

Seeing What You Said: Talking Face Generation Guided by a Lip Reading Expert

Jiadong Wang^{1*}, Xinyuan Qian^{2†}, Malu Zhang³, Robby T. Tan¹, Haizhou Li^{4,1}

¹National University of Singapore, ²University of Science and Technology Beijing

³University of Electronic Science and Technology of China

⁴SRIBD, School of Data Science, The Chinese University of Hong Kong, Shenzhen, China

Abstract

*Talking face generation, also known as speech-to-lip generation, reconstructs facial motions concerning lips given coherent speech input. The previous studies revealed the importance of lip-speech synchronization and visual quality. Despite much progress, they hardly focus on the content of lip movements i.e., the visual intelligibility of the spoken words, which is an important aspect of generation quality. To address the problem, we propose using a lip-reading expert to improve the intelligibility of the generated lip regions by penalizing the incorrect generation results. Moreover, to compensate for data scarcity, we train the lip-reading expert in an audio-visual self-supervised manner. With a lip-reading expert, we propose a novel contrastive learning to enhance lip-speech synchronization, and a transformer to encode audio synchronically with video, while considering global temporal dependency of audio. For evaluation, we propose a new strategy with two different lip-reading experts to measure intelligibility of the generated videos. Rigorous experiments show that our proposal is superior to other State-of-the-art (SOTA) methods, such as Wav2Lip, in reading intelligibility i.e., over **38%** Word Error Rate (WER) on LRS2 dataset and **27.8%** accuracy on LRW dataset. We also achieve the SOTA performance in lip-speech synchronization and comparable performances in visual quality.*

1. Introduction

Talking Face Generation (TFG) aims at generating high-fidelity talking heads which are temporally synchronized with the input speech. It plays a significant role in many Human Robot Interaction (HRI) applications, such as film dubbing [1], video editing, face animation [2, 3], and communication with people who has hearing loss but master in lip-reading. Thanks to its various practical usage, TFG has

also received an increasing attention in both industrial and research community over the past decades [4, 5].

In TFG, there are two major aspects of concern: lip-speech synchronization and visual quality. While humans are sensitive to subtle abnormalities in asynchronized speech and facial motions [6], the mechanism of speech production highly relies on lip movements [7]. As a result, the main challenge of TFG exists in temporal alignment between the input speech and synthesized video streams. One solution to this problem is to place an auxiliary embedding network at the end of the generator to analyze the audio-visual coherence. For example, [8] uses a pre-trained embedding network [9] as the lip-sync discriminator, while [10] investigates an asymmetric mutual information estimator. Another solution is to compute a sync loss between visual lip features of ground truth and generated video sequences [11]. Other attempts use the encoder-decoder structure to facilitate TFG by improving audio and visual representations, such as [12, 13], which disentangle the visual features to enhance audio representations into a shared latent space; or, [14], which disentangles audio factors, i.e., emotional and phonetic content, to remove sync-irrelevant features.

Apart from lip-speech synchronization, blurry or unrealistic visual quality also penalizes generation performance. To preserve defining facial features of target persons, skip connections [15] are applied [5]. Others [12, 16] employ Generative Adversarial Nets (GAN) to distinguish real and synthesized results by modelling the temporal dynamics of visual outputs. In this way, the resulting models can generate more plausible talking heads that can be qualitatively measured by subjective evaluations.

In addition, reading intelligibility should be indispensable, but it has not been emphasized. Reading intelligibility indicates how much text content can be interpreted from face videos by humans' lip reading ability, which is especially significant for hearing-impaired users. However, image quality and lip-speech synchronization do not explicitly reflect reading intelligibility. Specifically, a well-qualified

* jiadong.wang@u.nus.edu

† corresponding author (qianxy@ustb.edu.cn)

image may contain fine-grained lip-sync errors [8] while precise synchronization may convey incorrect text contents. According to the McGurk effect [17], when people listen and see an unpaired but synchronized sequence of speech and lip movements, they may recognize a phoneme from audio or video, or a fused artifact.

In this paper, we propose a TalkLip net to synthesize talking faces by focusing on reading intelligibility. Specifically, we employ a lip-reading expert which transcribes image sequences to text to penalize the generator for incorrectly generated face images. We replace some images of a face sequence with the generated ones, and feed it to the lip-reading expert during training to supervise the face generator.

However, lip reading is hard even for humans. In [18], four people with equal gender distribution are invited to read lip movements. However, the average error rate is as high as 47%. Therefore, a reliable lip-reading model relies on a great amount of data. We employ AV-Hubert [19], a self-supervised method, which has yielded SOTA performance in lip-reading, speech recognition, and audio-visual speech recognition. The encoders of lip-reading and speech recognition systems are highly synchronized since they are supervised by the same pseudo label during pre-training.

Leveraging the lip-reading expert from the AV-Hubert, we propose a new method to enhance lip-speech synchronization. Particularly, we conduct contrastive learning between audio embeddings for face generation and visual context features from the lip-reading encoder. Besides, the AV-Hubert also provides a synchronized speech recognition system whose encoder considers long-term temporal dependency, we adopt this encoder to encode audio inputs, instead of encoders in [8, 11, 13] only rely on short-term temporal dependency (0.2s audio), or the single-modality (audio) pre-trained encoder in [20]. Our contributions are summarized as follows:

- We tackle the reading intelligibility problem of speech-driven talking face generation by leveraging a lip-reading expert.
- To enhance lip-speech synchronization, we propose a novel cross-modal contrastive learning strategy, assisted by a lip-reading expert.
- We employ a transformer encoder trained synchronically with the lip-reading expert to consider global temporal dependency across the entire audio utterance.
- We propose a new strategy to evaluate reading intelligibility for TFG and make the benchmark code publicly available*.

