



Incorporating Ultrasound Tongue Images for Audio-Visual Speech Enhancement through Knowledge Distillation

Rui-Chen Zheng, Yang Ai, Zhen-Hua Ling*

National Engineering Research Center of Speech and Language Information Processing,
University of Science and Technology of China, Hefei, P. R. China

zhengruichen@mail.ustc.edu.cn, {yangai, zhling}@ustc.edu.cn

Abstract

Audio-visual speech enhancement (AV-SE) aims to enhance degraded speech along with extra visual information such as lip videos, and has been shown to be more effective than audio-only speech enhancement. This paper proposes further incorporating ultrasound tongue images to improve lip-based AV-SE systems' performance. Knowledge distillation is employed at the training stage to address the challenge of acquiring ultrasound tongue images during inference, enabling an audio-lip speech enhancement student model to learn from a pre-trained audio-lip-tongue speech enhancement teacher model. Experimental results demonstrate significant improvements in the quality and intelligibility of the speech enhanced by the proposed method compared to the traditional audio-lip speech enhancement baselines. Further analysis using phone error rates (PER) of automatic speech recognition (ASR) shows that palatal and velar consonants benefit most from the introduction of ultrasound tongue images.

Index Terms: speech enhancement, audio-visual, ultrasound tongue image, knowledge distillation

1. Introduction

Speech enhancement (SE) is a crucial research problem in speech signal processing that aims to improve the quality and intelligibility of speech signals corrupted by various types of distortions [1]. Numerous SE methods have been proposed to assist the application in several areas, such as speech or speaker recognition, hearing aids, and mobile communication [2]. Classical audio-only speech enhancement (AO-SE) approaches have been successful in estimating underlying target speech signals using either knowledge-based modeling [3, 4], or data-driven paradigms such as deep learning [5, 6]. Since speech perception is inherently multimodal, particularly audio-visual, recent studies have investigated using visual cues as well as acoustic signals for SE [7]. This approach, known as audio-visual speech enhancement (AV-SE) [8–10], is more effective than simple AO-SE methods as visual cues are essentially unaffected by the acoustic environment.

Lip videos are the most commonly used visual cues for AV-SE due to their easy availability. Besides, they can also help disambiguate phonetically similar sounds since they record the movement of external articulators involved in speech production. A deep AV-SE model based on convolutional neural networks (CNNs) [11] was proposed to separate a speaker's voice by predicting both the magnitude and the phase of the target signal given lip regions. A novel framework [12] incorporated

lip information for speech enhancement by integrating a generative adversarial network (GAN) to generate high-quality speech. Another study [13] proposed an AV-SE system that achieved impressive performance using a multi-layer fusion model with a multi-head cross-attention mechanism to fuse audio and lip features.

Speech production is a complex process relying on multiple articulators, including the jaw, lips, teeth, and tongue. Using lip videos for speech processing tasks like SE often has limitations due to the lack of descriptions on internal articulators, e.g., tongue and velar. To address this issue, some studies have suggested employing internal articulation features captured using medical imaging techniques such as magnetic resonance imaging (MRI) [14] and ultrasound tongue imaging (UTI) [15], to provide complementary data. Compared to MRI, UTI is relatively cheap, non-invasive, and can provide high-resolution images. UTI uses a real-time B-mode ultrasound transducer placed under the speaker's chin to visualize a midsagittal or coronal view of the tongue during speech production. Ultrasound tongue images have been used in various speech processing tasks, such as speech recognition [16, 17] and speech reconstruction [18, 19]. However, their potential for speech enhancement tasks has yet to be fully explored, creating a research gap that could be addressed by leveraging the advantages of ultrasound tongue images.

However, obtaining ultrasound tongue images is more complex than collecting lip videos due to the requirement of extra equipment. Therefore, this paper proposes introducing ultrasound tongue information to an audio-lip SE model through knowledge distillation. Specifically, an audio-lip-tongue SE teacher model with U-Net-based [20] structure is first proposed to incorporate ultrasound tongue images with lip videos to assist speech enhancement. An audio-lip SE student model is then guided to learn ultrasound tongue information from a pre-trained audio-lip-tongue SE teacher model via multiple loss functions. At inference time, the proposed audio-lip SE model is fed with noisy speech and lip videos, and the ultrasound tongue images are not necessary anymore. Experimental results demonstrate that the proposed method can generate speech with improved quality and intelligibility compared to conventional audio-lip baselines trained solely with clean speech supervision. Moreover, a notable reduction in phoneme error rate (PER), specifically for palatal and velar consonants, can be witnessed while applying automatic speech recognition (ASR) for transcription.

