

# Facetron: A Multi-speaker Face-to-Speech Model based on Cross-Modal Latent Representations

Seyun Um, Jihyun Kim, Jihyun Lee, Hong-Goo Kang  
*Dept. of Electrical and Electronic Engineering Yonsei University Seoul, Korea*

**Abstract**—In this paper, we propose a multi-speaker face-to-speech waveform generation model that also works for unseen speaker conditions. Using a generative adversarial network (GAN) with linguistic and speaker characteristic features as auxiliary conditions, our method directly converts face images into speech waveforms under an end-to-end training framework. The linguistic features are extracted from lip movements using a lip-reading model, and the speaker characteristic features are predicted from face images using a face encoder trained through cross-modal learning with a pre-trained acoustic model. Since these two features are uncorrelated and controlled independently, we can flexibly synthesize speech waveforms whose speaker characteristics vary depending on the input face images. We show the superiority of our proposed model over conventional methods in terms of objective and subjective evaluation results.<sup>1</sup>

**Index Terms**—GAN-based audio-visual model, Lip-reading, Multi-speaker speech generator

## I. INTRODUCTION

Automated lip-reading (ALR) is the task of predicting linguistic information from a sequence of visual lip movements. It has become an essential technology for communications for hearing-impaired people, silent communications for keeping privacy, and so on [1], [2]. Especially, there have been several attempts to generate high-quality speech waveforms directly from lip movements and face images based on deep learning frameworks [3]–[6].

To synthesize human voices from silent lip movements and face images, one simple approach is to concatenate a text-to-speech (TTS) module to a module that converts either face or lip images to contextual variables. However, this approach has a limitation in that unavoidable recognition errors at the contextual conversion module degrade the synthesis performance of the following speech synthesis module. Thus, there have been several attempts to estimate acoustic information rather than phonetic information as intermediate features for waveform synthesis, which have been shown to improve the robustness of the whole synthesis framework. These acoustic features include vocoding parameters such as pitch, voicing information, and spectral parameters. Meanwhile, Vougioukas *et al.* [3] proposed a direct generation method of speech waveforms from a sequence of lip images based on GAN framework. Though they were able to generate more realistic and intelligible speech in both speaker-dependent and speaker-independent settings thanks to the GAN approach, there is still room for improvements of the quality of the generated speech.

<sup>1</sup>The demo samples of the proposed and other models are available at <https://sam-0927.github.io/>

In this paper, we use a GAN-based method to generate raw speech waveforms from lip and face image sequences. We obtain two intermediate features: contextual information from lip movements and speaker characteristic-related features from face images. To estimate the speaker characteristic-related acoustic features from face images, we introduce a cross-modal learning method that utilizes the cosine similarity between embeddings obtained from face images and prosody embeddings obtained from a pre-trained prosody-extracting module [7]. Both contextual and speaker-related acoustic embeddings are used as conditional features for a GAN-based waveform generation module. Based on various experiments using small and large-scale datasets, we show that our model successfully generates high-quality speech even for unseen speakers.

Our main contributions of this paper are as follows: 1) We introduce various training criteria to improve the quality of synthesized speech in a multi-speaker face-to-speech generation scenario. 2) Our model extracts linguistic and speaker characteristic-related embeddings independently; thus, it is possible to modify speaker characteristic information without changing the linguistic information. 3) To reliably extract speaker characteristic-related information, we propose an effective cross-modal learning technique using a pre-trained prosody model. Since the prosody model was pre-trained on a database with a large number of speakers, we can successfully predict the acoustic information for many speakers even if they were unseen during training. 4) We use a well-established GAN-based generator to synthesize high-fidelity speech from facial movements. Notably, the discriminator module includes an effective discrimination criterion motivated by a perceptual quality metric, which induces the generated speech to become more natural.

