

Tune-In: Training Under Negative Environments with Interference for Attention Networks Simulating Cocktail Party Effect

Jun Wang¹, Max W. Y. Lam¹, Dan Su¹, Dong Yu²

¹ Tencent AI Lab, Shenzhen, China

² Tencent AI Lab, Bellevue WA, USA

{joinerwang, maxwylam, dansu, dyu}@tencent.com

Abstract

We study the cocktail party problem and propose a novel attention network called Tune-In, abbreviated for training under negative environments with interference. It firstly learns two separate spaces of speaker-knowledge and speech-stimuli based on a shared feature space, where a new block structure is designed as the building block for all spaces, and then cooperatively solves different tasks. Between the two spaces, information is cast towards each other via a novel cross- and dual-attention mechanism, mimicking the bottom-up and top-down processes of a human’s cocktail party effect. It turns out that substantially discriminative and generalizable speaker representations can be learnt in severely interfered conditions via our self-supervised training. The experimental results verify this seeming paradox. The learnt speaker embedding has superior discriminative power than a standard speaker verification method; meanwhile, Tune-In achieves remarkably better speech separation performances in terms of SI-SNRi and SDRi consistently in all test modes, and especially at lower memory and computational consumption, than state-of-the-art benchmark systems.

Introduction

The cocktail party effect (Narayan et al. 2009; Evans et al. 2015; Getzmann, Jasny, and Falkenstein 2016; Getzmann 2015) is the phenomenon of the ability of a human brain to focus one’s auditory attention to “tune in” a single voice and “tune out” the competing others. Although how humans solve the cocktail party problem (Cherry 1953) remains unknown, remarkable progress on speech separation (SS) has been made thanks to recent advances of deep-learning models. Ground-breaking successful models include the high-dimensional embedding-based methods, originally proposed as the deep clustering network (DPCL) (Hershey et al. 2016), and its extensions such as DANet (Chen, Luo, and Mesgarani 2017) and DENet (Wang et al. 2018). More recently, the performance has been improved by the time-domain audio separation networks (TasNet) (Luo and Mesgarani 2018) and the time convolutional networks (TCNs) based Conv-TasNet (Bai, Kolter, and Koltun 2018; Luo and Mesgarani 2019; Lam et al. 2020). Until very recently, State-of-the-art (SOTA) performance record of TCNs have been further advanced on

several benchmark datasets by the dual-path RNN (DPRNN) (Luo, Chen, and Yoshioka 2019) and the dual-path gated RNN (DPGRNN) (Nachmani, Adi, and Wolf 2020).

Modern deep-learning models also evolve rapidly for speaker identification and verification tasks (Snyder et al. 2018; Yi Liu 2019; Mirco Ravanelli 2018). A separately trained speaker identification network is adopted in (Nachmani, Adi, and Wolf 2020) to prevent channel swap for SS. Other emerging studies (Zeghidour and Grangier 2020; Shi et al. 2018; Xu et al. 2018; Shi, Xu, and Xu 2019; Wang et al. 2019; Shi, Liu, and Han 2020) suggest bridging the speaker identification and the SS tasks in a unified network. Especially, Wavesplit in (Zeghidour and Grangier 2020) jointly trains the speaker stacks and the separation stack to outperform DPRNN and DPGRNN. Similar to (Drude, von Neumann, and Haeb-Umbach 2018), it is based on K-means clustering, which results in significantly higher computational cost (Pariente et al. 2020). According to its reported parameter setting, the estimated model size is approximately 42.5M, about 18 times larger than ours. To the best of our knowledge, higher SOTA scores than DPRNN’s were all achieved by remarkably larger models. For another example, (Nachmani, Adi, and Wolf 2020) proposed a model of 7.5M parameters, about 3 times larger than ours; still, our system outperforms theirs in terms of SI-SNRi (Le Roux et al. 2019) and SDRi (Shi et al. 2018) on a benchmark dataset. We believe that it is only meaningful when the scores are compared under a fair constraint on the model sizes, as what we followed carefully and convincingly throughout our experiments.

A neurobiology study in (Evans et al. 2015) examines the modulation of neural activity associated with the interference properties. It demonstrates that a competing speech is processed predominantly within the same pathway as the target speech in the left hemisphere but is not treated equivalently within that stream. Another inspiring research (Getzmann, Jasny, and Falkenstein 2016) suggests an acceleration in attention switching and context updating when people have semantically cued changes (e.g., a word) in target speaker settings than uncued changes. These recent neurobiology studies (Evans et al. 2015; Getzmann, Jasny, and Falkenstein 2016; Getzmann 2015; Narayan et al. 2009) together suggest that, when listening with interference, the neural processes are more sophisticated than the simplified presumption in the previous attention-based separation (Shi et al. 2018).

Inspired by the above, we propose a novel attention network that entails fundamentally different bottom-up and top-down mechanisms in stark contrast to the prior arts. Our contribution lies mainly in four folds:

- a “globally attentive, locally recurrent” (GALR) architecture (Sec.) that breaks the memory and computation bottlenecks of self-attention and permits its usage over very long sequences (e.g., raw waveforms of speech mixtures);
- a novel cross- and dual-attention mechanism (Sec.), which enables the bottom-up and top-down casting of signal-level and speaker-level information in parallel;
- a theoretically grounded contrastive estimation loss, namely the Tune-InCE loss (Sec.) that can significantly enhance the robustness and generalizability of the learnt speaker embedding;
- a new training and inference framework (Sec.) that efficiently solves the speaker embedding and speech separation problems and achieves reciprocity between the two tasks. In particular, our system can perform speech verification (SV) directly on overlapped mixtures with a performance comparable to using clean utterances, as discussed in Sec. .

Conventional SV approaches require a complicated pipeline including first a speech activity detection module to remove noisy or silent parts, followed by a segmentation module, of which the output short segments are then grouped into corresponding speakers by a clustering module. Besides, to deal with interfered signals, an overlap detector and classifier are needed to remove the overlapped parts. Recent efforts have been made (Huang et al. 2020) to simplify such complicated pipelines into one stage. Still, performances are severely hurt in highly interfered scenarios (Fujita et al. 2019). In contrast, our framework has a quite different principle so that it can be relieved from the long pipeline. We call it “Tune-In”, abbreviated for “Training under negative environments with Interference” to mimic the cocktail party effect, where a cross-attention mechanism automatically retrieves and summarizes the most salient, relevant, and reliable speaker embedding from a non-purified embedding pool, while discarding the noisy, vacant, unreliable, or non-relevant parts. Therefore, the summarized embedding is expected to capture the speaker’s characteristics that are most salient and relevant to the current observation.

