, the editing result without the latent warping module (+Ex w/o warping), inversion of warped latent codes of the exemplar image, and the editing result with the latent warping module (+Ex). Our warping module improves the visual quality when the pose misalignment happens between the edited image and the exemplar.

but faces challenges in reconstructing 3D facial geometry. Combining the strengths of both methods could contribute to a more advanced multimodal facial editing framework.

E. Ablation study

1) Ablation study on multimodal inputs: We also conducted experiments to evaluate the effects of multimodal inputs on our FACEMUG framework. Fig. 12 showcases editing examples of FACEMUG using a total of 32 input configurations of five modalities. Generally, when multimodalities are consistent with each other, incorporating more modalities achieves much better visual quality. "None" means that only the masked image is utilized for image inpainting. When using sketches or semantic maps, FACEMUG clearly exhibits the structure of inputs. Incorporating colors enhances texture details and produces faithful edited results. Example images further help transfer the styles and facial identity to the edited regions. With the guidance of the text, more detailed facial attributes can be manipulated. Utilizing all modalities (five modalities) information allows FACEMUG to achieve the best overall performance. It also demonstrates that our method can work well on all the subsets of five modalities.

2) Ablation study on latent warping module: We conducted a study to assess the effectiveness of our latent warping module. Various irregular masks with different mask ratios [86] and a fixed center 25% (128×128) rectangular mask were used to simulate different editing situations. We show visual examples in Fig. 13. We can find that our latent warping module effectively aligns the exemplar pose to the edited image while preserving the facial attributes and identity of the exemplar images. As shown in Fig. 14, the quantitative performance results demonstrate that our latent warping module



Fig. 14. Quantitative comparison of the ablation study on the latent warping module on the FFHQ dataset. "+Ex w/o warping": the exemplar image without the latent warping module used; "+Ex": the exemplar image with the latent warping module used. Other modalities were not used.

TABLE VI Ablation study of the image generation performance of other main components on the FFHQ dataset with 50 – 60% mask ratios. (A) represents our full model (FACEMUG). (B) replaces our style fusion blocks with element-wise addition operations. (C) replaces our style fusion blocks with gated convolution blocks [88]. (D) replaces our diversity-enhanced attribute loss with the attribute loss [24]. (E) removes the facial feature bank. **Bold**: top-1 quantity.

Method	FID↓	U-IDS↑	LPIPS↓
А	3.274	28.14%	0.1917
В	3.829	24.05%	0.2010
С	3.478	26.70%	0.1954
D	3.599	26.44%	0.1991
E	3.725	25.26%	0.1971

achieves better quantitative performance (+Ex) compared to not using it (+Ex w/o warping). Since the edited and exemplar images contain different poses, directly transferring features to target regions may cause obvious artifacts. Our latent warping module mitigates this issue by adapting the exemplar pose to align with the edited image in the style latent space, thereby avoiding the boundary issue during editing.

3) Ablation study on other main components: Here, we explored the image generation performance of other main components by comparing FACEMUG to its five variants on the FFHQ dataset with 50 - 60% mask ratios, as shown in Table VI. We tested to replace our fusion blocks with element-wise addition operations (B) and gated convolution blocks [88] (C). The quantitative performance of models (B) and (C) exhibited a certain degree of decline, compared to our full model (A). When replacing our diversity-enhanced attribute loss with the attribute loss [24] (D), the quantitative measures also show a decrease. We also show FACEMUG's effectiveness by removing the facial feature bank (E). The quantitative scores dropped significantly, underscoring the importance of the facial feature bank for high-quality results.

IV. CONCLUSION

In this paper, we have explored a novel multimodal generative and fusion framework (FACEMUG) for globallyconsistent local facial editing, which generates realistic facial features in response to multimodal inputs on the edited regions while maintaining visual coherence with unedited parts. FACEMUG takes advantage of diverse input modalities (e.g., sketches, semantic maps, color maps, exemplar images, text, and attribute labels) to perform fine-grained and semantic facial editing on geometry, color, expressions, attributes, and identity within edited regions, and allows for incremental edits. Extensive experiments have demonstrated the effectiveness of the proposed method.