2. Proposed Method

Under most circumstances, only a noisy speech and its corresponding lip video can be obtained as reference data since acquiring ultrasound tongue images is not as straightforward as

* Corresponding Author. This work was partially funded by the Fundamental Research Funds for the Central Universities.

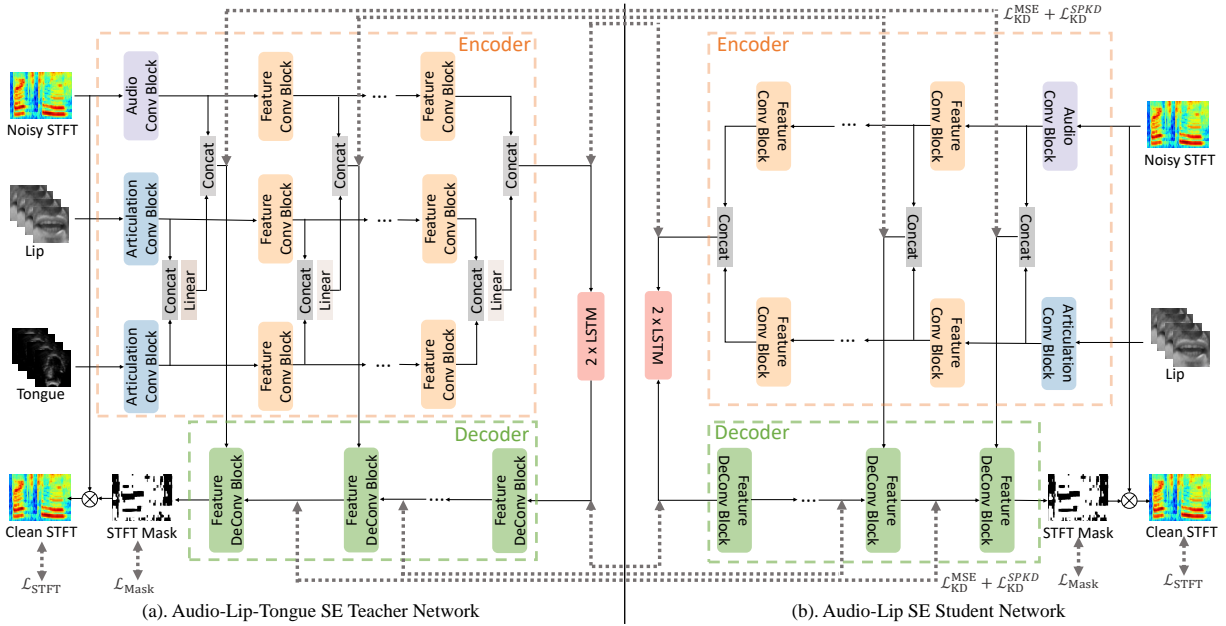


Figure 1: *Details of the proposed method. The dashed arrows indicate the loss function. (a) shows the structure of the Audio-Lip-Tongue SE teacher model; (b) shows the structure of the Audio-Lip SE student model introducing ultrasound knowledge with knowledge distillation method.*

capturing lip video. Hence, we propose enabling an audio-lip SE student model to assimilate ultrasound tongue information from a pre-trained audio-lip-tongue SE teacher model through knowledge distillation. Details are illustrated in Fig. 1 and will be introduced in this section.

2.1. Audio-Lip-Tongue SE Teacher Model

Inspired by previous study [12, 21], we propose a U-Net [20] style teacher model with noisy speech, lip videos, and ultrasound tongue images as input as shown in Fig.1a. The model can be roughly divided into a multi-modal encoder and a decoder connected by a long short-term memory (LSTM) based embedding block.

The encoder contains an articulation stream and an audio stream. For the articulation stream, we process both lip and tongue image sequences to obtain fused representations for describing articulation. Specifically, pixel-wise mean and standard deviation are computed for each utterance, repeated, and then appended as extra channels to the ultrasound and lip sequences [19]. The resulting input is of dimension $3 \times T \times H \times W$, where H and W are the height and width of the lip and tongue images. The input is first processed by an articulation convolutional block consisting of three strided 3D-CNN layers, each followed by batch normalization and Leaky-ReLU. The resulting features are then flattened along the time axis and processed by seven feature convolutional blocks. Each block involves a series of 2D-CNN layers with frequency pooling layers to reduce the frequency while preserving the time dimension. The lip and tongue features are concatenated at each layer and then passed through a linear layer to obtain the fused articulation representations with reduced size.

