

USEV: Universal Speaker Extraction With Visual Cue

Zexu Pan , Graduate Student Member, IEEE, Meng Ge , and Haizhou Li , Fellow, IEEE

I. INTRODUCTION

Abstract— A speaker extraction algorithm seeks to extract the target speaker’s speech from a multi-talker speech mixture. The prior studies focus mostly on speaker extraction from a highly overlapped multi-talker speech mixture. However, the target-interference speaker overlapping ratios could vary over a wide range from 0% to 100% in natural speech communication, furthermore, the target speaker could be absent in the speech mixture, the speech mixtures in such universal multi-talker scenarios are described as *general speech mixtures*. The speaker extraction algorithm requires an auxiliary reference, such as a video recording or a pre-recorded speech, to form top-down auditory attention on the target speaker. We advocate that a visual cue, i.e., lip movement, is more informative than an audio cue, i.e., pre-recorded speech, to serve as the auxiliary reference for speaker extraction in disentangling the target speaker from a *general speech mixture*. In this paper, we propose a universal speaker extraction network with a visual cue, that works for all multi-talker scenarios. In addition, we propose a scenario-aware differentiated loss function for network training, to balance the network performance over different target-interference speaker pairing scenarios. The experimental results show that our proposed method outperforms various competitive baselines for *general speech mixtures* in terms of signal fidelity.

Index Terms—*General speech mixture, multi-modal, scenario-aware differentiated loss, sparsely overlapped speech, target speaker extraction.*

Manuscript received 7 March 2022; revised 30 June 2022 and 2 August 2022; accepted 30 August 2022. Date of publication 12 September 2022; date of current version 30 September 2022. This work was supported in part by the Internal Project of Shenzhen Research Institute of Big Data under the Grant T00120220002, in part by the Guangdong Provincial Key Laboratory of Big Data Computing under the Grant B10120210117-KP02, in part by University Development Fund under Grants UDF01002333 and UF02002333, in part by The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen), in part by the Science and Engineering Research Council, Agency for Science, Technology and Research (A*STAR), Singapore, through the National Robotics Program under Human-Robot Interaction Phase 1 under Grant 192 25 00054, and in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy (University Allowance, EXC 2077, University of Bremen). The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Christine Evers. (Corresponding author: Meng Ge).

Zexu Pan is with the Integrative Sciences and Engineering Programme, Institute of Data Science, Department of Electrical and Computer Engineering, National University of Singapore, Singapore 119077 (e-mail: e0474123@u.nus.edu).

Meng Ge is with the Tianjin Key Laboratory of Cognitive Computing and Application, College of Intelligence and Computing, Tianjin University, Tianjin 300072, China, and also with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore 119077 (e-mail: gemeng@tju.edu.cn).

Haizhou Li is with the School of Data Science, Chinese University of Hong Kong, Shenzhen 518172, China, also with the University of Bremen, 28359, Germany, and also with the Kriston AI, Xiamen 361021, China (e-mail: haizhou.li@u.nus.edu).

Digital Object Identifier 10.1109/TASLP.2022.3205759

AT A cocktail party, human has the inherent ability to selectively listen to a speaker in the presence of interference speakers and background noises, this is also known as selective auditory attention [1], [2], [3], [4], [5]. Decades of efforts have been spent to emulate such attention with engineering solutions, which is also referred to as the *cocktail party* problem [6]. It is non-trivial but highly demanded in real-world applications such as hearing aids [7], active speaker detection and verification [8], [9], speaker localization [10], and automatic speech recognition [11].

Speech separation represents one way to address the *cocktail party* problem, which seeks to separate a multi-talker speech mixture into individual streams, each for one speaker. The traditional approaches, e.g., non-negative matrix factorization [12], [13], factorial hidden Markov models and Gaussian mixture models [14], and computational auditory scene analysis [15], [16], [17], explore the idea of spectro-temporal masking based on harmonic and pitch analysis to filter a speaker out from a speech mixture. The prior studies have laid the foundation for recent progress. With the advent of deep learning, speech separation has seen major progress [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28], even with different speaker overlapping ratios [29], [30], [31], [32], [33], [34], [35], [36], [37]. In a neural architecture, multiple speaker streams compete and segregate either with a masking or regression mechanism. As the number of speakers in the speech mixture is typically required in advance, speech separation has a limited scope of applications.

Speaker extraction is different from speech separation. It only extracts the speech of a target speaker from a multi-talker speech mixture, thus invariant to the number of speakers. Speaker extraction emulates a human’s selective auditory attention at a *cocktail party*, which typically requires an auxiliary reference, that provides the information of the target speaker [2], [38], [39], [40], [41], [42], [43]. Therefore, the quality of the auxiliary reference matters. A pre-recorded speech from the target speaker has been well studied to serve as such an auxiliary reference [2], [3], [4], [44], [45], [46], [47], [48], [49], the speaker extraction algorithm is expected to extract only the speech that has a similar voice signature to the pre-recorded speech.

The state-of-the-art speaker extraction algorithms perform remarkably well on highly overlapped speech, in which the speakers overlap almost 100%, e.g., WSJ0-2mix-extr dataset [2], [50]. However, in natural speech communication, a target speaker may speak and pause intermittently, interspersed with the speech of interference speakers, resulting in sparsely overlapped speech,

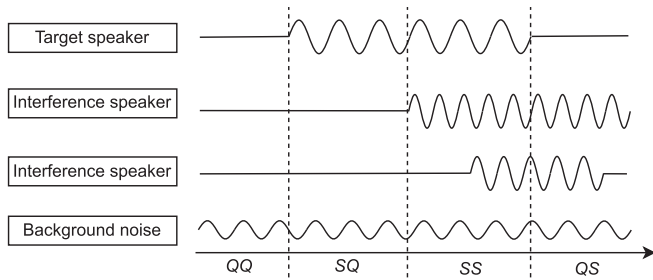


Fig. 1. An illustration of the four target-interference speaker pairing scenarios, namely QQ , SQ , SS , QS , that can occur within a *general speech mixture* clip.

i.e., the speaker overlapping ratio is only around 20% in meetings [51], and several datasets have been recorded to simulate such natural conversational speech with various speaker overlapping ratios, e.g., CHiME [52] and LibriCSS [53] datasets. Speaker extraction algorithms for such sparsely overlapped speech have less been studied.

In the task of speaker extraction, the target speaker can be absent or present in a speech mixture [54], [55]. Target absent (*TA*) refers to the scenario where the target speaker is quiet throughout the speech mixture clip. Target present (*TP*) refers to the scenario where the target speaker is present in the speech mixture clip, and the target-interference speaker overlapping ratio could range from 0% to 100%. Speech mixtures at a varying speaker overlapping ratio, with either absent or present target speaker, are referred to as *general speech mixtures* in this paper.

We argue that the *TA* and *TP* categorization is too broad to be technically meaningful as far as speaker extraction is concerned, and advocate four possible target-interference speaker pairing scenarios, that could occur in a *general speech mixture*, namely QQ , SQ , SS , QS , in Fig. 1. QQ denotes the scenario of *quiet* target speaker with *quiet* interference speakers; SQ denotes that of *speaking* target speaker with *quiet* interference speakers; SS denotes that of *speaking* target speaker with *speaking* interference speakers; QS denotes that of *quiet* target speaker with *speaking* interference speakers. A *general speech mixture* clip in practice could be a random combination of any of the four scenarios, where SS represents the scenario of highly overlapped speech in many prior studies [3], [4], [44], [50], while QQ , SQ , QS , and the scenario transition in between are not well studied yet. In Fig. 2, we illustrate some examples of the *TA* and *TP* speech clips. It is worth noting that the *TA* speech clips can be in either QQ or QS scenarios, or both. But QQ and QS scenarios can also possibly take place in the *TP* speech clips.

A speaker extraction algorithm trained on highly overlapped speech may falsely generate undesired output when the target speaker is quiet [54], which is called the ‘false-extraction’ problem. By simply fine-tuning a system on speech mixture at a variety of overlapping ratios, one may address the ‘false-extraction’ problem to some extent. However, speech mixtures of different overlapping ratios have different training expectations and difficulties for each target-interference speaker pairing scenario, thus

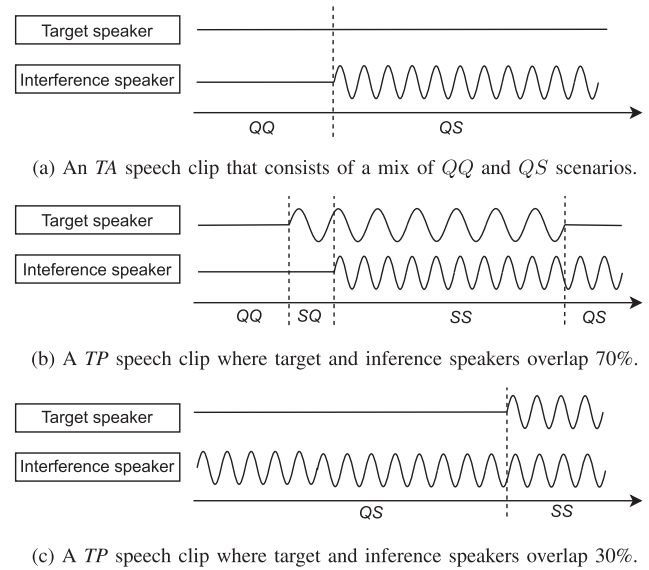


Fig. 2. Illustration of some speech mixture clips. (a) is an *TA* example. (b) and (c) are two *TP* examples.

a differentiated training strategy is needed. For example, for a highly overlapped speech mixture with a talkative target speaker, the algorithm just needs to focus on speech disentanglement between speakers, otherwise, the algorithm needs to do well in verifying whether the target speaker is active as well.

Humans perceive the world through various simultaneous sensory systems [56]. We attend to a speaker not only by reference to a registered voice signature but also through other means such as observing the lip movement [57] or understanding the contextual relevance [58]. Neuroscience studies suggest that hearing is improved by observing the lip movement in a conversation [59], [60], leveraging on the temporal correlation between visemes and speech [61]. A viseme is a basic unit of visual speech, which corresponds to a set of phonemes for acoustic speech [62], [63].

Although visemes and phonemes do not have a one-to-one correspondence, visemes do provide a fine-grained cue about underlying phonetic units being spoken [64]. Therefore, visemes have been widely adopted as the auxiliary reference for speaker extraction, which performs exceptionally well to disentangle the target speaker from high overlapped speech [38], [65], [66], [67]. Visemes also provide a high-level cue that discriminates between speech and non-speech signals [64], thus possibly alleviating the ‘false-extraction’ problem. We advocate that visemes are more informative than a pre-recorded speech to discriminate the speaking status of the target speaker at frame-level for a *general speech mixture*, thus serving as a robust auxiliary reference for speaker extraction algorithms. Next, we summarize our contributions of this work.

- 1) We address a unique research problem, i.e., speaker extraction from a *general speech mixture* with visual auxiliary reference, to deal with all multi-talker scenarios with one universal solution. This study brings audio-visual speaker extraction a step closer to solving the *cocktail party* problem in real-world applications.