*Code link: <https://github.com/Sxjdwang/TalkLip>

- Extensive experiments demonstrate the feasibility of our proposal and its superiority over other prevailing methods in reading intelligibility (over 38% WER on LRS and 27.8% accuracy on LRW). Additionally, our approach performs comparably to or better than other SOTA methods in terms of visual quality and lip-speech synchronization.

2. Related work

2.1. Speech-driven talking face generation

TFG was first studied in the 1990s [21] which aims at mapping acoustic features to time-aligned realistic facial motions. Traditional methods employ hidden Markov models (HMM) [22], such as [23] predicts facial motions through HMM learnt from time-aligned audio-visual streams with the minimized entropy. Recently, deep learning achieved great success which also brought benefits to the TFG task. Indeed, the majority of those TFG works can be categorized into intermediate representation-based and reconstruction-based methods [11], as elaborated next.

Intermediate representation-based TFG methods include two cascaded modules which learn intermediate facial representations i.e., landmarks or 3D meshes for face synthesis from the coherent speech input. For example, in [16], facial landmarks are predicted from input speeches which are then used to generate face videos conditioned on a reference image. In [24], landmarks dependent on speech content and speaker identity are disentangled to generate faces of unseen subjects. In [25], 3D facial meshes are reconstructed from the extracted expression, geometry and pose parameters to facilitate TFG. However, these methods strongly depend on the quality of 3D face model construction, and are hard to be applied to arbitrary identities [12]. Thus has limited generalization capability. Moreover, fine details i.e., teeth regions are difficult to be captured using 3D models [26] which may degrade the generation performance.

Reconstruction-based methods mostly adopt the encoder-decoder structure to generate talking face videos end-to-end. In [8], a pre-trained discriminator conditioned on the lip-speech sync is employed which emphasizes its impact on the sync quality. By disentangling speech- and identity-related factors from video, [12] learns better visual representations through the proposed associative and adversarial training strategy. Differently, [14] improves the audio representations by disentangling phonetic and emotional content from speech to generate better TFG results. Reconstruction-based methods are free of generating intermediate ground truth and contain fine details of mouth [11]. Therefore, we rely on the reconstruction strategy to synthesize photo-realistic talking faces from speech directly.

To be noted, our concept of employing a lip-reading ex-

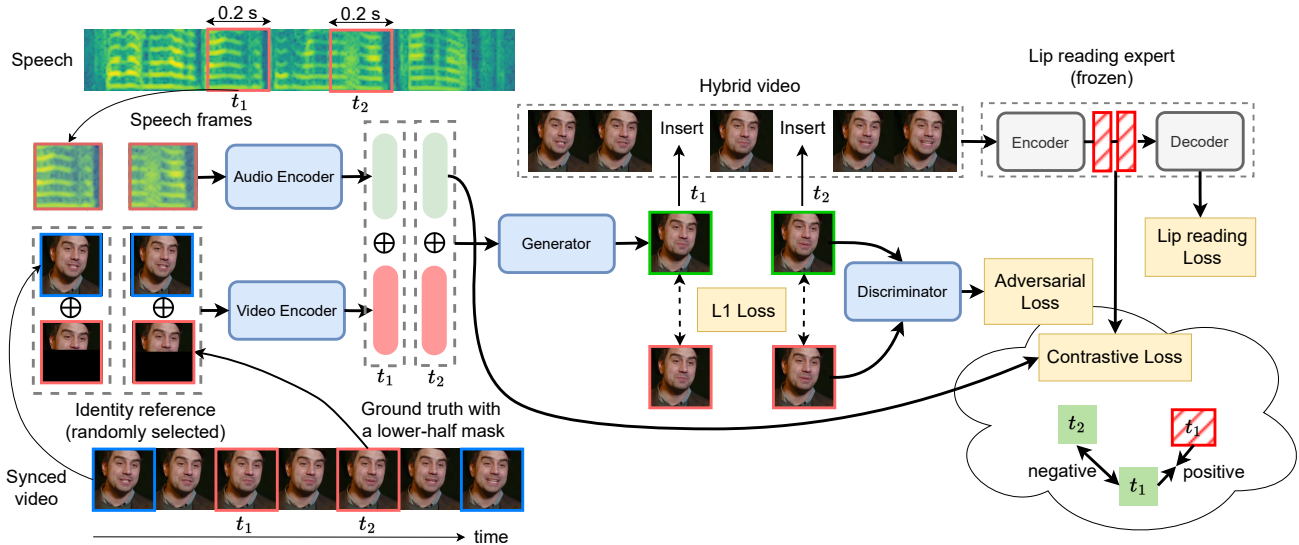


Figure 1. Block diagram of our proposed TalkLip net which produces talking faces given the coherent speech inputs, identity and pose inputs. The network consists of (1) an audio encoder for phoneme-level feature extraction; (2) a video encoder to extract the identity and pose information; (3) a GAN conditioned on the real/synthesized face images and (4) a novel back-end lip-reading expert to penalize the lip-reading error. The whole network is optimized with the four-fold losses, which are marked as yellow rectangles (\oplus denotes pairwise concatenation).

pert to penalize the incorrectly generated images can be easily extended to both categories while we select the reconstruction strategy to set an example.

2.2. Contrastive learning

Contrastive learning is a strategy to obtain more representative features and has widespread usage in various tasks, such as text-image retrieval [27], image classification [28] and acoustic event detection [29]. Its core idea is minimizing the distance between samples of the same class while maximizing the ones from different classes. To this end, one solution is by mapping all features (either homogeneous or heterogeneous) into a common representation space where *positive* feature pairs are attracted and *negative* pairs are repelled. Indeed, selecting positive and negative pairs is of particular significance, which guides the model to learn contrastive representations. It also promotes self-supervised learning without the requirement of labeled data, such as SimCLR [30] and SwAV [31].

In talking face generation, contrastive learning can assist lip-speech synchronization. Concretely, PC-AVS [12] uses audio features that are timely aligned with the visual lip feature as positive samples and adopt others as negative ones. However, they extract visual lip features only from single images without considering lip movement. Unlike PC-AVS, our contrastive learning focuses on the dynamic property of lip movement and extracts visual features from 5 frames.

