

Self-Supervised Disentangled Representation Learning for Robust Target Speech Extraction

Zhaoxi Mu¹, Xinyu Yang¹, Sining Sun², Qing Yang²

¹Xi'an Jiaotong University

²Du Xiaoman

wsmzxxh@stu.xjtu.edu.cn, xyphd@mail.xjtu.edu.cn, {sunsining,yangqing}@duxiaoman.com

Abstract

Speech signals are inherently complex as they encompass both global acoustic characteristics and local semantic information. However, in the task of target speech extraction, certain elements of global and local semantic information in the reference speech, which are irrelevant to speaker identity, can lead to speaker confusion within the speech extraction network. To overcome this challenge, we propose a self-supervised disentangled representation learning method. Our approach tackles this issue through a two-phase process, utilizing a reference speech encoding network and a global information disentanglement network to gradually disentangle the speaker identity information from other irrelevant factors. We exclusively employ the disentangled speaker identity information to guide the speech extraction network. Moreover, we introduce the adaptive modulation Transformer to ensure that the acoustic representation of the mixed signal remains undisturbed by the speaker embeddings. This component incorporates speaker embeddings as conditional information, facilitating natural and efficient guidance for the speech extraction network. Experimental results substantiate the effectiveness of our meticulously crafted approach, showcasing a substantial reduction in the likelihood of speaker confusion.

Introduction

The human auditory system excels in extracting the speech of a target speaker from a complex acoustic environment. Consequently, a longstanding objective of speech-processing research has been to develop machines capable of emulating similar auditory abilities. Target speech extraction (TSE) draws inspiration from human top-down selective auditory attention (Mesgarani and Chang 2012; Kaya and Elhilali 2017), which employs cues to selectively attend to specific auditory stimuli based on relevance. Discriminative cues utilized in TSE include spatial cues indicating the target speaker's direction (Gu et al. 2019), video recordings of the target speaker's mouth movements (Afouras et al. 2020; Gao and Grauman 2021), and pre-recorded reference speech (Wang et al. 2019; Xu et al. 2020). Reference speech holds particular value as it provides essential information about the target speaker's voice characteristics and is easily accessible. Accordingly, this paper aims to enhance the performance of monaural TSE methods driven by reference speech.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Schematic depicting information disentanglement of reference speech.

TSE methods typically comprise two main components: a reference speech encoding network (RSEN) and a speech extraction network (SEN) (Zhao, Gao, and Shinozaki 2020; Xu et al. 2020). The RSEN extracts speaker embeddings from the target speaker's reference speech, while the SEN predicts the target speaker's speech within the mixed speech guided by these embeddings. However, the performance of TSE methods often exhibits a long-tail distribution (Zhao et al. 2022), indicating that the extracted speech may suffer from the speaker confusion (SC) problem (Elminshawi et al. 2022), also known as target confusion (Zhao et al. 2022). The SC problem arises when the RSEN extracts ambiguous speaker embeddings that provide misleading guidance to the SEN. This confusion can cause the SEN to focus on the wrong speaker, leading to inaccurate extraction outcomes.

The SC problem primarily stems from embedding bias (Zhao et al. 2022). This bias arises when speaker embeddings extracted from the RSEN fail to accurately represent speaker cues. While these embeddings typically contain discriminative acoustic information associated with the target speaker, they can also be entangled with irrelevant interference information. As speech signals inherently carry both static global acoustic information and dynamic local semantic information, we argue that the reference speech can be disentangled into three distinct components: local information, helpful global information, and harmful global information, as illustrated in Figure 1. Local information pertains to semantic information, while helpful global information refers to speaker identity information. Harmful global information encompasses paralinguistic variables such as emotion, speaking rate, prosody and intonation, which can vary even for the same speaker (Deng et al. 2021; Zhao et al.

2022). If semantic information leaks into the embeddings, the SEN may prioritize the source that aligns with the reference speech’s semantic information. Likewise, if harmful global information leaks into the embeddings, the SEN may pay more attention to the source that aligns with the reference speech’s emotion, speaking rate, prosody, and intonation rather than the source that aligns with the speaker identity of the reference speech. Both of these scenarios can lead to the SC problem.

Previous studies often employ pre-trained speaker recognition networks to extract reference embeddings (Zhang, He, and Zhang 2020; Liu et al. 2023b). However, this approach may yield suboptimal embeddings for TSE due to pattern mismatch. Alternatively, the RSEN can be trained using a multi-class cross-entropy loss with speaker identity labels and jointly optimized with the SEN (Zhao, Gao, and Shinzaki 2020; Xu et al. 2020). However, a limitation of this method is its reliance on speaker identity labels in the training data, which may not be available in real-world scenarios.

To address the challenges above, we propose a novel two-phase self-supervised **disentangled representation learning** (DRL) method for robust **target speech extraction**, called SDR-TSE. Our approach involves explicitly disentangling the reference speech’s semantic and global information using the RSEN, followed by the implicit disentanglement of the speaker identity information within the global information using the global information disentanglement network (GIDN). By only utilizing the disentangled information for guidance, our TSE pipeline avoids any leakage of harmful information from the reference speech. Moreover, the RSEN and GIDN are trained in a self-supervised manner, negating the dependence on speaker identity labels and enhancing the applicability across diverse real-world scenarios.

Previous methods for integrating speaker embeddings and acoustic representations in TSE typically relied on simplistic summation and concatenation methods (Ge et al. 2020; Deng et al. 2021). However, these methods are susceptible to information overload, as speaker embeddings can overwhelm the acoustic representation information. To overcome this limitation, we introduce a natural fusion method by replacing the layer normalization in the Transformer with adaptive modulation layer normalization (AMLN). AMLN integrates speaker embeddings as conditional information to enhance the SEN’s perception capability for the target speaker without interfering with the acoustic representation.

In summary, this paper makes several notable contributions: (i) We propose a novel two-phase self-supervised DRL policy to effectively tackle the issue of speaker confusion in TSE. (ii) We propose an approach for incorporating speaker identity information into the SEN naturally and efficiently. (iii) We conduct comprehensive experiments to validate the significance of information disentanglement, and our method defines new state-of-the-art performance.

