

StyleSync: High-Fidelity Generalized and Personalized Lip Sync in Style-based Generator

Jiazhi Guan^{1,2*} Zhanwang Zhang^{1*} Hang Zhou^{1†} Tianshu Hu^{1†} Kaisiyuan Wang³
 Dongliang He¹ Haocheng Feng¹ Jingtuo Liu¹ Errui Ding¹ Ziwei Liu⁴ Jingdong Wang¹
¹Department of Computer Vision Technology (VIS), Baidu Inc. ²Tsinghua University
³The University of Sydney ⁴S-Lab, Nanyang Technological University
 guanjz20@mails.tsinghua.edu.cn {zhangzhanwang, zhouhang09, hutianshu01}@baidu.com



Figure 1. **Personalized lip-sync results generated by our StyleSync framework.** Our method not only supports high-fidelity modification to any target template video according to conditional audio but can further adapt to specific styles with personalized optimization. In this figure, our lip-sync results should have the same mouth shapes as the lip-synced video of the conditional audio.

Abstract

Despite recent advances in syncing lip movements with any audio waves, current methods still struggle to balance generation quality and the model’s generalization ability. Previous studies either require long-term data for training or produce a similar movement pattern on all subjects with low quality. In this paper, we propose StyleSync, an effective framework that enables high-fidelity lip synchronization. We identify that a style-based generator would sufficiently enable such a charming property on both one-shot

and few-shot scenarios. Specifically, we design a mask-guided spatial information encoding module that preserves the details of the given face. The mouth shapes are accurately modified by audio through modulated convolutions. Moreover, our design also enables personalized lip-sync by introducing style space and generator refinement on only limited frames. Thus the identity and talking style of a target person could be accurately preserved. Extensive experiments demonstrate the effectiveness of our method in producing high-fidelity results on a variety of scenes. Resources can be found at <https://hangz-nju-cuhk.github.io/projects/StyleSync>.

*Equal contribution.

†Corresponding authors.

1. Introduction

The problem of generating lip-synced videos according to conditional audio is of great importance to the field of digital human creation, audio dubbing, film-making, and entertainment. While the rapid development of this area has been witnessed within recent years, most methods [6,8,9,16,19,24,27,28,38,39,46,50,54,57–60] focus on generating a whole dynamic talking head. Results created under such settings can hardly be blended into an existing scene. Under real-world scenarios like audio dubbing, one crucial need is to seamlessly alter the mouth or facial area while preserving other parts of the scene unchanged, making these methods non-feasible.

Previous methods take two different paths for achieving seamless mouth modification. A number of studies [38,44] pursue realistic results on person-specific settings, which require long-term clips for target modeling. Moreover, they rely on prior 3D facial structural information. The uncertainty and errors accumulated in the 3D fitting procedure would greatly influence their performances. On the other hand, it is desired to build models that break the data limitation on more generalized scenes. As a result, a few methods [31,32,41] design person-agnostic models without relying on 3D or 2D structural priors. Nevertheless, such a setting is extremely challenging.

In order to produce high-fidelity lip-synced results on any-length videos, two essential challenges need to be addressed. **1)** How to efficiently design a powerful generative backbone network that supports both accurate audio information expression and seamless local alternation. Intuitively, the lip-sync quality naturally contradicts the preservation of the original target frame information [32,59]. **2)** How to effectively leverage as much provided information as possible and involve the personalized properties to a generalized model. Though few-shot meta-learning has been proven effective in generating talking heads [5,7,56,60], how to involve such ability into a lip-syncing pipeline has not been explored.

In this paper, we propose a highly concise and comprehensive framework named **StyleSync**, which produces high-fidelity lip-sync results on both generalized and personalized scenarios. The key is our *simple but lip-sync-oriented modifications to style-based generator*. Though style-based generators [22,23] have been leveraged in various talking head generation methods [5,55,59], their successes are only partially instructive. They aim at producing the whole head, which leads to unstable background and distortions which are non-acceptable in our scenarios.

By revisiting the details of style-based generators, we identify a few simple but essential modifications that make our framework suitable for lip-syncing. Different from the above methods, we adopt a masked mouth modeling protocol [32,41] and delicately design a *Mask-based Spa-*

tial Information Encoding strategy, where both the target and reference frames’ information is encoded into a noise space [52,53] of the generator according to different masking schemes. While the information on audio dynamics and high-level reference frame is injected into the style-modulated convolution in a similar manner as [25,59]. In this way, our method can be benefited from the strong generative power of style-based generators and also keeps the advantage of easy implementation and fast training.

Moreover, our network modification enables personalized information preserving (e.g., speaking styles and details of the mouth and jaw). We take inspiration from the recent studies of inverting StyleGAN priors [1,2,35,45] and propose a *Personalized Optimization* scheme. As audio dubbing is normally performed on speaking videos, our model can make use of only a few seconds of the person’s information and optimize additional person-specific parameters including the W^+ and the generator. Extensive experiments show that our framework clearly outperforms previous state of the arts on the one-shot setting by a large margin, and the target-specific optimization further enhances the fidelity of our results.

Our contributions can be summarized as follows: **1)** We present the **StyleSync** framework, which adopts simple but effective modifications including the *Mask-based Spatial Information Encoding* to a style-based generator. **2)** We propose the *Personalized Optimization* procedure which involves few-shot person-specific optimization into our framework. **3)** Extensive experiments demonstrate that our framework can directly produce accurate and high-fidelity one-shot lip-sync results. Moreover, our proposed personalized optimization further improves the generation quality. Our method outperforms previous methods by a clear margin.

2. Related Work

2.1. Audio-Driven Facial Animation

The topic of audio-driven facial animation has long been a research interest in both the computer vision and graphics community. Studies have been carried out on both digital 3D human faces [15,20,33,61] and realistic human heads [14,17,42]. We focus on human heads in the real world.

