Multi-Modal Multi-Correlation Learning for Audio-Visual Speech Separation

Xiaoyu Wang[1]*, Xiangyu Kong[2], Xiulian Peng[2], Yan Lu[2]

[1]Xi'an Jiaotong University
[2]Microsoft Research Asia, Beijing, China
wxystudio@stu.xjtu.edu.cn, {xiakon,xipe,yanlu}@microsoft.com

Abstract

In this paper we propose a multi-modal multi-correlation learning framework targeting at the task of audio-visual separation. Although previous efforts have been extensively put on combining audio and visual modalities, most of them solely adopt a straightforward concatenation of audio and visual features. To exploit the real useful information behind these two modalities, we define two key correlations which are: (1) *identity correlation* (between timbre and facial attributes); (2) *phonetic correlation* (between phoneme and lip motion). These two correlations together comprise the complete information, which shows a certain superiority in separating target speaker's voice especially in some hard cases, such as the same gender or similar content. For implementation, contrastive learning or adversarial training approach is applied to maximize these two correlations. Both of them work well, while adversarial training shows its advantage by avoiding some limitations of contrastive learning. Compared with previous research, our solution demonstrates clear improvement on experimental metrics without additional complexity. Further analysis reveals the validity of the proposed architecture and its good potential for future extension.

Index Terms: cross-modal learning, audio-visual speech separation, contrastive learning, adversarial training

1. Introduction

Conventional speech separation aims to extract specified speaker's voice in multi-speaker environment against other speakers' interference [1, 2], which is also known as the "Cocktail Party" problem. Our work focuses on audio-visual speech separation. In such scenario where the video of target speaker is given, recent studies have shown that using auxiliary visual information in video can significantly improve speech separation performance [3–29].

But using what kind of auxiliary visual information is the crux of difficulty. The most intuitive method regards static image as reference [3–6], which usually misses a great deal of useful information. Some other architectures take all the video frames as input to capture their temporal correspondence with audio [7–29]. Due to the high redundancy and excessive size of video frames, merely extracting lip motion from video is a more appropriate choice [30–32].

Combining audio with lip motion feature can greatly improve speech separation performance [30–32]. The underlying intuition is that a strong cross-modal correlation exists between aligned visual and audio signals (e.g. lip motion is closely related to speech content). Another crucial cross modal information lies in the speaker identity. Based on some previous efforts [29, 33, 34], here we summarize the major correlations

between audio and visual modalities into two aspects. The first is *speaker identity*, which corresponds to the speaker's facial attributes in video and timbre in audio respectively. The second is *phonetic information*, which is closely relevant to lip motion in video and phoneme in audio. Intuitively they can help to deal with some hard cases in speech separation task. For example, when two speakers' characteristics are quite similar, e.g. of the same gender, usually it would be more difficult to separate the two voices. The phonetic information (content) of speech can help with this issue. In another case, if two speakers say similar or even the same words simultaneously, we can turn to exploit the subtle difference in speaker characteristics as a hint.

Although some prior work already consider the above mentioned two audio-visual correlations [29, 34], they fail to make full use of them. Most of previous methods simply concatenate mixed audio with visual feature (image, lip motion, frames) as input [3–32], which allows the model to implicitly learn the correlation between audio and visual representation. [35] utilizes a contrastive learning framework to directly learn the audio-visual correlation. They just assume an underlying common bond between audio and the corresponding video without explicitly modelling the identity or phonetic correlation between them. [33] adopts the same method and makes a further step to target at the identity correlation, while they ignore to explicitly model the phonetic correlation. Unlike them, our work put forward a more comprehensive architecture that explicitly models both identity and phonetic correlations between audio and visual modalities and shows its superiority by experimental results.