Limitations and future work: Our approach comes with certain limitations. First, although the inference of FACE-MUG is fast, the training of FACEMUG is time-intensive and requires approximately one month to complete on a V100 GPU. We plan to develop a more lightweight model to expedite FACEMUG's training process in the future. Second, because of the limited training data, the pre-trained StyleGAN struggles to generate relatively extreme expressions, poses, and appearances, potentially leading to FACEMUG failures in these domains. A more expressive and powerful pre-trained StyleGAN will improve our model. Third, like most multimodal editing algorithms [7], FACEMUG may not generate satisfying results when different modalities contain contradictory guiding information. Designing a more sophisticated technique to reweight the contribution of each modality would be a promising step forward. Finally, an interesting future direction is to incorporate more modalities, such as text and audio, into our multimodal aggregation module to achieve more diverse editing. A possible solution is to employ facial landmarks as an intermediate motion representation [89].

ACKNOWLEDGEMENTS

The authors would like to thank the anonymous reviewers for their constructive comments. This work was supported by the Zhejiang Provincial Natural Science Foundation of China (Grant No. LZ21F020001). Xiaogang Jin was supported by Key R&D Program of Zhejiang (Grant No. 2023C01047).

REFERENCES

- Y.-R. Jiang, S.-Y. Chen, H. Fu, and L. Gao, "Identity-aware and shapeaware propagation of face editing in videos," *TVCG*, pp. 1–12, 2023.
- [2] L. Qu, J. Shang, X. Han, and H. Fu, "Reenactartface: Artistic face image reenactment," *TVCG*, pp. 1–13, 2023.
- [3] X. Han, K. Hou, D. Du, Y. Qiu, S. Cui, K. Zhou, and Y. Yu, "Caricatureshop: Personalized and photorealistic caricature sketching," *TVCG*, vol. 26, no. 7, pp. 2349–2361, 2020.
- [4] C. Wang, M. Chai, M. He, D. Chen, and J. Liao, "Cross-domain and disentangled face manipulation with 3d guidance," *TVCG*, vol. 29, no. 4, p. 2053–2066, 2023.