Talking Head Generation. Most studies on lip-syncing target generating the whole head of a talking person [6,8,9,16–18,25,26,40,42,54,58–60]. Specifically, a number of studies leverage structural information such as 2D [9] landmarks, 3D landmarks [60] and 3D meshes [6]. The uncertainty and inaccuracy of such representations would lead to error accumulation during the talking head generation procedure. Most person-specific methods [19,27,42] rely on them and produce results with a poor generalization or lip-

sync quality. Recently, NeRF is also used in person-specific modeling [16, 26, 43], but they also perform poorly when driven by audio from a different person.

Lip-Syncing on Faces. The other type of study focuses on lip-syncing the mouth part while keeping other information untouched in videos [31, 32, 39, 44]. Our work lies within the same scope. While Thies et al. [44] and Song et al. [38] produce realistic results, they rely on person-specific training on more than 2 minutes of data.

Specifically, Wav2Lip [32] generates person-agnostic lip-sync results that are highly accurate. However, they build the generative model on low-resolution images, leading to blurry results. We identify that it would be easier for the model to learn the mouth motions if less information about the image quality needs to be recovered. Moreover, their methods cannot capture the personalized patterns given the target template video.

2.2. Style-based Generator for Faces

StyleGAN models [21–23] have shown great success on image generation tasks, particularly on facial image generation and editing [1, 2, 34, 36, 37, 45].

StyleGAN Inversion. StyleGAN inversion is the practice of recovering the latent space of a given image with a pre-trained StyleGAN. Abdal *et al.* [1, 2] propose to invert images not only using the style space W , but also expand it to the W^+ , which better preserves details. Recent studies propose to learn encoders [3, 34, 45] for better editability and faster optimization results. Furthermore, recent studies try to refine the generator’s parameters through pivotal finetuning [35]. In our work, we take inspiration of previous StyleGAN inversion studies and propose a personalized optimization procedure. We encode a W^+ space from reference frames following [45] which enables personalized lip movements learning.

Applications. StyleGAN architecture has also been leveraged in face restorations [48, 53] and face swapping [51, 52]. These tasks require preserving the original facial emotion and expressions. On the other hand, style-based generators have also been applied to create facial animations [5, 25, 59], which are highly related to our task. However, they rely on the style vectors in the W or W^+ space for controlling both the appearances and motion dynamics. One major drawback of this setting is that W^+ space cannot easily account for the spatial consistency of backgrounds, leading to non-realistic results or visible artifacts.

3. Methodology

In this section, we introduce the details of our StyleSync framework as depicted in Fig. 2. It is simply designed by leveraging several successes of previous style-based generators. We introduce our modifications to the style-based generator that make a successful backbone in Sec. 3.1

and the training objectives of our generalized model in Sec. 3.2. Specifically, we illustrate our *Personalized Optimization* procedure which further strengthens our results with person-specific optimization in Sec. 3.3.

Task Formulation. Our goal is to sync the lip motion of a target person with any provided audio and seamlessly blend it into the original target video. We formulate our training in a typical self-reconstruction manner. The training setting is similar to Wav2Lip [32]. Given a target video $\mathbf{V} = \{I_1, \dots, I_T\}$ with T frames, we mask out the lower half of the face including the jaw and cheeks with a mask M . The goal is to recover the whole face with its corresponding audio $\mathbf{a} = \{a_1, \dots, a_T\}$ (\mathbf{a} is processed to spectrograms). As no information about the mouth and jaw shape is provided on the masked target frame $I_t^m = (\mathbf{1} - M) * I_t$, we leverage a random reference frame $I_{ref} \in \mathbf{V}$ during training to provide the desired context.

3.1. Modifying Style-based Generator for Lip Sync

Style-based Generator Overview. We perform lip-sync-oriented modifications to build a successful backbone for lip sync with the StyleGAN2 [23] architecture. The original StyleGAN2 takes a style vector w as input and uses it to modulate convolution operations in a total of L generative layers during training. During inference, different w s can be used at different layers, formulating the W^+ space $\mathbf{W} = \{w_1, \dots, w_L\}$. The basic elements of the original StyleGAN2, including the W^+ space and the StyleGAN generator G ’s layers are depicted in **blue** on Fig. 2. For simplification, certain detailed structures are omitted.

Mask-based Spatial Information Encoding. Our goal is to seamlessly blend a lip-synced mouth into the target frame with the assistance of a reference image. As discussed in Sec. 2.2, previous methods leveraging the style-based generator are not directly applicable to our scene, we seek a different way to encode the spatial information on faces.

It is proven by recent studies on face restoration [53] and swapping [52] that content-guided feature maps can be attached to StyleGAN layers in a similar status as *noise* without affecting the expressive power of the generative model. These contextual-rich *noises* $\mathbf{N} = \{N_1, \dots, N_L\}$ could well preserve the original spatial structure and attributes of the encoded visual information. This property provides an option for leveraging the attributes information lying in the reference image. Thus following Wav2Lip [32], we concatenate the masked and reference input together as the visual input of the framework and encode the corresponding feature maps $\mathbf{F}_f = \{F_1, \dots, F_L\}$. The simplest setting is to regard $\mathbf{N} = \mathbf{F}$.

However, the goal of generating lip motions needs to create dynamics that are different from the original facial structure but keeps facial identity, which is substantially different from previous methods that aim to keep the structure

