Imitating Arbitrary Talking Style for Realistic Audio-Driven Talking Face Synthesis

Haozhe Wu¹, Jia Jia^{1*}, Haoyu Wang¹, Yishun Dou², Chao Duan², Qingshan Deng²

¹Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China

Beijing National Research Center for Information Science and Technology (BNRist)

The Institute for Artificial Intelligence, Tsinghua University

²HiSilicon Company, Shenzhen, China

 $wuhz 19 @mails.tsinghua.edu.cn,\, jj ia @tsinghua.edu.cn,\, wang-hy 18 @mails.tsinghua.edu.cn,\, wang-h$

{douyishun,duanchao15,dengqingshan}@hisilicon.com

ABSTRACT

People talk with diversified styles. For one piece of speech, different talking styles exhibit significant differences in the facial and head pose movements. For example, the "excited" style usually talks with the mouth wide open, while the "solemn" style is more standardized and seldomly exhibits exaggerated motions. Due to such huge differences between different styles, it is necessary to incorporate the talking style into audio-driven talking face synthesis framework. In this paper, we propose to inject style into the talking face synthesis framework through imitating arbitrary talking style of the particular reference video. Specifically, we systematically investigate talking styles with our collected Ted-HD dataset and construct style codes as several statistics of 3D morphable model (3DMM) parameters. Afterwards, we devise a latent-style-fusion (LSF) model to synthesize stylized talking faces by imitating talking styles from the style codes. We emphasize the following novel characteristics of our framework: (1) It doesn't require any annotation of the style, the talking style is learned in an unsupervised manner from talking videos in the wild. (2) It can imitate arbitrary styles from arbitrary videos, and the style codes can also be interpolated to generate new styles. Extensive experiments demonstrate that the proposed framework has the ability to synthesize more natural and expressive talking styles compared with baseline methods.

CCS CONCEPTS

• **Computing methodologies** → *Animation*.

KEYWORDS

photorealistic talking face, talking style, style imitation

ACM Reference Format:

Haozhe Wu, Jia Jia, Haoyu Wang, Yishun Dou, Chao Duan, Qingshan Deng. 2021. Imitating Arbitrary Talking Style for Realistic Audio-Driven Talking Face Synthesis. In *Proceedings of the 29th ACM International Conference on*

MM '21, October 20-24, 2021, Virtual Event, China

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8651-7/21/10...\$15.00 https://doi.org/10.1145/3474085.3475280 Multimedia (MM '21), October 20-24, 2021, Virtual Event, China. ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3474085.3475280

1 INTRODUCTION



Figure 1: Different talking styles have significant differences in facial and head pose movements when pronouncing the word "significant".

The talking face synthesis is an eagerly anticipated technique for several applications, such as filmmaking, teleconference, virtual/mix reality, and human-computer interaction. One of the key essences behind the talking face synthesis is the stylization of facial and head pose movements. Different from the talking emotion which reflects in the short-term facial motions, the talking style is a crucial factor which affects long-term facial and head pose movements. People usually talk with diversified talking styles such as "excited", "solemn", "communicational", "storytelling", et al. Given one piece of speech, different talking styles exhibit significant differences in the facial and head pose movements. For example, as shown in Figure 1, people with the "excited" style usually talk aloud, and thus the facial movement of the mouth wide open occurs frequently. Meanwhile, the "solemn" talking style usually appears in formal occasions, and thus the exaggerated motion seldom occurs. Considering such huge differences between different styles, in order to synthesize diversified and realistic talking faces with respect to one piece of speech, it is necessary to incorporate talking style into the audio-driven talking face synthesis framework.

^{*}Corresponding author.

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



Figure 2: The overall framework of our method contains two stages: the stylized 3D talking face synthesis stage and the photorealistic render stage. The first stage synthesizes stylized 3D talking faces by imitating talking styles from the style codes. The second stage synthesizes photorealistic videos with deferred neural render and neural texture generation model. The stylized 3D talking face synthesis stage and the photorealistic render stage are trained separately in our framework. *T* denotes the time dimension of DeepSpeech features. For the calculation of style codes, *even style reference videos from one identity would yield diversified style codes*.

Previous efforts [10, 36] have shown the rationality of synthesizing stylized talking faces. Yi et al. [36] proposed Memory Augmented GAN model to synthesize stylized talking face with the wild training data. Cudeiro et al. [10] proposed Voice Operated Character Animation (VOCA) model to learn the identity-wise talking style. The VOCA model captures the talking style of each identity by injecting the one-hot identity vector into the audio-motion predicting network. Overall, these methods have the following two disadvantages which constrain the expressiveness of the synthesized talking styles: (1) These methods assume that each identity has only one talking style. However, in the real scenario, the talking style is only relatively stable inside one video clip, one person can talk with significantly different styles among different video clips. (2) These methods require substantial labour to collect adequate synchronized audio-visual data for each identity, which is unavailable for the wild scenarios.

To address the aforementioned problems, we propose to synthesize stylized talking faces by imitating styles from arbitrary clips of videos. With such motivation, the talking videos with stable and diversified talking styles are in need. Therefore, we collect the *Ted-HD* dataset which has 834 clips of videos with 60 identities. Each video clip has an average length of 23.5 seconds. The talking styles of the *Ted-HD* dataset are diversified, and the talking style inside each video clip is stable.

