



VCSE: Time-Domain Visual-Contextual Speaker Extraction Network

Junjie Li¹, Meng Ge^{1,2,*}, Zexu Pan², Longbiao Wang^{1,*}, Jianwu Dang^{1,3}

¹ Tianjin Key Laboratory of Cognitive Computing and Application,

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

² Department of Electrical and Computer Engineering, National University of Singapore, Singapore

³ Japan Advanced Institute of Science and Technology, Ishikawa, Japan

{mrjunjieli, gemeng, longbiao.wang}@tju.edu.cn, pan.zexu@u.nus.edu

Abstract

Speaker extraction seeks to extract the target speech in a multi-talker scenario given an auxiliary reference. Such reference can be auditory, i.e., a pre-recorded speech, visual, i.e., lip movements, or contextual, i.e., phonetic sequence. References in different modalities provide distinct and complementary information that could be fused to form top-down attention on the target speaker. Previous studies have introduced visual and contextual modalities in a single model. In this paper, we propose a two-stage time-domain visual-contextual speaker extraction network named VCSE, which incorporates visual and self-enrolled contextual cues stage by stage to take full advantage of every modality. In the first stage, we pre-extract a target speech with visual cues and estimate the underlying phonetic sequence. In the second stage, we refine the pre-extracted target speech with the self-enrolled contextual cues. Experimental results on the real-world Lip Reading Sentences 3 (LRS3) database demonstrate that our proposed VCSE network consistently outperforms other state-of-the-art baselines.

Index Terms: Speaker extraction, time-domain, phonetic sequence, visual-contextual

1. Introduction

Speaker extraction aims to separate the speech of the target speaker from a multi-talker mixture signal, which is also known as the cocktail party problem [1]. This is a fundamental but crucial problem to solve in signal processing that benefits a wide range of downstream applications, such as hearing aids [2], active speaker detection [3], speaker localization [4], and automatic speech recognition (ASR) [5, 6]. Although human can easily do that, it is a huge challenge to realize it in machines.

Before the deep learning era, popular techniques were computational auditory scene analysis [7] and non-negative matrix factorization [8]. The prior studies have laid the foundation for recent progress. With the advent and success of deep learning, speech separation algorithms such as permutation invariant training (PIT) [9], deep clustering [10], wavesplit [11] and Conv-TasNet [12], tackle the cocktail party problem by separating every speaker out in the mixture signal. Although a great success, there is inherent ambiguity in speaker labeling of the separated signals. An auxiliary reference, such as a pre-recorded speech signal [13–16] or video frame sequence [17–20], can be used to solve the speaker ambiguity. The speaker extraction algorithm employs such auxiliary reference to form top-down attention on the target speaker and extracts its speech.

* Corresponding author.

Neuroscience studies [21, 22] suggest that human perception is multimodal. According to the *reentry* theory [23], multimodal information is processed in our brain in an interactive manner, which educates each other. At the cocktail party, we hear the voice of the person, observe the lip motions and understand the contextual relevance. The lip motions that are synchronized with the target speech help us better capture the places of articulation, and it is robust against acoustic noise. The contextual information connects the preceding or following parts of the speech, which helps fill up the current severely corrupted speech by inferring from the context. The information from different modalities complements each other and together forms effective communication [24–27]. Inspired by these prior studies, we aim to emulate such a human perception process, to utilize the visual and contextual cues.

A number of audio-visual speaker extraction works explore the visual and contextual information by using the viseme-phoneme mapping cues [28–30]. They encode the lip images into visemes using a visual encoder pre-trained on the lip reading task, in which each viseme maps to multiple phonemes. However, such phoneme information derived from visual images only is weak, and the network may be using more of the lip motions cues derived from the visemes.

There are also studies incorporating ASR derived phonemes to explore contextual information in the speech separation [31–33] algorithm. In [31], the authors propose a two-stage method. The first stage is to obtain separated speech signals from a mixture signal with PIT. Then, they estimate the contextual embedding and guide the model to obtain the final speech in the second stage. In [32], the authors use the speech mixture and visual cues to get contextual information. This work utilizes visual and contextual modalities for the first time. Both of these two works are frequency-domain methods. In this paper, we aim to find a solution in the time-domain, as time-domain methods usually outperform frequency-domain counterparts by avoiding the difficult phase estimation problem [12]. Instead of incorporating multiple modalities in a single model, we aim to introduce one modality in each single stage.