Specifically, we follow the training scheme in [33] and [35] to take advantage of the two correlations defined above under a contrastive learning framework. The key point is how to construct a triplet $(\mathbf{a}, \mathbf{b}, \mathbf{c})$ from training data, which contains a positive pair $(\mathbf{a}, \mathbf{b})$ with correlation and a negative pair $(\mathbf{a}, \mathbf{c})$ without correlation. Our goal is to "pull" positive (similar) pairs relatively close under a specific metrics, while "pushing" negative (dissimilar) pairs far away. For speaker identity correlation, we regard sample of the same speaker as positive and treat pair of different speakers as negative. For phonetic correlation, we sample positive pairs from audio and video of the same sentence, and negative pairs from audio and video containing different sentences. Note that contrastive learning usually uses cosine distance to measure the gap between positive and negative pairs. But it may be hard to precisely measure the distance between audio and visual embeddings in latent space due to the limitation of cosine distance. For example, it is required that the two embeddings are of the same shape, which is not always the case in practice. Therefore a different correlation learning approach is proposed in this work, which is based on adversarial training. Unlike previous contrastive learning, a classification neural network is trained as the discriminator to tell audio and visual embeddings apart. In experiments we find the proposed

---

method outperforms conventional audio-visual speech separation approaches. Note that our method won't affect the model complexity and efficiency during inference stage since correlation learning is only applied in the training process.

In summary, our main contributions are threefold:

1. Besides the speaker identity correlation modelled in [33], we propose to explicitly model the phonetic correlation between audio (phoneme) and video (lip motion).

2. In addition to contrastive learning used in previous work, we propose an adversarial training approach to learn identity and phonetic audio-visual correlations.

3. We launch detailed experiments in public datasets and visualize the correlation score after the proposed correlation learning. This verifies the effectiveness of the proposed framework.

2. Approach

2.1. Pipeline

The conventional audio-visual speech separation approach is an end-to-end system. We choose it as our baseline model (AV baseline). The whole train and test process can be formulated as a mix-and-separate pipeline (as shown in Fig. 1). In the training process, we randomly collect two audio segments $\mathbf{a}_1$ and $\mathbf{a}_2$, Then the mixed audio $\mathbf{x}$ can be simply obtained by adding them together. All the audio signal we use in experiment is in the STFT (Short-time Fourier transform) domain $\mathbb{C}^{T_a \times F}$, where T_a and F are time and frequency dimension respectively.

The input to model is two-stream signal, including the mixed audio $\mathbf{x}$ and one target speaker's video frames $\mathbf{v}$. $\mathbf{v}$ composes face tracks on rgb space $\mathbb{R}^{W \times H \times C}$ and lip motion frames on black and white space $\mathbb{R}^{W \times H}$, where W and H are the width and height of the image. C denotes the channel number. The number of video frames is T_v. The audio input is processed by an audio encoder to obtain the audio embedding. The face tracks and lip motion frames are processed by two video encoders respectively to get the video embeddings. The two kinds of extracted embeddings are reshaped to be aligned in temporal dimension T for concatenation (see Fig. 1 left part).

Subsequently, the audio-visual embedding is fed into a backbone neural network. Instead of outputting the target audio directly, we predict the mask $\mathbf{m}$ within $0 \sim 1$. If the target audio is $\mathbf{a}_1$, the predicted audio can be obtained by element-wise multiplying mixed audio with the mask:

$$\mathbf{a}_1 = \mathbf{m}_1 * \mathbf{x} \tag{1}$$

2.2. Audio-visual correlation enhancement

The baseline model simply concatenates the audio and video information as input, without formally modelling the correlation between audio and video. Intuitively, there are two kinds of available correlations as mentioned in Sec. 1. The first one is speaker *identity correlation*, which is the connection between timbre and facial attributes. The second one is *phonetic correlation* between phoneme and lip motion. In order to utilize these correlations in latent space, we use several pretrained models to extract the target embeddings. Specifically, we use a speaker identification model [33] for audio and a face identification model [33] for video frames to obtain the audio and visual identity embeddings. Then we apply a speech recognition model [36] for audio and a lip reading model [37] for video

frames to extract phonetic and lip motion embeddings. After that, our goal is to pull the correlated embeddings closer and keep uncorrelated embeddings away from each other (correlation enhancement), which can be achieved by either contrastive learning or adversarial training based approach. The two correlation enhancement methods are described as below in detail.