For the audio stream, the encoder takes as input the real and imaginary parts of the noisy complex spectrogram with dimension $2 \times F \times T$, where F and T are the frequency and time dimensions of the spectrogram. Each time-frequency bin contains the real and imaginary parts of the corresponding complex spectrogram value [21]. The input is first processed by an audio convolutional block comprising two strided 2D-CNN layers

and then by seven feature convolutional blocks as used in the articulation stream. Finally, the audio representation are concatenated with the fused articulation representations layer by layer to obtain the ultimate multi-modal representations for SE.

Two LSTM layers are inserted between the encoder and decoder to better model temporal dependencies. The decoder with skip connection exhibits a symmetric structure concerning the encoder, whereby the convolutional layer is substituted with an upconvolutional layer, and the frequency pooling layer is replaced by a frequency upsampling layer. The input of each decoder layer is the concatenation of the output of the previous layer and the multi-modal representations given by the corresponding encoder layer. The final output feature map is fed through an activation layer to predict a complex mask with the same dimensions as the input noisy spectrogram. The resulting mask is applied to the noisy input through complex multiplication, yielding an enhanced complex spectrogram which is transformed back into the time domain via an inverse short-time fourier transform (iSTFT).

The training criterion of the teacher model is to minimize the following loss

$$\mathcal{L}_{Teacher} = \mathcal{L}_{Mask} + \alpha \mathcal{L}_{STFT}, \quad (1)$$

where \mathcal{L}_{Mask} and \mathcal{L}_{STFT} denote the mean square error (MSE) losses of the mask and the complex spectrogram, respectively. The hyper-parameter α is utilized to ensure both \mathcal{L}_{Mask} and \mathcal{L}_{STFT} are scaled to the similar magnitude.

2.2. Audio-Lip SE Student Model

An audio-lip SE student model is proposed with the architecture depicted in Fig.1b, where the input of tongue images is removed compared with the teacher. To utilize tongue information even in the absence of ultrasound tongue images during reference, we employ knowledge distillation to train the student. Specifically, in addition to the supervised training with the backbone losses \mathcal{L}_{Mask} and \mathcal{L}_{STFT} described in Eq.(1), the student model also receives supervision signals from the teacher. During forward inference of the student and teacher networks, the output of each layer is saved to compute the knowledge distillation loss

Table 1: Evaluation results of the proposed model compared with baseline models. Best results are highlighted in **bold**.

Method	SNR	2.5dB			-2.5dB			-7.5dB		
	Metrics	SegSNR	PESQ	STOI	SegSNR	PESQ	STOI	SegSNR	PESQ	STOI
Noisy		0.2414	1.4285	0.8503	-3.1703	1.2283	0.7527	-5.9847	1.1298	0.6229
Gabby et al. 2018 [10]		4.8731	1.8260	0.8871	3.8721	1.6502	0.8566	2.6978	1.4590	0.8005
Hegde et al. 2021 [22]		3.6086	1.8756	0.8882	2.9790	1.6890	0.8603	2.0987	1.4863	0.8052
Hussain et al. 2021 [23]		6.1654	2.0388	0.9245	4.5055	1.7750	0.8867	2.8217	1.5424	0.8177
Proposed		9.0061	2.1073	0.9335	6.8343	1.8241	0.8911	4.6944	1.5858	0.8221

for each layer of the encoder, decoder, and LSTM embedding blocks. For each layer, the MSE loss between the output features of the teacher and the student is calculated following

$$\mathcal{L}_{KD}^{MSE} = \sum_{l=1}^N \|\mathbf{F}_T^l - \mathbf{F}_S^l\|_2^2, \quad (2)$$

where N represents the number of layers, and \mathbf{F}_T^l and \mathbf{F}_S^l denote the output features of the l^{th} layer of the teacher and student model, respectively. We guarantee that the feature dimensions of the corresponding layers in the teacher and student models are identical.

