

SegTalker: Segmentation-based Talking Face Generation with Mask-guided Local Editing

Lingyu Xiong
South China University of Technology
Guangzhou, China
xiongly01@gmail.com

Xize Cheng*
Zhejiang University
Hangzhou, China
chengxize@zju.edu.cn

Jintao Tan
South China University of Technology
Guangzhou, China

Xianjia Wu
Huawei Cloud Computing
Technologies Co., Ltd
Shenzhen, China

Xiandong Li
Huawei Cloud Computing
Technologies Co., Ltd
Shenzhen, China

Lei Zhu
Peking University
Shenzhen, China

Fei Ma
Guangdong Laboratory of Artificial
Intelligence and Digital Economy (SZ)
Shenzhen, China

Minglei Li
Huawei Cloud Computing
Technologies Co., Ltd
Shenzhen, China

Huang Xu
Huawei Cloud Computing
Technologies Co., Ltd
Shenzhen, China

Zhihu Hu*
South China University of Technology
Guangzhou, China
eezhhu@scut.edu.cn

Abstract

Audio-driven talking face generation aims to synthesize video with lip movements synchronized to input audio. However, current generative techniques face challenges in preserving intricate regional textures (skin, teeth). To address the aforementioned challenges, we propose a novel framework called SegTalker to decouple lip movements and image textures by introducing segmentation as intermediate representation. Specifically, given the mask of image employed by a parsing network, we first leverage the speech to drive the mask and generate talking segmentation. Then we disentangle semantic regions of image into style codes using a mask-guided encoder. Ultimately, we inject the previously generated talking segmentation and style codes into a mask-guided StyleGAN to synthesize video frame. In this way, most of textures are fully preserved. Moreover, our approach can inherently achieve background separation and facilitate mask-guided facial local editing. In particular, by editing the mask and swapping the region textures from a given reference image (e.g. hair, lip, eyebrows), our approach enables facial editing seamlessly when generating talking face video. Experiments demonstrate that our proposed approach can effectively preserve texture details and generate temporally consistent video

while remaining competitive in lip synchronization. Quantitative and qualitative results on the HDTF and MEAD datasets illustrate the superior performance of our method over existing methods.

CCS Concepts

• **Computing methodologies** → **Computer vision tasks**; • **Information systems** → **Multimedia content creation**.

Keywords

Video Generation, Talking Face Generation, Attribute Editing

ACM Reference Format:

Lingyu Xiong, Xize Cheng, Jintao Tan, Xianjia Wu, Xiandong Li, Lei Zhu, Fei Ma, Minglei Li, Huang Xu, and Zhihu Hu. 2024. SegTalker: Segmentation-based Talking Face Generation with Mask-guided Local Editing. In *Proceedings of Proceedings of the 32nd ACM International Conference on Multimedia (MM '24)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3664647.3681108>

1 Introduction

Talking face generation, which aims to synthesize facial imagery precisely synchronized with input speech, has garnered substantial research attention in the field of computer vision and multimedia [10, 15, 16] due to its numerous applications including digital human, virtual conference and video dubbing [8, 9, 22, 44, 48].

There are many attempts to realize high-fidelity talking face. Early approaches first predict mouth shapes from speech using recurrent neural networks, then generate the face conditioned on the shapes [37]. Recent end-to-end methods directly map speech spectrograms to video frames leveraging different intermediate representations [22, 36, 38, 57, 59]. Zhang et al. [57] takes advantage of 3D Morphable Models (3DMMs), a parametric model that

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MM '24, October 28–November 1, 2024, Melbourne, VIC, Australia

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0686-8/24/10
<https://doi.org/10.1145/3664647.3681108>

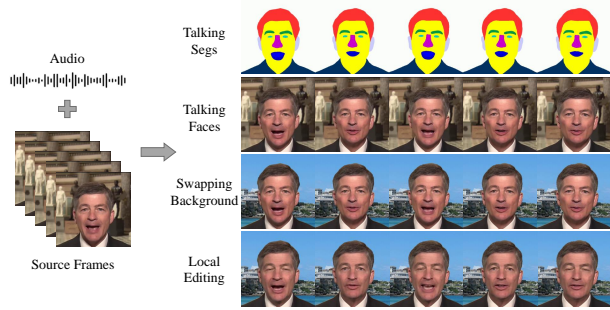


Figure 1: Given a talking video and another speech, SegTalker can produce high-fidelity and synchronized video with rich textures (row 2), enabling swapping background (row 3) and local editing such as blinking (row 4).

decomposes expression, pose, and identity, to transfer facial motions. Zhou et al. [59] employs the landmark as the representation. Meshry et al. [22] factorize the talking-head synthesis process into spatial and style components through coarse-grained masks, but they do not facilitate texture disentanglement and facial editing. More recently, Kicanaoglu et al. [19] perform unsupervised vector quantization on intermediate feature maps of StyleGAN to generate abundant semantic regions for local editing. Despite improvements in photo-realism, current talking face methods still face challenges in preserving identity-specific details such as hair, skin textures and teeth. Furthermore, within the current landscape of talking face generation methods, there is no single technique that can concurrently accomplish facial editing and background replacement. Our method elegantly incorporates facial editing into talking face generation in an end-to-end manner through the intermediate representation of segmentation.

In this paper, we aim to design a unified approach that realizes the controllable talking face synthesis and editing. We propose a novel framework termed SegTalker that explicitly disentangles textural details with lip movements by utilizing segmentation. Our framework consists of an audio-driven talking segmentation generation (TSG) module, followed by a segmentation-guided GAN injection (SGI) network to synthesize animation video. We utilize a pre-trained network [53] to extract segmentation as prior information to decompose semantic regions and enhance textural details, seamlessly enabling fine-grained facial local editing and background replacement. Specifically, given the input image and speech, we first conduct face parsing to obtain the segmentation. Subsequently, TSG module extracts image and speech embedding and then combines these embeddings to synthesize new segmentation with lips synchronized to the input speech. After that, SGI module employs a multi-scale encoder to project the input face into the latent space of StyleGAN [18]. Each facial region has a set of style codes for different layers of the StyleGAN generator. Then We inject the synthesized mask and style codes into the mask-guided generator to obtain the talking face. In this way, the structural information and textures of facial components are fully disentangled. Furthermore, facial local editing can be accomplished by simply modifying the synthesized mask or swapping the region textures from a given reference image, achieving seamless integration with talking face synthesis. In summary, our contributions are:

- We propose a novel framework that utilizes segmentation as an intermediate representation to disentangle the lip movements with image reconstruction for talking face generation, achieving consistent lip movements and preserving fine-grained textures.
- We employ a multi-scale encoder and mask-guided generator to realize the local control for different semantic regions. By manipulating the masks and smoothly swapping the textures, we can seamlessly integrate the facial local editing into the talking face pipeline and conduct swapping background.
- Experiments on HDTF and MEAD datasets demonstrate our superiority over state-of-the-art methods in visual quality, ID preservation and temporal consistency.

