

StyleLipSync: Style-based Personalized Lip-sync Video Generation

Taekyung Ki^{1*} Dongchan Min^{2*}

¹AITRICS, ²Graduate School of AI, KAIST

tkki@aitrics.com, alsehdcks95@kaist.ac.kr

<https://stylelipsync.github.io>

Abstract

In this paper, we present StyleLipSync, a style-based personalized lip-sync video generative model that can generate identity-agnostic lip-synchronizing video from arbitrary audio. To generate a video of arbitrary identities, we leverage expressive lip prior from the semantically rich latent space of a pre-trained StyleGAN, where we can also design a video consistency with a linear transformation. In contrast to the previous lip-sync methods, we introduce pose-aware masking that dynamically locates the mask to improve the naturalness over frames by utilizing a 3D parametric mesh predictor frame by frame. Moreover, we propose a few-shot lip-sync adaptation method for an arbitrary person by introducing a sync regularizer that preserves lip-sync generalization while enhancing the person-specific visual information. Extensive experiments demonstrate that our model can generate accurate lip-sync videos even with the zero-shot setting and enhance characteristics of an unseen face using a few seconds of target video through the proposed adaptation method.

1. Introduction

In the past few years, advances in deep learning have altered the dynamics of video creation. Now, users can easily make and edit videos with the help of deep learning. In particular, the task of generating a talking head video has received great interest due to its various practical uses. It can be applied in many applications such as film dubbing into a different language, face-to-face live chats, and virtual avatars in games and videos. Thus, a lot of prior works [21, 34, 28, 46, 45, 27] have been studied to generate a talking head video that has accurate lip shapes according to arbitrary audio inputs.

Most of the prior works mainly focus on enhancing synchronization between lip shapes and audio input. Some of the previous methods [46, 9, 34] use intermediate structural

representations such as landmarks and 3D models. They predicted the representations from the audio input and synthesized a talking head video of a target person. However, they suffered from inaccurate lip-sync results since such representations are too sparse to produce fine-grained details in lip-syncing. Recently, another line of methods [28, 27] mapped input audio to latent space and leveraged it to construct the mouth region of the target identity. While it achieves satisfactory results in lip-syncing, it generated blurry lower faces which are visually implausible. Furthermore, most methods only consider synthesizing frame-by-frame, lacking temporal consistency at the video level.

In this paper, we propose StyleLipSync, a style-based lip-sync video generative model that can generate identity-agnostic lip-synchronizing video from the arbitrary audio input. Our model consists of the following components. First, different from a previous masking method [21, 28, 27, 6] which masks the entire lower half face, we propose Pose-aware Masking. We analyze that the previous masking method cause unpleasant artifacts and unnatural jaw moving in the generated videos. To circumvent this, we utilize a 3D face mesh predictor [11, 23] and generate lip masks with consideration of pose information and facial semantics such as jaw shape. Second, our image decoder is based on a style-based generator, namely StyleGAN [18, 19, 17]. StyleGANs have demonstrated their effectiveness in various facial generative tasks, including face editing[1, 2], face enhancement [42], and video generation [35, 24]. As a pre-trained StyleGAN already contains expressive and diverse face priors in style latent space [1], we leverage it to synthesize the high-fidelity lip region of the target person. Furthermore, thanks to the continuous and linear nature of the latent space [18, 13, 32], we linearly manipulate the style codes using the audio input to generate lip-synced video frames. Additionally, we propose Style-aware Masked Fusion to effectively adopt a skip-connection to our decoder, which helps to preserve the 2D structure of the image and improves lip fidelity. Finally, we propose a Moving-average based Latent Smoothing module that makes the latent trajectory smoother for enhancing the

*Equal contribution.

temporal consistency in the synthesized talking head video.

While our model can synthesize a talking head video of the target person, there is a slight identity gap between a generated video and the target person. The gap can be noticeable, for example, in racial faces, which are relatively scarce in the training data. One approach to addressing this issue is to fine-tune the generator on a few seconds of video of the target person to create a personalized model. Several fine-tuning methods [30, 25, 37] have already been demonstrated and widely adopted by the industry to achieve product-level quality. However, we analyze that simply fine-tuning the generator loses its ability to generalize to arbitrary audio inputs, which is critical for generating talking head videos. Therefore, to minimize the side effect, we propose a sync regularizer enforcing the audio generalization performance. The key idea is to leverage the audio from the training data, not from the target video. Specifically, we not only optimize the generator to reconstruct the target video but also synthesize the video corresponding to the randomly sampled audio from the training data and maintain a sync correlation between the synthesized video and the audio. As a result, we obtain a personalized lip-sync generative model that can synthesize a video of the target person for arbitrary audio. Our contributions are summarized as follows:

- We present StyleLipSync, a lip-sync video generative model which generates lip-synchronizing video in-the-wild of 256×256 resolution with accurate and natural lip movement from a given masked video frames, audio segment, and single reference image.
- We additionally propose a few-shot adaptation method for entirely unseen faces, which uses only a few seconds of video by introducing a sync regularizer to maintain audio generalization.
- Experimental results demonstrate that StyleLipSync achieves state-of-the-art performance in terms of lip-sync and visual quality, even with the zero-shot setting.

2. Related Works

2.1. Lip-sync Video Generation

Lip-sync video generation aims to generate a talking face video with lip motion synchronized with the given input audio. Early works [21, 28] generate lip from the lower half masked face image corresponding to the input audio. Specifically, Wav2Lip [28] uses a pre-trained SyncNet [8] as a lip-sync expert which maximizes the correlation between the generated lip and the input audio. Similarly, we use lip-sync expert for audio-visual alignment, which is trained with a contrastive manner proposed in [24]. VideoRetalking [6] improves Wav2Lip [28] a two-stage manner, where it first generates low-resolution (96×96) video and then increases the resolution by using a single identity specific super-resolution network. In contrast, we propose

a zero-shot model that directly generates lip-sync video of 256×256 resolution and also propose a unseen face adaptation method then enhances the personal characteristics from a few-shot target video. SyncTalkFace [27] introduces a lip memory network, which encodes lip motion features into a discrete space in the training phase and retrieves a lip feature from query audio at the inference phase. In contrast to [27], we utilize continuous and linear latent space from a pre-trained style-based generator [19] to generate lip images with high fidelity and take video consistency into account in the latent space.

