

SPEAKER ACTIVITY DRIVEN NEURAL SPEECH EXTRACTION

Marc Delcroix¹, Katerina Zmolikova^{2}, Tsubasa Ochiai¹, Keisuke Kinoshita¹, Tomohiro Nakatani¹*

¹ NTT Corporation, Japan,

² Brno University of Technology, Speech@FIT and IT4I Center of Excellence, Czechia

ABSTRACT

Target speech extraction, which extracts the speech of a target speaker in a mixture given auxiliary speaker clues, has recently received increased interest. Various clues have been investigated such as pre-recorded enrollment utterances, direction information, or video of the target speaker. In this paper, we explore the use of speaker activity information as an auxiliary clue for single-channel neural network-based speech extraction. We propose a speaker activity driven speech extraction neural network (ADENet) and show that it can achieve performance levels competitive with enrollment-based approaches, without the need for pre-recordings. We further demonstrate the potential of the proposed approach for processing meeting-like recordings, where the speaker activity is obtained from a diarization system. We show that this simple yet practical approach can successfully extract speakers after diarization, which results in improved ASR performance, especially in high overlapping conditions, with a relative word error rate reduction of up to 25 %.

Index Terms— Speech extraction, Speaker activity, Speech enhancement, Meeting recognition, Neural network

1. INTRODUCTION

Recognizing speech in the presence of interfering speakers remains one of the challenges for automatic speech recognition (ASR). A conventional approach to tackle this problem consists of separating the observed speech mixture into all of its source speech signals before ASR. Single-channel speech separation has greatly progressed with the introduction of deep learning [1, 2]. However, most separation approaches suffer from two limitations, (1) they require knowing or estimating the number of speakers in the mixture, and (2) they suffer from a global permutation ambiguity issue, i.e., an arbitrary mapping between source speakers and outputs.

Target speech extraction [3] has been proposed as an alternative to speech separation to alleviate the above limitations. It focuses on extracting only the speech signal of a speaker of interest or target speaker by exploiting auxiliary clues about that speaker. The problem formulation becomes thus independent of the number of speakers in the mixture. Besides, the global permutation ambiguity is naturally solved thanks to the use of auxiliary clues. Several target speech extraction schemes have been proposed, exploiting different types of auxiliary clues such as pre-recorded enrollment utterances of the target speaker [4, 5], direction information [6], video of the target speaker [7, 8], or electroencephalogram (EEG) signals [9].

For example, SpeakerBeam [4, 10] is an early approach for enrollment utterance-based target speech extraction. It exploits a

*Katerina Zmolikova was partly supported by Czech Ministry of Education, Youth and Sports from the National Programme of Sustainability (NPU II) project "IT4Innovations excellence in science - LQ1602"

speaker embedding vector derived from an enrollment utterance to inform a speech extraction network which speaker to extract from the mixture. With SpeakerBeam, the speaker embedding vectors are obtained from an auxiliary network that is jointly trained with the speech extraction network. Consequently, the speaker embedding vector may capture the voice characteristics of the target speaker, which are needed to distinguish that speaker in a mixture.

Alternative approaches that exploit video clues have been proposed subsequently [7, 8]. In these works, a sequence of face embedding vectors are extracted from the video of the face of the target speaker speaking in the mixture and passed to the extraction network. Different from SpeakerBeam, the auxiliary clue is time-synchronized with the mixture signal. What information the face embedding vectors actually capture remains unclear, but we can naturally assume that they capture the lip movements of the target speaker and that the mouth opening and closing regions may be important information provided to the speech extraction network.

Although target speech extraction can achieve a high level of performance [7, 11], it is not always possible to have access to the enrollment utterance or video. In this paper, inspired by the works on visual speech extraction, we investigate the use of another clue, which consists of the speech activity of a speaker. The speaker activity consists of a signal that indicates for each time frame if the speaker is speaking or not. It can be a general and practical clue for speech extraction as there are various ways to obtain it.