“Tune-In” together with the “Contrastive Estimation” loss (that is the “Tune-InCE” loss) views the commonplace masking routine in SS from a novel perspective, as a self-supervised technique for robust representation learning, like the masking in BERT (Devlin et al. 2019). As far as we know, our model is also the smallest among all SOTA SS models, with an 11.5% model size reduction relative to the previously smallest model — DPRNN. Our model also achieved the SOTA SI-SNRi and SDRi scores at a much lower cost of up to 62.9% less run-time memory and 66.4% fewer computational operations according to results in Sec. , where more insights are discussed in the ablation studies. We also provide theoretical and empirical analysis in the Appendix.

Tune-In Framework

Overview

As shown in Fig. 1 (a), the goal of the “Speech-stimuli space” is to separate the C sources from an input mixture $\mathbf{x} = \sum_{c=1}^C \mathbf{c}_i \cdot \mathbf{s}_c$, where \mathbf{c}_i is a scaling factor to generate various signal-to-interference ratio (SIR) conditions (e.g., randomly from 0 to 5dB), and the sources are by C speakers from the overall N -speaker training corpus; the goal of the “Speaker-knowledge space” is to extract speaker embedding to discriminate the N speakers; the goal of the “generic-feature space” is to extract shared generic features with discriminative and separable power among speakers and sources for both down-stream tasks.

The Encoder and the Decoder are the same as those in (Luo and Mesgarani 2019; Luo, Chen, and Yoshioka 2019) for comparison purposes, and we ask our readers to refer to these prior works for the detailed structure. The Encoder transforms the mixture \mathbf{x} of the C waveform sources into a sequential input $\tilde{\mathbf{X}} \in \mathbb{R}^{D \times I}$, where D is the feature dimension, and I is the time-sequence length. We split $\tilde{\mathbf{X}}$ into S half-overlapping segments each of length K , and pack them into a 3-D tensor $\mathbf{X} \in \mathbb{R}^{D \times S \times K}$, and then pass it to a stack of blocks. These blocks share the novel structure namely the GALR structure. These GALR blocks stack together to accomplish tasks in all the spaces: 1) B blocks in the generic-feature space constitutes the basis part that works on generic feature extraction. 2) B_1 and B_2 GALR blocks are fine-tuned towards different downstream tasks of speaker identification and speech separation, respectively.

GALR Block

As shown in Fig. 1 (b), we propose the GALR block structure to break the memory and computational bottleneck of self-attention mechanisms on tasks dealing with very long sequences. Each GALR block receives a 3-D tensor $\mathbf{X}^{(b-1)} \in \mathbb{R}^{D \times S \times K}$ and outputs $\mathbf{X}^{(b)}$ of the same shape, for $b = 1, \dots, B$, where B denotes the number of blocks in the generic-feature space. To avoid confusion, we denote $\mathbf{X}_{\text{SK}}^{(b)}$ as the output of GALR blocks in the speaker-knowledge (SK) space with $b = 1, \dots, B_1$, and correspondingly $\mathbf{X}_{\text{SS}}^{(b)}$ for the speech-stimuli (SS) space with $b = 1, \dots, B_2$. Each GALR block consists of two primary layers: a locally recurrent (LR) layer that performs intra-segment processing, and a globally attentive (GA) layer that performs inter-segment processing.

Locally Recurrent Layer Similar to DPRNN, we use RNNs for modeling local dependencies at the lower intra-segment context level, e.g., signal continuity, signal structure, etc, which are inherently important for waveform reconstruction. Consequently, a BiLSTM layer is responsible for modeling short-term dependencies within each segment and outputs $\mathbf{L}^{(b)} \in \mathbb{R}^{D \times S \times K}$, where $\mathbf{L}_s^{(b)} = \text{LN} \left(\text{Lin}^{(b)} (\text{BiLSTM}^{(b)} (\mathbf{X}_{:,s,:}^{(b-1)})) \right) + \mathbf{X}_{:,s,:}^{(b-1)}$, for $s = 1, \dots, S$ refers to the local sequence within the s^{th} segment, and $\text{BiLSTM}(\cdot)$, $\text{Lin}(\cdot)$, $\text{LN}(\cdot)$ denote a Bi-directional LSTM layer, a linear projection, and a layer normalization module (Ba, Kiros, and Hinton 2016), respectively.

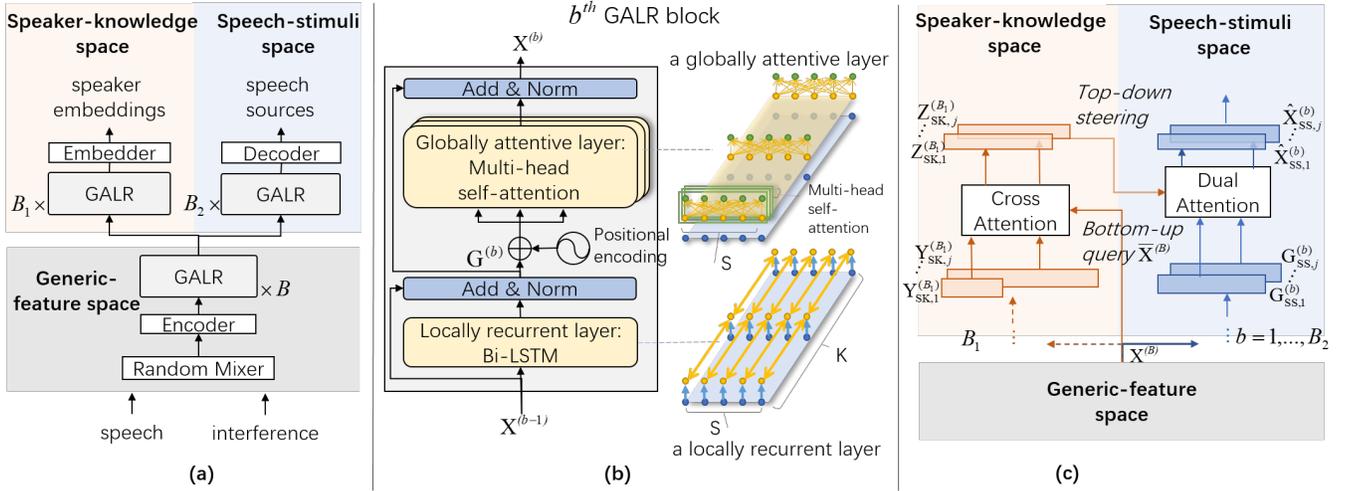


Figure 1: (a) An overview of the proposed Tune-In framework, comprised of three spaces: a generic-feature space, a speaker-knowledge space, and a speech-stimuli space. (b) the structure inside a Globally Attentive, Locally Recurrent (GALR) block. (c) the bottom-up and top-down information casting between the two spaces via a cross- and a dual- attention mechanism.