Based on the constructed dataset, we devise a two stage talking face synthesis framework as shown in Figure 2. The first stage imitates talking styles from arbitrary videos and synthesizes 3D talking faces according to the driven speech. Afterwards, in the second stage, we render the 3D face model photo-realistically from one static portrait of the speaker. Overall, the key idea of our framework is to construct style codes from video clips in the wild, and then synthesize talking faces by imitating the talking styles from the constructed style codes. Specifically, for the style codes construction, we conduct exhaustive observations on the Ted-HD dataset, which on one hand verify that even one identity has multiple talking styles, on the other hand find that the talking style is closely related to the variance of facial and head pose movements inside each video. With our observations, we define the style codes as several interpretable statistics of 3D morphable model (3DMM) [2] parameters. Having obtained the style codes of each talking video, we devise a latent-style-fusion (LSF) model to synthesize stylized 3D talking faces by imitating talking styles from the style codes. Detailedly, the LSF model first dropouts [27] information from the audio stream to prevent the audio from dominating the synthesis process. Further, the LSF model fuses the style codes with the latent audio representation frame-by-frame to synthesize 3D talking face with corresponding talking style. The overall implementation of the LSF model is simple but effective. Our model not only circumvents the annotation for talking styles and avoids collecting substantial training data for each identity, but also enables new talking style generation.

With our proposed framework, talking faces with diversified styles can thus be synthesized. We perform experiments on the *Ted-HD* dataset for evaluation. Compared with baseline methods, our framework synthesizes more expressive and diversified talking styles. We conduct extensive user studies to investigate the facial motion naturalness and audio-visual synchronization. With the mean opinion score (MOS) of 20 participants, our framework outperforms baseline methods by 0.67 on average in terms of facial motion naturalness and 0.11 on average in terms of audio-visual synchronization.

To summarize, our contributions are summarized as three-fold:

- We propose to synthesize stylized talking faces by imitating talking styles of arbitrary videos. The incorporation of style imitation leads to more diversified talking styles.
- We formalize the style codes of each talking video and devise the latent-style-fusion (LSF) model to synthesize stylized 3D talking faces from style codes and driven audio. Our framework do not require any annotation of talking styles for stylized talking face synthesis.
- We collect *Ted-HD* dataset, which contains 834 clips of wild talking videos with stable and diversified talking styles. Based on the *Ted-HD* dataset, we conduct extensive style observations and synthesize expressive talking styles. Code and dataset are publicly available ¹.

2 RELATED WORK

Talking face synthesis has received significant attention in previous literatures. Related work in this area can be grouped into two categories: the unimodal talking face synthesis [5, 18, 24, 29, 35, 39, 40] and the multimodal talking face synthesis [10, 13, 19, 32, 36, 38]. For one piece of driven speech, the unimodal talking face synthesis generates unique motion, while the multimodal talking face synthesis generates diversified facial and head pose movements.

Most of the prior works focused on the unimodal talking face synthesis. Karras *et al.* [21] proposed to synthesize 3D talking face with driven audio and emotion state. Suwajanakorn *et al.* [28] synthesized high quality videos of talking Obama through hours of Obama's weekly address footage. Since that Suwajanakorn's method requires hours of data for each identity, several methods [4– 6, 24, 29, 37, 39, 41] proposed to simultaneously reduce the required duration of training data and guarantee the photorealism of the synthesized video. The ATVG framework proposed by Chen *et al.* [6] and the DAVS framework proposed by Zhou *et al.* [39] synthesize talking face with only one image. Despite that these unimodal talking face synthesis methods can synthesize photo-realistic videos, the lack of style results in the stiffness of the synthesized results.

To synthesize diversified facial and head pose movements, some recent literatures addressed the multimodal talking face synthesis. Wang et al. [32] and Eskimez et al. [13] achieved multimodal synthesis through incorporating emotion condition vector, which enables the generation of diversified facial expressions. However, the synthesized results of these methods still lack in personality because of the overlook of talking style. To address this issue, some methods [10, 36] proposed to incorporate talking style into the synthesis framework. Yi et al. [36] proposed Memory-Augmented GAN model to synthesize stylized talking face with wild training data. However, only the head pose is synthesized with multiple styles, while the facial movements still lack personality. Further, Cudeiro et al. [10] proposed the Voice Operated Character Animation (VOCA) model to learn the identity-wise talking style. The VOCA model injects one-hot identity vector into the audio-motion predicting network, leading to discriminative styles of facial and head pose movements. However, the VOCA model on one hand requires substantial data for each identity, on the other hand forces one identity to have only one talking style, limiting its capacity on synthesizing

diversified styles. To resolve this problem, in this work, we propose to imitate talking styles from arbitrary wild talking videos.

3 PROBLEM FORMULATION

In this paper, we propose a two-stage talking face synthesis framework which synthesizes stylized talking videos with the following three inputs: one static portrait of the speaker, the driven audio and the style reference video. We formalize the first stage of the framework as the 3D talking face synthesis stage and the second stage as the photorealistic render stage. Between two stages, we apply the 3DMM face model [2] as a crucial bridge. Therefore, before formally defining the two stages, we firstly give a brief introduction of the face model we use.

