

Selective Listening by Synchronizing Speech With Lips

Zexu Pan , Ruijie Tao, Chenglin Xu , *Member, IEEE*, and Haizhou Li , *Fellow, IEEE*

Abstract—A speaker extraction algorithm seeks to extract the speech of a target speaker from a multi-talker speech mixture when given a cue that represents the target speaker, such as a pre-enrolled speech utterance, or an accompanying video track. Visual cues are particularly useful when a pre-enrolled speech is not available. In this work, we don't rely on the target speaker's pre-enrolled speech, but rather use the target speaker's face track as the speaker cue, that is referred to as the auxiliary reference, to form an attractor towards the target speaker. We advocate that the temporal synchronization between the speech and its accompanying lip movements is a direct and dominant audio-visual cue. Therefore, we propose a self-supervised pre-training strategy, to exploit the speech-lip synchronization cue for target speaker extraction, which allows us to leverage abundant unlabeled in-domain data. We transfer the knowledge from the pre-trained model to the attractor encoder of the speaker extraction network. We show that the proposed speaker extraction network outperforms various competitive baselines in terms of signal quality, perceptual quality, and intelligibility, achieving state-of-the-art performance.

Index Terms—Multi-modal, self-enrollment, speaker embedding, speech-lip synchronization, target speaker extraction, time-domain.

I. INTRODUCTION

HUMANS have a remarkable ability to focus attention on a particular speech signal in the presence of multiple noise

Manuscript received May 27, 2021; revised September 23, 2021 and November 22, 2021; accepted January 21, 2022. Date of publication February 24, 2022; date of current version May 12, 2022. This work was supported in part by the Agency for Science, Technology and Research (A*STAR) through its AME Programmatic Funding Scheme under Grant A18A2b0046, in part by the Science and Engineering Research Council, A*STAR, Singapore, through the National Robotics Program under Human-Robot Interaction Phase 1 under Grant 192 25 00054, in part by Deutsche Forschungsgemeinschaft (DFG) (German Research Foundation) through Germany's Excellence Strategy (University Allowance, University of Bremen) under Grant EXC 2077, and in part by The Chinese University of Hong Kong, Shenzhen under Grants UDF01002333 and UF02002333. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Mathew Magimai Doss. (*Corresponding author: Chenglin Xu.*)

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

Ruijie Tao is with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore 119077 (e-mail: ruijie.tao@u.nus.edu).

Chenglin Xu is with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore 119077, and also with Kuaishou Technology, Shenzhen 518063, China (e-mail: xuchenglin28@gmail.com).

Haizhou Li is with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore 119077, also with The Chinese University of Hong Kong, Shenzhen 518172, China, and also with the University of Bremen, 28359 Bremen, Germany (e-mail: haizhou.li@u.nus.edu).

Digital Object Identifier 10.1109/TASLP.2022.3153258

sources and competing background speakers [1]. The speaker extraction algorithm mimics human selective attention to extract only the target speaker's speech in such an adverse acoustic environment, which is also referred to as the cocktail party problem [2].

The human brain has a limited capacity to process all sensory stimuli perceived. Instead, human cognition relies on an inherent attention mechanism to focus the neural resources according to the contingencies of the moment [3]. Such attention can be categorized into a top-down and bottom-up process. Top-down attention refers to the voluntary allocation of neural resources based on guidance such as a specific goal or a reference. Bottom-up attention is driven by external stimuli that are salient because of their inherent properties relative to the background [4]. In this paper, we aim to mimic human visual top-down attention during listening to solve the single-channel target speaker extraction problem. This is a non-trivial, but crucial step for other downstream tasks such as speaker recognition [5]–[9], automatic speech recognition [10], sound source localization [11], and voice conversion [12]–[14].

Humans separate a speech mixture by attending to the salient prosody of individual speech tracks. Speech separation represents one direction of research towards machine intelligence of such bottom-up auditory attention, which recently has seen major progress [15]–[23]. The formulation of speech separation usually requires the number of speakers to be known in advance, which limits the scope of real-world applications. The formulation of speaker extraction avoids such limitation, that uses a reference signal to form the top-down attention, which is also referred to as the attractor in this paper. It proves to be effective in dealing with an unknown number of speakers. Most speaker extraction algorithms make use of auditory stimuli such as a reference speech signal to characterize the speakers. The reference speech signal is encoded as a speaker embedding in advance. In this way, the speaker extraction algorithm extracts speech that sounds similar to the reference speech signal [24]–[36].