- 2) We propose to categorize a *general speech mixture* into four different target-interference speaker pairing scenarios, and devise a scenario-aware differentiated loss function, to moderate the training of the four scenarios. We confirm the effectiveness through experiments on various speaker extraction models.
- 3) We evaluate the state-of-the-art speaker extraction algorithms on *general speech mixtures* and confirm the effectiveness of visual reference over speech reference. We also study the use of the dual-path recurrent neural network architecture in the speaker extractor, instead of the dilated temporal convolutional neural network, to capture global temporal dependencies in speech.

The rest of the paper is organized as follows. In Section II, we discuss related work. In Section III, we formulate the proposed method. In Section IV, we present the experimental setup. In Section V, we report the experimental results. Finally, we conclude the study in Section VI.

II. RELATED WORK

The speaker extraction algorithm usually requires an auxiliary reference to form the top-down attention on the target speaker. Either an auditory or a visual cue may serve as a such auxiliary reference. Next, we discuss the prior studies for three related tasks to set the stage for our study.

A. Highly Overlapped Speech Mixture With Audio Cue

Each speaker has a unique voice signature, which can be characterized by a speaker embedding, such as i-vector [68], x-vector [69], and d-vector [70]. Given a speech sample, the speaker embedding can be derived by a speaker encoder that is trained with a speaker recognition task [71].

VoiceFilter [4] employs the d-vector encoded from a pre-recorded speech, that serves as an audio cue, to form the top-down attention for speaker extraction. The study explores the discriminative power of the d-vector to find the speech track in the multi-talker speech that has a similar voice signature to the d-vector, and extracts that speech track. The idea of speaker embedding is extended to the context of speaker extraction, where the speaker encoder is jointly trained with the speaker extractor network in a multi-task learning framework [2], [44], [72]. The speaker extractor network is trained by a signal reconstruction loss while the speaker encoder is trained with a speaker recognition loss. In this way, the resulting speaker embedding is not only discriminative between speakers, but also optimized for speaker extraction.

These prior studies assumed that speakers highly overlap in the multi-talker speech, and the target speaker is always present. The case of highly overlapped speech puts the speaker extraction algorithm in a stress test. However, it only represents one of many possible multi-talker scenarios. There are other scenarios, for example, the speaker of interest is absent, or the speaker of interest is the only speaker in the speech. To work in real-world applications, a speaker extraction model needs to perform across all scenarios.

B. Absent Target Speaker Speech Mixture With Audio Cue

In natural speech communication, interlocutors typically take turns to speak, resulting in speech mixtures of varying extents of overlapping. In particular, a speaker of interest could remain quiet throughout while some others are speaking, resulting in *TA* mixture clips. There have been studies to include *TA* clips into the training data to overcome ‘false-extraction’ [54], [55]. It is noted that the commonly used objective function scale-invariant signal-to-noise ratio (SI-SDR) [73] or signal-to-noise ratio (SDR) is undefined, or defined as a constant, for *TA* clips. It was proposed [30], [54], [55], [74] to minimize the energy of *TA* clips instead, at the same time maximizing the SI-SDR or SDR for *TP* clips. We refer to such a loss as the scenario-aware uniform loss in this paper, it is scenario-aware because it differentiates *TA* with *TP* clips, it is uniform as it applies a single loss function to each of an entire clip. These prior studies [54], [55] focus on the case in which the target speaker is either absent or highly overlapped with the interference speaker in a speech clip. We consider the more realistic *general speech mixtures* in this paper as illustrated in Fig. 1, where sparsely overlapped speech is included.

A recent speech separation study [31] proposed the source-aggregated SDR, which computes the mean over the SDR values of each output channel, this works as long as the target speaker is not absent in at least one output channel. Other studies [36], [37] used the mean squared error loss on the estimated mask or spectrogram to avoid the absent target speaker problem, while it remains unclear whether this is better as it is not directly optimizing the signal quality of the extracted speech waveform. In another study of the VAD-SE network [75], a target speaker voice activity detection (VAD) module and a speaker extraction network [75] are jointly trained. During run-time inference, the VAD mutes the speaker extraction output when the target speaker is quiet. In other words, the speaker extraction performance highly depends on the VAD accuracy.

Both the scenario-aware uniform loss [30], [54], [55], [74] and the VAD-SE [75] network use an audio cue, i.e., a pre-recorded speech, as the auxiliary reference. However, for a *general speech mixture*, the audio cue is not only used for disentanglement but also used to verify whether the target speaker is speaking or quiet for every frame, which is difficult. In this paper, instead of an audio cue, we explore the use of a visual cue, i.e., lip movement, that doesn’t require the pre-recording of speech for the speaker extraction task. We believe the visual cue is a more direct cue verifying the speaking status of the target speaker compared to the audio cue.

C. Highly Overlapped Speech Mixture With Visual Cue

In the *cocktail party* problem, the visual reference, if present, is not corrupted by either acoustic noise or interference speech, which can be employed to form the top-down attention on the target speaker. The visual reference usually takes the form of a single face image from the target speaker, or a sequence of lip images of the target speaker that is synchronized with the speech.

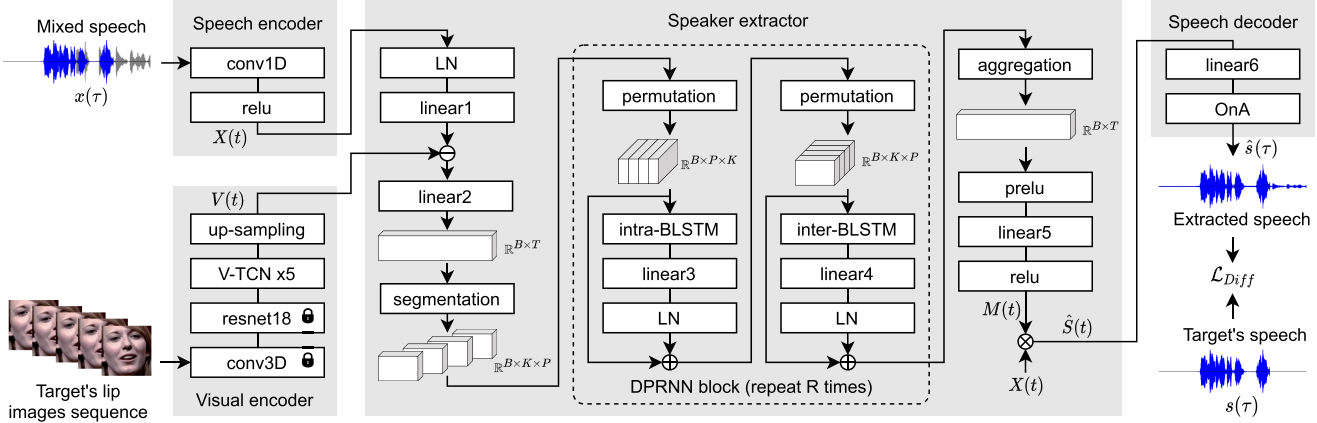


Fig. 3. The proposed universal speaker extraction network, or USEV network in short. It consists of a speech encoder, a visual encoder, a speaker extractor, and a speech decoder. The symbol \ominus refers to frame-wise concatenation of features; the symbols \oplus and \otimes refer to element-wise addition and multiplication of features. The network layers are represented by rectangles, the intermediate features are represented by 3D blocks. The lock sign in the visual encoder means that the weights of the network layers are frozen during the USEV network training.

FaceFilter [76] makes use of a single face image and explores the general relationship between the facial appearance and the voice signature, such as age, gender, and ethnicity. The network encodes the face image into a speaker embedding to form the attention on the target speaker.

Many audio-visual speaker extraction algorithms use the visemes encoded from a sequence of lip images of the target speaker as the visual cue. The conversation [57], and Time-domain speaker extraction network (TDSE) [66] are such examples, where they pre-train a visual encoder on a visual speech recognition task [77]. The visual encoder encodes a sequence of lip images, that are synchronized with the target’s speech, into a sequence of lip embeddings, also known as visemes. The lip embeddings, that are time-aligned with speech frames, are then used to model the temporal synchronization and interaction between the lip movement and speech.

The visemes are shown effective in disentangling a target speaker from highly overlapped speech, but have not been studied for the *general speech mixtures*. In this paper, we employ the visemes to form the top-down attention on the target speaker.

III. UNIVERSAL SPEAKER EXTRACTION NETWORK WITH VISUAL CUE

Let $x(\tau)$ be a multi-talker speech clip, consisting of the target speaker’s speech $s(\tau)$ and interference speaker’s speech $b_i(\tau)$,

$$x(\tau) = s(\tau) + \sum_{i=1}^I b_i(\tau), \quad (1)$$

where $i \in \{1, \dots, I\}$ denotes the index of interference speakers. The task of target speaker extraction aims to recover $\hat{s}(\tau)$ that is close to $s(\tau)$ from $x(\tau)$.

A. Network Architecture

We now formulate a Universal Speaker Extraction network with a Visual cue, named USEV (pronounced as ‘use v’) network, as depicted in Fig. 3. The network is universal because

it is expected to perform for a speech clip with all possible target-interference speaker pairing scenario combinations.

We adopt the time-domain approach for the network design, which is originally proposed in the time-domain speech separation network (Conv-TasNet) [78], to avoid the phase estimation problem arising from the frequency-domain approach. The USEV network has four components: 1) The speech encoder transforms the time-domain speech samples $x(\tau)$ into a sequence of speech embeddings $X(t)$. 2) The visual encoder encodes the target’s lip images sequence into a sequence of visual embeddings $V(t)$. 3) The speaker extractor estimates a mask $M(t)$, which only lets the target speaker pass in $X(t)$. 4) The speech decoder renders the masked speech embeddings $\hat{S}(t)$ into time-domain speech samples $\hat{s}(\tau)$.

1) *Speech Encoder*: The speech encoder consists of a 1-dimensional (1D) convolution $conv1D$ followed by a rectified linear activation $relu$. The speech encoder behaves like a frequency analyzer to convert the time-domain speech samples $x(\tau)$ into a spectrum-like frame-based embedding sequence $X(t)$ in the latent space,

$$X(t) = relu(conv1D(x(\tau), 1, N, L)) \in \mathbb{R}^{N \times T} \quad (2)$$

where the $conv1D$ has input channel size 1, output channel size N , kernel size L , and stride $L/2$. The output $X(t)$ is a T frame embedding sequence of dimension of N , where $t \in \{1, \dots, T\}$.

2) *Visual Encoder*: The visual encoder seeks to encode the target’s lip images sequence into a sequence of visual embeddings $V(t) \in \mathbb{R}^{N \times T}$, representing the target speaker’s visemes, and in sync with the target’s speech. We design the visual encoder with a structure similar to the visual encoder in MuSE [67], which consists of a 3-dimensional (3D) convolution $conv3D$, an 18 layer residual convolutional neural network $resnet18$, 5 repeated visual temporal convolutional network $V-TCN$, and an up-sampling layer.

