

Talking Head Generation with Audio and Speech Related Facial Action Units

Sen Chen
chensen@tju.edu.cn

Zhilei Liu*
zhileiliu@tju.edu.cn

Jiaxing Liu
jiaxingliu@tju.edu.cn

Zhengxiang Yan
zhengxiangyan@tju.edu.cn

Longbiao Wang*
longbiao_wang@tju.edu.cn

College of Intelligence and Computing
Tianjin University
Tianjin, 300072, China

Abstract

The task of talking head generation is to synthesize a lip synchronized talking head video by inputting an arbitrary face image and audio clips. Most existing methods ignore the local driving information of the mouth muscles. In this paper, we propose a novel recurrent generative network that uses both audio and speech-related facial action units (AUs) as the driving information. AU information related to the mouth can guide the movement of the mouth more accurately. Since speech is highly correlated with speech-related AUs, we propose an Audio-to-AU module in our system to predict the speech-related AU information from speech. In addition, we use AU classifier to ensure that the generated images contain correct AU information. Frame discriminator is also constructed for adversarial training to improve the realism of the generated face. We verify the effectiveness of our model on the GRID dataset and TCD-TIMIT dataset. We also conduct an ablation study to verify the contribution of each component in our model. Quantitative and qualitative experiments demonstrate that our method outperforms existing methods in both image quality and lip-sync accuracy.

1 Introduction

Recently, talking head generation has attracted more and more attention in the fields of academic and industry, which is essential in the applications of human-computer interaction, film making, virtual reality, computer games, etc. This research explores how to generate a talking head video by inputting anyone's image as an identity image and driving information related to mouth movement, e.g., speech audio, and text.

Before deep learning became popular, many researchers in early work relied on Hidden Markov Models (HMM) to capture the dynamic relationship between audio and lip motion

[8, 24, 35]. In recent years, Deep Neural Networks (DNN) is widely used in talking head generation. Some methods selected the best-matched lip region image from the specific person’s database by inputting audio information, then synthesize it into the target face [10, 30]. These methods are subject dependent and bring a lot of overhead when transferring to a new subject. Later, many works study how to generate arbitrary speaker instead of specific speaker [5, 22, 36]. However, these works ignore the temporal relationship of features, so the generated videos are accompanied by jitter. Some researchers used sequence model to learn the temporal relationship to reduce jitter [6, 29, 32]. Nevertheless, talking is also a kind of movement driven by facial muscles, especially in the mouth region. Existing methods ignore the local driving information of the mouth muscles, resulting in the lip movement are not very synchronized with the audio. Inspired by this limitation, we consider using both audio and local information of mouth muscles to drive the talking head generation.

Facial Action Coding System (FACS) is a comprehensive and objective description system of facial movements [7]. It defines a set of basic facial action units (AUs), and each AU represents a basic facial muscle movement. The value of AUs can either use binary classification to indicate whether these AUs are activated, or use intensity value to indicate activation intensity. FACS has attracted much attention in face editing [17, 18, 23], such as facial expression editing [23]. These works proved that AU information can be used to edit local facial regions. AU labels are usually extracted by face image detection [2, 27], but they can also be extracted through other modalities related to facial motion, such as speech audio. A few researchers have constructed the relationship between audio information and speech-related AU information in recent years [21, 24]. They proved that audio as information for speech-related AU recognition is feasible. Therefore, we propose an Audio-to-AU module to obtain speech-related AU representation from speech as local driving information.

To tackle the limitations of existing methods that ignore the local information of the mouth muscles, this paper proposes a novel talking head generation model using speech-related AUs as local information to drive the muscle movements in the mouth region. Specifically, we use both audio and AU representation as the driving information, so that audio information drives the whole mouth, and AU information focuses on the local muscles. In addition, we also use a recurrent generative network to capture the temporal dependence between consecutive frames. In summary, our contribution can be summarized as follows:

- A novel recurrent generative network is proposed by integrating both audio and speech-related facial AUs as driving information to drive talking head generation.
- An Audio2AU module is designed in our framework to predict speech-related AU information from speech. Moreover, we add pre-trained AU classifier to supervise that the generated images contain correct AU information.
- Extensive quantitative and qualitative experiments conducted on two datasets demonstrate that our framework achieves high-quality talking head generation of arbitrary identities and achieves significant improvement over existing methods.