2.2. GAN Prior

Style-based generators [18, 19, 17] demonstrate the power of their semantic latent spaces, namely \mathcal{W} , in image generation, image editing [13, 32], and video generation [35]. GAN-inversion [12, 1, 2, 29, 36] utilizes the pre-trained GANs to invert an image into corresponding latent code so one can manipulate the attributes of the image only within the latent space. Extended $\mathcal{W}+$ have been shown its much expressive power. For instance, pSp [29] adopts the feature pyramid networks (FPNs) [22] to use $\mathcal{W}+$, which follow the nature of the progressive generation [15] of StyleGAN [18] and achieves state-of-the-art performance in image-to-image translation (e.g., face inpainting). Similarly, we utilize $\mathcal{W}+$ for diverse and strong lip prior since we aim to build a lip-sync video generative model of arbitrary identity.

Recent works [42, 38, 41] not only adopt pre-trained GAN prior as their decoder but also introduce a skip-connection that concatenates the encoded and generated features, which helps the model preserve 2D spatial information. Specifically, GPEN [42] uses skip-connection that concatenates the encoded and generated features and achieves state-of-the-art performance in blind face restoration. StyleSwap [41] adopts it to face swap and introduces the ToMask branch that predicts the target facial attribute regions for swapping in a supervised manner. In contrast to those methods, we use an additive skip connection more efficient than the concatenation, along with the unsupervised predicted masked sum, which helps the decoder distinguish the target lip region from the whole face and therefore increases the lip image fidelity.

2.3. Personalization

Although GAN prior has been successful in various tasks [29, 36], it is still challenging to faithfully recover person-specific information that lies in the out-of-distribution [1, 2, 29, 36, 42]. Recently, a few-shot personalization has become an alternative to solving the problem [25]. Pivotal-tuning-inversion (PTI) [30] fine-tunes the image generator while freezing a single latent code, namely *pivot*, to compensate the person-specific information in the generative

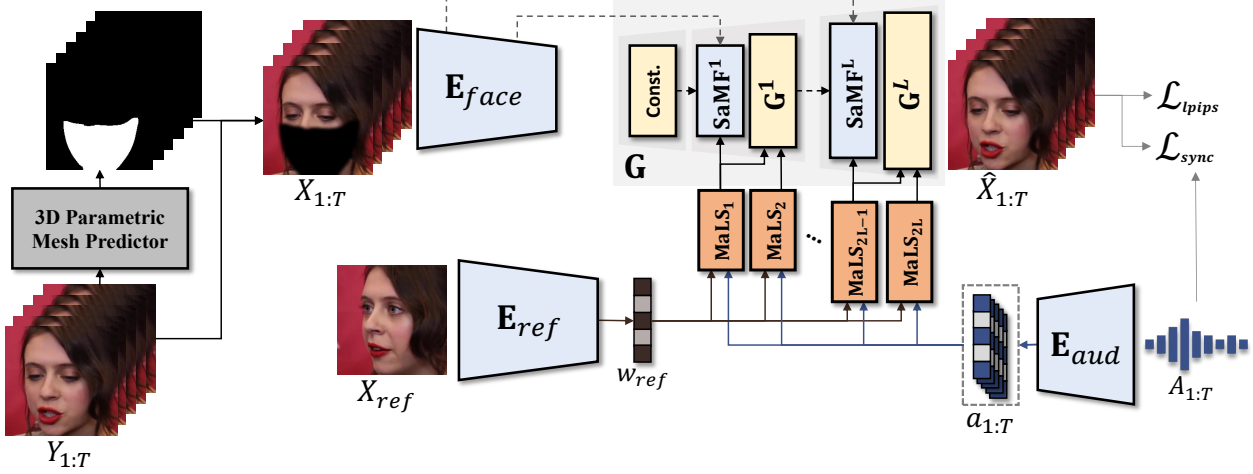


Figure 1. A framework of StyleLipSync. We leverage a 3D parametric mesh predictor [11, 23] to obtain pose-aware masked frames $X_{1:T}$, which inherits the facial pose of input frames. Face encoder E_{face} maps $X_{1:T}$ into 2D spatial features and then fed into the decoder G through *style-aware masked fusion* (SaMF). Single reference image X_{ref} and audio segments $A_{1:T}$ are mapped into latent space, followed by *Moving-average based Latent Smoothing* (MaLS). This module outputs smooth video latent codes $\tilde{w}_{1:T} \subseteq \mathcal{W}+$ that represent temporally consistent lip movement. With the guidance of SaMFs and the smooth video latent codes $\tilde{w}_{1:T}$, StyleLipSync can generate temporally consistent lip-synced videos.

process, not in the encoding process. MyStyle [25] adopts PTI [30] to image inpainting and semantic editing, by restricting the latent space to a subspace spanned by the multiple pivots from the few photos (roughly 100) of the target person. Stitch-it-in-time [37] adopts the pivotal-tuning to the video editing in a multi-stage manner, which leverages an off-the-shelf latent manipulation [31] to manipulate video in the latent space and then stitch it to the source video. Inspired by them, we propose a few-shot unseen face adaptation method that slightly fine-tunes the image decoder for a given latent code trajectory of a target identity and maintains the audio generalization by introducing a sync-regularizer.

3. Method

Given lip masked video frames $X_{1:T} = (X_t)_{t=1}^T$ and audio segments $A_{1:T} = (A_t)_{t=1}^T$, a lip-sync method generates video frames

$$\hat{X}_{1:T} = (\hat{X}_t)_{t=1}^T, \quad (1)$$

where $\hat{X}_{1:T}$ has a lip movement synchronized with the audio segments $A_{1:T}$. In contrast to the previous lip-sync methods [21, 28, 27, 6], we leverage 3D parametric facial mesh predictor [11, 23] to compute lip mask so that the generator can be aware of semantically meaningful facial pose information (Section 3.1). We utilize a pre-trained StyleGAN [19] as our decoder G . When audio encoder E_{aud} and reference encoder E_{ref} map their inputs into the latent space $\mathcal{W}+$ (Section 3.3), the decoder G generates lip-synced video frames $\hat{X}_{1:T}$ from these latent codes (Section 3.2) guided by the proposed *style-aware masked fusion*. For

enhancing the temporal consistency, we propose a *Moving-average based Latent Smoothing* module, which learns local motion between the latent codes, and makes video latent trajectory smoother. Finally, sync loss [8, 24] is used for the audio-lip synchronization. The overall framework of our model is described in Figure 1.