Contrastive learning approach

For identity correlation learning, we first follow [33] and [35] to adopt a triplet loss with cosine distance $d(.)$. In practice, a group of triplets denoted as $(\mathbf{a}_{\mathcal{A}_1}, \mathbf{a}_{\mathcal{A}_2}, \mathbf{a}_{\mathcal{B}})$ will be constructed from the dataset. Here $\mathbf{a}_{\mathcal{A}_1}$ and $\mathbf{a}_{\mathcal{A}_2}$ are two audio segments randomly sampled from speaker $\mathcal{A}$, $\mathbf{a}_{\mathcal{B}}$ is an audio segment sampled from speaker $\mathcal{B}$. We regard the $(\mathbf{a}_{\mathcal{A}_1}, \mathbf{a}_{\mathcal{A}_2})$ as positive pair and $(\mathbf{a}_{\mathcal{A}_1}, \mathbf{a}_{\mathcal{B}})$ as negative pair to construct the triplet. After extracting the corresponding speaker identity embedding $\mathbf{i}^a$ from audio and $\mathbf{i}^v$ from video through pretrained model, we apply the triplet loss $\mathcal{L}_1$ with margin m to speaker identity embedding as follows (as proposed by [33]).

$$\mathcal{L}_1 = \max\{d(\mathbf{i}^a_{\mathcal{A}_1}, \mathbf{i}^v_{\mathcal{A}_2}) - d(\mathbf{i}^a_{\mathcal{A}_1}, \mathbf{i}^v_{\mathcal{B}}) + m, 0\} \tag{2}$$

Then we seek to explicitly model the correlation between lip motion and phoneme. The phonetic embedding in video (lip motion) $\mathbf{p}^v$ is extracted by a pretrained lip reading model. And we regard the output of feature extraction layers in a speech recognition model as phoneme embedding $\mathbf{p}^v$ in audio (phoneme). The triplet loss $\mathcal{L}_2$ with margin m can also be applied as above. The difference is that only two audio segments A and B are needed for this loss:

$$\mathcal{L}_2 = \max\{d(\mathbf{p}^a_{\mathcal{A}}, \mathbf{p}^v_{\mathcal{A}}) - d((\mathbf{p}^a_{\mathcal{A}}, \mathbf{p}^v_{\mathcal{B}}) + m, 0\} \tag{3}$$

Via minimizing $\mathcal{L}_1$ and $\mathcal{L}_2$, we can pull the correlated embeddings closer while pushing others away from each other.

$$\mathcal{L} = \mathcal{L}_1 + \mathcal{L}_2 \tag{4}$$

Although triplet loss is able to deal with such a correlation learning task, it might not always be the best solution. Specifically, it relies on handcrafted metrics to learn the gap between positive and negative data pairs, which is usually a cosine distance as in [33]. One main limitation of cosine distance is that the magnitude of vectors is not taken into account, while merely their direction information is included. In practice, this implies that the difference of values is not fully considered, which might lead to severe inaccuracy in some cases. Here we are inspired by adversarial training method and propose a new correlation learning approach that can learn metrics automatically to evaluate a task.