Globally Attentive Layer RNNs are far more sensitive to the nearby elements than to the distant ones. For example, an LSTM-based language modeling (LM) is capable of using about 200 tokens of context on average, as revealed by (Khandelwal et al. 2018), but sharply distinguishes nearby context (recent 50 tokens) from the distant history. As for speech, high temporal correlation, continuous acoustic-signal structure generally exist more often in intra-segment sequences (0.25ms- 2ms granularity in our settings) than in inter-segment sequences (64ms- 512ms granularity). This suggests that RNNs are potentially more suited for intra-segment modeling than for inter-segment modeling. Moreover, (Ravanelli et al. 2018) discovered that RNNs reset the stored memory to avoid bias towards an “unrelated” history.

One strength of attention mechanisms over RNNs lies in a fully connected sequence processing strategy (Vaswani et al. 2017), where every element is connected to other elements in a sequence via a direct path without any recursively processing, memorizing reset, or update mechanisms like RNNs. However, a conventional self-attention layer remains impractical to be applied in speech separation tasks, as very long speech signal sequences are involved, thus run-time memory consumption becomes a bottleneck. In terms of computational complexity, self-attention layers are also slower than recurrent layers when the sequence length is larger than the representation dimensionality, like most in speech separation.

Our proposed GALR block resorts to the self-attention mechanism (Vaswani et al. 2017) for inter-segment computation to model context-aware global dependencies. It revises DPRNN to better address long-range context modeling for audio separation, leading to a lower-cost, higher-performance structure. Since the hyperparameter K in GALR controls the granularity of the local information modeled by the locally recurrent layer, the overall sequence length I that the attention model deals with is reduced by a factor of K , which considerably relieves the run-time memory consumption. Note that

self-attention has been used in spectrogram-based source enhancement by (Zadeh et al. 2019; Koizumi et al. 2020; Chen, Mao, and Liu 2020) recently, but never orthogonally across RNN-summarized frames, as in our proposed method. To validate that self-attention and RNN are better summarizers across segments and within the segment, respectively, we performed an ablation study in Appendix 1.3.

Moreover, the efficiency can be further improved by reducing the times of running the global attention mechanism, which is originally K times over the S -length sequences in $\mathbf{L}^{(b)} \in \mathbb{R}^{D \times S \times K}$. We first pool the output of the locally recurrent model $\mathbf{L}^{(b)}$ through a 1×1 convolution $\text{Conv2D}(\cdot)$ to map the K -dimension feature into Q dimensions with $Q \ll K$. Then a layer normalization $\text{LN}_D(\cdot)$ (Ba, Kiros, and Hinton 2016) is applied along the feature dimension D . Subsequently, a positional embedding vector, \mathbf{P} , is added to the features: $\mathbf{G}^{(b)} = \text{LN}_D(\text{Conv2D}(\mathbf{L}^{(b)})) + \mathbf{P}$, where $\mathbf{G}^{(b)} \in \mathbb{R}^{D \times S \times Q}$. Thus, the computation is reduced by a proportion of $(K - Q)/K$. Then, we apply the multi-head self-attention $\text{Attn}(\cdot)$ to $\mathbf{G}^{(b)}$ following the encoder in Transformer (Vaswani et al. 2017), of which the detailed equations are omitted here for their widely known operations. The output $\hat{\mathbf{X}}^{(b)} = \text{Attn}(\mathbf{G}^{(b)}) \in \mathbb{R}^{D \times S \times Q}$ goes through another $2D \ 1 \times 1$ convolution to inversely map the Q dimensions back to K dimensions, yielding the final output of a GALR block $\mathbf{X}^{(b)} \in \mathbb{R}^{D \times S \times K}$. Likewise, each following block continues on the local-to-global and global-to-local (also fine-to-coarse and coarse-to-fine across the time axis) interactions.

Self-supervised Learning with Cross- and Dual-Attention Mechanism

On top of the generic-feature space, the following GALR blocks are tied to each downstream representation space.

Speech-stimuli Space We want to learn deep speech representations for separating source signals in a mixture. The output of the last GALR block in this space, $\mathbf{X}_{SS}^{(B_2)}$, is passed to the Decoder that is the same as in DPRNN (Luo, Chen, and Yoshioka 2019)¹.

Speaker-knowledge Space The goal here is to learn speaker representations that are easily separable and discriminative for identifying different speakers, including new unseen ones. We pass the output of the last GALR block $\mathbf{X}_{SK}^{(B_1)} \in \mathbb{R}^{D \times S_j \times K}$ to the Embedder, where it is projected along the feature dimension from D dimensions to $D' = C \times D$ dimensions. The resultant $(C \times D) \times S_j \times K$ dimensions are then averaged over the K intra-segment positions and lead to $C \times (D \times S_j)$ dimensional outputs, which are considered as C speaker features $\mathbf{Y}_{SK,j}^{(B_1)} \in \mathbb{R}^{D \times S_j}$, $j = 1, \dots, C$, where S_j denotes the number of segments in $\mathbf{Y}_{SK,j}^{(B_1)}$.

Cross Attention As shown in Fig.1 (c), we design a cross-and dual-attention mechanism, where the two spaces interact only via queries. The bottom-up queries from the speech-stimuli space retrieve the most relevant information from the speaker-knowledge space and filter out non-salient, noisy, or redundant information, so that our system is spared from the long and complex pre-processing steps (VAD, segmentation, overlap cutting, etc.) to "purify" the speaker embedding. Inspired by the cross-stitch (Misra et al. 2016) and the cross-attention methods (Hudson and Manning 2018) used between text and image (knowledge base), we are the first to use a cross-attention strategy between speaker embedding and speech signal representations. Query(\cdot), Key(\cdot), and Value(\cdot) denote linear transformation functions, where the corresponding input vectors are linearly projected along the feature dimension ($\mathbb{R}^D \mapsto \mathbb{R}^D$) into query, key, and value vectors, respectively. Our detailed implementation of these functions is the same as (Sperber et al. 2018) with code.²

The bottom-up query $\bar{\mathbf{X}}^{(B)} \in \mathbb{R}^{D \times S}$ is generated by averaging the generic-feature space's output $\mathbf{X}^{(B)} \in \mathbb{R}^{D \times S \times K}$ over the K intra-segment positions. Next, we cast the query vector onto the speaker-knowledge space via the cross-attention approach. Specifically, this is achieved by computing the inner product between Query($\bar{\mathbf{X}}^{(B)} \in \mathbb{R}^{D \times S}$) and Key($\mathbf{Y}_{SK,j}^{(B_1)} \in \mathbb{R}^{D \times S_j}$):

$$\mathbf{a}_{\text{crs},j} = \text{softmax}(\text{Query}(\bar{\mathbf{X}}^{(B)})^\top \cdot \text{Key}(\mathbf{Y}_{SK,j}^{(B_1)})). \quad (1)$$

yielding a cross-attention distribution $\mathbf{a}_{\text{crs},j} \in \mathbb{R}^{S \times S_j}$.