We propose a single-channel speaker activity driven speech extraction neural network (ADENet). We hypothesize that a neural network can exploit the speaker activity information to identify and extract the speech of a speaker, assuming that the speakers in the mixture do not always fully overlap. We experimentally demonstrate that assuming the availability of oracle speaker activity, ADENet can achieve superior performance to SpeakerBeam, without requiring pre-recorded enrollment speech.

In practice, there are various possible approaches to obtain the speaker activity information including visual-based voice activity detection (VAD) [12], personal VAD [13] or diarization [14–16]. For example, speaker diarization has greatly progressed [14–17], and finding speaker activity regions in meetings has achieved a low diarization error rate (DER) even in overlapping conditions. It does not require pre-recorded enrollment utterances of the speakers or video. Diarization is often used as a pre-processing for ASR. However, diarization itself does not extract the speech signals, which limits ASR performance in overlapping conditions. As a practical potential use-case of ADENet, we investigate using the speaker activity information obtained from a diarization system as clues for ADENet. We show that this simple approach can improve ASR performance in meeting-like single-channel recordings, especially in severe overlapping conditions. Note that our proposed ADENet does not depend on how the speaker activity information is obtained, making it a versatile speech extraction method.

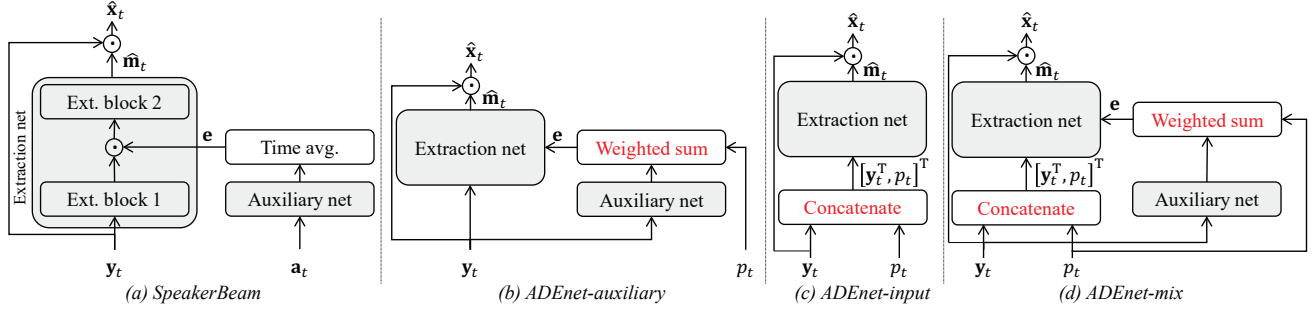


Fig. 1. Diagram of SpeakerBeam and different variations of ADENet. See Sect. 5.3 for details of the auxiliary and extraction blocks.

In the remainder of the paper, we first briefly review SpeakerBeam in Section 2, which serves as a basis for the proposed ADENet that we describe in Section 3. We discuss related works in Section 4. In Section 5, we present experimental results based on the LibriCSS corpus [18]. Finally, we conclude the paper in Section 6.

2. SPEAKERBEAM

Figure 1-(a) shows a schematic diagram of SpeakerBeam [10]. Let \mathbf{y}_t be an observed speech mixture, \mathbf{a}_t be an enrollment utterance, and \mathbf{x}_t be the speech signal of the target speaker. Here, we consider networks operating in the frequency domain, and the signals \mathbf{y}_t , \mathbf{x}_t , and \mathbf{a}_t are sequences of amplitude spectrum coefficients, and t represents the time-frame index. Note that we could derive the methods presented in this paper equivalently for time-domain networks [11].

SpeakerBeam is composed of two networks. First, an auxiliary network accepts the enrollment utterance and outputs a fixed dimension speaker embedding vector \mathbf{e} obtained as the time-average of the output of the auxiliary network,

$$\mathbf{e} = \frac{1}{T^a} \sum_{t=1}^{T^a} f(\mathbf{a}_t), \quad (1)$$

where $f(\cdot)$ represents the auxiliary network, and T^a is the number of frames of the enrollment utterance.