3.1. Pose-aware Masking

Dynamic head motion is an important factor in the natural talking style. However, existing methods [21, 28, 27, 6] employ the rectangular lower-half mouth masking method without consideration of the pose information. It often fails to detect appropriate masking regions when the head moves dynamically, which leads to unpleasant artifacts and unnatural jaw movement in the generated videos (see the first row in Figure 5 for examples). To address this limitation, we use the face meshes by leveraging a 3D face mesh predictor [11], which captures 3D parameters and predicts dense face geometry. We predict the 3D parameters [4] and the mesh from given video frames. Then, the predicted expression parameter $\delta \in \mathbb{R}^{64}$ is used to adjust the mesh to obtain naturally opened and closed mouth meshes. We normalize the mesh vertices using the predicted pose parameters $\tau \in \mathbb{R}^3$ (translation), $\gamma \in \text{SO}(3)$ (rotation) and leave only the lower frontal vertices. These meshes are combined and projected onto the original 2D plane to finally get our pose-aware lip masks. Figure 2 illustrates the framework of the pose-aware masking. This masking not only captures the pose information but also inherits facial semantics such as jaw shape. Ablation studies in Section 5.4 show that the pose-aware masking helps the model to increase visual quality along

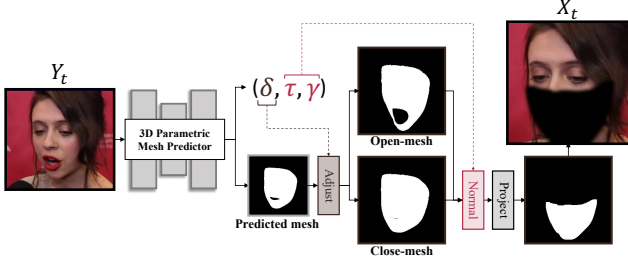


Figure 2. Illustration of pose-aware masking. The expression parameter $\delta \in \mathbb{R}^{64}$ and the pose parameter $\tau \in \mathbb{R}^3, \gamma \in \text{SO}(3)$ are used to compute the natural mask.

with dynamic pose.

3.2. Decoder

Lip Prior from Style-based Generator. Generating the lip-sync videos from scratch for an arbitrary person is very hard since the mapping from audio to lip is basically one-to-many. In this paper, we leverage a style-based generator as our image decoder \mathbf{G} for the following two reasons. First, a pre-trained StyleGAN already contains expressive and diverse face priors [18, 19, 17] in the form of latent code, namely *style code* in the latent space $\mathcal{W}+$ [2]. Those latent spaces enable us to better synthesize the lip region of the target person with the diverse lip prior. Second, the style codes form the continuous and linear [18, 13, 32] latent spaces, which enables us to design a high-level visual transformation, such as natural motion, only with a linear transformation of latent code [39]. Hence, we can generate a talking head video with smooth lip motion by simply manipulating the style codes using audio, which the previous lip-sync methods never take into account.

Style-aware Masked Fusion (SaMF). Recently, it has been explored that adopting skip-connections based on the concatenation to GAN-inversion helps to preserve the 2D spatial information of the input [41, 14, 42]. Similarly, we adopt an additive skip-connection to our model for effectively preserving the non-masked region of $X_{1:T}$ while faithfully utilizing the latent space.

Specifically, we propose *style-aware masked fusion* (SaMF) for efficiently preserving the 2D spatial feature and relieving the information gap between masked and non-masked regions. SaMFs are introduced at the beginning of the decoder blocks. The decoder \mathbf{G} consists of L decoder blocks, each of which takes 2 style codes to modulate 3 convolution weights, as illustrated in Figure 3. The first style code in each decoder block modulates 2 different convolution weights, one for the convolution in the original block and the other for the SaMF. SaMF learns to predict a 1-channel mask S_t^l of the current resolution from the encoded feature through the newly modulated convolution followed by the sigmoid, which is used for spatial weighted fusion of the encoded feature and the generated feature.

Formally, given encoded feature $\mathbf{E}_{face}^l(X_t)$ and gen-

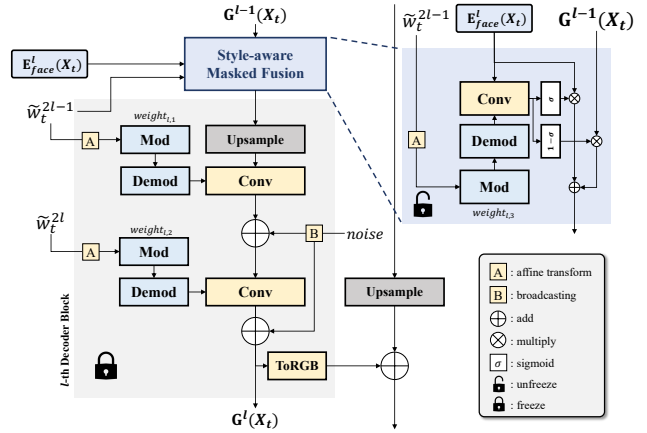


Figure 3. Illustration of the decoder block. The encoded feature $\mathbf{E}_{face}^l(X_t)$ is injected into l -th decoder block through Style-aware Masked Fusion (SaMF). Note that only the convolutions in SaMF are trainable, while the others are frozen during the training phase.

erated feature $\mathbf{G}^{l-1}(X_t)$ of same dimension $\mathbb{R}^{h \times w \times c}$, SaMF firstly predicts a spatial mask $S_t^l \in \mathbb{R}^{h \times w \times 1}$ from $\mathbf{E}_{face}^l(X_t)$ and then output fused feature as follows:

$$S_t^l \otimes \mathbf{E}_{face}^l(X_t) + (1 - S_t^l) \otimes \mathbf{G}^{l-1}(X_t). \quad (2)$$

Ablation studies (Section 5.4) show that SaMFs improve the fidelity of the mouth since they separate masked and non-masked regions.