Motivated by the previous works, we propose a two-stage time-domain visual-contextual speaker extraction (VCSE) network. The VCSE network is conditioned on auxiliary visual reference only, but it makes use of both the visual lip movement cues and self-enrolled contextual cues in the extraction process. In the first stage, the network pre-extracts the target speech and estimates the underlying phonemes using a pre-trained ASR system. In the second stage, the pre-extracted speech is refined with the contextual cues encoded from the self-enrolled phonetic sequence. Experimental results on real-world Lip Reading Sentences 3 (LRS3) database [34] show that our VCSE network consistently outperforms other state-of-the-art baselines.

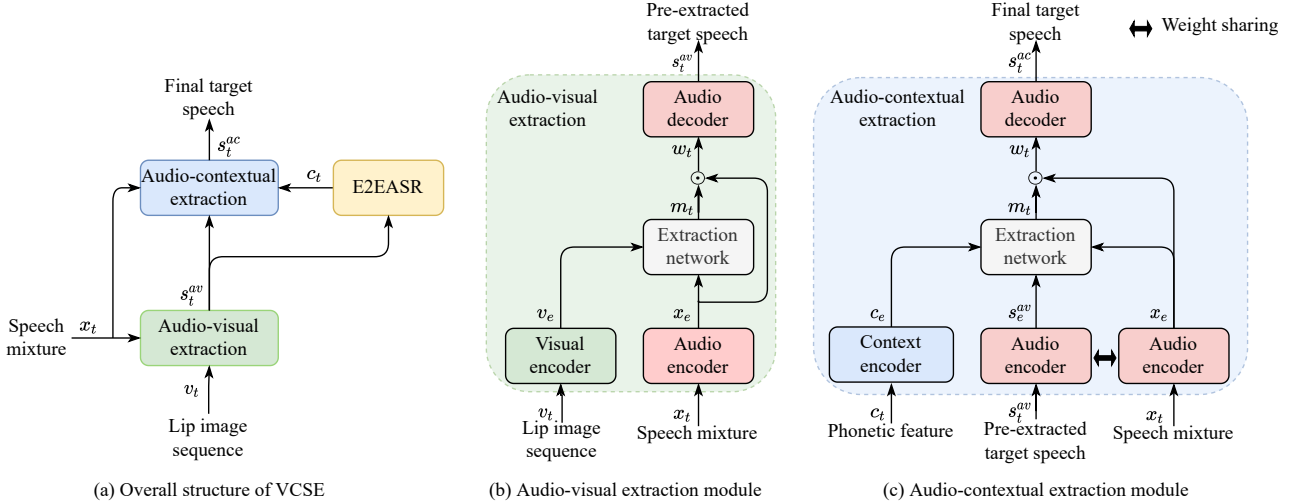


Figure 1: (a) The overall structure of VCSE. x_t, v_t, s_t^{av}, c_t and s_t^{ac} represent time-domain speech mixture, lip image sequence, pre-extracted target speech, phonetic feature, and final target speech, respectively. (b) The structure of the audio-visual extraction module. \odot denotes element-wise multiplication. (c) The structure of the audio-contextual extraction module.

2. VCSE network

We design a two-stage network to utilize visual and contextual cues completely. In the first stage, visual cues is introduced to pre-extract the target speech and estimate the underlying phonemes, because it is robust against acoustic noise. In the second stage, with the self-enrolled contextual cues encoded from phonetic embedding, the network can refine the current speech frame by inferring from the preceding and following parts of pre-extracted speech frames.

2.1. Network architecture

The VCSE network consists of an audio-visual extraction module, an end-to-end ASR (E2EASR) module, and an audio-contextual extraction module, as depicted in Fig. 1 (a). In the first stage, the pre-extracted speech s_t^{av} is extracted with the help of visual cues in the audio-visual extraction module, which takes the time-domain speech mixture x_t and lip image sequence v_t as inputs. The E2EASR module encodes s_t^{av} to produce phonetic embedding c_t . In the second stage, the audio-contextual module takes the phonetic embedding c_t , pre-extracted target speech s_t^{av} and speech mixture x_t as inputs to acquire final target speech s_t^{ac} .