2 Related Work

2.1 Audio-driven talking head generation

In early work, researchers relied on the modeling method of HMM. Bregler et al.[3] proposed a video rewrite method to match the lip movement to the phoneme sequence of the

new audio track. Xie et al. [54] used a coupled Hidden Markov Model (cHMM) to model the subtle relationship of audio-visual speech. Later, deep learning methods were introduced to talking head generation. Suwajanakorn et al. [60] used a database of speeches by the then US President Obama to retrieve the lip region image that best matches the input audio, and synthesized it into the target image. Chung et al. [6] proposed an encoder-decoder structure that uses audio frames to drive the mouth movement of the input face. Prajwal et al. [22] used a pre-trained SyncNet as lip-sync discriminator for adversarial training with the generator to supervise the correctness of mouth motion. Eskimez et al. [9] introduced an emotion encoder to generate talking face video with a specific emotion, while our method generates talking face video with specific AUs. Zhou et al. [58] disentangled the content and speaker information in the input audio signal, and used landmarks as an intermediate representation that reflect speaker-aware dynamics. Different from the method of Zhou et al. [58], we use speech-related AUs to indicate whether the specific muscles should be activated, instead of learning the location information represented by landmarks. Some researchers use the sequence modeling method to enhance the temporal dependence of generated frames. Song et al. [29] combined audio features and identity features into a Recurrent Neural Network (RNN) to enhance the temporal dependence. Vougioukas et al. [51] proposed a temporal Generative Adversarial Network (GAN) to generate more natural face image sequence. However, all the above methods did not consider the local information of the mouth muscles.

2.2 AU-based face editing

As one of the most comprehensive ways to describe facial movements, Facial Action Coding System (FACS) has recently attracted widespread attention [18, 19, 23]. Pumarola et al. [23] proposed an AU-based face editing system, which uses AU intensity labels to edit the input face to generate a face with specific facial muscles action. Liu et al. [18] utilized local AU regional rules to control the status of each AU and used an attentive mechanism to combine them into the whole facial expressions. Zhou et al. [59] proposed a conditional difference adversarial autoencoder for facial expression synthesis, which can generate a face image with a target AU label. In our method, we utilize speech-related AU information to edit the local region of the mouth.

2.3 Speech-related AU recognition

Since speech is highly correlated with speech-related AUs, a few researchers have tried to establish the relationship between audio information and speech-related AU information in recent years [20, 24]. Meng et al. [20] proposed a continuous-time Bayesian network (CTBN) to simulate the dynamic relationship between phoneme and AUs. Then AU recognition is performed by probabilistic inference via the CTBN model. Ringeval et al. [24] used LLD features of acoustic feature datasets to carry out action unit recognition through an RNN-LSTM model. In our work, we propose an Audio-to-AU module to predict speech-related AU information from every audio frame to guide more accurate mouth movement.

2.4 GAN-based image synthesis

Generative Adversarial Network (GAN) [10] is a framework for estimating generative models through adversarial processes. It includes a generator to produce realistic fake samples and a discriminator to distinguish real and fake images. Recently works have proved that

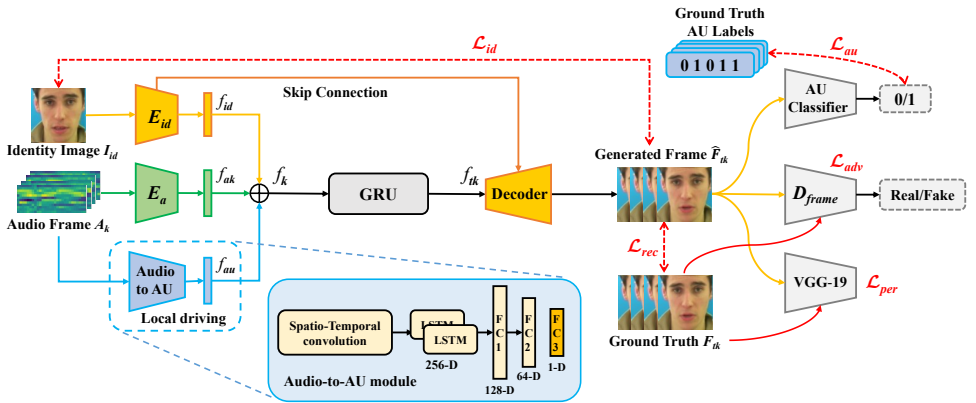


Figure 1: The pipeline of our proposed method

GAN can produce realistic images, and it has achieved excellent results in many fields, such as image translation [13], face generation [9, 29, 37]. Chen et al. [9] devised a cascade GAN approach to generate the talking face. Eskimez et al. [8] used GAN training to improve the image quality and mouth-speech synchronization. Here, we use the GAN training method to enforce the generated image distribution to approach real image distribution.

3 Proposed Method

The architecture of our proposed method is shown in Fig. 1. The generator contains an identity encoder, audio encoder, Audio-to-AU module, RNN module, and image decoder. In addition, we propose AU classifier to supervise whether the AU information of the generated image is accurate. We use the VGG-19 network [28] to extract high-level features of the generated frame and ground truth frame, and compare the differences between them. Frame discriminator is used for adversarial training [10] with the generator to make the generated frame more realistic. In the following, we describe each module in detail.