3. Proposed Methods

We propose a TalkLip net, illustrated in Fig. 1, to improve the reading intelligibility of the generated talking faces videos. Specifically, the TalkLip net takes images as the identity and pose references while taking the coherent speech as the lip movement reference. Given synthesized videos from the generator, the lip-reading expert is employed to penalize inaccurate lip movements via a lip-reading loss. This will backpropagate gradients to the generator, the audio and video encoders, as described in Sec. 3.5. Furthermore, assisted by the lip-reading expert, we perform contrastive learning between visual context features and audio embeddings to improve lip-speech synchronization, as depicted in Sec. 3.6. Finally, all losses are integrated in Sec. 3.6 to facilitate TFG.

3.1. Overview of lip-reading expert

The goal of a lip-reading system is to transcribe a word sequence (or something like that) from the lip movement in a video. To tackle data scarcity, we adopt an audio-visual self-supervised method, AV-Hubert [19] to have a reliable lip-reading expert.

The self-supervision is shown in Fig. 2. The visual front-end consists of a 3D Convolutional Neural Network (CNN) to capture local lip movement and a 2D ResNet-18 [32]. The modality selector determines the training of audio-only, video-only or audio-visual speech recognition systems by masking features of no-involved modalities. All systems share the same transformer encoder and are regressed by

the same frame-by-frame pseudo label, which is the clustering of Mel-Frequency Cepstral Coefficients (MFCC) or hidden audio-visual representation in the transformer encoder. Thus, audio and visual context features are synchronized.

Afterwards, a lip-reading network is constructed by combing the visual frontend and the transformer encoder pre-trained in self-supervision with a randomly initialized transformer decoder. This lip-reading network is fine-tuned under the supervision of text transcription. Once the fine-tuning is complete, the lip-reading network is frozen and serves as an expert in the training of talking face generation, as shown in Fig. 1.

3.2. Audio encoder

The audio encoder is to encode phoneme-level embeddings and provides the embeddings to the generator as the reference for mouth shape and lip movement.

We use two different audio encoders to extract the embeddings. We call them *local* and *global* audio embeddings. The local audio embedding is extracted using a CNN-based network. A 0.2-second audio segment whose centre is synchronized with the pose reference is fed as the input, as shown in Fig. 1. However, a 0.2-second audio segment does not contain temporal variations of the entire speech. The transformer encoder [33] is a popular method to address this limitation in some speech tasks [34, 35]. Herein, we employ a pre-trained transformer encoder as mentioned in Sec. 3.1, which is jointly trained with the lip-reading expert, as shown in Fig. 2. The transformer encoder takes the entire speech as input to generate audio context features of all frames. Then we choose a frame of the context features, which are time-aligned with the pose reference as the global audio embedding, as depicted in Fig. 2.

3.3. Video encoder

The video encoder is to extract identity and pose information from images to a united visual embedding and supplies the embedding to the generator to synthesize an image consistent with provided identity and pose.

The visual embedding is extracted from two images: an identity reference and a pose reference. The pose reference is the same as the target face image, except it masks the lower-half face to prevent our TalkLip net from learning the movements in the lip region. The identity reference is randomly picked within the same video sequence. Two references are concatenated in the channel dimension as visual input and fed to a CNN-based video encoder.

3.4. Video generation

The generator constructs a talking face image given an audio embedding and a video embedding. As Unet [15] have a great trade-off between a reconstruction effect and computational efficiency [36], we adopt a skip connection

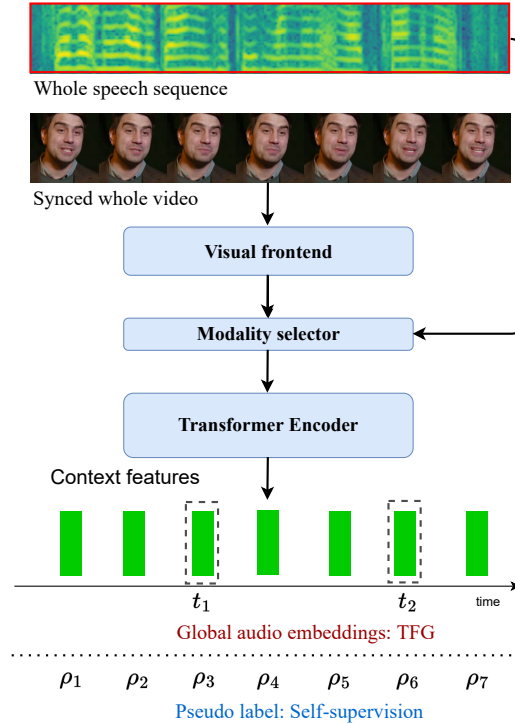


Figure 2. The self-supervised training of AV-Hubert and the process of capturing *global* audio embeddings using the pre-trained transformer encoder. ρ is frame-level pseudo labels for training.

like Unet between the video encoder and the generator, which is employed in [8] as well. The generator consists of transposed CNN blocks.

Given a synthesized talking face image and a ground truth image, two generation losses are applied to improve the video generation quality i.e., the reconstruction loss as Eq. (1) and the GAN loss as Eq. (2), Eq. (3).

We adopt L1 as the reconstruction loss between the synthesized and real face images due to its better convergence to a local minimum than that of L2 loss [37]. Given N_i pairs of synthesized faces v' and ground truth faces v , the reconstruction loss is formulated as:

$$\mathcal{L}_{\text{rec}} = \frac{1}{N_i} \sum_i |v_i - v'_i| \quad (1)$$

To increase the realism of synthesized talking face videos, many works adopt GAN [11, 38]. Considering its outstanding performance, in this paper, a GAN loss is also applied. Let us denote G and D as the generator and discriminator in Fig. 1, respectively. To leverage GANs [39], two losses are involved as follows:

$$\mathcal{L}_{\text{gen}} = \mathbb{E}[\log(1 - D(v'))] \quad (2)$$

$$\mathcal{L}_{\text{disc}} = \mathbb{E}[\log(1 - D(v))] + \mathbb{E}[\log(D(v'))] \quad (3)$$

where $\mathcal{L}_{\text{disc}}$ optimizes the discriminator D to distinguish synthesized face image v' from ground truth faces v , while

\mathcal{L}_{gen} improves the quality of synthesized faces v' to fool the discriminator.