3.3. Encoders

Our model has three different encoders: face encoder \mathbf{E}_{face} , reference encoder \mathbf{E}_{ref} , and audio encoder \mathbf{E}_{aud} .

Face encoder \mathbf{E}_{face} takes masked video frames $X_{1:T}$ as input and then outputs l 2D spatial features $\mathbf{E}_{face}^l(X_t) = \{\mathbf{E}_{face}^l(X_t) \mid l \in [1, 2, \dots, L]\}$ for each t . These features are injected into the decoder \mathbf{G} through the style-aware masked fusion to efficiently preserve 2D spatial structure, as described in Section 3.2.

Reference Encoder \mathbf{E}_{ref} maps a reference X_{ref} into $2L$ reference style codes, each of which has 512 dimensions. We simply denote the reference style code as $w_{ref} = [w_{ref}^1 | w_{ref}^2 | \dots | w_{ref}^{2L}] \in \mathbb{R}^{512 \times 2L}$.

Similar to \mathbf{E}_{ref} , audio encoder \mathbf{E}_{aud} maps a single audio segment into $2L$ audio style codes, each of which has 512 dimensions. As we use T consecutive audio segments $A_{1:T}$, \mathbf{E}_{aud} independently maps A_t into $a_t = [a_t^1 | a_t^2 | \dots | a_t^{2L}] \in \mathbb{R}^{512 \times 2L}$. We simply denote $a_{1:T} = (a_t)_{t=1}^T$ for T audio style codes. Please refer to our supplementary materials for the detailed encoder architectures.

From these style codes w_{ref} and $a_{1:T}$, we compute target video's style codes over frames $w_{1:T}$ by

$$w_{1:T} = (w_1, w_2, \dots, w_T), \quad (3)$$

where $w_t = w_{ref} + a_t \in \mathcal{W}+$. We compute the style codes by simply adding these two different codes based on the

linearity [13, 32] of the latent space $\mathcal{W}+$. The style codes $w_{1:T}$ are then fed into the decoder \mathbf{G} through the learned affine transformations to generate synced lip motion.

Temporal Consistency. Thanks to the semantically rich latent space $\mathcal{W}+$, our model can generate an accurate lip-sync video frame by frame. However, this frame-wisely independent encoding of style codes turns out to lead inconsistent mouth movements in the final results. To remedy this, we assume that the generated style codes $w_{1:T}$ form a trajectory of the target video [40] in $\mathcal{W}+$ and enforce the smooth local transition of the motion [24, 33, 39] into the trajectory. Toward this, we introduce *Moving-average based Latent Smoothing* (MaLS), each of which consists of a stack of the weighted moving-average [3] and 1D convolutions operating on the style codes along the time-axis. More precisely, we employ $2L$ MaLS for L resolutions, each of which takes the l -th component of w_{t-1}, w_t , and w_{t+1} as its inputs to learn the local difference between them, and then we inject the local motions into w_{ref} to compute the smooth style code \tilde{w}_t :

$$\begin{aligned}\tilde{w}_t^l &= w_{ref}^l + \text{MaLS}^l(w_{t-1:t+1}^l), \\ &= w_{ref}^l + \text{Conv1Ds} \left(\sum_{\tau=t-1}^{t+1} \alpha_\tau \cdot w_\tau^l \right),\end{aligned}\quad (4)$$

where MaLS^l denotes the l -th MaLS, and α_τ is the weight of the moving-average.

With this smooth latent codes $\tilde{w}_{1:T}$, we compute the video frames:

$$\hat{X}_{1:T} = (\mathbf{G}(\tilde{w}_t, \mathbf{E}_{face}(X_t)))_{t=1}^T. \quad (5)$$

For better initialization [29], we add the average code w_{avg} of the pre-trained generator to each \tilde{w}_t in (4).

3.4. Training Objective

We train StyleLipSync to reconstruct target video frames from corresponding audio. We randomly choose $T = 5$ consequent frames with corresponding audio segments and 1 single reference frame. Image perceptual loss [43] is used to minimize perceptual image distance between generated frames \hat{X} and ground-truth frames Y :

$$\mathcal{L}_{l_{lips}} = \sum_{i=1}^N \left\| \phi^i(\hat{X}) - \phi^i(Y) \right\|_2, \quad (6)$$

where N is the number of feature extractor, $\phi^i(\cdot)$ is the i th feature extractor, and $\|\cdot\|_2$ is the ℓ^2 loss. Similar to [33], we use the multi-scale perceptual loss with 3 levels.

For audio-visual alignment, we utilize SyncNet trained in a contrastive manner [24] that minimizes the cosine distance between generated frames \hat{X} and corresponding audio segment A :

$$\mathcal{L}_{sync} = 1 - \cos(f_v(\hat{X}), f_a(A)), \quad (7)$$

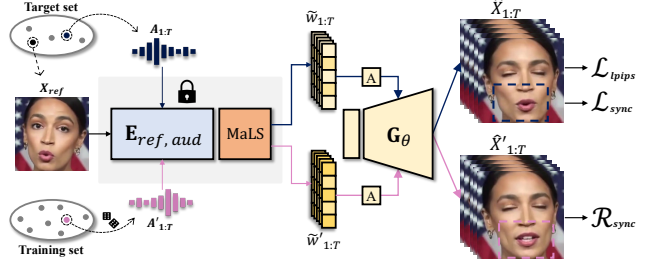


Figure 4. Adaptation for Unseen Face. We slightly tune the decoder \mathbf{G}_θ with the proposed sync regularizer \mathcal{R}_{sync} , while freezing all encoders' weight. Face encoder \mathbf{E}_{face} and SaMFs are omitted here for simplicity.

where $\cos(\cdot, \cdot)$ denotes the cosine similarity. $f_v(\cdot)$ and $f_a(\cdot)$ are the frame and audio feature extractor, respectively. The final objective \mathcal{L}_{train} is computed as:

$$\mathcal{L}_{train} = \lambda_1 \mathcal{L}_{lips} + \lambda_2 \mathcal{L}_{sync}, \quad (8)$$

where λ_1 and λ_2 are the balancing coefficients.