3.1 Network architecture

Identity Encoder: The identity encoder E_{id} contains four 2D convolution layers and a fully connected layer, each convolution layer is followed by ReLU activation function. The input I_{id} is a face image resized to 112×112 as identity image. It can be a randomly selected frame from the video. For convenience, we use the first frame of each video in this paper. The output of E_{id} is a 512-dimensional identity feature vector $f_{id} = E_{id}(I_{id})$.

Audio Encoder: The audio encoder E_a contains five 2D convolution layers and two fully connected layers, and each convolution layer is followed by batch normalization and ReLU activation function. For an audio sequence, we preprocess it before input. Specifically, we extract Mel-Frequency Cepstral Coefficient (MFCC) features, and then use a fixed size sliding window to crop MFCC segments continuously. Finally, the input are continuous MFCC frames $A = [A_1, A_2, \dots, A_k, \dots, A_n]$, and the outputs are a series of 512-dimensional audio feature vectors $f_a = [f_{a1}, f_{a2}, \dots, f_{ak}, \dots, f_{an}]$, where $f_{ak} = E_a(A_k)$.



Figure 2: Interpretation of speech-related AUs

Audio-to-AU module: To guide the muscle movements in the mouth region more accurately, we propose an Audio-to-AU module to extract speech-related AU information from the speech in real-time. It is pre-trained with the paired audio and AU data. The pre-trained Audio-to-AU module is shown in Fig. 1. In the pre-training stage, the input MFCC features are respectively convolved in the frequency domain and time domain. They then pass through two LSTM layers and three fully connected layers and use the sigmoid activation function to obtain the probability of AU occurrence. After pre-training, we remove the last fully connected layer of this network as the Audio-to-AU module and add it to our framework, because the multi-dimensional AU representation obtained by the penultimate full connected layer is more conducive to the model to learn AU information than a one-dimensional label. Since talking is a facial movement driven by multiple facial muscles in the mouth region, multiple Audio-to-AU modules are necessary to extract this local facial information. Based on the anatomy knowledge [19], five speech-related AUs are selected in this work, namely AU10, AU14, AU20, AU25, and AU26. The detailed interpretation is shown in Fig. 2. Each output of Audio-to-AU module is a 64-dimensional AU feature vector f_{au_i} , where i is the number of the i -th AU. We concatenate the five 64-dimensional vectors and finally get the 320-dimensional AU feature vector $f_{au} = f_{au_{10}} \oplus f_{au_{14}} \oplus f_{au_{20}} \oplus f_{au_{25}} \oplus f_{au_{26}}$.

Recurrent Neural Network module: The RNN module uses a Gate Recurrent Unit (GRU) layer to maintain the temporal dependence between frames. We concatenate the identity feature f_{id} , the audio feature f_{ak} and the speech-related AUs feature f_{au} to get $f_k = f_{id} \oplus f_{ak} \oplus f_{au}$, where k represents the k -th frame in the sequence. We input the feature sequence into GRU and get $f_i = [f_{i1}, f_{i2}, \dots, f_{ik}, \dots, f_{in}]$, where $f_{ik} = RNN(f_k)$.

Image Decoder: The image decoder is used to generate a talking head video. It consists of a fully connected layer and six transposed convolution layers. In order to preserve the input identity information and the facial texture, we utilize a structure similar to U-Net [25], which uses skip connection between the identity encoder and the image decoder. From the input feature $f_i = [f_{i1}, f_{i2}, \dots, f_{ik}, \dots, f_{in}]$, we can get the decoded image sequence $\hat{F}_i = [\hat{F}_{i1}, \hat{F}_{i2}, \dots, \hat{F}_{ik}, \dots, \hat{F}_{in}]$.

AU Classifier: The AU classifier is used to predict the occurrence probability of the speech-related AUs in the generated frames. It consists of four convolution layers and three fully connected layers. Every two convolution layers are followed by a max-pooling layer. Then the sigmoid function is used to get the occurrence probability of each AU. To focus on the mouth region, we only input the lower face of the generated image. Binary Cross-Entropy Loss is used to calculate the loss between the predicted AU and the ground truth AU label:

$$\mathcal{L}_{bce} = -\frac{1}{n_{au}} \sum_{i=1}^{n_{au}} w_i [y_i \log \hat{y}_i + (1 - y_i) \log(1 - \hat{y}_i)] \quad (1)$$

where y_i represents the ground truth label of the i -th AU, which is 1 if occurrence and 0 otherwise, and \hat{y}_i represents the corresponding predicted probability of occurrence. For most facial AU detection benchmarks, the occurrence rates of AUs are imbalanced [19]. Because

AUs are not independent of each other, the imbalance of training data will negatively impact the performance. Therefore, we add weight w_i in Eq. 1 to alleviate the data imbalance problem, where $w_i = \frac{(1/r_i)n_{au}}{\sum_{i=1}^{n_{au}} 1/r_i}$, r_i is the occurrence rate of the i -th AU in the training set [27].