## II. RELATED WORK

### A. Lip-to-speech synthesis

Lip-to-speech algorithms consist of two main stages: the extraction of intermediate features from input images and the generation of speech waveforms from them. Vid2Speech [6] and Lipper [8] extract line spectrum pairs (LSPs) from face images and synthesize speech waveforms using linear prediction (LP) synthesis. However, since low order LP coefficients do not include sufficient acoustic information for speech synthesis, the generated speech sounds are often not natural. Vid2voc [9] adopts the WORLD vocoder [10] to improve speech quality. To further improve intelligibility, this

method jointly trains a visual speech recognition network and the vocoder parameter prediction network using a multi-task learning framework.

Another line of approaches is predicting spectrograms or mel-spectrograms, which contain more detailed acoustic information, and using them as intermediate features for speech waveform generation. One example of this approach is an extension of Vid2Speech [6] which uses facial movements as additional inputs to predict such spectrograms. Another example, Lip2AudSpec [11], uses a pre-trained autoencoder to compress input embeddings and reconstruct spectrograms. By utilizing a well-established sequence-to-sequence architecture, Prajwal *et al.* [4] develop Lip2Wav that learns the speaking styles of a small number of target speakers. Vougioukas *et al.* [3] directly generate raw speech waveforms using adversarial training, but their improvements are not particularly significant because they use a simple architecture of a stack of transposed convolutions for a generator module.

Our proposed model defines linguistic and acoustic features as intermediate features. We extract those features using two feature extractors: a lip-reading network that predicts linguistic features and a face encoder that predicts speaker characteristics. We also utilize the high-synthesis ability of a GAN-based neural vocoder to generate natural speech waveforms.

### B. Face-speech cross-modal learning

Cross-modal features between face and speech have been used to improve the performance of various speech-related tasks such as speech separation, speech enhancement, and speaker recognition. Due to their matching correspondence, the features can be also used for a domain transformation task between face and speech. For example, Speech2Face [12] synthesizes a face image given speech segments as input. In Face2Speech [5], a pre-trained multi-speaker TTS system synthesizes speech given speaker embeddings extracted from face images. Our method uses a metric-learning approach to train our face encoder to predict similar speaker characteristics from the features extracted from face images.

## III. PROPOSED MODEL

Figure 1 illustrates the system architecture of our proposed model, which consists of a mel-spectrogram generator and a neural vocoder. Our mel-spectrogram generator, which we call Facetron, transforms face and lip images into mel-spectrograms. It consists of three modules: a lip encoder, a face encoder, and a decoder. We then use a high-quality neural vocoder HiFi-GAN [13] to synthesize waveforms from the generated mel-spectrogram. Unlike a TTS model which has a one-to-many mapping problem, Facetron doesn't suffer from the mismatch of time alignment because the frame rate of face images is fixed. Detailed descriptions of each module in Facetron are provided in the following paragraphs.

**Lip encoder.** Lip encoder is a feature extractor that outputs lip embeddings given a sequence of lip images with size  $144 \times 144$  cropped from face images. The lip encoder is able to predict contextual and linguistic information because it is

trained with additional full-connected layers which predict actual graphemes from the output of the lip encoder. The design of the lip encoder is based on LipNet [1]. It consists of a 3-layer 3D CNNs, a residual block with channel-wise dropout and spatial pooling, a 2-layer bidirectional gated recurrent unit (Bi-GRU), and 2 fully connected layers. The residual block consists of two 3D CNNs that have 96-32 channels with the same kernel size of (3, 3, 3), a stride size of (1, 1, 1), and a padding size of (1, 1, 1). These bottleneck blocks play an important role in improving performance by extracting linguistic features. The output is multiplied by a scale factor of 0.2. Then, all outputs of 3D CNNs blocks are fed to the Bi-GRU module. We use connectionist temporal classification (CTC) loss [14] as the criterion for recognizing characters.

**Face encoder.** Our proposed model targets a multi-speaker system, which can synthesize different speech depending on speaker identities even when the linguistic information is the same. Thus, we define face embedding, which is the output of face encoder, to represent speaker characteristics and use it as a speaker-dependent feature for multi-speaker synthesis. We design the face encoder based on FaceNet [15], which is composed of several residual blocks consisting of 2D CNNs with pooling layers, dropout, and batch normalization. We randomly select one image from the 75 face images and use it as input for the face encoder in order to avoid the extraction of contextual information from a sequence of lip images.