Related Work

Speech Separation. Speech separation (SS) refers to the process of isolating individual speech components from mixed speech signals (Wang and Chen 2018). Early advancements in SS primarily focused on techniques in the

time-frequency domain (Hershey et al. 2016; Yu et al. 2017; Kolbaek et al. 2017). To circumvent the explicit phase estimation problem, Luo and Mesgarani (2019) proposed Conv-TasNet, a time-domain SS model that employs a CNN to extract speech features. Additionally, to handle the separation of long speech sequences while reducing computational complexity, Luo, Chen, and Yoshioka (2020) proposed the DPRNN, consisting of three components: segmentation, chunk processing, and overlap-add. Recently, Subakan et al. (2021) proposed the Sepformer, which replaces the RNN in DPRNN with a Transformer architecture.

Target Speech Extraction. TSE tackles the challenges of unknown speaker numbers and speaker permutations, which are not fully resolved in SS methods, by leveraging the target speaker’s reference speech. Previous research has investigated the SC problem in TSE. For instance, Zhao et al. (2022) proposed a two-stage solution. In the training stage, they integrated metric learning methods to enhance the discriminability of the embeddings extracted by the RSEN. In the inference stage, they employed a post-filtering strategy to rectify erroneous results. However, this method is limited to scenarios involving two speakers in mixed speech. More recently, Liu et al. (2023b) introduced two novel loss functions to optimize performance metrics by defining the reconstruction quality at the chunk level. These loss functions make use of metric-correlated distribution information, enabling the SEN to focus on the chunks where SC occurs.

Disentangled Representation Learning. DRL aims to separate unique, independent, and informative factors present in the data. Disentangled latent variables exhibit sensitivity to changes in a single underlying factor while insensitive to other factors, thus ensuring statistical independence (Bengio, Courville, and Vincent 2013; Wang et al. 2022). DRL has found widespread applications in speech synthesis, conversion, and enhancement to enhance model interpretability and controllability. For instance, Choi et al. (2021) and Qian et al. (2022) utilized information perturbation as a data enhancement technique to learn speaker-independent feature representations and disentangle the speaker information from other information in speech. However, these approaches rely on pre-trained speech feature extractors. Hou et al. (2021) addressed the issue of noise type mismatch in speech enhancement by employing a noise type classifier with a gradient reversal layer (GRL) as the disentangler to learn noise-agnostic feature representations. Similarly, Nekvinda and Dusek (2020) and Liu et al. (2023a) also employed GRL to disentangle speaker identity information from speech signals. However, these methods rely on noise type or speaker identity labels for training, while our proposed DRL method is self-supervised.

Methodology

Notations and Problem Formulation

The mixed speech signal $y \in \mathbb{R}^T$ of length T , comprising multiple speakers and noise interference, can be expressed as the sum of the target speech $u \in \mathbb{R}^T$ and other interfering components $v \in \mathbb{R}^T$,

$$y = u + v \quad (1)$$

The objective is to separate the target speech u from other interfering components in y using the target speaker’s cues,

$$\hat{u} = \mathcal{F}(y, z_s; \theta_{\mathcal{F}}) \quad (2)$$

where \hat{u} represents the target speech predicted by the SEN \mathcal{F} , and $\theta_{\mathcal{F}}$ denotes the parameters of \mathcal{F} . The target speaker embedding z_s is utilized to guide \mathcal{F} and is extracted by encoding a reference speech x from the same speaker as u .

Model Overview

Figure 2 illustrates the architecture of our proposed model, which consists of the SEN, RSEN and GIDN. The SEN is built upon the state-of-the-art (SOTA) SS model Sepformer (Subakan et al. 2021), which follows an encoder-masker-decoder framework, as depicted in Figure 2(a). The waveform encoder of the SEN consists of a 1-D convolution that encodes y into a time-domain feature representation. This feature representation is subsequently segmented into overlapping chunks, concatenated and fed into stacked Transformer blocks to alternate between local and global modelling. The resulting output is then transformed back into sequential permutation using the overlap-add method, enabling the prediction of the target speaker mask. The mask is multiplied with the feature representation of y encoded by the waveform encoder to derive the feature representation of the target speaker. Ultimately, this feature representation is fed into the waveform decoder, composed of a 1-D transposed convolution, to generate the target speaker’s speech.

The RSEN is a vital component of the SDR-TSE, comprising three main components: a global information encoder E_g , a semantic information encoder E_c , and a spectrogram decoder D , as displayed in Figure 2(c). E_g and E_c encode the reference speech into global and semantic spaces, respectively. D reconstructs the spectrogram of the reference speech using the global and semantic information representations. The global information representation produced by E_g is then fed into the GIDN to extract the speaker embedding, as depicted in Figure 2(b). The speaker embedding serves as a guidance signal in the Intra- and Inter-AM-Transformer blocks of the SEN. These components will be described in more detail in the subsequent sections.

Reference Speech Encoding Network

Let $x \in \mathbb{R}^{T_x}$ represent the reference speech, and $X \in \mathbb{R}^{F_x \times T_x}$ denote the spectrogram of x generated via Short-Time Fourier Transform (STFT). The global and semantic latent representations of X are denoted as $z_g \in \mathbb{R}^{d_g \times T_g}$ and $z_c \in \mathbb{R}^{d_c \times T_c}$, respectively. To disentangle these two latent representations, we assume they are probabilistically independent. We define the joint latent representation as $z = [z_g, z_c]$ and factorize the prior distribution $p(z)$ and posterior distribution $p(z | X)$ of z_g and z_c by following the independence assumption:

$$p(z) = p(z_g)p(z_c) \quad (3)$$

$$p(z | X) = p(z_g | X)p(z_c | X) \quad (4)$$

E_g and E_c estimate the posterior distributions $p(z_g | X)$ and $p(z_c | X)$ as $q(z_g | X)$ and $q(z_c | X)$. The variational

autoencoder (VAE), known for its capability to disentangle information and model the semantic information of speech (Wang et al. 2022), is employed to construct the RSEN. To learn meaningful semantic information representation, we assume that the prior $p(z_c)$ of semantic information follows a standard normal distribution.