2.1.1. Audio-visual extraction module

Audio-visual extraction module employs the visual cues, i.e., lip image sequences as reference to pre-extract the target speech from the speech mixture.

The audio-visual extraction module shares a similar structure with AV-ConvTasNet proposed in [28], which consists of four parts: video encoder, audio encoder, extraction network and audio decoder, as shown in Fig. 1(b). The audio encoder and decoder perform 1D convolution and de-convolution, respectively. The extraction network is a stack of several temporal dilated convolutional blocks (TCN). The visual encoder consists of a 3D convolution layer and an 18-layer ResNet. Different from [28] where the visual encoder is pre-trained on the lip reading task to capture viseme-phoneme mapping cues, we do not pre-train the visual encoder. We join-train the visual en-

coder with the other parts to let the network decide on its own intrinsically what is the best visual representation.

Considering a time-domain audio mixture signal x_t and lip image sequence v_t , the visual encoder encodes v_t into visual embedding v_e , which is time-aligned with the audio signal at frame-level. Audio encoder transforms the speech mixture x_t into latent embedding x_e . The extraction network takes v_e and x_e as inputs and estimates the mask m_t . The m_t is element-wise multiplied with the x_e to obtain the latent representation of the pre-extracted target speech w_t . The audio decoder renders pre-extracted target speech s_t^{av} from w_t .

2.1.2. E2EASR module

We adopt the encoder of OpenTransformer network¹ as our E2EASR module. The OpenTransformer has a transformer architecture [35] and is trained with the connectionist temporal classification (CTC) loss function using a clean speech signal, as depicted in Fig. 2. The output of transformer encoder c_t is the phonetic feature.

2.1.3. Audio-contextual extraction module

The audio-contextual extraction module has a similar architecture with the audio-visual extraction module except for a context encoder replacing the visual encoder, as depicted in Fig. 1(c). Although sharing similar architecture, they do not share the model weights as the two modules perform different tasks, i.e., pre-extraction in the audio-visual extraction module and refinement in the audio-contextual extraction model.

The context encoder is a stack of five 1D convolutional blocks with the exponential growth dilation factor 2^d , where $d \in \{0, 1, 2, 3, 4\}$. A linear layer is followed to adjust the feature dimension. Each 1D conv block contains 1D conv with kernel size 5, channel size 256, and the exponential growth padding $\in \{2, 4, 8, 16, 32\}$, 1D batch normalization and ReLU activation function².

¹<https://github.com/ZhengkunTian/OpenTransformer>

²<https://github.com/pragyak412/Improving-Voice-Separation-by-Incorporating-End-To-End-Speech-Recognition>

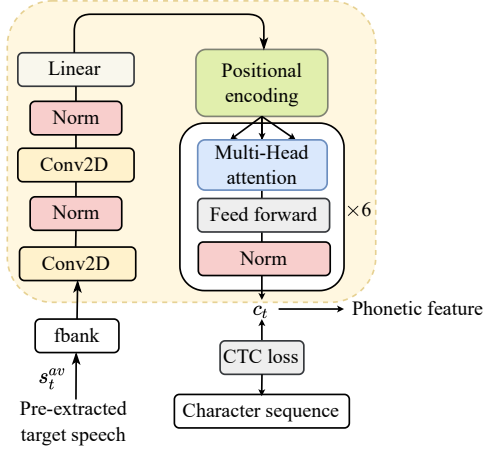


Figure 2: Illustration of E2EASR module.

The context encoder is designed to acquire the contextual embedding c_e from the phonetic feature c_t . Two audio encoders encode two audio signals into the latent representations s_e^{av} , x_e , and share the same weights. According to [13], it is helpful for separation to project the speech mixture and reference speech into the same latent space. These three features are concatenated and modeled by the extraction network to estimate mask m_t . The audio decoder decodes audio representation w_t to final target speech s_t^{ac} .