The second network computes a time-frequency mask, \mathbf{m}_t , given the mixture signal, \mathbf{y}_t , and the speaker embedding vector, \mathbf{e} , as $\mathbf{m}_t = g(\mathbf{y}_t, \mathbf{e})$, where $g(\cdot)$ represents the speech extraction network. The extracted speech signal, $\hat{\mathbf{x}}_t$, is then obtained as $\hat{\mathbf{x}}_t = \mathbf{m}_t \odot \mathbf{y}_t$, where \odot is the element-wise multiplication.

There are several approaches to combine the two networks using, e.g., feature concatenation, addition, or multiplication. We employ here the multiplicative combination [19], where the output of the first hidden layer of the speech extraction network $g(\cdot)$ is multiplied with the speaker embedding vector \mathbf{e} as proposed in [20]. We jointly train both networks to minimize the mean square error (MSE) between the estimated and reference target speech.

3. NEURAL ACTIVITY DRIVEN SPEECH EXTRACTION

SpeakerBeam has demonstrated high speech extraction performance. However, it requires a pre-recorded enrollment utterance that may not always be available. Moreover, there may be a mismatch between the enrollment utterance and the target speech in the mixture due to, e.g., different recording conditions.

Here, we assume that we can have access to the speaker activity information instead of the pre-recorded utterance. Let $p_t \in \{0, 1\}$ be a signal representing the activity of a speaker. p_t takes a value of 1 when the speaker is speaking and 0 otherwise. We consider

two cases, i.e. (1) speaker activity *with overlap*, where p_t contains speaker active regions, including regions where the interference speakers overlap with the target speaker, and (2) speaker activity *without overlap*, where we remove the overlap regions from the speaker active regions based on the activity information of the interference speakers.

We first describe three configurations of ADENet to exploit the speaker activity for target speech extraction shown in Fig. 1(b),(c),(d). We then discuss how to obtain target speaker activity in practice. Finally, we discuss a training strategy to make ADENet robust to estimation errors in the speaker activity signals.

3.1. Three configurations of the proposed ADENet

The first configuration, shown in Fig. 1-(b), exploits a similar architecture as SpeakerBeam, but replaces the enrollment utterance by the input mixture and the time-averaging operation of Eq. (1) by a weighted-sum as,

$$\mathbf{e} = \frac{1}{\sum_t p_t} \sum_t p_t f(\mathbf{y}_t). \quad (2)$$

This formulation is similar to SpeakerBeam using instead of pre-recorded enrollment utterance the regions of the mixture signal where the target speaker is active. If we can remove the overlapping regions based on the speaker activity information (i.e., activity without overlap), the enrollment utterance consists of the regions where the target speaker is the single active speaker. This naturally avoids any mismatch in the recording conditions between the enrollment and the mixture. We call this approach *ADENet-auxiliary*.

An alternative option consists of simply concatenating the speaker activity p_t and the speech mixture \mathbf{y}_t at the input of the network. This approach does not compute any explicit speaker embedding vector but expects that the extraction network can learn to track and identify the target speaker internally based on its activity. Consequently, it does not use any auxiliary network. This approach is shown in Fig. 1-(c). We call it *ADENet-input*.

The third configuration is shown in Fig. 1-(d). It consists of a combination of both approaches, where we use the speaker activity at the input and also to compute the speaker embedding vector with Eq. (2). We call this approach *ADENet-mix*.

3.2. Speaker activity information

ADENet requires the speech activity of the target speaker in the mixture. We aim at building a versatile system that can perform speech extraction independently of the way we obtain the speaker activity.