E_g is responsible for encoding X into a global feature space, producing the global information representation z_g . This process can be expressed as:

$$q(z_g | X) = E_g(X; \theta_{E_g}) \quad (5)$$

θ_{E_g} represents the parameters of E_g . E_g aims to capture all global features of the reference speech. To enable E_c to encode X into a semantic space and acquire a meaningful semantic representation z_c , we assume $q(z_c | X)$ follows a conditionally independent Gaussian distribution with unit variance to reduce complexity, motivated by Liu, Breuel, and Kautz (2017) and Chou and Lee (2019). Formally, we express this as follows:

$$q(z_c | X) = \mathcal{N}(E_c(X; \theta_{E_c}), I) \quad (6)$$

where θ_{E_c} represents the parameters of E_c . The role of D is to reconstruct the spectrogram \hat{X} of the reference speech x using the encoded representations z_g and z_c . The reconstruction process can be formulated as follows:

$$\hat{X} = D(z_g, z_c; \theta_D) \quad (7)$$

θ_D represents the parameters of D . The reconstruction step allows E_g and E_c to effectively encode global and semantic information of the reference speech. To achieve this, we apply average pooling on the temporal dimension of z_g and leverage the technique of adaptive instance normalization (AdaIN) (Chou and Lee 2019) to conditionally reconstruct \hat{X} using the pooled embedding vector.

The RSEN is optimized by maximizing the Evidence Lower Bound (ELBO):

$$\max_{\theta_{E_c}, \theta_{E_g}, \theta_D} \text{ELBO} = \mathbb{E}_{q(z_c|X)} \mathbb{E}_{q(z_g|X)} (\log p(X | z_c, z_g)) - D_{KL}(q(z_c | X) \| p(z_c)) \quad (8)$$

D_{KL} represents the Kullback-Leibler (KL) divergence. Equivalently, Eq.(8) can be optimized by minimizing the reconstruction loss \mathcal{L}_{REC} and the KL divergence loss \mathcal{L}_{KL} :

$$\mathcal{L}_{\text{REC}} = \|\hat{X} - X\|_1 \quad (9)$$

$$\mathcal{L}_{\text{KL}} = \|z_c\|_2^2 \quad (10)$$

$\|\cdot\|_1$ and $\|\cdot\|_2$ represent the L^1 norm and L^2 norm. To alleviate computational complexity, we calculate \mathcal{L}_{REC} in the time-frequency domain of the speech signal. \mathcal{L}_{KL} encourages the posterior $q(z_c | X)$ to align with the prior $p(z_c) = \mathcal{N}(z_c | 0, I)$, where I denotes the identity matrix. Notably, the training process is self-supervised and does not rely on speaker identity labels. E_c and D are exclusively utilized during training and discarded during inference.

Disentangled Representation Learning. In the RSEN, we employ various techniques to disentangle the latent variables z_g and z_c . Firstly, we introduce instance normalization (IN) in E_c as an information bottleneck to filter out

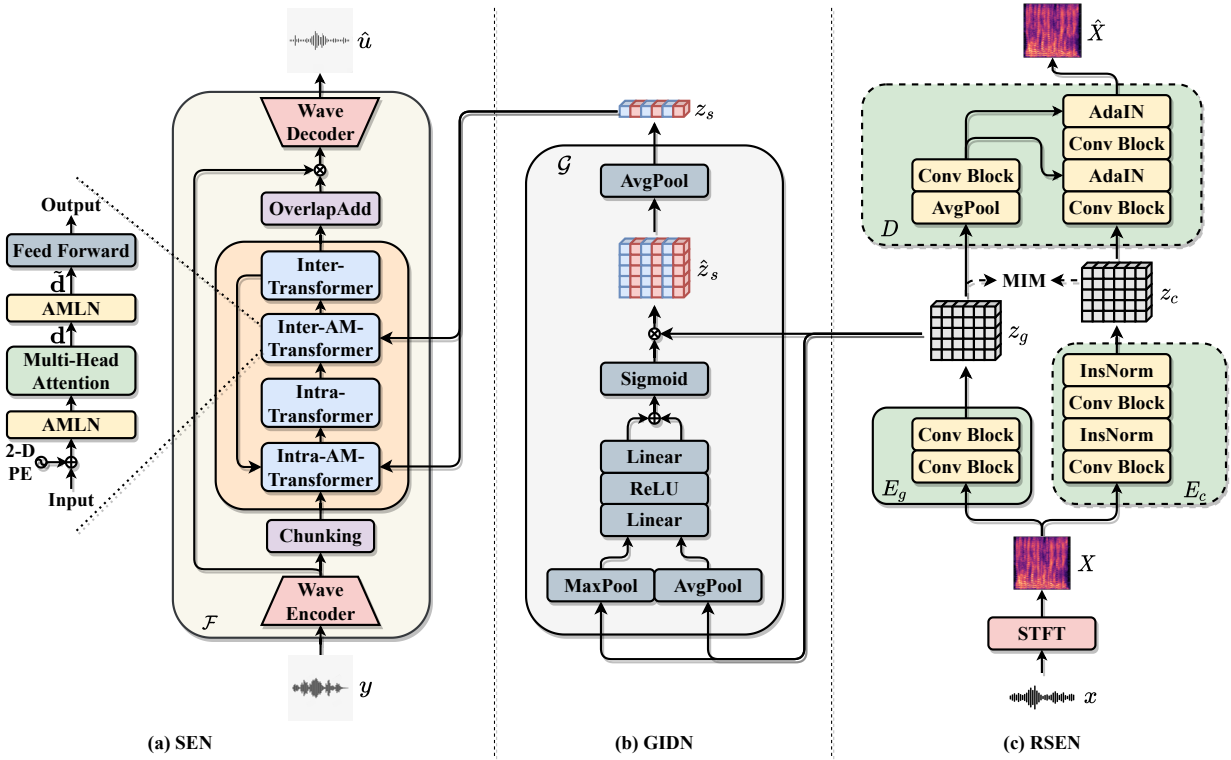


Figure 2: The architecture of the SDR-TSE. (a), (b) and (c) depict the speech extraction network, global information disentanglement network and reference speech encoding network. The semantic information encoder E_c and spectrogram decoder D within the dashed box are utilized solely for training purposes to facilitate disentanglement and discarded during inference. MIM refers to mutual information minimization. The red channels in the feature map of the GIDN indicate activated channels containing speaker identity information, while blue channels represent suppressed channels containing harmful information.