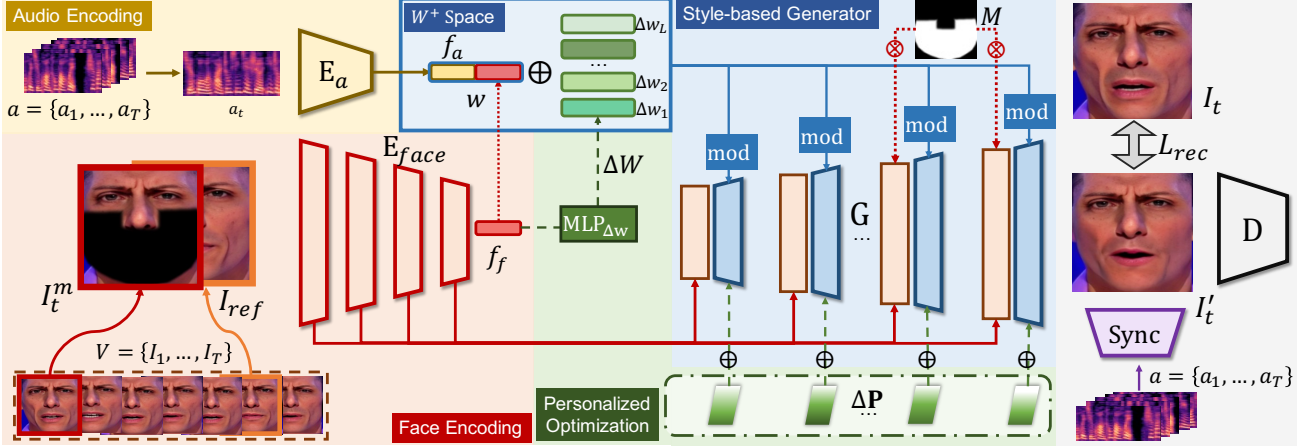


Figure 2. **Our StyleSync framework.** The building blocks in **Blue** indicate the style-based generator. The masked target frame I_t^m is concatenated with a reference I_{ref} and encoded to f_f by E_{face} (**Red**). The audio information is encoded to an audio feature f_a (**Yellow**). The features are concatenated together to form the style (W) space. Specifically, we devise a Personalized Optimization procedure (**Green**) including the learning of the ΔW and the ΔP . This part is not trained during the initial backbone learning.

fixed. The mouth shapes would be greatly influenced by the reference image I_{ref} when adopting the same protocol. We identify that more than desired facial structure information is infused in the low-level information of \mathbf{F} (layers with higher resolutions). As a result, we intuitively mask out the low-level part of the information in the face features. Here we simply define

$$F'_l = (\mathbf{1} - M) * F_l, \text{ for } l > \lfloor \frac{L}{2} \rfloor. \quad (1)$$

The facial information is colored in **red** on Fig. 2. The $\mathbf{1}$ is an all-one matrix that has the same size as M .

Style Information Encoding. We follow previous studies [5, 40, 59] to encode both audio dynamics and facial information into the style space (W space) of the style-based generator. The audio feature f_a is encoded from the encoder E_a and the face feature f_f is the bottleneck of the face encoder E_{face} , $w = \text{concat}(f_a, f_f)$. As the information in the reference frames is already fused to the generator through spatial noise encoding, the fusing of the face feature is less necessary. Experiments show that the difference is subtle with or without f_f , but we still keep this design in accordance with common practice.

Ingredients for Personalization. Our above designs are not only simple but also possess external potentials of personalized lip-sync modeling. We leave the details to Sec. 3.3 and briefly discuss the additional ingredients that make our model comprehensive.

Specifically, our extension is inspired by the advances in StyleGAN inversion [1, 2, 35]. It has pointed out that extending w to the W^+ space which contains $\mathbf{W} = \{w + \Delta w_1, \dots, w + \Delta w_L\}$ is essential to recovering specific images. Moreover, recent studies [4, 35] also explore tuning

the parameters of the generator to improve the network’s fitting ability on a specific target. Thus our generator has the potential of learning limited parameter shifts ΔP in order to fit a specific person.

These personalized ingredients are depicted in **green** on Fig. 2 and are not optimized during the generalized training procedure. In order to keep the pipeline consistent across both settings, we set $\Delta W = \mathbf{0}$ and $\Delta P = \mathbf{0}$ in our generalized model.

3.2. Backbone Training Objectives

During training, the whole backbone networks take the 6-channel concatenation of $[I_t^m, I_{ref}]$ and the audio clip a_t at the same time step as input. It predicts $I_t' = G(\mathbf{N}, \mathbf{W})$, which aims at recovering the unmasked target image I_t . In order to keep our design simple, the training objectives are mostly aligned with StyleGAN2 [23] and Wav2Lip [32].

Reconstruction Loss. The pixel-level reconstruction loss \mathcal{L}_{rec} is fundamental for training our task. Here we leverage the commonly used L_1 loss and perceptual loss [47].

$$\mathcal{L}_{rec} = \|I_t' - I_t\|_1 + \sum_{m=1}^{N_{vgg}} \|\text{VGG}_m(I_t') - \text{VGG}_m(I_t)\|_1, \quad (2)$$

where VGG_m is the m th layer’s output of a pre-trained VGG19 network.

Adversarial Loss. We directly adopt the same discriminator D from StyleGAN2 [23] for the adversarial training:

$$\mathcal{L}_{adv} = \min_G \max_D (\mathbb{E}_{I_t} [\log D(I_t)] + \mathbb{E}_{I_t'} [\log(1 - D(I_t'))]). \quad (3)$$

Particularly, we initialize the discriminator’s weights with the pre-trained version in StyleGAN2 [23]. This practice produces visibly sharper results.

Lip-Sync Loss. We additionally train a SyncNet [12] which consists of a visual encoder S_v and an audio encoder S_a . It identifies whether visual and audio clips are timely aligned with contrastive loss. When training the generator, we predict 5 consecutive frames $I'_{t:t+4}$ within one batch at each inference step and supervise the training with SyncNet’s assistance. The objective is:

$$\mathcal{L}_{sync} = -\frac{S_v(I'_{t:t+4})^T \cdot S_a(a_{t:t+4})}{\|S_v(I'_{t:t+4})\|_2 \|S_a(a_{t:t+4})\|_2} \quad (4)$$

The overall loss functions across the generalized training can be written as:

$$\mathcal{L}_g = \mathcal{L}_{adv} + \lambda_r \mathcal{L}_{rec} + \lambda_s \mathcal{L}_{sync}. \quad (5)$$

3.3. Personalized Optimization

After the training of the backbone networks, our method can readily generate high-fidelity lip sync results for arbitrary subject. However, as the model is trained in a generalized manner, the generated mouth motion patterns across different people are basically the same. It has been verified that different identities possess different talking styles [50, 54]. Here we pursue to capture such personalized property with the original template video. Notably, we focus on the few-shot setting that only less than one minute of the original video is given, which cannot be handled by previous person-specific models [16, 19, 38, 44]. Below, we illustrate how we successfully design our personalized optimization module based on the above discussions.

