

Target Confusion in End-to-end Speaker Extraction: Analysis and Approaches

Zifeng Zhao, Dongchao Yang, Rongzhi Gu, Haoran Zhang, Yuexian Zou*

ADSPLAB, School of ECE, Peking University, Shenzhen, China

{zhaozifeng, dongchao98}@stu.pku.edu.cn, {1701111335, haoranzhang, zouyx}@pku.edu.cn

Abstract

Recently, end-to-end speaker extraction has attracted increasing attention and shown promising results. However, its performance is often inferior to that of a blind speech separation (BSS) counterpart with a similar network architecture, due to the auxiliary speaker encoder may sometimes generate ambiguous speaker embeddings. Such ambiguous guidance information may confuse the separation network and hence lead to wrong extraction results, which deteriorates the overall performance. We refer to this as the *target confusion problem*. In this paper, we conduct an analysis of such an issue and solve it in two stages. In the training phase, we propose to integrate metric learning methods to improve the distinguishability of embeddings produced by the speaker encoder. While for inference, a novel post-filtering strategy is designed to revise wrong results. Specifically, we first identify these confusion samples by measuring the similarities between output estimates and enrollment utterances, after which the true target sources are recovered by a subtraction operation. Experiments show that performance improvement of more than 1 dB SI-SDRi can be brought, which validates the effectiveness of our methods and emphasizes the impact of the *target confusion problem*¹.

Index Terms: speech separation, end-to-end speaker extraction, target confusion problem, metric learning, post-filtering

1. Introduction

Speech separation, also referred to as the cocktail-party problem, is considered to be one of the fundamental problems in speech processing areas [1]. Although easy for human beings, the same task is still challenging for machines.

A target speaker extraction (TSE) model normally consists of two parts: a speaker encoder, which maps the enrollment utterance of the target speaker to an embedding, and a separation network, which extracts target speaker's speech from the mixture under the guidance of the injected speaker embedding. In particular, these two components are jointly trained from scratch in end-to-end speaker extraction. Many studies developed their models based on state-of-the-art separation network architectures from speech separation (e.g. TCN [2][3][4] and DPRNN [4][5]), and achieved considerable performance.

However, our preliminary experiments as well as recent research [5] show that, the performances of end-to-end speaker extraction are prone to long-tail distributions, which is depicted in Figure 1. As a result, end-to-end speaker extraction is often slightly inferior to its BSS counterpart when a similar separation network is adopted [3][5], despite the assistance of an additional speaker encoder and enrollment utterances. Such a gap



Figure 1: The target confusion problem in a two-speaker scenario. The scatter diagram on the right is the performance distribution in terms of SI-SDRi on the test set of Libri2Mix, each axis corresponds to a speaker. $s_1, s_2 > 5dB$ denotes that SI-SDRi of both s_1 and s_2 are above 5 dB, while $s_1/s_2 < 5dB$ means that either of them is below 5 dB. Spectrograms on the left are from a data sample where target confusion happens. On the top is the observed mixture. Spectrograms in the second row are the ground truths for each target speaker, and those in the last row are the output estimates of each target source.

originates from an issue which we term as the *target confusion problem* in this paper, where the speaker embedding provides an ambiguous guidance, and thus the separation network targets at a wrong speaker (i.e. the interferer). This is illustrated in the red dashed box in Figure 1. Intuitively, there are two possible causes for this phenomenon. One is the *utterance bias*, that is, the target speaker's speech (either the source or the enrollment) in the data sample deviates from its speaker cluster; The other is the *embedding bias*, which means the output of speaker encoder does not represent the guidance information accurately.

In previous studies, multi-class cross-entropy (CE loss) was proposed specially for the speaker encoder in end-toend speaker extraction, and joint training is carried out together with the reconstruction loss through a multi-task learning [2][3][6][7]. However, such a classification paradigm does not optimize similarities explicitly, which may produce suboptimal embeddings for end-to-end speaker extraction.