global information while preserving the semantic information, motivated by Chou and Lee (2019) and Chen et al. (2021). Furthermore, during training, we minimize the mutual information (MI) between z_g and z_c to prevent any mutual leakage between the semantic and global information. Incorporating MI as a regularization term can enhance the disentanglement capability of the VAE and constrain the dependencies between z_g and z_c . Specifically, in the previous section, we assumed that the random variables z_c and z_g are mutually independent. To achieve this, we minimize the KL divergence between their joint distribution and the product of their marginal distributions, which can be expressed as:

$$\min_{\theta_{E_c}, \theta_{E_g}} \mathcal{I}(z_c, z_g) = \mathbb{E}_{p(z_c, z_g)} \left[\log \frac{p(z_c, z_g)}{p(z_c)p(z_g)} \right] \quad (11)$$

$\mathcal{I}(z_c, z_g)$ represents the MI between z_c and z_g . Eq.(11) can also function as a regulariser of Eq.(8). We employ the vCLUB method (Cheng et al. 2020) to estimate an upper bound on the MI, which is given by:

$$\mathcal{I}_{\text{vCLUB}}(z_c, z_g) = \mathbb{E}_{p(z_c, z_g)} [\log q(z_c | z_g)] - \mathbb{E}_{p(z_c)} \mathbb{E}_{p(z_g)} [\log q(z_c | z_g)] \quad (12)$$

$q(z_c | z_g)$ serves as a variational approximation to the true posterior distribution $p(z_c | z_g)$, and is parameterized by

a neural network \mathcal{V} . An unbiased estimator of the vCLUB between z_c and z_g can be given by:

$$\hat{\mathcal{I}}_{\text{vCLUB}}(z_c, z_g) = \frac{1}{N} \sum_{i=1}^N \left[\log q(z_c^{(i)} | z_g^{(i)}) - \frac{1}{N} \sum_{j=1}^N \log q(z_c^{(j)} | z_g^{(i)}) \right] \quad (13)$$

N denotes the number of samples. $\hat{\mathcal{I}}_{\text{vCLUB}}(z_c, z_g)$ provides a reliable upper bound on MI with a well-performing variational approximation. To improve the accuracy of the vCLUB MI estimator, we train the variational approximation network \mathcal{V} to maximize the log-likelihood:

$$\mathcal{L}_{\text{LL}} = \frac{1}{N} \sum_{i=1}^N \log q(z_c^{(i)} | z_g^{(i)}) \quad (14)$$

\mathcal{V} and RSEN are alternately optimized during training.

Global Information Disentanglement Network

We consider that each channel of the output z_g from E_g encompasses distinct types of information (Mu, Yang, and Zhu 2023; Mu et al. 2023). Some channels may contain valuable speaker identity information, while others may contain harmful global information. To disentangle the useful

speaker identity information, the GIDN \mathcal{G} is utilized to refine z_g further. The GIDN extracts the speaker embedding $z_s \in \mathbb{R}^{d_s}$, exclusively containing speaker identity information for guiding the SEN. Specifically, we employ channel attention (Woo et al. 2018) to adjust the importance of individual channels within z_g to focus on channels that contribute to the TSE while suppressing irrelevant channels. The weight assigned to each channel of z_g can be calculated as:

$$\phi(z_g) = \sigma(W_1 f(W_0(z_g^{\text{avg}})) + W_1 f(W_0(z_g^{\text{max}}))) \quad (15)$$

z_g^{avg} and z_g^{max} are derived from z_g by performing average and max pooling operations. Afterwards, they are fed into two linear layers that share weights W_0 and W_1 . The activation function f used is ReLU, and σ denotes the sigmoid function. The resulting output \hat{z}_s of the channel attention is:

$$\hat{z}_s = \phi(z_g) \otimes z_g \quad (16)$$

\otimes denotes element-wise multiplication. We apply average pooling along the temporal dimension of \hat{z}_s to obtain the time-independent speaker embedding z_s to guide the SEN.

To ensure that \mathcal{G} selectively activates the channels in z_g that contain speaker identity information while suppressing those with harmful global information, we incorporate the concept of contrastive learning for training \mathcal{G} . The objective is to guarantee that the reference speech x belongs to the same speaker as \hat{u} extracted by the SEN, while the interference signal \hat{v} , computed as $y - \hat{u}$, belongs to other speakers. Thus, the speaker embedding $z_s^{(x)}$ of the reference speech x , extracted by \mathcal{G} , should exhibit similarities to the embedding $z_s^{(\hat{u})}$ of \hat{u} while being distinct from the embedding $z_s^{(\hat{v})}$ of \hat{v} . To enforce this constraint, we employ a similarity discriminative loss \mathcal{L}_{SIM} to train \mathcal{G} , defined as follows:

$$\mathcal{L}_{\text{SIM}} = \langle z_s^{(x)}, z_s^{(\hat{u})} \rangle - \langle z_s^{(x)}, z_s^{(\hat{v})} \rangle \quad (17)$$

$\langle \cdot, \cdot \rangle$ is the dot product operator, and cosine similarity is employed to quantify the similarity between two embeddings. Notably, when back-propagating \mathcal{L}_{SIM} , the parameters of the SEN are held constant, while the parameters of the GIDN and E_g are optimized.

Speech Extraction Network

To effectively utilize z_s to guide the SEN \mathcal{F} , we propose the Adaptive Modulation Transformer (AM-Transformer) as a replacement for the Transformer module in Sepformer, motivated by Min et al. (2021) and Wu et al. (2022). The AM-Transformer is capable of naturally incorporating speaker identity information as a condition. As depicted in Figure 2(a), we substitute the layer normalization (Ba, Kiros, and Hinton 2016) in Transformer with adaptive modulation layer normalization (AMLN). In contrast to the fixed gain and bias in layer normalization, we leverage $z_s \in \mathbb{R}^{d_s}$ as a condition to predict the gain and bias of the input acoustic representation. Specifically, given the input acoustic representation $\mathbf{d} \in \mathbb{R}^{H \times T_d}$ for the AMLN, we calculate its mean $\mu \in \mathbb{R}^{T_d}$ and standard deviation $\sigma \in \mathbb{R}^{T_d}$. The normalized vector $\mathbf{h} \in \mathbb{R}^{H \times T_d}$ of \mathbf{d} is defined as:

$$\mathbf{h} = \frac{\mathbf{d} - \mu}{\sigma} \quad (18)$$

Algorithm 1: SDR-TSE Optimization

Require: The training data D^* containing mixed-target-reference speech triplets (y, u, x) .