2.2. Loss function

We adopt the scale-invariant signal-to-noise ratio (SI-SNR) [12,36] as loss function to train the audio-visual extraction module and the audio-contextual extraction module in speech signal reconstruction, which is defined as:

$$\begin{cases} s_{target} := \frac{\hat{s}^T s}{\|\hat{s}\|^2} s \\ e_{noise} := \hat{s} - s_{target} \\ \mathcal{L}_{SI-SNR}(s, \hat{s}) := -10 \log_{10} \frac{\|s_{target}\|^2}{\|e_{noise}\|^2} \end{cases} \quad (1)$$

where s and \hat{s} denote the clean target speech and an estimated speech, respectively.

2.3. Training strategy

Our training process is divided into 5 steps.

1. We train the audio-visual extraction module alone with the SI-SNR loss function $\mathcal{L}_{SI-SNR}(s, s_t^{av})$.
2. We pre-train the OpenTransformer network, which forms the E2EASR module in our VCSE network. The pre-trained OpenTransformer network reaches 9.5% word error rate on the LRS3 test sets.
3. We fix the weights of the audio-visual extraction module and the E2EASR module, and train the audio-contextual extraction module. It is worth noting that in this step of training, we use the phonetic feature c_t encoded from the clean target speech instead of the pre-extracted speech. This enables the network to converge faster with such oracle phonetic feature. The SI-SNR loss function $\mathcal{L}_{SI-SNR}(s, s_t^{ac})$ is used in this step.
4. We repeat the training in step 3, except that the phonetic feature c_t in this step is encoded from the pre-extracted speech s_t^{av} . The SI-SNR loss function $\mathcal{L}_{SI-SNR}(s, s_t^{ac})$ is used in this step.

5. We release all of the fixed module weights and fine-tune the whole system end-to-end. The SI-SNR loss function $\mathcal{L}_{SI-SNR}(s, s_t^{ac})$ is used in this step.

3. Experimental setup

3.1. Datasets

In this paper, all modules are trained on the LRS3 [34] dataset. This dataset contains thousands of spoken sentences from TED and TEDx videos. There are 118,516 (252h), 31,982 (30h) and 1,321 (0.85h) utterances in pre-train, trainval and test sets, respectively, and there is no overlap between the videos used to create the test set and the ones used for the pre-train and trainval sets. The sampling rate of the audio signal is 16kHz.

Automatic speech recognition: The original LRS3 pre-train set is used to pre-train the OpenTransformer network. We select utterances that are less than 20 seconds to pre-train the module to fit into GPU memory.

Speaker extraction: We simulate a two-speaker mixture sets to train and evaluate our VCSE network, using the LRS3 datamix script³. The long utterances are truncated to 3 seconds, and the utterances less than 3 seconds are dropped. The two-speaker mixtures are generated by selecting different utterances and mixed at various signal-to-noise ratios (SNR) between -5dB and 5dB. The corresponding video stream is sampled in 25 frames per second (fps). The lip region of each video frame is detected by face recognition algorithm, and we resize it to 120*120 pixels. The numbers of utterances for training, validation and test are 50,000, 5,000, 3,000 respectively.

3.2. Training setup

For audio-visual extraction module and audio-contextual extraction module, adam is used as an optimizer. And the initial learning rate is set to 0.001. Besides, the learning rate is halved if the validation loss increases consecutively for 3 epochs. The training process stops when validation loss increases consecutively for 6 epochs.

For E2EASR module, we use CTC as loss function. We select adam as an optimizer. Warmup strategy [35] is adopted to adjust learning rate, which first increases the learning rate linearly and then decreases thereafter proportionally to the inverse square root of the step number. The length of each phonetic embedding is 73, and the dimension of channel is 256.

4. Experimental results

We evaluate the system's performance using scale-invariant signal-to-noise ratio improvement (SI-SNRi) and signal-to-distortion ratio improvement (SDRi) [37], the higher the better for both metrics.

4.1. Comparison with speech separation baselines

We compare the VCSE with two speech separation algorithms, Conv-TasNet (PIT) [12] and Multi-stage-AC model [31] as shown in Table 1, in which the latter model makes use of the self-enrolled contextual cues. The speech separation algorithms require knowing the number of speakers in the mixture.