In practice, there are several possible options to obtain the speaker activity. Naturally, a conventional VAD [21], would not work in this context because it could not discriminate the target speaker from the interference in a mixture. However, if we have

access to video recordings of the target speaker, we could use a visual-based VAD to detect the speaker activity [12, 22]. If enrollment utterances are available, we can obtain the speaker activity with a personal VAD [13]. If none of the above are available, we can use diarization [14, 16, 17] to obtain the speech activity of the speakers in a recording.

As an example use-case, we investigate using ADEnet in a meeting scenario. We use diarization to obtain the speech activity of all speakers in a meeting. We then apply ADEnet to each speaker identified by the diarization system. This builds upon the great progress of recent diarization systems that can achieve a low DER in meeting-like situations [15–17]. Combining diarization with ADEnet offers a practical yet relatively simple approach for enhancing ASR performance in overlapping conditions.

However good recent VAD or diarization systems are, we cannot expect to obtain oracle speaker activity information. Besides, depending on the scenario, it may not be possible to obtain speaker activity without overlap. It is thus essential to develop a system robust to errors in the speaker activity detection. In the next section, we explain a training strategy making ADEnet robust to errors in the activity signals.

3.3. Training of AEDnet

We require the triplet of the observed mixture \mathbf{y}_t , the reference target speech \mathbf{x}_t , and the target speaker activity p_t to train ADEnet. We use simulated speech mixtures and extract the *oracle speaker activity* by applying a conventional VAD approach [21] on the reference target speech. Besides, we use speaker activity without overlap, to allow the system to better capture the characteristics of the target speaker.

Using oracle speaker activity information obtained from the reference signals allows us to build a system independent of the method used for estimating the speaker activity at test time. However, it is unrealistic to assume the availability of oracle error-free speaker activity information. Therefore, we include noise to the speaker activity information at training time to make the system robust. Practically, for each oracle speech segment found by the VAD, we modify the segment boundaries by adding to the start and end times a value uniformly sampled between -1 and 1 seconds.

4. RELATED WORK

There have been several works that integrated speech activity into non-neural multi-channel speech separation systems to extract speech. [23–26]. For example, supervised independent vector analysis (IVA) modifies the objective function of IVA by including an extra-term related to the target speaker, such as the source activity, allowing the system to extract that speaker [25].

The diarization information has also been combined with spatial clustering-based separation in the guided source separation (GSS) framework [26]. GSS introduces the speaker activity as a prior for the mask estimation of the complex angular central Gaussian mixture model (cACGMM) based source separation, enabling identifying the target source from the cACGMM outputs. GSS has been applied successfully to severe recording conditions [26]. Note that supervised IVA and GSS require multi-microphone recordings, while our proposed ADEnet can work with a single microphone.

If multi-microphone recordings are available, we could extend GSS with ADEnet by using time-frequency masks of the sources obtained from ADEnet instead of the speaker activity as priors to cACGMM as proposed in [27] for speech separation. Such an extension will be part of our future works.

5. EXPERIMENTS

We performed two types of experiments, with (1) simulated two-speaker mixtures to show the potential of using speaker activity for speech extraction and (2) meeting-like recordings to illustrate a potential use-case of ADEnet as a pre-processor for ASR.

5.1. Dataset

We used the same training data for all experiments. It consisted of 63000 mixtures of reverberant two-speaker mixtures with background noise at SNR between 10 and 20 dB. The speech signals were taken from the LibriSpeech corpus [28]. We further created a validation and test set with the same conditions. The test set consisted of 1364 mixtures with an averaged overlapping ratio of 38.5 %.

In the second experiment, we used the meeting-like LibriCSS corpus [18], which consists of 8-speaker meeting-like recordings sessions of 10 minutes, obtained by re-recording LibriSpeech utterances played through loudspeakers in a meeting room. The overlap ratio varies from 0 to 40 %.

5.2. Auxiliary information

For training, we used the “oracle” speaker activity without overlap obtained by applying the webrtc VAD [21] to the reference signals. We trained the system with or without *noisy activity training* (adding noise to the speaker activity signals, as described in 3.3). In the first experiment with simulated two-speaker mixtures, we used the oracle speaker activity with and without noise. In the second experiment with LibriCSS, we used target speaker VAD (TS-VAD) based diarization [16] to obtain the speech activity regions of each speaker.