2 Related Work

2.1 Audio-driven Talking Face Generation

Talking face generation, which aims to synthesize photo-realistic video of a talking person giving a speech as input, has garnered increasing research attention in recent years. With the emergence of generative adversarial networks (GANs) [12], many methods [22, 25, 36, 57, 59] have been proposed for synthesizing animation video. In terms of the intermediate representations, the existing works can be categorized into landmark-based, 3D-based and others. In the landmark-based methods, Suwajanakorn et al. [37] use recurrent neural network (RNN) to build the mapping from the input speech to mouth landmark, and then generate mouth texture. Zhou et al. [59] combine LSTM and self-attention to predict the locations of landmarks. Zhong et al. [58] utilizes transformer to predict landmarks, then combines multi-source features (prior information, landmarks, speech) to synthesize talking face. Recently, DiffTalk [32] takes speech and landmarks as conditioned inputs and utilizes a latent diffusion model [28] to generate talking faces. For 3DMM-based method, SadTalker [56] learns realistic 3D motion coefficients for stylized audio-driven single-image talking face animation, achieving high-quality results by explicitly modeling audio-motion connections. Some styleGAN-based method [5, 39, 52] such as StyleHEAT [52] leverages a pre-trained StyleGAN to achieve high-resolution editable talking face generation from a single portrait image, allowing disentangled control via audio. More recently, the emergence of neural radiance field (NeRF) provides a new perspective for 3D-aware talking face generation [13, 31]. However, these intermediate representations have difficulty in capturing fine-grained details and preserving identity i.e., teeth and skin textures which degrade the visual quality heavily. Wav2Lip [25] adopts the encoder-decoder architecture to synthesize animation videos. However, there are conspicuous artifacts with a low resolution in the synthesized videos. In this work, we employ a novel representation, segmentation, to disentangle lip movement with image reconstruction, and further extract per-region features to preserve texture details.

2.2 GAN Inversion

GAN inversion aims to invert real images into the latent space of a generator for reconstruction and editing. Several StyleGAN inversion methods have been proposed, they can typically be divided

into three major groups of methods: 1) gradient-based optimization of the latent code [1, 2, 17, 30], 2) encoder-based [3, 26, 40, 50, 51] and 3) fine-tune methods [4, 23, 27, 54]. The gradient-based optimization methods directly optimize the latent code using gradient from the loss between the real image and the generated one. The encoder-based methods train an encoder network over a large number of samples to directly map the RGB image to latent code. The gradient-based optimization methods always give better performance while the encoder-based cost less time. The fine-tuning methods make a trade-off between the above two and use the inverted latent code from encoder as the initialization code to further optimization. However, existing works focus on global editing and cannot make fine-grained control of the local regions. Our method uses a variation of [26] to realize local editing via manipulating a novel \mathcal{W}^{c+} [21] latent space.

3 Proposed Methods

3.1 Overview

To tackle the lack of regional textures in talking face generation, we explicitly disentangle semantic regions by introducing segmentation mechanism. Leveraging segmentation as an intermediate representation, our approach decouples audio-driven mouth animation and image texture injection. The speech is solely responsible for driving the lip contours, while the injection module focuses on extracting per-region textures to generate the animation video. The overall framework of our proposed model, termed SegTalker, is illustrated in fig. 2. The pipeline consists of two sub-networks: (1) talking segmentation generation (TSG) and (2) segmentation-guided GAN injection network (SGI), which are elaborated in section 3.2 and section 3.3, respectively.

3.2 Talking Segmentation Generation

The first proxy sub-network is the talking segmentation generation (TSG) module. Given speech and image frame, this network first employs parsing network [53] to extract mask and then synthesizes talking segmentation. The original network generates 19 categories in total. For the sake of simplicity, we merge the same semantic class (e.g. left and right eyes), resulting in 12 final classes. During pre-processing, video is unified to 25fps with speech sampled at 16kHz. To incorporate temporal information, Following [43], the global and local features are extracted as speech embedding. We employ **mask encoder** to extract visual embedding from two masks: a pose source and an identity reference. The two masks are concatenated in the channel dimension. The pose source aligns with the target segmentation but with the lower half occluded. The identity reference provides facial structural information of the lower half to facilitate training and convergence. Without concerning textural information, the model only focuses on learning the structural mapping from speech to lip movements.

We employ a CNN-based network to extract the embedding of a 0.2-second audio segment whose center is synchronized with the pose source. Similar to text, speech always contains sequential information. To better capture temporally relevant features, we employ the pre-trained AV-Hubert [34, 35] as a part of **audio encoder** to extract long-range dependencies. AV-Hubert has conducted pre-training for audio-visual alignment [46, 47], so the

extracted embedding is very close to the semantic space of the video. When using AV-Hubert to extract audio embedding, we only need to feed speech and the visual signal is masked. Specifically, given a 3s speech chunk, we feed it into the Transformer-based AV-Hubert to produce contextualized speech features. We then extract the feature embedding corresponding to the given image segment. Given the mixed speech embedding and visual embedding, **Generator** synthesizes the final talking segmentation. We adopt U-Net [29] as the backbone architecture. In addition, skip connections and transposed convolutions are utilized for feature fusion and up-sampling.

Given a mask synthesized by the model and the ground truth mask, we employ two types of losses to improve generation quality i.e., the reconstruction loss and the syncnet loss.