4. Unseen Face Adaptation

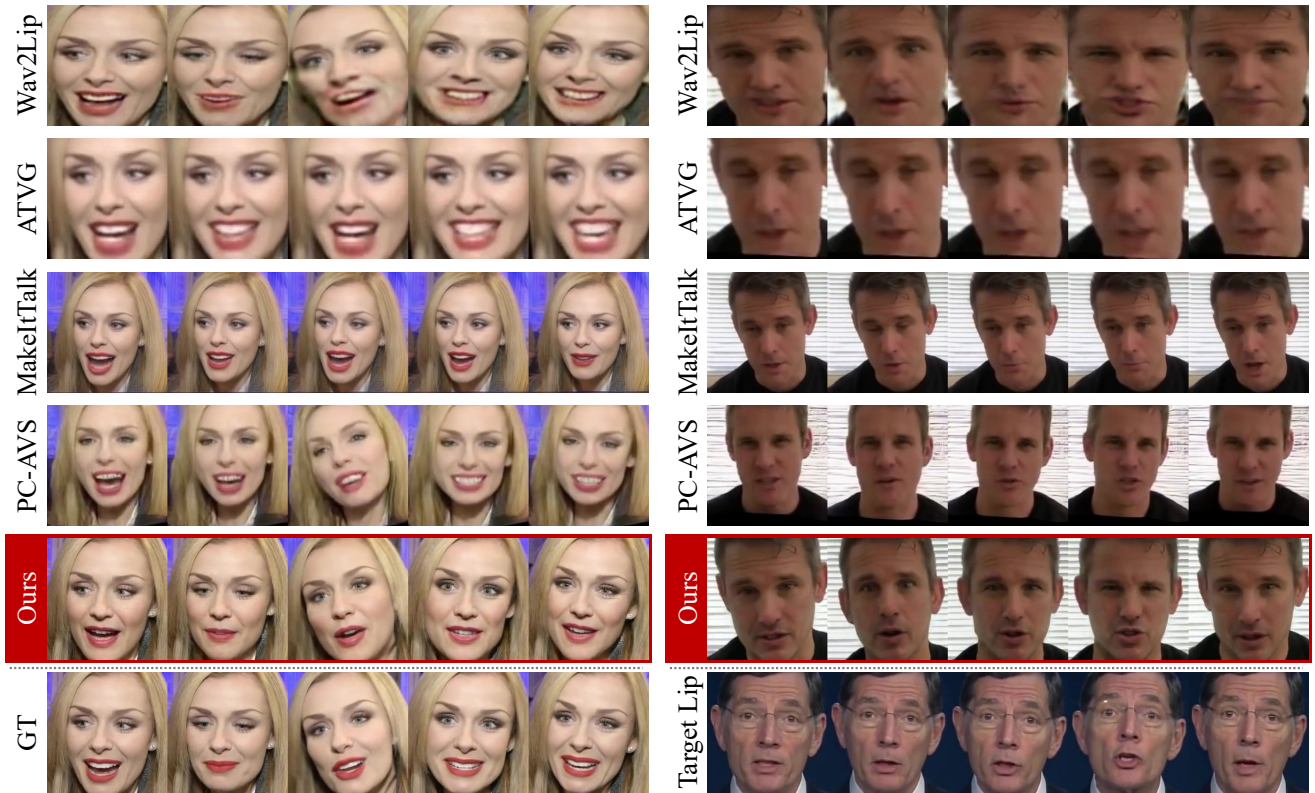
Although StyleLipSync successfully generates accurately lip-synced videos with high fidelity, the model would fail to exactly synthesize unseen faces lying in the out-of-distribution. This problem refers to the *failure of id preservation*. Therefore, to handle this, we fine-tune our decoder \mathbf{G} on a target person video to make a personal model that is better able to synthesize toward the target person.

Let $X_{1:T}$ be masked video frames of *unseen face*, which corresponds to the audio segments $A_{1:T}$, and X_{ref} be a reference frame. The frozen encoders convert each input into the intermediate representations, $\tilde{w}_{1:T}$ and $\mathbf{E}_{face}(X_t)$. Then we fine-tune the decoder \mathbf{G}_θ , now parameterized by θ , by minimizing the distance between $(\mathbf{G}_\theta(\tilde{w}_t, \mathbf{E}_{face}(X_t)))_{t=1}^T$ and target frames, as same as (8). However, fine-tuning the decoder on a short video of a single identity leads to over-fitting and losing the lip-sync generality as the generator can memorize the target video [26]. To prevent the model from this scenario, we introduce a sync regularizer \mathcal{R}_{sync} to enforce audio generality to the decoder \mathbf{G}_θ by leveraging the audio from the training dataset, not from the target video. Formally, given audio segments $A'_{1:T}$ randomly chosen from the training data (Voxceleb2 [7]) and w_{ref} , we compute smooth style codes $\tilde{w}'_{1:T}$, and then decode them to a synced video $X'_{1:T}$. The sync regularizer \mathcal{R}_{sync} is defined as

$$\mathcal{R}_{sync} = 1 - \cos(f_v(\hat{X}'), f_a(A')), \quad (9)$$

which enforces \mathbf{G}_θ to generate $\hat{X}'_{1:T}$ aligning with $A'_{1:T}$. The final object for a single person adaptation is given as follows:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \mathcal{L}_{train} + \lambda_R \mathcal{R}_{sync}, \quad (10)$$



(a) Reconstruction results on Voxceleb2 [7] test data.

(b) Cross-id results on HDTF [44].

Figure 5. Comparison with state-of-the-art methods. The different field of view comes from the pre-processing strategy of each model.

where λ_R is the regularizer coefficient. Ablation studies in Section 5.4 show that \mathcal{R}_{sync} regularizes the audio generalization even if we use the audio from the training set.

