



Strategies to Improve Robustness of Target Speech Extraction to Enrollment Variations

Hiroshi Sato, Tsubasa Ochiai, Marc Delcroix, Keisuke Kinoshita, Takafumi Moriya,
Naoki Makishima, Mana Ihori, Tomohiro Tanaka, Ryo Masumura

NTT Corporation, Japan

hiroshi.satou.bh@hco.ntt.co.jp

Abstract

Target speech extraction is a technique to extract the target speaker's voice from mixture signals using a pre-recorded enrollment utterance that characterize the voice characteristics of the target speaker. One major difficulty of target speech extraction lies in handling variability in "intra-speaker" characteristics, i.e., characteristics mismatch between target speech and an enrollment utterance. While most conventional approaches focus on improving *average performance* given a set of enrollment utterances, here we propose to guarantee the *worst performance*, which we believe is of great practical importance. In this work, we propose an evaluation metric called worst-enrollment source-to-distortion ratio (SDR) to quantitatively measure the robustness towards enrollment variations. We also introduce a novel training scheme that aims at directly optimizing the worst-case performance by focusing on training with difficult enrollment cases where extraction does not perform well. In addition, we investigate the effectiveness of auxiliary speaker identification loss (SI-loss) as another way to improve robustness over enrollments. Experimental validation reveals the effectiveness of both worst-enrollment target training and SI-loss training to improve robustness against enrollment variations, by increasing speaker discriminability.

Index Terms: target speech extraction, tail performance, noise robust speech recognition, speaker embeddings, SpeakerBeam

1. Introduction

Although recent advances in deep learning technology improved automatic speech recognition (ASR) accuracy drastically, it remains difficult to recognize the overlapping speech of multiple speakers, which naturally occurs in realistic situations. Speech enhancement front-end has widely been investigated in combination with ASR back-end to recognize speech in overlapping conditions [1–5].

There are two major approaches explored as speech enhancement front-end. One approach is speech separation where the mixture of speech signals are separated into signals from each source speaker [6–9]. Speech separation enables ASR of all speakers in the speech mixture and has shown promising performance even with a single channel setup. However, this type of approach may suffer from the global permutation ambiguity problem i.e. arbitrary mapping between source and output signals. The other speech enhancement approach, target speech extraction (TSE), naturally solves the permutation issue by focusing only on the signal from the target speaker [10–19]. TSE extracts the target speaker's speech from a mixture by using auxiliary information about the target. Pre-recorded target's speech is commonly used for the auxiliary information [10–15], which we call an enrollment utterance hereafter. TSE is suitable for applications that only require recognizing a particular target

speaker as those deployed in personalized devices like smartphones and smart speakers.

Despite the advantage of obtaining the target speaker-specific signal, TSE has the potential challenge of performance deviation, depending on the choices of enrollment utterances. In general, speaker-characteristics are different between an enrollment utterance and the target speech in the mixture since the tone or intonation of speech varies depending on when we speak. This "intra-speaker" speaker-characteristic variance can result in wrong extraction of the non-target speech. This is a specific problem for TSE and not for speech separation, because TSE has to perform a more challenging task of separation and identification.

In previous works [10], extraction performance was typically evaluated by performing extraction with a *single* enrollment utterance prepared for each mixture signal. Thus, the performance variation with respect to enrollment choices for the same mixture was not fully considered. However, the evaluation of the performance variation over enrollment choices is an important issue, because it is practically essential to perform TSE robustly against every possible choice of enrollment. Especially, preventing the extraction failure caused by a specific enrollment choice is a significant challenge.

To this end, we propose a novel evaluation metric *worst-enrollment source-to-distortion ratio (SDR)* as a quantitative measure of the worst performance over enrollment choices. To evaluate the worst-enrollment SDR, we prepare multiple enrollment utterances as enrollment candidates for each mixture in evaluation data, and perform TSE with every candidate. The worst-enrollment SDR is defined as the worst SDR value over enrollment candidates.