The $conv3D$ and $resnet18$ are pre-trained from visual speech recognition task [77], that are denoted by a lock sign in Fig. 3. With the pre-trained weights fixed during speaker extraction training, we seek to retain the pre-trained knowledge to encode

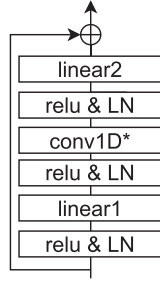


Fig. 4. The architecture of the V -TCN in the visual encoder. It consists of a linear layer $linear1$ with input and output sizes N and $N * 2$, respectively; a group convolution layer $conv1D^*$ with input channel size $N * 2$, output channel size $N * 2$, and group size $N * 2$, kernel size 3; and a linear layer $linear2$ with input and output sizes $N * 2$ and N , respectively. Each of the $linear1$, $linear2$, and $conv1D^*$ layers are preceded with a $relu$ and a layer normalization LN .

visemes that synchronize with the phonetic sequence of the target’s speech.

The visual features extracted from the pre-trained $conv3D$ and $resnet18$ are different from the speech embeddings [79], that are not optimized for speaker extraction directly. We design an adaptation network with 5 repeated V -TCN similar to the reentry model [5], with non-shared network weights, to adapt the visual embeddings towards the speaker extraction task. The architecture of a V -TCN is shown in Fig. 4. The 5 V -TCN is followed by an up-sampling layer to match the temporal resolution of the visual embeddings to the same as the speech embeddings $X(t)$.

3) *Speaker Extractor*: In computational auditory scene analysis [80], the selective filter is well studied to emulate the human’s selective auditory attention in attentive listening, which is usually modeled as a receptive mask. The speaker extractor in the USEV network adopts the masking method [2] to estimate a mask $M(t)$ which only lets the target speaker pass in the speech embeddings $X(t)$. The masked speech embeddings $\hat{S}(t)$ are obtained by element-wise multiplication between $X(t)$ and $M(t)$,

$$\hat{S}(t) = X(t) \otimes M(t) \in \mathbb{R}^{N \times T} \quad (3)$$

where $X(t)$ is the output of the speech encoder.

The speaker extractor requires a reference to form the top-down auditory attention on the target speaker. The visual embeddings $V(t)$ are trained just for that by providing the target speaker’s visemes, that are synchronized with the target’s speech. The inputs to the speaker extractor are the speech embeddings $X(t)$ and the visual embeddings $V(t)$. The studies on the *reentry* model [5], TDSE [66], and MuSE [67] suggest that, by concatenating the time-aligned visual embeddings with their corresponding speech embeddings, the speaker extractor is able to effectively estimate the mask $M(t)$. We adopt the concatenation approach at the start of the speaker extractor.

Before the concatenation operation, $X(t)$ is passed through a layer normalization LN , followed by a bottleneck linear layer $linear1$ [78] with input and output sizes N and B respectively.

The concatenated embeddings are passed through a linear projection layer $linear2$ with $B + N$ and B as the input and output sizes respectively.

In view that a speech utterance is usually encoded into a long sequence of speech embeddings. It is important to effectively model the long-term dependencies in the mask estimation process. Dilated temporal convolutional neural network (TCN) with a large receptive field has been widely used, such as WaveNet [81], Conv-TasNet [78], and TDSE [66]. However, the TCN has a fixed receptive field, and thus has difficulty in learning the long-term dependencies.

DPRNN [82] is a dual-path recurrent neural network for speech separation, which has a dynamic receptive field to capture global dependencies. DPRNN segments a long sequence of embeddings into short chunks and applies intra- and inter-chunk operations with an interlacing structure. In view of its success, we use the dual-path structure [82] in the context of audio-visual speaker extraction. which consists of a segmentation layer, R repeated DPRNN blocks, and an aggregation layer in the speaker extractor.

After the linear projection layer $linear2$, the 2-dimensional (2D) vector is passed through a segmentation layer, which is segmented into chunks with a window size K and hop size $K/2$. K is selected such that $K \approx \sqrt{2T}$ [82]. The resultant P chunks are concatenated to form a 3D tensor, where $P = \frac{2 \times T}{K} + 1$.

After segmentation, the 3D tensor is passed through R repeated DPRNN blocks with non-shared network weights, each with a chain of intra-chuck processing and inter-chuck processing with residual connections. The intra-chuck processing has an intra-bidirectional long short-term memory *intra-BLSTM*, with input and hidden size B and $B \times 2$ respectively, which is applied to the intra-chuck sequence (K dimension) of the 3D tensor. The *intra-BLSTM* is followed by a linear layer $linear3$ with input and output sizes $B \times 4$ and B respectively, and a LN . Similar to the intra-chuck processing, the inter-chuck has a similar network architecture and parameters, except that the *inter-BLSTM* is applied to the inter-chuck sequence (P dimension) of the 3D tensor.

The aggregation layer is an inverse operation of the segmentation layer, which transforms the 3D tensor back to 2D. It is followed by a parametric rectified linear activation *prelu*, a linear layer $linear5$ with input and output sizes B and N , and a $relu$ layer.

4) *Speech Decoder*: The speech decoder renders the masked speech embeddings $\hat{S}(t)$ into time-domain speech samples $\hat{s}(t)$. It consists of a linear layer $linear6$ and an overlap-and-add operation *OnA*,

$$\hat{s}(\tau) = OnA(Linear6(\hat{S}(t), N, L), L/2) \quad (4)$$

where the *OnA* operation has a frame shift of $L/2$.

B. Scenario-Aware Differentiated Loss

In time-domain source separation or extraction algorithms, SDR is widely used as the objective function [83], [84], [85], which performs very well in the case of highly overlapped sound sources. However, for *TA* mixture clips, the SDR is a constant

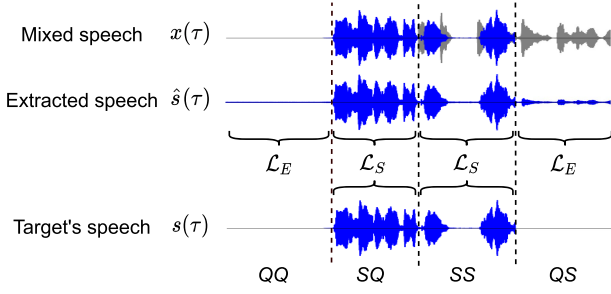


Fig. 5. An illustration of the scenario-aware differentiated loss. We segment an extracted speech clip according to the different target-interference speaker pairing scenarios (QQ , SQ , SS , QS) and apply separate loss functions for different segments.

value. The scenario-aware uniform loss studies [30], [54], [55], [74], which the loss function is shown in (5), point to a direction that, by adding an ϵ at the numerator, the SDR loss reduces to a form that minimizes the energy of the TA speech clips, at the same time, still maximizes the SDR for TP speech clips.

$$\mathcal{L}_{Uni} = -10 \log_{10} \left(\frac{\|s\|^2 + \epsilon}{\|\hat{s} - s\|^2 + \epsilon} \right) \quad (5)$$

where ϵ is a small value of $1e^{-8}$. We omit the subscript (τ) for the clean $s(\tau)$ and extracted speech $\hat{s}(\tau)$ in the formula for brevity.

In reality, a TP speech clip could be a combination of QQ , SQ , SS , QS scenario segments. Such a multi-scenario clip is called a heterogeneous clip, as shown in Figs. 1 and 2, while a single scenario clip is called a homogeneous clip. It is understood that we have different expectations for each of the scenarios, e.g., low energy silence outputs for QS and QQ scenarios, and high SDR outputs for SS and SQ ones. On the other hand, the difficulty of training for different scenarios is expected to be different, e.g., SS scenario is more complex than SQ , and QS is more complex than QQ . Therefore, it is best that we measure the output speech quality of each scenario differently. However, the loss functions, either SDR or energy loss, are segmental in nature. A single uniform loss reflects the average loss quantity indiscriminately over the entire heterogeneous clip, that doesn't seek to optimize for each scenario, which could result in unbalanced training between scenarios. We argue that a scenario-specific loss for each scenario-homogeneous segment should be more appropriate.

Furthering the study of scenario-aware uniform loss, we propose a training strategy with a scenario-aware differentiated loss to deal with the four scenarios as summarized in Fig. 5. During training, the speaker extraction network takes the entire speech clip as input and the scenario labels as the supervision signals. We segregate the extracted speech clip into the respective scenario segments according to the labels, i.e., QQ , SQ , SS , QS .

For the SQ and SS segments, where the target speaker is acoustically active, we apply the SDR as the loss \mathcal{L}_S between the ground-truth and the extracted speech,

$$\mathcal{L}_S = -10 \log_{10} \left(\frac{\|s\|^2}{\|\hat{s} - s\|^2 + \epsilon} + \epsilon \right) \quad (6)$$

For the QQ and QS segments where the target speaker is acoustically inactive, we apply the energy of the segments as the loss \mathcal{L}_E ,

$$\mathcal{L}_E = 10 \log_{10} (\|\hat{s}\|^2 + \epsilon) \quad (7)$$

We define the total scenario-aware differentiated loss \mathcal{L}_{Diff} as a weighted sum of the 4 differentiated loss values as follows,

$$\mathcal{L}_{Diff} = \alpha \mathcal{L}_E^{QQ} + \beta \mathcal{L}_S^{SQ} + \gamma \mathcal{L}_S^{SS} + \delta \mathcal{L}_E^{QS} \quad (8)$$

where $\alpha, \beta, \gamma, \delta$ are the weights for the loss, also called the hyperparameters.

With the scenario-aware differentiated loss \mathcal{L}_{Diff} , we are able to optimize the network with a scenario-specific loss for each scenario-homogeneous segment, yet having an overall objective across the entire heterogeneous clip, i.e., a random combination of different scenarios. The weights in \mathcal{L}_{Diff} moderate the contributions of the individual scenarios towards the overall objective.

One may argue that we may optimize the homogeneous clips i.e., clips of a single scenario, scenario by scenario as in [54]. It is noted that such a training strategy will not take care of the temporal transition between scenarios. The scenario-aware differentiated loss \mathcal{L}_{Diff} for an entire heterogeneous clip is necessary.

The labels of QQ, SQ, SS, QS scenarios and the segmentation process are only required during training, and not during inference. The training on heterogeneous clips with the proposed loss \mathcal{L}_{Diff} allows for optimization of the network for unknown speaker pairing scenarios at run-time.

It is worth mentioning that we do not normalize the extracted speech during inference, because this may potentially amplify the unwanted noise in the QQ and QS scenarios. We choose the SDR loss over the SI-SDR loss for (6) in the proposed scenario-aware differentiated loss, because the SDR loss is scale-sensitive, thus keeping the extracted speech at the same scale as the ground truth clean speech. It is noted that the SI-SDR loss is scale-invariant [73], thus the scale of the extracted speech may run away without a proper normalization.

C. Training Procedure

The overall training of the USEV network consists of three stages.