Some AUs rarely appear in the training set, so the network prediction tends to be absent. To alleviate this problem, we introduce a weighted multi-label Dice coefficient loss [28]:

$$\mathcal{L}_{dice} = \frac{1}{n_{au}} \sum_{i=1}^{n_{au}} w_i \left[1 - \frac{2y_i \hat{y}_i + \varepsilon}{y_i^2 + \hat{y}_i^2 + \varepsilon} \right] \quad (2)$$

where ε is the smooth term. Thus, our AU loss is defined as:

$$\mathcal{L}_{au} = \mathcal{L}_{bce} + \mathcal{L}_{dice} \quad (3)$$

Frame Discriminator: The frame discriminator is used for adversarial training to improve the realism of generated frame \hat{F}_{tk} . The adversarial loss is shown in Eq. 4:

$$\mathcal{L}_{adv} = \mathbb{E}_{F_{tk} \sim P_F} [\log D_{frame}(F_{tk})] + \mathbb{E}_{\hat{F}_{tk} \sim P_{\hat{F}}} [\log(1 - D_{frame}(\hat{F}_{tk}))] \quad (4)$$

where P_F and $P_{\hat{F}}$ denote the distributions of ground truth images and generated images respectively, and F_{tk} is the ground truth of the k -th frame.

VGG-19 Network: In order to improve the quality of generated images, we use perceptual loss [29] to reflect perceptual-level similarity of images. The pre-trained VGG-19 network [28] is adopted as the perceptual feature extractor. The perceptual loss is defined as:

$$\mathcal{L}_{per}(F_{tk}, \hat{F}_{tk}) = \frac{1}{n} \sum_{i=1}^n \|\phi_i(F_{tk}) - \phi_i(\hat{F}_{tk})\|_1 \quad (5)$$

where ϕ_i denotes the i -th feature extraction layer of the VGG-19 network.

3.2 Loss functions

In addition to AU loss, adversarial loss and perceptual loss mentioned above, reconstruction loss and identity loss are also considered. The reconstruction loss L_{rec} is used to minimize the pixel-level difference between the generated image \hat{F}_{tk} and the real image F_{tk} :

$$\mathcal{L}_{rec}(\hat{F}_{tk}, F_{tk}) = \|\hat{F}_{tk} - F_{tk}\|_1 \quad (6)$$

The identity loss L_{id} is used to maintain the identity information, and to reduce the jitter effect of generated video. Specifically, it penalizes the difference between the upper face of \hat{F}_{tk} and the upper face of I_{id} :

$$\mathcal{L}_{id}(\hat{F}_{tk_p}, I_{id_p}) = \frac{1}{W \times \frac{H}{2}} \sum_{p \in [0, W] \times [0, \frac{H}{2}]} \|\hat{F}_{tk_p} - I_{id_p}\|_1 \quad (7)$$

where W and H are the width and height of each frame respectively, \hat{F}_{tk_p} is the upper half of generated image \hat{F}_{tk} , and I_{id_p} is the upper half of input image I_{id} .

Finally, the overall loss of our proposed framework is defined as Eq. 8:

$$\mathcal{L}_{total} = \lambda_{rec} \mathcal{L}_{rec} + \lambda_{id} \mathcal{L}_{id} + \lambda_{per} \mathcal{L}_{per} + \lambda_{au} \mathcal{L}_{au} + \lambda_{adv} \mathcal{L}_{adv} \quad (8)$$

where λ_{rec} , λ_{id} , λ_{per} , λ_{au} , λ_{adv} are trade-off parameters.

4 Experimental Results

4.1 Dataset

We conduct extensive experiments on the GRID dataset [6] and TCD-TIMIT dataset [12]. The GRID dataset is a large audio-visual corpus, which consists of high-quality audio and video recordings of 1000 sentences spoken by each of 33 speakers. The TCD-TIMIT dataset has high-quality audio and video footage of 59 speakers uttering approximately 100 phonetically rich sentences each. In our experiments, the GRID dataset is divided into the training set and test set at the ratio of 8: 2, with 27 speakers as the training set and 6 speakers as the test set. For the TCD-TIMIT dataset, we divide 49 speakers into the training set and 10 speakers into the test set. We extract all the video frames and use the Dlib toolkit [16] of the HOG-based face detection algorithm to detect all the faces. Then we crop the face regions and resize them to 112×112 . OpenFace [11, 12] is used to detect the AU labels of each real face image as ground truth labels. For the audio inputs, we extract 12-dimensional MFCC features (excluding energy dimension). We try different lengths of audio frames and find that 280ms performs best. Then we align the middle of each audio frame with a corresponding video frame. The audio sliding window slides synchronously with the video frame.