The similarity-preserving knowledge distillation (SPKD) loss [24, 25] is utilized to further supervise the student model. The SPKD loss aims to achieve dimensional compression and to transmit similarity information simultaneously by computing pairwise similarity matrices. Given a mini-batch input, the feature map of the l^{th} layer is represented as $\mathbf{F}^l \in \mathbb{R}^{b \times c \times t \times f}$, where b is the batch size, c is the number of output channels, t is the number of frames, and f is the dimension of the feature space. To account for potential interference of information from different frames, the features are first segmented on frame-level and then flattened into two dimensions. The transformed feature of the j^{th} frame is denoted as $\mathbf{F}^{(l,j)} \in \mathbb{R}^{b \times f'}$, where $f' = c \times f$. The similarity matrix of each frame for the teacher $\mathbf{G}_T^{(l,j)}$ and the student $\mathbf{G}_S^{(l,j)}$ is calculated independently following [25]. The SPKD loss for the l^{th} layer is then calculated as the sum of the similarity distances across all frames. The overall SPKD loss is determined by summing the SPKD loss of each layer as follows

$$\mathcal{L}_{KD}^{SPKD} = \sum_{l=1}^N \mathcal{L}_{SPKD}^l = \frac{1}{b^2} \sum_{l=1}^N \sum_{j=1}^T \|\mathbf{G}_T^{(l,j)} - \mathbf{G}_S^{(l,j)}\|_F^2, \quad (3)$$

where $\|\cdot\|_F$ is the Frobenius norm.

Therefore, to train the proposed audio-lip SE model, the overall loss function can be written as

$$\mathcal{L}_{Student} = \mathcal{L}_{Mask} + \alpha \mathcal{L}_{STFT} + \gamma_1 \mathcal{L}_{KD}^{MSE} + \gamma_2 \mathcal{L}_{KD}^{SPKD}, \quad (4)$$

where, α , γ_1 and γ_2 are all hyper-parameters used to ensure that the loss functions are scaled to the similar magnitude.

3. Datasets and Implementation Details

3.1. Datasets

For our experiments, we utilized the TaL corpus [26], a multi-speaker dataset containing ultrasound tongue imaging, optical lip video, and audio for each utterance. Specifically, we employed 10,271 utterances from 73 speakers for training, and 810 utterances from 81 speakers for validation, where eight speakers from the validation sets were unseen in the training set. The test set included 1749 utterances from 73 seen speakers and 1407 utterances from 8 unseen speakers. The content of the three sets was mutually exclusive from each other.

Noisy speech was generated by mixing noise with the clean

speech in the TaL corpus. During training, we introduced ten different noise types as in [27]: eight noise recordings from the DEMAND database [28] and two artificially generated noises, namely speech-shaped noise and babble noise. These noise types were added to the speech signal at three different signal-to-noise (SNR) values (0dB, 5dB, and 10dB). For validation and testing, we added five additional noises: living room, office, bus, street cafe, and a public square. We used slightly higher SNR values (-2.5dB, 2.5dB and 7.5dB) than the ones used during training following previous work [27].

3.2. Implementation Details

The lip videos and ultrasound tongue images were processed following the pipeline described in [19]. The resulting lip videos and ultrasound tongue images were resized to 64×128 and had a frame rate of 81.5 fps. For audio preprocessing, the audio was downsampled to 16 kHz, and the STFT was computed using a Hann window with a length of 512, a hop size of 196 and an FFT point number of 512 to match the frame rate of the ultrasound. The resulting complex spectrogram had a dimension of $2 \times 257 \times T$.

We employed the Adam optimizer with an initial learning rate of 1e-3 for backpropagation. The learning rate decreased by 0.1 once learning stagnated, i.e., the validation error did not improve for ten epochs. The model was trained at utterance level by randomly cropping all the samples in the mini-batch to have the same number of frames as the shortest one.

4. Experimental Results

4.1. Overall Performance

Three recent AV-SE methods [10, 22, 23] whose source codes were available online were used as baselines for comparison with our proposed audio-lip SE model. When evaluating the method [22], as a generated pseudo lip stream was adopted for AV-SE in its original paper, we used the natural lip stream for a fair comparison. Each baseline was trained on our constructed dataset following the training strategy in its original paper. Segmental signal-to-noise ratio (SegSNR), perceptual evaluation of speech quality (PESQ), and short-term objective intelligibility (STOI) were used as evaluation metrics. The results in Table 1 demonstrate that our proposed method outperforms all three baselines in all evaluation metrics, indicating improved speech quality and intelligibility¹.