**Decoder.** Decoder generates mel-spectrograms given the lip and face embeddings. We use a light version of HiFi-GAN [13] as the decoder. The three main modules of HiFi-GAN form the backbone of our decoder architecture: multi-receptive field fusion (MRF) for a generator, multi-period discriminator (MPD), and multi-scale discriminator (MSD). We use modified configurations of the generator and the discriminator architecture to fit the dimension of the input embedding and the mel-spectrograms, respectively. The upsampling rates of the generator are set to (1, 1, 2), the kernel sizes to (16, 16, 4), and the initial channel dimension to 640. The periods used in the MPD are (2, 3, 5) and the sizes of the convolutional layers in the MSD are (80, 160, 240, 480, 960, 960).

To sum up, the generation process of Facetron is represented as follows:

$$l_k = \mathbb{E}_{lip}(X_k), \quad f = \mathbb{E}_{face}(X) \quad (1)$$

$$c_k = \text{concat}(l_k, f) \quad (2)$$

$$\hat{M} = \mathbb{G}_{decoder}(C), \quad C \ni c_1, \dots, c_k, \dots, c_L \quad (3)$$

The lip embedding  $l_k$  of the  $k$ -th frame is extracted from the lip encoder  $\mathbb{E}_{lip}(\cdot)$ . The face embedding  $f$  is obtained by the face encoder  $\mathbb{E}_{face}(\cdot)$  using a face image  $X$ . The input of the decoder  $C$  is a sequence of  $c_k$ 's along the total  $L$  frames, each of which is a concatenation of  $l_k$  and  $f$ . Finally, the decoder predicts mel-spectrogram  $\hat{M}$ .

## IV. TRAINING AND INFERENCE

In this section, we explain a reliable end-to-end training method for the aforementioned modules and a criterion for

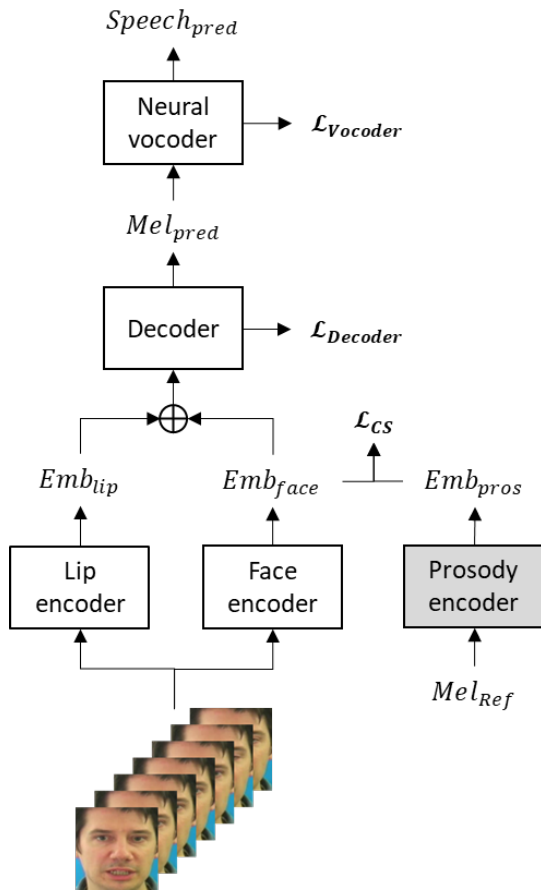


Fig. 1. System architecture of the proposed model. The input to the waveform generator (HiFi-GAN) is a concatenation of face embedding and lip embedding generated by the face encoder and the lip encoder, respectively. Gray block refers to pre-trained module.

selecting face embeddings in inference.