4.2 Training details

Our network is implemented using Pytorch and trained on a single NVIDIA Titan V GPU. The Audio-to-AU module and the AU classifier are pre-trained on the GRID dataset, and when experimenting on the TCD-TIMIT dataset, they are refined to fit the new dataset. The VGG-19 network is pre-trained on the ImageNet dataset [26]. We adopt Adam optimizer with $\beta_1 = 0.5$ and the learning rate of 0.0002 during training. The weights of λ_{rec} , λ_{id} , λ_{per} , λ_{au} , λ_{adv} are 1.5, 1.5, 0.07, 0.02, 0.002 respectively. We use \mathcal{L}_{au} to fine-tune the Audio-to-AU module and the AU classifier during training, and adopt Adam optimizer with the learning rate of $1e-6$ and $1e-7$, respectively. We first train our network without frame discriminator for 20 epochs, then add it to fine-tune the network for another 20 epochs.

4.3 Quantitative results

To evaluate the quality of generated images, we adopt the reconstruction metrics Peak SNR (PSNR) and Structural Similarity (SSIM) [33]. For the lip-sync performance, we verify the recognition accuracy and F1 score of the five selected speech-related AUs in generated frames. Specifically, we use the OpenFace toolkit [11, 12] to detect the state of the five selected AUs (activated or not) in each generated frame, then compare them with ground truth labels.

Method	GRID				TCD-TIMIT			
	PSNR \uparrow	SSIM \uparrow	Avg. F1	Avg. Acc.	PSNR \uparrow	SSIM \uparrow	Avg. F1	Avg. Acc.
CRAN[24]	28.041	0.694	0.710	78.71	24.381	0.650	0.562	81.41
Speech2Vid[14]	28.863	0.755	0.738	78.97	25.132	0.685	0.556	81.91
Baseline (\mathcal{L}_{rec})	28.681	0.746	0.725	77.77	25.003	0.717	0.545	80.96
Proposed method	29.838	0.769	0.751	80.92	26.201	0.745	0.590	84.92

Table 1: Quantitative results on the GRID test set and TCD-TIMIT test set. Avg. F1 and Avg. Acc. are average F1 score and average accuracy (%) of speech-related AUs respectively.

Table 1 shows the quantitative results on the GRID test set and TCD-TIMIT test set. To compare with other methods, we implement the methods proposed by Song et al. [29] and Jamaludin et al. [14] in the same conditions, and use the same training-testing data split as our proposed method. Our baseline only uses reconstruction loss. We can see that our proposed method has significant improvements in image quality and lip-sync accuracy in both datasets. Compared with Song et al. [29], our proposed method improves PSNR by about 1.8 and average F1 score of AUs by 4.1% on the GRID dataset. Our method also improves PSNR by 1.82 and average accuracy of AUs by about 3.5% on the TCD-TIMIT dataset. Similarly, our method is significantly higher than Jamaludin et al. [14] in all metrics.

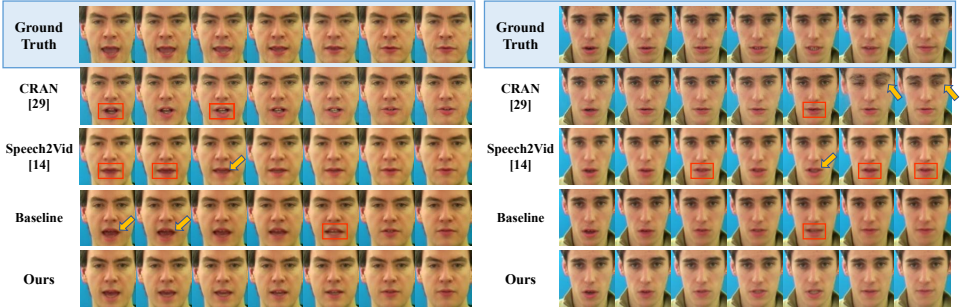


Figure 3: Qualitative results produced by our proposed method and other methods on the GRID test set. Each speaker is a continuous sequence of talking faces. Red boxes are used to mark the examples of lip out of sync, and arrows indicate the examples of blurring.