- 1: Initialize the entire system randomly.
 - 2: **while** not converged **do**
 - 3: Sample $\{(y_i, u_i, x_i)\}_{i=1}^N$ from D^* .
 - 4: **Forward-Propagation**
 - 5: Reconstruct the spectrogram $\{\hat{X}_i\}_{i=1}^N$ of $\{x_i\}_{i=1}^N$ and predict the target speech $\{\hat{u}_i\}_{i=1}^N$.
 - 6: **Back-Propagation**
 - 7: Update $\theta_{\mathcal{V}}$ by maximizing \mathcal{L}_{LL} .
 - 8: Update θ_{E_g} , θ_{E_c} , θ_D , $\theta_{\mathcal{G}}$ and $\theta_{\mathcal{F}}$ by minimizing \mathcal{L}_{KL} , \mathcal{L}_{REC} , $\mathcal{I}_{\text{VCLUB}}$ and $\mathcal{L}_{\text{SI-SNR}}$.
 - 9: Update $\theta_{\mathcal{G}}$ and θ_{E_g} by minimizing \mathcal{L}_{SIM} .
 - 10: **end while**
-

The output of AMLN, denoted as $\tilde{\mathbf{d}} \in \mathbb{R}^{H \times T_d}$, is given by:

$$\tilde{\mathbf{d}} = \gamma(z_s) \cdot \mathbf{h} + \beta(z_s) \quad (19)$$

$\gamma(z_s) \in \mathbb{R}^H$ and $\beta(z_s) \in \mathbb{R}^H$ represent two affine transformations of z_s , implemented by two fully connected layers. They adaptively scale and shift \mathbf{h} based on the condition z_s . Additionally, we employ 2-D position encoding (PE) (Raisi et al. 2020; Lin et al. 2023) instead of the original PE. This modification enables more effective utilization of intra- and inter-chunk positional information. The entire model is optimized using the SI-SNR loss $\mathcal{L}_{\text{SI-SNR}}$ (Roux et al. 2019). The training process is shown in Algorithm 1.

Experiments

Datasets and Implementation Details

The TSE model is trained and evaluated using the widely-used two-speaker mixed dataset WSJ0-2mix (Hershey et al. 2016) and its derivative dataset WSJ0-2mix-extr (Xu et al. 2020). WSJ0-2mix-extr is utilized for all ablation studies.

In the SEN, each AM-Transformer layer is accompanied by 7 Transformer layers. The Intra-Inter iteration is repeated twice. For the RSEN, we employ convolutional blocks stacked by 1-D convolutional layers to construct E_g , E_c , and D . The reference speech x is processed using STFT with a window length of 512, a hop length of 128, and an STFT window size of 512. The resulting spectrogram is then converted to a magnitude spectrogram so that F_x is 257. The dimensions d_g , d_c , d_s , and H are all set to 256. The variational approximation network \mathcal{V} is implemented using two four-layer fully connected networks to predict the mean and variance of the posterior distribution, respectively. The model encompasses a total of 45M parameters. The weights of $\mathcal{L}_{\text{SI-SNR}}$, \mathcal{L}_{REC} , \mathcal{L}_{KL} , $\mathcal{I}_{\text{VCLUB}}$, \mathcal{L}_{LL} and \mathcal{L}_{SIM} are set to 1, 10^{-3} , 10^{-4} , 10^{-4} , 10^{-3} and 10^{-3} , respectively, determined through a grid search.

Evaluation Metrics

To facilitate comparison with other methods, the performance of TSE is evaluated by SI-SNRi and SDRi (Vincent, Gribonval, and Févotte 2006), while speech quality is assessed by PESQ (Rix et al. 2001). During our experiment,

we observed that the occurrence probability of SC (SI-SNRi is negative) is not high for the entire extracted speech, but SC often occurs in specific speech chunks. Therefore, it is more appropriate to utilize the negative SI-SNRi ratio (NSR) for segmented speech rather than the entire speech (Zhang, He, and Zhang 2020; Zhao et al. 2022) to measure the probability of SC occurrence. To quantify this, we employ the chunk-wise SC measure metric r_{scr} (Liu et al. 2023b), which is defined as the ratio of the number of speech chunks with SC to the total number of active speech chunks:

$$M = \left\lceil \frac{T-L}{O} + 1 \right\rceil \quad (20)$$

$$S(k) = s(\hat{u}(k), u(k)) - s(y(k), u(k)) \quad (21)$$

$$N_{\text{sc}} = \sum_{k=1}^M \mathbb{I}(S(k) < 0) \quad (22)$$

$$N_{\text{valid}} = \sum_{k=1}^M \mathbb{I}(E(u(k)) > \eta) \cdot \mathbb{I}(E(\hat{u}(k)) > \eta) \quad (23)$$

$$r_{\text{scr}} = \frac{N_{\text{sc}}}{N_{\text{valid}}} \times 100\% \quad (24)$$

To calculate r_{scr} , the ground truth target speech u , predicted target speech \hat{u} , and mixture y need to be segmented into chunks. The total number of chunks M is determined by the speech length T , chunk length L , and hop length O . $\lceil \cdot \rceil$ denotes the ceiling function. For each chunk $k = 1, 2, \dots, M$, we calculate the chunk-level SI-SNR improvement $S(k)$, where s represents the SI-SNR. If $S(k) < 0$, it is considered that SC occurred in the k th chunk. \mathbb{I} denotes the indicator function. E represents the energy of the speech, and η is the energy-related threshold. L , O , and η are set to 250 ms, 125 ms, and 5% of the maximum energy, respectively.

Investigation of Disentangled Representations

To intuitively illustrate the information captured in the embedding vectors of reference speech at each phase, we employ t-SNE to visualize the spatial distribution of z_c , z_g , and z_s , as depicted in Figure 3(a), (b), and (c) respectively. It is apparent that z_c contains substantial overlapping information across different speakers, indicating its representation of speaker-independent semantic information. Although z_g exhibits speaker clustering to some extent, there is still a small overlap among different speakers, implying that the global information of speech represented by z_g may share similar characteristics across various speakers. In contrast, z_s exhibits distinct speaker clustering, with well-defined boundaries separating each speaker and notable distribution differences between male and female speakers. These results suggest that z_s effectively represents speaker identity.