- [5] W. Su, H. Ye, S.-Y. Chen, L. Gao, and H. Fu, "Drawinginstyles: Portrait image generation and editing with spatially conditioned stylegan," *TVCG*, vol. 29, no. 10, pp. 4074–4088, 2023.
- [6] W. Xia, Y. Yang, J.-H. Xue, and B. Wu, "Tedigan: Text-guided diverse face image generation and manipulation," in *CVPR*, 2021, pp. 2256– 2265.
- [7] X. Huang, A. Mallya, T.-C. Wang, and M.-Y. Liu, "Multimodal conditional image synthesis with product-of-experts gans," in *ECCV*, 2022, p. 91–109.
- [8] Z. Huang, K. C. Chan, Y. Jiang, and Z. Liu, "Collaborative diffusion for multi-modal face generation and editing," in *CVPR*, 2023, pp. 6080– 6090.
- [9] N. G. Nair, W. G. C. Bandara, and V. M. Patel, "Unite and conquer: Plug & play multi-modal synthesis using diffusion models," in *CVPR*, 2023, pp. 6070–6079.
- [10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *NeurIPS*, vol. 27, 2014, pp. 2672–2680.
- [11] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of stylegan," in *CVPR*, 2020, pp. 8107–8116.
- [12] K. C. Chan, X. Wang, X. Xu, J. Gu, and C. C. Loy, "Glean: Generative latent bank for large-factor image super-resolution," in *CVPR*, 2021, pp. 14245–14254.
- [13] G. Lample, N. Zeghidour, N. Usunier, A. Bordes, L. DENOYER, and M. A. Ranzato, "Fader networks:manipulating images by sliding attributes," in *NeurIPS*, vol. 30, 2017.
- [14] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," *TPAMI*, vol. 43, no. 12, pp. 4217– 4228, 2021.
- [15] T. Karras, M. Aittala, S. Laine, E. Härkönen, J. Hellsten, J. Lehtinen, and T. Aila, "Alias-free generative adversarial networks," in *NeurIPS*, vol. 34, 2021, pp. 852–863.
- [16] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *ICML*, vol. 37, 2015, pp. 2256–2265.
- [17] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *NeurIPS*, vol. 33, 2020, pp. 6840–6851.
- [18] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," in *NeurIPS*, vol. 34, 2021, pp. 8780–8794.
- [19] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "Highresolution image synthesis with latent diffusion models," in *CVPR*, 2022, pp. 10684–10695.
- [20] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," in *ICLR*, 2021.
- [21] A. Vahdat, K. Kreis, and J. Kautz, "Score-based generative modeling in latent space," in *NeurIPS*, vol. 34, 2021, pp. 11287–11302.
- [22] Y. Jo and J. Park, "Sc-fegan: Face editing generative adversarial network with user's sketch and color," in *ICCV*, 2019, pp. 1745–1753.
- [23] Y. Wu, Y. Yang, Q. Xiao, and X. Jin, "Coarse-to-fine: facial structure editing of portrait images via latent space classifications," *TOG*, vol. 40, no. 4, p. Article 46, 2021.
- [24] W. Lu, H. Zhao, X. Jiang, X. Jin, Y. Yang, M. Wang, J. Lyu, and K. Shi, "Do inpainting yourself: Generative facial inpainting guided by exemplars," *arXiv preprint*, vol. arXiv:2202.06358, 2022.
- [25] Z. Yi, Q. Tang, S. Azizi, D. Jang, and Z. Xu, "Contextual residual aggregation for ultra high-resolution image inpainting," in *CVPR*, 2020, pp. 7508–7517.
- [26] W. Lu, X. Jiang, X. Jin, Y.-L. Yang, M. Gong, T. Wang, K. Shi, and H. Zhao, "Grig: Few-shot generative residual image inpainting," *arXiv* preprint, vol. arXiv:2304.12035, 2023.
- [27] S. Zhao, J. Cui, Y. Sheng, Y. Dong, X. Liang, E. I. Chang, and Y. Xu, "Large scale image completion via co-modulated generative adversarial networks," in *ICLR*, 2021.
- [28] T. Portenier, Q. Hu, A. Szabó, S. A. Bigdeli, P. Favaro, and M. Zwicker, "Faceshop: Deep sketch-based face image editing," *TOG*, vol. 37, no. 4, 2018.
- [29] W. Xiong, J. Yu, Z. Lin, J. Yang, X. Lu, C. Barnes, and J. Luo, "Foreground-aware image inpainting," in CVPR, 2019, pp. 5833–5841.
- [30] Y. Ren, X. Yu, R. Zhang, T. H. Li, S. Liu, and G. Li, "Structureflow: Image inpainting via structure-aware appearance flow," in *ICCV*, 2019, pp. 181–190.
- [31] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo, "Stargan: Unified generative adversarial networks for multi-domain image-toimage translation," in *CVPR*, 2018, pp. 8789–8797.
- [32] Y. Choi, Y. Uh, J. Yoo, and J.-W. Ha, "Stargan v2: Diverse image synthesis for multiple domains," in CVPR, 2020, pp. 8185–8194.