In this paper, we first conduct an analysis of the *target* confusion problem, emphasizing the importance of distinctive speaker embeddings for end-to-end speaker extraction. Then we explore three different metric learning methods, namely, triplet loss, prototypical loss and generalized end-to-end loss, and integrate them into the end-to-end training of a deep speaker extraction model. The key behind this is that speaker extraction is an open-set setting and we need speaker embeddings with large inter-speaker and small intra-speaker distances. Finally, to further eliminate target confusion during inference, we propose a post-filtering strategy to revise the wrong results. To be specific, we first identify confusion samples by comparing the similarities between the target source estimate and the enrollment utterance, then the true target source can be recovered by a subtraction operation. Experiments show that our methods improve the baseline by more than 1 dB in terms of SI-SDRi.

^{*} Corresponding author.

This paper is partially supported by Shenzhen Science & Technology Fundamental Research Program (No:GXWD20201231165807007-20200814115301001) and National Natural Science Foundation of China (NSFC 62176008).

¹A demo is available at https://zhazhafon.github.io/demo-confusion/



Figure 2: Similarity analysis in embedding space. $cos(e_i, s_j)$ denotes the cosine distance between the speaker embedding of enrollment utterance e_i from speaker *i* and that of the source s_j from speaker *j*. Speaker *l* is set to be the target by default in (a)(c), and speaker 2 the interferer. Dashed line in red denotes a border where two similarities are equal in (a)(c), while in (b)(d) it delimits a margin. Samples indicated by pink stars, whose SI-SDRi metrics are significant negative values, are very likely to have confused the target, which is confirmed by our listening test as well as a previous research [8]. All data are from the test set of Libri2Mix with no discard.

2. Target confusion problem

2.1. Target speaker extraction

Given an enrollment utterance e, speaker extraction is to extract target speaker's voice \hat{s} out of a speech mixture y. To make it simple, a two-speaker anechoic setup is considered in the following (i.e. $y = s_1 + s_2$). Either of the speakers in the mixture can be set as the target speaker, and the other the interferer.

2.2. Target confusion problem

As depicted in Figure 1, end-to-end speaker extraction model tends to come across with the *target confusion problem* during inference, that is, the model extracts the interfering speaker instead of the target speaker, and hence generate a wrong result. This leads to a situation that end-to-end speaker extraction even performs inferior to its BSS counterpart when a similar separation network is used, despite the assistance of an additional speaker encoder and enrollment utterances. Intuitively, *target confusion problem* may originate from two aspects:

utterance bias Considering the variability of speech, an utterance may deviate from its speaker cluster where it belongs to, and even tend to an interfering speaker cluster. Such variability comes from many uncontrollable factors of the speech like emotion, intonation, prosody and even speed. We refer to this as *utterance bias*. Note that this may occur in the source signal s as well as the enrollment utterance e.

embedding bias On one hand, the network architecture of speaker encoders for end-to-end speaker extraction are generally simpler compared with those used in speaker recognition tasks [9][10][11], which results in limited capability in speaker characteristics modeling. On the other hand, SI-SDRi is usually set as the only loss function in the end-to-end training, which does not guarantee well distinguishable speaker embeddings; These bring about the *embedding bias*, that is, the speaker embedding is not distinctive enough, such that it does not represent the target speaker accurately, or it fails to distinguish the target speaker from the interferer.

We conduct an experiment for a further analysis and comparison. In Figure 2(a)(b), similarities between speakers are measured by a pretrained ECAPA-TDNN [11], which is a stateof-the-art speaker embedder used in speaker verification with an equal error rate (EER) of less than 1%. As shown in Figure 2(a), 99.7% of the enrollment utterances are closer to their target sources instead of the interfering speech in the embedding space. And in Figure 2(b), for more than 98.9% of our test cases, the two aforementioned similarities have a margin of more than 0.1. Most interestingly, only 2.4% of target confusion samples, which is denoted with pink stars, lie beyond the border in Figure 2(a), and only 6.4% of them out of the 0.1 margin in

Figure 2(b).

Things turn out to be very different when we come to the speaker encoder in an end-to-end trained TSE model. TD-SpeakerBeam was adopted for the evaluation, in which the speaker encoder is composed of an encoder layer and a convolution block [3]. For a certain amount of samples, the enrollment utterances are much closer to the interferers instead of their target speakers, as shown in Figure 2(c). It is worth noting that more than 45.1% of target confusion samples lie above the border where two similarities are equal. And in Figure 2(d), 65% of target confusion samples lie out of the 0.1 margin.