#### A. Training objectives

The overall loss function of the proposed model is as follows:

$$\mathcal{L}_{total} = \mathcal{L}_{Decoder} + \mathcal{L}_{CS} + \mathcal{L}_{Vocoder}, \quad (4)$$

where  $\mathcal{L}_{Decoder}$ ,  $\mathcal{L}_{CS}$ , and  $\mathcal{L}_{Vocoder}$  are loss functions for the decoder, the face encoder, and the neural vocoder, respectively. Note that the lip encoder is trained beforehand because the lip-reading task requires too much time for training.  $\mathcal{L}_{Decoder}$  is used to generate high-quality mel-spectrogram and includes generator loss, discriminator loss, and feature mapping loss as in [13].

One important issue in training our model is the way of providing speaker-related information to face embeddings. We introduce prosody information as a target for face embeddings because the speaker identity is strongly related to the speaker’s prosody. We extract prosody embeddings from a pre-trained prosody encoder in Mellotron [7] after fine-tuning it to utilize its general prosody modeling capability. We adopt a cosine similarity loss  $\mathcal{L}_{CS}$  between the prosody embedding and the

face embedding as a criterion for the face-prosody matching task:

$$\mathcal{L}_{CS} = \text{CS}(f, p), \quad (5)$$

$$f = \mathbb{E}_{face}(X), p = \mathbb{E}_{prosody}(M),$$

where  $f$  is the face embedding obtained by passing a face image  $X$  through the face encoder  $\mathbb{E}_{face}(\cdot)$ , and  $p$  is the prosody embedding obtained by passing mel-spectrogram  $M$  through the fine-tuned prosody encoder  $\mathbb{E}_{prosody}(\cdot)$ . For stable training of the face encoder, we train the face encoder separately from Facetron before the end-to-end training the whole model until the cosine similarity loss does not decrease.

#### B. Face embedding selection

During inference, we choose face embeddings using the Inter-to-Intra category distance over I2I method [16] instead of random selection to accurately represent the target speaker’s voice. First, we calculate the average of each face embedding from the target speaker generated by the face encoder, and also set the average of the face embeddings from a different (negative) speaker whose face embeddings are located closest to the target speaker’s embedding. Then, we apply the I2I method as follows:

$$\mathbf{f}_T^{I2I} = \frac{1}{2} \arg \max_{\mathbf{f}} \frac{\|\mathbf{f}_N^{Ave} - \mathbf{f}_T^{Ave}\|_2}{\mathbb{E}_{\mathbf{x} \in X_T} \|\mathbf{f} - \mathbf{x}\|_2}, \quad (6)$$

where  $\mathbf{f}_T^{I2I}$  denotes the selected face embedding from the target speaker.  $X_T$ ,  $\mathbf{f}_N^{Ave}$ , and  $\mathbf{f}_T^{Ave}$  denote the face embeddings belonging to the negative and target speakers, respectively. Using this method, Facetron is able to generate high quality mel-spectrograms not only for seen speakers’ voices, but also those of unseen speakers’.

## V. EXPERIMENTS AND RESULTS

In this section, we compare the performance of the proposed model to that of conventional models [3], [4], [9] in terms of synthesized speech quality, linguistic information accuracy, and speaker matching accuracy. We use the GRID dataset [17] for experiments, which contains high-quality audio and facial recordings of 1,000 sentences per speaker, spoken by 33 speakers (17 males, 16 females). The sampling rate of the audio signals is 16,000 Hz, and the duration of each sentence is 3 seconds. We train and evaluate our model on a small dataset (4 speakers) scenario and a large dataset (29 speakers) scenario to show the effectiveness of our method when trained on different amounts of data.

#### A. Small dataset

A small dataset consists of four speakers (s1, s2, s4, s29), which is the same dataset setup as reference models such as Voc-based [9], GAN-based approach [3], and Lip2Wav model [4]. We split the small dataset into training, validation, and test sets at ratios of 90%, 5%, and 5%, respectively, which is an identical setup to previous works. The output of the Voc-based and GAN-based approach are downsampled to 16 kHz

TABLE I  
OBJECTIVE AND SUBJECTIVE RESULTS OF THE PROPOSED AND REFERENCE MODELS FOR SMALL DATASET. THE BEST RESULTS ARE SHOWN IN BOLD.