To further demonstrate the significance of information disentanglement, we employ each disentangled information individually to guide the same speech extraction network, as depicted in Table 1. The AMLN method is utilized to modulate the SEN with z_g and z_s . For z_g , average pooling is applied along its temporal dimension, resulting in an embedding vector that serves as a condition. Cross-attention mechanism (Liu et al. 2023b; Lin et al. 2023) is employed to fuse

ID	CI (z_c)	GI (z_g)	SI (z_s)	SI-SNRi \uparrow	SDRi \uparrow	PESQ \uparrow	$r_{\text{scr}}\downarrow$
1	×	×	✓	19.9	20.2	3.85	6.89
2	×	✓	×	19.3	19.6	3.81	7.22
3	✓	×	×	15.6	15.9	3.50	12.14
4	✓	✓	×	17.5	17.8	3.61	9.06
5	✓	×	✓	17.9	18.3	3.63	8.35

Table 1: Experimental results involve using various information representations and their combination to guide the SEN on the WSJ0-2mix-extr dataset. CI, GI, and SI represent semantic, global, and speaker identity information.

z_c with the input of both the Intra- and Inter-Transformer to avoid the disruption of content information with dynamic local features caused by global pooling.

The findings in Table 1 demonstrate that utilizing only the disentangled z_s to guide the SEN yields the most favourable outcomes with the lowest SC ratio (ID-1). When employing z_g as the guidance without the GIDN (ID-2), all performance metrics deteriorate, confirming the channel attention’s capability to suppress harmful global information. Using solely z_c as the guidance (ID-3) yields the poorest results, particularly with r_{scr} increasing to 12.14, indicating a severe SC problem. This signifies that the semantic information in the reference speech offers little guidance in aiding TSE. When combining z_c with either z_g or z_s (ID-4 and ID-5), the performance metrics achieve suboptimal levels compared to ID-2 and ID-1, further affirming the necessity of information disentanglement. The results in Table 1 align with the distribution of each embedding illustrated in Figure 3.

Ablation Study

To assess the effectiveness of each module and loss function in our proposed method, we conducted ablation experiments on the IN layer of E_c , MI minimization, \mathcal{L}_{SIM} , \mathcal{L}_{REC} , \mathcal{L}_{KL} , and 2-D PE. Each ablation experiment was conducted independently while keeping all other components unchanged. The results are summarized in Table 2. It is evident that all the components contribute to improving the overall performance. The IN layer acts as an information bottleneck for E_c to help filter out global information. MI minimization filters out semantic and global information from the output of E_g and E_c , respectively. \mathcal{L}_{SIM} filters out harmful global information from the speaker identity information. These three components complement each other, and all contribute positively to information disentanglement. The inclusion of \mathcal{L}_{REC} and \mathcal{L}_{KL} ensures that E_g and E_c learn meaningful global and semantic features, respectively. The utilization of 2-D PE also enhances the performance of the SEN by facilitating better comprehension of the relative positional relationships at both intra- and inter-chunk time steps.

Next, we compared various modulation policies for acoustic representations in SEN and speaker embeddings, as outlined in Table 3. The results indicate that simplistic addition and concatenation methods yield the worst performance. This is because such approaches fail to ensure the preservation of undisturbed acoustic representation information during fusion. Compared to Gated Conv (Liu and




Figure 3: 2-D visualization of the spatial distribution of z_c , z_g , and z_s for the reference speech of five different speakers on the WSJ0-2mix-extr test set. Speakers are labelled as M (male) and F (female).

Method	SI-SNRi \uparrow	SDRi \uparrow	PESQ \uparrow	r_{scr} \downarrow
SDR-TSE	19.9	20.2	3.85	6.89
w/o IN	18.2	18.4	3.69	8.55
w/o MIM	17.3	17.6	3.59	8.60
w/o \mathcal{L}_{SIM}	18.4	18.7	3.72	8.36
w/o \mathcal{L}_{REC}	16.2	16.5	3.50	9.92
w/o \mathcal{L}_{KL}	18.8	19.1	3.78	7.61
w/o 2-D PE	19.4	19.6	3.81	7.27

Table 2: Results of ablation experiments on the WSJ0-2mix-extr dataset.

Modulation Policy	SI-SNRi \uparrow	SDRi \uparrow	PESQ \uparrow	r_{scr} \downarrow
Summation	18.7	18.9	3.72	8.23
Concatenation	19.0	19.3	3.74	8.07
Gated Conv (2022)	19.3	19.5	3.80	7.63
ConSM (2023)	19.2	19.5	3.77	7.88
AMLN (Ours)	19.9	20.2	3.85	6.89

Table 3: Experimental results of employing various modulation policies on the WSJ0-2mix-extr dataset.

Xie 2022) and ConSM (Chen et al. 2023), our method possesses the advantage of applying layer normalization to the acoustic representation prior to adaptive modulation. This sequential process enhances the perception of speaker embeddings by preserving their inherent identity information, as we observed that normalizing after modulation could disrupt speaker identity information.

Comparison with the State-of-the-Art

We compared our method with SOTA TSE methods on the WSJ0-2mix-extr and WSJ0-2mix datasets, and the results are summarized in Table 4 and Table 5. Our method outperforms other methods across all metrics. This performance improvement can be attributed to two key factors. Firstly, we employ the SOTA SS framework Sepformer (Subakan et al. 2021) as the backbone, effectively leveraging the dual-path framework’s ability to capture both long- and short-term dependencies in sequences, along with the powerful sequential modelling capability of Transformer. Secondly, our method benefits from the utilization of the DRL. Notably, our RSEN and GIDN are trained self-supervised. Despite the absence

Method	SI-SNRi \uparrow	SDRi \uparrow	PESQ \uparrow	r_{scr} \downarrow
TseNet (2019)	12.2	12.6	3.14	-
SpEx (2020)	14.2	14.6	3.36	9.68
SpEx+ (2020)	15.7	15.9	3.49	9.29
DPRNN-Spe-IRA (2021)	17.5	17.7	3.62	-
SDR-TSE (Ours)	19.9	20.2	3.85	6.89

Table 4: Performance comparison with SOTA methods on the WSJ0-2mix-extr dataset.