Basic Learning Settings. The basic learning procedure of the personalized optimization is similar to backbone training. On a template video $\bar{\mathbf{V}} = \{\bar{I}_1, \dots, \bar{I}_T\}$ with audio $\bar{\mathbf{a}} = \{\bar{a}_1, \dots, \bar{a}_T\}$, we randomly select two frames as target \bar{I}_t and reference \bar{I}_{ref} for training as one sample. During training, the generalized loss \mathcal{L}_g continuously supervises the whole procedure.

As the encoder E_a encodes person-agnostic speech content information, it is fixed during the personalized optimization step. While the face encoder E_{face} encodes the attribute information of the target person, we do not optimize its parameters on a single person to avoid overfitting.

Style Space Optimization. According to our formulation, the W^+ space contains both the lip motion and the high-level facial style information. Thus optimizing it would ideally lead to person-specific lip motion pattern. E4E [45] verifies that learning the small displacements around the original W space would both extends the inversion quality of the image and enables strong editing ability. Thus we propose

to learn a set of $\Delta \mathbf{W} = MLP_{\Delta w}(f_f)$ in a similar way according to the encoded face feature f_f with a few layers of MLPs $MLP_{\Delta w}$.

Generator Tuning. The generator most accounts for the synthesizing and the blending quality. We allow the parameters \mathbf{P} of our generator G to shift a little margin to $\Delta \mathbf{P}$ with the personalized data we use.

Learning Objectives. The final learning objectives of the personalized optimization procedure can be summarized as:

$$\mathcal{L}_p = \mathcal{L}_g + \lambda_p \left(\sum_{i,j} |\Delta \mathbf{W}_i^j|_2^2 + \sum_{m,n} |\Delta \mathbf{P}_n^m|_2^2 \right), \quad (6)$$

where \mathbf{W}_i^j denote the j th parameter of the i th element in \mathbf{W} . This stands the same for \mathbf{P}_n^m . We restrict the displacements in a limited step.

4. Experiments

Datasets. We conduct experiments on two commonly used audio-visual datasets, LRW [11] and VoxCeleb2 [10]. We follow the datasets’ original train/test split and train our model on a mixture of these two datasets.

- **LRW [11]** is an audio-visual dataset collected from BBC news for lip reading. It is one of the earliest large-scale audio-visual datasets with high quality. It consists of 1,000 one-second utterances in 500 words.
- **VoxCeleb2 [10]** VoxCeleb [30] and VoxCeleb2 are large-scale audio-visual datasets collected for the speaker verification task. We re-download around one-fifth of VoxCeleb2 for training, which are of high quality. The official test set is processed with GPEN [53] for evaluation.

Implementation Details. We process all videos at 25 fps and align all faces according to pre-detected landmarks at the eyes. All faces are cropped to the size of 256×256 . A same U-shape mask is adopted as shown in Fig. 2 to erase the mouth, cheeks and jaws at the target frame. All audios are processed in the same way as Wav2Lip [32]. In order to keep our design simple and avoid extra hyperparameter tuning, we adopt most settings from the previous studies [23, 32, 53]. Specifically, the generator has a total of $L = 14$ style-convolution layers. λ_r is empirically set to 10 and all other λ s are selected to 1. For personalized optimization, we train the person’s data for 5 epochs. Long training would not lead to better results. Particularly, we find that training the personalized model together with the general training data leads to more stable results, the portion of personal data and general data is set as 1:1 in our experiments.