5.3. Experimental settings

In all experiments, we used a similar architecture for the extraction network, which consisted of 3 bidirectional long short-term memory (BLSTM) layers with 1200 units, and two fully connected (FC) layers with rectified linear unit (ReLU) activation functions. The auxiliary network consisted of two FC layers with 64 hidden units and ReLU activation functions. We used the same configuration for the SpeakerBeam network. The input features consisted of 257-dimensional log amplitude spectrum coefficients obtained using a short-time Fourier transform with a window of 32 msec and a shift of 8 msec. The training loss consisted of the MSE between the spectra of the reference signal and the extracted speech. We trained the networks using the ADAM optimizer using the padertorch toolkit [29].

5.4. Experiments on simulated two-speaker mixtures

We first performed experiments on the simulated two-speaker mixtures. Table 1 shows the signal to distortion ratio (SDR) [30] obtained with ADEnet with and without noisy activity training and for different activity signals at test time. At test time, we used oracle speaker activity signals without overlap as during training (i.e. “w/o overlap”) without, and with adding noise, i.e. “Oracle” and “+Noise”, respectively. We also explored training and test mismatch by using speaker activity with overlap at test time (i.e. “w/ overlap”). The SDR of the mixture signals is -0.2 dB. We compare the proposed ADEnet with SpeakerBeam using pre-recorded enrollment utterances. We used a randomly selected utterance of the target speaker from the LibriSpeech corpus as the enrollment utterance. SpeakerBeam achieved an SDR of 9.4 dB. Besides, masking out the regions

Table 1. Speech enhancement results in terms of SDR [dB] for simulated 2-speaker data with ADENet without and with noisy activity training. The SDR of the mixture signal is -0.2 dB, and the SDR of SpeakerBeam is 9.4 dB.

	Noisy activity training	Activity signal at test time				Avg
		w/o overlap Oracle	+Noise	w/ overlap Oracle	+Noise	
ADENet-aux	-	9.7	6.0	6.0	4.3	6.5
ADENet-in	-	10.2	4.3	3.7	2.8	5.3
ADENet-mix	-	10.1	3.9	3.7	3.0	5.2
ADENet-aux	✓	8.4	8.5	7.9	8.0	8.2
ADENet-in	✓	9.0	9.2	7.7	8.6	8.6
ADENet-mix	✓	9.2	9.4	7.9	8.6	8.8

where the speaker is inactive from the observed signal (“inactivity-masking”) leads to an SDR or 3.8 dB with oracle speaker activity, and -1.6 dB when using noisy speaker activity.

The upper part of Table 1 shows the results without noisy activity training. The first column confirms that ADENet with oracle speaker activity can achieve higher SDR than SpeakerBeam (10.2 dB vs. 9.4 dB) without the need for pre-recorded enrollment utterance. Moreover, it greatly outperforms the simple “inactivity-masking”. These results demonstrate that ADENet can effectively extract speakers based on the speaker activity. However, the SDR degrades significantly when precise speaker activity is not available as revealed by the poor results when including noise or overlap regions.

The lower part of Table 1 shows the results with noisy activity training. The oracle performance degrades, but the systems become much more robust to errors in the activity signals. Overall, the performance in noisy conditions becomes a little worse than SpeakerBeam but still much better than using “inactivity-masking”. ADENet-mix with noisy training achieves superior performance on average, and we use this configuration in the following experiments. Figure 2 shows an example of extracted speech, showing that ADENet can extract a speaker even in the overlapping regions.

This experiment demonstrates that single-channel speech extraction is possible without explicit use of enrollment utterances or video. Moreover, the relatively stable performance with various quality of the activity signals indicates that ADENet could be used with various diarization or VAD methods that provide or not speaker regions without overlap.