	Objective			Subjective
	NISQA $\uparrow$	CER $\downarrow$	WER $\downarrow$	MOS $\uparrow$
GAN-based	2.79	9.1	18.3	1.42
Voc-based	2.63	<b>6.2</b>	<b>11</b>	2.19
Lip2wav	3.39	8	15.4	3.57
Facetron	<b>3.47</b>	7.1	13.7	<b>3.69</b>

TABLE II  
OBJECTIVE AND SUBJECTIVE RESULTS OF THE PROPOSED AND REFERENCE MODELS FOR LARGE SEEN/UNSEEN DATASET. THE (1<sup>ST</sup>, 2<sup>ND</sup>) AND (3<sup>RD</sup>, 4<sup>TH</sup>, 5<sup>TH</sup>, 6<sup>TH</sup>) ROWS ARE RESULTS FOR SEEN AND UNSEEN SPEAKER CASE, RESPECTIVELY. SMOS OF SEEN SPEAKER CASE IS ABOUT SPEAKER SIMILARITY AND UNSEEN SPEAKER CASE IS FOR GENDER SIMILARITY.

	Objective			Subjective	
	NISQA $\uparrow$	CER $\downarrow$	WER $\downarrow$	MOS $\uparrow$	SMOS $\uparrow$
Lip2wav	<b>3.60</b>	7.4	11.6	3.20	3.53
Facetron	3.44	<b>6.9</b>	<b>10.7</b>	<b>3.74</b>	<b>4.16</b>
Lip2wav	3.19	12.5	24.1	2.83	94.56%
Facetron	<b>3.24</b>	<b>7.7</b>	<b>14.7</b>	<b>3.87</b>	<b>99.22%</b>
VCVTS	-	-	-	2.31	-
Facetron	-	-	-	<b>4.19</b>	-

TABLE III  
OBJECTIVE AND SUBJECTIVE RESULTS OF THE PROPOSED MODEL FOR DISENTANGLEMENT.

	Objective			Subjective	
	NISQA $\uparrow$	CER $\downarrow$	WER $\downarrow$	MOS $\uparrow$	SMOS $\uparrow$
Facetron	3.41	8.3	14.3	3.57	3.8

(from 50KHz) to match the frequency range of the others. As in our model, we employ a HiFi-GAN vocoder for Lip2Wav.

In lip-to-speech applications, it is crucial to synthesize speech waveforms that convey linguistic information accurately. To evaluate this performance, we performed an automatic speech recognition (ASR) task with the synthesized speech waveforms under the framework used in [18]. The ASR model achieves 5.6% character error rate (CER) and 8.3% word error rate (WER) for original recorded speech. Table I shows the recognition accuracy. Our model achieves the WER of 13.7%, which is lower than the WER of the GAN-based and Lip2Wav. Although the Voc-based model shows a higher ASR performance than Facetron, its subjective quality, i.e. mean opinion score (MOS), is significantly lower than Facetron, which means that our model can synthesize natural speech while maintaining the linguistic contents. Overall, the results show that the speech synthesized by our model contains more accurate linguistic information than the baseline models.

To objectively measure overall speech quality, we compute non-intrusive objective speech quality assessment (NISQA-TTS) [19] scores. Our model outperforms the baseline models with a score of 3.47, showing that it synthesizes much higher speech quality than others when trained on a small dataset. We

also conducted MOS tests on the speech signals synthesized by Facetron and other baseline models. We randomly selected ten audio samples per speaker, that is, a total of 40 evaluation samples per model. Fifteen subjects participated in the MOS test; they were asked to evaluate the perceptual quality of the audio samples. Compared to the baseline models, Facetron achieves significantly higher MOS scores than baselines, which is the same result to NISQA.

### B. Large dataset

Large dataset utilizes the full dataset with all 33 speakers. During training, we excluded four speakers (s1, s2, s4, s29). With these four speakers' data, we evaluated the performance in unseen speakers' conditions. Table II shows the results. Facetron achieves the MOS scores of 3.74 and 3.87 for seen and unseen data, respectively. Notably, the results are almost identical between the seen and unseen cases, while Lip2Wav demonstrates degraded quality in the unseen case. These results show that Facetron generates high-quality speech signals for both seen and unseen speakers. In addition, Facetron is better than Lip2Wav in terms of WER; Facetron achieves 10.7% and 14.7% WER for seen and unseen speakers while Lip2Wav achieves 11.6% and 24.1%, respectively.