3.5. Lip reading loss

The lip-reading expert is applied on hybrid videos i.e., the ground truth videos with some frames substituted by corresponding synthesized images. Therefore, the gradients on generated images will further back-propagate to optimize the generator, audio and video encoder.

Let us denote $V = [v_1, v'_2, \dots, v_T] \in \mathbb{R}^{T \times 3 \times H \times W}$ as a hybrid video, where v and v' are the real and synthesized images, respectively, and T, H, W are the number of frames, height, and width. We follow the masking strategy in [40] to choose which frames of v are replaced by v' . In particular, we randomly sample a proportion p of all frames as a set of starting frames whose subsequent M frames will be replaced with v' . In this way, lip movements within v' and between v and v' are considered. We assume $Y \in \mathbb{R}^{L \times C}$ as the text content of the ground truth video, where L and C denote the text length and the output class number.

The lip-reading expert [19] contains (1) a 3D convolutional layer followed by a ResNet-18 [41] to extract features of lip movement; (2) a transformer encoder to produce a context feature $R \in \mathbb{R}^{T \times f}$ by calculating global temporal dependency where f is the feature dimension and (3) a transformer decoder to predict the text $\hat{Y} \in \mathbb{R}^{L \times C}$. Note that the lip-reading expert is frozen. Finally, we use cross-entropy to compute the lip-reading loss:

$$\mathcal{L}_{\text{lip}} = -Y \log P(\hat{Y}|V) \quad (4)$$

3.6. Contrastive loss

We use the contrastive strategy to enhance lip-speech synchronization by attracting audio embeddings and their time-aligned visual context features while repelling the audio embeddings from different frames. To achieve such a goal, we select infoNCE [42].

Since the lip-reading expert is frozen, we do not use visual context features as anchors. Instead, given audio embeddings E^a from the audio encoder and visual context features R from the lip-reading expert, the contrastive loss is computed as:

$$\theta(x, x') = \exp(\mathcal{F}(x) \cdot \mathcal{F}(x') / \tau), \quad (5)$$

$$\mathcal{L}_{\text{av}} = - \sum_{i \in \Upsilon} \log \frac{\theta(E_i^a, R_i)}{\theta(E_i^a, R_i) + \sum_{\substack{j \in \Upsilon \\ i \neq j}} \theta(E_i^a, E_j^a)} \quad (6)$$

where Υ denotes all masked frames described in Sec. 3.5. We follow [43] to employ Eq. (5) to calculate the similarity of the two features, where τ is a pre-defined temperature parameter, \mathcal{F} is a linear layer, and x and x' generally indicate two feature vectors. Besides, x and x' before and after \mathcal{F} are normalized.

Finally, we combine all the aforementioned losses to optimize the network:

$$\mathcal{L}_t = \lambda_l \cdot \mathcal{L}_{\text{lip}} + \lambda_c \cdot \mathcal{L}_{\text{av}} + \lambda_r \cdot \mathcal{L}_{\text{rec}} + \lambda_g \cdot \mathcal{L}_{\text{gan}} \quad (7)$$

where $\lambda_l, \lambda_c, \lambda_r$ and λ_g scale the contributions of different loss items. \mathcal{L}_{gan} is the sum of \mathcal{L}_{gen} and $\mathcal{L}_{\text{disc}}$.

4. Experiments

4.1. Experimental settings

Dataset We train the proposed method in the train set (29 hours) of LRS2 and evaluate it with other competitive approaches in the test sets of two publicly available datasets: LRW [44], and LRS2 [45]. LRW is a dataset for audio-visual word classification. Each sample in the LRW contains a video with 29 frames (about 1 second), synchronized audio, and a target word in the middle of the sample. There are 500 different words in the whole dataset. LRS2 is an open-world audio-visual speech recognition dataset. Each sample in it consists of video, audio and target text sequence. There are two subsets for training with a 60k vocabulary: pretrain (195 hours), and train (29 hours). Audio and video sample rates of both datasets are 16 KHz and 25 fps, respectively.

Metrics Comparisons are evaluated on the following widely used metrics: PSNR [11, 46], SSIM [38, 47], LSE-C [8, 13]. PSNR and SSIM are metrics to measure visual quality, which are calculated by comparing generated images with ground truth. LSE-C is proposed in [8] to quantify lip-speech synchronization through a pre-trained syncnet [9]. The higher their scores, the better their performances.

Additionally, we propose a new strategy to evaluate reading intelligibility, including the Word Error Rate (WER) of synthesized video on LRS2, and the Word Accuracy (ACC) on the unseen dataset LRW. To evaluate the WER, we employ the AV-Hubert (large) [19], which achieves SOTA performance on the lip-reading task. We fine-tune it on the whole LRS2 dataset as one lip-reading observer. Its performance on the test set of LRS2 is 23.82% WER. Note that the lip-reading expert in our proposed method is different from the lip-reading observer since the lip-reading expert is AV-Hubert (base). Besides, To further increase the difference between lip-reading expert and observer, a lip-reading network adopting Conformer [48, 49] serves as another lip-reading observer. We run Conformer in teacher-forcing mode [50] to calculate WER. In other words, ground truth text is involved in inference. For measuring ACC on LRW, we choose MVM [51], whose prediction accuracy on LRW is 88.5%.