Reconstruction Loss Unlike previous generative tasks that often synthesize RGB image and adopt L1 loss for reconstruction, talking segmentation involves generating segmentation where each pixel denotes a particular class. To stay consistent with semantic segmentation task, we employ cross entropy loss as our reconstruction loss. Given N_i generated masks y_i and ground truth \hat{y}_i with a categories of M regions, the cross entropy loss is defined as:

$$\mathcal{L}_{ce} = \frac{1}{N_i} \sum_i \sum_{c=1}^M y_{ic} \log(y_{ic}) \quad (1)$$

where y_{ic} denotes the one-hot encoded vector for the i -th generated segmentation belonging to the c -th category. For generated segmentation, different classes occupy varying proportions of areas. Semantically important regions like lips and eyes constitute small fractions, while background dominates most areas. To mitigate the class imbalance issue, a weighted cross entropy is formulated as:

$$\mathcal{L}_{w-ce} = \frac{1}{N_i} \sum_i \sum_{c=1}^M w_c y_{ic} \log(y_{ic}) \quad (2)$$

where w_c denotes the weight for the corresponding category and is determined on the inverse proportionality of the areas of different regions on the whole dataset.

SyncNet Loss The reconstruction loss mainly restores images at the pixel level without effective semantic supervision. Therefore, we train a segmentation-domain SyncNet from [25] to supervise lip synchronization. During training, a speech chunk is randomly sampled from speech sequences, which can be either synchronized (positive example) or unsynchronized (negative example). The SyncNet consists of a speech encoder and a mask encoder. For the mask, we use one-hot encoding as input and concatenate T_o masks along the channel dimension. Specifically, the SyncNet takes inputs of a window T_o of consecutive lower-half frames and a speech segment S . After passing through the speech encoder and mask encoder, 512-dim embeddings $s = E_{\text{speech}}(S)$ and $m = E_{\text{mask}}(T_o)$ are obtained respectively. Cosine similarity distance and binary cross entropy loss are then calculated between the embeddings. The losses are formally defined as:

$$P_{\text{sync}} = \frac{s \cdot m}{\max(\|s\|_2 \cdot \|m\|_2, \epsilon)} \quad (3)$$

$$\mathcal{L}_{\text{sync}} = \frac{1}{N} \sum_i -\log(P_{\text{sync}}^i) \quad (4)$$

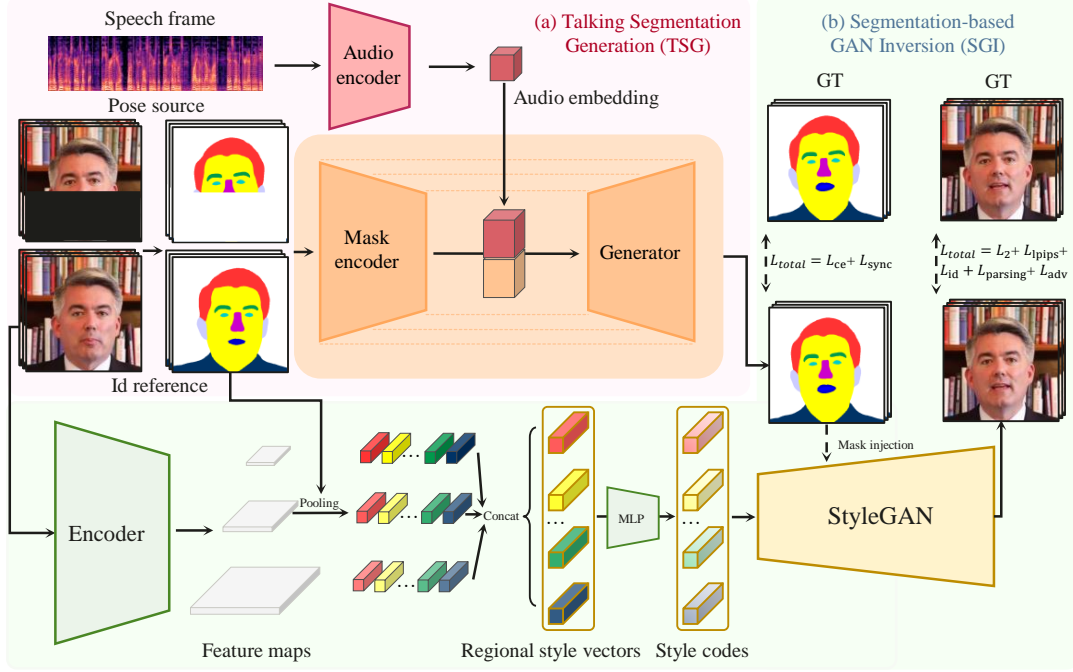


Figure 2: Overview of the proposed SegTalker framework for talking face generation. (a) talking segmentation generation (TSG) module takes mel and mask as inputs, then synthesizes the talking segmentation with lip synchronized to input speech. (b) Given reference image and mask from TSG, segmentation-guided GAN injection (SGI) network utilizes a mask-guided multi-scale encoder to extract different semantic region codes, then injects the style codes and synthesized mask from TSG into the mask-guided generator to obtain the final talking face image.

where P_{sync} is a single value between $[0, 1]$ and N is the batch size. ϵ is used to prevent division by zero. We train the lip-sync expert on the HDFT dataset [57] with a batch size of 8, $T_v = 5$ frames, $S = 0.2s$ segment, using the Adam optimizer with a learning rate of $1e-4$. After approximately one day of training, the model converges. Our expert network eventually achieves 81% accuracy on the test set.

3.3 Segmentation-guided GAN Injection

The second sub-network illustrated in fig. 2 is the segmentation-guided GAN injection (SGI) network. Given a portrait and its corresponding mask, SGI first encodes the image into the latent space to obtain the latent code, then inverts the generated latent code back to the image domain through style injection.