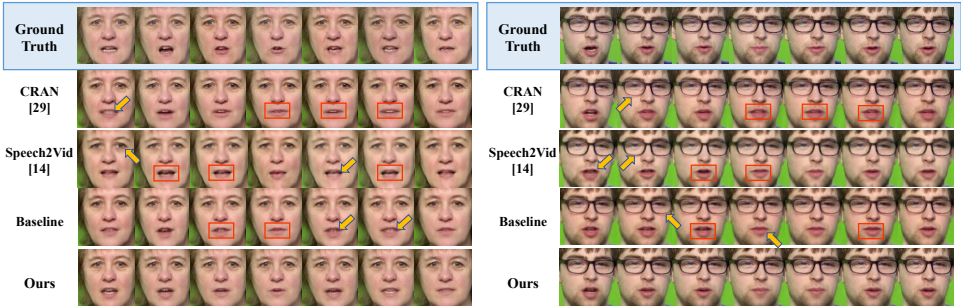


Figure 4: Qualitative results were produced by our proposed method and other methods on the TCD-TIMIT test set.

4.4 Qualitative results

The qualitative results on the GRID test set are shown in Fig. 3. We use two different audio sequences to generate talking head for two speakers. Compared with others, our method performs better in both image quality and lip-sync accuracy. For example, the images generated by Song et al. [29] has the problem of lip out of sync, such as the mouths marked with red boxes in Fig. 3. Besides, there are subtle artifacts in the eye region indicated by the arrows on the right picture, and the generated faces are blurry. Similarly, our baseline and Jamaludin

et al. [24] also has the problem of lip out of sync, and the teeth indicated by the arrows are very blurry. Our proposed method overcomes these shortcomings and performs well.

To further prove the effectiveness of our model, we also show the qualitative results on the TCD-TIMIT test set in Fig. 4. The face images generated by our method are obviously clearer than others. We still use the red boxes to mark the examples of mouth movement out of sync, and use arrows to indicate the examples of blurring, such as teeth. In the left picture, the eyes generated by Song et al. [29] are a bit dull, and the eyes generated by Jamaludin et al. [24] are blurry. Besides, the movement of the mouth generated by Jamaludin et al. [24] is very mismatched. In the right picture, except for the blurred teeth and lips, the arrows also indicate that the glasses generated by other methods are blurry. Our proposed method performs best in both facial texture and mouth movement.

4.5 Ablation study

To quantify the contribution of each component of our method, we conduct an ablation study on the GRID dataset. As shown in Table 2, all the metrics have been improved after adding perceptual loss, identity loss, adversarial loss, and AU loss progressively. Especially after adding identity loss, PSNR has improved significantly. We think that when AU loss is used together with the Audio-to-AU module, it can achieve the maximum effect. Finally, when we add the Audio-to-AU module, which can promote each other with AU classifier, all the metrics have achieved the best results, especially the average F1 score and average accuracy of AUs have been significantly improved. These experiments prove the effectiveness of our proposed method. To observe the results of each AU, we also show the F1 score and accuracy of the five speech-related AUs respectively in Table 2. Most of the results of our proposed model perform best.

Method	PSNR \uparrow	SSIM \uparrow	Avg. F1	Avg. Acc.	AU10		AU14		AU20		AU25		AU26	
					F1	Acc.	F1	Acc.	F1	Acc.	F1	Acc.	F1	Acc.
Baseline (\mathcal{L}_{rec})	28.681	0.746	0.725	77.77	0.745	71.14	0.477	82.01	0.762	74.86	0.939	89.21	0.703	71.65
$\mathcal{L}_{rec,per}$	28.897	0.748	0.732	78.76	0.771	75.99	0.501	82.95	0.742	74.26	0.940	89.10	0.708	71.50
$\mathcal{L}_{rec,per,id}$	29.403	0.753	0.735	79.27	0.757	74.87	0.516	82.20	0.774	79.13	0.943	89.75	0.685	70.41
$\mathcal{L}_{rec,per,id,adv}$	29.528	0.755	0.736	79.55	0.756	74.40	0.511	83.30	0.788	80.37	0.945	90.15	0.682	69.51
$\mathcal{L}_{rec,per,id,adv,au}$	29.618	0.764	0.740	80.20	0.761	75.71	0.481	83.80	0.800	80.70	0.946	90.17	0.710	70.63
Full model(+Audio2AU)	29.838	0.769	0.751	80.92	0.763	76.37	0.531	83.96	0.805	81.64	0.949	90.79	0.707	71.83

Table 2: Quantitative ablation study on the GRID test set. \mathcal{L}_{rec} , \mathcal{L}_{per} , \mathcal{L}_{id} , \mathcal{L}_{adv} , \mathcal{L}_{au} are reconstruction loss, perceptual loss, identity loss, adversarial loss and AU loss, respectively. Avg. F1 and Avg. Acc. are average F1 score and average accuracy (%) of AUs respectively.