4.2. Ablation Studies

A series of ablation studies were conducted to provide further evidence of the efficacy of incorporating ultrasound tongue knowledge through the knowledge distillation method. The outcomes are presented in Table 2. Specifically, the effectiveness of the SPKD loss was demonstrated by removing the loss function \mathcal{L}_{KD}^{SPKD} in Eq.4 (“w/o SPKD”). To demonstrate the effectiveness of introducing ultrasound tongue information

¹Speech samples are available at: <https://zhengrachel.github.io/UTIforAVSE-demo/>

Table 2: Evaluation results of the proposed model in ablation studies. Best results are highlighted in **bold**. Underline characters indicate the sub-optimal results.

Method	SNR	2.5dB			-2.5dB			-7.5dB		
	Metrics	SegSNR	PESQ	STOI	SegSNR	PESQ	STOI	SegSNR	PESQ	STOI
Proposed		9.0061	2.1073	0.9335	6.8343	1.8241	0.8911	4.6944	1.5858	0.8221
w/o SPKD		8.9927	2.1133	0.9328	6.8251	1.8218	0.8892	4.6442	1.5802	0.8194
w/o KD (Audio-Lip)		8.9741	2.0841	0.9304	6.8073	1.8030	0.8854	4.6314	1.5604	0.8142
Audio-Lip-Tongue		9.0315	2.1122	0.937	6.9267	1.8322	0.8987	4.8707	1.5973	0.8394
Audio-Tongue		8.8509	2.0990	0.9321	6.7223	1.8133	0.8901	4.6153	1.5631	0.822
Audio-Only		8.6748	1.9872	0.9215	6.4224	1.7088	0.8661	4.0888	1.4773	0.7739

Table 3: PERs (%) of different phoneme categories for clean speech, noisy speech, and the speech enhanced by different methods. Consonant phonemes are divided according to the place of articulation. Best results are highlighted in **bold**. Underline characters indicate the sub-optimal results. The numbers in parentheses show relative PER reduction (%) compared with the audio-only method.

Methods	Silence	Vowels	Consonants							
			Labial	Labio-dental	Dental	Alveolar	Alveo-palatal	Palatal	Velar	Glottal
Clean	0.58	3.09	2.54	1.24	2.18	2.96	1.94	1.88	2.69	2.46
Noisy	25.10	46.10	41.58	37.44	43.43	41.74	33.03	37.53	37.10	46.35
Audio-Only	10.54	31.97	33.54	34.03	31.02	33.46	14.8	30.78	27.37	33.76
Audio-Lip	7.71	27.67	27.97	<u>31.82</u>	24.78	30.29	12.48	30.33	26.48	30.42
	(26.85)	(13.45)	(16.61)	(6.49)	(20.12)	(9.47)	(15.68)	(1.46)	(3.25)	(9.89)
Audio-Lip-Tongue	6.34	20.22	24.95	32.15	22.40	25.61	11.13	23.55	17.35	25.09
	(39.84)	(36.75)	(25.61)	(5.52)	(27.79)	(23.46)	(24.80)	(23.49)	(36.61)	(25.68)
Proposed	<u>7.34</u>	<u>25.31</u>	<u>26.42</u>	30.56	25.44	<u>29.35</u>	<u>11.17</u>	<u>26.59</u>	<u>22.85</u>	<u>29.87</u>
	(30.36)	(20.83)	(21.23)	(10.20)	(17.99)	(12.28)	(24.53)	(13.61)	(16.51)	(11.52)

via the knowledge distillation method, an audio-lip SE model was trained solely with \mathcal{L}_{Mask} and \mathcal{L}_{STFT} (“w/o KD (Audio-Lip)”). Additionally, for reference, we included results from the well-trained audio-lip-tongue SE teacher model together with the audio-tongue and audio-only SE models built by removing the lip or the whole articulation inputs in the teacher model.

Table 2 shows that while the audio-lip-tongue SE teacher model achieved the best performance, the proposed method won the second place on most metrics. “w/o SPKD” achieved results inferior to the proposed method except for SNR=2.5dB, indicating the effectiveness of introducing the SPKD loss particularly in low SNR scenarios. The outcomes of the proposed method surpassed those of the “w/o KD (Audio-Lip)”, suggesting that the knowledge distillation method effectively incorporates a portion of the articulation knowledge derived from the ultrasound tongue. Furthermore, the performance of the proposed method outperformed that of both the audio-tongue and audio-only models, highlighting the effectiveness of combining multi-modal sources of articulation knowledge.