5. Experiments

5.1. Dataset

We train our model on Voxceleb2 [7], which consists of in-the-wild talking face videos collected from YouTube. It contains more than 145,000 videos of about 6,100 identities in the train set and more than 4,900 videos of more than 110 identities in the test set. We convert all videos into frames with 25 fps and then crop and resize those frames into 256×256 resolution, following the method of [33]. For more semantically rich face priors, we only use videos where the detected face bounding boxes' height (and width) is longer than 256. After pre-processing, the remaining 11051 videos and 340 videos are used as the training set and the test set, respectively. All audios are re-sampled to 16kHz and converted into a mel-spectrogram to be used as our audio representation similar to the method in [45, 24]. We also use HDTF [44] to further test our model for cross-id experiments. It is widely used to evaluate high-resolution talking face generative models, where the head pose dynamics are not as significant as the Voxceleb2.

5.2. Implementation Details

We pre-train StyleGAN2 [19] on the Voxceleb2 [7] following the implementation of [16]. For training StyleLip-Sync and SyncNet, we set the frame length $T = 5$, and employ Adam [20] optimizer with learning rate 10^{-4} throughout the training phase in both cases. All experiments are performed on 2 TITAN RTX GPUs. We set $\lambda_1 = 10$, $\lambda_2 = 0.1$, $\lambda_R = 0.1$ for the objective (10), and $\alpha_{t-1} = 0.25$, $\alpha_t = 0.5$, $\alpha_{t+1} = 0.25$ for MaLS (4). For all inference, we use the first frame as the reference image.

For image quality metric, we use PSNR, and SSIM. We also calculate CSIM [10] in the cross-id experiment that measures the face similarity between the images in the pre-trained face embedding space. For lip-sync quality metric, we use LMD, LSE-D, and LSE-C. LMD [5] is the absolute distance of facial landmarks between the target and generated frames. LSE-D and LSE-C are proposed in [28], where LSE-D measures the distance between the lip and audio representations, and LSE-C measures the lip-sync confidence, respectively.

5.3. Evaluation

Reconstruction Results. We compare the other state-of-the-art methods in the lip-sync (Wav2Lip [28]) and talking

face generation (ATVG [5], MakeItTalk [46], and PC-AVS [45]) on reconstruction results of Voxceleb2 test dataset. Table 1 shows that StyleLipSync outperforms all metrics except the LSE-C score. Wav2Lip [28] achieves the highest LSE-C score, however, it achieves low image quality metrics since it generates 96×96 low-resolution videos. PC-AVS [45] achieves comparable lip-sync scores with StyleLipSync, however, achieves low image quality metrics than our model since it highly relies on its specific pre-processing and fails to generate in many cases. We also illustrate the qualitative results in Figure 4. Wav2Lip [28] produces a lip-synchronized video with considerable visual artifacts since it cannot adapt to the video of dynamic head pose. MakeItTalk [46] generates low lip-sync video since it uses sparse facial landmarks. PC-AVS [45] generates an accurate lip-synchronized video following the input head pose. However, it struggles to preserve the unseen identity involving visual artifacts. StyleLipSync generates the natural lip motion with high fidelity and preserves the input identity, which is comparable to the ground truth video.

Cross Identity Results. To evaluate lip-sync generalization, we conduct cross-id experiments settings using unseen videos and audio. We randomly sample 10 videos of different identities and 150 audio without duplication from HDTF [44]. We use the first 10 seconds of videos and audio. For each video, we generate 15 lip-synced videos from 15 different audios, where those face-audio pairs are not from the same source. In Table 2, we report the LSE-D and LSE-C for lip-sync quality and the CSIM for face similarity. Wav2Lip [28] achieves a higher LSE-C score than StyleLipSync, however, it achieves a low CSIM score since it produces a low-resolution video with visual artifacts. MakeItTalk [46] achieves the best CSIM score, while the lip-sync quality is the worst since it uses sparse facial landmarks. PC-AVS [45] outperforms LSE-C while achieving the lowest CSIM since it can't preserve an unseen face's identity. StyleLipSync achieves the best LSE-D score and comparable CSIM score to MakeItTalk [46]. We show qualitative results of cross-id experiments with target lip references in Figure 4. PC-AVS [45] can generate accurate lip-synchronizing video compared to the target lip, while it fails to preserve facial details of the unseen face. MakeItTalk [46] produces a high-resolution and identity-preserving video, however, it is out-of-sync compared to the target. StyleLipSync generates a high-resolution lip synchronizing video compared to the target lip, without any visual artifacts.

5.4. Ablation Studies

Ablation Studies on Zero-shot Model. Figure 6 and Table 3 summarize the ablation studies on our zero-shot method of reconstruction on Voxceleb2 test data. If we replace the pose-aware masking with standard rectangular

Table 1. Quantitative comparison of reconstruction on Voxceleb2 test data. The best score for each metric is in **bold**.

Method	Voxceleb2 (Reconstruction)				
	Image			Lip-Sync	
	SSIM \uparrow	PSNR \uparrow	LMD \downarrow	LSE-D \downarrow	LSE-C \uparrow
Wav2Lip $_{96 \times 96}$ [28]	0.448	13.534	6.422	6.999	8.329
ATVG $_{128 \times 128}$ [5]	0.461	13.349	7.165	8.821	5.421
MakeItTalk $_{256 \times 256}$ [46]	0.419	12.868	3.649	10.895	3.624
PC-AVS $_{224 \times 224}$ [45]	0.369	13.210	2.812	7.278	7.699
Ours $_{256 \times 256}$	0.631	19.607	2.696	6.628	8.056

Table 2. Quantitative comparison of cross-identity results for unseen face. We report CSIM [10] as the image quality metric since there is no ground truth frames for the cross-id experiments. The best score for each metric is in **bold**.