Comparison methods Five SOTA methods of talking face generation are used to compare: ATVGnet [16], Wav2Lip [8], Faceformer [20], PC-AVS [13],

SyncTalkFace [11]. ATVGnet belongs to the intermediate representation-based method, which first predicts a face landmark and then generates a realistic image. Wav2Lip is a representative of reconstruction methods and claims SOTA performance in lip sync synchronization. Faceformer represents an emerging manner, namely sequence-to-sequence manner, which employs a transformer to take the whole audio and output a video sequence. Originally, in [20], this method is applied to generate a vertex sequence (3D scan data) from an audio sequence. It does not work if we simply replace the vertex sequence by the image sequence. Thus, we adopt the video encoder and corresponding visual input to replace the style encoder in [20] and add a skip connection between the visual input and the output image. PC-AVS is also a reconstruction method, which disentangles identity, speech content, and poses. SyncTalkFace addresses lip-speech synchronization via a memory network. The last two methods do not publicly share their code or training code. Thus, their performances are collected from [11].

Implementation details Below preprocessing is applied on raw audio and videos. Following Wav2Lip [8], we detect faces on each image as Region of interest (ROI) and then crop and resize the ROI to 96×96 . Audio wavforms are preprocessed to mel-spectrogram with hop and window lengths, and mel bins are 12.5 ms, 50 ms, and 80.

We only fine-tune the last 3 transformer blocks of the pre-trained transformer audio encoder during the TFG training. The detailed architectures of our method can be found in the supplementary. During inference, we always use the first image of a video as the identity reference.

4.2. Experimental result

Quantitative results Performance comparisons on visual quality, lip-speech sync and reading intelligibility are shown in Tab. 1. LRS2 and unseen LRW are involved. With the help of the global audio encoder and contrastive learning, the TalkLip ($g + c$) achieves the SOTA performance on lip-speech sync. Although Wav2Lip is superior to other TalkLip nets in lip-speech sync (LSE-C), its performance in reading intelligibility (WER_1 and WER_2) on LRS2 lags behind all the TalkLip nets. This suggests that improving lip-speech sync does not lead to better reading intelligibility.

On LRS2, it is observed that our methods outperform others in reading intelligibility while they are close to or even better than real videos. Such advantage remains in the unseen LRW dataset, which suggests that our method is robust across different datasets and lip-reading tasks e.g., lip-reading of a continuous utterance and a single word. We note that the TalkLip ($g + c$) even has better reading intelligibility than the ground truth. We hypothesize that the lip-reading expert is trained to capture relative standard and full lip movements. And the audio transformer encoder is

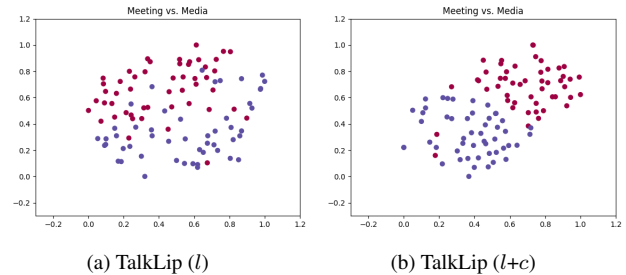


Figure 3. The t-SNE visualization of audio embeddings correspond to ‘Meeting’ (red) and ‘Media’ (blue).

trained well to extract phoneme-level information. Thus TalkLip nets learn from the lip-reading expert and tend to generate more standard lip movement than natural lip movements. Note that WER_1 of some methods are higher than 100%. The reason is that the lip-reading observer of WER_1 does not know the length of the target text and may predict a much longer sentence. In this case, the edit distance algorithm can output a WER exceeding 100%. Differently, the teacher-forcing mode knows the length so that $WER_2 < 100\%$.

PSNR and SSIM show that the visual quality of our methods is worse than SyncTalkFace. We note that SyncTalkFace is trained on the whole LRS2 dataset, whose training data is 8 times more than ours. Despite the limited training data, our results are close to the second-best method, Wav2Lip.

Qualitative results To qualitatively evaluate different methods, we provide the uniformly sampled snapshots of two generated talking face videos in Fig. 4. Specifically, the ground truth video is provided in the first row where synthesized images of different methods follow the next. From Fig. 4 (a), we can see that our proposal provides the most similar image frames as the ground truth video with smooth variations in mouth shape. For Fig. 4 (b), we are also of the best performance in video quality and the naturalness in mouth movements.

4.3. User study

We invite 15 people to evaluate the generated talking face videos of our method with the other 3 SOTA methods. 6 videos from LRS2 and 9 videos from LRW are involved in this user study with human measures in visual quality, speech-lip sync and realness. Each metric ranges from 1 to 5. Moreover, to evaluate reading intelligibility, we randomly collect 5 videos from LRW for participants to score the reading intelligibility of the corresponding ground-truth words. We do not use samples from LRS2 since lip reading on continuous long videos is hard for people without professional lip-reading training. The results from all participants are averaged and listed in Tab. 2.

From the results, we can see that our proposal achieves

Table 1. Quantitative results on LRW and LRS2 (N.A. indicates the result is not applicable). Performances of methods with * are collected from [11] which are trained using the whole LRS2 dataset (224 hours) while our methods are trained by the LRS2 dataset (29 hours). Thus, we can not evaluate their performances in reading intelligibility. WER₁, WER₂ are scored by the AV-Hubert (large) [19] and Conformer [48, 49]. ACC is evaluated by MVM [51]. *g*, *l*, and *c* in parentheses behind **TalkLip** indicate global audio embedding, local audio embedding, and contrastive learning, respectively. **Base** refers to the **TalkLip** (*l*).