Comparing above observations we can draw some conclusions. First, while *utterance bias* may exist in some situations where speakers' voices are very similar, it is much less significant than expected, at least in our test case; Second, a considerable amount of target confusion samples are caused by speaker embeddings that are not distinguishable enough. We argue that *embedding bias* is underestimated.

3. Methods

3.1. Metric learning for end-to-end speaker extraction

In this section, we introduce how to integrate practical metric learning methods with the end-to-end training of a speaker extraction model. The essence behind this is to generate speaker embeddings with large inter-speaker and small intra-speaker distance through explicit optimization in the metric space, so that it does not confuse the target and the interferer. Three different metric learning methods are explored in the following, including triplet loss [12], prototypical loss [13][14] and generalized end-to-end loss [15][16].

Multi-task learning A multi-task learning framework is adopted to combine the reconstruction loss and the metric learning loss:

$$L = \beta L_{ML} + \frac{1}{N} \sum_{n=1}^{N} L_n \tag{1}$$

where N denotes the batch size, L_n and L_{ML} are loss functions for waveform reconstruction and metric learning. β is a hyperparameter. The negative scale-invariant signal-to-distortion ratio is used as reconstruction loss [17].

Triplet loss (TL) A triplet (u, v, w) consists of an anchor u, a positive v and a negative w. The triplet loss forces the encoder to reserve a margin between the distance of the anchor-positive pair (u, v) and that of the anchor-negative pair (u, w):

$$l_{TL}(u, v, w) = max(0, d(u, v) - d(u, w) + \alpha)$$
(2)

where d(a, b) denotes the L2 distance between L2-normed embeddings of utterance a and b. α is the margin, which is a hyperparameter. We propose two schemes to form the triplet. In the first scheme $TL_1(s_t, e_t, e_f)$, target source s_t is set as the anchor, while the enrollment utterance of target speaker and interferer are intuitively set as the positive and negative respectively; In the second scheme $TL_2(s_t, \hat{s}_t, e_f)$, target estimate \hat{s}_t replaces the enrollment e_t as the positive. At last, the loss is averaged over the batch: $L_{TL} = \frac{1}{N} \sum_{n=1}^{N} l_{TL}(u_n, v_n, w_n)$. **Prototypical loss (PL)** In prototypical loss, utterances of speaker k are divided into a support set S_k and a query set Q_k . The prototype \mathbf{r}_k , i.e. the speaker centroid, is calculated as the mean of speaker embeddings from S_k :

$$\mathbf{r}_{k} = \frac{1}{|S_{k}|} \sum_{x_{s} \in S_{k}} Enc(x_{s})$$
(3)

where Enc() denotes the speaker encoder which maps an utterance x_s to an embedding. The likelihood that an utterance x_n in the batch belongs to its speaker z_n is calculated with a softmax over all I speakers in the training set:

$$p_{PL}(x_n, z_n) = \frac{e^{-d(Enc(x_n), \mathbf{r}_{z_n})}}{\sum_i^I e^{-d(Enc(x_n), \mathbf{r}_i)}}$$
(4)

Following the setup in TL, two different schemes are investigated: $PL_1(x_n = e_t)$ and $PL_2(x_n = \hat{s}_t)$. Speech for S_k are from the whole training set, while those for Q_k are all from the current batch. Finally, negative logarithm is applied to Eq. (4) for a maximum likelihood estimation (MLE):

$$L_{PL} = \frac{1}{|Q_k|} \sum_{x_q \in Q_k} -log(p_{PL}(x_q, z_q))$$
(5)