Method	HDTF (Cross-id)		
	Image	Lip-Sync	
	CSIM \uparrow	LSE-D \downarrow	LSE-C \uparrow
Wav2Lip $_{96 \times 96}$ [28]	0.656	7.047	8.576
ATVG $_{128 \times 128}$ [5]	0.287	8.668	6.040
MakeItTalk $_{256 \times 256}$ [46]	0.770	10.641	4.725
PC-AVS $_{224 \times 224}$ [45]	0.238	6.921	8.858
Ours $_{256 \times 256}$	0.737	6.825	8.209

Table 3. Ablation study on zero-shot model. The best score for each metric is in **bold**.

Method (ours)	Voxceleb2 (Reconstruction)				
	Image			Lip-Sync	
	SSIM \uparrow	PSNR \uparrow	LMD \downarrow	LSE-D \downarrow	LSE-C
w/o Pose Mask	0.602	18.867	3.057	6.771	7.748
w/o MaLS	0.593	18.186	2.740	6.994	7.577
w/o SaMF	0.591	18.181	2.764	6.838	7.780
Full	0.631	19.607	2.696	6.628	8.056

masking (w/o Pose Mask) in [28, 6, 27], considerable visual artifacts occur around the masked region since it is insufficient to capture the pose difference between the reference and the target. To validate SaMF, we replace the modulated convolutions in SaMFs with the standard convolutions. Figure 6(c) shows that SaMFs significantly improve lip region's fidelity since the modulated convolution helps to be aware of the lip style. As shown in Table 3, MaLS significantly improves lip-sync quality, which cannot be reflected in a single image. Please refers to our project page for ablation studies on MaLS.

Ablation Studies on Unseen Face Adaptation. We conduct ablation studies on the proposed unseen face adaptation following the same setting in Table 2. Additionally, we use 60 seconds of video for each 10 personalized models and 15 audios of 10 seconds from different identities for inference. Figure 7(b) shows the lip-sync metrics according to the adaptation step. In the cases without the sync regularizer, the models lose the lip-sync generality, in other words, it memorizes the short target video as the adaptation phase proceeds. Introducing the sync regularizer with the sync




Figure 6. Qualitative comparison of zero-shot model.

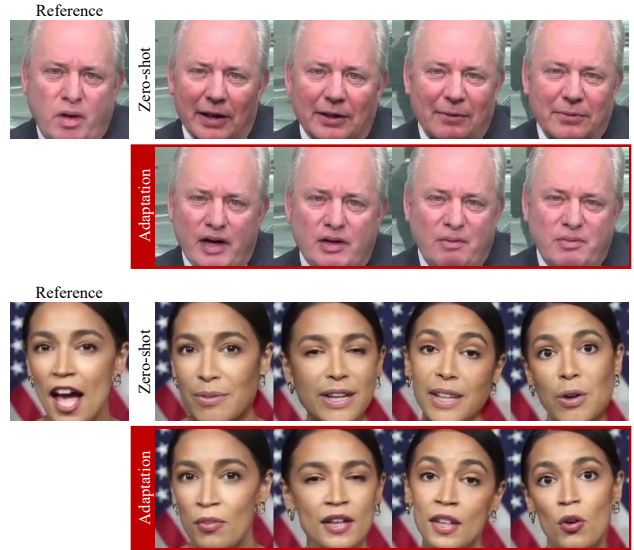
Table 4. Mean Opinion score (MOS) user study results with 95% confidence interval on cross-id setting. The score ranges in 1 to 5. The best score for each metric is in **bold**.

Method	User Study (MOS)		
	Lip-sync Accuracy	Face Similarity	Visual Quality
Wav2Lip [28]	3.76 ± 0.20	2.98 ± 0.25	2.03 ± 0.21
ATVG [5]	2.19 ± 0.21	2.45 ± 0.25	1.54 ± 0.17
MakeItTalk [46]	2.32 ± 0.24	3.47 ± 0.23	2.95 ± 0.24
PC-AVS [45]	3.28 ± 0.21	2.55 ± 0.22	2.51 ± 0.22
Ours	4.01 ± 0.16	3.42 ± 0.22	3.55 ± 0.20
Ours (Personalized)	3.52 ± 0.21	4.03 ± 0.17	3.64 ± 0.19

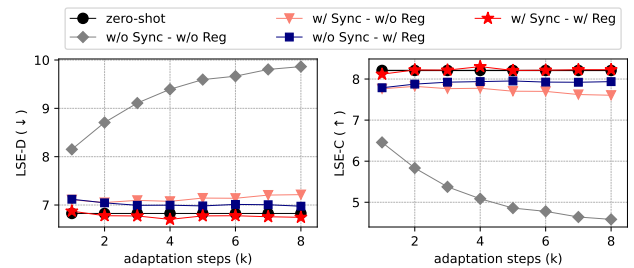
loss stabilizes the lip-sync metrics and even improves them compared to the zero-shot results. If we use the sync regularizer without the sync loss, lip-sync quality is stabilized, however slightly lower than the zero-shot results due to the lack of ground-truth audio-visual correlation. Audio generalization for unseen face data is maintained even though we used audio from the learning data. Figure 7(a) supports the validity of the adaptation method. It shows that visual difference between the zero-shot results and the adaptation results. The zero-shot model can generate an accurate lip motion for the target audio, while it shows a slight difference in person-specific details compared to the reference images. Through the proposed adaptation method, personal-specific lip shape, teeth, and wrinkles are faithfully recovered while maintaining the lip motion of the zero-shot results.

5.5. User Study

We further conduct a user study based on MOS (Mean opinion score) to compare the perceptual quality of each model, including our zero-shot and personalized model. 5 videos generated by each model in cross-id setting are used for this study. 20 participants scored lip-sync accuracy, face similarity, and visual quality of each video in the range of 1 to 5. As shown in Table 4, our models outperform all metrics. Specifically, our zero-shot model achieves the highest lip-sync accuracy, and our adaptation model achieves the highest score in face similarity and visual quality with competitive lip-sync accuracy.



(a) Comparison of cross-id results of zero-shot and adaptation.



(b) Ablation study on the proposed sync regularizer.

Figure 7. Experimental results of the proposed adaptation method using the cross-id setting. It improves the person specific visuals while maintaining the lip-sync generality.