As the above-mentioned speaker extraction techniques use the speaker embedding as the attractor in the neural solutions, they require pre-enrollment of a reference speech signal, which sometimes is not available in practice, for example, when we would like a robot to pay auditory attention to a passer-by. Unlike the reference speech signal, the real-time video recording can be easily obtained during human-computer interactions. In this work, we seek to develop a speaker extraction solution, which doesn't require a pre-enrolled speech signal as the reference, but assumes that the target speaker's face is visually present. We use

the target speaker’s face track as the auxiliary reference for the target speaker extraction.

Human attention is multi-modal [37], through a variety of sensory systems such as audition, vision, touch, and smell. These multi-modal stimuli are processed in human nervous systems in an interactive manner, which is referred to as the *reentry* [38]. The *reentry* theory describes the bidirectional exchange of signals along reciprocal axonal fibers linking two or more brain areas, that are temporally correlated, and can educate each other. The human brain has a cortex to process the audio and visual stimuli together [39]. In a cocktail party, visual cues are not corrupted by background noise, reverberation, and interference speech thus providing a robust attractor for auditory attention [40]. The studies in neuroscience show that speech comprehension is improved by looking at the speaker in adverse conditions [41], [42], leveraging on the long-term cross-modal temporal integration. Inspired by these findings, we aim to emulate the human visual top-down attention in a computational model for speaker extraction, named the *reentry* model after the *reentry* theory by Edelman [38].

In human selective listening, there are typically three types of visual cues, namely face-voice association [43]–[45], viseme-phoneme mapping [46]–[50], and speech-lip synchronization [51]–[55]. The face-voice association explores the general correlation between a human face and its voice signatures such as gender, age, and nationality. The viseme-phoneme mapping defines the relationship between the shapes of lips and mouth with respect to the sounds they represent, i.e., phonemes [56]. Among the visual cues, we advocate that the temporal synchronization between speech and lip movements would be the most informative, which will be the focus of this paper. The study on speech-lip temporal synchronization is a departure from the time-domain audio-visual speaker separation [47] and the multi-modal speaker extraction [57] model, where a supervised pre-training strategy is employed to explore the viseme-phoneme mapping cue for speaker extraction. To emulate human visual top-down attention during listening in the cocktail party scenario, we explore the speech-lip synchronization in a multi-talker setting with a pre-trained network, named the speech-lip synchronization (SLSyn) network. We propose a self-supervised training strategy for the SLSyn network such that the learning of speech-lip synchronization could leverage abundant unlabeled training data that are in the same domain as the speaker extraction task. We transfer the pre-trained knowledge from the SLSyn network to the *reentry* model, which performs ‘listening with ears and eyes’. The proposed *reentry* model outperforms various competitive baselines and achieves the state-of-the-art performance on VoxCeleb2 [58] dataset mixtures, in terms of signal quality, perceptual quality, and intelligibility. The *reentry* model also shows its robustness in cross-dataset evaluations and visual occlusion evaluations.

The rest of the paper is organized as follows. In Section II, we introduce the related works on target speaker extraction. In Section III, we formulate the proposed *reentry* model. In Section IV, we describe the experimental setup. In Section V, we report the results. Finally, Section VI concludes the study.

II. REFERENCE SIGNALS FOR SPEAKER EXTRACTION

Neural speaker extraction typically requires a reference signal to form the top-down attention towards the voice of the target speaker in a speech mixture. Studies in neuroscience suggest that either auditory [59] or visual stimulus [41], [42] can serve as a reference cue in human selective listening.

A. Pre-Enrolled Auditory Stimuli

Each speaker has a unique voice signature, which can be characterized by a fixed dimensional vector, referred to as speaker embedding, such as i-vector [60], x-vector [61], d-vector [62], and other similar feature representations [63].