There exist various latent spaces such as \mathcal{W} , $\mathcal{W}+$ and \mathcal{S} space. Many works [1, 26, 33, 40, 49] have investigated their representational abilities from the perspectives of distortion, perception, and editability. Here, we choose \mathcal{W}^{c+} space, a variation of $\mathcal{W}+$ space originated from [21] as representation of latent code. To leverage this representation, a powerful encoder is required to accurately map each input image to a corresponding code. Although many encoders [26, 40, 49] have been proposed, they focus on extracting global latent code for global editing, such as age, emotion, making them unsuitable for textures disentanglement and local editing. To this end, we adopt a variation of [26] for latent code extraction. The encoder utilizes a feature pyramid network (FPN) [20] for feature

fusion, ultimately generating fine-grained, medium-grained, and coarse-grained feature maps at three different scales. The mask is then resized to match each feature map. Subsequently, a global average pooling (GAP) is employed to extract semantic region features according to the segmentation, resulting in multi-scale style vectors. These are concatenated and further passed through an MLP to obtain the \mathcal{W}^{c+} style codes.

Specifically, given a source image I and its corresponding mask M , we first utilize a multi-scale encoder E_ϕ to obtain the feature maps $F = [F_i]_{i=1}^N$ at different resolutions:

$$F = E_\phi(I) \quad (5)$$

Here N is equal to three. We then aggregate per-region features based on the mask M and features F . Specifically, for each feature map F_i , we first downsample the mask to match the feature map size, then perform global average pooling (GAP) to aggregate features for different regions:

$$u_{ij} = \text{GAP}(F_i \circ (\text{Down}(M)_i = j)), \{j = 1, 2, \dots, C\} \quad (6)$$

Where u_{ij} denotes the averaged feature of region j in feature map i , C is the number of semantic regions, $\text{Down}(\dots)$ is the downsampling operation to align with F_i , and $\langle \circ \rangle$ is the element-wise product. Subsequently, the multi-scale feature vectors $\{u_{ij}\}_{i=1}^N$ of region j are concatenated and passed through a multi-layer perceptron (MLP) to obtain the style codes:

$$s_j = \text{MLP}([\{u_{ij}\}_{i=1}^N]) \quad (7)$$

Table 1: Quantitative comparisons on the HDTF dataset. The best performance is highlighted in red (1st best) and blue (2nd best). Local Editing and Swapping Background demonstrate the method’s capabilities for facial attribute editing and swapping background. For better visualization, we scale up the FVD by a factor of 0.1.

Method	Local Editing	Swapping Background	Synchronization			Visual Quality				
			Sync \uparrow	F-LMD \downarrow	M-LMD \downarrow	FID \downarrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	FVD $\downarrow \times 0.1$
Real Video	N/A	N/A	7.305	0	0	2.623	0.0109	-	1.000	1.048
Wav2Lip [25]	N/A	N/A	8.413	3.727	4.528	16.299	0.1031	31.821	0.921	17.423
SadTalker [56]	Blink Only	N/A	6.725	34.970	36.172	59.059	0.7567	15.967	0.496	85.638
DiffTalk [32]	N/A	N/A	4.615	2.004	1.614	23.498	0.111	31.654	0.913	16.606
StyleHEAT [52]	Blink Only	N/A	6.296	29.672	31.390	111.229	0.7969	14.968	0.465	91.2356
AD-NeRF [13]	N/A	Yes	5.216	3.574	3.825	18.614	0.1013	30.640	0.923	18.438
Ours	Yes	Yes	6.872	3.405	3.173	10.348	0.0494	33.590	0.934	9.205

where s_j denotes the style code of j -th categories. Then, the mask and style codes $s \in \mathbb{R}^{C \times 18 \times 512}$ are fed into the mask-guided StyleGAN generator to synthesize the talking face. For the detailed architecture of the Mask-guided StyleGAN, please refer to the supplement.

Prior Learning To seamlessly integrate SGI into the overall framework, we randomly select a mask from the images within a 15-frame range of the input image. Through such a training strategy, the model can learn the priors of semantic regions like teeth and eyes. Specifically, when given an image with closed mouth but a randomly selected mask corresponds to a visible-teeth state, it learns to model the teeth’ prior information and can naturally connect with the TSG module.

Loss Functions SGI is trained with a series of weighted objectives: pixel-wise and LPIPS loss for perpetual quality, id loss to prevent identity drift, a face parsing loss and an adversarial loss. The details are described in our Appendix.

4 Experiments

4.1 Experimental settings

Dataset Since StyleGAN [18] typically generates high resolution images, e.g. 512 or 1024, while most existing talking face datasets have a lower resolution of 256 or below, we opt to train on the HDTF dataset [57] for high-quality talking face synthesis. The HDTF dataset is collected from YouTube website published in the last two years, comprising around 16 hours of videos ranging from 720P to 1080P resolution. It contains over 300 subjects and 10k distinct sentences. We collected a total of 392 videos, with 347 used for training and the remaining 45 for testing. The test set comprises videos with complex backgrounds and rich textures, thereby offering a comprehensive evaluation of the model performance.

Metrics We conduct quantitative evaluations on several widely used metrics. To evaluate the lip synchronization, we adopt the confidence score of SyncNet [11] (**Sync**) and Landmark Distance around mouths (**M-LMD**) [6]. To evaluate the accuracy of generated facial expressions, we adopt the Landmark Distance on the whole face (**F-LMD**). To evaluate the quality of generated talking face videos, we adopt **PSNR** [43], **SSIM** [45], **FID** [14] and **LPIPS** [55]. To measure the Temporal coherence of generated videos, we employ **FVD** [41]. Higher scores indicate better performance for Sync,

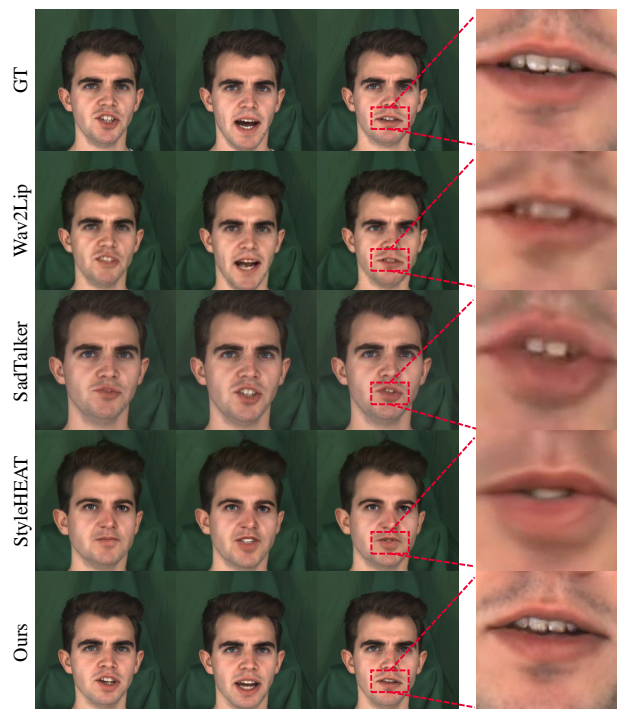


Figure 3: Qualitative results on MEAD dataset. Our model demonstrates superior overall performance, particularly in preserving fine-grained textures (e.g. teeth)

PSNR, and SSIM, while lower scores are better for F-LMD, M-LMD, FID, LPIPS and FVD.