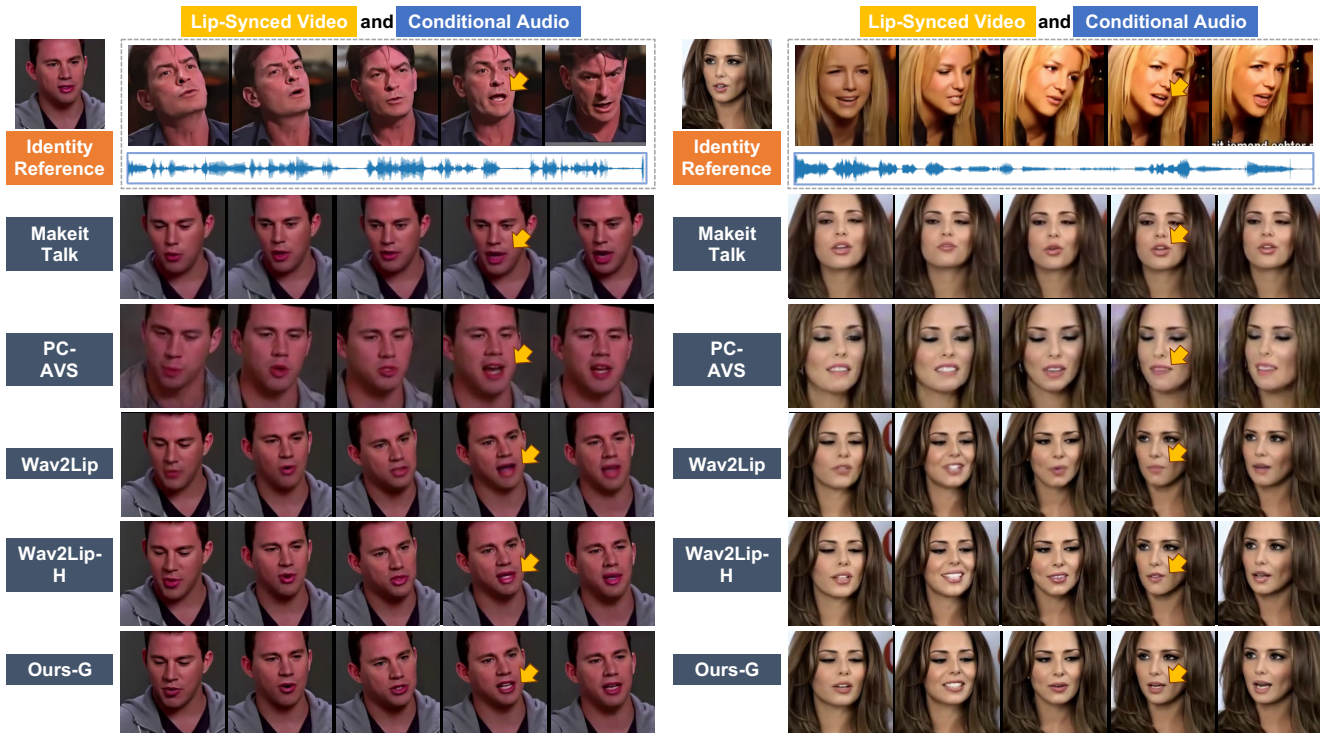


Figure 3. **Qualitative Results.** The top row shows the lip-synced videos of the conditional driving audio. MakeitTalk [60] fails to generate accurate mouth shape and lacks head dynamics. PC-AVS [59] and Wav2Lip [32] generate lip-synced results aligned well with audios. But they produce visibly blurry results. Particularly, our reproduced Wav2Lip-H generates high-quality results with good lip-sync. However, they still produce certain artifacts. While our generalized model directly produces high-fidelity results that have the same lip motion as the original synced video.

Comparing Methods. As our model is built based on the person-agnostic setting, we compare our StyleSync framework with three state-of-the-art methods including, MakeitTalk [60], PC-AVS [59] and **Wav2Lip** [32]. **MakeitTalk** [60] synthesizes talking head videos with natural head pose. **PC-AVS** [59] a pose source video to achieve pose-controllable talking face generation. While MakeitTalk and PC-AVS produce a whole talking head, Wav2Lip which focuses on the mouth part is our main competing method. As Wav2Lip is originally trained on a low-resolution with small networks, we carefully reproduce it on the same data and even loss functions as our methods. We denote this model as **Wav2Lip-H** (High-Quality). For fair comparison, all basic comparisons are carried out on the generalized setting **Ours-G**.

We also perform personalized optimization on the test set, where each frame lasts less than 10 seconds for both datasets. The results are denoted as **Ours-P**. We equally perform the generator finetuning on the above methods, and all of them perform worse than their original models, thus the results are not listed below. Additionally, we compare our personalized setting with ADNeRF [16], a recent talking head advance using neural radiance fields [29].

4.1. Quantitative Evaluation

As quantitative evaluations can only be performed on the self-construction setting, we avoid directly leveraging all information within the frame for reconstruction. We uniformly select a same random frame as the reference frame for PC-AVS, Wav2Lip, Wav2Lip-H and ours.

Evaluation Metrics. We follow previous studies [8, 9, 58, 59] to adopt the popularly used SSIM [49], PSNR metrics for evaluating the generation quality, and the landmark distances around the mouth (LMD) and SyncNet’s [12] confidence score to show the lip-sync quality. Please be noted that the SyncNet metric is computed on the officially released version, which is different from our implementation in the loss functions.

Particularly, as the details of the target person cannot be fully recovered, we leverage the ArcFace [13] network and compute the frame-by-frame feature cosine distances \mathcal{D}_{ID} for evaluating whether the generated results preserve the identity well.

Evaluation Results. The quantitative experiments are carried out on the test set of LRW [11] and VoxCeleb2 [10] datasets. Please refer to Table 1 for the results. It can be seen that on the LRW dataset, the generalized version

Method	LRW					VoxCeleb2				
	SSIM \uparrow	PSNR \uparrow	LMD \downarrow	Sync _{conf} \uparrow	\mathcal{D}_{ID} \uparrow	SSIM \uparrow	PSNR \uparrow	LMD \downarrow	Sync _{conf} \uparrow	\mathcal{D}_{ID} \uparrow
MakeitTalk	0.69	29.83	2.75	3.88	0.79	0.63	28.38	6.94	2.15	0.71
PC-AVS	0.79	30.26	1.84	7.19	0.82	0.71	29.53	2.75	8.16	0.74
Wav2Lip	0.79	30.54	1.28	7.39	0.90	0.80	30.53	1.92	8.90	0.90
Wav2Lip-H	0.80	31.38	1.20	7.19	0.88	0.81	30.53	1.87	8.35	0.90
GT	1.00	100.00	0.00	7.65	1.00	1.00	100.00	0.00	7.71	1.00
Ours-G	0.85	31.78	1.18	7.25	0.89	0.79	31.00	1.47	8.25	0.90
Ours-P	0.88	32.66	0.86	6.35	0.93	0.82	31.54	1.15	7.26	0.93

Table 1. **Quantitative results on LRW and VoxCeleb.** For LMD the lower the better, and the higher the better for other metrics.

of our StyleSync already outperforms previous methods by a large margin for generation quality. While on the VoxCeleb2 dataset, the gap is less obvious. The reason might be that most training data of our methods are basically frontal view faces selected from LRW, while VoxCeleb2 contains more complicated scenes. On the other hand, after our personalized optimization, the performance of our model advances again. This shows the generated results are clearly more similar to the targets at this stage, which can also be verified by looking into the identity distances.

Meanwhile, our method also achieves comparable performance on the lip-sync metrics on both datasets. Our LMD score is slightly better than the competing methods. As for the SyncNet score, we achieve comparable results on LRW that are closer to the ground truth’s SyncNet score. We argue it is meaningless to refer to the SyncNet scores for verifying the lip-sync quality once the metric has outperformed the ground truth. The $Sync_{conf}$ only reflects how well an audio-visual pair fits the learned SyncNet model rather than the true perceptual quality. Thus though generated results might outperform ground truth on the metric, it does not mean better sync quality.

After the personalization, the lip-sync score degrades. We assume that the specificity of the speaking style reduces the mouth opening level. This is also shown in Fig. 5. The woman tends not to open her mouth wide, leading to smaller mouth movements

4.2. Qualitative Evaluation

Subjective evaluation is crucial in identifying the ability of generative models, particularly on videos. We strongly recommend readers to watch our **supplementary video** at <https://hangz-nju-cuhk.github.io/projects/StyleSync>.

Two cross-driven examples and their comparisons with the SOTA methods are shown in Fig. 3. We use an conditional audio from an arbitrary person selected from the test set to drive the template. It can be seen that MakeitTalk [60] cannot produce accurate lip movements. Moreover, both PC-AVS and Wav2Lip produce visible artifacts or blurry



Figure 4. Self-driven results compared with personalized method.

results. Our reproduced Wav2Lip-H adopts the same training setting as ours, thus generates plausible results for most cases. Nevertheless, our generalized model still produces the most accurate mouth shapes with the highest fidelity.