We leverage the 3DMM face model to represent each 3D face. With 3DMM, the face shape **S** is represented as an affine model of facial expression and facial identity:

$$\mathbf{S} = \mathbf{S}(\alpha, \beta) = \bar{\mathbf{S}} + \mathbf{B}_{id}\alpha + \mathbf{B}_{exp}\beta,\tag{1}$$

where $\mathbf{\bar{S}} \in \mathbf{R}^{N \times 3}$ is the average face shape; N is the number of the vertexes in the face model; \mathbf{B}_{id} and \mathbf{B}_{exp} are the PCA bases of identity and expression; α and β are the identity parameters and the expression parameters. Following deng *et al.* [12], we adopt the 2009 Basel Face Model [23] for $\mathbf{\bar{S}}$, \mathbf{B}_{id} , and use expression bases \mathbf{B}_{exp} of Guo *et al.* [14] built from Facewarehouse [3], resulting in $\alpha \in \mathbb{R}^{80}$, $\beta \in \mathbb{R}^{64}$. Afterwards, the 3D face shape is projected on the 2D plane according to the head pose and translation $p \in \mathbb{R}^7$, where 4 elements represent the pose quaternion and 3 elements represent the translation. Overall, the parameter set (α, β, p) controls the appearance of each face. In our framework, the facial movements are the time series of parameter β , which we denote as $\beta(t)$, while the head pose movements are the time series of parameter p, which we denote as p(t). With $\beta(t)$ and p(t), we formalize the two stages of our framework as follows.

3D Talking Face Synthesis Stage. In this stage, given driven audio X_a , the facial and head pose movements of style reference video $\beta_{sty}(t)$, $p_{sty}(t)$, we aim to generate corresponding facial and head pose movements $\beta_{pred}(t)$, $p_{pred}(t)$.

Photorealistic Render Stage. In this stage, given the predicted movements $\beta_{pred}(t)$, $p_{pred}(t)$ and the input portrait X_p , we aim to generate photorealistic videos Y.

4 TALKING STYLE OBSERVATION

In this section, we systematically investigate how different talking styles reflect in the facial and head pose movements $\beta(t), p(t)$. Afterwards, we formally define interpretable style codes for each video based on our observation.

In order to observe talking styles of each video, we should firstly collect a suitable dataset for observation. The video for the style observation requires the following characteristics: (1) It should be of high resolution, (2) it should contain natural and expressive facial and head pose movements, (3) each clip of video cannot be too short, or the talking style can hardly be observed, (4) the talking style should both be stable inside the clip and be diversified across different clips, (5) the camera pose and location should be static inside each clip, otherwise the head pose parameters will be influenced by the camera movements. Considering the characteristics

¹https://github.com/wuhaozhe/style_avatar



Figure 3: Correlation between talking styles and 3DMM parameter series. We investigate 8 different statistics and find that $\sigma(\beta(t)), \sigma(\frac{\partial \beta(t)}{\partial t}), \sigma(\frac{\partial p(t)}{\partial t})$ are mostly correlated with talking styles.

above, current publicly available wild datasets like VoxCeleb2 [7] and LRS3 [1] are much too noisy and thus do not meet the requirements. Meanwhile, several in-lab datasets like MEAD [33] and GRID [9] do not have natural facial expressions and head poses, which dissatisfy the requirements either.

To address this issue, we manually collect a wild dataset *Ted-HD* which is suitable for style observation and further style synthesis. The *Ted-HD* dataset selects several speech videos from Ted website. Each video in the dataset has one person giving a speech, which focuses on the facial part of each person and is of high resolution. We cut each video into several clips according to the scene change. Totally, *Ted-HD* dataset contains 834 clips of videos with 60 identities. The average length of each clip is 23.5 seconds, and the total duration of the dataset is 6 hours. The talking style of these videos are diversified across different clips. Even for the same identity, there could be different talking styles.

Having obtained the dataset, for each video, we reconstruct the facial and head pose movements $\beta(t)$, p(t). We conduct exhaustive data observations on the correlation between talking style and $\beta(t)$, p(t), with the aim of answering the following questions:

- Q1: Whether one identity has multiple talking styles.
- **Q2**: How the talking style is reflected in time series $\beta(t)$, p(t).

To answer the **Q1**, we verify the talking style diversity of each identity through user study with A/B test. Specifically: we first randomly construct 100 triplets, each triplet contains two talking videos v_1, v_2 from the same identity and one talking video v_3 from the other identity. Next, we reconstruct $\beta(t), p(t)$ of v_1, v_2, v_3 , retarget $\beta(t), p(t)$ to the same identity and render the retargeted faces to videos. The retargeted videos are signified as v'_1, v'_2, v'_3 . Afterwards, we show v'_1, v'_2, v'_3 and their transcripts to users, asking the following question: which pair of (v'_1, v'_2) and (v'_1, v'_3) has more similar talking styles. Statistics demonstrate that among 100 triplets, (v'_1, v'_3) is more similar in 30 triplets, while (v'_1, v'_2) is more similar in 70 triplets. With the statistics that 30% of videos inside each identity

have dissimilar talking styles, we conclude that one identity has multiple talking styles.

Since that the talking style is not consistent for each identity, formulating style codes is necessary. Therefore, in Q2, we conduct experiments to find how talking style is reflected in time series $\beta(t), p(t)$. Specifically, we first randomly select 300 talking videos and annotate the talking style of these videos to three categories: tedious, solemn, and excited. Afterwards, for the motion series $\beta(t), p(t)$ of each video, we calculate its derivative series with respect to time *t*: $(\frac{\partial \beta(t)}{\partial t}, \frac{\partial p(t)}{\partial t})$. Next, we calculate the mean value $\mu(\cdot)$ and the standard deviation $\sigma(\cdot)$ of $(\beta(t), p(t), \frac{\partial \beta(t)}{\partial t}, \frac{\partial p(t)}{\partial t})$ along time, yielding 8 feature vectors (4 for mean, 4 for standard deviation). To observe the relationship between these feature vectors and talking styles, we leverage t-SNE algorithm [31] to visualize each feature vector and plot points from different style categories with different colors, as shown in Figure 3. Figure 3 demonstrates that the talking style is closely related with $\sigma(\beta(t))$, $\sigma(\frac{\partial \beta(t)}{\partial t}), \sigma(\frac{\partial p(t)}{\partial t})$, especially $\sigma(\frac{\partial p(t)}{\partial t})$. Meanwhile the talking style is less related with $\mu(\cdot)$, which denotes that the talking style is mostly reflected in the fluctuation of movements, rather than the idle state of movements.