Method	SI-SNRi \uparrow	SDRi \uparrow	PESQ \uparrow	r_{scr} \downarrow
SpEx (2020)	15.8	16.3	3.14	10.42
SpEx+ (2020)	16.9	17.2	3.45	9.68
DPRNN-Spe-IRA (2021)	17.3	17.6	3.43	-
X-SepFormer (S_{sc}) (2023b)	19.1	19.7	3.75	8.56
X-SepFormer (S_{wt}) (2023b)	18.8	19.3	3.74	8.03
SDR-TSE (Ours)	19.6	19.9	3.82	6.67

Table 5: Performance comparison with SOTA methods on the WSJ0-2mix dataset.

of speaker identity labels, our method still achieves SOTA performance. As illustrated in Table 5, when compared to X-SepFormer (Liu et al. 2023b), which also adopts Sepformer as the underlying architecture, our method attains better performance by incorporating DRL and AM-Transformer. Our approach showcases a significant decrease of 1.36 in r_{scr} , indicating its efficacy in addressing the SC problem.

Conclusion

This paper introduces SDR-TSE, a novel approach to tackle the speaker confusion problem in TSE from the perspective of information disentanglement. Our self-supervised DRL policy disentangles the speaker identity information from the reference speech in two phases, providing effective guidance for TSE. Additionally, we propose the AM-Transformer, which integrates the AMLN to preserve the acoustic representation’s information in the SEN and enhance its perception of speaker embeddings. Through extensive experiments, we showcase our meticulously designed method’s strong information disentanglement capability and its exceptional performance in TSE.

References

- Afouras, T.; Owens, A.; Chung, J. S.; and Zisserman, A. 2020. Self-Supervised Learning of Audio-Visual Objects from Video. In *ECCV (18)*, volume 12363 of *Lecture Notes in Computer Science*, 208–224. Springer.
- Ba, L. J.; Kiros, J. R.; and Hinton, G. E. 2016. Layer Normalization. *CoRR*, abs/1607.06450.
- Bengio, Y.; Courville, A. C.; and Vincent, P. 2013. Representation Learning: A Review and New Perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(8): 1798–1828.
- Chen, J.; Rao, W.; Wang, Z.; Lin, J.; Ju, Y.; He, S.; Wang, Y.; and Wu, Z. 2023. MC-SpEx: Towards Effective Speaker Extraction with Multi-Scale Interfusion and Conditional Speaker Modulation. *CoRR*, abs/2306.16250.
- Chen, Y.; Wu, D.; Wu, T.; and Lee, H. 2021. Again-VC: A One-Shot Voice Conversion Using Activation Guidance and Adaptive Instance Normalization. In *ICASSP*, 5954–5958. IEEE.
- Cheng, P.; Hao, W.; Dai, S.; Liu, J.; Gan, Z.; and Carin, L. 2020. CLUB: A Contrastive Log-Ratio Upper Bound of Mutual Information. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, 1779–1788. PMLR.
- Choi, H.; Lee, J.; Kim, W.; Lee, J.; Heo, H.; and Lee, K. 2021. Neural Analysis and Synthesis: Reconstructing Speech from Self-Supervised Representations. In *NeurIPS*, 16251–16265.
- Chou, J.; and Lee, H. 2019. One-Shot Voice Conversion by Separating Speaker and Content Representations with Instance Normalization. In *INTERSPEECH*, 664–668. ISCA.
- Deng, C.; Ma, S.; Sha, Y.; Zhang, Y.; Zhang, H.; Song, H.; and Wang, F. 2021. Robust Speaker Extraction Network Based on Iterative Refined Adaptation. In *Interspeech*, 3530–3534. ISCA.
- Elminshawi, M.; Mack, W.; Chakrabarty, S.; and Habets, E. A. P. 2022. New Insights on Target Speaker Extraction. *CoRR*, abs/2202.00733.
- Gao, R.; and Grauman, K. 2021. VisualVoice: Audio-Visual Speech Separation With Cross-Modal Consistency. In *CVPR*, 15495–15505. Computer Vision Foundation / IEEE.
- Ge, M.; Xu, C.; Wang, L.; Chng, E. S.; Dang, J.; and Li, H. 2020. SpEx+: A Complete Time Domain Speaker Extraction Network. In *INTERSPEECH*, 1406–1410. ISCA.
- Gu, R.; Chen, L.; Zhang, S.; Zheng, J.; Xu, Y.; Yu, M.; Su, D.; Zou, Y.; and Yu, D. 2019. Neural Spatial Filter: Target Speaker Speech Separation Assisted with Directional Information. In *INTERSPEECH*, 4290–4294. ISCA.
- Hershey, J. R.; Chen, Z.; Roux, J. L.; and Watanabe, S. 2016. Deep Clustering: Discriminative Embeddings for Segmentation and Separation. In *ICASSP*, 31–35. IEEE.
- Hou, N.; Xu, C.; Chng, E. S.; and Li, H. 2021. Learning Disentangled Feature Representations for Speech Enhancement Via Adversarial Training. In *ICASSP*, 666–670. IEEE.
- Kaya, E. M.; and Elhilali, M. 2017. Modelling Auditory Attention. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1714): 20160101.
- Kolbaek, M.; Yu, D.; Tan, Z.; and Jensen, J. 2017. Multitalker Speech Separation With Utterance-Level Permutation Invariant Training of Deep Recurrent Neural Networks. *IEEE ACM Trans. Audio Speech Lang. Process.*, 25(10): 1901–1913.
- Lin, J.; Cai, X.; Dinkel, H.; Chen, J.; Yan, Z.; Wang, Y.; Zhang, J.; Wu, Z.; Wang, Y.; and Meng, H. 2023. AV-SepFormer: Cross-Attention SepFormer for Audio-Visual Target Speaker Extraction. *CoRR*, abs/2306.14170.
- Liu, G.; Zhang, Y.; Lei, Y.; Chen, Y.; Wang, R.; Li, Z.; and Xie, L. 2023a. PromptStyle: Controllable Style Transfer for Text-to-Speech with Natural Language Descriptions. *CoRR*, abs/2305.19522.
- Liu, K.; Du, Z.; Wan, X.; and Zhou, H. 2023b. X-SepFormer: End-to-End Speaker Extraction Network with Explicit Optimization on Speaker Confusion. *CoRR*, abs/2303.05023.
- Liu, M.; Breuel, T. M.; and Kautz, J. 2017. Unsupervised Image-to-Image Translation Networks. In *NIPS*, 700–708.
- Liu, W.; and Xie, C. 2022. Gated Convolutional Fusion for Time-Domain Target Speaker Extraction Network. In *INTERSPEECH*, 5368–5372. ISCA.
- Luo, Y.; Chen, Z.; and Yoshioka, T. 2020. Dual-Path RNN: Efficient Long Sequence Modeling for Time-Domain Single-Channel Speech Separation. In *ICASSP*, 46–50. IEEE.
- Luo, Y.; and Mesgarani, N. 2019. Conv-TasNet: Surpassing Ideal Time-Frequency Magnitude Masking for Speech Separation. *IEEE ACM Trans. Audio Speech Lang. Process.*, 27(8): 1256–1266.
- Mesgarani, N.; and Chang, E. F. 2012. Selective Cortical Representation of Attended Speaker in Multi-Talker Speech Perception. *Nature*, 485(7397): 233–236.
- Min, D.; Lee, D. B.; Yang, E.; and Hwang, S. J. 2021. MetaStyleSpeech: Multi-Speaker Adaptive Text-to-Speech Generation. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, 7748–7759. PMLR.
- Mu, Z.; Yang, X.; Yang, X.; and Zhu, W. 2023. A Multi-Stage Triple-Path Method for Speech Separation in Noisy and Reverberant Environments. *CoRR*, abs/2303.03732.
- Mu, Z.; Yang, X.; and Zhu, W. 2023. Multi-Dimensional and Multi-Scale Modeling for Speech Separation Optimized by Discriminative Learning. *CoRR*, abs/2303.03737.
- Nekvinda, T.; and Dusek, O. 2020. One Model, Many Languages: Meta-Learning for Multilingual Text-to-Speech. In *INTERSPEECH*, 2972–2976. ISCA.
- Qian, K.; Zhang, Y.; Gao, H.; Ni, J.; Lai, C.; Cox, D. D.; Hasegawa-Johnson, M.; and Chang, S. 2022. ContentVec: An Improved Self-Supervised Speech Representation by Disentangling Speakers. In *ICML*, volume 162 of *Proceedings of Machine Learning Research*, 18003–18017. PMLR.
- Raisi, Z.; Naiel, M. A.; Fieguth, P.; Wardell, S.; and Zelek, J. 2020. 2D Positional Embedding-Based Transformer for Scene Text Recognition. *Journal of Computational Vision and Imaging Systems*, 6(1): 1–4.