In Fig. 5, we also show the visualized examples after adding different components on the GRID test set. Our baseline only uses reconstruction loss, and the generated images are blurry. After adding perceptual loss and identity loss, the image quality is improved, such as facial texture and teeth indicated by the arrows. With the addition of adversarial loss, the facial texture and teeth are further refined. After adding AU loss, lip synchronization has been improved. Finally, the full model with the Audio-to-AU module achieves the best results in both image quality and lip-sync accuracy. We also use red boxes to mark some examples where the lips are not very synchronized in Fig. 5.

4.6 User study

To subjectively evaluate our model, we further conduct a user study. 15 participants are required to compare the 30 videos generated by Song et al. [29], Jamaludin et al. [24]

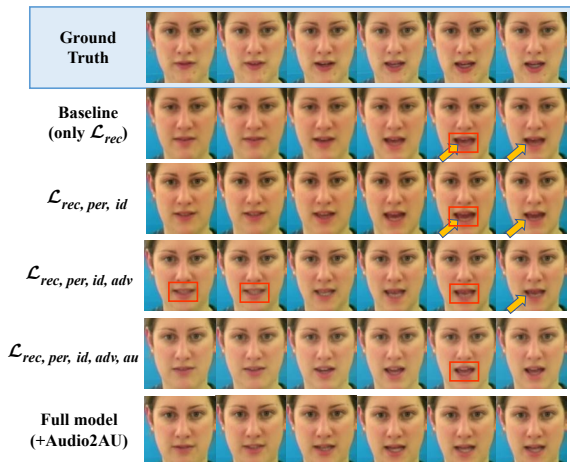


Figure 5: Qualitative results of ablation study on the GRID test set. Each row is a continuous sequence of talking faces. Audio2AU means our proposed Audio-to-AU module.

Method	Lip-sync Accuracy	Image Quality	Video Fluency
Ours vs. CRAN [49]	0.77 / 0.23	0.73 / 0.27	0.92 / 0.08
Ours vs. Speech2Vid [42]	0.70 / 0.30	0.69 / 0.31	0.86 / 0.14

Table 3: User study on the quality of videos generated by different methods. All the users are required to compare the videos generated by our method and others.

and our method. The identity images and audios are randomly selected from the GRID and TCD-TIMIT datasets. Users are required to rank videos generated by different methods from the following three perspectives: (1) Lip-sync accuracy; (2) Image quality; and (3) Video fluency, whether the video jitters. We summarize the evaluation results in Table 3.

As shown in Table 3, our method is superior to the methods of Song et al. [49] and Jamaludin et al. [42] in all aspects, especially in video fluency. The results of the user study further verify the effectiveness of our method in both image quality and lip-sync accuracy.

5 Conclusion

In this paper, we propose a novel talking head generation system, which uses both audio and speech-related facial action units (AUs) as driving information. The proposed Audio-to-AU module is used to obtain speech-related AU information. AU classifier is used to supervise that the generated frames contain correct AU information. We also utilize frame discriminator for adversarial training to improve the realism of generated frames. We conduct extensive experiments for quantitative and qualitative evaluation, and we use the ablation study to verify each component’s contribution to our model. The experimental results on the GRID dataset and TCD-TIMIT dataset demonstrate the superiority over state-of-the-arts in both image quality and lip-sync accuracy. In the future, we will focus on multimodal representation fusion to promote complementation.

Acknowledgments

This work is supported by the National Natural Science Foundation of China under Grants of 61503277 and 61771333.

References

- [1] Tadas Baltrušaitis, Marwa Mahmoud, and Peter Robinson. Cross-dataset learning and person-specific normalisation for automatic action unit detection. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, volume 6, pages 1–6. IEEE, 2015.
- [2] Tadas Baltrušaitis, Amir Zadeh, Yao Chong Lim, and Louis-Philippe Morency. Openface 2.0: Facial behavior analysis toolkit. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 59–66. IEEE, 2018.
- [3] Christoph Bregler, Michele Covell, and Malcolm Slaney. Video rewrite: Driving visual speech with audio. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, pages 353–360, 1997.
- [4] 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 (CVPR)*, pages 7832–7841, 2019.
- [5] Joon Son Chung, Amir Jamaludin, and Andrew Zisserman. You said that? In *British Machine Vision Conference (BMVC)*, pages 1–12, 2017.
- [6] Martin Cooke, Jon Barker, Stuart Cunningham, and Xu Shao. An audio-visual corpus for speech perception and automatic speech recognition. *The Journal of the Acoustical Society of America*, 120(5):2421–2424, 2006.
- [7] P. Ekman and W. Friesen. Facial action coding system: A technique for the measurement of facial movement. Consulting Psychologists Press, 1978.
- [8] Sefik Emre Eskimez, Ross K Maddox, Chenliang Xu, and Zhiyao Duan. End-to-end generation of talking faces from noisy speech. In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1948–1952. IEEE, 2020.
- [9] Sefik Emre Eskimez, You Zhang, and Zhiyao Duan. Speech driven talking face generation from a single image and an emotion condition. *IEEE Transactions on Multimedia*, 2021.
- [10] Bo Fan, Lijuan Wang, Frank K Soong, and Lei Xie. Photo-real talking head with deep bidirectional lstm. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4884–4888, 2015.
- [11] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2672–2680, 2014.