In this work, we investigate two approaches to improve worst-enrollment performance. First, to directly optimize the TSE model for the worst-enrollment SDR, we propose *Worst-enrollment target training* that focuses on the difficult enrollment cases. Specifically, worst-enrollment target training calculates the training loss for the enrollment utterance whose extraction performance is the worst within a set of several enrollment utterances selected for each mixture. As another approach, we investigated speaker identification loss (SI-loss) [20, 21] as well as its combination with the worst-enrollment target training. The SI-loss was initially proposed to increase speaker discriminability by introducing an auxiliary speaker identification task to TSE. They may help learn to extract more robust embeddings with less intra-speaker variance, which could also result in improving performance on worse enrollment cases, although prior works [20, 21] have not evaluated this.

In our results, both the proposed worst SDR-loss and SI-loss are shown to be effective to increase speaker discriminability and improve worst-enrollment SDR by 0.9 dB on average, by the combination of these approaches. Moreover, the proposed

approach drastically reduced the overall extraction failure ratio (defined as the ratio of utterances whose SDR improvement is less than 5 dB) by 34% relative.

Our contribution is threefold:

- We propose an evaluation metric of worst-enrollment SDR to quantitatively measure the worst performance over enrollment choices.
- We propose a novel worst-enrollment target training to directly optimize the TSE model for the worst-enrollment SDR criterion.
- We evaluate the effectiveness of SI-loss as well as our proposal on the worst-enrollment SDR and revealed the effectiveness of both methods for improving robustness toward enrollment variations.

2. Conventional Method

In this section, we describe the TSE model adopted in this work as well as the conventional way of training it.

2.1. Overview of target speech extraction (TSE)

In this work, we assume that a mixture signal $\mathbf{Y} \in \mathbb{R}^T$ is captured by a single channel microphone as $\mathbf{Y} = \mathbf{S} + \mathbf{I} + \mathbf{N}$ where $\mathbf{S} \in \mathbb{R}^T$, $\mathbf{I} \in \mathbb{R}^T$, and $\mathbf{N} \in \mathbb{R}^T$ denote the target speaker's speech, interfering speaker's speech, and background noise, respectively, and T denotes the number of time samples.

Target speech extraction $\text{TSE}(\cdot)$ aims at extracting target speech from the observed mixture signal with auxiliary information about the target to inform which speaker to extract. In this work, we use an enrollment utterance of the target speaker as the auxiliary information, which is denoted as $\mathbf{C}^S \in \mathbb{R}^{T_s}$ where T_s is the number of samples. We assume that one utterance of enrollment is available in the decoding phase. Given \mathbf{C}^S , the target signal $\hat{\mathbf{S}} \in \mathbb{R}^T$ is estimated from the observed mixture as:

$$\hat{\mathbf{S}} = \text{TSE}(\mathbf{Y}, \mathbf{C}^S; \Lambda) \quad (1)$$

where Λ denotes learnable parameters of the TSE network.

In this work, we adopted time-domain SpeakerBeam as a TSE architecture [20]. Time-domain SpeakerBeam is a deep neural network based speech enhancement method that is composed of two sub-networks. *Speaker embedding network* $\text{Embed}(\cdot)$ is designed to extract the latent representation of the target speaker $\mathbf{e}^S \in \mathbb{R}^D$ from the waveform signal of an enrollment utterance \mathbf{C}^S as follows:

$$\mathbf{e}^S = \text{Embed}(\mathbf{C}^S; \theta_{\text{emb}}) \quad (2)$$

where θ_{emb} denotes the learnable parameters of the speaker embedding network and D is the dimension of the representation. *Extraction network* $\text{Extract}(\cdot)$ extracts target speech from a mixture signal given the speaker embedding as follows:

$$\hat{\mathbf{S}} = \text{Extract}(\mathbf{Y}, \mathbf{e}^S; \theta_{\text{ext}}) \quad (3)$$

where θ_{ext} denotes the learnable parameters of the extraction network.