The Conv-TasNet (PIT) model is a time-domain speech separation model that originally proposes the use of temporal convolutional neural network in estimating the intrinsic masks

³<https://github.com/JusperLee/LRS3-For-Speech-Separation>

Table 1: Results of various models using visual information (V), speaker embedding (A_S) or contextual information (C) as reference in terms of SI-SNRi (dB) and SDRi (dB). “BSS” and “SE” denote blind source separation and speaker extraction, respectively. “Reference” denotes the reference stimulus or auxiliary stimulus. Star (*) marks that the model is established by ourselves according to original papers. “Oracle” means oracle phonetic feature.

Task	Model	Reference	SI-SNRi	SDRi
BSS	Conv-TasNet (PIT)* [12]	-	11.4487	11.7377
	Multi-stage-AC* [31]	C	14.4009	14.6527
SE	A_S -ConvTasNet	A_S	11.2973	11.7457
	AV-ConvTasNet* [28]	V	14.5356	14.7627
	AC-ConvTasNet	C (Oracle)	15.6915	15.9126
	AVC-ConvTasNet	V+C (Oracle)	14.8687	15.0889
SE	VCSE	V+C	15.8527	16.0800

for every speaker. The Multi-stage-AC model is the two-stage method, which uses the Conv-TasNet (PIT) model to pre-separate the speech mixture and use an ASR model to extract the corresponding contextual features in the first stage. In the second stage, a TASNet incorporates the contextual feature and speech mixture to obtain the final estimated speech.

It is seen that the Multi-stage-AC model outperforms the Conv-TasNet (PIT) model due to the use of additional contextual cues. Our proposed VCSE outperforms the Multi-stage-AC model, as we utilize additional visual cues in pre-extracting the target speech in stage 1, which boosts the quality of the pre-extracted speech and self-enrolled contextual cues, therefore better higher signal quality for the final estimated speech. In addition, the VCSE network does not require prior knowledge of the number of speakers.

4.2. Comparison with speaker extraction baselines

We compare the VCSE network with other models using reference from different modalities as shown in Table 1. The A_S -ConvTasNet uses audio cues, i.e., a pre-recorded speech signal, to perform speaker extraction task. The audio cues mean an utterance taken from the target speaker that is not present in the current speech mixture. The AV-ConvTasNet follows [28], except that the visual encoder is not pre-trained on the lip reading task. Same as the VCSE, the AV-ConvTasNet model takes lip image sequence as reference to guide model extracting corresponding speech. The AC-ConvTasNet takes oracle phonetic embedding as reference for speaker extraction. This oracle phonetic embedding is comes from the E2EASR module which is applied on the clean speech. The AVC-ConvTasNet concatenates the representation of lip image sequence and oracle phonetic embedding as reference to extract target speech.

It’s worth noting that A_S -ConvTasNet performs worse than Conv-TasNet (PIT) method. In our training set, there are about 5,000 speakers, and the total number of simulated utterances is 50,000. There isn’t enough utterance from every speaker, which causes the network unable to distinguish each speaker’s characteristic under the current amount of dataset. On the other hand, the visual/contextual cues are not affected by the number of speakers, showing the advantage of visual/contextual cues are more robust against dataset variations.

Table 2: Comparisons of VCSE and VCSEv in terms of SI-SNRi (dB) and SDRi (dB). “V” and “C” represent visual and contextual auxiliary information. “Input to AC” denotes the input variable to Audio-Contextual extraction module.

Model	Reference	Input to AC	SI-SNRi	SDRi
VCSE	V+C	x_t, c_t, s_t^{av}	15.8527	16.0800
VCSEv	V+C	x_t, c_t	15.3576	15.5834

Our method significantly outperforms the A_S -ConvTasNet, as visual cues are much more robust against noise compared to the reference speech, not to mention that we also use the self-enrolled contextual cues. The VCSE provides 1.3 dB improvement over AV-ConvTasNet that uses visual cues alone. This indicates the usefulness of our self-enrolled contextual cues. Our proposed VCSE even outperforms the AC-ConvTasNet which makes use of oracle contextual information, which shows the importance of the visual cues in speaker extraction. The result of AVC-ConvTasNet shows combining multiple modalities in a single model can not take full advantage of multimodal information. Our proposed VCSE takes both advantage of visual and contextual information through a combination of an audio-visual extraction module and an audio-contextual extraction module.