Based on such observation, we define the style codes as the standard deviation of facial and head pose movements. Formally, given arbitrary video with the reconstructed parameter series $\beta(t)$, p(t), the style codes **sty** are defined as:

$$\mathbf{sty} = \sigma(\beta(t)) \oplus \sigma(\frac{\partial \beta(t)}{\partial t}) \oplus \sigma(\frac{\partial p(t)}{\partial t}), \tag{2}$$

where \oplus signifies the vector concatenation.

To summarize, we make the following two conclusions: (1) one identity have multiple talking styles, (2) the talking style is closely related to the variance of facial and head pose movements inside each video, following which we define the style codes in Equation 2. The style codes are utilized to synthesize diversified talking styles, details will be illustrated in the section 5.

5 METHODOLOGY

Following the style codes defined in section 4, we propose a twostage talking face synthesis framework to imitate arbitrary talking styles, as shown in Figure 2. Our framework synthesizes stylized talking videos with the following three inputs: one static portrait of the speaker, the driven audio and the style reference video. In the first stage of our framework, we devise a latent-style-fusion (LSF) model to synthesize stylized 3D talking faces through imitating arbitrary talking styles. Based on the synthesized 3D talking faces, in the second stage, we leverage the deferred neural render [30] and few-shot neural texture generation model to generate video frames photo-realistically. In the next two subsections, we will introduce the two stages respectively.

Stylized 3D Talking Face Synthesis 5.1

In the first stage of our framework, we propose the latent-stylefusion (LSF) model for talking face synthesis. Overall, the input of the LSF model is the driven audio and the reference talking video for style imitation. The LSF model learns motion related information from audio, and then combines latent audio representation with style information to synthesize 3D talking meshes with target talking style. Details will be illustrated as follows.

For the driven input audio X_a of T seconds, we firstly leverage the DeepSpeech [15] model to extract speech features. Deepspeech is a deep neural model for automatic speech recognition (ASR). The extracted features from DeepSpeech not only contains rich speech information but also is robust to background noise and generalizes well to different identities. Feeding the input audio X_a to the DeepSpeech model, yields latent representations $X_d \in \mathbb{R}^{50T \times D_a}$, where D_a is the dimension of the DeepSpeech features, and 50T denotes that the DeepSpeech features have 50 frames per second. Afterwards, for the reference talking video, we calculate its style codes sty $\in \mathbb{R}^{D_s}$ as illustrated in Section 4 for style imitation, where D_s is the dimension of the style codes.

Having obtained the X_d and sty, we now elaborate the 3D talking face synthesis process. We devise a latent-style-fusion (LSF) model, which takes $\mathbf{X}_d \in \mathbb{R}^{50T \times D_a}$ and $\mathbf{sty} \in \mathbb{R}^{D_s}$ as input, outputs facial movements $\beta_{pred}(t) \in \mathbb{R}^{25T \times 64}$ and head pose movements $p_{pred}(t) \in \mathbb{R}^{25T \times 7}$ with 25 frames per second. Based on $\beta_{pred}(t)$ and $p_{pred}(t)$, we reconstruct the 3D talking meshes with the 3DMM face model [2].

The LSF model leverages a latent fusion mechanism to both synthesize stylized faces and guarantee the synchronization between audio and motion. Specifically, as shown in Figure 2, the LSF model firstly takes audio X_d as input and encodes X_d with the bottom part of ResNet-50 [16], yielding latent audio representation X1. Afterwards, the LSF model fuses the latent audio representation X_I and style codes sty to acquire mixed representation for synthesis. During the fusion process, the LSF model first dropouts the latent audio representation \mathbf{X}_l to obtain \mathbf{X}'_l , while the information from sty remains unchanged. Next, the LSF model concats each frame of X'_{l} with style codes sty, leading to the mixed representation. Further, the LSF model leverages the top part of ResNet-50 to predict facial movements $\beta_{pred}(t)$ and head pose movements $p_{pred}(t)$ from the mixed representation. It's noteworthy to stress that the fusion between latent audio representation and style codes

enables the synthesis of more expressive talking styles. Meanwhile, the dropout of audio information prevents from discarding style information for synthesis. The overall implementation of the LSF model is simple but effective.

For the training stage of the LSF model, we adopt the parameter series $\beta(t), p(t)$ reconstructed from the 3D face reconstruction algorithm [12] as ground truth. For each training video, we calculate its style codes sty and randomly clip the the input audio X_d and the ground truth $\beta(t)$, p(t) to fixed length. Afterwards, we feed X_d and corresponding sty to the LSF model, yielding the predicted expression parameter series $\beta_{pred}(t)$ and $p_{pred}(t)$. Based on the predicted $\beta_{pred}(t)$ and $p_{pred}(t)$, we apply the L₁ loss as follows:

$$\mathcal{L}_{L_1} = ||\beta(t) - \beta_{pred}(t)||_1 + ||p(t) - p_{pred}(t)||_1.$$
(3)

It is worth emphasizing that the training of the LSF model doesn't require any additional annotation on the speaking identity. Only by training on wild videos with stable talking styles, we obtain expressive style space.