We also evaluated our model by conducting subjective speaker and gender recognition on synthesized speech samples. We conducted similarity MOS (SMOS), where subjects were asked to estimate how similar the synthesized and recorded audio samples were from the perspective of speaker and gender, respectively. The performance of these similarity measurement tasks is directly related to the capability of the face encoder. For seen speakers, our model obtains an SMOS score of 4.16 and Lip2Wav obtains 3.53 on the speaker similarity task. Therefore, we can say that our face encoder encodes speaker characteristic information for seen speakers much better than Lip2Wav. For unseen speakers, we performed a gender recognition task because it was not possible to recognize unknown speakers' identities. Our model achieves much higher accuracy than Lip2Wav (99.22% vs. 94.56%).

We also tried to compare the performance of our model with those of more recently published models such as VCVTS [20]. However, it was challenging to conduct this comparison because there was no official code available. Therefore, we performed a simple MOS test using a small number of speech samples taken from the official VCVTS demo page.<sup>2</sup> We downloaded three generated samples of unseen speakers and compared their quality with our generated samples of corresponding speakers. The speaker split setting of VCVTS is the same as ours as described in Section V-B. Eleven subjects were guided to evaluate the perceptual quality of the generated audio samples. Our model achieved an average MOS score of 4.19 while VCVTS only obtained 2.31, as shown in Table II. Although this was a listening test with only a small number of samples, the result supports the claim that our proposed model has significantly superior performance to VCVTS.

<sup>2</sup><https://wendison.github.io/VCVTS-demo/>




Fig. 2. Face embeddings extracted from Facetron with CS loss (left) and without it (right). Face embeddings of the same speaker are represented using the same type of color and marker and connected by a translucent line. Unseen speaker’s face embedding has a black border.

## VI. ABLATION STUDY

### A. Disentanglement

To verify whether linguistic and speaker identity features are successfully disentangled, we randomly synthesized 30 speech samples using lip and identity features from different speakers on the large dataset. As shown in Table III, our model achieves 8.3% CER, 14.3% WER, and 3.41 NISQA on objective metrics, and 3.57 MOS and 3.8 SMOS on subjective metrics. This result means that our model still shows high performance even though we used linguistic and speaker identity features from different speakers. This verifies that our method is able to effectively extract and disentangle linguistic and acoustic information.

### B. Effect of CS loss

The face encoder learns encoding speaker-related information from face images by training with the cosine similarity (CS) to prosody embeddings, as described in Section IV. Figure 2 illustrates the t-SNE plots to demonstrate the advantage of using CS loss. Face embeddings generated by Facetron with CS loss form clear speaker clusters while embeddings from the model without CS loss form distributed clusters. The effectiveness of CS loss is verified when synthesizing unseen speakers’ speech. Without CS loss, the model can fail to synthesize a consistent speaker identity in a single speech sample. For example, the front part of the sentence can be spoken in a male voice while the back part can be spoken in a female voice, which is not desirable. Examples of these samples has been uploaded to the demo page.

## VII. CONCLUSION

In this paper, we proposed Facetron, a model for synthesizing speech from images of facial movements. Our model combines a lip embedding that contains linguistic information and a face-driven embedding that contains speaker information and is trained using a cross-modal learning framework. Experiments show that our model is able to synthesize more natural and intelligible speech compared to previously proposed methods.

## VIII. ACKNOWLEDGEMENTS

This research was supported by the Yonsei Signature Research Cluster Program of 2022 (2022-22-0002).