It is common that a speaker embedding derived from the target speaker is employed to form the attractor for speaker extraction. VoiceFilter [26] and TseNet [64] are such examples, where they pre-train a speaker encoder, which is then used to derive d-vector or i-vector from the reference speech signal, to extract the target speaker. SpEx [24] and SpEx+ [25] further the idea by introducing a multi-task learning strategy, to jointly train the speaker encoder with the speaker extractor network, thus benefiting from the improved speaker characterization [26], [64] as well as the task-oriented optimization [28], [65], [66].

B. Self-Enrolled Auditory Stimuli

Despite much success, neural speaker extraction with speaker pre-enrollment is not practical in some scenarios. This prompts us to look into solutions without the need for a pre-enrolled reference speech signal. A recent study [67] on audio-visual speaker extraction seeks to use a visual reference to replace an audio reference. In this study, the visual reference is employed to extract the target speech first, which is subsequently encoded as a speaker embedding to serve as the reference whenever the visual cues are occluded. While such a two-pass mechanism improves in the absence of visual cues, the self-enrolled speaker embedding does not contribute when visual cues are adequately available. This suggests that the self-enrolled speaker embedding is not as informative as the visual reference, partly due to the fact that the two-pass mechanism is not jointly optimized.

MuSE [57] furthers the idea of using a visual reference in audio-visual speaker extraction. It self-enrolls the speaker’s voice signature from the intermediate extracted target speech on the fly. MuSE leverages the visual reference as well as the self-enrolled speaker embedding in an integrated manner, which shows that both of them actively contribute to the speaker extraction task. AVSE [68] is another method that self-enrolls the speaker’s voice signature directly from the speech mixture. AVSE is a frequency-domain implementation while MuSE takes a time-domain approach. This paper is motivated by the findings in the prior studies on the effective use of visual reference to self-enroll a speaker.

C. Face-Voice Association

There exist general correlations between the human face and voice. Prior studies [43], [44] have explored such knowledge for

speaker extraction. They employ a speaker encoder that is pre-trained to learn the cross-modal correspondence and to encode a face image into a speaker embedding as the attractor. It is noted that speaker extraction systems that rely on such general correlation knowledge generally don't outperform those with the pre-enrolled auditory stimuli.

D. Viseme-Phoneme Mapping

Phonemes are the smallest units of spoken language, and the visual equivalent are the visemes. Generally speaking, a viseme corresponds to a set of phonemes that have identical appearances on the lips. Therefore, one is able to associate a phoneme with a viseme class, but a viseme may map to multiple phonemes. The viseme-phoneme mapping is usually trained from a visual speech recognition task, which aligns the input viseme sequence with the output phoneme sequence or word sequence [69]–[71].

Some prior studies have effectively made use of such viseme-phoneme mapping in speaker extraction, i.e. TDSE [47] and MuSE [57]. They transfer part of the pre-trained visual speech recognition network as the visual encoder of the speaker extraction network, where the visual encoder encodes the input viseme sequence to a viseme embedding sequence. The viseme embeddings are used to form the attention attractor. As the visual encoder is optimized for the phoneme recognition task, therefore, not the best for speech reconstruction, we are prompted to look into a better way to train the visual encoder for speaker extraction. Furthermore, the training of visual speech recognizer requires video data with text transcription, which is scarce. We are motivated to leverage the audio-visual information that exists in abundantly available parallel audio-visual data for self-supervised learning, without the need for text transcription.

Motivated by a study that face embeddings encoded by face recognition models can recover facial expressions [72], 'looking to listen at the cocktail party' [73] employs viseme embeddings which are encoded by a face recognition model as the attractor for speaker extraction. However, the face embeddings are trained to discriminate faces. It is neither an optimal representation for viseme-phoneme mapping nor audio-visual synchronization.

E. Audio-Visual Synchronization

Some studies have implemented the use of audio-visual synchronization cue for speaker extraction [52], [53], where they encode the raw visual frames or facial landmarks into a sequence of visual embeddings as the reference signal.