During the inference stage, feeding style codes of arbitrary talking videos to LSF model not only yields desired talking styles but also maintains the synchronization between the driven audio and talking faces. Meanwhile, we can interpolate among different talking styles to acquire new talking styles. For the audio representation X_d with arbitrary duration, since that the trained LSF model only digests the audio with fixed length, we apply the slide window strategy to synthesize the corresponding facial movements $\beta_{pred}(t)$ and head pose movements $p_{pred}(t)$.

So far, we have obtained the untextured talking 3D faces. In the next subsection, we'll introduce how we render these 3D faces photo-realistically.

5.2 **Photorealistic Render**

Conventional deferred neural render [30] requires substantial training data for each identity. In order to both synthesize photorealistic results and guarantee the few-shot capacity, we devise a few-shot neural texture generation model and combine the generated neural texture with the deferred neural render, which enables synthesizing photorealistic videos with only one source portrait. As shown in Figure 2, the deferred neural render incorporates the generated neural texture, conducts UV texture sampling on the neural texture, and translates the sampled image to photorealistic domain.

Algorithm 1 UV Texture Sampling
1: $\mathbf{X}_{uv} \in \mathbb{R}^{2 \times H \times W}$
2: $\mathbf{Y}_t \in \mathbb{R}^{D_t \times H_t \times W_t}$
3: $\mathbf{X}_{s} \in \mathbb{R}^{D_{t} \times H \times W}$
4: for $i \leftarrow 1$ to H do
5: for $j \leftarrow 1$ to W do
$u = \mathbf{X}_{uv}[0, i, j]$
7: $v = \mathbf{X}_{uv}[1, i, j]$
8: $\mathbf{X}_{s}[:, i, j] = \text{BilinearSample}(\mathbf{Y}_{t}, u, v)$
9: return X _s

Detailedly, for the input 3D talking faces, we firstly leverage the UVAtlas tool ² to obtain the UV coordinate of each vertex in

²https://github.com/microsoft/UVAtlas

the 3D model. Afterwards, we rasterize the 3D face model to 2D image $X_{uv} \in \mathbb{R}^{2 \times H \times W}$, of which each pixel represents the UV coordinate. Subsequently, for the input portrait $X_p \in \mathbb{R}^{3 \times H \times W}$ and the 3D face model, we extract the RGB texture $X_t \in \mathbb{R}^{3 \times H_t \times W_t}$, where H_t, W_t signifies the height and width of the texture. Based on X_t , we leverage the pix2pix [17] model to generate neural texture $Y_t \in \mathbb{R}^{D_t \times H_t \times W_t}$, where D_t signifies the dimension of the neural texture. With the neural texture, we conduct UV texture sampling on X_{uv} to obtain the sampled image $X_s \in \mathbb{R}^{D_t \times H \times W_t}$, the details of the sampling algorithm are illustrated in the Algorithm 1. Finally, we translate the sampled image X_s to photorealistic image through the pix2pixHD [34] model.

During the training stage, the few-shot texture generation model and the deferred neural render model are trained simultaneously. Given the rasterized input X_{uv} , we denote the rendered image as Y' and the ground truth image as Y. We combine the perceptual loss [20] and L₁ loss together as \mathcal{L} to optimize the neural texture and pix2pixHD model. Formally:

$$\mathcal{L} = ||\mathbf{Y} - \mathbf{Y}'||_1 + ||\phi(\mathbf{Y}) - \phi(\mathbf{Y}')||_1,$$
(4)

where $\phi(\cdot)$ is the first few layers of VGGNet [26] pretrained on the ImageNet [11]. Due to the limitation of the 3DMM face model, in our render, we only synthesize the facial part of the image, without considering the hair and background rendering.

6 EXPERIMENTS

In this section, we conduct extensive experiments to demonstrate the effectiveness of our framework. We evaluate our framework on the collected *Ted-HD* dataset. Our method has acquired better synthesis results both qualitatively and quantitatively.

6.1 Dataset

As illustrated in Section 4, the currently available datasets are either collected in lab which have constrained talking styles or collected in the wild of which the talking styles are unstable and noisy. Therefore for the training and testing of the LSF model, we leverage the *Ted-HD* dataset described in Section 4. Totally, the *Ted-HD* dataset has 834 clips of videos, we select 799 clips for training, and hold out the remained 35 for testing. The training set and test set have no overlap on identities. Additionally, for the training of the deferred neural render and the few-shot neural texture generation model, we leverage the Lip Reading in the Wild (LRW) [8] dataset.

6.2 Implementation Details

During the training of the LSF model, the input DeepSpeech audio features have 50 frames per second (FPS), while each frame has the dimension D_a of 29. The input style codes **sty** have the dimension D_s of 135 (64 for $\sigma(\beta(t))$, 64 for $\sigma(\frac{\partial\beta(t)}{\partial t})$, 7 for $\sigma(\frac{\partial p(t)}{\partial t})$). The predicted facial movements $\beta_{pred}(t)$ and the head pose movements $p_{pred}(t)$ have 25 frames per second. For the convenience of training, we randomly clip the input to 80 frames, and clip the output to 32 frames. For the implementation of the LSF model, we apply the ResNet-50 as the backbone. We leverage the first 16 layers of ResNet-50 to encode the DeepSpeech features, combine the encoded

features with style codes, and leverage the last 34 layers of ResNet-50 to predict the motion series. When optimizing, we adopt the Adam optimizer [22] to train the LSF model with the initial learning rate of 5×10^{-4} . We train for 50000 iterations with a mini-batch size of 128 samples.