5.5. Experiments on LibriCSS meeting-like data

We present results on the LibriCSS meeting-like corpus, using the TS-VAD-based diarization to obtain the speech activity regions of all speakers in the meeting. We use TS-VAD as an example of diarization, but other diarization approaches could be used such as, e.g., [14, 15]. TS-VAD achieved a DER of 7.3 % [31]. We use the ESPnet toolkit and the transformer-based recipe developed for LibriCSS [32, 33] as ASR back-end, which provides a strong baseline for the task [31].

Diarization finds the speech activity for the speakers in the recording. From the diarization output, we process each continuous speech segment separately, by adding 2 seconds of context at the beginning and end to create extended segments. We used the speaker activity estimated by the diarization within the extended segment as speaker activity signal. Besides, we also experimented with speaker activity without overlap by removing from the activity signal the regions where multiple speakers were detected. ASR is performed separately on each speech segments found by the diarization and

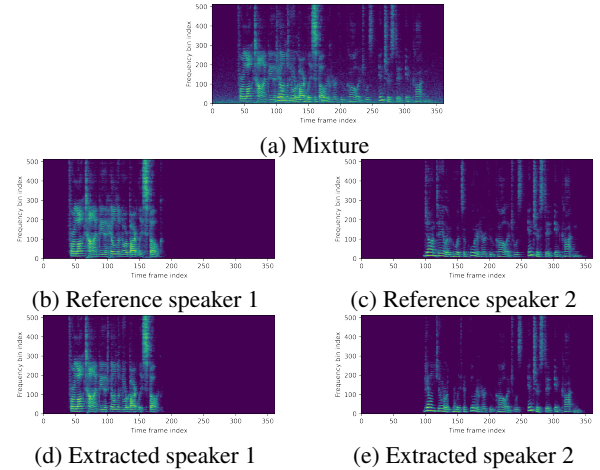


Fig. 2. Example of extracted speech with ADENet-mix.

Table 2. cpWER using TS-VAD-based diarization, the proposed ADENet and a Transformer-based ASR backend.

	w/o overlap	Overlap ratio in %						Avg
		0L	0S	OV10	OV20	OV30	OV40	
no proc	na	11.2	9.4	16.2	23.1	33.6	41.1	23.9
ADENet	-	10.5	8.5	13.8	18.7	26.5	30.7	19.2
ADENet	✓	10.5	8.6	13.6	18.3	25.8	29.8	18.8

evaluated using the concatenated minimum-permutation word error rate (cpWER) [34] that includes the errors caused by diarization.

Table 2 shows the cpWER obtained with the diarized output with and without ADENet. Note that ADENet operates on single-channel recordings. Even if we used a powerful diarization, we see the degradation caused by the overlapping speech. We confirm that of ADENet consistently improves cpWER, especially for the high overlapping regions, with relative WER improvement ranging from 6.5 % to 25 %. These results should be put into perspective with those obtained with a multi-channel system using GSS [26] with seven microphones, which achieved an average WER of 11.2 %. Compared to GSS, ADENet has the advantage that it can work with single-channel recordings. The integration of ADENet with GSS is possible and will be part of our future works.

6. CONCLUSION

We investigated the potential of using speaker activity to develop a simple and versatile neural speech extraction. We showed experimentally that the proposed ADENet with precise speaker activity information could achieve high single-channel speech extraction performance. Besides, combined with diarization, it can significantly improve the ASR performance in meeting-like situations.

For further improvements, we plan to investigate more powerful network configurations [11] and combine ADENet with multi-microphone schemes [26, 27]. Besides, we would like to compare ADENet with visual clue based speech extraction approaches.

Acknowledgment. This work was started at JSALT 2020 at JHU, with support from Microsoft, Amazon, and Google. We thank Desh Raj, Pavel Denisov, and Shinji Watanabe for providing the LibriCSS baseline.