Generalized end-to-end loss (GL) Different from *PL*, *GL* utilizes two kinds of speaker centroids:

$$\mathbf{c}_{k}(x_{n}) = \begin{cases} \frac{1}{|C_{k}|} \sum_{x_{c} \in C_{k}} Enc(x_{c}), x_{n} \notin C_{k} \\ \frac{1}{|C_{k}|-1} \sum_{x_{c} \in C_{k}, x_{c} \neq x_{n}} Enc(x_{c}), x_{n} \in C_{k} \end{cases}$$

$$(6)$$

where C_k is an utterance bank of speaker k where speech is from the whole training set, x_n is an utterance in the batch whose similarity to be measured. Similar to *PL*, likelihood is calculated with a softmax:

$$p_{GL}(x_n, z_n) = \frac{e^{w \cdot cos(Enc(x_n), \mathbf{c}_{z_n}(x_n)) + b}}{\sum_i^I e^{w \cdot cos(Enc(x_n), \mathbf{c}_i(x_n)) + b}}$$
(7)

where w and b are learnable weights. Following the previous, we investigate two settings: $GL_1(x_n = e_t)$ and $GL_2(x_n = \hat{s}_t)$. At last, negative logarithm is applied for a MLE:

$$L_{GL} = \frac{1}{N} \sum_{n}^{N} -log(p_{GL}(x_n, z_n))$$
(8)

3.2. Post-filtering strategy

To further improve the robustness of the system, we propose a post-filtering strategy (PF) to first identify and then rectify those target confusion samples during inference. Specifically, our pipeline has three steps. The model trained with the aforementioned methods first consumes the speech mixture y_m and an enrollment utterance e_m to generate a target source estimate \hat{s}'_m ; Secondly, \hat{s}'_m is evaluated in two dimensions: one is its similarity with the target speaker, denoted as π , and the other is that with the interferer, denoted as ϕ . Considering that groundtruth sources are not available during inference, speaker clusters



Figure 3: Configuring the post-filtering strategy on validation set. (a) Distribution of data samples in a space spanned by $\pi = L2(\hat{s}_1, e_1)$ and $\phi = L2(\hat{s}_1, e_2)$, where L2(a, b) denotes the L2 distance between two L2-normed embeddings of utterance a and b. Rectangular and linear decision borders are depicted by black and blue dashed lines respectively. (b) Target confusion samples identified by our methods are illustrated with purple stars. All data are from the dev set of Libri2Mix with no discard.

are estimated by their enrollment speech, as illustrated in Figure 3(a). Then a decision border to classify target confusion samples can be easily obtained by solving an optimization problem in a *M*-sample discrete space spanned by π and ϕ . For this we propose two objective functions. One is a rectangular border PF^{rec} :

$$\max_{\Pi,\Phi} \quad \sum_{m=1}^{M} g_m^{rec}(\Pi,\Phi) \tag{9}$$

$$g_m^{rec}(\Pi, \Phi) = \begin{cases} l(s_m, y - f(y_m | e_m)), \pi > \Pi, \phi < \Phi\\ l(s_m, f(y_m | e_m)), otherwise \end{cases}$$
(10)

where Π and Φ are threshold parameters for π and ϕ respectively. l is a SI-SDRi metric [18], and f is the TSE model which produces a target source estimate normalized w.r.t. the input mixture y_m , given an enrollment utterance e_m . The other is a linear border PF^{lin} , such that:

r

$$\max_{\mu,\lambda} \quad \sum_{m=1}^{M} g_m^{lin}(\mu,\lambda) \tag{11}$$

$$g_m^{lin}(\mu,\lambda) = \begin{cases} l(s_m, y - f(y_m|e_m)), \phi < \mu\pi + \lambda\\ l(s_m, f(y_m|e_m)), otherwise \end{cases}$$
(12)

where μ and λ are parameters to be tuned. There are many training-free ways to solve Eq. (9)-(12), e.g. brute-force searching. Note that the test set is assumed to be inaccessible for the tuning, and parameters are configured only using the validation set where ground-truth sources are available. Lastly, target confusion samples can be identified during inference, by thresholding on π and ϕ with tuned parameters Π and Φ (or with μ and λ), as depicted in Figure 3(b). Under the two-speaker anechoic setup introduced in Section 2.1, those identified samples are inverted by being subtracted from the mixture y, after which the true target source is recovered:

$$\hat{s}_m = y_m - \hat{s}'_m \tag{13}$$

where \hat{s}_m is the final output after the post-filtering. More complex scenarios like multiple speakers ($\# spk \geq 3$) or noisy environment will be explored in the future.