4.3. Comparative study with variant of VCSE

To verify the complementary effect of s_t^{av} , we design a variant of VCSE that does not utilize pre-extracted speech in the audio-contextual extraction module, referred to as VCSEv. Table 2 shows the comparisons of VCSE and VCSEv. It shows that in the second stage, refining the s_t^{av} is a better option compared to re-extraction from the speech mixture. The final target speech s_t^{ac} of VCSE is derived from both visual and contextual cues, while s_t^{ac} of VCSEv actually derives from self-enrolled contextual cues. This indicates our effective use of both the visual and contextual cues.

5. Conclusions and future work

In this paper, we propose a novel time-domain visual-contextual speaker extraction (VCSE) network to leverage visual cues and self-enrolled contextual cues for speaker extraction. Unlike the previous methods incorporating multiple modalities in a single model, the VCSE network introduces visual information and contextual knowledge stage by stage to take full advantage of different modalities. The experimental results show that our model achieves significant improvement over the previous methods. In the future, we are considering to investigate the complementary effect in depth between visual and contextual modalities.

6. Acknowledgements

This research is supported by the National Key RD Program of China under Grant 2018YFB1305200, the internal project of Shenzhen Research Institute of Big Data under the Grant No. T00120220002, the Guangdong Provincial Key Laboratory of Big Data Computing under the Grant No. B10120210117- KP02, The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen) and the University Development Fund, CUHK-Shenzhen, under the Grant No. UDF01002333 and UF02002333.