- [33] E. Richardson, Y. Alaluf, O. Patashnik, Y. Nitzan, Y. Azar, S. Shapiro, and D. Cohen-Or, "Encoding in style: a stylegan encoder for image-toimage translation," in *CVPR*, 2021, pp. 2287–2296.
- [34] S.-Y. Chen, F.-L. Liu, Y.-K. Lai, P. L. Rosin, C. Li, H. Fu, and L. Gao, "DeepFaceEditing: Deep face generation and editing with disentangled geometry and appearance control," *TOG*, vol. 40, no. 4, pp. 90:1–90:15, 2021.
- [35] L. Zhang, A. Rao, and M. Agrawala, "Adding conditional control to text-to-image diffusion models," in *ICCV*, 2023, pp. 3836–3847.
- [36] C.-H. Lee, Z. Liu, L. Wu, and P. Luo, "Maskgan: Towards diverse and interactive facial image manipulation," in CVPR, 2020, pp. 5548–5557.
- [37] A. Chen, R. Liu, L. Xie, Z. Chen, H. Su, and J. Yu, "Sofgan: A portrait image generator with dynamic styling," *TOG*, vol. 41, no. 1, p. Article 1, 2022.
- [38] P. Zhu, R. Abdal, Y. Qin, and P. Wonka, "Sean: Image synthesis with semantic region-adaptive normalization," in CVPR, 2020.
- [39] Y. Shi, X. Yang, Y. Wan, and X. Shen, "Semanticstylegan: Learning compositional generative priors for controllable image synthesis and editing," in *CVPR*, 2022, pp. 11254–11264.
- [40] W. Wang, J. Bao, W. Zhou, D. Chen, D. Chen, L. Yuan, and H. Li, "Semantic image synthesis via diffusion models," *arXiv preprint* arXiv:2207.00050, 2022.
- [41] J. Sun, X. Wang, Y. Shi, L. Wang, J. Wang, and Y. Liu, "Ide-3d: Interactive disentangled editing for high-resolution 3d-aware portrait synthesis," *TOG*, vol. 41, no. 6, 2022.
- [42] T. Xiao, J. Hong, and J. Ma, "Elegant: Exchanging latent encodings with gan for transferring multiple face attributes," in *ECCV*, 2018, pp. 172–187.
- [43] J. Guo, Z. Qian, Z. Zhou, and Y. Liu, "Mulgan: Facial attribute editing by exemplar," *arXiv preprint*, vol. arXiv:1912.12396, 2019.
- [44] L. Li, J. Bao, H. Yang, D. Chen, and F. Wen, "Advancing high fidelity identity swapping for forgery detection," in CVPR, 2020, pp. 5073–5082.
- [45] X. Liu, Y. Yang, and P. Hall, "Learning to warp for style transfer," in CVPR, 2021, pp. 3701–3710.
- [46] R. Chen, X. Chen, B. Ni, and Y. Ge, "Simswap: An efficient framework for high fidelity face swapping," in ACM MM, 2020, pp. 2003–2011.
- [47] X. Li, S. Zhang, J. Hu, L. Cao, X. Hong, X. Mao, F. Huang, Y. Wu, and R. Ji, "Image-to-image translation via hierarchical style disentanglement," in *CVPR*, 2021, pp. 8639–8648.
- [48] Y. Zhao, B. Price, S. Cohen, and D. Gurari, "Guided image inpainting: Replacing an image region by pulling content from another image," in WACV, 2019, pp. 1514–1523.
- [49] R. Zhang, W. Li, Y. Zhang, H. Zhang, J. Yu, R. Yang, and W. Xu, "Image re-composition via regional content-style decoupling," in ACM MM, 2021, pp. 3–11.
- [50] H. Kim, Y. Choi, J. Kim, S. Yoo, and Y. Uh, "Exploiting spatial dimensions of latent in gan for real-time image editing," in *CVPR*, 2021, pp. 852–861.
- [51] Y. Wu, Y. Yang, and X. Jin, "Hairmapper: Removing hair from portraits using gans," in CVPR, 2022, pp. 4227–4236.
- [52] Y. Shen, C. Yang, X. Tang, and B. Zhou, "Interfacegan: Interpreting the disentangled face representation learned by gans," *TPAMI*, vol. 44, no. 4, pp. 2004–2018, 2022.
- [53] R. Abdal, P. Zhu, N. J. Mitra, and P. Wonka, "Styleflow: Attributeconditioned exploration of stylegan-generated images using conditional continuous normalizing flows," *TOG*, vol. 40, no. 3, may 2021.
- [54] R. Abdal, P. Zhu, J. Femiani, N. Mitra, and P. Wonka, "Clip2stylegan: Unsupervised extraction of stylegan edit directions," in *SIGGRAPH*, 2022, pp. 48:1–48:9.
- [55] Y. Shen and B. Zhou, "Closed-form factorization of latent semantics in gans," in CVPR, 2020, pp. 1532–1540.
- [56] E. Härkönen, A. Hertzmann, J. Lehtinen, and S. Paris, "Ganspace: Discovering interpretable gan controls," in *NeurIPS*, vol. 33, 2020, pp. 9841–9850.
- [57] O. Patashnik, Z. Wu, E. Shechtman, D. Cohen-Or, and D. Lischinski, "Styleclip: Text-driven manipulation of stylegan imagery," in *ICCV*, 2021, pp. 2085–2094.
- [58] Y. Zhu, H. Liu, Y. Song, Z. Yuan, X. Han, C. Yuan, Q. Chen, and J. Wang, "One model to edit them all: Free-form text-driven image manipulation with semantic modulations," in *NeurIPS*, vol. 35, 2022, pp. 25146–25159.
- [59] A. C. Baykal, A. B. Anees, D. Ceylan, E. Erdem, A. Erdem, and D. Yuret, "Clip-guided stylegan inversion for text-driven real image editing," *TOG*, vol. 42, no. 5, 2023.
- [60] R. Abdal, Y. Qin, and P. Wonka, "Image2stylegan: How to embed images into the stylegan latent space?" in *ICCV*, 2019, pp. 4431–4440.