4. Experiments

4.1. Preparation

TD-SpeakerBeam [3] is adopted for our experiments. It is chosen such that we can fairly compare it with a blind speech sep-



Figure 4: Joint distributions of two speakers' SI-SDRi performance on the test set of Libri2Mix. In each mixture audio, two speakers are set to be the target in turns. (a)(b)(c) are the joint distributions, and (d)(e)(f) are their corresponding statistical histograms.

aration counterpart Conv-TasNet [17]. These two models have similar network architectures, except for that TD-SpeakerBeam has an additional speaker encoder and an embedded adaptation layer. We validate our methods on the popular LibriMix [19] dataset. The *train-100* subset is used for training, *dev* subset for configuring the post-filtering as well as for the validation set during training, while *test* is used for the final evaluation. All speech audios are in 8 kHz. During training, both input mixtures and enrollment speech are randomly truncated to 3 seconds, while full-length audios are used for testing.

4.2. Results

We compare the proposed training methods with three baselines on the *sep_clean* task of Libri2Mix: (1) NS: negative SI-SDR as the only training target; (2) CE: multi-task learning with a mulit-class cross-entropy loss for speaker classification and a negative SI-SDR loss for waveform approximation; (3) BSS: a blind speech separation model (Conv-TasNet) trained with permutation-invariant training [20][21].

	SI-SDRi(dB)	PESQ	params
NS	12.86	2.75	-
CE	13.05	2.78	β=0.2
BSS	13.40	2.74	-
TL_1	13.31	2.82	$\beta = 0.2, \alpha = 1$
TL_2	13.36	2.83	β =0.2, α =1
PL_1	13.46	2.85	$\beta = 0.2, S_k = 5$
PL_2	13.46	2.85	$\beta = 0.1, S_k = 5$
GL_1	13.47	2.85	β=0.1
GL_2	13.44	2.83	$\beta=0.1$
$NS + PF^{rec}$	13.13	2.76	П=0.4, Ф=0.4
$NS + PF^{lin}$	13.14	2.76	$\mu = 0.4, \lambda = 0.2$
$CE + PF^{rec}$	13.31	2.79	Π=0.5, Φ=0.5
$CE + PF^{lin}$	13.32	2.79	$\mu = 0.4, \lambda = 0.2$
$PL_2 + PF^{rec}$	13.82	2.85	Π=0.8, Φ=1.0
$PL_2 + PF^{lin}$	13.88	2.86	μ =0.6, λ =0.3

Table 1: Comparing the overall performance.

Results are presented in terms of SI-SDRi [18] and PESQ [22] in Table 1. For the sake of space, only the best results are reported, together with their hyperparameters. As illustrated in the first and the third row, the TSE model is inferior to its BSS counterpart by 0.54 dB in terms of SI-SDRi, which is consistent with our statements in Section 1. By observing row four to row nine, we can see that all metric learning methods promote the performance and outperform the CE baseline. The performance difference between scheme 1 and scheme 2 is minor. Among proposed training methods, TL performs the worst

(13.36 dB SI-SDRi). PL and GL achieve similar results, improving the performance by 0.6 dB and 0.61 dB SI-SDRi respectively, and both of them outperform the BSS baseline.

For the post-filtering strategy, threshold parameters are tuned on dev set in advance and set to be constant during inference. Note that the threshold parameters should not be too precise (e.g. one decimal place would be fair enough) to avoid overfitting on the validation set. As shown in the last six rows in Table 1, both PF^{rec} and PF^{lin} further improve the performance of our methods as well as baselines. An example² is depicted in Figure 5. Interestingly, applying PF on the basis of proposed training methods brings more gain in SI-SDRi than simply applying it to the baselines, and it further advances our results by 0.36 dB and 0.42 dB. This is due to that proposed training methods form a more distinctive decision border in the subspace spanned by π and ϕ , which is vital for the post-filtering.

We visualize some of the results in Figure 4. The proposed methods significantly alleviate the long-tail distribution in end-to-end speaker extraction. The best performance is achieved by $PL_2 + PF^{lin}$, with a SI-SDRi of 13.88 dB and a PESQ of 2.86.