Additionally, we show a case of personalized optimization on a 50-second video clip in Fig. 1. It can be seen that our method preserves the speaking style of the target person with accurate lip sync. Comparing our personalized results (Ours-P) with AD-NeRF [16] (Fig. 4), it is evident that our method maintains more person-specific details with higher fidelity, even though we use significantly less personal data (10s for tuning v.s. 5min for training in [16]).

User Study. We also invite 15 participants to conduct a user study for further subjective evaluation and the results are reported in Table 2. Specifically, we randomly select 63 videos from LRW and VoxCeleb2 datasets as test sets and generate videos by using our generalized StyleSync and the comparison methods accordingly. By adopting the commonly used Mean Opinion Scores (MOS) rating protocol, we request all the participants to provide their ratings (from 1 to 5, the higher the better) on three aspects for each generated video: (1) Lip-Sync quality; (2) Generation quality; (3) Video realism.

As shown in Table 2, MakeitTalk [60] achieves the lowest scores in all the aspects due to lacking reasonable lip movements and head pose. PC-AVS [59], Wav2Lip [32] and Wav2Lip-H achieve relatively higher lip-sync score, but only Wav2Lip-H is able to synthesize videos with less blurry textures around the mouth. Overall, our StyleSync

MOS on \ Approach	MakerTalk	PC-AVS	Wav2Lip	Wav2Lip-H	Ours-G
Lip-Sync Quality	2.06	3.00	3.49	3.67	4.24
Generation Quality	2.63	2.16	1.87	3.42	4.52
Video Realness	1.89	2.16	2.17	2.98	4.06

Table 2. User study measured by Mean Opinion Scores. The scores are ranged from 1 (worst) to 5 (best).

Method	SSIM \uparrow	PSNR \uparrow	LMD \downarrow	Sync _{conf} \uparrow	\mathcal{L}_{ID} \uparrow
w/o mask	0.81	31.20	1.80	7.89	0.90
w/o sync	0.80	31.34	1.55	7.41	0.90
Ours-G	0.79	31.00	1.47	8.25	0.90
P w/o $\Delta\mathbf{W}$	0.82	31.51	1.32	7.31	0.92
Ours-P	0.82	31.54	1.15	7.26	0.93

Table 3. Quantitative ablations results on VoxCeleb2.

outperforms its counterparts in all the three aspects by a large margin, indicating the effectiveness of our approach.

4.3. Ablation Study

To further demonstrate the contributions of our novel designs, we perform an ablation study on VoxCeleb2 dataset under both generalized setting and personalized optimization setting (denoted as “Ours-G” and “Ours-P”, respectively). Concretely, we construct two variants for “Ours-G” (denoted as “w/o mask” and “w/o sync”) by removing the masking operation in Eq. 1 and the lip-sync loss during training, respectively. While for “Ours-P”, we form one variant (denoted as “Ours-P w/o $\Delta\mathbf{W}$ ”) by setting $\Delta\mathbf{W} = \mathbf{0}$. We have also experimented personalized optimization with fixed generated, however, this would lead to blurry results. This ablation is omitted here.

The quantitative and qualitative results are shown in Table 3 and Fig. 5. As shown in Table 3, we observe that both “w/o mask” and “w/o sync” achieve comparable performance on the image quality metrics but leads to a performance drop on the lip-sync metrics. “w/o mask” achieves worse scores when compared with “Ours-G” in terms of LMD and Sync, which indicates that the proposed masking strategy can obviously alleviate the influence from the reference frames. While “w/o sync” reasonably suffers from lip-sync degradation and the results in Fig. 5 illustrate inconsistent lip movements when compared with the lip-synced video. The results demonstrate the effectiveness of the additional supervision from SyncNet.

It can be seen from Fig. 5 that with the personalized optimization, both the identity and the pattern of the mouth opening on the generated frames become more similar to the original template video. It is also clear that after the personalized training, the mouth opening is less obvious than in the generalized model. We analyze that this is part of the

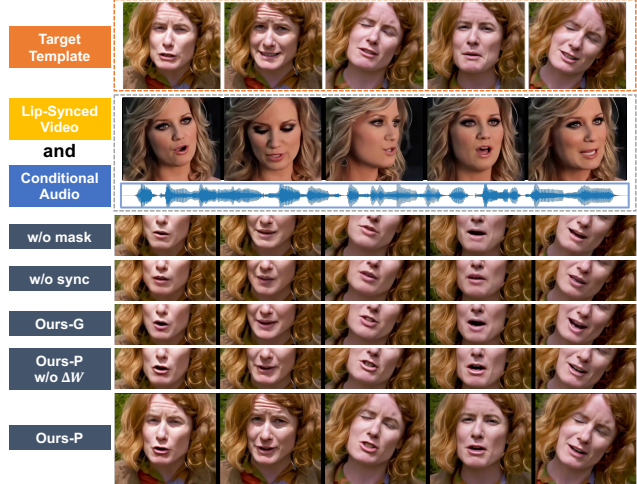


Figure 5. Ablation study with visual results. Zoom in for details.

talking style of this target person. This also leads to poorer metric values on SyncNet.

In terms of the comparisons between our personalized model and “Ours-P w/o $\Delta\mathbf{W}$ ”, it achieves similar scores on the image quality as well as lip-sync quality. However, the identity score without $\Delta\mathbf{W}$ is slightly degraded, which indicates that it is beneficial to involve $\Delta\mathbf{W}$ into the training procedure for better preserving identity.

5. Conclusion and Discussions

Conclusion. In this paper, we propose **StyleSync**, which produces high-fidelity lip sync results for both the one-shot and the few-shot settings. We highlight the unique properties of our method: **1)** Video results produced by our generalized model clearly outperform previous state-of-the-arts. **2)** Our model is built upon the success of recent style-based generators with simple modification. It is easy to implement and friendly to train. **3)** By involving StyleGAN-inspired personalized optimization procedure, our model can further be improved on a specific person given only a few clips.