7. REFERENCES

- [1] M. Kolbaek, D. Yu, Z. Tan, and J. Jensen, “Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks,” *IEEE/ACM Trans. ASLP*, vol. 25, no. 10, pp. 1901–1913, 2017.
- [2] Y. Luo and N. Mesgarani, “TasNet: Surpassing ideal time-frequency masking for speech separation,” in *Proc. of ICASSP’18*, 2018.
- [3] K. Zmolikova, M. Delcroix, K. Kinoshita, T. Higuchi, A. Ogawa, and T. Nakatani, “Speaker-aware neural network based beamformer for speaker extraction in speech mixtures,” in *Proc. of Interspeech’17*, 2017, pp. 2655–2659.
- [4] K. Zmolikova, M. Delcroix, K. Kinoshita, T. Higuchi, A. Ogawa, and T. Nakatani, “Learning speaker representation for neural network based multichannel speaker extraction,” in *Proc of ASRU’17*, 2017, pp. 8–15.
- [5] J. Chen, D. Su, L. Chen, M. Yu, Y. Qian, and D. Yu, “Deep extractor network for target speaker recovery from single channel speech mixtures,” in *Proc. of Interspeech’18*, 2018, pp. 307–311.
- [6] R. Gu, L. Chen, S.-X. Zhang, J. Zheng, Y. Xu, M. Yu, D. Su, Y. Zou, and D. Yu, “Neural spatial filter: Target speaker speech separation assisted with directional information,” in *Proc. Interspeech 2019*, 2019, pp. 4290–4294.
- [7] 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,” *ACM Trans. on Graphics (TOG)*, vol. 37, no. 4, pp. 1–11, 2018.
- [8] T. Afouras, J. S. Chung, and A. Zisserman, “The conversation: Deep audio-visual speech enhancement,” in *Proc. Interspeech 2018*, 2018, pp. 3244–3248.
- [9] E. Ceolini, J. Hjortkjær, D. D. Wong, J. O’Sullivan, V. S. Raghavan, J. Herrero, A. D. Mehta, S.-C. Liu, and N. Mesgarani, “Brain-informed speech separation (biss) for enhancement of target speaker in multitalker speech perception,” *NeuroImage*, vol. 223, pp. 117282, 2020.
- [10] K. Zmolikova, M. Delcroix, K. Kinoshita, T. Ochiai, T. Nakatani, L. Burget, and J. Cernocky, “SpeakerBeam: Speaker aware neural network for target speaker extraction in speech mixtures,” *IEEE JSTSP*, vol. 13, no. 4, pp. 800–814, 2019.
- [11] M. Delcroix, T. Ochiai, K. Zmolikova, K. Kinoshita, N. Tawara, T. Nakatani, and S. Araki, “Improving speaker discrimination of target speech extraction with time-domain speakerbeam,” in *Proc. of ICASSP’20*, 2020.
- [12] P. Liu and Z. Wang, “Voice activity detection using visual information,” in *Proc. of ICASSP’04*, 2004, vol. 1, pp. I–609.
- [13] S. Ding, Q. Wang, S.-Y. Chang, L. Wan, and I. Lopez Moreno, “Personal VAD: Speaker-Conditioned Voice Activity Detection,” in *Proc. of Odyssey’20*, 2020, pp. 433–439.
- [14] D. Garcia-Romero, D. Snyder, G. Sell, D. Povey, and A. McCree, “Speaker diarization using deep neural network embeddings,” *Proc. of ICASSP’17*, pp. 4930–4934, 2017.
- [15] Z. Huang, S. Watanabe, Y. Fujita, P. García, Y. Shao, D. Povey, and S. Khudanpur, “Speaker diarization with region proposal network,” *Proc. of ICASSP’20*, pp. 6514–6518, 2020.
- [16] I. Medennikov, M. Korenevsky, T. Prisyach, Y. Y. Khokhlov, M. Korenevskaya, I. Sorokin, T. V. Timofeeva, A. Mitrofanov, A. Andrusenko, I. Podluzhny, A. Laptev, and A. Romanenko, “Target-speaker voice activity detection: a novel approach for multi-speaker diarization in a dinner party scenario,” *ArXiv*, vol. abs/2005.07272, 2020.