4.3. Analysis on Phoneme Categories

To further study the gains of incorporating ultrasound tongue images, we employed an ASR engine to analyze the enhanced speech’s PERs of different phoneme categories. An ASR API provided in ESPNet² [29] was utilized to transcribe the enhanced speech. Each enhanced utterance was forcibly aligned with its transcription by the Montreal Forced Aligner (MFA)³ [30] tool, consequently obtaining a time-aligned phoneme sequence. Ground truth phoneme sequences were obtained by aligning clean utterances with their corresponding texts using the MFA tool. The recognized and ground truth phoneme sequences were further aligned by minimizing the edit distance, and the PERs for different phoneme categories could be cal-

culated. Phonemes were categorized according to the English (UK) MFA dictionary v2.0.0⁴.

The results are shown in Table 3. Comparing audio-lip and audio-only methods, we can see that after introducing lip information for SE, PERs reduced significantly for some lip-related phonemes, such as silence, labials and dentals. However, for palatal and velar consonants, the relative PER reductions were quite small which indicates the limitations of using only lip videos as articulation information. After further incorporating ultrasound tongue images, the audio-lip-tongue method achieved much more uniform PER reductions among phoneme categories than the audio-lip method. For palatals, velars and vowels whose articulations were mainly determined by internal tongue movement, their relative PER reductions improved significantly from 1.46% to 23.49%, from 3.25% to 36.61% and from 13.45% to 36.75%, respectively. Although not using tongue images at inference time, our proposed method also achieved lower PERs than the audio-lip method for most phoneme categories, especially for palatals, velars and vowels, indicating the effectiveness of incorporating tongue information through knowledge distillation.

5. Conclusion

This paper proposes incorporating ultrasound tongue images via knowledge distillation for AV-SE systems. An audio-lip-tongue SE model is first trained and used as a teacher to guide the training of an audio-lip SE student model through knowledge distillation. Our proposed method demonstrates superior performance over the conventional AV-SE baselines. Furthermore, the evaluated PER results indicate that lip and tongue modalities provide valuable articulation knowledge. Investigating the feasibility of constructing pseudo ultrasound tongue features from audio and lip video modalities to assist speech processing tasks will be our future work.