- Rix, A. W.; Beerends, J. G.; Hollier, M. P.; and Hekstra, A. P. 2001. Perceptual Evaluation of Speech Quality (PESQ)-a New Method for Speech Quality Assessment of Telephone Networks and Codecs. In *ICASSP*, 749–752. IEEE.
- Roux, J. L.; Wisdom, S.; Erdogan, H.; and Hershey, J. R. 2019. SDR - Half-baked or Well Done? In *ICASSP*, 626–630. IEEE.
- Subakan, C.; Ravanelli, M.; Cornell, S.; Bronzi, M.; and Zhong, J. 2021. Attention Is All You Need in Speech Separation. In *ICASSP*, 21–25. IEEE.
- Vincent, E.; Gribonval, R.; and Févotte, C. 2006. Performance Measurement in Blind Audio Source Separation. *IEEE Trans. Speech Audio Process.*, 14(4): 1462–1469.
- Wang, D.; and Chen, J. 2018. Supervised Speech Separation Based on Deep Learning: An Overview. *IEEE ACM Trans. Audio Speech Lang. Process.*, 26(10): 1702–1726.
- Wang, Q.; Muckenhirn, H.; Wilson, K. W.; Sridhar, P.; Wu, Z.; Hershey, J. R.; Saurous, R. A.; Weiss, R. J.; Jia, Y.; and Lopez-Moreno, I. 2019. VoiceFilter: Targeted Voice Separation by Speaker-Conditioned Spectrogram Masking. In *INTERSPEECH*, 2728–2732. ISCA.
- Wang, X.; Chen, H.; Tang, S.; Wu, Z.; and Zhu, W. 2022. Disentangled Representation Learning. *CoRR*, abs/2211.11695.
- Woo, S.; Park, J.; Lee, J.; and Kweon, I. S. 2018. CBAM: Convolutional Block Attention Module. In *ECCV (7)*, volume 11211 of *Lecture Notes in Computer Science*, 3–19. Springer.
- Wu, Y.; Tan, X.; Li, B.; He, L.; Zhao, S.; Song, R.; Qin, T.; and Liu, T. 2022. AdaSpeech 4: Adaptive Text to Speech in Zero-Shot Scenarios. In *INTERSPEECH*, 2568–2572. ISCA.
- Xu, C.; Rao, W.; Chng, E. S.; and Li, H. 2019. Time-Domain Speaker Extraction Network. In *ASRU*, 327–334. IEEE.
- Xu, C.; Rao, W.; Chng, E. S.; and Li, H. 2020. SpEx: Multi-Scale Time Domain Speaker Extraction Network. *IEEE ACM Trans. Audio Speech Lang. Process.*, 28: 1370–1384.
- Yu, D.; Kolbæk, M.; Tan, Z.; and Jensen, J. 2017. Permutation Invariant Training of Deep Models for Speaker-Independent Multi-Talker Speech Separation. In *ICASSP*, 241–245. IEEE.
- Zhang, Z.; He, B.; and Zhang, Z. 2020. X-TaSNet: Robust and Accurate Time-Domain Speaker Extraction Network. In *INTERSPEECH*, 1421–1425. ISCA.
- Zhao, J.; Gao, S.; and Shinozaki, T. 2020. Time-Domain Target-Speaker Speech Separation with Waveform-Based Speaker Embedding. In *INTERSPEECH*, 1436–1440. ISCA.
- Zhao, Z.; Yang, D.; Gu, R.; Zhang, H.; and Zou, Y. 2022. Target Confusion in End-to-End Speaker Extraction: Analysis and Approaches. In *INTERSPEECH*, 5333–5337. ISCA.