During the training of the deferred neural render and the fewshot neural texture generation model, the input UV image X_{uv} has a size of $2 \times 224 \times 224$, while the neural texture has a size of $16 \times 64 \times 64$. The dimension D_t of the texture is set to 16, which enables each pixel to contain richer texture information. Meanwhile, the texture size is set to be smaller than the UV image size, which avoids oversampling in the sample process. The output image X_s is conventional RGB image with size of $3 \times 224 \times 224$. When optimizing, we adopt the Adam optimizer to simultaneously train the neural render and the texture generation model. The learning rate is set to 2×10^{-4} . We train for 1000000 iterations with mini-batch size of 6 samples.

6.3 Comparison with VOCA on Style Synthesis

To the best of our knowledge, the VOCA model [10] is the only available method that captures diversified talking styles of facial movements. Therefore, in this section, we systematically compare our method with the VOCA model [10] to demonstrate the effectiveness of the LSF model. Different from our method, the VOCA model learns identity-level talking styles. Specifically, the VOCA model injects a one-hot identity code to the time convolution network, and directly predicts face model vertexes from the DeepSpeech features and the identity code. By adjusting the one-hot identity code, the VOCA model outputs different talking styles.

Table 1: The mean opinion scores (MOS) of different metrics, higher signifies better. TSE denotes talking style expressiveness, FMN denotes facial movement naturalness, and AVS denotes audio-visual synchronization.

	TSE	FMN	AVS
VOCA [10]	3.28	3.13	3.56
LSF	3.41	3.21	3.64

We compare the style space learned from the VOCA model and the style space learned from the LSF model with extensive user studies. Specifically, we randomly select 10 clips of driven audio, each with a duration from 10 to 20 seconds. Afterwards, we randomly sample 5 talking styles from the VOCA style space and 5 talking styles from the style space of our method. With the sampled talking styles and the driven audio, we synthesize corresponding talking faces and retarget the synthesized faces to the same identity. Afterwards, we subsume videos with the same driven audio and synthesis model to the same group. For each group of video, we invite 20 participants to rate (1) the expressiveness of the talking style, (2) the naturalness of facial movements, (3) the audio-visual synchronization between driven audio and talking face. We ask participants to rate the mean opinion score (MOS) [25] in the range 1-5 (higher MOS score denotes better results). When showing videos to participants, since that the VOCA model only gives untextured 3D faces, we also just give the 3D talking faces synthesized from the



Figure 4: The interpolation results between excited talking style and solemn talking style with the same driven audio. The first row has solemn talking style, and the last row has excited talking style, while the middle row has the average style between the first and the last. Among these rows, we observe smooth transition of facial and head pose movements.

LSF model, without utilizing the deferred neural render to generate photorealistic results for fair comparison.

Table 1 demonstrates the results of user studies. Through the experimental results, we observe that our LSF model has acquired higher talking style expressiveness, which verifies that the style space learned by taking style imitation is more expressive than the identity-level style space. Meanwhile, our LSF model outperforms the VOCA model by 0.08 in terms of the facial movement naturalness and audio-visual synchronization. Such results confirm the effectiveness of style imitation in LSF model further. Additionally, compared with the VOCA model which requires substantial training data for each identity, our method is trained on wild dataset which does not require any annotation on identity or talking style.

6.4 Study on the Style Space

In this section, we extensively investigate the style space learned in our method. We conduct two qualitative experiments, which verify that not only the synthesized talking styles are diversified, but new talking styles can also be generated from the interpolation of different talking styles.



Figure 5: The distance between the lower lip and the upper lip for the same driven audio conditioned on different talking styles. Different color denotes different talking styles.

To verify the diversification of the talking styles, we visualize the distance between lower and upper lip as a function of time. Specifically, we randomly select one piece of driven audio and 10 different talking styles, and then synthesize 10 facial movements corresponding to each talking style and driven audio. Afterwards, we calculate the lip distance as shown in Figure 5. Through Figure 5 we observe that the lip distance significantly varies among different talking styles, which verifies that the LSF model is able to synthesize diversified talking styles. Meanwhile, different talking styles demonstrate similar trends of fluctuation, the peak and the valley of the distance curve highly overlap, which confirms that the LSF model also guarantees the synchronization between audio and synthesized motions.

To confirm that the style space in the LSF model is expressive and interpolative, we visualize the interpolation results of different talking styles, as shown in Figure 4. Detailedly, we select two representative talking styles: excited and solemn, and conduct linear interpolation between the two styles to generate new talking styles. From each row of Figure 4, we observe that the facial and head pose movements translate smoothly from excited talking style to solemn style. For the excited talking style, the lip motion is exaggerated and the head frequently shakes. Meanwhile, for the solemn talking style, the lip motion and head pose are stable. We provide more synthesis results in the supplementary materials.

6.5 Comparison with One-Shot Synthesis

In this section, we conduct experiments to demonstrate that our method synthesizes more natural and expressive talking faces compared with several baseline methods. Specifically, we compare our method with the following baseline methods: (1) the ATVG framework [6], (2) the MakeltTalk framework [41], (3) the Wav2Lip framework [24]. For Wav2Lip framework which requires few seconds of videos as input, we repeat input portrait as videos for fair comparison. Figure 6 shows some synthesis results We observe that our method has more expressive facial movements and head pose movements while also guarantees the synthesized results to be photorealistic.