Others [74], [75] further the idea by pre-training the visual encoder of the speech separation network on audio-visual synchronization tasks. [74] pre-trains the visual encoder on the scene synchronization task which finds that the scene motion is temporally synchronized with audio events. [75] pre-trains the visual encoder on the object synchronization task which detects the object that is responsible for the sound in a video. All studies point to the direction that the scene and object synchronization features encoded by the visual encoder could be employed as a visual attractor in speech separation networks. However, the audio-visual synchronization for speech to lips in a multi-talker scenario is not as prominent as that for whistling sound to a train.

Furthermore, the audio-object synchronization features are not pre-trained to disentangle overlapped speech.

The success of the effective use of the audio-visual synchronization cue prompts us to look into speech-lip synchronization that is a more direct and informative cue for speaker extraction. Humans have the experience to improve listening by observing the lips when someone is talking. In this paper, we would like to focus on the use of the speech-lip synchronization cue for speaker extraction.

III. AUDIO-VISUAL SPEAKER EXTRACTION NETWORK

We propose an audio-visual speaker extraction network using the speech-lip synchronization cue, as depicted in Fig. 1, that is called the *reentry* model. Let $x(\tau)$ be a multi-talker speech waveform in the time-domain, which consists of the target speaker's speech $s(\tau)$ and other interference speech $b_i(\tau)$:

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

We would like to extract $\hat{s}(\tau)$ that approximates $s(\tau)$.

In the *reentry* model: 1) The speech encoder transforms $x(\tau)$ into a sequence of spectrum-like frame-based embeddings $X(t)$, referred to as the mixture speech embeddings.¹ 2) The attractor encoder encodes $x(\tau)$ and its accompanying face track images $v(t)$ into a sequence of speech-lip synchronization embeddings $V(t)$ to form an attractor signal. 3) The speaker encoders encode the intermediate estimated speech embeddings into a speaker embedding A^r , which represents the attended speaker's voice signature. 4) The speaker extractor takes in $V(t)$ and A^r as top-down attention attractors to estimate a mask $M(t)$ for $X(t)$ which only lets the target speaker pass. The extracted speech embeddings $\hat{S}(t)$ are estimated by element-wise multiplication between the estimated mask $M(t)$ and the mixture speech embeddings $X(t)$:

$$\hat{S}(t) = X(t) \otimes M(t) \quad (2)$$

5) Finally, the speech decoder transforms the extracted speech embeddings $\hat{S}(t)$ into a time-domain waveform $\hat{s}(\tau)$.

A. Self-Supervised Speech-Lip Synchronization Learning

Traditional approaches have been devoted to using statistical theory to learn audio-visual synchronization [76], [77]. With deep learning, audio-visual synchronization learning has seen major progress with self-supervised learning, where a network learns from positive and negative samples easily derived from found data. The positive samples could be the videos with the original soundtrack, while the negative samples are the videos with a random time-shifted soundtrack [74], [75], [78].

During learning, one common approach is to encode audio and visual signals into embedding sequences separately, and compare the similarity of the two embedding sequences [75], [79], [80]. It makes the decision on short video segments, and

¹In this paper, a variable with τ as the index represents a sequence of samples in the time-domain, while a variable with t as the index represents a sequence of frame-based embeddings.