Adversarial training approach

In particular, we train a discriminator D whose loss function $\mathcal{L}_{\mathcal{D}}$ is to distinguish between the audio and visual speaker identity embeddings. And we design the generator (backbone network) trained with loss function $\mathcal{L}_{\mathcal{G}}$ to inhibit the two embeddings from being classified by the discriminator:

$$\mathcal{L}_{\mathcal{G}} = \min_G \mathbb{E}_{\mathbf{x} \sim \mathbf{i}^v} \log(D(\mathbf{x})) + \mathbb{E}_{\mathbf{x} \sim \mathbf{i}^a} \log(1 - D(\mathbf{x})) \tag{5}$$

$$\mathcal{L}_{\mathcal{D}} = \max_D \mathbb{E}_{\mathbf{x} \sim \mathbf{i}^v} \log(D(\mathbf{x})) + \mathbb{E}_{\mathbf{x} \sim \mathbf{i}^a} \log(1 - D(\mathbf{x})) \tag{6}$$

We train the backbone network and discriminator alternately. Finally the discriminator is unable to classify the two
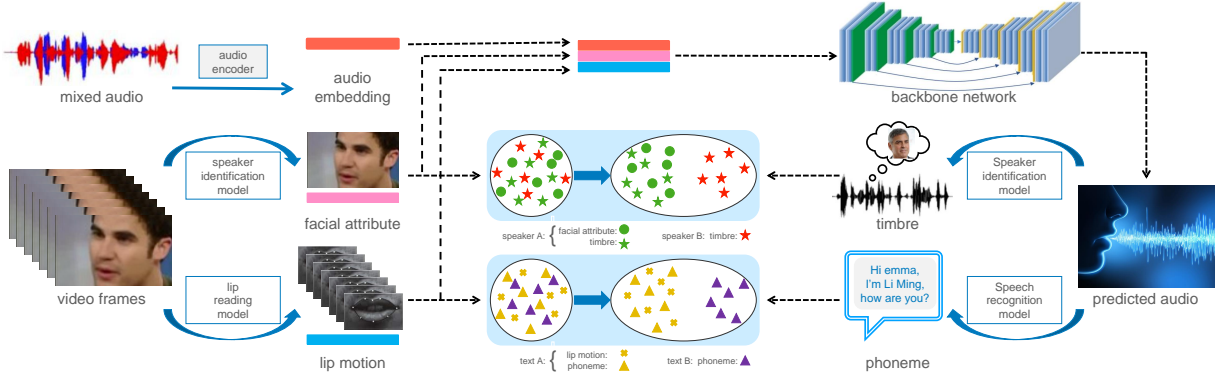
Figure 1: *The pipeline of the proposed method. A concatenation of audio and visual embeddings is put into the backbone neural network. Two pretrained models are used to extract the speaker identity and phoneme embedding from target speaker's predicted audio. These two pairs of embedding will be refined in latent space with contrastive learning or adversarial training method to maximize the correlation between the same speaker and text.*

embedding in latent space, indicating that the speaker identity in two modalities tend to be consistent. The phonetic correlation between lip motion and phoneme embedding can be enhanced in the same way.

We propose this method based on a very simple but clear intuition: any hand-crafted loss function has its limitation.The experimental results confirm our intuition: for all the dataset, the proposed adversarial training method outperforms contrastive learning method.

In addition, it should be emphasized that all our approaches are tested under the same complexity as baseline method. Although we use a number of pretrained models in our method, in test stage, only the baseline model is needed. For example, we use wave2vec2.0 pretrained model [36] in contrastive learning method (or use a discriminator to discriminate two embeddings in adversarial training method). Both of the correlation enhancement modules will be discarded during inference. Therefore our method won't increase model complexity compared with previous work.

3. Experiment

3.1. Setup

The whole network structure can be divided into the following parts: a backbone network (U-net) for separation, a lip reading network [37] for extracting lip motion embedding, a speech recognition network wav2vec2.0 [36] for recovering phoneme information, and two resnet18 classification network [33] for speaker identification. All the models are pretrained following the above work. In adversarial training framework, all the discriminators are MLP-based. We adopt SI-SNR [38] as the main loss function and metrics. Signal-to-Distortion Ratio (SDR), Signal-to-Interference Ratio (SIR), Signalto-Artifacts Ratio (SAR), PESQ (Perceptual Evaluation of Speech Quality) [39] and STOI (Short-Time Objective Intelligibility) [40] are also measured. We cut audio segments into 2.56s with a sampling rate of $16000Hz$, the corresponding video clip contains 64 frames with fps=25. We apply PyTorch official STFT and iSTFT function with nfft=512, hop size=160, window size=512 and center=True. We use 4 Nvidia V100 GPUs with 32GB for train and test.