Finally, we compute the sum of Value($\mathbf{Y}_{SK,j}^{(B_1)} \in \mathbb{R}^{D \times S_j}$) weighted by the cross-attention distribution $\mathbf{a}_{\text{crs},j}$, and then average over S segments, to produce for speaker j a vector $\mathbf{Z}_{SK,j}^{(B_1)} \in \mathbb{R}^D$, which we call a steering vector in contrast to the following "autopilot" mode that does not rely this vector:

$$\mathbf{Z}_{SK,j}^{(B_1)} = \frac{1}{S} \sum_S \sum_{S_j} \mathbf{a}_{\text{crs},j} \cdot \text{Value}(\mathbf{Y}_{SK,j}^{(B_1)})^\top. \quad (2)$$

¹<https://github.com/ShiZiqiang/dual-path-RNNs-DPRNNs-based-speech-separation>

²<http://msperber.com/research/self-att>

For notation simplicity, we denote $\mathbf{Z}_{SK,j}^{(B_1)}$ as \mathbf{Z}_j from now on. It is expected to capture the speaker's characteristics that are most salient and relevant to the currently observed content. Note that the generalization power of the steering vector \mathbf{Z}_j is vital for maintaining the separation performance over unseen speakers. Therefore, based on Eq. (2), we investigated two techniques to increase the robustness of the steering vectors: (1) apply embedding dropout Dropout(0.1), and (2) add element-wise Gaussian embedding noise $\mathcal{N}(0, 0.1)$, to the steering vector \mathbf{Z}_j only at training time. We study the ablations of each part of these techniques in Sec. .

Dual Attention A top-down steering is raised from the speaker-knowledge space by projecting the steering vector onto the speech-stimuli space via the dual-attention approach. It mimics the neural process that uses specific object representations in auditory memory to enhance the perceptual precision during top-down attention (Lim, Westmann, and Obleser 2015). This is simulated here by firstly passing \mathbf{Z}_j through two linear mappings $\text{r}(\cdot)$ and $\text{h}(\cdot)$ to modulate the original self-attention input $\mathbf{G}_{SS,j}^{(b)}$ and to generate keys and values that capture the top-down speaker information. Then we compute the dual-attention distribution $\mathbf{a}_{\text{dual},j}$ as below:

$$\mathbf{a}_{\text{dual},j} = \text{softmax}(\text{Query}(\mathbf{G}_{SS,j}^{(b)})^\top \cdot \text{Key}(\text{r}(\mathbf{Z}_j) \odot \mathbf{G}_{SS,j}^{(b)} + \text{h}(\mathbf{Z}_j))) \quad (3)$$

where $\mathbf{G}_{SS,j}^{(b)}$ is the input of each globally attentive layer for $b = 0, \dots, B_2$ in the speech-stimuli space. Note that $\mathbf{G}_{SS,j}^{(0)}$ is equal for all $j = 1, \dots, C$, as they share the value computed from the last output $\mathbf{X}^{(B)}$ from the generic-feature space.

Finally, we sum the values weighted by $\mathbf{a}_{\text{dual},j}$ to produce the modified self-attention output:

$$\hat{\mathbf{X}}_{SS,j}^{(b)} = \sum_S \mathbf{a}_{\text{dual},j} \cdot \text{Value}(\text{r}(\mathbf{Z}_j) \odot \mathbf{G}_{SS,j}^{(b)} + \text{h}(\mathbf{Z}_j))^\top \quad (4)$$

Notably, Eq. (3-4) together can be seen as a replacement of the original self-attention operation in each globally attentive layer in the speech-stimuli space. In Appendix 1.4, we empirically validate and analyze the effectiveness of dual attention by comparing it to another successful approach — feature-wise linear modulation (FiLM) (Perez et al. 2017).

Losses

In the speech-stimuli space, we use the scale-invariant signal-to-noise ratio (SI-SNR) (Le Roux et al. 2019) loss, $\mathcal{L}_{\text{SI-SNR}}$, to reconstruct the source signals using an utterance-level permutation invariant training (u-PIT) method (Yu et al. 2017). However, losses such as SI-SNR expend heavy computation at reconstructing every detail of the signals, while often ignoring the global context. Instead, in the speaker-knowledge space, we want to learn robust representations that entail the underlying shared-speaker context among one speaker's different utterances of speech signals under different interfering signals. Therefore, we design more appropriate losses ($\mathcal{L}_{\text{InCE}} + \mathcal{L}_{\text{reg}}$) below for extracting the underlying shared information between the extended sequence context. All in all,

the joint loss is $\mathcal{L}_{\text{SI-SNR}} + \lambda(\mathcal{L}_{\text{InCE}} + \mathcal{L}_{\text{reg}})$, where λ is a weighting factor. Appendix 1.7 elaborates the speaker-embedding-based permutation computation for training speedup.

Self-supervised Loss $\mathcal{L}_{\text{InCE}}$ is our proposed self-supervised loss called Tune-InCE loss defined as follows:

$$\mathcal{L}_{\text{InCE}} = -\mathbb{E}_{\mathcal{D}} \left[\log \left(f(\mathbf{Z}_j, \mathbf{E}_{i_j}) / \sum_{i=1}^N f(\mathbf{Z}_j, \mathbf{E}_i) \right) \right], \quad (5)$$

$$f(\mathbf{Z}_j, \mathbf{E}_i) = \exp(-\alpha \|\mathbf{Z}_j - \mathbf{E}_i\|_2^2), \quad (6)$$

where \mathbf{E}_i is the speaker embedding vector defined for all $i = 1, \dots, N$ speakers in the training set, i_j indicates the corresponding speaker index for speaker $j = 1, \dots, C$ in the current mixture, $\alpha > 0$ is a learnable scaler, and $\mathbb{E}_{\mathcal{D}} [\cdot]$ denotes the expectation over the training set \mathcal{D} containing all input mixture utterances and corresponding speaker IDs.