Additionally, we conduct both user studies and quantitative evaluations on the *Ted-HD* dataset to verify the effectiveness of our method. Specifically, for the user studies, we firstly synthesize



Figure 6: Comparison with several baseline methods (ATVG, MakeItTalk and Wav2Lip). Our method not only yields expressive facial and head pose movements, but also synthesizes photorealistic videos.

videos with the randomly selected 20 clips of driven audios and 5 different identities. Afterwards, for each video, we invite 20 participants to rate (1) the naturalness of facial movements, (2) the audio-visual synchronization between driven audio and talking face. The mean opinion score (MOS) is rated in the range 1-5. Besides, we also evaluate the synthesized video quality with the signal-noiseratio (SNR) metric. We do not leverage the PSNR metric since that the ground truth talking videos with arbitrary talking styles are not available.

Table 2: Comparison with baseline methods on the *Ted-HD* dataset, where FMN denotes facial movement naturalness, and AVS denotes audio-visual synchronization. The Pre-Fusion method removes latent style fusion in the LSF model, details are illustrated in Section 6.6

	FMN	AVS	SNR (dB)
ATVG [6]	2.71	3.51	2.98
MakeItTalk [41]	3.08	3.13	3.01
Wav2Lip [24]	2.97	4.28	2.78
Pre-Fusion	2.19	2.25	5.70
Ours	3.59	3.75	5.76

Table 2 shows the comparison results. From Table 2 we observe that our method has achieved the most expressive facial motion and best video quality. We also notice that the Wav2Lip method achieves unsatisfying motion naturalness since that it cannot resolve the oneshot synthesis scenario. Meanwhile, we observe that the AVS of our method is slightly lower than Wav2Lip, that is because the talking style synthesis in our LSF model mildly sacrifices the performance of audio-visual synchronization, but our method still has better AVS performance compared with ATVG and MakeItTalk.

6.6 Effectiveness of Latent Style Fusion

To verify the rationality of the latent style fusion mechanism in LSF model, we conduct the following ablation experiments. For comparison, we remove the latent style fusion mechanism, and directly concatenate DeepSpeech representation with style codes as the input of ResNet-50. Afterwards, we compare the synthesized results with LSF model through similar user studies as Section 6.5 does. The experimental results are shown in the last two rows of Table 2. From the results we observe that the motion naturalness and audio-visual synchronization would degrade significantly once we remove the latent style fusion mechanism, which verifies the effectiveness of the latent style fusion. Meanwhile, the video quality remains constant since that the motion synthesis does not influence the photorealistic render stage.

7 CONCLUSION

In this paper, we propose the concept of style imitation for audiodriven talking face synthesis. To imitate arbitrary talking styles, we firstly formulate the style codes of each talking video as several interpretable statistics of 3DMM parameters. Afterwards, we devise a latent-style-fusion (LSF) model to synthesize stylized talking faces according to the style codes and driven audio. The incorporation of style imitation not only circumvents the annotation for talking style during the training phase, but also endows the capacity of arbitrary style synthesis and new talking style generation. Additionally, to synthesize expressive talking styles, we collect Ted-HD dataset with 834 clips of talking videos, which contains stable and diversified talking styles. We conduct extensive experiments on the constructed dataset and obtain expressive synthesis results with our Ted-HD dataset and LSF model. The constructed Ted-HD dataset will be made publicly available in the future. We hope that the proposal of talking style imitation and the construction of Ted-HD dataset pave a new way for audio-driven talking face synthesis.

8 ACKNOWLEDGMENTS

This work is supported by the National Key R&D Program of China under Grant No. 2020AAA0108600, the state key program of the National Natural Science Foundation of China (NSFC) (No.61831022), Beijing Academy of Artificial Intelligence No. BAAI2019QN0302.

REFERENCES

- T. Afouras, J. S. Chung, and A. Zisserman. 2018. LRS3-TED: a large-scale dataset for visual speech recognition. In arXiv preprint arXiv:1809.00496.
- [2] Volker Blanz and Thomas Vetter. 1999. A morphable model for the synthesis of 3D faces. In Proceedings of the 26th annual conference on Computer graphics and interactive techniques. 187–194.
- [3] Chen Cao, Yanlin Weng, Shun Zhou, Yiying Tong, and Kun Zhou. 2013. Facewarehouse: A 3d facial expression database for visual computing. *IEEE Transactions* on Visualization and Computer Graphics 20, 3 (2013), 413–425.
- [4] Lele Chen, Guofeng Cui, Celong Liu, Zhong Li, Ziyi Kou, Yi Xu, and Chenliang Xu. 2020. Talking-head Generation with Rhythmic Head Motion. In *European Conference on Computer Vision*. Springer, 35–51.
- [5] Lele Chen, Zhiheng Li, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. 2018. Lip movements generation at a glance. In Proceedings of the European Conference on Computer Vision (ECCV). 520–535.
- [6] Lele Chen, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. 2019. Hierarchical cross-modal talking face generation with dynamic pixel-wise loss. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 7832– 7841.
- [7] J. S. Chung, A. Nagrani, and A. Zisserman. 2018. VoxCeleb2: Deep Speaker Recognition. In INTERSPEECH.
- [8] J. S. Chung and A. Zisserman. 2016. Lip Reading in the Wild. In Asian Conference on Computer Vision.
- [9] Martin Cooke, Jon Barker, Stuart Cunningham, and Xu Shao. 2006. An audiovisual corpus for speech perception and automatic speech recognition. The Journal of the Acoustical Society of America 120, 5 (2006), 2421–2424.
- [10] Daniel Cudeiro, Timo Bolkart, Cassidy Laidlaw, Anurag Ranjan, and Michael J Black. 2019. Capture, learning, and synthesis of 3D speaking styles. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 10101– 10111.
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition. Ieee, 248–255.
- [12] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. 2019. 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. 0–0.
- [13] Sefik Emre Eskimez, You Zhang, and Zhiyao Duan. 2020. Speech Driven Talking Face Generation from a Single Image and an Emotion Condition. arXiv e-prints (2020), arXiv-2008.
- [14] Yudong Guo, Jianfei Cai, Boyi Jiang, Jianmin Zheng, et al. 2018. Cnn-based real-time dense face reconstruction with inverse-rendered photo-realistic face images. *IEEE transactions on pattern analysis and machine intelligence* 41, 6 (2018), 1294–1307.
- [15] Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, et al. 2014. Deep speech: Scaling up end-to-end speech recognition. arXiv preprint arXiv:1412.5567 (2014).
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition. 770–778.
- [17] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-Image Translation with Conditional Adversarial Networks. CVPR (2017).
- [18] Amir Jamaludin, Joon Son Chung, and Andrew Zisserman. 2019. You said that?: Synthesising talking faces from audio. *International Journal of Computer Vision* 127, 11 (2019), 1767–1779.
- [19] Xinya Ji, Hang Zhou, Kaisiyuan Wang, Wayne Wu, Chen Change Loy, Xun Cao, and Feng Xu. 2021. Audio-driven emotional video portraits. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 14080–14089.
- [20] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual losses for realtime style transfer and super-resolution. In *European conference on computer* vision. Springer, 694–711.
- [21] Tero Karras, Timo Aila, Samuli Laine, Antti Herva, and Jaakko Lehtinen. 2017. Audio-driven facial animation by joint end-to-end learning of pose and emotion.