2.2. Conventional training method

The learnable parameters $\Lambda = \{\theta_{\text{emb}}, \theta_{\text{ext}}\}$ are jointly optimized to get optimal speaker representation space with respect to TSE. The optimization target is a signal level distance between estimated target speech $\hat{\mathbf{S}}$ and ground truth target speech \mathbf{S} . Specifically, we adopt scale-dependent SDR $\text{SDR}(\mathbf{S}, \hat{\mathbf{S}}) = 10 \log_{10}(\|\hat{\mathbf{S}}\|^2 / \|\mathbf{S} - \hat{\mathbf{S}}\|^2)$ as the distance measure. The SDR loss is denoted as $\mathcal{L}_{\text{sdr}}(\mathbf{S}, \mathbf{Y}, \mathbf{C}^S) =$

$-\text{SDR}(\mathbf{S}, \text{TSE}(\mathbf{Y}, \mathbf{C}^S; \Lambda))$. Thus, we need paired data of mixture signal, target speech, and enrollment utterance $(\mathbf{S}, \mathbf{Y}, \mathbf{C}^S)$ for the optimization.

To increase the variety of the combination of mixture signal and enrollment speech, multiple enrollment utterances can be prepared for each mixture in training data, and one of them is selected for training at every epoch [10]. Formally, the enrollment utterance \mathbf{C}^S is randomly or iteratively chosen from a set of candidates $\mathcal{C} = \{\tilde{\mathbf{C}}_1^S, \dots, \tilde{\mathbf{C}}_N^S\}$ where N denotes the number of enrollments prepared for each mixture signal. In this case, the loss function \mathcal{L} is given by

$$\mathcal{L} = \mathcal{L}_n \text{ with } n \sim \mathcal{U}\{1, \dots, N\} \quad (4)$$

where \mathcal{L}_n denotes loss calculated for n -th enrollment defined as $\mathcal{L}_n = \mathcal{L}_{\text{sdr}}(\mathbf{S}, \mathbf{Y}, \tilde{\mathbf{C}}_n^S)$ and $\mathcal{U}\{1, \dots, N\}$ denotes the uniform distribution that takes integer values between 1 to N .

3. Strategies to improve the worst performance

In this section, we first define the worst-enrollment SDR to quantitatively evaluate the robustness against enrollment variation, and then explain training strategies to improve it.

3.1. Worst-enrollment SDR

To evaluate extraction performance focusing on the robustness over enrollment variability, we propose worst-enrollment SDR metrics. We prepare N enrollment utterances \mathcal{C} for each mixture for evaluation, as in training. SDR are evaluated for each extracted speech using each enrollment utterance. n -th worst-enrollment SDR is defined as n -th worst value within N enrollment choices, averaged over mixtures. Formally, the worst-enrollment SDR l is described as follows:

$$l = \min_{n \in \mathcal{C}} \text{SDR}(\mathbf{S}, \text{TSE}(\mathbf{Y}, \tilde{\mathbf{C}}_n^S; \Lambda)). \quad (5)$$

The same definition holds for other metrics than SDR such as character error rate (CER).

3.2. Worst-enrollment target training

In the previous training method, enrollment utterances are sampled from candidates randomly or iteratively as in Eq. (4). Thus, all prepared enrollments are treated equally in training. However, it may not be optimal in terms of the worst-enrollment SDR. To directly optimize TSE performance on worst-enrollment SDR, we propose worst-enrollment target training that trains the model with weight placed on difficult enrollments.

For worst-enrollment target training, $K (\leq N)$ enrollment utterances are randomly sampled from enrollment candidates \mathcal{C} to make the subset \mathcal{C}' for each mixture at each epoch. In *hard* version of worst-enrollment target training, the loss function $\mathcal{L}_{\text{worst,hard}}$ is calculated as the maximum, thus worst, value of minus SDR loss over K enrollments as follows:

$$\mathcal{L}_{\text{worst,hard}} = \max_{n \in \mathcal{C}'} \mathcal{L}_n. \quad (6)$$

The sampling of \mathcal{C}' from \mathcal{C} is introduced to reduce computational cost in model forwarding. It also prevents focusing only on a certain enrollment utterance and decreasing enrollment variation in the training.