Implementation Details We use PyTorch [24] to implement our framework. We train TSG module on a single NVIDIA A100 GPU with 40GB, while SGI module is trained on 4 NVIDIA A100 GPUs. In stage 1, We crop and resize face to 512 \times 512. Speech waveforms are pre-processed to mel-spectrogram with hop and window lengths, and mel bins are 12.5ms, 50ms, and 80. The batch size is set to 20 and the Adam solver with an initial learning rate of 1e-4 ($\beta_1 = 0.5, \beta_2 = 0.999$) is utilized for optimization. In stage 2, we set the batch size to 4 for each GPU and initialize the learning rate as 1e-4 with the Adam optimizer ($\beta_1 = 0.9, \beta_2 = 0.999$). The generator



Figure 4: Qualitative comparisons of our results with several state-of-the-art methods for talking face synthesis. our method produces high-fidelity video frames with rich textural details, while other methods struggle to preserve identity and contain artifacts. It is worth noting that AD-NeRF needs to train on these two identities respectively to produce the results.

is initialized with StyleGAN weights [18]. In the first stage, we train the model with only the cross entropy loss for approximately 50K iterations, then incorporate the expert SyncNet to supervise lip movements for an extra 50k iterations. In the second stage, we train the model for 400K iterations.

4.2 Experimental Results

Qualitative Talking Segmentation Results In the first sub-network, we visualize the talking segmentation results illustrated in fig. 5. It can be observed that the generated segmentations effectively delineate distinct facial regions, even elaborating details such as earrings. Additionally, the synthesized lips exhibit strong synchronization with the ground truth. Subsequently, the high-quality segmentations produced by TSG are utilized as guidance for the SGI to deliver the final output.

Quantitative Results We compare several state of the art methods: Wav2Lip [25], SadTalker [56] (3DMM-based), DiffTalk [32] (diffusion-based), StyleHEAT [52] (styleGAN-based) and AD-NeRF [13] (NeRF-based). We conduct the experiments in the self-driven setting on the test set, where the videos are not seen during training. In these methods, the head poses of Wav2Lip, DiffTalk, and SegTalker are fixed in their samples. For other methods, head poses

are randomly generated. The results of the quantitative evaluation are reported in table 1.

Our method achieves better visual quality, temporal consistency, and also shows comparable performance in terms of lip synchronization metrics. Since DiffTalk takes ground truth landmark as conditional input, it is reasonable for DiffTalk to achieve the lowest LMD in self-driven sets. However, DiffTalk performs poorly in frame-to-frame coherence, especially with significant jitter in the mouth region (see supplementary video). In synchronization, despite scoring slightly lower on metrics relative to Wav2lip, our method achieves a similar score with ground truth videos. Furthermore, Our method outperforms existing state-of-the-art approaches on both pixel-level metrics such as PSNR, as well as high-level perceptual metrics including FID and LPIPS, thereby achieving enhanced visual quality. We additionally measure the FVD metric and Our FVD score is the best. This means that our method is able to generate temporal consistency and visual-satisfied videos. This is largely attributed to the implementation of SGI module. By explicitly disentangling different semantic regions via segmentation, SGI can better preserve texture details during image reconstruction. Moreover, our method is the only approach that can simultaneously



Figure 5: Visualization of synthesized segmentation(row 1, row 2) and real images(row 2, row 4).

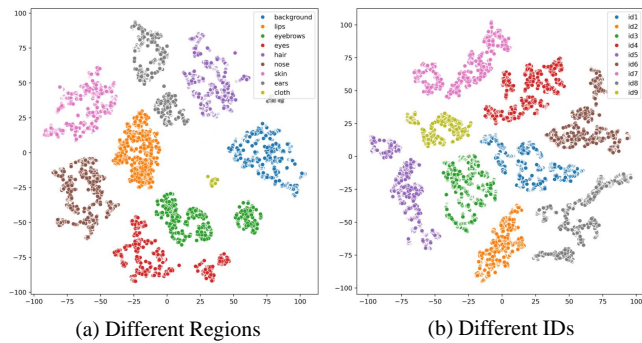


Figure 6: (a) Visualization of the regional style codes of a speaker. (b) Visualization of the style codes of 8 speakers in a particular region (here hair for example).

achieve facial editing and background replacement which will be discussed in the following section.

Qualitative Results To qualitatively evaluate the different methods, we perform uniformly sampled images from two synthesized talking face videos which are shown in fig. 4. Specifically, the ground truth videos are provided in the first row where synthesized images of different methods follow the next and ours are illustrated in the bottom row. In comparison to Wav2lip [25], our results exhibit enhanced detail in the lip and teeth regions. For SadTalker [56], It employs single-frame animation, which inevitably causes background movement and generates artifacts when wrapping motion sequences. Additionally, it also cannot handle the scenarios with changing background. The incorporation of segmentation in our approach allows high-quality background replacement. DiffTalk [7] can generate visually satisfying results; however, diffusion-based methods still face significant challenges in terms of temporal consistency. The mouth area of DiffTalk is prone to shaking and leads to poor lip synchronization performance. StyleHEAT [52] is also a StyleGAN-based approach, but it cannot directly drive speech to generate a talking face video. Instead, it requires the assistance of SadTalker to extract features from the first stage, then warps the