	SDR	PESQ	STOI
[35](AV Baseline)	8.46	2.27	0.843
[35](CMC loss)	8.85	2.39	0.854
Ours(AV baseline)	9.392	2.536	0.851
Ours(triplet)	**9.623**	**2.545**	**0.855**
Ours(adversarial)	**9.982**	**2.584**	**0.861**

Table 1: *Results on LRS3. The first and second rows are baseline version without correlation learning and their proposed CMC loss in [35].*

3.2. Dataset

VoxCeleb2 [41] contains over 1 million utterances for 6,112 celebrities in YouTube. **LRS3** [42] consists of thousands of spoken sentences from TED and TEDx videos. We split train and validation set from development set with the ratio of 4:1. For testing our model, we randomly synthesize 500 and 200 audio mixtures from test set respectively.

3.3. Results

The result on VoxCeleb2 and LRS3 datasets is given by Table 1 and Table 2. Note that since the train & test split of VoxCeleb2 dataset are not given by [33], we make a split on our own. For fairness and clarity, we show all the results including the one reported by [33] based on their train & test split, the result on our split with [33]'s officially released model and the result of our implemented version based on [33]'s official code. We can see our proposed approach achieves consistent improvement on all the metrics. In comparison with [35] on LRS3 dataset, without any detail given for their split, we hold the official partition for our experiments. Without explicitly modelling identity/phonetic correlations, the CMC loss [35] apply triplet loss directly on audio and video embeddings. As far as we know, our method can outperform the State-of-the-Art speech separation method implemented on VoxCeleb2 and LRS3 datasets [33] [35]. Also, adversarial training proves to be more effective than triplet loss.

In both VoxCeleb2 and LRS3, we achieve better performance with adversarial training than contrastive learning, which proves our former intuition. With an automatically learned metric instead of a handcrafted one, adversarial training shows its

	SDR	SIR	SAR	PESQ	STOI	SI-SNR
[33](Reported)	10.2	17.2	11.3	2.83	0.87	-
[33](Released)	7.023	13.708	9.546	2.569	0.792	6.471
[33](Our impl.)	7.962	14.347	10.195	2.579	0.791	7.467
Ours(triplet)	**8.178**	**14.692**	**10.38**	**2.6**	**0.793**	**7.676**
Ours(adversarial)	**8.949**	**16.012**	**10.79**	**2.687**	**0.811**	**8.477**

Table 2: *Results on VoxCeleb2. The first row is the result reported in [33]. Because they don't show how to construct the test set, we use their officially released model to test on our test set (the second row). The third row is our implementation of [33] using official code based on our device. The last two rows are our method's results using the same hyper-parameters as above.*

potential and flexibility compared with triplet loss used in contrastive learning.
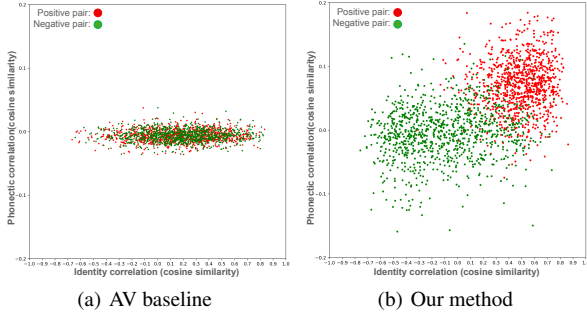

(a) AV baseline   (b) Our method

Figure 2: *Visualization of positive and negative pairs' correlations. (a) shows the two correlations of AV baseline (without correlation learning). (b) shows the correlations after joint identity & phonetic correlation learning.*