- 1) The *conv3D* and *resnet18* layers in the visual encoder are pre-trained together with a LSTM back-end on a visual speech recognition tasks.¹ The *conv3D* and *resnet18* layers are then kept frozen during subsequent USEV network training.
- 2) The entire USEV network is pre-trained on highly overlapped speech clips, with the SDR loss as shown in (6). This allows the USEV network to focus on the SS scenario as it is the most difficult scenario among the four [54].
- 3) The entire USEV network is trained on the general speech mixture clips, with the scenario-aware differentiated loss as shown in (8).

¹The *conv3D* and *resnet18* structures follow [86], and their pre-trained weights on visual speech recognition task are taken from https://github.com/lordmartian/deep_avsr

IV. EXPERIMENTAL SETUP

A. Dataset

1) *Highly Overlapped Speech Mixtures (VoxCeleb2-Mix Dataset)*: We use the VoxCeleb2 [87] dataset to simulate a highly overlapped speech dataset, denoted as VoxCeleb2-mix. The dataset is used for model pre-training.

The VoxCeleb2 dataset is an audio-visual dataset derived from YouTube videos. It has over 1 million videos from 6,112 celebrities. The videos are pre-processed with face detection and tracking algorithms. The resulting face tracking sequences are used as auxiliary references both for training and testing. We sample the audios at 16 kHz, the videos are synchronized with the audios and sampled at 25 FPS.

To create the VoxCeleb2-mix dataset, we randomly select 320,000 videos from 3,200 speakers in the original train set to create a training set (160,000 speech mixture clips), and 36,237 videos from 118 speakers in the original test set to form a test set (3,000 speech mixture clips). In either case, we only include videos that have at least 4 seconds of duration. To simulate a highly overlapped speech clip, an interference speech is mixed with the target speech at a Signal-to-Noise ratio (SNR) randomly set between 10 dB to -10 dB. Between the two mixing speech clips, the longer clip is truncated to the length of the shorter one to maximize the overlap. In addition, a noise from the WHAM! noise dataset [88] is also mixed with the target speech at an SNR randomly set between 15 dB to -5 dB.

2) *General Speech Mixtures (IEMOCAP-Mix Dataset)*: We use the Interactive Emotional Dyadic Motion Capture (IEMOCAP) [89] dataset to simulate the *general speech mixture* clips, named IEMOCAP-mix dataset. It is used to train and evaluate the contrastive baselines and the proposed USEV network.

The IEMOCAP dataset is an acted, multi-modal dataset. It consists of 12 hours of 150 dyadic conversations, each conversation has 2 speakers, with a total of 10 speakers in the dataset. The speakers' faces are always visible throughout, either speaking or not. We sample the audios at 16 kHz, the videos are synchronized with the audios and sampled at 25 FPS.

We first obtain video clips with clean utterances from the conversations. The utterances are then used to simulate the *general speech mixture* clips that cover all scenarios as described in Fig. 1.

From the original IEMOCAP dataset, we detect and track the faces in the videos following the VoxCeleb2 procedure [87], and obtain 300 face tracking videos. Based on the available speaker diarization label in the IEMOCAP dataset, we mute the audio when the face in the video is not speaking. In this way, each face tracking video consists of alternating quiet and speech segments, where a speech segment is always associated with the face in the video. We randomly select 240, 30, and 30 face tracking videos to simulate IEMOCAP-mix training, validation, and test sets respectively. A speaker may appear in multiple sets, but the speech does not. We split a face tracking video into multiple utterances by using a random quiet position between adjacent utterances as the delimiter.

To simulate the *general speech mixture* clips, we drop the utterances that are shorter than 3 seconds. We use a random segment from a target speaker's utterance, and mix it with random

segments from 1 or 2 interference speaker's utterances at an SNR ratio between 10 dB to -10 dB. Each speech clip could contain speech under one or more of the four scenarios as described in Fig. 1. In this way, the dataset covers all four scenarios. We simulate 400,000, 10,000, and 6,000 speech mixture clips for IEMOCAP-mix training, validation and test sets. The average length of the speech mixture clips is 5 seconds.

Besides the clean IEMOCAP-mix dataset simulated, we simulated a noisy version of the IEMOCAP-mix dataset, by adding noise from the WHAM! noise dataset [88] to the clean version of the IEMOCAP-mix dataset with an SNR randomly set between 15 dB to -5 dB.

3) *Composition of IEMOCAP-Mix*: A *general speech mixture* contains speech under one or more of the four scenarios, namely QQ , SQ , SS , QS . We record the start and endpoints of the scenarios in each speech clip. With such scenario labels, the training data are ready for network training with the scenario-aware differentiated loss, while the test data can be used for reporting evaluation performance by scenarios. We report the total duration in hours for the 4 scenarios in the last four columns of Table I.

As a *general speech mixture* could also be categorized as TA and TP speech clips, we report the total number of TP speech clips in groups by the target-interference speaker overlapping ratio, i.e., 0%, (0,20]%, (20,40]%, (40,60]%, (60,80]%, and (80,100]% in Table I. The target-interference speaker overlapping ratio is defined as the ratio of the duration of the overlapped segment SS to the total duration of non-silence speech $SS+SQ+QS$, i.e., $(SS \text{ duration}) / (SS+SQ+QS \text{ duration})$. We also report the total number of TA speech clips, the TA speech clips have no target speaker, therefore the target-interference speaker overlapping ratio does not apply. To simulate real-world data, we allow multiple target-interference speaker pairing scenarios to take place in a single speech clip.

B. USEV and Contrastive Baselines

We select two time-domain target speaker extraction networks as the contrastive baselines, namely SpEx+ [44] and TDSE [66]. The SpEx+ network employs a pre-recorded speech as the auxiliary reference, that shows the state-of-the-art performance. The TDSE network represents the recent advances of using the target's visemes as the auxiliary reference. The USEV network is a departure from the SpEx+ and TDSE networks. It is worth noting that the original SpEx+ and TDSE networks are trained and evaluated for highly overlapped speech in the previous works, which is neither a universal nor realistic acoustic scenario. We train the SpEx+, TDSE, and USEV network on general speech mixture clips, and seek to address all possible multi-talker scenarios in *cocktail party*, which represents an important step toward real-world applications.

To compare the proposed scenario-aware differentiated loss in (8) with the scenario-aware uniform loss [74], we develop contrastive baselines with the scenario-aware uniform loss function for the SpEx+, TDSE, and USEV networks respectively.

- 1) *SpEx+*: The SpEx+ network has a similar architecture to the USEV network, except that the SpEx+ has a speaker encoder, which encodes the audio cue, in place of a

TABLE I
A SUMMARY OF THE NUMBER OF SPEECH CLIPS FOR DIFFERENT TARGET-INTERFERENCE SPEAKER OVERLAPPING RATIOS, AND TOTAL DURATION FOR 4 TARGET-INTERFERENCE PAIRING SCENARIOS IN THE SIMULATED IEMOCAP-MIX DATASET

	TA clips	TP clips with different target-interference speaker overlapping ratio						Duration (hour)			
		0 %	(0,20] %	(20,40] %	(40,60] %	(60,80] %	(80,100] %	QQ	SQ	SS	QS
#training	20 437	34 859	24 256	46 846	73 533	84 946	115 123	48.73	83.19	266.34	133.41
#validation	295	585	496	1112	1869	2129	3514	0.85	1.79	7.53	3.01
#test	169	379	282	649	1081	1492	1948	0.53	1.05	4.19	1.74

visual encoder. In addition, the speaker extractor of the SpEx+ consists of repeated stacks of TCN blocks instead of repeated DPRNN blocks in USEV. We re-implement the SpEx+ by using scenario-aware uniform loss as the training objective, which is referred to as SpEx+(S). We also re-implement another variant of the SpEx+ network by using the scenario-aware differentiated loss as the training objective, which is referred to as SpEx+(D). Both the SpEx+(S) and SpEx+(D) serve as the contrastive baselines of the USEV network.

- 2) *TDSE*: The TDSE network has a similar architecture to the USEV network, except that the speaker extractor in the TDSE consists of repeated stacks of TCN blocks instead of repeated DPRNN blocks in USEV. We re-implement the TDSE by using scenario-aware uniform loss as the training objective, which is referred to as TDSE(S). We also re-implement another variant of the TDSE network by using the scenario-aware differentiated loss as the training objective, which is referred to as TDSE(D).
- 3) *USEV*: The USEV network with the scenario-aware differentiated loss is denoted as USEV(D). We also implement a variant of the USEV network with scenario-aware uniform loss, i.e., USEV(S) for comparison.

C. Implementation Details

The SpEx+, TDSE, and USEV networks have the same training procedure. They are pre-trained on highly overlapped speech clips first (VoxCeleb2-mix dataset), then trained and evaluated on the *general speech mixture* clips (IEMOCAP-mix dataset).

For pre-training on the VoxCeleb2-mix dataset, we use the adam optimizer with an initial learning rate of 0.001. The learning rate is decreased by 2% for every epoch, we train the networks for 30 epochs. During pre-training, the speech clips are truncated to 6 seconds to fit into the GPU memory.

For training on the IEMOCAP-mix dataset, we use the adam optimizer with an initial learning rate of 0.0001. The learning rate is decreased by 2% for every epoch, the training stops when the best validation loss does not improve for 8 consecutive epochs. During training, the speech clips are truncated to 6 seconds to fit into the GPU memory, during inference, the full speech clips are evaluated.² For the SpEx+ and TDSE networks, the L , B , N , and R are set to 40, 256, 256, and 4 according to [66]. For the USEV network, the L , B , N , R , and K are set to 40, 64, 256, 6, and 100 according to [82].

²The codes for the data generation and the USEV network are available at <https://github.com/zexupan/USEV>.

D. Evaluation Metrics

We use SI-SDR (dB) or power (dB/s) to evaluate our proposed method, which the two metrics are shown as follows:

$$SI-SDR = 10 \log_{10} \left(\frac{\| \frac{\langle \hat{s}, s \rangle s}{\|s\|^2 + \epsilon} \|^2}{\| \hat{s} - \frac{\langle \hat{s}, s \rangle s}{\|s\|^2 + \epsilon} \|^2 + \epsilon} + \epsilon \right) \quad (9)$$

$$Power = 10 \log_{10} \left(\frac{\| \hat{s} \|^2}{T_s} + \epsilon \right) \quad (10)$$

where T_s is the duration of \hat{s} in seconds.

V. RESULTS

We report three groups of experiments on the IEMOCAP-mix dataset. First, we empirically study the hyper-parameters settings and the training strategy, which are reported in Section V-A. Second, we present 3 comparative studies, that are reported in Section V-B. We compare the visual cue with the audio cue in experiment 1, the scenario-aware differentiated loss with the scenario-aware uniform loss in experiment 2, and the DPRNN structure with the TCN structure in experiment 3. Third, we study how the quality of the visual cue affects the speaker extraction performance of the USEV and TDSE networks, which are reported in Section V-C.