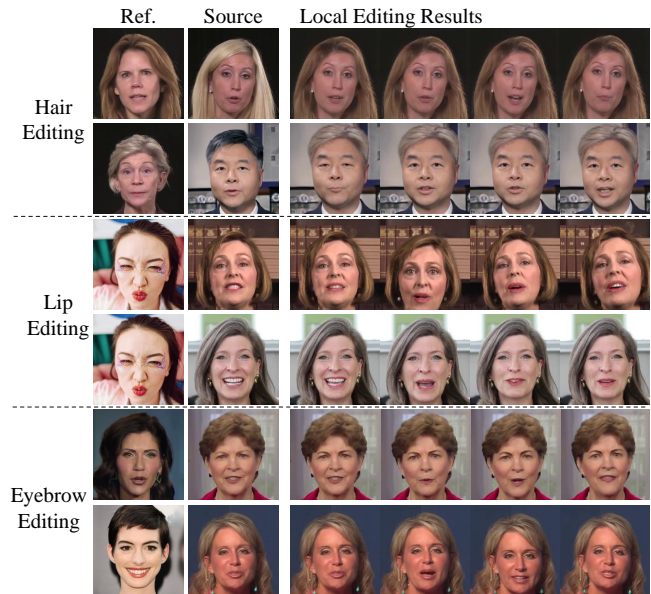


Figure 7: Qualitative results of local editing. Our method produces more high-fidelity results while maintaining the details and identity information of other regions.

features to generate video. Therefore, the quality of the video generated by StyleHEAT is limited by the quality of the output generated by SadTalker. AD-NeRF [13] is a NeRF-based method capable of generating high-quality head part but consistently exists artifact in the connection between the head and neck. Moreover, its inference is time-consuming (10s per image) and requires fine-tuning for each speaker (about 20 hours). In addition, to demonstrate the generalization ability of the proposed method, we conduct validation on another dataset, as illustrated in fig. 3. Compared to other methods, our method excels in preserving texture details, particularly in fine-grained structural regions such as teeth. In contrast, our method can produce more realistic and high-fidelity results while achieving accurate lip sync, satisfactory identity preservation and rich facial textures. For more comparison results, please refer to our demo videos in the supplement materials.

Disentangled Semantics Visualization To demonstrate the Disentanglement of the model across different semantic regions, we employ t-distributed stochastic neighbor embedding(t-SNE) [42] visualization to illustrate the per-region features, as depicted in fig. 6(a). For Clarity, We select eight sufficiently representative regions (appear in all videos) and utilize the mask-guided encoder to extract style codes from these semantic regions. In fig. 6(a), each region is marked with a distinct color. As shown, the style codes of same region cluster in the style space and different semantic regions are explicitly separated substantially. This demonstrates that our mask-guided encoder can accurately disentangle different region features. Furthermore, in fig. 6(b), we visualize the features of different IDs within a particular region to demonstrate the capability of the encoder. It can be seen that the style codes of different IDs are fully disentangled and our model can learn meaningful features.

Facial Editing and Swapping Results Our method also supports facial editing and background swaps while generating video.

Table 2: Ablation results under different configurations. P.L. and C.S. denote prior learning and cross entropy, respectively.

Conf.	Sync \uparrow	FID \downarrow	PSNR \uparrow	SSIM \uparrow
w/o P.L.	5.314	28.254	28.685	0.847
w/o C.S.	3.126	63.264	14.257	0.479
w/o SyncNet	4.174	10.647	32.036	0.913
All pipeline	6.872	10.348	33.590	0.934

Table 3: User Study

Method	Lip Sync.	Face Quality	ID Preserve.	Overall Naturalness
Wav2Lip [25]	32.4%	5.6%	20.4%	15.3%
SadTalker [56]	16.3%	13.7%	22.1%	12.9%
StyleHEAT [52]	9.8%	10.5%	5.4%	6.4%
Ours	41.5%	70.2%	52.1%	65.4%

Given a reference image and a sequence of source images, our method can transfer the candidate region texture to the source images. As depicted in fig. 7, we illustrate three local editing tasks, including fine-grained hair editing, lip makeup, and eyebrow modifications. Besides, we also can manipulate blinking in a controllable manner by simply editing the eye regions of mask, illustrated in fig. 1. Compared with existing blink methods, our method does not design a specialized module for blinking editing, as well as enables other types of local editing, substantially enhancing model applicability and scalability. Additionally, our model intrinsically disentangles the foreground and background, allowing for seamless background swapping and widening the application scenarios of talking faces. As shown in fig. 8, with a provided reference background image and a video segment, we can not only generate synchronized talking face video but also achieve video background swapping, resulting in high-fidelity and photo-realistic video.

4.3 Ablation Study

In this section, we perform an ablation study to evaluate the 2 sub-agent networks, which are shown in table 2. We develop 3 variants with the modification of the framework corresponding to the 2 sub-network: 1) *w/o prior learning*, 2) *w/o cross entropy* and 3) *w/o SyncNet*.

The first component is the implementation of *priors learning*. Without *priors learning*, the method produces poor visual quality. This mechanism offers structural prior information for the mouth and teeth regions, which facilitates the model learning personalized details of these areas.

The second component is the *cross entropy*. Without *cross entropy*, the method exhibits very poor performance whenever on both lip synchronization and visual quality. By employing cross-entropy loss instead of L1 loss, we overcome the issue of erroneous segmentation predictions around region boundaries, improving the model’s control over different semantic areas. Furthermore, *cross entropy* also facilitates learning lip movements from speech, exhibiting a certain extent of lip synchronization.

**Figure 8: Example of Swapping background. Given a video and a background image, our method can produce natural and photo-realistic swapping videos.**

The last component is the *SyncNet*, which is performed to reinforce the model learning the mapping from speech to lip. The performance of visual quality is comparable to the baseline when we do not apply *SyncNet*. However, without *SyncNet* would lead to poor lip synchronization performance, which is demonstrated in table 2.

4.4 User Study

We conduct a user study to evaluate the performance of all the methods. We choose 20 videos with 10 seconds clips as our test set. These samples contain different poses, ages, backgrounds and expressions to show the generalization of our method. We invite 15 participants and let them choose the best method in terms of face quality, lip synchronization, identity preservation and overall naturalness. The results are shown in table 3. It demonstrates that our model outperforms other methods across multiple dimensions.