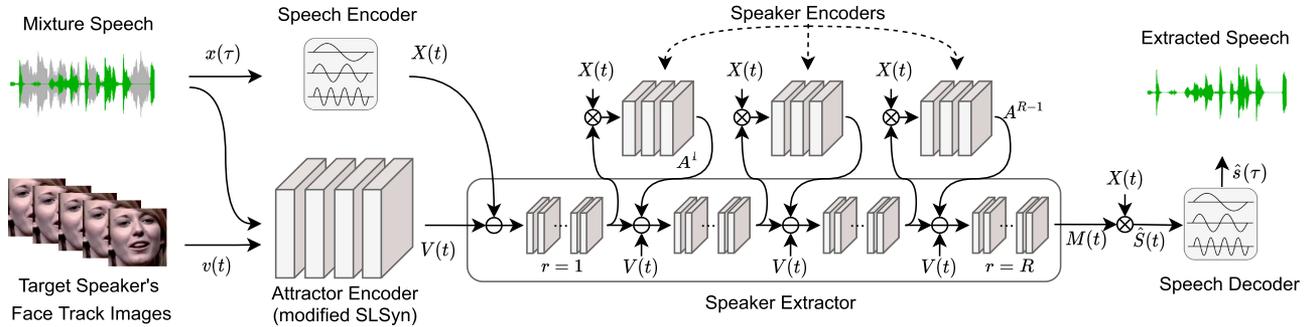


Fig. 1. The proposed audio-visual speaker extraction network, named the *reentry* model. The attractor encoder derives feature representations for speech-lip synchronization; the speaker extractor produces a receptive mask for the target speaker dynamically; the speaker encoders characterize the target speaker and provide the top-down attention to the speaker extractor. The symbol \oplus represents the concatenation of features along channel dimension in convolutional layers; the symbol \otimes represents element-wise multiplication.

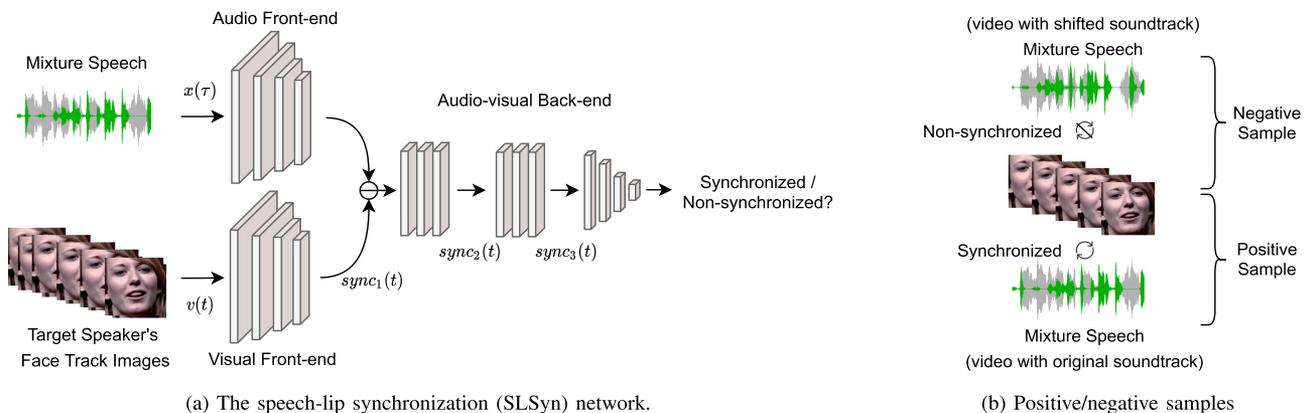


Fig. 2. The proposed speaker-independent speech-lip synchronization network, that learns to extract speech-lip synchronization embeddings. In (a), the SLSyn network takes the face track and the soundtrack as inputs and makes a synchronization decision. Part of the pre-trained SLSyn network is used as the attractor encoder in Fig. 1. In (b), we illustrate a positive (synchronized) and a negative (non-synchronized) audio-visual sample. The green waveform for a target speaker in the synchronized sample is shifted to form a negative sample.

doesn't make good use of the temporal correlation between modalities. Another idea is to fuse the audio and visual embedding sequences, that are extracted separately, through an audio-visual network such as a convolutional neural network [74] or a long short-term memory (LSTM) network [81] to capture their temporal dependencies.

Motivated by the second idea, we propose a speech-lip synchronization (SLSyn) network to detect the presence of synchronization between a multi-talker speech signal and a face track. As the SLSyn network is optimized for speech-lip synchronization detection, it is expected to derive speech-lip synchronization embeddings that characterize the audio-visual joint events. The SLSyn network is pre-trained with multiple speakers to be a speaker-independent model.

As shown in Fig. 2(a), the SLSyn network consists of an audio front-end to extract audio features from the mixture speech $x(\tau)$, a visual front-end to extract visual features from the target speaker's face track images $v(t)$, and an audio-visual back-end to fuse the concatenated audio and visual features for an utterance level binary classification. As the temporal convolutional neural network (TCN) is effective in capturing long-term dependencies [16], [82] in speech processing [83], [84], we adopt a TCN

structure as the audio front-end. We design a visual front-end with a 3D convolutional layer and a ResNet to encode raw images into visemes similar to [85], representing the appearance of the lips. The audio and visual features are concatenated at the frame level, and taken by the TCN-based audio-visual back-end.