A. System Tuning

1) *Weights for the Scenario-Aware Differentiated Loss*: We evaluate the USEV(D) on the clean and noisy IEMOCAP-mix validation set to empirically find the appropriate weights $\alpha, \beta, \gamma, \delta$ for the differentiated loss in Table II. We compare the SI-SDR for the SQ and SS scenario segments, and the power of the QQ and QS scenario segments. It is worth mentioning that the proposed scenario-aware differentiated loss for training is scale-sensitive, thus the extracted speech for the systems are on the same scale. When comparing between systems, the higher the better for the SI-SDR, and the lower the better for the power.

In Table II, systems 1 to 4 are trained and evaluated on the clean IEMOCAP-mix dataset, and systems 5 to 7 are trained and evaluated on the noisy IEMOCAP-mix dataset. It is seen that systems 3 and 7 achieves reasonably low power for the QQ and QS scenarios and high SI-SDR for the SQ and SS scenarios for clean and noisy dataset respectively. We therefore set $\alpha, \beta, \gamma, \delta$ to be 0.005, 1, 1, 0.005 respectively as in system 3 and 7 for future system evaluations. It is worth noting that the selected weights are optimized for the IEMOCAP-mix dataset, for which the total duration of the four scenario segments are shown in Table I. If the total duration of the four scenario

TABLE II

EXPERIMENTS OF USEV(D) WITH DIFFERENT HYPER-PARAMETERS ON THE CLEAN AND NOISY IEMOCAP-MIX VALIDATION SET RESPECTIVELY. WE REPORT THE SI-SDR (DB) FOR THE SQ AND SS SCENARIO SEGMENTS, AND THE POWER (DB/S) OF THE QQ AND QS SCENARIO SEGMENTS. THE HIGHER THE BETTER FOR THE SI-SDR AND THE LOWER THE BETTER FOR THE POWER. N.A. REPRESENTS NOT APPLICABLE

Noise Condition	System	$\alpha - \beta - \gamma - \delta$	QQ Power	SQ SI-SDR	SS SI-SDR	QS Power
Clean	Mixture	N.A.	-79.75	61.99	-0.86	8.34
	1	0.01-1-1-0.01	-73.80	34.91	6.11	-55.88
	2	0.01-0.1-1-0.01	-75.91	27.30	6.75	-68.71
	3	0.005-1-1-0.005	-72.84	35.45	6.58	-37.87
	4	0.001-0.1-1-0.001	-67.87	30.36	7.19	-16.02
Noisy	Mixture	N.A.	13.01	3.62	-2.12	21.18
	5	0.01-1-1-0.01	-69.26	0.15	3.57	-59.08
	6	0.005-1-1-0.01	-62.89	1.10	3.51	-56.12
	7	0.005-1-1-0.005	-50.17	4.29	4.13	-37.71
	8	0.001-1-1-0.001	-17.58	6.74	4.50	-13.76

TABLE III

EXPERIMENTS OF USEV(D) WITH DIFFERENT TRAINING STRATEGY ON THE NOISY IEMOCAP-MIX VALIDATION SET. WE REPORT THE SI-SDR (DB) FOR THE SQ AND SS SCENARIO SEGMENTS, AND THE POWER (DB/S) OF THE QQ AND QS SCENARIO SEGMENTS. THE HIGHER THE BETTER FOR THE SI-SDR AND THE LOWER THE BETTER FOR THE POWER

System	PT	FT	QQ Power	SQ SI-SDR	SS SI-SDR	QS Power
Mixture	N.A.	N.A.	13.01	3.62	-2.12	21.18
7	✓	✓	-50.17	4.29	4.13	-37.71
9	✓	✗	-6.87	3.28	0.75	6.58
10	✗	✓	-65.24	-1.75	-0.31	-53.67

segments changes when another dataset is used, these weights need to be adjusted accordingly.

2) *Training Strategy*: We evaluate the USEV(D) on the noisy IEMOCAP-mix validation set to empirically find the appropriate training strategy as shown in Table III. We compare the SI-SDR for the SQ and SS scenario segments, and the power of the QQ and QS scenario segments. The higher the better for the SI-SDR and the lower the better for the power. Systems 9 and 10 are trained and evaluated on the noisy IEMOCAP-mix dataset. System 7 in Table III is the same as system 7 in Table II.

The USEV(D) in system 7 is pre-trained (PT) on highly overlapped speech first (VoxCeleb2-mix dataset), and further trained (FT) on the *general speech mixtures* (IEMOCAP-mix dataset). We conduct two experiments to justify this training procedure. In system 9, USEV(D) is pre-trained on the VoxCeleb2-mix dataset, but it is not further trained on the IEMOCAP-mix dataset. It is seen that system 9 has lower SI-SDR for the SQ and SS scenarios, and higher power for the QQ and QS scenarios, as compared to system 7 which is further trained on the IEMOCAP-mix dataset. This could be due to the fact that system 9 is only trained on highly overlapped speech.

In system 10, we do not pre-train the USEV(D) on the VoxCeleb2-mix dataset, but rather train it directly on the IEMOCAP-mix dataset from scratch. System 10 significantly lags behind system 7 on the SQ and SS scenarios in terms

of SI-SDR, where the target speaker is highly overlapped with the interference speakers or noise. We note that the SQ and SS scenarios are two very challenging scenarios among the four, the pre-training on highly overlapped speech is effective because it well prepares the network for the two scenarios.

B. Comparative Studies

The SpEx+(S), TDSE(S), and USEV(S) adopt the uniform loss as the training objective, while the SpEx+(D), TDSE(D), and USEV(D) adopt the differentiated loss. For a fair comparison between the two loss functions, we report the performance on an entire speech clip with a single evaluation metric, i.e., SI-SDR for TP speech clips or power for TA speech clips in Table IV.

In Table IV, we first categorize the speech clips into two categories: i) TA speech clips: speech clips with the absent target speaker. ii) TP speech clips: speech clips with the present target speaker. We report the average power of TA speech clips, the lower the better. For TP speech clips, we report the SI-SDR by their target-interference speaker overlapping ratios. The average SI-SDR is also reported in the last column, the higher the better for the SI-SDR.

In addition, we report the results using differentiated metrics in terms of QQ , SQ , SS , QS scenarios in Table V, to examine the performance over individual scenarios. The models reported in Table IV and Table V are both trained on the IEMOCAP-mix dataset, but are evaluated using different metrics, i.e., a single SI-SDR or power over an entire speech clip in Table IV, and differentiated metrics for every scenario segment in Table V.

1) *Visual vs Audio Cue*: We compare the use of visemes with the use of a pre-recorded speech as the auxiliary reference on the IEMOCAP-mix test set. The SpEx+ network employs a pre-recorded speech as the audio cue, while the TDSE and USEV networks employ the visemes as the visual cue. The pre-recorded speech used for the SpEx+ is the target's other clean speech utterance that is not present in the speech mixture, the reference utterance is about 5 seconds long on average.

The SpEx+(S), TDSE(S), and USEV(S) are trained using the uniform loss. As shown in Table IV, for both clean and noisy conditions, the TDSE(S) and USEV(S) outperform the SpEx+(S) in terms of average SI-SDR, and the TA power. As shown in Table V, for both clean and noisy conditions, the TDSE(S) and USEV(S) outperform the SpEx+(S) for all scenarios. Results show that when the uniform loss is used, the visemes outperform the pre-recorded speech as the auxiliary reference.

The SpEx+(D), TDSE(D), and USEV(D) are trained using the differentiated loss. The weights of the differentiated loss are selected the same as system 7 in Table II. As shown in Table IV, for both clean and noisy conditions, the TDSE(D) and USEV(D) consistently outperform the SpEx+(D) in terms of average SI-SDR, and the TA power. As shown in Table V, for both clean and noisy conditions, the USEV(D) and TDSE(D) outperform the SpEx+(D) for all scenarios, except for the SQ scenario under noisy conditions. Results show that when the differentiated loss is used, the visemes outperform the pre-recorded speech as the auxiliary reference.

TABLE IV

COMPARATIVE STUDY AMONG VARIANTS OF THE SpEx+ [54], TDSE [66], AND USEV NETWORKS ON IEMOCAP-MIX TEST SET. WE REPORT THE AVERAGE POWER (dB/S) OF TA SPEECH CLIPS, THE SI-SDR (dB) FOR TP SPEECH CLIPS BY TARGET-INTERFERENCE OVERLAPPING RATIOS, AND THE AVERAGE SI-SDR (dB) FOR TP SPEECH CLIPS. THE HIGHER THE BETTER FOR THE SI-SDR AND THE LOWER THE BETTER FOR THE POWER

Noise Condition	Model	Reference	Extractor	Loss	TA Power	0% SI-SDR	(0,20]% SI-SDR	(20,40]% SI-SDR	(40,60]% SI-SDR	(60,80]% SI-SDR	(80,100]% SI-SDR	Average SI-SDR
Clean	Mixture	N.A.	N.A.	N.A.	9.96	33.41	-0.98	-1.32	-1.52	-1.53	-1.99	4.34
	SpEx+(S) SpEx+(D)	Speech	TCN	Uniform Differentiated	-40.87 3.13	12.36 20.18	8.05 10.31	5.36 8.34	4.80 6.19	3.88 5.44	3.07 4.40	6.25 9.14
	TDSE(S) TDSE(D)	Visual	TCN	Uniform Differentiated	-76.86 -22.82	23.98 29.80	13.85 15.72	7.58 9.70	7.94 8.29	6.75 6.95	5.82 6.13	10.99 12.76
	USEV(S) USEV(D)	Visual	DPRNN	Uniform Differentiated	-77.48 -43.00	30.53 33.83	15.23 15.20	7.14 9.91	7.37 8.06	5.96 6.89	5.46 5.95	11.95 13.31
Noisy	Mixture	N.A.	N.A.	N.A.	22.93	1.9	-2.53	-2.47	-2.76	-2.79	-2.91	-1.93
	SpEx+(S) SpEx+(D)	Speech	TCN	Uniform Differentiated	-32.27 2.93	1.76 6.66	0.11 0.96	1.84 3.02	1.67 2.43	0.85 2.16	1.13 1.54	1.22 2.79
	TDSE(S) TDSE(D)	Visual	TCN	Uniform Differentiated	-75.46 -32.93	4.92 8.56	3.79 6.64	3.79 6.61	4.56 5.80	4.17 4.90	3.58 4.29	4.14 6.13
	USEV(S) USEV(D)	Visual	DPRNN	Uniform Differentiated	-74.04 -22.23	5.01 8.97	3.16 6.05	3.78 6.22	4.11 5.30	3.69 4.46	3.36 4.14	3.85 5.86

TABLE V

COMPARATIVE STUDY AMONG VARIANTS OF THE SpEx+ [54], THE TDSE [66], AND USEV NETWORKS ON IEMOCAP-MIX TEST SET. WE REPORT THE SI-SDR FOR THE SQ AND SS SCENARIO SEGMENTS, AND THE POWER OF THE QQ AND QS SCENARIO SEGMENTS. THE HIGHER THE BETTER FOR THE SI-SDR AND THE LOWER THE BETTER FOR THE POWER