We also propose a *soft* version of the worst-enrollment target training to prevent the model from excessively adapting to difficult cases e.g. outlier enrollments that is too noisy or dif-

Table 1: *Data generation setup.*

	dataset	mixture type	SIR [dB]	SNR [dB]	#speakers (train set / dev set)	#mixtures or utterances (train set / dev set)
(a)	training data for TSE	2 speakers and noise	-5 - 5	0 - 20	777 / 49	50,000 / 5,000
(b)	training data for ASR	1 speaker and noise	-	0 - 20	1,367 / 29	403,072 / 4,000
(c)	evaluation data	2 speakers and noise	-5 - 5	5 - 15	30	4,000




Figure 1: *Extraction performance as a function of the enrollment, from worst to best, for the baseline TSE system. The number in blue indicate the failure rate, i.e. the ratio of samples whose SDRi is below 5dB. Each box plot indicates 5th, 25th, 50th, 75th, 95th percentile values aggregated over 4,000 mixtures in evaluation dataset.*

ferent from the typical speaking style of the target speaker. The loss function $\mathcal{L}_{\text{worst,soft}}$ softly weight difficult cases over K enrollments using a softmax function with temperature as follows:

$$\mathcal{L}_{\text{worst,soft}} = \sum_{n \in \mathcal{C}'} w_n \mathcal{L}_n \quad (7)$$

$$w_n = \frac{\exp(\mathcal{L}_n / \tau)}{\sum_{n' \in \mathcal{C}'} \exp(\mathcal{L}_{n'} / \tau)} \quad (8)$$

where w_n denotes the weights and τ denotes the temperature parameter of the softmax function, respectively. The temperature parameter controls the degree of smoothing over candidates. The hard version can be seen as $\tau \rightarrow 0$ limit.

3.3. Auxiliary speaker identification loss

Auxiliary SI-loss [20] is an approach that was proposed to help speaker embedding network learn to extract more invariant latent expression to intra-speaker variance. It may also improve robustness toward enrollment variation by better handling difficult enrollment utterances e.g. those close to speaker boundaries. In SI-loss training, an auxiliary speaker identification task given a speaker embedding \mathbf{e}^s is introduced to improve speaker discriminability. The multitask loss for SI-loss training $\mathcal{L}_{\text{multitask}}$ is given by,

$$\mathcal{L}_{\text{multitask}} = \mathcal{L}_{\text{sdr}} + \alpha \text{CE}(\mathbf{1}^s, \text{Softmax}(\mathbf{W}\mathbf{e}^s)) \quad (9)$$

where $\mathbf{1}^s$ denotes a one-hot vector representing the target speaker ID, \mathbf{W} denotes an additional learnable projection matrix, and $\text{CE}(\cdot)$ and $\text{Softmax}(\cdot)$ are cross entropy and softmax function, respectively. α is a multitask weight parameter.

We also investigated the combination of the hard version of worst-enrollment target training and auxiliary SI-loss. For its training, we replaced the \mathcal{L}_{sdr} in Eq. (9) with $\mathcal{L}_{\text{worst,hard}}$ and calculate CE loss only for the worst enrollment within \mathcal{C}' in terms of the SDR loss.