The activations at the intermediate layers of the SLSyn network represent different levels of abstraction. We study three different feature representations as the speech-lip synchronization embeddings, namely $sync_1(t)$, $sync_2(t)$, and $sync_3(t)$, as shown in Fig. 2(a). The three feature representations have the same time-resolution as the target speaker's face track.

As illustrated in Fig. 2(b), the original videos with speech-lip synchronized soundtracks serve as the positive samples, while videos with randomly shifted soundtracks serve as the negative samples. An additional interference speech signal is added to both the positive and negative samples to simulate the cocktail party scenario.

We minimize the following binary cross-entropy (BCE) loss for the SLSyn network training:

$$\mathcal{L}_{BCE} = -y \log(\hat{y}) - (1 - y) \log(1 - \hat{y}) \quad (3)$$

where $y \in \{0, 1\}$ indicates whether the input video has a shifted soundtrack while \hat{y} is the predicted probability of the video and soundtrack being synchronized at the end of the SLSyn network.

This learning is similar to the noise-contrastive estimation [86], where the noisy samples in the SLSyn network training are asynchronous videos. As human-annotated labels are not needed, the pre-training can be done on abundantly available videos, also in the same domain as the speaker extraction algorithms to minimize the domain mismatch.

B. Speaker Extraction Using the Speech-Lip Synchronization Cue

We now formulate the *reentry* model, as illustrated in Fig. 1. The *reentry* model consists of a speech encoder, an attractor encoder, a speaker extractor, a number of speaker encoders, and a speech decoder. The speaker extractor has an interlace structure between a number of speaker encoders and speaker extraction modules.

1) *Speech Encoder*: The speech encoder takes a time-domain approach [16] to encode the input speech waveform $x(\tau)$ into a spectrum-like speech embedding sequence $X(t)$. It consists of a 1D convolution *Conv1D*, with channel size N , kernel size L , and stride $L/2$. The *Conv1D* is followed by a rectified linear activation *ReLU*,

$$X(t) = \text{ReLU}(\text{Conv1D}(x(\tau), N, L, L/2)) \quad (4)$$

where N is also the speech embedding dimension. As $x(\tau)$ is a time-domain signal, the 1D convolution behaves like a frequency analyzer [16].

2) *Attractor Encoder*: The attractor encoder seeks to encode the synchronized sequence of images $v(t)$ with the mixture speech waveform $x(\tau)$ into a sequence of embeddings $V(t)$ representing the speech-lip synchronization. The SLSyn network is trained just to do that.

We adopt part of the pre-trained SLSyn network, followed by several adaptation layers to form the attractor encoder. The speech-lip synchronization embeddings extracted from the pre-trained SLSyn network, namely $\text{sync}_1(t)$, $\text{sync}_2(t)$, and $\text{sync}_3(t)$, are not optimized for speaker extraction directly, therefore, they have different properties than those from the speaker extractor [87]. The adaptation layers adapt the speech-lip synchronization embeddings towards the speaker extraction task. We follow [46], [47] to design a stack of TCNs in the adaptation layers.

3) *Speaker Extractor*: Receptive masks have been well studied in the source separation literature to mimic human selective attention, such as ideal binary mask [88], ideal ratio mask [89], ideal amplitude mask [90], Wiener-filter like mask [91], and phase sensitive mask [91]. We adopt the masking methods [16], [25], [47] to estimate a receptive mask $M(t)$ for the speech of the target speaker as shown Fig. 1.

Similar to [16], [24], [47], we use the TCN structure for the speaker extractor to capture the long-range dependency of the speech signal. The speaker extractor consists of R repeated stacks of TCNs in a speaker extraction pipeline. We illustrate the reentry model with $R = 4$ in Fig. 1. An intermediate mask

$M^r(t)$ is estimated after the r th stack. Each stack consists of B TCNs with an exponentially growing dilation factor 2^b ($b = 0, \dots, B - 1$). With the pipelined stacks, the receptive field of the network increases exponentially, and better models the long-term temporal dependencies of speech signals [16]. Hence, the masks produced by the TCN stacks are expected to be progressively refined along the speaker extraction pipeline.