Limitations. As our method blends a lip-synced face into an existing video with a fixed mask, the head pose and expressions of the target person cannot be changed. Additionally, under certain extreme cases where the target’s jaw is extremely large, it might get out of our masked area.

Ethical Considerations. Our method could be used to create non-existing talks and speeches, which might be maliciously used. We will issue our core models strictly to research institutions.

Acknowledgement. This study is supported by the Ministry of Education, Singapore, under its MOE AcRF Tier 2 (MOE-T2EP20221-0012), NTU NAP, and under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4432–4441, 2019. 2, 3, 4
- [2] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan++: How to edit the embedded images? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8296–8305, 2020. 2, 3, 4
- [3] Yuval Alaluf, Or Patashnik, and Daniel Cohen-Or. Restyle: A residual-based stylegan encoder via iterative refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6711–6720, 2021. 3
- [4] Yuval Alaluf, Omer Tov, Ron Mokady, Rinon Gal, and Amit Bermano. Hyperstyle: Stylegan inversion with hypernetworks for real image editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18511–18521, 2022. 4
- [5] Egor Burkov, Igor Pasechnik, Artur Grigorev, and Victor Lempitsky. Neural head reenactment with latent pose descriptors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13786–13795, 2020. 2, 3, 4
- [6] Lele Chen, Guofeng Cui, Ziyi Kou, Haitian Zheng, and Chenliang Xu. What comprises a good talking-head video generation?: A survey and benchmark. *arXiv preprint arXiv:2005.03201*, 2020. 2
- [7] Lele Chen, Guofeng Cui, Celong Liu, Zhong Li, Ziyi Kou, Yi Xu, and Chenliang Xu. Talking-head generation with rhythmic head motion. In *European Conference on Computer Vision*, pages 35–51. Springer, 2020. 2
- [8] Lele Chen, Zhiheng Li, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. Lip movements generation at a glance. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 520–535, 2018. 2, 6
- [9] Lele Chen, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. Hierarchical cross-modal talking face generation with dynamic pixel-wise loss. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7832–7841, 2019. 2, 6
- [10] J. S. Chung, A. Nagrani, and A. Zisserman. Voxceleb2: Deep speaker recognition. In *INTERSPEECH*, 2018. 5, 6
- [11] Joon Son Chung and Andrew Zisserman. Lip reading in the wild. In *Asian conference on computer vision*, pages 87–103. Springer, 2016. 5, 6
- [12] Joon Son Chung and Andrew Zisserman. Out of time: automated lip sync in the wild. In *ACCV*, 2016. 5, 6
- [13] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 6
- [14] Bo Fan, Lijuan Wang, Frank K Soong, and Lei Xie. Photo-real talking head with deep bidirectional lstm. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4884–4888. IEEE, 2015. 2
- [15] Yingruo Fan, Zhaojiang Lin, Jun Saito, Wenping Wang, and Taku Komura. Faceformer: Speech-driven 3d facial animation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18770–18780, 2022. 2
- [16] Yudong Guo, Keyu Chen, Sen Liang, Yongjin Liu, Hujun Bao, and Juyong Zhang. Ad-nerf: Audio driven neural radiance fields for talking head synthesis. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2, 3, 5, 6, 7
- [17] Amir Jamaludin, Joon Son Chung, and Andrew Zisserman. You said that?: Synthesising talking faces from audio. *International Journal of Computer Vision*, 127(11):1767–1779, 2019. 2
- [18] Xinya Ji, Hang Zhou, Kaisiyuan Wang, Qianyi Wu, Wayne Wu, Feng Xu, and Xun Cao. Eamm: One-shot emotional talking face via audio-based emotion-aware motion model. *SIGGRAPH*, 2022. 2
- [19] Xinya Ji, Hang Zhou, Kaisiyuan Wang, Wayne Wu, Chen Change Loy, Xun Cao, and Feng Xu. Audio-driven emotional video portraits. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14080–14089, 2021. 2, 5
- [20] Tero Karras, Timo Aila, Samuli Laine, Antti Herva, and Jaakko Lehtinen. Audio-driven facial animation by joint end-to-end learning of pose and emotion. *ACM Transactions on Graphics (TOG)*, 36(4):1–12, 2017. 2
- [21] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *Proc. NeurIPS*, 2021. 3
- [22] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. 2, 3
- [23] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020. 2, 3, 4, 5
- [24] Avisek Lahiri, Vivek Kwatra, Christian Frueh, John Lewis, and Chris Bregler. Lipsync3d: Data-efficient learning of personalized 3d talking faces from video using pose and lighting normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2755–2764, 2021. 2
- [25] Borong Liang, Yan Pan, Zhizhi Guo, Hang Zhou, Zhibin Hong, Xiaoguang Han, Junyu Han, Jingtuo Liu, Errui Ding, and Jingdong Wang. Expressive talking head generation with granular audio-visual control. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3387–3396, 2022. 2, 3
- [26] Xian Liu, Yinghao Xu, Qianyi Wu, Hang Zhou, Wayne Wu, and Bolei Zhou. Semantic-aware implicit neural audio-driven video portrait generation. *ECCV*, 2022. 2, 3