- [61] ——, "Image2stylegan++: How to edit the embedded images?" in *CVPR*, 2020, pp. 8293–8302.
- [62] D. Roich, R. Mokady, A. H. Bermano, and D. Cohen-Or, "Pivotal tuning for latent-based editing of real images," *TOG*, vol. 42, no. 1, 2022.
- [63] J. Zhu, Y. Shen, D. Zhao, and B. Zhou, "In-domain gan inversion for real image editing," in ECCV, 2020, pp. 592–608.
- [64] O. Tov, Y. Alaluf, Y. Nitzan, O. Patashnik, and D. Cohen-Or, "Designing an encoder for stylegan image manipulation," *TOG*, vol. 40, no. 4, p. Article 133, 2021.
- [65] T. Wang, Y. Zhang, Y. Fan, J. Wang, and Q. Chen, "High-fidelity gan inversion for image attribute editing," in CVPR, 2022, pp. 11 379–11 388.
- [66] B. Li, T. Ma, P. Zhang, M. Hua, W. Liu, Q. He, and Z. Yi, "Reganie: Rectifying gan inversion errors for accurate real image editing," AAAI, vol. 37, no. 1, pp. 1269–1277, 2023.
- [67] Y. Shi, N. Siddharth, B. Paige, and P. H. S. Torr, "Variational mixtureof-experts autoencoders for multi-modal deep generative models," in *NeurIPS*, 2019.
- [68] T. M. Sutter, I. Daunhawer, and J. E. Vogt, "Generalized multimodal ELBO," in *ICLR*, 2021.
- [69] M. Suzuki, K. Nakayama, and Y. Matsuo, "Joint multimodal learning with deep generative models," in *ICLRW*, 2017.
- [70] F. Zhan, Y. Yu, R. Wu, J. Zhang, S. Lu, L. Liu, A. Kortylewski, C. Theobalt, and E. Xing, "Multimodal image synthesis and editing: The generative ai era," *TPAMI*, vol. 45, no. 12, pp. 15098–15119, 2023.
- [71] S. Bounareli, V. Argyriou, and G. Tzimiropoulos, "Finding directions in gan's latent space for neural face reenactment," *BMVC*, 2022.
- [72] A. K. Tewari, M. A. Elgharib, G. Bharaj, F. Bernard, H.-P. Seidel, P. Pérez, M. Zollhöfer, and C. Theobalt, "Stylerig: Rigging stylegan for 3d control over portrait images," *CVPR*, pp. 6141–6150, 2020.
- [73] Y. Shen, J. Gu, X. Tang, and B. Zhou, "Interpreting the latent space of gans for semantic face editing," in CVPR, 2020.
- [74] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *ICML*, 2021, pp. 8748–8763.
- [75] F. Shamshad, M. Naseer, and K. Nandakumar, "Clip2protect: Protecting facial privacy using text-guided makeup via adversarial latent search," in CVPR, 2023, pp. 20595–20605.
- [76] N. Popović, R. Chakraborty, D. P. Paudel, T. Probst, and L. Van Gool, "Spatially multi-conditional image generation," in WACV, 2023, pp. 734–743.
- [77] S. Zhao, Z. Liu, J. Lin, J.-Y. Zhu, and S. Han, "Differentiable augmentation for data-efficient gan training," in *NeurIPS*, vol. 33, 2020, pp. 7559–7570.
- [78] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in CVPR, 2019, pp. 4685–4694.
- [79] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *CVPR*, 2018, pp. 586–595.
- [80] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *ICLR*, 2015.
- [81] X. Hou, L. Shen, O. Patashnik, D. Cohen-Or, and H. Huang, "Feat: Face editing with attention," *arXiv preprint*, vol. arXiv:2202.02713, 2022.
- [82] L. Mescheder, S. Nowozin, and A. Geiger, "Which training methods for gans do actually converge?" in *ICML*, vol. 80, 2018, pp. 3481–3490.
- [83] Zllrunning, "Face parsing in pytorch," 2023, [Online; accessed 14-Feb-2023]. [Online]. Available: https://github.com/zllrunning/face-parsing. PyTorch
- [84] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of GANs for improved quality, stability, and variation," in *ICLR*, 2018.
- [85] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "Gans trained by a two time-scale update rule converge to a local nash equilibrium," in *NeurIPS*, vol. 30, 2017, pp. 6629–6640.
- [86] G. Liu, F. A. Reda, K. J. Shih, T.-C. Wang, A. Tao, and B. Catanzaro, "Image inpainting for irregular holes using partial convolutions," in *ECCV*, 2018, pp. 89–105.
- [87] Y. Jiang, Z. Huang, X. Pan, C. C. Loy, and Z. Liu, "Talk-to-edit: Finegrained facial editing via dialog," in *ICCV*, 2021, pp. 13799–13808.
- [88] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. Huang, "Free-form image inpainting with gated convolution," in *ICCV*, 2019, pp. 4470–4479.
- [89] H. Hu, X. Wang, J. Sun, Y. Fan, Y. Guo, and C. Jiang, "Vectortalker: Svg talking face generation with progressive vectorisation," *arXiv preprint*, vol. arXiv:2312.11568, 2023.