In Appendix 1.1, we prove that our proposed form of Eq. (6) about $f(\mathbf{Z}_j, \mathbf{E}_i)$ corresponds to treating each speaker embedding \mathbf{E}_{i_j} as a cluster centroid of different steering vectors \mathbf{Z}_j of utterances generated by the same speaker i_j while $\alpha > 0$ controls the cluster size. $f(\mathbf{Z}_j, \mathbf{E}_i)$ can also be regarded as a regularized variant of the contrastive predictive coding (CPC) loss (Oord, Li, and Vinyals 2018; Pascual et al. 2019; Ravanelli and Bengio 2019; Chorowski et al. 2019; Ravanelli et al. 2020). We also prove that by minimizing $\mathcal{L}_{\text{InCE}}$ we maximize a lower bound of the mutual information between the steering vector \mathbf{Z}_j and the speaker vector \mathbf{E}_{i_j} , which encourages learning separable inter-speaker utterance embeddings and compact intra-speaker utterance embeddings. During training, whenever a steering vector of source j by the speaker i_j is generated, the corresponding speaker embedding vector is updated using exponential moving average (EMA): $\Delta \mathbf{E}_{i_j} = \varepsilon(\mathbf{Z}_j - \mathbf{E}_{i_j})$, where $\varepsilon \in [0, 1]$ is a small hyper-parameter scalar that controls the update step size.

\mathcal{L}_{reg} is a regularization loss that avoids collapsing to a trivial solution of all zeros:

$$\mathcal{L}_{\text{reg}} = -\frac{1}{\gamma C} \sum_{j=1}^C \min_{i \in \{1, \dots, N\} \setminus \{i_j\}} \log \|\mathbf{E}_{i_j} - \mathbf{E}_i\|_1, \quad (7)$$

where $\gamma > 0$ is a weighting factor. See Table ?? for \mathcal{L}_{reg} 's ablation study.

Training and Inference Modes

We train and test the proposed Tune-In system sequentially in different modes.

1) “Autopilot” mode: performing a standalone speech separation without any speaker knowledge, when the training graph only traverses the generic-feature space and the speech-stimuli space, i.e., the forward and backward path goes only along the blue lines as shown in Fig. 1 (c), and only the SI-SNR loss is used. Since in “Autopilot” mode there is no top-down steering information from the speaker-knowledge space, the outputs of $r(\cdot)$ and $h(\cdot)$ in Eq. (3-4) are an all-one vector and an all-zero vector, respectively.

2) “Online” mode: performing concurrent speaker-knowledge extraction and speech separation, based on the same online input signals, i.e., $S_j = S$. Now the training

forward path goes along both the blue and the brick-red lines in Fig. 1 (c). In this mode, the computation in the generic-feature space is shared by both down-streaming spaces.

3) “Offline” mode: performing offline speaker-knowledge extraction that utilizes a given enrollment (i.e., $S_j \neq S$) with the target speaker known a priori. The enrollment was not restricted to be a mixture with another random interfering speaker (with no prior knowledge) or a clean utterance. Since the speaker embedding could be pre-computed offline, this part of the inference time is saved. In Fig. 1 (c), we depict this property with different segment lengths and dotted lines connecting the generic-feature space and the speaker-knowledge space. Unlike traditional SV systems, we can extract reliable speaker embedding from not only clean utterances but also challenging overlapped mixtures. This is significant merit particularly important to realistic industrial deployment.

Until now, it is free to collect speaker enrollments from unconstrained various environments. Our proposed bottom-up query approach makes the model only attend to the most relevant parts of variable-length sequences of either online signals (analog to recent short-term memory) or pre-collected enrollments (analog to long-term persistent memory and experience). The top-down query then integrates the retrieved information to iteratively compute the speech representation through the stack of GALR blocks in the speech-stimuli space. All the above attentive operations are directly inferred from the input data without resorting to supervision.

In both “online” and “offline” modes, the model components in the speaker-knowledge space and the speech-stimuli space are trained cooperatively by casting information towards each other via the cross- and dual-attention mechanism. Our joint training method is very different from existing approaches that generally fuse the speaker embeddings and speech features into the same vector space through linear combinations, multiplication, or concatenation (Ji et al. 2020; Perez et al. 2017; Zeghidour and Grangier 2020). Specifically, we facilitate the communication of the two spaces only through soft-attention mechanisms. Consequently, the interaction between the two spaces is mediated through probability distributions only. The casting of speech features onto speaker embedding restricts the space of the valid top-down steering operations by anchoring the steering vectors \mathbf{Z}_j back in the original speaker-knowledge space. This serves as a form of regularization that restricts \mathbf{Z}_j to be a compact vector in the unique speaker-knowledge space, which enables the self-supervised techniques as discussed in Sec. .

Moreover, our network achieves better transparency and interpretability than the existing approaches. Unlike the black-box network in prior works, we can easily interpret the cross attention and the consequent top-down steering process. In Appendix 1.5, we illustrate an intriguing phenomenon that selective bottom-up cross attention similar to humans’ behavior could be automatically learnt by our network.

External Speaker Augmentation We find that data augmentation with external speaker datasets can enhance the generalizability of learnt speaker embedding over unseen speakers during a test. It is particularly useful when the number of speakers is small in the original training dataset. For

example, we used Librispeech for the speaker augmentation for WSJ0-2mix in Table ???. Our framework is flexible to utilize external speaker datasets to continue training the blocks in the speaker-knowledge space, where the forward and backward path goes only along the brick-red lines as shown in Fig. 1, and only the speaker loss $\mathcal{L}_{\text{InCE}} + \mathcal{L}_{\text{reg}}$ is used. After that, the blocks in the speech-stimuli space are fine-tuned using the original training set with the extracted \mathbf{Z}_j by the more generalized speaker-knowledge component. Table ?? provides an ablation study regarding this technique.

Evaluation and Analysis

Datasets and Model Setup

We describe the datasets and model setup in brief as below. Additional information needed to reproduce our experiments is stated in detail in Appendix 1.2.

Datasets We used a benchmark 8kHz dataset WSJ0-2mix (Hershey et al. 2016) for comparison with state-of-the-art source separation systems. Moreover, to evaluate performance on separating speech from music interference, we created another dataset WSJ0-music by mixing the speech corpus of WSJ0-2mix with music clips. We also used a large-scale publicly available 16kHz benchmark dataset Librispeech (Panayotov et al. 2015).