Method	LRW				LRS2				
	PSNR \uparrow	SSIM \uparrow	LSE-C \uparrow	ACC (%) \uparrow	PSNR \uparrow	SSIM \uparrow	LSE-C \uparrow	WER ₁ (%)	WER ₂ \downarrow
Ground Truth	N.A.	1.000	6.88	88.51	N.A.	1.000	8.25	23.82	40.9
ATVGnet	30.71	0.791	5.64	18.10	30.42	0.751	5.05	113.69	91.8
Wav2Lip	31.52	0.874	7.18	59.98	31.36	0.854	8.40	82.06	73.9
Faceformer	29.19	0.856	5.58	53.43	29.47	0.840	6.42	97.64	79.0
PC-AVS*	30.44	0.778	6.42	-	29.89	0.747	6.73	-	-
SyncTalkFace*	33.13	0.893	6.62	-	32.59	0.876	7.93	-	-
TalkLip (<i>l</i>)	31.24	0.867	6.44	79.78	31.38	0.849	7.58	45.74	55.7
TalkLip (<i>l</i> + <i>c</i>)	31.52	0.867	6.51	83.17	31.14	0.850	7.76	38.00	49.2
TalkLip (<i>g</i>)	30.78	0.871	7.01	86.57	30.86	0.854	8.38	25.31	36.5
TalkLip (<i>g</i> + <i>c</i>)	31.18	0.866	7.28	87.81	31.19	0.850	8.53	23.43	35.1
Base w.o. \mathcal{L}_{lip}	31.22	0.865	6.01	48.58	31.08	0.852	7.09	103.57	82.2
Base w.o. $\mathcal{L}_{lip, gan}$	30.64	0.864	5.03	30.80	30.70	0.851	5.93	116.26	89.3



(a) ‘Another two hundreds’

(b) ‘More than a third of the world’

Figure 4. Snapshots of the generated talking face videos correspond to a part of the underlined text.

the best performance in all subjective measures. Specifically, for visual quality, the results are consistent with that in Tab. 1. The performance of our methods approximates Wav2Lip and is better than others. For lip-speech synchronization and realism, our methods outperform others by a distinguishable margin. This verifies the capability of our method in generating videos of human naturalness. Finally, for reading intelligibility of the generated face videos,

our proposal obtains the score of 3.76, beats the second-best method i.e., Wav2Lip by 0.38 (11.2%) and the worst method i.e., Faceformer by 0.72 (23.7%) thanks to the assistance of the lip-reading expert.

4.4. Ablation study on the GAN loss

Although the GAN loss is designed to make synthesized videos more realistic, we find it also lead to better

Table 2. User study on face videos generated by different methods in visual quality (VQ), lip-speech sync (Sync), realness (Real), and reading intelligibility (RI).

Method	VQ \uparrow	Sync \uparrow	Real \uparrow	RI \uparrow
ATVGnet [16]	3.27	2.96	2.73	3.20
Wav2Lip [8]	3.75	3.77	3.75	3.38
Faceformer [20]	2.84	2.77	2.56	3.04
TalkLip ($g + c$)	3.78	4.00	3.89	3.76

lip-speech sync and reading intelligibility. Compared with **Base** w.o. $\mathcal{L}_{\text{lip, gan}}$, **Base** w.o. \mathcal{L}_{lip} improves LSE-C by about 19.5 % on LRW and LRS2. Besides, reading intelligibility is also improved, especially on LRW.

4.5. Ablation study on the lip-reading loss

We propose to employ a lip-reading expert to improve reading intelligibility of synthesized talking faces. Compared with the TalkLip (l) without the lip-reading loss, the TalkLip (l) with the lip-reading loss shows at least 26.5% improvement on LRS2 and 31.2% improvement on LRW. These advantages confirm the positive effect of the lip-reading loss on improving reading intelligibility.

4.6. Ablation study on contrastive learning

We propose contrastive learning to enhance lip-speech synchronization. From Tab. 1, the contribution of contrastive learning is proven. LSE-C of the TalkLip ($l + c$) is higher than that of the TalkLip (l) on both LRS2 and LRW datasets. The only difference between them is the former employs contrastive learning. Besides, the TalkLip ($g + c$) also has this advantage over the TalkLip (g).

Moreover, it is observed that reading intelligibility is also improved with contrastive learning. In particular, as shown in Tab. 1, the TalkLip ($l + c$) gains an improvement of at least 6.5% WER on LRS2 and 3.39% ACC on LRW by involving contrastive learning. The improvement may benefit from the fact that contrastive learning pushes audio embeddings from different frames more apart. To validate this guess, we pick a pair of words in LRW and visualize the distributions of audio embeddings using t-SNE [52] in Fig. 3. In detail, LRW provides the start and end times of ground-truth words. Thus, for each sample, we calculate audio embeddings at the ground-truth interval and sum them up as the feature to be visualized. The pair are ‘Meeting’ and ‘Media’ which have a similar prefix.

It is shown that audio embeddings of ‘Meeting’ and ‘Media’ are mixed without contrastive learning, while there are only three outliers with contrastive learning.

With these results, we conclude that contrastive learning is capable of enhancing lip-speech synchronization and repelling audio embeddings at different frames.

4.7. Ablation study on global audio embedding

We propose to use global audio embedding because of two merits: (1) global audio embeddings are calculated from entire utterances by a transformer encoder, which benefits from the global temporal dependency to represent phoneme-level information and (2) the transformer encoder is pre-trained synchronically with the lip-reading expert and therefore its outputs are synchronized with visual content features.

From Tab. 1, the first merit is proven because the TalkLip ($g + c$) outperforms the TalkLip ($l + c$) both in WER on LRS2 and ACC on LRW. It is observed that two TalkLip nets with the global audio encoder outperform those with the local audio encoder on LSE-C, as shown in Tab. 1. This improvement confirms the benefit of audio-visual synchronical pre-training to lip-speech synchronization.