- [12] Naomi Harte and Eoin Gillen. TCD-TIMIT: An audio-visual corpus of continuous speech. *IEEE Transactions on Multimedia*, 17(5):603–615, 2015.
- [13] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1125–1134, 2017.
- [14] Amir Jamaludin, Joon Son Chung, and Andrew Zisserman. You said that?: Synthesising talking faces from audio. *International Journal of Computer Vision*, 127(11): 1767–1779, 2019.
- [15] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision (ECCV)*, pages 694–711, 2016.
- [16] Davis E King. Dlib-ml: A machine learning toolkit. *The Journal of Machine Learning Research*, 10:1755–1758, July 2009.
- [17] Le Li, Zhilei Liu, and Cuicui Zhang. Face inpainting with dynamic structural information of facial action units. In *International Conference on Image and Graphics*, pages 668–679, 2019.
- [18] Zhilei Liu, Diyi Liu, and Yunpeng Wu. Region based adversarial synthesis of facial action units. In *International Conference on Multimedia Modeling*, pages 514–526, 2020.
- [19] Brais Martinez, Michel F Valstar, Bihan Jiang, and Maja Pantic. Automatic analysis of facial actions: A survey. *IEEE Transactions on Affective Computing*, 10(3):325–347, 2017.
- [20] Zibo Meng, Shizhong Han, and Yan Tong. Listen to your face: Inferring facial action units from audio channel. *IEEE Transactions on Affective Computing*, 10(4):537–551, 2017.
- [21] Fausto Milletari, Nassir Navab, and Seyed Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 565–571. IEEE, 2016.
- [22] 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.
- [23] Albert Pumarola, Antonio Agudo, Aleix M Martinez, Alberto Sanfeliu, and Francesc Moreno-Noguer. Ganimation: Anatomically-aware facial animation from a single image. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 818–833, 2018.
- [24] Fabien Ringeval, Erik Marchi, Marc Mehu, Klaus Scherer, and Björn Schuller. Face reading from speech—predicting facial action units from audio cues. In *Sixteenth Annual Conference of the International Speech Communication Association (INTER-SPEECH)*, pages 1977–1981, 2015.

- [25] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 234–241, 2015.
- [26] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [27] Zhiwen Shao, Zhilei Liu, Jianfei Cai, and Lizhuang Ma. Deep adaptive attention for joint facial action unit detection and face alignment. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 705–720, 2018.
- [28] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [29] Yang Song, Jingwen Zhu, Dawei Li, Andy Wang, and Hairong Qi. Talking face generation by conditional recurrent adversarial network. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 919–925, 2019.
- [30] Supasorn Suwajanakorn, Steven M. Seitz, and Ira Kemelmacher-Shlizerman. Synthesizing obama: learning lip sync from audio. *ACM Trans. Graph.*, 36(4):95:1–95:13, 2017.
- [31] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. End-to-end speech-driven facial animation with temporal gans. In *British Machine Vision Conference 2018 (BMVC)*, 2018.
- [32] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Realistic speech-driven facial animation with GANs. *International Journal of Computer Vision (IJCV)*, pages 1–16, 2019.
- [33] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [34] Lei Xie and Zhiqiang Liu. A coupled HMM approach to video-realistic speech animation. *Pattern Recognition*, 40(8):2325–2340, 2007.
- [35] Eli Yamamoto, Satoshi Nakamura, and Kiyohiro Shikano. Lip movement synthesis from speech based on hidden markov models. *Speech Communication*, 26(1-2):105–115, 1998.
- [36] Hang Zhou, Yu Liu, Ziwei Liu, Ping Luo, and Xiaogang Wang. Talking face generation by adversarially disentangled audio-visual representation. In *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI)*, pages 9299–9306, 2019.
- [37] 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. *arXiv preprint arXiv:2104.11116*, 2021.

- [38] Yang Zhou, Xintong Han, Eli Shechtman, Jose Echevarria, Evangelos Kalogerakis, and Dingzeyu Li. Makelttalk: speaker-aware talking-head animation. *ACM Transactions on Graphics (TOG)*, 39(6):1–15, 2020.
- [39] Yuqian Zhou and Bertram Emil Shi. Photorealistic facial expression synthesis by the conditional difference adversarial autoencoder. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 370–376. IEEE, 2017.