- [27] Yuanxun Lu, Jinxiang Chai, and Xun Cao. Live speech portraits: real-time photorealistic talking-head animation. *ACM Transactions on Graphics (TOG)*, 40(6):1–17, 2021. 2
- [28] Yifeng Ma, Suzhen Wang, Zhipeng Hu, Changjie Fan, Tangjie Lv, Yu Ding, Zhidong Deng, and Xin Yu. Styletalk: One-shot talking head generation with controllable speaking styles. *AAAI*, 2023. 2
- [29] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pages 405–421. Springer, 2020. 6
- [30] A. Nagrani, J. S. Chung, and A. Zisserman. Voxceleb: a large-scale speaker identification dataset. In *INTER-SPEECH*, 2017. 5
- [31] Se Jin Park, Minsu Kim, Joanna Hong, Jeongsoo Choi, and Yong Man Ro. Synctalkface: Talking face generation with precise lip-syncing via audio-lip memory. In *AAAI Conference on Artificial Intelligence*. Association for the Advancement of Artificial Intelligence, 2022. 2, 3
- [32] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar. 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*, pages 484–492, 2020. 2, 3, 4, 5, 6, 7
- [33] Alexander Richard, Michael Zollhöfer, Yandong Wen, Fernando de la Torre, and Yaser Sheikh. Meshtalk: 3d face animation from speech using cross-modality disentanglement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [34] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2287–2296, 2021. 3
- [35] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Trans. Graph.*, 2021. 2, 3, 4
- [36] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. Interfacegan: Interpreting the disentangled face representation learned by gans. *TPAMI*, 2020. 3
- [37] Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in gans. In *CVPR*, 2021. 3
- [38] Linsen Song, Wayne Wu, Chen Qian, Ran He, and Chen Change Loy. Everybody’s talkin’: Let me talk as you want. *arXiv preprint arXiv:2001.05201*, 2020. 2, 3, 5
- [39] Yang Song, Jingwen Zhu, Dawei Li, Xiaolong Wang, and Hairong Qi. Talking face generation by conditional recurrent adversarial network. *arXiv preprint arXiv:1804.04786*, 2018. 2, 3
- [40] Yasheng Sun, Hang Zhou, Ziwei Liu, and Hideki Koike. Speech2talking-face: Inferring and driving a face with synchronized audio-visual representation. In *IJCAI*, 2021. 2, 4
- [41] Yasheng Sun, Hang Zhou, Kaisiyuan Wang, Qianyi Wu, Zhibin Hong, Jingtuo Liu, Errui Ding, Jingdong Wang, Ziwei Liu, and Koike Hideki. Masked lip-sync prediction by audio-visual contextual exploitation in transformers. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022. 2
- [42] Supasorn Suwajanakorn, Steven M Seitz, and Ira Kemelmacher-Shlizerman. Synthesizing obama: learning lip sync from audio. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017. 2
- [43] Jiayang Tang, Kaisiyuan Wang, Hang Zhou, Xiaokang Chen, Dongliang He, Tianshu Hu, Jingtuo Liu, Gang Zeng, and Jingdong Wang. Real-time neural radiance talking portrait synthesis via audio-spatial decomposition. *arXiv preprint arXiv:2211.12368*, 2022. 3
- [44] Justus Thies, Mohamed Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. Neural voice puppetry: Audio-driven facial reenactment. In *European Conference on Computer Vision*, pages 716–731. Springer, 2020. 2, 3, 5
- [45] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)*, 2021. 2, 3, 5
- [46] Suzhen Wang, Lincheng Li, Yu Ding, Changjie Fan, and Xin Yu. Audio2head: Audio-driven one-shot talking-head generation with natural head motion. *arXiv preprint arXiv:2107.09293*, 2021. 2
- [47] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018. 4
- [48] Xintao Wang, Yu Li, Honglun Zhang, and Ying Shan. Towards real-world blind face restoration with generative facial prior. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [49] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [50] Haozhe Wu, Jia Jia, Haoyu Wang, Yishun Dou, Chao Duan, and Qingshan Deng. Imitating arbitrary talking style for realistic audio-driven talking face synthesis. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 1478–1486, 2021. 2, 5
- [51] Chao Xu, Jiangning Zhang, Miao Hua, Qian He, Zili Yi, and Yong Liu. Region-aware face swapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7632–7641, 2022. 3
- [52] Zhiliang Xu, Hang Zhou, Zhibin Hong, Ziwei Liu, Jiaming Liu, Zhizhi Guo, Junyu Han, Jingtuo Liu, Errui Ding, and Jingdong Wang. Styleswap: Style-based generator empowers robust face swapping. In *European Conference on Computer Vision*, pages 661–677. Springer, 2022. 2, 3
- [53] Tao Yang, Peiran Ren, Xuansong Xie, and Lei Zhang. Gan prior embedded network for blind face restoration in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 672–681, 2021. 2, 3, 5
- [54] Ran Yi, Zipeng Ye, Juyong Zhang, Hujun Bao, and Yong-Jin Liu. Audio-driven talking face video generation with

- learning-based personalized head pose. *arXiv e-prints*, pages arXiv–2002, 2020. [2](#), [5](#)
- [55] Fei Yin, Yong Zhang, Xiaodong Cun, Mingdeng Cao, Yanbo Fan, Xuan Wang, Qingyan Bai, Baoyuan Wu, Jue Wang, and Yujiu Yang. Styleheat: One-shot high-resolution editable talking face generation via pre-trained stylegan. *ECCV*, 2022. [2](#)
- [56] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-shot adversarial learning of realistic neural talking head models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9459–9468, 2019. [2](#)
- [57] Chenxu Zhang, Yifan Zhao, Yifei Huang, Ming Zeng, Saifeng Ni, Madhukar Budagavi, and Xiaohu Guo. Facial: Synthesizing dynamic talking face with implicit attribute learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3867–3876, 2021. [2](#)
- [58] Hang Zhou, Yu Liu, Ziwei Liu, Ping Luo, and Xiaogang Wang. Talking face generation by adversarially disentangled audio-visual representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9299–9306, 2019. [2](#), [6](#)
- [59] Hang Zhou, Yasheng Sun, Wayne Wu, Chen Change Loy, Xiaogang Wang, and Ziwei Liu. Pose-controllable talking face generation by implicitly modularized audio-visual representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4176–4186, 2021. [2](#), [3](#), [4](#), [6](#), [7](#)
- [60] Yang Zhou, Xintong Han, Eli Shechtman, Jose Echevarria, Evangelos Kalogerakis, and Dingzeyu Li. Makeltalk: speaker-aware talking-head animation. *ACM Transactions on Graphics (TOG)*, 39(6):1–15, 2020. [2](#), [6](#), [7](#)
- [61] Yang Zhou, Zhan Xu, Chris Landreth, Evangelos Kalogerakis, Subhansu Maji, and Karan Singh. Visemenet: Audio-driven animator-centric speech animation. *ACM Transactions on Graphics (TOG)*, 37(4):1–10, 2018. [2](#)