During training, each utterance is online mixed (masked) with an interfering utterance from another randomly chosen speaker. For testing, we pre-generated a fixed set of interfered utterances using the same SNR conditions as in training.

Model Setup The encoder and decoder structure, as well as the model’s hyper-parameter settings, were directly inherited from DPRNN’s setup (Luo, Chen, and Yoshioka 2019) for comparison purposes. Note that no model hyper-parameter has been fine-tuned towards our proposed structure, otherwise more improvement could be reasonably expected for ours. Tune-In loss hyper-parameters were set empirically.

Results and Discussion

Speaker Verification Performance This experiment investigated a downstream SV task to verify the discriminative and generalization property of our learnt representations. Note that we did not aim to achieve a SOTA SV score, but to examine the effectiveness of representation learning. An advanced SincNet-based SV model (Ravanelli and Bengio 2018) was considered a qualified reference system, for its reproducible high performance and speaker-embedding based network.³ One of the reference models, namely “SincNet-clean”, was trained on the clean Librispeech training set like conventional SV settings; the other reference model, namely “SincNet-masked”, was trained on the online-masked (mixed) Librispeech training set same as our proposed system. Fig. 2 plots the true positive rate (TPR) against the false positive rate (FPR), and gives the Receiving Operating Characteristic (ROC) curves by different models. The blue line indicates the standard performance by the SincNet SV system on clean test data (“SincNet-clean_test_on_clean”)

³<https://github.com/mravanelli/SincNet>.

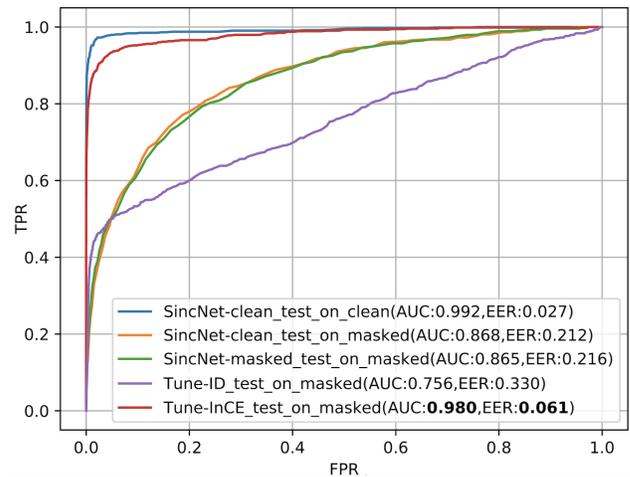


Figure 2: ROC curves of speaker verification by different models.

and can serve as an upper-bar reference. The yellow and green lines indicate that SincNet, either trained with clean (“SincNet-clean_test_on_masked”) or masked data (“SincNet-masked_test_on_masked”), can not handle masked test data well. The red line denoting the “Tune-InCE_test_on_masked” was by our proposed self-supervised model on the corrupted test data in the adverse conditions with 0-5dB SIR. Surprisingly, it achieves an AUC as high as 0.980 and an EER of 0.061, and its ROC line is approaching the “SincNet-clean_test_on_clean”, i.e., the performance on clean data by the reference SV system.

We also trained our model using a fully supervised approach of learnable token embedding, which has been widely applied in both NLP (Devlin et al. 2019) and speech domain (Shi et al. 2018; Zeghidour and Grangier 2020). The learnable embedding (Shi et al. 2018) maintained the target embedding of the inventory of the N speaker IDs. We call the fully supervised model “Tune-ID”, where the token embeddings for speaker IDs replaced our EMA based speaker embedding \mathbf{E}_{i_j} in “Tune-InCE”, and the conventional token embedding loss (Shi et al. 2018), replaced our proposed self-supervised Tune-InCE loss. All other training conditions for “Tune-ID”, such as Gaussian noise, dropout, online mixing strategy, etc., are the same as “Tune-InCE”. We can observe that the fully supervised model “Tune-ID_test_on_masked” performs poorly in generalization. To analyze the above performances, in Appendix 1.6, we study the generalization ability of the deep representations learnt by our proposed self-supervised approach and compare it to the fully supervised approach. The experimental result suggests that SincNets, even the one trained with the masked, are fragile when the speech inputs are not clean enough. In contrast, the self-supervised learning prevented our model from learning a trivial task of only predicting speaker identity, but enforced the model to learn deep representations with essential discriminative power.

Speech Separation Performance We then evaluated our proposed self-supervised Tune-In model on a speech separa-

Dataset	Model	Parameters	Memory	GFLOPs	SI-SNRi	SDRi
WSJ0-2mix	TDAA (Shi et al. 2018)	14.8M	-	-	-	12.6
	Conv-TasNet	8.8M	-	-	15.5	15.9
	‡DPRNN	2.6M	1,970MB	84.6	18.8	19.1
	Tune-In Autopilot	2.3M	730MB	28.4	20.3	20.5
	Tune-In Offline	2.4M	799MB	28.9	20.4	20.6
	Tune-In Online	3.2M	1,309MB	37.3	20.8[†]	21.0[†]
WSJ0-music	‡DPRNN	2.6M	231MB	10.7	14.5	14.8
	Tune-In Autopilot	2.3M	186MB	8.3	15.9	16.2
	Tune-In Online	3.2M	332MB	10.7	15.9	16.2
LibriSpeech	‡DPRNN	2.6M	1,152MB	40.5	12.0	12.5
	Tune-In Autopilot	2.3M	1,013MB	31.1	12.2	12.7
	Tune-In Offline	2.4M	1,112MB	31.7	12.6	13.1
	Tune-In Online	3.2M	1,912MB	39.9	13.0	13.6

Table 1: Performance comparison with conventional separation systems. Models with ‡ were reproduced in our own implementation for performance and runtime costs evaluation. Memory and GFLOPs were calculated on a 1s audio input.

tion task. As shown in Table ??, all systems were assessed in terms of SI-SNRi and SDRi. The ‘‘Tune-In Online’’ model inferred the speaker embeddings and steering vectors online and performed speech separation simultaneously; thus the additional task brought the extra complexity and around 0.8M model parameter size more than that of ‘‘Tune-In Autopilot’’ and ‘‘Tune-In Offline’’. The results with [†] indicate speaker augmentation was applied during training (see Table ?? for ablation study). As discussed in Sec., ‘‘Tune-In Offline’’ uses enrollments from target speakers. These enrollments were 16s long and collected from challenging conditions as online, i.e., the same SIR range but with different interferences.