Figure 2: *Worst-enrollment SDRi for conventional and proposed methods. Each box plot indicates 5th, 25th, 50th, 75th, 95th percentile values. All the training strategies improve the lower limit performance.*

4. Experiments

4.1. Experimental Setup

4.1.1. Evaluation details

Five systems were prepared for evaluation: a conventional system trained with loss function in Eq. (4), a system trained with hard worst-enrollment loss in Eq. (6), a system trained with soft worst-enrollment loss in Eq. (7), a system trained with an auxiliary SI-loss in Eq. (9), and a system trained with the combination of the worst-enrollment training and an auxiliary SI-loss. We evaluated the performance in terms of SDR improvement (SDRi), character error rate (CER), and the failure ratio. Here we define the failure ratio as the ratio of data whose SDRi was less than 5 dB in evaluation data. We computed scale-dependent SDR for evaluation using BSS Eval toolbox [22].

4.1.2. Dataset

All the training and evaluation were performed on simulated mixtures generated from speech recordings from the Corpus of Spontaneous Japanese (CSJ) [23] and noise recordings from the CHiME-4 corpus [24], sampled at 16 kHz. Table 1 shows the data generation setup for training and evaluation. Signal-to-interference ratio (SIR) and signal-to-noise ratio (SNR) in the table stands for the energy ratio between the signals from target and interfering speakers, and between the speech mixture without noise and the noise. For training and evaluation of the TSE model, we prepared $N = 10$ auxiliary utterance candidates for each mixture. Candidates were selected randomly from the target speaker’s utterances except those used as target speech and shorter than two seconds.

4.1.3. System Configuration and Training Procedure

We adopted time-domain SpeakerBeam architecture for TSE [20]. We used blocks of eight stacked 1-D convolution layers as proposed in [25] for the auxiliary and extraction networks, repeated three times for the extraction network. The other parameters of the model follow our prior work [20] and are omitted here because of space limitations. For ASR evaluation, we adopted transformer-based encoder-decoder ASR model with connectionist temporal classification objective [26] trained ac-

Table 2: Performance of each system for worst, 2nd-worst and best enrollment choices, as well as average over enrollments. Both worst-enrollment target training and auxiliary SI-loss improve the worst-enrollment performance in terms of SDRi and CER.

Metrics	SDRi [dB] (\uparrow)				CER [%] (\downarrow)		Failure ratio [%] (\downarrow) (SDRi < 5dB)		
	mean	worst	2nd worst	best	mean	worst	mean	worst	best
Conventional	15.1 \pm 3.8	13.3	14.3	15.9	13.3	19.6	2.0	7.7	0.2
Worst loss (hard)	15.2 \pm 3.6	14.0	14.8	15.8	13.1	18.1	1.4	4.8	0.2
Worst loss (soft)	15.2 \pm 3.4	14.1	14.8	15.8	13.1	18.2	1.3	4.5	0.2
SI-loss	15.3\pm3.5	14.1	14.8	15.9	13.0	17.6	1.3	4.7	0.1
Worst loss (hard) + SI-loss	15.3\pm3.5	14.2	14.8	15.8	13.0	17.5	1.3	4.4	0.2

ording to the recipe of the ESPnet open-source toolkit [27].

For optimization, we adopted the Adam optimizer [28]. The initial learning rate was set as $5e-4$ and was halved if the loss on the development set did not decrease for 3 epochs. The models were trained for 120 epochs. SI-loss was adopted from scratch, whereas worst-enrollment target training was adopted from 100 epochs. The best-performing model based on a development set was chosen for evaluation. For worst-enrollment target training, the number of selected enrollments from candidates was $K = 3$ for both hard and soft versions. The temperature parameter in soft worst-enrollment target training was set as $\tau = 2.0$. Multitask weight in auxiliary SI-loss training was $\alpha = 1.0$.

4.2. Results

4.2.1. Effect of enrollment variations on the performance

In this section, we discuss how much enrollment choices affect performance. Fig. 1 shows the performance distribution of n -th worst-enrollment SDRi obtained with the conventional system. The ratio of data whose SDRi was less than 5 dB is also presented as the failure ratio. The figure shows that extraction performance varies depending on enrollment choices even though each plot shares the same mixtures. Especially, the 5 and 25 percentile values differ greatly between n -th worst-enrollment SDRi. In other words, enrollment choices affect the lower bound performance (failure cases) rather than the upper bound performance. In addition, the failure ratio of 7.2 % for the worst-enrollment SDRi is much greater than other enrollment choices. This indicates that enrollment choices have a high possibility of determining whether TSE succeeds. Thus, the robustness toward enrollment choices is an issue to be addressed.