3.4. Correlation visualization

To prove whether the proposed approach learns to enhance audio-visual correlations, a visualization of identity & phonetic correlations before (AV baseline) and after adding correlation learning modules is shown by Fig. 2. Particularly, we randomly sample positive and negative audio segment pairs from the test set. The same as defined in Sec. 2.2, for identity correlation the positive pairs belong to the same speaker, for phonetic correlation the positive pairs represent two audio segments with the same phoneme information. And the negative pairs are vice versa. Fig. 2's horizontal and vertical axes respectively indicate the identity and phonetic correlations (measured by cosine similarity) of the sample pairs. If a sampled pair shows higher similarity, it implies the model learns to discover stronger correlation. Hence the positive pairs are expected to get a higher score while negative pairs get lower. As can be seen from Fig. 2 (a), in AV baseline the positive and negative pairs are hard to distinguished along either the identity or phonetic correlation axis. On the other hand, as shown in Fig. 2 (b), the positive pairs show much stronger correlations on both axes after correlation learning while negative pairs are still as before. This indicates that our proposed correlation learning surely enhances phonetic correlation between correlated audio-visual representations besides the identity correlation learning of [33]. Obviously this will make the positive samples more separable from those negative ones and thus benefit speech separation task.

To further validate the effectiveness of our phonetic correlation learning, we plot the SI-SNR score and phonetic correlation for all positive pairs in Fig. 3. Vertical and horizontal axes rep-

resent the SI-SNR score and the phonetic correlation measured by cosine similarity respectively. Fig. 3(a) shows the results obtained by AV baseline vs. only adding phonetic correlation learning (Ph). Fig. 3(b) shows identity correlation learning [33] (Id) vs. joint identity & phonetic correlation learning as proposed by this paper (Id+Ph). The more points in the right-top area, the better separation performance and correlation score is achieved. We can see [33] surely bring an improvement compared with the AV baseline (concatenation), and our proposed method can further acquire a more significant one. Such an improvement is in accordance with the enhanced phonetic correlation as the horizontal axis shows. In Fig. 3(b) there are denser blue points in the right-top area, which verifies the benefit of phonetic correlation learning.
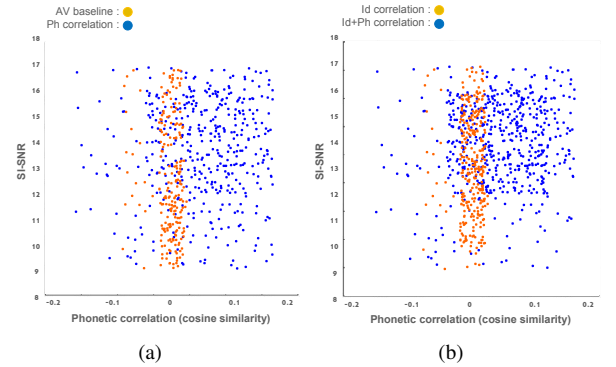

(a)   (b)

Figure 3: *Visualize the relationship between positive pairs' phonetic correlation and SI-SNR. (a) AV baseline (without correlation learning) vs. learning phonetic correlation (Ph). (b) learning identity correlation (Id) [33] and jointly learning both identity and phonetic correlation (Id+Ph).*

4. Conclusion

In this paper, we propose a novel audio-visual speech separation framework with explicit correlation learning. In addition to previous identity correlation learning method [33], we further propose a scheme with explicit phonetic correlation learning between audio phoneme and visual lip motion. Experimental results show our approach outperforms previous work with a clear margin. We conduct a detailed analysis of the audio-visual correlation and speech separation performance, which proves the effectiveness of explicitly learning phonetic correlation. Adversarial training also shows its potential and performs more effectively than contrastive learning. This work validates the necessity of explicit correlation learning in audio-visual speech separation, which also bears great potential for other tasks.