²https://github.com/espnet/espnet_model_zoo

³<https://github.com/MontrealCorpusTools/Montreal-Forced-Aligner>

⁴<https://mfa-models.readthedocs.io/en/latest/dictionary/English/>

6. References

- [1] A. Kumar and D. Florencio, "Speech enhancement in multiple-noise conditions using deep neural networks," in *Proc. Interspeech*, 2016, pp. 3738–3742.
- [2] A. Chaudhari and S. Dhonde, "A review on speech enhancement techniques," in *Proc. ICPC*, 2015, pp. 1–3.
- [3] Y. Ephraim and H. L. Van Trees, "A signal subspace approach for speech enhancement," *IEEE Transactions on speech and audio processing*, vol. 3, no. 4, pp. 251–266, 1995.
- [4] L.-P. Yang and Q.-J. Fu, "Spectral subtraction-based speech enhancement for cochlear implant patients in background noise," *The journal of the Acoustical Society of America*, vol. 117, no. 3, pp. 1001–1004, 2005.
- [5] Y. Xu, J. Du, L.-R. Dai, and C.-H. Lee, "A regression approach to speech enhancement based on deep neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 1, pp. 7–19, 2014.
- [6] A. Défossez, G. Synnaeve, and Y. Adi, "Real time speech enhancement in the waveform domain," in *Proc. Interspeech*, 2020, pp. 3291–3295.
- [7] D. Michelsanti, Z.-H. Tan, S.-X. Zhang, Y. Xu, M. Yu, D. Yu, and J. Jensen, "An overview of deep-learning-based audio-visual speech enhancement and separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1368–1396, 2021.
- [8] I. Almajai and B. Milner, "Visually derived wiener filters for speech enhancement," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 6, pp. 1642–1651, 2010.
- [9] J.-C. Hou, S.-S. Wang, Y.-H. Lai, Y. Tsao, H.-W. Chang, and H.-M. Wang, "Audio-visual speech enhancement using multi-modal deep convolutional neural networks," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 2, no. 2, pp. 117–128, 2018.
- [10] A. Gabbay, A. Shamir, and S. Peleg, "Visual speech enhancement," in *Proc. Interspeech*, 2018, pp. 1170–1174.
- [11] T. Afouras, J. S. Chung, and A. Zisserman, "The Conversation: Deep audio-visual speech enhancement," in *Proc. Interspeech*, 2018, pp. 3244–3248.
- [12] X. Xu, Y. Wang, D. Xu, Y. Peng, C. Zhang, J. Jia, and B. Chen, "VSEGAN: visual speech enhancement generative adversarial network," in *Proc. ICASSP*, 2022, pp. 7308–7311.
- [13] X. Xu, Y. Wang, J. Jia, B. Chen, and D. Li, "Improving visual speech enhancement network by learning audio-visual affinity with multi-head attention," in *Proc. Interspeech*, 2022.
- [14] A. D. Scott, M. Wylezinska, M. J. Birch, and M. E. Miquel, "Speech MRI: morphology and function," *Physica Medica*, vol. 30, no. 6, pp. 604–618, 2014.
- [15] M. Stone, "A guide to analysing tongue motion from ultrasound images," *Clinical linguistics & phonetics*, vol. 19, no. 6-7, pp. 455–501, 2005.
- [16] L. Tóth, G. Gosztolya, T. Grósz, A. Markó, and T. G. Csapó, "Multi-task learning of speech recognition and speech synthesis parameters for ultrasound-based silent speech interfaces," in *Proc. Interspeech*, 2018, pp. 3172–3176.
- [17] M. S. Ribeiro, A. Eshky, K. Richmond, and S. Renals, "Silent versus modal multi-speaker speech recognition from ultrasound and video," in *Proc. Interspeech*, 2021, pp. 641–645.
- [18] D. Porras, A. Sepúlveda-Sepúlveda, and T. G. Csapó, "DNN-based acoustic-to-articulatory inversion using ultrasound tongue imaging," in *Proc. IJCNN*, 2019, pp. 1–8.
- [19] J.-X. Zhang, K. Richmond, Z.-H. Ling, and L. Dai, "TaLnet: Voice reconstruction from tongue and lip articulation with transfer learning from text-to-speech synthesis," in *Proc. AAAI*, vol. 35, no. 16, 2021, pp. 14402–14410.
- [20] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. MICCAI*, 2015, pp. 234–241.
- [21] R. Gao and K. Grauman, "VisualVoice: Audio-visual speech separation with cross-modal consistency," in *Proc. CVPR*, 2021, pp. 15490–15500.
- [22] S. B. Hegde, K. Prajwal, R. Mukhopadhyay, V. P. Namboodiri, and C. Jawahar, "Visual speech enhancement without a real visual stream," in *Proc. WACV*, 2021, pp. 1926–1935.
- [23] T. Hussain, M. Gogate, K. Dashtipour, and A. Hussain, "Towards intelligibility-oriented audio-visual speech enhancement," *Clarity Workshop on Machine Learning Challenges for Hearing Aids*, 2021.
- [24] F. Tung and G. Mori, "Similarity-preserving knowledge distillation," in *Proc. ICCV*, 2019, pp. 1365–1374.
- [25] J. Cheng, R. Liang, Y. Xie, L. Zhao, B. Schuller, J. Jia, and Y. Peng, "Cross-layer similarity knowledge distillation for speech enhancement," in *Proc. Interspeech*, 2022, pp. 926–930.
- [26] M. S. Ribeiro, J. Sanger, J.-X. Zhang, A. Eshky, A. Wrench, K. Richmond, and S. Renals, "TaL: a synchronised multi-speaker corpus of ultrasound tongue imaging, audio, and lip videos," in *Proc. SLT*, 2021, pp. 1109–1116.
- [27] C. Valentini-Botinhao and J. Yamagishi, "Speech enhancement of noisy and reverberant speech for text-to-speech," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 8, pp. 1420–1433, 2018.
- [28] J. Thiemann, N. Ito, and E. Vincent, "The diverse environments multi-channel acoustic noise database (DEMAND): A database of multichannel environmental noise recordings," in *Proc. ICA*, vol. 19, no. 1, 2013, p. 035081.
- [29] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Yalta Soplin, J. Heymann, M. Wiesner, N. Chen *et al.*, "ESP-Net: End-to-end speech processing toolkit," in *Proc. Interspeech*, 2018, pp. 2207–2211.
- [30] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal Forced Aligner: Trainable text-speech alignment using kaldif," in *Proc. Interspeech*, 2017, pp. 498–502.