4.2.2. Evaluation on speech extraction performance

Here we discuss the performance of each training method. Table 2 shows SDRi, CER and failure ratio of each system with worst, 2nd-worst and best-enrollment choices as well as averaged values over enrollments. The table shows that both worst-enrollment loss and SI-loss improved performance for worst and 2nd-worst enrollment choices. Hard/soft worst-enrollment loss training, SI-loss training, and the combination of these two approaches improved worst-enrollment SDR by 0.7, 0.8, 0.8, and 0.9 dB and CER by 1.5, 1.4, 2.0, and 2.1 % points, respectively. As a deeper evaluation of worst-enrollment performance, Fig. 2 shows the performance distribution of the worst-enrollment SDRi for each system. Hard/soft worst-enrollment loss training, SI-loss training, and the combination of them improved 5 percentile values largely from 1.84 dB to 5.32, 5.66, 5.46, and 5.88 dB, respectively. From the observations above, these strategies successfully improve robustness towards enrollment variation by especially improving lower bound performances in worst enrollment cases.

Table 3: Intra-speaker inter-speaker variance ratio of speaker embeddings. Higher values indicate higher speaker discriminability learned by the embedding network.

	variance ratio
Conventional	2.96
Worst loss (soft)	3.35
SI-loss	3.49
Worst loss (hard) + SI-loss	3.61

As for the failure ratio of each system, as shown in table 2, all the training methods reduce the failure ratio, especially for the worst enrollment cases. The failure ratio was also improved on average over enrollments, by 34 % relative from conventional training method with the combination of hard worst-enrollment loss training and auxiliary SI-loss. This seems to be because the failure of extraction most often occurs in worst enrollments cases as seen in Fig. 1. From this observation, we can say that these approaches contribute to reducing the overall failure that often occurs for difficult enrollment utterances by improving the worst performance.

4.2.3. Analysis on speaker discriminability

To further understand the effect of these two types of strategies, we conducted analysis on the speaker discriminability of each system. Table 3 shows the speaker discriminability evaluated with between-class within-class variance ratio of speaker embeddings e^s [29, 30]. The training with SI-loss improves the speaker variance ratio as expected, by introducing an auxiliary speaker identification task. It should be noted that worst-enrollment loss also improves the speaker discriminability, and the combination of the worst-enrollment loss with auxiliary SI-loss shows the largest improvement. This may explain the superior performance of the combination of these approaches to the worst-enrollment performance shown in Table 2. The evaluation with a more difficult setup, where the speaker characteristics of the interfering speaker are more similar to those of the target speaker e.g. recordings from the speakers belonging to the same family will be a part of our future work.

5. Conclusion

In this work, we focused on the performance variance of TSE depending on the choice of enrollment utterances. We introduced a new performance metric, worst-enrollment SDRi, to quantitatively evaluate the robustness toward enrollment variability. To directly optimize the network for worst-enrollment SDRi, we proposed worst-enrollment target training that trains the model with placing weight on difficult enrollments. We also evaluated auxiliary speaker identification loss in terms of worst-enrollment performance, which was previously proposed for improving speaker discriminability. Experimental validation showed the effectiveness of both approaches to increase the robustness toward enrollment variability.