5. References

[1] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *2017 ICASSP*. IEEE, 2017, pp. 241–245.

[2] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *2016 ICASSP*. IEEE, 2016, pp. 31–35.

[3] S. Majumder, Z. Al-Halah, and K. Grauman, "Move2hear: Active audio-visual source separation," *arXiv:2105.07142*, 2021.

[4] L. Qu, C. Weber, and S. Wermter, "Multimodal target speech separation with voice and face references," *arXiv:2005.08335*, 2020.

[5] S.-W. Chung, S. Choe, J. S. Chung, and H.-G. Kang, "Facefilter: Audio-visual speech separation using still images," *arXiv:2005.07074*, 2020.

[6] H. Zhou, X. Xu, D. Lin, X. Wang, and Z. Liu, "Sep-stereo: Visually guided stereophonic audio generation by associating source separation," in *ECCV*. Springer, 2020, pp. 52–69.

[7] E. Tzinis, S. Wisdom, A. Jansen, S. Hershey, T. Remez, D. P. Ellis, and J. R. Hershey, "Into the wild with audioscope: Unsupervised audio-visual separation of on-screen sounds," *arXiv:2011.01143*, 2020.

[8] J. Lee, S.-W. Chung, S. Kim, H.-G. Kang, and K. Sohn, "Looking into your speech: Learning cross-modal affinity for audio-visual speech separation," in *CVPR*, 2021, pp. 1336–1345.

[9] Y. Tian, D. Hu, and C. Xu, "Cyclic co-learning of sounding object visual grounding and sound separation," in *CVPR*, 2021, pp. 2745–2754.

[10] T.-D. Truong, C. N. Duong, T. De Vu, H. A. Pham, B. Raj, N. Le, and K. Luu, "The right to talk: An audio-visual transformer approach," *arXiv:2108.03256*, 2021.

[11] K. Tan, Y. Xu, S.-X. Zhang, M. Yu, and D. Yu, "Audio-visual speech separation and dereverberation with a two-stage multimodal network," *JSTSP*, vol. 14, no. 3, pp. 542–553, 2020.

[12] C. Li and Y. Qian, "Listen, watch and understand at the cocktail party: Audio-visual-contextual speech separation." in *Interspeech*, 2020, pp. 1426–1430.

[13] ——, "Deep audio-visual speech separation with attention mechanism," in *2020 ICASSP*. IEEE, 2020, pp. 7314–7318.

[14] Y. Li, Z. Liu, Y. Na, Z. Wang, B. Tian, and Q. Fu, "A visual-pilot deep fusion for target speech separation in multitalker noisy environment," in *2020 ICASSP*. IEEE, 2020, pp. 4442–4446.

[15] C. Gan, D. Huang, H. Zhao, J. B. Tenenbaum, and A. Torralba, "Music gesture for visual sound separation," in *CVPR*, 2020, pp. 10 478–10 487.

[16] L. Zhu and E. Rahtu, "Visually guided sound source separation using cascaded opponent filter network," in *ACCV*, 2020.

[17] A. Rouditchenko, H. Zhao, C. Gan, J. McDermott, and A. Torralba, "Self-supervised audio-visual co-segmentation," in *2019 ICASSP*. IEEE, 2019, pp. 2357–2361.

[18] P. Zhang, J. Xu, Y. Hao, B. Xu *et al.*, "Audio-visual speech separation with adversarially disentangled visual representation," *arXiv:2011.14334*, 2020.

[19] T. Ochiai, M. Delcroix, K. Kinoshita, A. Ogawa, and T. Nakatani, "Multimodal speakerbeam: Single channel target speech extraction with audio-visual speaker clues." in *INTERSPEECH*, 2019, pp. 2718–2722.

[20] R. Lu, Z. Duan, and C. Zhang, "Audio–visual deep clustering for speech separation," *TASLP*, vol. 27, no. 11, pp. 1697–1712, 2019.