ACM Transactions on Graphics (TOG) 36, 4 (2017), 1-12.

- [22] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, Yoshua Bengio and Yann LeCun (Eds.). http://arxiv.org/abs/1412.6980
- [23] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. 2009. A 3D face model for pose and illumination invariant face recognition. In 2009 sixth IEEE international conference on advanced video and signal based surveillance. Ieee, 296–301.
- [24] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar. 2020. A lip sync expert is all you need for speech to lip generation in the wild. In Proceedings of the 28th ACM International Conference on Multimedia. 484–492.
- [25] ITUT Recommendation. 2006. Vocabulary for performance and quality of service.
 [26] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, Voshua Bengio and Yann LeCun (Eds.). http://arxiv.org/abs/1409.1556
- [27] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
- [28] Supasorn Suwajanakorn, Steven M Seitz, and Ira Kemelmacher-Shlizerman. 2017. Synthesizing obama: learning lip sync from audio. ACM Transactions on Graphics (ToG) 36, 4 (2017), 1–13.
- [29] Justus Thies, Mohamed Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. 2020. Neural voice puppetry: Audio-driven facial reenactment. In European Conference on Computer Vision. Springer, 716–731.
- [30] Justus Thies, Michael Zollhöfer, and Matthias Nießner. 2019. Deferred neural rendering: Image synthesis using neural textures. ACM Transactions on Graphics (TOG) 38, 4 (2019), 1–12.
- [31] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. Journal of machine learning research 9, 11 (2008).
- [32] Kaisiyuan Wang, Qianyi Wu, Linsen Song, Zhuoqian Yang, Wayne Wu, Chen Qian, Ran He, Yu Qiao, and Chen Change Loy. 2020. Mead: A large-scale audiovisual dataset for emotional talking-face generation. In *European Conference on Computer Vision*. Springer, 700–717.
- [33] Kaisiyuan Wang, Qianyi Wu, Linsen Song, Zhuoqian Yang, Wayne Wu, Chen Qian, Ran He, Yu Qiao, and Chen Change Loy. 2020. MEAD: A Large-scale Audio-visual Dataset for Emotional Talking-face Generation. In ECCV.
- [34] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. 2018. High-resolution image synthesis and semantic manipulation with conditional gans. In Proceedings of the IEEE conference on computer vision and pattern recognition. 8798–8807.
- [35] Olivia Wiles, A Koepke, and Andrew Zisserman. 2018. X2face: A network for controlling face generation using images, audio, and pose codes. In Proceedings of the European conference on computer vision (ECCV). 670–686.
- [36] Ran Yi, Zipeng Ye, Juyong Zhang, Hujun Bao, and Yong-Jin Liu. 2020. Audiodriven talking face video generation with learning-based personalized head pose. arXiv e-prints (2020), arXiv-2002.
- [37] Lingyun Yu, Jun Yu, Mengyan Li, and Qiang Ling. 2020. Multimodal inputs driven talking face generation with spatial-temporal dependency. *IEEE Transactions on Circuits and Systems for Video Technology* (2020).
- [38] Dan Zeng, Han Liu, Hui Lin, and Shiming Ge. 2020. Talking Face Generation with Expression-Tailored Generative Adversarial Network. In Proceedings of the 28th ACM International Conference on Multimedia. 1716–1724.
- [39] Hang Zhou, Yu Liu, Ziwei Liu, Ping Luo, and Xiaogang Wang. 2019. Talking face generation by adversarially disentangled audio-visual representation. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 33. 9299–9306.
- [40] Hang Zhou, Yasheng Sun, Wayne Wu, Chen Change Loy, Xiaogang Wang, and Ziwei Liu. 2021. Pose-controllable talking face generation by implicitly modularized audio-visual representation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 4176–4186.
- [41] Yang Zhou, Xintong Han, Eli Shechtman, Jose Echevarria, Evangelos Kalogerakis, and Dingzeyu Li. 2020. MakeltTalk: speaker-aware talking-head animation. ACM Transactions on Graphics (TOG) 39, 6 (2020), 1–15.