6. References

- [1] J. Barker, R. Marxer, E. Vincent, and S. Watanabe, "The third 'chime' speech separation and recognition challenge: Dataset, task and baselines," in *IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, 2015, pp. 504–511.
- [2] J. Barker, S. Watanabe, E. Vincent, and J. Trmal, "The fifth 'chime' speech separation and recognition challenge: dataset, task and baselines," *arXiv preprint arXiv:1803.10609*, 2018.
- [3] P. Comon and C. Jutten, *Handbook of Blind Source Separation: Independent component analysis and applications*. Academic press. Elsevier, 2010.
- [4] J. Li, L. Deng, Y. Gong, and R. Haeb-Umbach, "An overview of noise-robust automatic speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 4, pp. 745–777, 2014.
- [5] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 10, pp. 1702–1726, 2018.
- [6] Y. Luo and N. Mesgarani, "Tasnet: time-domain audio separation network for real-time, single-channel speech separation," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 696–700.
- [7] 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 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 241–245.
- [8] M. Kolbæk, D. Yu, Z.-H. Tan, and J. Jensen, "Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks," *IEEE ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 10, pp. 1901–1913, 2017.
- [9] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 31–35.
- [10] K. Žmolíková, M. Delcroix, K. Kinoshita, T. Ochiai, T. Nakatani, L. Burget, and J. Černocký, "Speakerbeam: Speaker aware neural network for target speaker extraction in speech mixtures," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 4, pp. 800–814, 2019.
- [11] M. Delcroix, K. Žmolíková, K. Kinoshita, A. Ogawa, and T. Nakatani, "Single channel target speaker extraction and recognition with speaker beam," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5554–5558.
- [12] J. Wang, 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. Interspeech*, 2018, pp. 307–311. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1205>
- [13] M. Delcroix, K. Žmolíková, T. Ochiai, K. Kinoshita, S. Araki, and T. Nakatani, "Compact network for speakerbeam target speaker extraction," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6965–6969.
- [14] Q. Wang, I. L. Moreno, M. Saglam, K. Wilson, A. Chiao, R. Liu, Y. He, W. Li, J. Pelecanos, M. Nika *et al.*, "Voicefilter-lite: Streaming targeted voice separation for on-device speech recognition," in *Proc. Interspeech*, 2020, pp. 2677–2681.
- [15] C. Xu, W. Rao, E. S. Chng, and H. Li, "Time-domain speaker extraction network," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 327–334.
- [16] 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 Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 1–11, 2018.
- [17] T. Afouras, J. S. Chung, and A. Zisserman, "The conversation: Deep audio-visual speech enhancement," in *Proc. Interspeech*, 2018, pp. 3244–3248. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1400>
- [18] H. Sato, T. Ochiai, K. Kinoshita, M. Delcroix, T. Nakatani, and S. Araki, "Multimodal attention fusion for target speaker extraction," in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 778–784.
- [19] 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 *Interspeech*, 2019, pp. 4290–4294.
- [20] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 691–695.
- [21] S. Mun, S. Choe, J. Huh, and J. S. Chung, "The sound of my voice: Speaker representation loss for target voice separation," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7289–7293.
- [22] E. Vincent, R. Gribonval, and C. Févotte, "Performance measurement in blind audio source separation," *IEEE transactions on audio, speech, and language processing*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [23] K. Maekawa, "Corpus of spontaneous japanese: Its design and evaluation," in *ISCA & IEEE Workshop on Spontaneous Speech Processing and Recognition*, 2003.
- [24] "Chime4 challenge," http://spandh.dcs.shef.ac.uk/chime_challenge/CHIME4, Cited May 28 2022.
- [25] Y. Luo and N. Mesgarani, "Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [26] S. Karita, N. E. Y. Soplin, S. Watanabe, M. Delcroix, A. Ogawa, and T. Nakatani, "Improving transformer-based end-to-end speech recognition with connectionist temporal classification and language model integration," in *Proc. Interspeech*, 2019, pp. 1408–1412.
- [27] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen *et al.*, "Espnet: End-to-end speech processing toolkit," in *Proc. Interspeech*, 2018, pp. 2207–2211.
- [28] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, 2015.
- [29] C. M. Bishop and N. M. Nasrabadi, *Pattern recognition and machine learning*. Springer, 2006, vol. 4, no. 4.
- [30] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2010.