The *reentry* theory suggests that in the human brain, temporally correlated signals educate each other through bi-directional exchanges. Following the *reentry* theory, we propose to make use of the interactive evidence between the speech and lips on the fly as we decode the target speech.

The speaker extractor takes the mixture speech embeddings $X(t)$ as input, and the top-down attention cues as the condition, to extract the clean speech from the target speaker. Besides the speech-lip synchronization cue $V(t)$, we also consider the self-enrolled speaker embedding A^r as another top-down attention cue. The top-down attention cues modulate with the input mixture speech to estimate the receptive masks. The study in SpEx [24] shows that concatenating the attention stimulus at the start of every TCN stack effectively guides the mask estimation. Bringing the idea forward to the audio-visual context, we concatenate A^r and $V(t)$ with $M^r(t)$ to form the input of each TCN stack, except for the first TCN stack which we concatenate $V(t)$ with $X(t)$ as input. The A^r and $V(t)$ are up-sampled along the time dimension to match the temporal resolution of $M^r(t)$ before concatenation. We will discuss the self-enrolled speaker encoders next.

4) *Speaker Encoders*: A speaker encoder derives an utterance level speaker embedding A^r a.k.a self-enrolled speaker embedding from the intermediate estimated speech embeddings $\hat{S}^r(t)$ after the r th TCN stack in the speaker extraction pipeline. The $\hat{S}^r(t)$ is defined as:

$$\hat{S}^r(t) = X(t) \otimes M^r(t) \quad (5)$$

Since $M^r(t)$ becomes available only after the first TCN stack, we apply the speaker encoders starting from the second TCN stack. We concatenate A^r with $M^r(t)$ in the same way as $V(t)$. As $\hat{S}^r(t)$ varies along the speaker extraction pipeline, we design $R - 1$ speaker encoders of the same architecture, but with distinctive weights that are trained independently of each other, as opposed to weight sharing in SpEx [24].

5) *Speech Decoder*: The speech decoder is an inverse function of the speech encoder, that takes the extracted speech embeddings $\hat{S}(t)$ as input and reconstructs the time-domain waveform $\hat{s}(\tau)$ as output. We formulate the speech decoder as a linear layer followed by an overlap and add *OnA* operation [92]:

$$\hat{s}(\tau) = \text{OnA}(\text{Linear}(\hat{S}(t), N, L), L/2) \quad (6)$$

where the *Linear* layer has input and output feature dimension of N and L respectively, and the overlap and add operation has a frame shift of $L/2$.

C. Multi-Task Learning

To ensure that the speaker embedding A^r is encoded for the purpose of speaker extraction, we train the $R - 1$ speaker

encoders together with the speaker extractor network, with two training objectives, i.e., the scale-invariant signal-to-noise ratio (SI-SDR) loss [93] to measure the signal quality, and the cross-entropy (CE) loss for speaker classification.

The loss \mathcal{L}_{SI-SDR} measuring signal quality is shown in (7), which is applied to the output of the *reentry* model between the extracted speech $\hat{s}(\tau)$ and clean speech $s(\tau)$. We omit the subscript (τ) for brevity:

$$\mathcal{L}_{SI-SDR} = -10 \log_{10} \left(\frac{\|\langle \hat{s}, s \rangle_s\|^2}{\|\hat{s} - \langle \hat{s}, s \rangle_s\|^2} \right) \quad (7)$$

The loss \mathcal{L}_{CE}^r that accounts for speaker classification is shown in (8), which is applied to every speaker encoder's output A^r of the $R - 1$ speaker encoders:

$$\mathcal{L}_{CE}^r = - \sum_{c=1}^C p[c] \log \hat{p}^r[c] \quad (8)$$

where $p[c] = 1$ for target speaker, and $p[c] = 0$ otherwise, $\hat{p}^r[c] = \text{softmax}(W^r A^r)[c]$ is the predicted probability for speaker c , C is the number of speakers in the dataset, and W^r is a learnable weight matrix specific to each speaker encoder for speaker classification.

The overall loss \mathcal{L}_{total} is defined as:

$$\mathcal{L}_{total} = \mathcal{L}_{SI-SDR} + \gamma \sum_{r=1}^{R-1} \mathcal{L}_{CE}^r \quad (9)$$

where γ is a scaling factor.