## REFERENCES

- [1] Yannis M Assael, Brendan Shillingford, Shimon Whiteson, and Nando De Freitas, “Lipnet: End-to-end sentence-level lipreading,” *arXiv preprint arXiv:1611.01599*, 2016.
- [2] Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Senior, “Lip reading sentences in the wild,” in *Proc. CVPR*. IEEE, 2017, pp. 3444–3453.
- [3] Konstantinos Vougioukas, Pingchuan Ma, Stavros Petridis, and Maja Pantic, “Video-Driven Speech Reconstruction Using Generative Adversarial Networks,” in *Proc. Interspeech 2019*, 2019, pp. 4125–4129.
- [4] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Nambodiri, and CV Jawahar, “Learning individual speaking styles for accurate lip to speech synthesis,” in *Proc. CVPR*, 2020, pp. 13796–13805.
- [5] Shunsuke Goto, Kotaro Onishi, Yuki Saito, Kentaro Tachibana, and Koichiro Mori, “Face2speech: Towards multi-speaker text-to-speech synthesis using an embedding vector predicted from a face image,” *Proc. Interspeech 2020*, pp. 1321–1325, 2020.
- [6] Ariel Ephrat and Shmuel Peleg, “Vid2speech: speech reconstruction from silent video,” in *Proc. ICASSP*. IEEE, 2017, pp. 5095–5099.
- [7] Rafael Valle, Jason Li, Ryan Prenger, and Bryan Catanzaro, “Mellotron: Multispeaker expressive voice synthesis by conditioning on rhythm, pitch and global style tokens,” in *Proc. ICASSP*. IEEE, 2020, pp. 6189–6193.
- [8] Yaman Kumar, Rohit Jain, Khwaja Mohd Salik, Rajiv Ratn Shah, Yifang Yin, and Roger Zimmermann, “Lipper: Synthesizing thy speech using multi-view lipreading,” in *Proc. Assoc. Adv. Artif. Intell.*, 2019, vol. 33, pp. 2588–2595.
- [9] Daniel Michelsanti, Olga Slizovskaia, Gloria Haro, Emilia Gómez, Zheng-Hua Tan, and Jesper Jensen, “Vocoder-Based Speech Synthesis from Silent Videos,” in *Proc. Annu. Conf. Int. Speech Commun.*, 2020, pp. 3530–3534.
- [10] Masanori Morise, Fumiya Yokomori, and Kenji Ozawa, “WORLD: a vocoder-based high-quality speech synthesis system for real-time applications,” *IEICE Trans. on Inf. Syst.*, vol. 99, no. 7, pp. 1877–1884, 2016.
- [11] Hassan Akbari, Himani Arora, Liangliang Cao, and Nima Mesgarani, “Lip2audspec: Speech reconstruction from silent lip movements video,” in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.* IEEE, 2018, pp. 2516–2520.
- [12] Tae-Hyun Oh, Tali Dekel, Changil Kim, Inbar Mosseri, William T Freeman, Michael Rubinstein, and Wojciech Matusik, “Speech2face: Learning the face behind a voice,” in *Proc. CVPR*, 2019, pp. 7539–7548.
- [13] Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae, “Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis,” *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [14] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proc. ICML*, 2006, pp. 369–376.
- [15] Florian Schroff, Dmitry Kalenichenko, and James Philbin, “Facenet: A unified embedding for face recognition and clustering,” in *Proc. CVPR*, 2015, pp. 815–823.
- [16] Se-Yun Um, Sangshin Oh, Kyunguen Byun, Inseon Jang, ChungHyun Ahn, and Hong-Goo Kang, “Emotional speech synthesis with rich and granularized control,” in *Proc. ICASSP*. IEEE, 2020, pp. 7254–7258.
- [17] 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*, vol. 120, no. 5, pp. 2421–2424, 2006.
- [18] William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *Proc. ICASSP*. IEEE, 2016, pp. 4960–4964.
- [19] Gabriel Mittag and Sebastian Möller, “Deep learning based assessment of synthetic speech naturalness,” *arXiv preprint arXiv:2104.11673*, 2021.
- [20] Disong Wang, Shan Yang, Dan Su, Xunying Liu, Dong Yu, and Helen Meng, “Vcvtvs: Multi-speaker video-to-speech synthesis via cross-modal knowledge transfer from voice conversion,” in *Proc. ICASSP*. IEEE, 2022, pp. 7252–7256.