5. Conclusion

In this paper, we analyse the importance of reading intelligibility and propose a novel TalkLip net to synthesize talking face videos with great reading intelligibility. Particularly, a TalkLip net employs a pre-trained lip-reading expert to penalize the incorrect lip-reading predictions from synthesized videos. To evaluate reading intelligibility, we propose a new strategy that contains single-word ACC measurement and WER on videos of continuous utterances and prepares a benchmark code. From the results, TalkLip nets are superior to the other competitive methods on reading intelligibility by a large margin both subjectively and objectively. Besides, based on the lip-reading expert, we propose a novel contrastive learning to enhance lip-speech synchronization, and a synchronically pre-trained audio transformer encoder to leverage global temporal dependency. Our extensive experiments confirm that the contrastive learning is effective to lip-speech synchronization while it also improves the reading intelligibility. Our results also prove that the audio transformer encoder generates better phoneme-level audio embeddings and helps lip-speech synchronization. Therefore, our TalkLip net can synthesize talking face videos with superior reading intelligibility, lip-speech synchronization and competitive visual quality.

Acknowledgement This work is supported by 1) Huawei Noah’s Ark Lab; 2) National Natural Science Foundation of China (Grant No. 62271432); 3) Internal Project Fund from Shenzhen Research Institute of Big Data (Grant No. T00120220002); 4) Guangdong Provincial Key Laboratory of Big Data Computing, The Chinese University of Hong Kong, Shenzhen (Grant No. B10120210117-KP02); 5) Human-Robot Collaborative AI for Advanced Manufacturing and Engineering (Grant No. A18A2b0046), Agency of Science, Technology and Research (A*STAR), Singapore. 6) MOE AcRF Tier, A-0009455-01-00.

References

- [1] H. Kim, P. Garrido, A. Tewari, W. Xu, J. Thies, M. Niessner, P. Pérez, C. Richardt, M. Zollhöfer, and C. Theobalt, “Deep video portraits,” *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 1–14, 2018. [1](#)
- [2] Y. Song, J. Zhu, D. Li, X. Wang, and H. Qi, “Talking face generation by conditional recurrent adversarial network,” *arXiv preprint arXiv:1804.04786*, 2018. [1](#)
- [3] K. Gu, Y. Zhou, and T. Huang, “Flnet: Landmark driven fetching and learning network for faithful talking facial animation synthesis,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, no. 07, 2020, pp. 10 861–10 868. [1](#)
- [4] L. Chen, R. K. Maddox, Z. Duan, and C. Xu, “Hierarchical cross-modal talking face generation with dynamic pixel-wise loss,” in *Proc. of Int. Conf. on Computer Vision and Pattern Recognition*, 2019, pp. 7824–7833. [1](#)
- [5] A. Jamaludin, J. S. Chung, and A. Zisserman, “You said that?: Synthesising talking faces from audio,” *International Journal of Computer Vision*, vol. 127, no. 11, pp. 1767–1779, 2019. [1](#)
- [6] L. Yu, J. Yu, M. Li, and Q. Ling, “Multimodal inputs driven talking face generation with spatial–temporal dependency,” *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 31, no. 1, pp. 203–216, 2020. [1](#)
- [7] Z.-H. Ling, K. Richmond, and J. Yamagishi, “An analysis of hmm-based prediction of articulatory movements,” *Speech Communication*, vol. 52, no. 10, pp. 834–846, 2010. [1](#)
- [8] K. Prajwal, R. Mukhopadhyay, V. P. Nambodiri, and C. 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*, 2020, pp. 484–492. [1](#), [2](#), [4](#), [5](#), [6](#), [8](#)
- [9] J. S. Chung and A. Zisserman, “Out of time: automated lip sync in the wild,” in *ACCV 2016 International Workshops*. Springer, 2017, pp. 251–263. [1](#), [5](#)
- [10] H. Zhu, H. Huang, Y. Li, A. Zheng, and R. He, “Arbitrary talking face generation via attentional audio-visual coherence learning,” *arXiv preprint arXiv:1812.06589*, 2018. [1](#)
- [11] S. J. Park, M. Kim, J. Hong, J. Choi, and Y. M. R., “ync-talk-face: Talking face generation with precise lip-syncing via audio-lip memory,” in *Association for the Advancement of Artificial Intelligence*, 2022, pp. 234–778. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#)
- [12] H. Zhou, Y. Liu, Z. Liu, P. Luo, and X. Wang, “Talking face generation by adversarially disentangled audio-visual representation,” in *Association for the Advancement of Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 9299–9306. [1](#), [2](#), [3](#)
- [13] H. Zhou, Y. Sun, W. Wu, C. C. Loy, X. Wang, and Z. Liu, “Pose-controllable talking face generation by implicitly modularized audio-visual representation,” in *Proc. of Int. Conf. on Computer Vision and Pattern Recognition*, 2021, pp. 4176–4186. [1](#), [2](#), [5](#)
- [14] G. Mittal and B. Wang, “Animating face using disentangled audio representations,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2020, pp. 3290–3298. [1](#), [2](#)
- [15] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241. [1](#), [4](#)
- [16] L. Chen, R. K. Maddox, Z. Duan, and C. Xu, “Hierarchical cross-modal talking face generation with dynamic pixel-wise loss,” in *Proc. of Int. Conf. on Computer Vision and Pattern Recognition*, 2019, pp. 7832–7841. [1](#), [2](#), [5](#), [8](#)
- [17] H. McGurk and J. MacDonald, “Hearing lips and seeing voices,” *Nature*, vol. 264, no. 5588, pp. 746–748, 1976. [2](#)
- [18] A. B. Hassanat, “Visual speech recognition,” *Speech and Language Technologies*, vol. 1, pp. 279–303, 2011. [2](#)
- [19] B. Shi, W.-N. Hsu, K. Lakhota, and A. Mohamed, “Learning audio-visual speech representation by masked multimodal cluster prediction,” *arXiv preprint arXiv:2201.02184*, 2022. [2](#), [3](#), [5](#), [7](#)
- [20] Y. Fan, Z. Lin, J. Saito, W. Wang, and T. Komura, “Faceformer: Speech-driven 3d facial animation with transformers,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 18 770–18 780. [2](#), [5](#), [6](#), [8](#)
- [21] H. Yehia, P. Rubin, and E. Vatikiotis-Bateson, “Quantitative association of vocal-tract and facial behavior,” *Speech Communication*, vol. 26, no. 1-2, pp. 23–43, 1998. [2](#)
- [22] C. Bregler, M. Covell, and M. Slaney, “Video rewrite: Driving visual speech with audio,” in *Proc. of the 24th annual conf. on Computer graphics and interactive techniques*, 1997, pp. 353–360. [2](#)
- [23] M. Brand, “Voice puppetry,” in *Proc. of the 26th annual conf. on Computer graphics and interactive techniques*, 1999, pp. 21–28. [2](#)
- [24] Y. Zhou, X. Han, E. Shechtman, J. Echevarria, E. Kalogerakis, and D. Li, “Makeltalk: speaker-aware talking-head animation,” *ACM Transactions On Graphics (TOG)*, vol. 39, no. 6, pp. 1–15, 2020. [2](#)
- [25] L. Song, W. Wu, C. Qian, R. He, and C. C. Loy, “Everybody’s talkin’: Let me talk as you want,” *IEEE Trans. on Information Forensics and Security*, vol. 17, pp. 585–598, 2022. [2](#)
- [26] Z. Zhang, L. Li, Y. Ding, and C. 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*, 2021, pp. 3661–3670. [2](#)
- [27] M. Li, R. Xu, S. Wang, L. Zhou, X. Lin, C. Zhu, M. Zeng, H. Ji, and S.-F. Chang, “Clip-event: Connecting text and images with event structures,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 420–16 429. [3](#)