Our Tune-In model outperforms the DPRNN model by a very large margin in terms of SI-SNRi and SDRi, which is achieved at a lower cost of up to 62.9% less run-time memory and 66.4% fewer computational operations. Even compared to the most recent SOTA model (DPGRNNs) (Nachmani, Adi, and Wolf 2020) with parameters as large as 7.5M, our Tune-In model still can maintain a significant advantage of 0.7 in terms of SI-SNRi and SDRi. For WSJ0-music, we used a window length of 16 samples instead of 4 samples to speed up the training for both Tune-In and DPRNN systems. By comparing results on WSJ0-music by Tune-In Autopilot v.s. Online, we noticed that the speaker embedding was not as beneficial for the speech-music separation task. This is reasonable because speech-music separation tends to rely on other key factors, e.g., rhythm, spectral continuity, pitch, etc., rather than speaker embedding. Nevertheless, results on WSJ0-music demonstrate the consistent advantage of our proposed GALR structure over DPRNN by a large margin.

Ablation Study In Table ??, we presented the test SI-SNRi for variants of our best-performing Tune-In Online system. Each variant removes a technique that is proposed for presumably better speaker embedding learning. The first two ablations experimented regarding the two different regularization techniques proposed after Eq. (2). When the dropout replaces the element-wise Gaussian noise, we observed a notable degradation. This empirically suggests that adding Gaussian noise rather than dropout is effective for building

a more robust speaker-knowledge-steered separation system. The removal of speaker augmentation also degrades the performance. Meanwhile, we found that the removal of regularization loss would bring down SI-SNRi. The ablation results together suggest that the combination of the proposed techniques is pivotal in achieving state-of-the-art performance.

Conclusions

‘‘Tune-In’’ attempts to practice the cocktail party effect in a new attentive way. We find it is critical to learn and maintain a speaker-knowledge and a speech-stimuli space separately. A novel cross- and dual-attention mechanism is proposed for information exchange between the two spaces, mimicking the bottom-up and top-down processes of a human neural system. The GALR structure as the building block in Tune-In breaks the memory and computation bottleneck of conventional self-attention layers. We reveal that substantially discriminative and generalizable representations can be learnt in severely interfered conditions via our self-supervised training. Tune-In outperforms a standard SV system and meanwhile achieves SOTA SS performances consistently in all test modes, particularly at a lower cost. Future work includes 1) a multi-channel extension as human’s cocktail party effect works best as a binaural effect, 2) the effect revealed by prior neurobiology study about an acceleration in attention switching and context updating, and 3) Tune-In for speech recognition by replacing the speech-stimuli space with a speech-phoneme space.

Model	SI-SNRi
Tune-In Online	20.8
w/ Embedding Dropout	19.8
w/o Embedding Noise	20.0
w/o Speaker Augmentation	20.6
w/o Regularization Loss	20.2

Table 2: Model ablations for our Tune-In Online model reporting results over the test set on WSJ0-2mix.

References

- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer Normalization. *arXiv preprint arXiv:1607.06450* .
- Bai, S.; Kolter, J. Z.; and Koltun, V. 2018. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. *arXiv preprint arXiv:1803.01271* .
- Chen, J.; Mao, Q.; and Liu, D. 2020. Dual-Path Transformer Network: Direct Context-Aware Modeling for End-to-End Monaural Speech Separation. *arXiv preprint arXiv:2007.13975* .
- Chen, Z.; Luo, Y.; and Mesgarani, N. 2017. Deep Attractor Network for Single-microphone Speaker Separation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 246–250. IEEE.
- Cherry, E. C. 1953. Some Experiments on The Recognition of Speech, with One and with Two Ears. *The Journal of the acoustical society of America* 25(5): 975–979.
- Chorowski, J.; Weiss, R. J.; Bengio, S.; and Oor, A. 2019. Un-supervised Speech Representation Learning Using Wavenet Autoencoders. In *IEEE TASLP*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, 4171–4186.
- Drude, L.; von Neumann, T.; and Haeb-Umbach, R. 2018. Deep Attractor Networks for Speaker Re-identification and Blind Source Separation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE.
- Evans, S.; McGettigan, C.; Agnew, Z. K.; Rosen, S.; and Scott, S. K. 2015. Getting the Cocktail Party Started: Masking Effects in Speech Perception. *Journal of Cognitive Neuroscience* 28(3): 483–500.
- Fujita, Y.; Kanda, N.; Horiguchi, S.; Nagamatsu, K.; and Watanabe, S. 2019. End-to-end Neural Speaker Diarization with Permutation-free Objectives. In *Proc. INTERSPEECH*.
- Getzmann, S.; Naatanen, R. 2015. The Mismatch Negativity as a Measure of Auditory Stream Segregation in a Simulated "Cocktail-party" Scenario: Effect of Age. *Neurobiology of Age* .
- Getzmann, S.; Jasny, J.; and Falkenstein, M. 2016. Switching of Auditory Attention in "Cocktail-party" Listening: ERP Evidence of Cueing Effects in Younger and Older Adults. *Brain and Cognition* 111: 1–12.
- Hershey, J. R.; Chen, Z.; Le Roux, J.; and Watanabe, S. 2016. Deep Clustering: Discriminative Embeddings for Segmentation and Separation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 31–35. IEEE.
- Huang, Z.; Watanabe, S.; Fujita, Y.; Garcia, P.; Shao, Y.; Povey, D.; and Khudanpur, S. 2020. Speaker Diarization with Region Proposal Network. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- Hudson, D. A.; and Manning, C. D. 2018. Compositional Attention Networks for Machine Reasoning. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Ji, X.; Yu, M.; Zhang, C.; Su, D.; Yu, T.; Liu, X.; and Yu, D. 2020. Speaker-Aware Target Speaker Enhancement by Jointly Learning with Speaker Embedding Extraction. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7294–7298. IEEE.
- Khandelwal, U.; He, H.; Qi, P.; and Jurafsky, D. 2018. Sharp Nearby, Fuzzy Far Away: How Neural Language Models Use Context. *arXiv preprint arXiv:1805.04623* .
- Koizumi, Y.; Yatabe, K.; Delcroix, M.; Masuyama, Y.; and Takeuchi, D. 2020. Speech Enhancement Using Self-adaptation and Multi-head Self-attention. *arXiv preprint arXiv:2002.05873* .
- Lam, M. W.; Wang, J.; Su, D.; and Yu, D. 2020. Mixup-breakdown: a Consistency Training Method for Improving Generalization of Speech Separation Models. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* .
- Le Roux, J.; Wisdom, S.; Erdogan, H.; and Hershey, J. R. 2019. SDR– Half-baked or Well Done? In *2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 626–630. IEEE.
- Lim, S. J.; Wostmann, M.; and Obleser, J. 2015. Selective Attention to Auditory Memory Neurally Enhances Perceptual Precision. *Journal of Neuroscience* 35(49): 16094–16104.
- Luo, Y.; Chen, Z.; and Yoshioka, T. 2019. Dual-path RNN: Efficient Long Sequence Modeling for Time-domain Single-channel Speech Separation. *arXiv preprint arXiv:1910.06379* .
- Luo, Y.; and Mesgarani, N. 2018. Tasnet: Time-domain Audio Separation Network for Real-time, Single-channel Speech Separation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 696–700. IEEE.
- Luo, Y.; and Mesgarani, N. 2019. Conv-tasnet: Surpassing Ideal Time-frequency Magnitude Masking for Speech Separation. *IEEE/ACM transactions on audio, speech, and language processing* 27(8): 1256–1266.
- Mirco Ravanelli, Y. B. 2018. Speaker Recognition from Raw Waveform with SincNet. In *Proceedings of SLT*.
- Misra, I.; Shrivastava, A.; Gupta, A.; and Hebert, M. 2016. Cross-stitch Networks for Multi-task Learning. *CVPR* .
- Nachmani, E.; Adi, Y.; and Wolf, L. 2020. Voice Separation with an Unknown Number of Multiple Speakers. *ICML* .
- Narayan, R.; Best, V.; Ozmeral, E.; McClaine, E.; Dent, M.; Shinn-Cunningham, B.; and Sen, K. 2009. Cortical Interference Effects in the Cocktail Party Problem. *Nature Neuroscience* 10(12): 1601–1607.
- Oord, A.; Li, Y.; and Vinyals, O. 2018. Representation Learning with Contrastive Predictive Coding. *arXiv preprint arXiv:1807.03748* .