5 Conclusion

In this paper, we present a new framework SegTalker for talking face generation, which disentangles the lip movements and textures of different facial components by employing a new intermediate representation of segmentation. The overall framework consists of two sub-networks: TSG and SGI network. TSG is responsible for the mapping from speech to lip movement in the segmentation domain and SGI employs a multi-scale encoder to project source image into per-region style codes. Then, a mask-guided generator integrates the style codes and synthesized segmentation to obtain the final frame. Moreover, By simply manipulating different semantic regions of segmentation or swapping the different textures from a reference image, Our method can seamlessly integrate local editing and support coherent swapping background.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. 2019. Image2stylegan: How to embed images into the stylegan latent space?. In *Proceedings of the IEEE/CVF international conference on computer vision*. 4432–4441.
- [2] Rameen Abdal, Yipeng Qin, and Peter Wonka. 2020. Image2stylegan++: How to edit the embedded images?. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 8296–8305.
- [3] Yuval Alaluf, Or Patashnik, and Daniel Cohen-Or. 2021. Restyle: A residual-based stylegan encoder via iterative refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 6711–6720.
- [4] Yuval Alaluf, Omer Tov, Ron Mokady, Rinon Gal, and Amit Bermano. 2022. Hyperstyle: Stylegan inversion with hypernetworks for real image editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 18511–18521.
- [5] Mohammed Mousa Alghamdi, He Wang, Andrew J. Bulpitt, and David C. Hogg. 2022. Talking Head from Speech Audio using a Pre-trained Image Generator. *Proceedings of the 30th ACM International Conference on Multimedia* (2022).
- [6] Lele Chen, Zhiheng Li, Ross K. Maddox, Zhiyao Duan, and Chenliang Xu. 2018. Lip Movements Generation at a Glance. In *European Conference on Computer Vision*. 520–535.
- [7] Kun Cheng, Xiaodong Cun, Yong Zhang, Menghan Xia, Fei Yin, Mingrui Zhu, Xuan Wang, Jue Wang, and Nannan Wang. 2022. Videoretalking: Audio-based lip synchronization for talking head video editing in the wild. In *SIGGRAPH Asia 2022 Conference Papers*. 1–9.
- [8] Xize Cheng, Rongjie Huang, Linjun Li, Tao Jin, Zehan Wang, Aoxiong Yin, Minglei Li, Xinyu Duan, Zhou Zhao, et al. 2023. TransFace: Unit-Based Audio-Visual Speech Synthesizer for Talking Head Translation. *arXiv preprint arXiv:2312.15197* (2023).
- [9] Xize Cheng, Tao Jin, Rongjie Huang, Linjun Li, Wang Lin, Zehan Wang, Ye Wang, Huadai Liu, Aoxiong Yin, and Zhou Zhao. 2023. Mixspeech: Cross-modality self-learning with audio-visual stream mixup for visual speech translation and recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 15735–15745.
- [10] Xize Cheng, Tao Jin, Linjun Li, Wang Lin, Xinyu Duan, and Zhou Zhao. 2023. Opensr: Open-modality speech recognition via maintaining multi-modality alignment. *arXiv preprint arXiv:2306.06410* (2023).
- [11] Joon Son Chung and Andrew Zisserman. 2016. Out of Time: Automated Lip Sync in the Wild. In *ACCV Workshops*.
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems* 27 (2014).
- [13] Yudong Guo, Keyu Chen, Sen Liang, Yong-Jin Liu, Hujun Bao, and Juyong Zhang. 2021. Ad-nerf: Audio driven neural radiance fields for talking head synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5784–5794.
- [14] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* 30 (2017).
- [15] Rongjie Huang, Huadai Liu, Xize Cheng, Yi Ren, Linjun Li, Zhenhui Ye, Jinzheng He, Lichao Zhang, Jinglin Liu, Xiang Yin, et al. 2023. Av-transpeech: Audio-visual robust speech-to-speech translation. *arXiv preprint arXiv:2305.15403* (2023).
- [16] Tao Jin, Xize Cheng, Linjun Li, Wang Lin, Ye Wang, and Zhou Zhao. 2023. Re-thinking missing modality learning from a decoding perspective. In *Proceedings of the 31st ACM International Conference on Multimedia*. 4431–4439.
- [17] Kyoungkook Kang, Seongtae Kim, and Sunghyun Cho. 2021. Gan inversion for out-of-range images with geometric transformations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 13941–13949.
- [18] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 8110–8119.
- [19] Berkay Kicanaoglu, Pablo Garrido, and Gaurav Bharaj. 2023. Unsupervised Facial Performance Editing via Vector-Quantized StyleGAN Representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2371–2382.
- [20] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2117–2125.
- [21] Zhian Liu, Maomao Li, Yong Zhang, Cairong Wang, Qi Zhang, Jue Wang, and Yongwei Nie. 2023. Fine-Grained Face Swapping via Regional GAN Inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8578–8587.
- [22] Moustafa Meshry, Saksham Suri, Larry S Davis, and Abhinav Shrivastava. 2021. Learned spatial representations for few-shot talking-head synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 13829–13838.
- [23] Yotam Nitzan, Kfir Aberman, Qiuwei He, Orly Liba, Michal Yarom, Yossi Gandelsman, Inbar Mosseri, Yael Pritch, and Daniel Cohen-Or. 2022. Mystyle: A personalized generative prior. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–10.
- [24] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019).
- [25] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Nambodiri, and CV Jawahar. 2020. A lip sync expert is all you need for speech to lip generation in the wild. In *Proceedings of the 28th ACM international conference on multimedia*. 484–492.
- [26] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. 2021. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2287–2296.
- [27] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. 2022. Pivotal tuning for latent-based editing of real images. *ACM Transactions on graphics (TOG)* 42, 1 (2022), 1–13.
- [28] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10684–10695.
- [29] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention*. 234–241.
- [30] Rohit Saha, Brendan Duke, Florian Shkurti, Graham W Taylor, and Parham Aarabi. 2021. Loho: Latent optimization of hairstyles via orthogonalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1984–1993.
- [31] Shuai Shen, Wanhua Li, Zheng Zhu, Yueqi Duan, Jie Zhou, and Jiwen Lu. 2022. Learning dynamic facial radiance fields for few-shot talking head synthesis. In *European Conference on Computer Vision*. 666–682.
- [32] Shuai Shen, Wenliang Zhao, Zibin Meng, Wanhua Li, Zheng Zhu, Jie Zhou, and Jiwen Lu. 2023. DiffTalk: Crafting Diffusion Models for Generalized Audio-Driven Portraits Animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1982–1991.
- [33] Yujun Shen, Jinjin Gu, Xiaou Tang, and Bolei Zhou. 2020. Interpreting the latent space of gans for semantic face editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 9243–9252.
- [34] Bowen Shi, Wei-Ning Hsu, Kushal Lakhotia, and Abdelrahman Mohamed. 2022. Learning Audio-Visual Speech Representation by Masked Multimodal Cluster Prediction. *arXiv preprint arXiv:2201.02184* (2022).
- [35] Bowen Shi, Wei-Ning Hsu, and Abdelrahman Mohamed. 2022. Robust Self-Supervised Audio-Visual Speech Recognition. *arXiv preprint arXiv:2201.01763* (2022).
- [36] Linsen Song, Wayne Wu, Chen Qian, Ran He, and Chen Change Loy. 2022. Everybody’s talkin’: Let me talk as you want. *IEEE Transactions on Information Forensics and Security* 17 (2022), 585–598.
- [37] Supasorn Suwajanakorn, Steven M Seitz, and Ira Kemelmacher-Shlizerman. 2017. Synthesizing obama: learning lip sync from audio. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–13.
- [38] Jintao Tan, Xize Cheng, Lingyu Xiong, Lei Zhu, Xiandong Li, Xianjia Wu, Kai Gong, Minglei Li, and Yi Cai. 2024. Landmark-guided Diffusion Model for High-fidelity and Temporally Coherent Talking Head Generation. *arXiv preprint arXiv:2408.01732* (2024).
- [39] Shuai Tan, Bin Ji, and Ye Pan. 2024. Style2Talker: High-Resolution Talking Head Generation with Emotion Style and Art Style. In *AAAI Conference on Artificial Intelligence*.
- [40] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. 2021. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–14.
- [41] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. 2019. FVD: A new Metric for Video Generation. In *ICLR Workshop on Deep Generative Models for Highly Structured Data*.
- [42] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research* 9, 86 (2008), 2579–2605.
- [43] Jiadong Wang, Xinyuan Qian, Malu Zhang, Robby T Tan, and Haizhou Li. 2023. Seeing What You Said: Talking Face Generation Guided by a Lip Reading Expert. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14653–14662.
- [44] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. 2021. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10039–10049.
- [45] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- [46] Zehan Wang, Ziang Zhang, Xize Cheng, Rongjie Huang, Luping Liu, Zhenhui Ye, Haifeng Huang, Yang Zhao, Tao Jin, Peng Gao, et al. 2024. Molecule-Space:

- Free Lunch in Unified Multimodal Space via Knowledge Fusion. *arXiv preprint arXiv:2405.04883* (2024).
- [47] Zehan Wang, Yang Zhao, Haifeng Huang, Jiageng Liu, Aoxiong Yin, Li Tang, Linjun Li, Yongqi Wang, Ziang Zhang, and Zhou Zhao. 2023. Connecting multimodal contrastive representations. *Advances in Neural Information Processing Systems* 36 (2023), 22099–22114.
- [48] Yihan Wu, Junliang Guo, Xu Tan, Chen Zhang, Bohan Li, Ruihua Song, Lei He, Sheng Zhao, Arul Menezes, and Jiang Bian. 2023. Videodubber: Machine translation with speech-aware length control for video dubbing. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 13772–13779.
- [49] Zongze Wu, Dani Lischinski, and Eli Shechtman. 2021. Stylespace analysis: Disentangled controls for stylegan image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12863–12872.
- [50] Yangyang Xu, Shengfeng He, Kwan-Yee K Wong, and Ping Luo. 2023. RIGID: Recurrent GAN Inversion and Editing of Real Face Videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 13691–13701.
- [51] Xu Yao, Alasdair Newson, Yann Gousseau, and Pierre Hellier. 2022. Feature-style encoder for style-based gan inversion. *arXiv preprint arXiv:2202.02183* (2022).
- [52] Fei Yin, Yong Zhang, Xiaodong Cun, Mingdeng Cao, Yanbo Fan, Xuan Wang, Qingyan Bai, Baoyuan Wu, Jue Wang, and Yujiu Yang. 2022. StyleHEAT: One-Shot High-Resolution Editable Talking Face Generation via Pre-trained StyleGAN. *arxiv:2203.04036* (2022).
- [53] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. 2018. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Proceedings of the European conference on computer vision*. 325–341.
- [54] Libing Zeng, Lele Chen, Yi Xu, and Nima Kalantari. 2023. MyStyle++: A Controllable Personalized Generative Prior. *arXiv preprint arXiv:2306.04865* (2023).
- [55] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *CVPR*.
- [56] Wenxuan Zhang, Xiaodong Cun, Xuan Wang, Yong Zhang, Xi Shen, Yu Guo, Ying Shan, and Fei Wang. 2023. SadTalker: Learning Realistic 3D Motion Coefficients for Stylized Audio-Driven Single Image Talking Face Animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8652–8661.
- [57] Zhimeng Zhang, Lincheng Li, Yu Ding, and Changjie Fan. 2021. Flow-guided one-shot talking face generation with a high-resolution audio-visual dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3661–3670.
- [58] Weizhi Zhong, Chaowei Fang, Yinqi Cai, Pengxu Wei, Gangming Zhao, Liang Lin, and Guanbin Li. 2023. Identity-Preserving Talking Face Generation with Landmark and Appearance Priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9729–9738.
- [59] Yang Zhou, Xintong Han, Eli Shechtman, Jose Echevarria, Evangelos Kalogerakis, and Dingzeyu Li. 2020. Makeltalk: speaker-aware talking-head animation. *ACM Transactions On Graphics (TOG)* 39, 6 (2020), 1–15.