[21] J. Wu, Y. Xu, S.-X. Zhang, L.-W. Chen, M. Yu, L. Xie, and D. Yu, "Time domain audio visual speech separation," in *2019 ASRU*. IEEE, 2019, pp. 667–673.

[22] X. Xu, B. Dai, and D. Lin, "Recursive visual sound separation using minus-plus net," in *ICCV*, 2019, pp. 882–891.

[23] H. Zhao, C. Gan, W.-C. Ma, and A. Torralba, "The sound of motions," in *ICCV*, 2019, pp. 1735–1744.

[24] R. Gao and K. Grauman, "Co-separating sounds of visual objects," in *ICCV*, 2019, pp. 3879–3888.

[25] A. Gabbay, A. Ephrat, T. Halperin, and S. Peleg, "Seeing through noise: Visually driven speaker separation and enhancement," in *2018 ICASSP*. IEEE, 2018, pp. 3051–3055.

[26] M. Gogate, A. Adeel, R. Marxer, J. Barker, and A. Hussain, "Dnn driven speaker independent audio-visual mask estimation for speech separation," *arXiv:1808.00060*, 2018.

[27] F. U. Khan, B. P. Milner, and T. Le Cornu, "Using visual speech information in masking methods for audio speaker separation," *TASLP*, vol. 26, no. 10, pp. 1742–1754, 2018.

[28] R. Gao, R. Feris, and K. Grauman, "Learning to separate object sounds by watching unlabeled video," in *ECCV*, 2018, pp. 35–53.

[29] A. Ephrat, I. Mosseri, O. Lang, T. Dekel, K. Wilson, A. Hassidim, W. T. Freeman, and M. Rubinstein, "Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation," *arXiv:1804.03619*, 2018.

[30] M. Gogate, K. Dashtipour, P. Bell, and A. Hussain, "Deep neural network driven binaural audio visual speech separation," in *2020 IJCNN*. IEEE, 2020, pp. 1–7.

[31] R. Gu, S.-X. Zhang, Y. Xu, L. Chen, Y. Zou, and D. Yu, "Multimodal multi-channel target speech separation," *JSTSP*, vol. 14, no. 3, pp. 530–541, 2020.

[32] R. Lu, Z. Duan, and C. Zhang, "Listen and look: Audio–visual matching assisted speech source separation," *IEEE Signal Processing Letters*, vol. 25, no. 9, pp. 1315–1319, 2018.

[33] R. Gao and K. Grauman, "Visualvoice: Audio-visual speech separation with cross-modal consistency," *arXiv:2101.03149*, 2021.

[34] T. Afouras, J. S. Chung, and A. Zisserman, "The conversation: Deep audio-visual speech enhancement," *arXiv:1804.04121*, 2018.

[35] N. Makishima, M. Ihori, A. Takashima, T. Tanaka, S. Orihashi, and R. Masumura, "Audio-visual speech separation using cross-modal correspondence loss," in *2021 ICASSP*. IEEE, 2021, pp. 6673–6677.

[36] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *arXiv:2006.11477*, 2020.

[37] N. Ma, X. Zhang, H.-T. Zheng, and J. Sun, "Shufflenet v2: Practical guidelines for efficient cnn architecture design," in *ECCV*, 2018, pp. 116–131.

[38] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "Sdr–half-baked or well done?" in *2019 ICASSP*. IEEE, 2019, pp. 626–630.

[39] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," in *2001 ICASSP*, vol. 2. IEEE, 2001, pp. 749–752.

[40] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, "An algorithm for intelligibility prediction of time–frequency weighted noisy speech," *IEEE/ACM TASLP*, vol. 19, no. 7, pp. 2125–2136, 2011.

[41] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," in *INTERSPEECH*, 2018.

[42] T. Afouras, J. S. Chung, and A. Zisserman, "Lrs3-ted: a large-scale dataset for visual speech recognition," in *arXiv:1809.00496*, 2018.