- [28] P. Wang, K. Han, X.-S. Wei, L. Zhang, and L. Wang, “Contrastive learning based hybrid networks for long-tailed image classification,” in *Proc. of Int. Conf. on Computer Vision and Pattern Recognition*, 2021, pp. 943–952. 3
- [29] H.-H. Wu, P. Seetharaman, K. Kumar, and J. P. Bello, “Wav2CLIP: Learning robust audio representations from clip,” in *Proc. of IEEE Int. Conf. on Audio, Speech and Signal Processing*, 2022, pp. 4563–4567. 3
- [30] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *Proc. of Int. Conf. on Machine Learning*. PMLR, 2020, pp. 1597–1607. 3
- [31] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, “Unsupervised learning of visual features by contrasting cluster assignments,” *Proc. of Int. Conf. on Neural Information Proc. Systems*, vol. 33, pp. 9912–9924, 2020. 3
- [32] T. Stafylakis and G. Tzimiropoulos, “Combining residual networks with lstms for lipreading,” in *Proc. of the Conference International Speech Communication Association (INTERSPEECH)*, 2017. 3
- [33] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017. 4
- [34] J. Ao, R. Wang, L. Zhou, C. Wang, S. Ren, Y. Wu, S. Liu, T. Ko, Q. Li, Y. Zhang, Z. Wei, Y. Qian, J. Li, and F. Wei, “SpeechT5: Unified-modal encoder-decoder pre-training for spoken language processing,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, 2022, pp. 5723–5738. 4
- [35] R. Tao, Z. Pan, R. K. Das, X. Qian, M. Z. Shou, and H. Li, “Is someone speaking? exploring long-term temporal features for audio-visual active speaker detection,” in *Proceedings of the 29th ACM International Conference on Multimedia*, 2021, pp. 3927–3935. 4
- [36] M. Li, W. Zuo, and D. Zhang, “Deep identity-aware transfer of facial attributes,” *arXiv preprint arXiv:1610.05586*, 2016. 4
- [37] H. Zhao, O. Gallo, I. Frosio, and J. Kautz, “Loss functions for image restoration with neural networks,” *IEEE Transactions on computational imaging*, vol. 3, no. 1, pp. 47–57, 2016. 4
- [38] B. Liang, Y. Pan, Z. Guo, H. Zhou, Z. Hong, X. Han, J. Han, J. Liu, E. Ding, and J. Wang, “Expressive talking head generation with granular audio-visual control,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 3387–3396. 4, 5
- [39] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial networks,” *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020. 4
- [40] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 449–12 460, 2020. 5
- [41] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778. 5
- [42] A. v. d. Oord, Y. Li, and O. Vinyals, “Representation learning with contrastive predictive coding,” *arXiv preprint arXiv:1807.03748*, 2018. 5
- [43] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, “Supervised contrastive learning,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 18 661–18 673, 2020. 5
- [44] J. S. Chung and A. Zisserman, “Lip reading in the wild,” in *Asian conference on computer vision*. Springer, 2016, pp. 87–103. 5
- [45] T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman, “Deep audio-visual speech recognition,” *IEEE transactions on pattern analysis and machine intelligence*, 2018. 5
- [46] K. Vougioukas, S. Petridis, and M. Pantic, “Realistic speech-driven facial animation with gans,” *International Journal of Computer Vision*, vol. 128, no. 5, pp. 1398–1413, 2020. 5
- [47] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: from error visibility to structural similarity,” *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004. 5
- [48] J. Wang, X. Qian, and H. Li, “Predict-and-update network: Audio-visual speech recognition inspired by human speech perception,” *arXiv preprint arXiv:2209.01768*, 2022. 5, 7
- [49] P. Ma, S. Petridis, and M. Pantic, “End-to-end audio-visual speech recognition with conformers,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 7613–7617. 5, 7
- [50] R. J. Williams and D. Zipser, “A learning algorithm for continually running fully recurrent neural networks,” *Neural computation*, vol. 1, no. 2, pp. 270–280, 1989. 5
- [51] M. Kim, J. H. Yeo, and Y. M. Ro, “Distinguishing homophones using multi-head visual-audio memory for lip reading,” in *Proceedings of the 36th AAAI Conference on Artificial Intelligence, Vancouver, BC, Canada*, vol. 22, 2022. 5, 7
- [52] L. Van der Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of machine learning research*, vol. 9, no. 11, 2008. 8