Noise Condition	Model	Reference	Extractor	Loss	QQ Power	SQ SI-SDR	SS SI-SDR	QS Power	Parameters (million)
Clean	Mixture	N.A.	N.A.	N.A.	-80	63.02	-0.8573	10.23	-
	SpEx+(S) SpEx+(D)	Speech	TCN	Uniform Differentiated	-45.56 -42.25	14.56 22.95	3.50 5.09	-17.15 -14.65	11.9
	TDSE(S) TDSE(D)	Visual	TCN	Uniform Differentiated	-65.17 -69.29	23.42 33.88	5.75 6.33	-29.37 -30.24	22.1
	USEV(S) USEV(D)	Visual	DPRNN	Uniform Differentiated	-65.29 -71.87	24.26 35.40	5.35 6.19	-29.17 -35.58	15.3
Noisy	Mixture	N.A.	N.A.	N.A.	13.33	3.88	-2.05	21.30	-
	SpEx+(S) SpEx+(D)	Speech	TCN	Uniform Differentiated	-27.45 -27.10	2.69 4.87	1.08 2.07	-13.66 -7.91	11.9
	TDSE(S) TDSE(D)	Visual	TCN	Uniform Differentiated	-35.91 -59.38	4.68 3.27	3.03 4.04	-26.12 -40.14	22.1
	USEV(S) USEV(D)	Visual	DPRNN	Uniform Differentiated	-34.37 -50.04	4.45 4.85	2.76 3.83	-25.28 -34.00	15.3

2) *Differentiated Loss vs Uniform Loss*: We compare the use of differentiated loss with the uniform loss on the baselines and our proposed USEV network in disentangling the general speech mixture clips from the IEMOCAP-mix test set.

As shown in Table IV, for both clean and noisy conditions, the SpEx+(D) performs better in terms of SI-SDR evaluation, but the SpEx+(S) performs better in terms of TA power. Further look into Table V, for both clean and noisy conditions, the SpEx+(D) performs better for the SQ and SS scenarios, but the SpEx+(S) has lower power for the QQ and QS scenarios. For the SpEx+ network, the uniform loss focuses more on muting the network for inactive target speaker scenarios, while the differentiated loss focuses more on the signal extraction quality for active target speaker scenarios, this could be caused by the large weights

used for the SQ and SS scenarios used by the differentiated loss function.

As shown in Table IV, for both clean and noisy conditions, the TDSE(D) outperforms the TDSE(S) in terms of the SI-SDR evaluation, but the TDSE(S) performs better for TA power. As shown in Table V, for both clean and noisy conditions, the TDSE(D) outperforms the TDSE(S) for scenarios, except for the SQ scenario under noisy conditions.

As shown in Table IV, for both clean and noisy conditions, the USEV(D) outperforms USEV(S) in terms of the SI-SDR evaluation, but the USEV(S) performs better for TA power. As shown in Table V, for both clean and noisy conditions, the USEV(D) outperforms USEV(S) for all scenarios.

In Table IV, the uniform loss consistently performs better than the differentiated loss in terms of power for TA speech

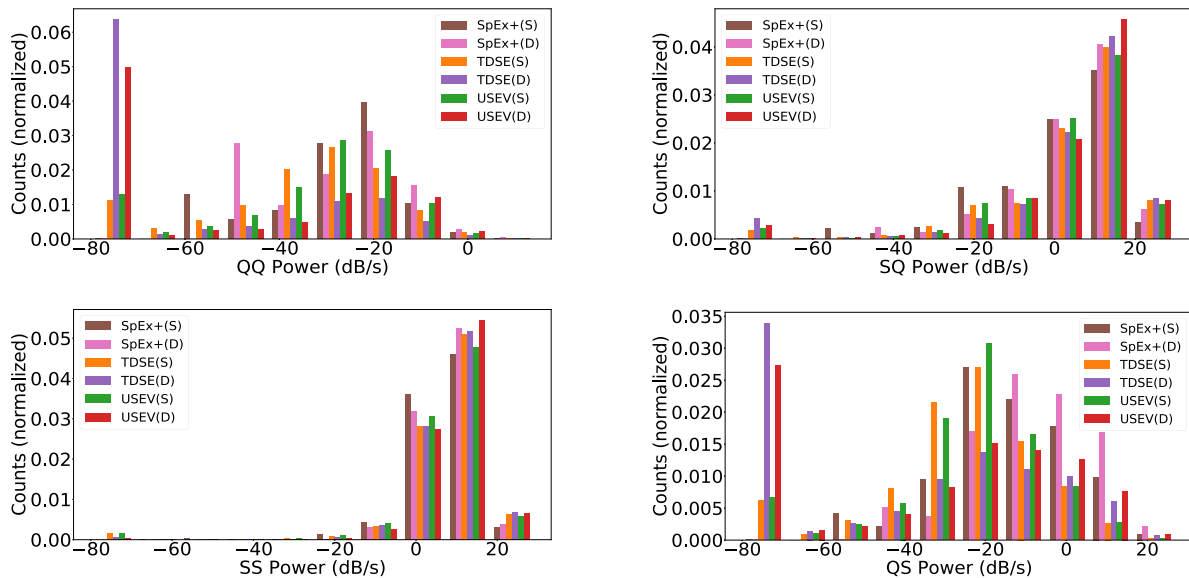


Fig. 6. The histogram of the power by the SpEx+(S), SpEx+(D), TDSE(S), TDSE(D), USEV(S), and USEV(D) on the noisy IEMOCAP-mix test sets. The lower the better for the power of the QQ and QS scenario, and the power of the SQ and SS scenario are expected to be similar to that of the clean target speech, which is on average 10 dB/s.

clips. This is because the uniform loss puts a high weight of 1 on the TA speech clips during training, as opposed to 0.005 for the differentiated loss (QQ and QS). It is worth noting that the TA speech clips are only a subset of all QQ and QS scenario segments. If we compare the overall power of the QQ and QS scenarios in Table V, the differentiated loss consistently outperforms the uniform loss in terms of power for TDSE and USEV models.

We also present the power distribution of the models on the noisy IEMOCAP-mix test sets in Fig. 6. It is seen that for the SQ and SS scenarios in which the target speaker is active, all models perform similar, with most of the extracted speech signals have a power value that is close to that of the clean target speech, which is on average 10 dB/s. For the QQ and QS scenarios in which the target speaker is inactive, it is seen that the USEV(D) and TDSE(D) have more samples with lower power compared to the USEV(S) and TDSE(S), showing that the differentiated loss outperforms the uniform loss for audio-visual models.

3) *DPRNN vs TCN*: We compare the DPRNN with TCN in the speaker extractor when the visual cue is employed as the auxiliary reference. The TDSE network employs a repeated stack of TCN in the speaker extractor to estimate the receptive mask, while the USEV network employs a repeated DPRNN block to estimate the receptive mask.

As shown in Table IV, when the uniform loss is used for network training, the USEV(S) outperforms the TDSE(S) for clean conditions, but the TDSE(S) outperforms the USEV(S) for noisy conditions. As shown in Table V, for clean conditions, the USEV(S) performs better for the QQ and SQ scenarios, but the TDSE(S) performs better for the SS and QS scenarios, for noisy condition, the TDSE(S) performs better than the USEV(S) for all scenarios.

As shown in Table IV, when the differentiated loss is used for network training, the USEV(S) outperforms the TDSE(D) for

clean conditions, but the TDSE(D) outperforms the USEV(D) for noisy conditions. As shown in Table V, the USEV(D) outperforms the TDSE(D) in all scenarios except for the SS scenario under clean condition, the TDSE(D) outperforms the USEV(D) in all scenarios except for the SQ scenario under noisy condition.

The USEV tends to perform better when the overlapping ratio is small and the TDSE tends to perform better when the overlapping ratio is large. Overall, the USEV and TDSE networks achieve comparable results on both clean and noisy IEMOCAP-mix test sets. However, the USEV network has 6.8 million fewer parameters compared with the TDSE network due to the smaller size of the DPRNN blocks as shown in Table V. In addition, when the kernel size L in the speech encoder of the two networks is set the same, which is 40 in this paper, the training of the USEV network is 4 times faster than that of the TDSE network. The TDSE and USEV network have more network parameters compared with the SpEx+ network due to the large visual encoder module. The no. of network parameters reported for the TDSE and USEV networks do not include the face detection and tracking module.

C. Effect of Visual Occlusion

When face detection and tracking algorithms fail to detect the target speaker or the lip is occluded for some reason, the visual cue is absent or incomplete. We simulate the above scenario, referred to as visual occlusion, by setting a random duration of the lip images sequence to zero signals for each speech clip in the clean IEMOCAP-mix dataset, while keeping the audio signal intact. We study the impact of such visual occlusion on the audio-visual models, namely the TDSE(S), TDSE(D), USEV(S), and USEV(D), and present their results in Table VI.

The models are first pre-trained on the VoxCeleb2-mix dataset, where there is no occlusion in the training data. In the

TABLE VI

COMPARATIVE STUDIES AMONG VARIANTS OF THE TDSE [66], AND USEV NETWORKS ON IEMOCAP-MIX TEST SET FOR A VISUAL CUE WITH OCCLUSION. THE VALUES REPORTED ARE THE AVERAGE OVER ALL VISUAL OCCLUSION SETTINGS IN ITS GROUP, FOR MODELS THAT ARE TRAINED ON IEMOCAP-MIX TRAINING DATA WITH AND WITHOUT OCCLUSION (OCCL.)