Figure 5: Spectrograms from a data sample. In the first column are the original estimates from the deep model, in the second column are final outputs after post-filtering, and the ground truths are in the last column. As shown in the blue dashed box, target confusion is rectified by proposed methods.

5. Conclusions

In this paper, we conduct an analysis of what we refer to as the *target confusion problem* in end-to-end speaker extraction, and proposed to solve it with metric learning methods and a post-filtering strategy. Experiments show that our methods promote the performance by more than 1 dB SI-SDRi. Our methods are compatible with any off-the-shelf TSE models since they add no modification to the network architecture. In future work, we plan to extend our methods to more complicated scenarios, for example, multi-talker ($\#spk \ge 3$) and noisy extraction.

 $^{^2 \}rm More$ audio examples are available at our demo webpage: https://zhazhafon.github.io/demo-confusion/

6. References

- 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.
- [2] C. Xu, W. Rao, E. S. Chng, and H. Li, "Spex: Multi-scale time domain speaker extraction network," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1370–1384, 2020.
- [3] 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)*. IEEE, 2020, pp. 691–695.
- [4] Y. Hao, J. Xu, J. Shi, P. Zhang, L. Qin, and B. Xu, "A unified framework for low-latency speaker extraction in cocktail party environments." in *Proc. Interspeech*, 2020, pp. 1431–1435.
- [5] M. Elminshawi, W. Mack, S. Chakrabarty, and E. A. Habets, "New insights on target speaker extraction," *arXiv preprint arXiv:2202.00733*, 2022.
- [6] M. Ge, C. Xu, L. Wang, E. S. Chng, J. Dang, and H. Li, "Spex+: A complete time domain speaker extraction network," *Proc. Interspeech*, pp. 1406–1410, 2020.
- [7] —, "Multi-stage speaker extraction with utterance and framelevel reference signals," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6109–6113.
- [8] Z. Zhang, B. He, and Z. Zhang, "X-tasnet: Robust and accurate time-domain speaker extraction network," *Proc. Interspeech*, pp. 1421–1425, 2020.
- [9] W. Cai, J. Chen, and M. Li, "Exploring the encoding layer and loss function in end-to-end speaker and language recognition system," in *Proc. Odyssey 2018 The Speaker and Language Recognition Workshop*, 2018, pp. 74–81.
- [10] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," *Proc. Interspeech*, pp. 1086–1090, 2018.
- [11] B. Desplanques, J. Thienpondt, and K. Demuynck, "Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification," in *Proc. Interspeech*, 2020, pp. 3830–3834.
- [12] C. Zhang, K. Koishida, and J. H. L. Hansen, "Text-independent speaker verification based on triplet convolutional neural network embeddings," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 9, pp. 1633–1644, 2018.
- [13] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," *Advances in neural information processing* systems, vol. 30, 2017.
- [14] J. Wang, K.-C. Wang, M. T. Law, F. Rudzicz, and M. Brudno, "Centroid-based deep metric learning for speaker recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP).* IEEE, 2019, pp. 3652–3656.
- [15] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, "Generalized endto-end loss for speaker verification," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4879–4883.
- [16] J. S. Chung, J. Huh, S. Mun, M. Lee, H.-S. Heo, S. Choe, C. Ham, S. Jung, B.-J. Lee, and I. Han, "In defence of metric learning for speaker recognition," *Proc. Interspeech*, pp. 2977–2981, 2020.
- [17] Y. Luo and N. Mesgarani, "Conv-tasnet: Surpassing ideal timefrequency magnitude masking for speech separation," *IEEE/ACM transactions on Audio, Speech, and Language Processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [18] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "Sdrhalf-baked or well done?" in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 626–630.

- [19] J. Cosentino, M. Pariente, S. Cornell, A. Deleforge, and E. Vincent, "Librimix: An open-source dataset for generalizable speech separation," arXiv preprint arXiv:2005.11262, 2020.
- [20] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP).* IEEE, 2017, pp. 241–245.
- [21] 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.
- [22] A. Rix, J. Beerends, M. Hollier, and A. Hekstra, "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," in *IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (ICASSP)*, vol. 2, 2001, pp. 749–752 vol.2.