6. Conclusion

We proposed StyleLipSync, a lip-sync video generative model for arbitrary identity, which leverages expressive lip priors from a pre-trained style-based generator. In contrast to existing lip-sync generative models, we introduce pose-aware masking for lip region by utilizing a 3D parametric mesh predictor, which inherits the pose information in the mask itself. Designing a smooth lip motion by using the moving-average based latent smoothing in the continuous and linear latent space, StyleLipSync can generate temporally consistent lip motion. Furthermore, we propose a few-shot lip-sync adaptation method for a single person who lies in the out-of-distribution, which uses a few seconds of the target person’s video. Experimental results show that our StyleLipSync can generate realistic lip-sync video from arbitrary audio even with the zero-shot setting, and the proposed adaptation method enhances the person-specific visual information without losing the lip-sync generality.

Limitation and Future Works. Since learning audio-visual representation requires a large number of different identities (e.g., Voxceleb2 [7]), extending our method to higher resolution ($512 \times 512 \uparrow$) is very challenging. We consider it our future work to develop an effective encoding system for a limited number of identities using a style-based generator [18, 19, 17]. Designing an effective reference encoder to improve lip identity preservation in a zero-shot setting can be another future work.

Ethical Considerations. Since our method can generate a video of a specific person talking specific words only with a few seconds of video source, it has the potential for misuse. As discussed in [6], attaching visual watermarks on the generated videos can be a solution to it.

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] 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.
- [3] Yuval Alaluf, Or Patashnik, Zongze Wu, Asif Zamir, Eli Shechtman, Dani Lischinski, and Daniel Cohen-Or. Third time’s the charm? image and video editing with stylegan3. *arXiv preprint arXiv:2201.13433*, 2022.
- [4] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, pages 187–194, 1999.
- [5] 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.
- [6] Kun Cheng, Xiaodong Cun, Yong Zhang, Menghan Xia, Fei Yin, Mingrui Zhu, Xuan Wang, Jue Wang, and Nannan Wang. Videoretalking: Audio-based lip synchronization for talking head video editing in the wild. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022.
- [7] J. S. Chung, A. Nagrani, and A. Zisserman. Voxceleb2: Deep speaker recognition. In *INTERSPEECH*, 2018.
- [8] Joon Son Chung and Andrew Zisserman. Out of time: automated lip sync in the wild. In *Asian Conference on Computer Vision*, pages 251–263, 2016.
- [9] Dipanjan Das, S. Biswas, Sanjana Sinha, and Brojeshwar Bhowmick. Speech-driven facial animation using cascaded gans for learning of motion and texture. In *European Conference on Computer Vision*, 2020.
- [10] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019.
- [11] 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*, 2019.
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2014.
- [13] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. *Advances in Neural Information Processing Systems*, 33:9841–9850, 2020.
- [14] Jingwen He, Wu Shi, Kai Chen, Lean Fu, and Chao Dong. Gcfsr: a generative and controllable face super resolution method without facial and gan priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1889–1898, 2022.
- [15] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- [16] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems*, 33:12104–12114, 2020.
- [17] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34:852–863, 2021.
- [18] 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.
- [19] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020.
- [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [21] Prajwal KR, Rudrabha Mukhopadhyay, Jerin Philip, Abhishek Jha, Vinay Namboodiri, and CV Jawahar. Towards automatic face-to-face translation. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 1428–1436, 2019.
- [22] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2117–2125, 2017.
- [23] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chu-Ling Chang, Ming Guang Yong, Juhyun Lee, et al. Mediapipe: A framework for building perception pipelines. *arXiv preprint arXiv:1906.08172*, 2019.

- [24] Dongchan Min, Minyoung Song, and Sung Ju Hwang. Styletalker: One-shot style-based audio-driven talking head video generation. *arXiv preprint arXiv:2208.10922*, 2022.
- [25] Yotam Nitzan, Kfir Aberman, Qiurui He, Orly Liba, Michal Yarom, Yossi Gandelsman, Inbar Mosseri, Yael Pritch, and Daniel Cohen-Or. Mystyle: A personalized generative prior. *arXiv preprint arXiv:2203.17272*, 2022.
- [26] Utkarsh Ojha, Yijun Li, Jingwan Lu, Alexei A Efros, Yong Jae Lee, Eli Shechtman, and Richard Zhang. Few-shot image generation via cross-domain correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10743–10752, 2021.
- [27] 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 *36th AAAI Conference on Artificial Intelligence (AAAI 22)*. Association for the Advancement of Artificial Intelligence, 2022.
- [28] 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.
- [29] 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.
- [30] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Transactions on Graphics (TOG)*, 42(1):1–13, 2022.
- [31] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9243–9252, 2020.
- [32] Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1532–1540, 2021.
- [33] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *Advances in Neural Information Processing Systems*, 32, 2019.
- [34] Justus Thies, Mohamed A. Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. Neural voice puppetry: Audio-driven facial reenactment. In *European Conference on Computer Vision*, 2019.
- [35] Yu Tian, Jian Ren, Menglei Chai, Kyle Olszewski, Xi Peng, Dimitris N Metaxas, and Sergey Tulyakov. A good image generator is what you need for high-resolution video synthesis. *arXiv preprint arXiv:2104.15069*, 2021.
- [36] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021.
- [37] Rotem Tzaban, Ron Mokady, Rinon Gal, Amit Bermano, and Daniel Cohen-Or. Stitch it in time: Gan-based facial editing of real videos. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022.
- [38] Xintao Wang, Yu Li, Honglun Zhang, and Ying Shan. Towards real-world blind face restoration with generative facial prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9168–9178, 2021.
- [39] Yaohui Wang, Di Yang, Francois Bremond, and Antitza Dantcheva. Latent image animator: Learning to animate images via latent space navigation. In *International Conference on Learning Representations*, 2022.
- [40] Weihao Xia, Yujiu Yang, and Jing-Hao Xue. Gan inversion for consistent video interpolation and manipulation. *arXiv preprint arXiv:2208.11197*, 2022.
- [41] 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, 2022.
- [42] 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.
- [43] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018.
- [44] Zhimeng Zhang, Lincheng Li, Yu Ding, and Changjie Fan. 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*, pages 3661–3670, 2021.
- [45] 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.
- [46] 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.