Wanglong Lu is a Ph.D. student in the Department of Computer Science at Memorial University of Newfoundland, Canada. He received his B.Sc. degree in digital media technology from the Communication University of Zhejiang, China, in 2018, and his M.Sc. degree in computer software and theory from Wenzhou University, China, in 2021. His research interests include image editing, image generation, and image recognition.



Jikai Wang received his B.Sc. degree from Shaoxing University Yuanpei College, China, in 2022. He is a graduate student in the College of Computer Science and Artificial Intelligence at Wenzhou University, China. His current research interests include computer graphics, computer vision, and deep learning.



Xianta Jiang is currently an Associate Professor in the Department of Computer Science at Memorial University of Newfoundland. He received a dual Ph.D. in computer science from Simon Fraser University and Zhejiang University (2015). He then took a post-doctoral fellowship in Engineering Science, Simon Fraser University and worked as a senior research associate in the Department of Surgery at the University of Alberta. His research interests include machine intelligence, HCI, and wearables.



Xiaogang Jin is a Professor in the State Key Laboratory of CAD&CG, Zhejiang University. He received a B.Sc. degree in computer science and M.Sc. and Ph.D degrees in applied mathematics from Zhejiang University, P.R. China, in 1989, 1992, and 1995, respectively. His current research interests include image processing, digital humans, traffic simulation, virtual try-on, digital faces, implicit surface modeling and applications, creative modeling, computergenerated marbling, sketch-based modeling, and vir-

tual reality. He received an ACM Recognition of Service Award in 2015 and two Best Paper Awards from CASA 2017 and CASA 2018. He is a member of the IEEE and ACM.



Hanli Zhao is a Professor in the College of Computer Science and Artificial Intelligence at Wenzhou University, China. He received his B.Sc. degree in software engineering from Sichuan University, China, in 2004, and his Ph.D. degree in computer science from the State Key Lab of CAD&CG, Zhejiang University in 2009. His current research interests include computer graphics, computer vision, pattern recognition, medical image analysis, and deep learning.