7. References

- [1] A. W. Bronkhorst, “The cocktail party phenomenon: A review of research on speech intelligibility in multiple-talker conditions,” *Acta Acustica united with Acustica*, vol. 86, no. 1, pp. 117–128, 2000.
- [2] D. Wang, “Time-frequency masking for speech separation and its potential for hearing aid design,” *Trends in amplification*, vol. 12, no. 4, pp. 332–353, 2008.
- [3] 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 *Proc. of the 29th ACM Int. Conf. on Multimedia*, 2021, pp. 3927–3935.
- [4] M. Ge, C. Xu, L. Wang, E. S. Chng, J. Dang, and H. Li, “L-spex: Localized target speaker extraction,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7287–7291.
- [5] A. Narayanan and D. Wang, “Investigation of speech separation as a front-end for noise robust speech recognition,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 22, no. 4, pp. 826–835, 2014.
- [6] T. Yoshioka, H. Erdogan, Z. Chen, and F. Alleva, “Multi-microphone neural speech separation for far-field multi-talker speech recognition,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 5739–5743.
- [7] D. Wang and G. J. Brown, *Computational auditory scene analysis: Principles, algorithms, and applications*. Wiley-IEEE press, 2006.
- [8] M. N. Schmidt and R. K. Olsson, “Single-channel speech separation using sparse non-negative matrix factorization,” in *Proc. INTERSPEECH*, vol. 2. Citeseer, 2006, pp. 2–5.
- [9] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, “Permutation invariant training of deep models for speaker-independent multi-talker speech separation,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2017, pp. 241–245.
- [10] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2016, pp. 31–35.
- [11] N. Zeghidour and D. Grangier, “Wavesplit: End-to-end speech separation by speaker clustering,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 2840–2849, 2021.
- [12] Y. Luo and N. Mesgarani, “Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [13] M. Ge, C. Xu, L. Wang, E. S. Chng, J. Dang, and H. Li, “Spex+: A complete time domain speaker extraction network,” *arXiv preprint arXiv:2005.04686*, 2020.
- [14] C. Han, Y. Luo, C. Li, T. Zhou, K. Kinoshita, S. Watanabe, M. Delcroix, H. Erdogan, J. R. Hershey, N. Mesgarani *et al.*, “Continuous speech separation using speaker inventory for long multi-talker recording,” *arXiv preprint arXiv:2012.09727*, 2020.
- [15] C. Xu, W. Rao, E. S. Chng, and H. Li, “Spex: Multi-scale time domain speaker extraction network,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 28, pp. 1370–1384, 2020.
- [16] Q. Wang, H. Muckenhirn, K. Wilson, P. Sridhar, Z. Wu, J. Hershey, R. A. Saurous, R. J. Weiss, Y. Jia, and I. L. Moreno, “Voice-filter: Targeted speech separation by speaker-conditioned spectrogram masking,” *arXiv preprint arXiv:1810.04826*, 2018.
- [17] Z. Pan, R. Tao, C. Xu, and H. Li, “Selective listening by synchronizing speech with lips,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 30, pp. 1650–1664, 2022.
- [18] Z. Aldeneh, A. P. Kumar, B.-J. Theobald, E. Marchi, S. Kajarekar, D. Naik, and A. H. Abdelaziz, “On the role of visual cues in audio-visual speech enhancement,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2021, pp. 8423–8427.
- [19] 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 preprint arXiv:1804.03619*, 2018.
- [20] Z. Pan, X. Qian, and H. Li, “Speaker extraction with co-speech gestures cue,” *IEEE Signal Processing Letters*, 2022.
- [21] D. Remley, *The neuroscience of multimodal persuasive messages: Persuading the brain*. Routledge, 2017.
- [22] N. Ward, E. Paul, P. Watson, G. Cooke, C. Hillman, N. J. Cohen, A. F. Kramer, and A. K. Barbey, “Enhanced learning through multimodal training: evidence from a comprehensive cognitive, physical fitness, and neuroscience intervention,” *Scientific reports*, vol. 7, no. 1, pp. 1–8, 2017.
- [23] G. M. Edelman, *Neural Darwinism: The theory of neuronal group selection*. Basic books, 1987.
- [24] L. D. Rosenblum, “Speech perception as a multimodal phenomenon,” *Current Directions in Psychological Science*, vol. 17, no. 6, pp. 405–409, 2008.
- [25] Z. Pan, Z. Luo, J. Yang, and H. Li, “Multi-modal attention for speech emotion recognition,” in *Proc. INTERSPEECH*, 2020, pp. 364–368.
- [26] Z. Chen, S. Wang, and Y. Qian, “Multi-modality matters: A performance leap on voxceleb,” in *Proc. INTERSPEECH*, 2020, pp. 2252–2256.
- [27] Q. Song, B. Sun, and S. Li, “Multimodal sparse transformer network for audio-visual speech recognition,” *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–11, 2022.
- [28] J. Wu, Y. Xu, S.-X. Zhang, L.-W. Chen, M. Yu, L. Xie, and D. Yu, “Time domain audio visual speech separation,” in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2019, pp. 667–673.
- [29] Z. Pan, R. Tao, C. Xu, and H. Li, “MuSE: Multi-modal target speaker extraction with visual cues,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2021, pp. 6678–6682.
- [30] Z. Pan, M. Ge, and H. Li, “USEV: Universal speaker extraction with visual cue,” *arXiv preprint arXiv:2109.14831*, 2021.
- [31] N. Takahashi, M. K. Singh, S. Basak, P. Sudarsanam, S. Ganapathy, and Y. Mitsufuji, “Improving voice separation by incorporating end-to-end speech recognition,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 41–45.
- [32] C. Li and Y. Qian, “Listen, watch and understand at the cocktail party: Audio-visual-contextual speech separation,” in *Proc. INTERSPEECH*, 2020, pp. 1426–1430.
- [33] S. Petridis, T. Stafylakis, P. Ma, F. Cai, G. Tzimiropoulos, and M. Pantic, “End-to-end audiovisual speech recognition,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 6548–6552.
- [34] T. Afouras, J. S. Chung, and A. Zisserman, “Lrs3-ted: a large-scale dataset for visual speech recognition,” *arXiv preprint arXiv:1809.00496*, 2018.
- [35] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [36] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “SDR—half-baked or well done?” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2019, pp. 626–630.
- [37] E. Vincent, R. Gribonval, and C. Févotte, “Performance measurement in blind audio source separation,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 14, no. 4, pp. 1462–1469, 2006.