- [17] Y. Fujita, S. Watanabe, S. Horiguchi, Y. Xue, and K. Nagamatsu, “End-to-end neural diarization: Reformulating speaker diarization as simple multi-label classification,” *ArXiv*, vol. abs/2003.02966, 2020.
- [18] Z. Chen, T. Yoshioka, L. Lu, T. Zhou, Z. Meng, Y. Luo, J. Wu, and J. Li, “Continuous speech separation: Dataset and analysis,” *Proc. of ICASSP’20*, pp. 7284–7288, 2020.
- [19] L. Samarakoon and K. C. Sim, “Subspace LHUC for fast adaptation of deep neural network acoustic models,” in *Proc. of Interspeech’16*, 2016, pp. 1593–1597.
- [20] M. Delcroix, K. Zmolikova, T. Ochiai, K. Kinoshita, S. Araki, and T. Nakatani, “Compact network for SpeakerBeam target speaker extraction,” in *Proc. of ICASSP’19*, 2019, pp. 6965–6969.
- [21] “<https://github.com/wiseman/py-webrtcvad>,”.
- [22] D. Soderoy, B. Rivet, L. Girin, J.-L. Schwartz, and C. Jutten, “An analysis of visual speech information applied to voice activity detection,” in *Proc. of ICASSP’06*, 2006, vol. 1, pp. I–I.
- [23] B. Rivet, L. Girin, and C. Jutten, “Visual voice activity detection as a help for speech source separation from convolutive mixtures,” *Speech Communication*, vol. 49, no. 7-8, pp. 667–677, 2007.
- [24] T. Ono, N. Ono, and S. Sagayama, “User-guided independent vector analysis with source activity tuning,” in *Proc. of ICASSP’12*, 2012, pp. 2417–2420.
- [25] F. Nesta, S. Mosayyebpour, Z. Koldovský, and K. Paleček, “Audio/video supervised independent vector analysis through multimodal pilot dependent components,” in *Proc. of EU-SIPCO’17*, 2017, pp. 1150–1164.
- [26] C. Bøddeker, J. Heitkaemper, J. Schmalenstroer, L. Drude, J. Heymann, and R. Haeb-Umbach, “Front-end processing for the CHiME-5 dinner party scenario,” in *Proc of Interspeech’18*, 2018.
- [27] T. Nakatani, N. Ito, T. Higuchi, S. Araki, and K. Kinoshita, “Integrating DNN-based and spatial clustering-based mask estimation for robust MVDR beamforming,” in *Proc. of ICASSP’17*, 2017, pp. 286–290.
- [28] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *Proc. of ICASSP’15*, 2015, pp. 5206–5210.
- [29] “<https://github.com/fgnt/padertorch>,”.
- [30] E. Vincent, R. Gribonval, and C. Févotte, “Performance measurement in blind audio source separation,” *IEEE trans. ASLP*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [31] D. Raj, P. Denisov, Z. Chen, H. Erdogan, Z. Huang, M. He, S. Watanabe, J. Du, T. Yoshioka, Y. Luo, N. Kanda, J. Li, S. Wisdom, and J. R. Hershey, “Integration of speech separation, diarization, and recognition for multi-speaker meetings: System description, comparison, and analysis,” in *Submitted to SLT’21*, 2021.
- [32] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. Enrique Yalta Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, “ESPnet: End-to-end speech processing toolkit,” in *Proc. of Interspeech’18*, 2018, pp. 2207–2211.
- [33] “https://github.com/espnet/espnet/tree/master/egs/libri_css/asr1,”.
- [34] S. Watanabe, M. Mandel, J. Barker, and E. Vincent, “CHiME-6 challenge: Tackling multispeaker speech recognition for unsegmented recordings,” *ArXiv*, vol. abs/2004.09249, 2020.