Model	Extractor	Loss	Train Occl.	TA Power	0% SI-SDR	(0,20]% SI-SDR	(20,40]% SI-SDR	(40,60]% SI-SDR	(60,80]% SI-SDR	(80,100]% SI-SDR	Average SI-SDR
Mixture	N.A.	N.A.	N.A.	9.96	33.41	-0.98	-1.32	-1.52	-1.53	-1.99	4.34
TDSE(S)	TCN	Uniform	✗	-76.98	-11.50	-3.28	-14.85	-20.36	-16.09	-12.53	-13.10
TDSE(D)	TCN	Differentiated		-12.72	15.19	6.82	3.22	0.54	-0.16	-0.51	4.18
USEV(S)	DPRNN	Uniform		-75.34	-8.88	-6.97	-20.40	-25.67	-21.76	-16.96	-16.77
USEV(D)	DPRNN	Differentiated		-29.03	15.88	5.50	2.62	0.04	-0.44	-0.58	3.84
TDSE(S)	TCN	Uniform	✓	-61.57	23.48	13.14	7.40	5.90	5.63	4.92	10.08
TDSE(D)	TCN	Differentiated		5.24	25.56	12.87	8.00	6.10	5.55	4.60	10.45
USEV(S)	DPRNN	Uniform		-67.28	24.40	12.08	4.07	4.48	3.63	3.77	8.74
USEV(D)	DPRNN	Differentiated		-6.03	31.11	12.45	7.57	5.73	5.00	4.37	11.04

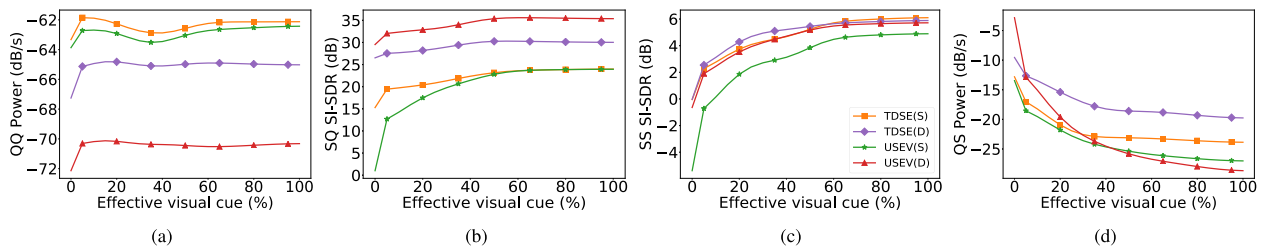


Fig. 7. Comparative studies among the TDSE(S), TDSE(D), USEV(S), and USEV(D) with partially occluded visual reference. The models are trained on the IEMOCAP-mix training data with visual occlusion. We report the performance in QQ , SQ , SS , QS scenarios separately as a function of the effective visual cue, i.e., the ratio of the effective visual duration to the total length of the speech clip. For each model, the lower the better for power in (a) QQ scenario and (d) QS scenario, and the higher the better for SI-SDR in (b) SQ scenario and (c) SS scenario.

training stage with the IEMOCAP-mix dataset, if the occlusion data are not involved (Train Occl. ✗), none of the models perform well on occlusion evaluations. If the occlusion data are involved in the training stage on the IEMOCAP-mix dataset (Train Occl. ✓), all models improve, with USEV(D) achieving the highest average SI-SDR value of 11.04 dB. The USEV tends to perform better when the overlapping ratio is small and the TDSE tends to perform better when the overlapping ratio is large, overall they achieve comparable performance for the visual occlusion analysis.

We also evaluate the models with the different effective visual cue and report the results in Fig. 7 for the four scenarios, to observe the trend of model performance with the different effective visual cue. We define the effective visual cue (%) as the ratio of the non-occluded visual duration to the total length of the speech clip. The average performance for each 5% interval of the effective visual cue is plotted, e.g., the first point is the average performance for 0-5% effective visual cue, and the last point is 96-100% effective visual cue. It is observed that the performance of SQ , SS , QS scenarios improves as the effective visual duration increases. When the effective visual cue reaches 20%, the USEV(D) performs reasonably well across all scenarios, when it goes beyond 50%, the USEV(D)'s performance starts to saturate.

VI. CONCLUSION

In this paper, we propose a categorization of a *general speech mixture* into four target-interference speaker pairing scenarios

namely QQ , SQ , SS , QS . We also proposed a universal speaker extraction network with a visual cue to disentangle the *general speech mixture* with a scenario-aware differentiated loss function. The experiments show that the proposed differentiated loss function is effective. This paper marks an important step towards solving a realistic *cocktail party* problem.

REFERENCES

- [1] E. C. Cherry, "Some experiments on the recognition of speech, with one and with two ears," *J. Acoustical Soc. Amer.*, vol. 25, no. 5, pp. 975–979, 1953.
- [2] C. Xu, W. Rao, E. S. Chng, and H. Li, "SpEx: Multi-scale time domain speaker extraction network," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 28, pp. 1370–1384, 2020.
- [3] K. Žmolíková et al., "SpeakerBeam: Speaker aware neural network for target speaker extraction in speech mixtures," *IEEE J. Sel. Topics Signal Process.*, vol. 13, no. 4, pp. 800–814, Aug. 2019.
- [4] Q. Wang et al., "VoiceFilter: Targeted voice separation by speaker-conditioned spectrogram masking," in *Proc. Interspeech*, 2019, pp. 2728–2732.
- [5] Z. Pan, R. Tao, C. Xu, and H. Li, "Selective listening by synchronizing speech with lips," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 30, pp. 1650–1664, 2022.
- [6] A. W. Bronkhorst, "The cocktail party phenomenon: A review of research on speech intelligibility in multiple-talker conditions," *Acta Acustica United With Acustica*, vol. 86, no. 1, pp. 117–128, 2000.
- [7] D. Wang, "Deep learning reinvents the hearing aid," *IEEE Spectr.*, vol. 54, no. 3, pp. 32–37, Mar. 2017.
- [8] R. Tao, Z. Pan, R. K. Das, X. Qian, M. Z. Shou, and H. Li, "Is someone speaking? Exploring long-term temporal features for audio-visual active speaker detection," in *Proc. 29th ACM Int. Conf. Multimedia*, 2021, pp. 3927–3935.
- [9] R. Tao, R. K. Das, and H. Li, "Audio-visual speaker recognition with a cross-modal discriminative network," in *Proc. Interspeech*, 2020, pp. 2242–2246.

- [10] X. Qian, M. Madhavi, Z. Pan, J. Wang, and H. Li, "Multi-target DoA estimation with an audio-visual fusion mechanism," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2021, pp. 4280–4284.
- [11] X. Yue, G. Lee, E. Yilmaz, F. Deng, and H. Li, "End-to-end code-switching ASR for low-resourced language pairs," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2019, pp. 972–979.
- [12] M. N. Schmidt and R. K. Olsson, "Single-channel speech separation using sparse non-negative matrix factorization," in *Proc. 9th Int. Conf. Spoken Lang. Process.*, 2006, pp. 2–5.
- [13] A. Cichocki, R. Zdunek, and S.-i. Amari, "New algorithms for non-negative matrix factorization in applications to blind source separation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, vol. 5, 2006, pp. V–V.
- [14] M. Stark, M. Wohlmayr, and F. Pernkopf, "Source-filter-based single-channel speech separation using pitch information," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 19, no. 2, pp. 242–255, Feb. 2010.
- [15] R. Lyon, "A computational model of binaural localization and separation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 1983, vol. 8, pp. 1148–1151.
- [16] D. Wang and G. J. Brown, *Computational Auditory Scene Analysis: Principles, Algorithms, and Applications*. Hoboken, NJ, USA: Wiley, 2006.
- [17] G. Hu and D. Wang, "Auditory segmentation based on onset and offset analysis," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 15, no. 2, pp. 396–405, Feb. 2007.
- [18] C. Xu, W. Rao, X. Xiao, E. S. Chng, and H. Li, "Single channel speech separation with constrained utterance level permutation invariant training using grid LSTM," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 6–10.
- [19] I. Kavaleroev et al., "Universal sound separation," in *Proc. Workshop Appl. Signal Process. Audio Acoust.*, 2019, pp. 175–179.
- [20] Z. Chen, Y. Luo, and N. Mesgarani, "Deep attractor network for single-microphone speaker separation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2017, pp. 246–250.
- [21] D. Stoller, S. Ewert, and S. Dixon, "Wave-U-Net: A multi-scale neural network for end-to-end audio source separation," in *Proc. 19th Int. Soc. Music Inf. Retrieval Conf.*, 2018, pp. 334–340.
- [22] M. Kolbaek, D. Yu, Z.-H. Tan, and J. Jensen, "Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 25, no. 10, pp. 1901–1913, Oct. 2017.
- [23] Y. Liu and D. Wang, "Divide and conquer: A deep CASA approach to talker-independent monaural speaker separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 12, pp. 2092–2102, Dec. 2019.
- [24] T. von Neumann et al., "End-to-end training of time domain audio separation and recognition," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 7004–7008.
- [25] T. von Neumann, C. Boeddeker, K. Kinoshita, M. Delcroix, and R. Haeb-Umbach, "Speeding up permutation invariant training for source separation," in *Proc. Speech Commun., 14th ITG Conf.*, 2021, pp. 1–5.
- [26] J. Han, Y. Long, L. Burget, and J. Černocký, "DPCCN: Densely-connected pyramid complex convolutional network for robust speech separation and extraction," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 7292–7296.
- [27] C. Subakan, M. Ravanelli, S. Cornell, and F. Grondin, "REAL-M: Towards speech separation on real mixtures," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 6862–6866.
- [28] H. Taherian, K. Tan, and D. Wang, "Location-based training for multi-channel talker-independent speaker separation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 696–700.
- [29] T. von Neumann, K. Kinoshita, M. Delcroix, S. Araki, T. Nakatani, and R. Haeb-Umbach, "All-neural online source separation, counting, and diarization for meeting analysis," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2019, pp. 91–95.
- [30] T. von Neumann, K. Kinoshita, C. Boeddeker, M. Delcroix, and R. Haeb-Umbach, "Graph-PIT: Generalized permutation invariant training for continuous separation of arbitrary numbers of speakers," in *Proc. Interspeech*, 2021, pp. 3490–3494.
- [31] T. von Neumann, K. Kinoshita, C. Boeddeker, M. Delcroix, and R. Haeb-Umbach, "SA-SDR: A novel loss function for separation of meeting style data," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 6022–6026.
- [32] C. Li et al., "Dual-path modeling for long recording speech separation in meetings," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2021, pp. 5739–5743.
- [33] T. Yoshioka, H. Erdogan, Z. Chen, and F. Alleva, "Multi-microphone neural speech separation for far-field multi-talker speech recognition," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 5739–5743.
- [34] S. Chen et al., "Continuous speech separation with conformer," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2021, pp. 5749–5753.
- [35] Z.-Q. Wang and D. Wang, "Count and separate: Incorporating speaker counting for continuous speaker separation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2021, pp. 11–15.
- [36] Y. Zhang et al., "Continuous speech separation with recurrent selective attention network," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 6017–6021.
- [37] Z. Zhang et al., "All-neural beamformer for continuous speech separation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 6032–6036.
- [38] T. Ochiai, M. Delcroix, K. Kinoshita, A. Ogawa, and T. Nakatani, "Multi-modal SpeakerBeam: Single channel target speech extraction with audio-visual speaker clues," in *Proc. Interspeech*, 2019, pp. 2718–2722.
- [39] H. Sato, T. Ochiai, K. Kinoshita, M. Delcroix, T. Nakatani, and S. Araki, "Multimodal attention fusion for target speaker extraction," in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2021, pp. 778–784.
- [40] M. Delcroix, K. Zmolikova, K. Kinoshita, A. Ogawa, and T. Nakatani, "Single channel target speaker extraction and recognition with speaker beam," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 5554–5558.
- [41] Z. Pan, X. Qian, and H. Li, "Speaker extraction with co-speech gestures cue," *IEEE Signal Process. Lett.*, vol. 29, pp. 1467–1471, 2022.
- [42] Z. Pan, M. Ge, and H. Li, "A hybrid continuity loss to reduce over-suppression for time-domain target speaker extraction," in *Proc. Interspeech*, 2022, pp. 1786–1790.
- [43] Y. Wu, C. Li, J. Bai, Z. Wu, and Y. Qian, "Time-domain audio-visual speech separation on low quality videos," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 256–260.
- [44] M. Ge, C. Xu, L. Wang, E. S. Chng, J. Dang, and H. Li, "SpEx: A complete time domain speaker extraction network," in *Proc. Interspeech*, 2020, pp. 1406–1410.
- [45] M. Ge, C. Xu, L. Wang, E. S. Chng, J. Dang, and H. Li, "Multi-stage speaker extraction with utterance and frame-level reference signals," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2021, pp. 6109–6113.
- [46] S. He, H. Li, and X. Zhang, "Speakerfilter: Deep learning-based target speaker extraction using anchor speech," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 376–380.
- [47] X. Xiao et al., "Single-channel speech extraction using speaker inventory and attention network," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2019, pp. 86–90.
- [48] J. Shi, J. Xu, Y. Fujita, S. Watanabe, and B. Xu, "Speaker-conditional chain model for speech separation and extraction," in *Proc. Interspeech*, 2020, pp. 2707–2711.
- [49] M. Delcroix et al., "Improving speaker discrimination of target speech extraction with time-domain SpeakerBeam," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 691–695.
- [50] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2016, pp. 31–35.
- [51] Ö. Çetin and E. Shriberg, "Analysis of overlaps in meetings by dialog factors, hot spots, speakers, and collection site: Insights for automatic speech recognition," in *Proc. 9th Int. Conf. Spoken Lang. Process.*, 2006, pp. 212–224.
- [52] J. Barker, S. Watanabe, E. Vincent, and J. Trmal, "The fifth 'CHiME' speech separation and recognition challenge: Dataset, task and baselines," in *Proc. Interspeech*, 2018, pp. 1561–1565.
- [53] Z. Chen et al., "Continuous speech separation: Dataset and analysis," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 7284–7288.
- [54] M. Borsdorf, C. Xu, H. Li, and T. Schultz, "Universal speaker extraction in the presence and absence of target speakers for speech of one and two talkers," in *Proc. Interspeech*, 2021, pp. 1469–1473.
- [55] Z. Zhang, B. He, and Z. Zhang, "X-TaSNet: Robust and accurate time-domain speaker extraction network," in *Proc. Interspeech*, 2020, pp. 1421–1425.
- [56] L. Smith and M. Gasser, "The development of embodied cognition: Six lessons from babies," *Artif. Life*, vol. 11, no. 1–2, pp. 13–29, 2005.
- [57] T. Afouras, J. S. Chung, and A. Zisserman, "The conversation: Deep audio-visual speech enhancement," in *Proc. Interspeech*, 2018, pp. 3244–3248.