- Panayotov, V.; Chen, G.; Povey, D.; and Khudanpur, S. 2015. Librispeech: an ASR Corpus Based on Public Domain Audio Books. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5206–5210. IEEE.
- Pariante, M.; Cornell, S.; Cosentino, J.; Sivasankaran, S.; Tzinis, E.; Heitkaemper, J.; Olvera, M.; Stoter, F.; Hu, M.; Donas, J. M. M.; Ditter, D.; Frank, A.; Deleforge, A.; and Vincent, E. 2020. Asteroid: the PyTorch-based Audio Source Separation Toolkit for Researchers. *arXiv preprint arXiv:2005.04132* .
- Pascual, S.; Ravanelli, M.; Serra, J.; Bonafonte, A.; and Bengio, Y. 2019. Learning Problem-agnostic Speech Representations from Multiple Self-supervised Tasks. In *Proc. INTERSPEECH*.
- Perez, E.; Strub, F.; Vries, H.; Dumoulin, V.; and Courville, A. 2017. FiLM: Visual Reasoning with a General Conditioning Layer. In *arXiv preprint arXiv:1709.07871*.
- Ravanelli, M.; and Bengio, Y. 2018. Speaker Recognition from Raw Waveform with SincNet. In *In Proceedings of SLT 2018*.
- Ravanelli, M.; and Bengio, Y. 2019. Learning Speaker Representations with Mutual Information. In *Proc. INTERSPEECH*.
- Ravanelli, M.; Brakel, P.; Omologo, M.; and Bengio, Y. 2018. Light Gated Recurrent Units for Speech Recognition. *IEEE Transactions on Emerging Topics in Computational Intelligence* 2(2): 92–102.
- Ravanelli, M.; Zhong, J.; Pascual, S.; Swietojanski, P.; Monteiro, J.; Trmal, J.; and Bengio, Y. 2020. Multi-task Self-supervised Learning for Robust Speech Recognition. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- Shi, J.; Xu, J.; Liu, G.; and Xu, B. 2018. Listen, Think and Listen Again: Capturing Top-down Auditory Attention for Speaker-independent Speech Separation. *Proc. IJCAI* .
- Shi, J.; Xu, J.; and Xu, B. 2019. Which Ones Are Speaking? Speaker-inferred Model for Multi-talker Speech Separation. *Proc. INTERSPEECH* .
- Shi, Z.; Liu, R.; and Han, J. 2020. Speech Separation Based on Multi-Stage Elaborated Dual-Path Deep BiLSTM with Auxiliary Identity Loss. *INTERSPEECH* .
- Snyder, D.; Garcia-Romero, D.; Sell, G.; Povey, D.; and Khudanpur, S. 2018. X-vectors: Robust DNN Embeddings for Speaker Recognition. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5329–5333.
- Sperber, M.; Niehues, J.; Neubig, G.; Stuker, S.; and Waibel, A. 2018. Self-Attentional Acoustic Models. *Proc. INTERSPEECH* .
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention Is All You Need. In *Advances in neural information processing systems*, 5998–6008.
- Wang, J.; Chen, J.; Su, D.; Chen, L.; Yu, M.; Qian, Y.; and Yu, D. 2018. Deep Extractor Network for Target Speaker Recovery from Single Channel Speech Mixtures. *arXiv preprint arXiv:1807.08974* .
- Wang, Q.; Muckenhirn, H.; Wilson, K.; Sridhar, P.; Wu, Z.; Hershey, J.; Saurous, R. A.; Weiss, R. J.; Jia, Y.; and Moreno, I. L. 2019. VoiceFilter: Targeted Voice Separation by Speaker-Conditioned Spectrogram Masking. *Proc. INTERSPEECH* .
- Xu, J.; Shi, J.; Liu, G.; Chen, X.; and Xu, B. 2018. Modeling Attention and Memory for Auditory Selection in a Cocktail Party Environment. *Proc. AAAI* .
- Yi Liu, Liang He, J. L. 2019. Large Margin Softmax Loss for Speaker Verification. In *Proc. INTERSPEECH*.
- Yu, D.; Kolbæk, M.; Tan, Z.-H.; and Jensen, J. 2017. Permutation Invariant Training of Deep models for Speaker-independent Multi-talker Speech Separation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 241–245. IEEE.
- Zadeh, A.; Ma, T.; Poria, S.; and Morency, P. 2019. Wild-Mix Dataset and Spectro-Temporal Transformer Model for Monaural Audio Source Separation. *arXiv preprint arXiv:1911.09783* .
- Zeghidour, N.; and Grangier, D. 2020. Wavesplit: End-to-end Speech Separation by Speaker Clustering. In *arXiv preprint arXiv:2002.08933*.