- [58] C. Li and Y. Qian, "Listen, watch and understand at the cocktail party: Audio-visual-contextual speech separation," in *Proc. Interspeech*, 2020, pp. 1426–1430.
- [59] E. Z. Golubovic, G. B. Cogan, C. E. Schroeder, and D. Poeppel, "Visual input enhances selective speech envelope tracking in auditory cortex at a "cocktail party";" *J. Neurosci.*, vol. 33, no. 4, pp. 1417–1426, 2013.
- [60] M. J. Crosse, G. M. Di Liberto, and E. C. Lalor, "Eye can hear clearly now: Inverse effectiveness in natural audiovisual speech processing relies on long-term crossmodal temporal integration," *J. Neurosci.*, vol. 36, no. 38, pp. 9888–9895, 2016.
- [61] G. M. Edelman, *Neural Darwinism: The Theory of Neuronal Group Selection*. New York, NY, USA: Basic books, 1987.
- [62] H. L. Bear and R. Harvey, "Phoneme-to-viseme mappings: The good, the bad, and the ugly," *Speech Commun.*, vol. 95, pp. 40–67, 2017.
- [63] D. W. Massaro and J. A. Simpson, *Speech Perception by Ear and Eye: A Paradigm for Psychological Inquiry*. London, U.K.: Psychol. Press, 2014.
- [64] D. Michelsanti et al., "An overview of deep-learning-based audio-visual speech enhancement and separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 1368–1396, 2021.
- [65] A. Ephrat et al., "Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation," *ACM Trans. Graph.*, vol. 37, no. 4, pp. 1–11, 2018.
- [66] J. Wu et al., "Time domain audio visual speech separation," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2019, pp. 667–673.
- [67] Z. Pan, R. Tao, C. Xu, and H. Li, "Muse: Multi-modal target speaker extraction with visual cues," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2021, pp. 6678–6682.
- [68] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 19, no. 4, pp. 788–798, May 2011.
- [69] D. Snyder, P. Ghahremani, D. Povey, D. Garcia-Romero, Y. Carmiel, and S. Khudanpur, "Deep neural network-based speaker embeddings for end-to-end speaker verification," in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2016, pp. 165–170.
- [70] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, "Generalized end-to-end loss for speaker verification," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 4879–4883.
- [71] T. Liu, R. K. Das, M. Madhavi, S. Shen, and H. Li, "Speaker-utterance dual attention for speaker and utterance verification," in *Proc. Interspeech*, 2020, pp. 4293–4297.
- [72] K. Žmolíková, M. Delcroix, K. Kinoshita, T. Higuchi, A. Ogawa, and T. Nakatani, "Learning speaker representation for neural network based multichannel speaker extraction," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2017, pp. 8–15.
- [73] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "SDR-half-baked or well done?," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2019, pp. 626–630.
- [74] S. Wisdom et al., "What's all the fuss about free universal sound separation data?," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2021, pp. 186–190.
- [75] Q. Lin, L. Yang, X. Wang, L. Xie, C. Jia, and J. Wang, "Sparsely overlapped speech training in the time domain: Joint learning of target speech separation and personal VAD benefits," in *Proc. Asia-Pacific Signal Inf. Process. Assoc. Annu. Summit Conf.*, 2021, pp. 689–693.
- [76] S.-W. Chung, S. Choe, J. S. Chung, and H.-G. Kang, "FaceFilter: Audio-visual speech separation using still images," in *Proc. Interspeech*, 2020, pp. 3481–3485.
- [77] T. Stafylakis and G. Tzimiropoulos, "Combining residual networks with LSTMs for lipreading," in *Proc. Interspeech*, 2017, pp. 3652–3656.
- [78] Y. Luo and N. Mesgarani, "Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 8, pp. 1256–1266, Aug. 2019.
- [79] Z. Pan, Z. Luo, J. Yang, and H. Li, "Multi-modal attention for speech emotion recognition," in *Proc. Interspeech*, 2020, pp. 364–368.
- [80] G. J. Brown and M. Cooke, "Computational auditory scene analysis," *Comput. Speech Lang.*, vol. 8, no. 4, pp. 297–336, 1994.
- [81] A. van den Oord et al., "WaveNet: A generative model for raw audio," in *Proc. 9th ISCA Speech Synth. Workshop*, 2016, pp. 125–125.
- [82] Y. Luo, Z. Chen, and T. Yoshioka, "Dual-path RNN: Efficient long sequence modeling for time-domain single-channel speech separation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 46–50.
- [83] N. Zeghidour and D. Grangier, "Wavesplit: End-to-end speech separation by speaker clustering," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 2840–2849, 2021.

- [84] M. Delcroix, J. B. Vázquez, T. Ochiai, K. Kinoshita, Y. Ohishi, and S. Araki, "SoundBeam: Target sound extraction conditioned on sound-class labels and enrollment clues for increased performance and continuous learning," 2022, *arXiv:2204.03895*.
- [85] S. Wisdom, E. Tzinis, H. Erdogan, R. Weiss, K. Wilson, and J. Hershey, "Unsupervised sound separation using mixture invariant training," *Adv. Neural Inf. Process. Syst.*, vol. 33, pp. 3846–3857, 2020.
- [86] T. Afouras, J. S. Chung, and A. Zisserman, "Deep lip reading: A comparison of models and an online application," in *Proc. Interspeech*, 2018, pp. 3514–3518.
- [87] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *Proc. Interspeech*, 2018, pp. 1086–1090.
- [88] G. Wichern et al., "Wham!: Extending speech separation to noisy environments," in *Proc. Interspeech*, 2019, pp. 1368–1372.
- [89] C. Busso et al., "IEMOCAP: Interactive emotional dyadic motion capture database," *Lang. Resour. Eval.*, vol. 42, no. 4, 2008, Art. no. 335.



Zexu Pan (Graduate Student Member, IEEE) received the B.Eng. degree in electrical and electronic engineering from Nanyang Technological University, Singapore, in 2019. He is currently working towards the Ph.D. degree with the Institute of Data Science, National University of Singapore. His research interests include multi-modal speech extraction, speech enhancement, speaker diarization, and multi-talker robust automatic speech recognition.



Meng Ge received the B.Eng. degree in software engineering from Tianjin Polytechnic University, Tianjin, China, in 2015, and the M.Eng. degree in software engineering from the Tianjin University, Tianjin, China, in 2017. He is currently a Ph.D. candidate at Tianjin Key Laboratory of Cognitive Computing and Application, College of Intelligence and Computing, Tianjin University, Tianjin, China. His research interests include source separation, speech extraction, speech enhancement, multi-talker robust automatic speech recognition.



Haizhou Li (Fellow, IEEE) received the B.Sc., M.Sc., and Ph.D. degrees in electrical and electronic engineering from South China University of Technology, Guangzhou, China in 1984, 1987, and 1990, respectively. He is currently the Presidential Chair Professor with the School of Data Science, The Chinese University of Hong Kong, Shenzhen, China. He is also an Adjunct Professor with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore. Prior to that, he was the Visiting Professor with CRIN,

France, during 1994–1995, Research Manager with the Apple-ISS Research Centre, during 1996–1998, Research Director with Lernout and Hauspie Asia Pacific, during 1999–2001, Vice President with InfoTalk Corp. Ltd., during 2001–2003, and the Principal Scientist and Department Head of Human Language Technology with the Institute for Infocomm Research, Singapore, during 2003–2016. His research interests include automatic speech recognition, speaker and language recognition, natural language processing. He was the recipient of the National Infocomm Award 2002, and President's Technology Award 2013 in Singapore. Dr Li was the Editor-in-Chief of *IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH AND LANGUAGE PROCESSING* (2015–2018), Member of the Editorial Board of *Computer Speech and Language* since 2012, Elected Member of IEEE SPEECH AND LANGUAGE PROCESSING TECHNICAL COMMITTEE (2013–2015), President of the *International Speech Communication Association* (2015–2017), President of *Asia Pacific Signal and Information Processing Association* (2015–2016), and President of *ASIAN FEDERATION OF NATURAL LANGUAGE PROCESSING* (2017–2018). He was the General Chair of *ACL* 2012, *INTERSPEECH* 2014, *ASRU* 2019 and *ICASSP* 2022. He is also the Fellow of the ISCA, and Fellow of the Academy of Engineering, Singapore. He was named Nokia Visiting Professor in 2009 by the Nokia Foundation, U Bremen Excellence Chair Professor in 2019, and DFG Mercator Fellow in 2022.