Now that each speaker encoder is constrained with a speaker classification loss, it is expected to derive a speaker embedding A^r that best represents the target speaker's voice signature, helping the speaker extractor in the next TCN stack. Besides, the gradient of the speaker classification loss is passed from the speaker encoders to the stacks of TCNs in the multi-task learning framework. Hence, the speaker extractor is also trained to extract the target speaker's speech containing the target speaker's voice signature, besides the signal extraction quality. Therefore, the speaker self-enrollment is progressively refined with the speaker extractor with the interlaced structure, ensuring the convergence of speaker extraction.

It is worth noting that A^r is derived from the mixture speech and visual signals without the supervision of speaker identity at run-time inference. Therefore, A^r is also called the self-enrolled speaker embedding.

D. Relationship With MuSE

MuSE [57] employs a pre-trained visual speech recognition network as the attractor encoder and exploits the viseme-phoneme mapping to form the attention attractor. The main difference between the *reentry* model and MuSE lies in two aspects: 1) The *reentry* model employs a self-supervised pre-training strategy for the SLSyn network that doesn't require text transcription of the soundtrack as the visual speech recognition network does. This brings us multiple advantages. We can now easily adapt the SLSyn network on abundantly available domain data to reduce domain mismatch. 2) The SLSyn network learns

the speech-lip synchronization rather than the viseme-phoneme mapping. The latter is not a direct cue for overlapped speech disentanglement and signal reconstruction.

While the *reentry* model is different from the MuSE framework, it shares some ideas that are studied in MuSE and other related work, that is to exploit a self-enrollment strategy, and to employ a multi-stage progressive refinement pipeline for speaker extraction. We expect that the progressive refinement architecture will outperform a single-stage extraction [94]. The findings in this paper certainly reinforce this line of thought.

IV. EXPERIMENTAL SETUP

A. Experiment Data

1) *Speaker Extraction*: We simulate a two-speaker mixture dataset (VoxCeleb2-2mix) and a three-speaker mixture dataset (VoxCeleb2-3mix) from the VoxCeleb2 dataset [58] for a series of systems evaluations. The VoxCeleb2 dataset contains over 1 million face track videos for 6112 celebrities extracted from YouTube videos. The face track videos are obtained using a face detector as in [95], and a tracker that is based on ROI overlap [96]. As the face track video comes with a soundtrack, each utterance is thus temporally synchronized with the speaker's face track.

To simulate the mixture datasets, we remove utterances that are shorter than 4 seconds, and randomly select 40000 clean utterances from 800 speakers in the original VoxCeleb2 train set to create our train set (20000 utterances); we randomly select another 8000 clean utterances from the same 800 speakers in the original train set to create our validation set (5000 utterances); we use the 36237 clean utterances from 118 speakers in the original test set to form our test set (3000 utterances). The interference speech is mixed with the target speech at a random Signal-to-Noise (SNR) ratio between 10 dB to -10 dB. The long utterance is truncated to the length of the short utterance when creating a mixture speech. The sampling rate of the audio data is 16 kHz, and the face track video stream is available at 25 FPS. We use a sequence of sampled face track images from the target speaker's utterance as the visual input to our proposed network.

The train set is used for the network training, and the validation set is used for optimizing the network configurations. The utterances and speakers in the test set do not overlap with those in the train set and validation set, which allows us to perform speaker-independent experiments using the test set.

2) *Synchronization Detection*: To pre-train the speaker-independent speech-lip synchronization network, we create a multi-talker dataset (VoxCeleb2-sync) from the original VoxCeleb2 dataset [58]. We create 4 million video clips to form a train set, and 40 thousand video clips to form a validation set. They range from 1 to 4 seconds. A quarter of the video clips are with clean speech and three-quarters of them are with simulated two-speaker mixture speech. The simulation follows the same protocol as that of the VoxCeleb2-2mix dataset.

The video clips with the original soundtrack are the positive samples that have synchronized speech and lip movements. The negative samples have a randomly shifted soundtrack by 0.2 to 1 seconds relative to the visual sequence. Besides the

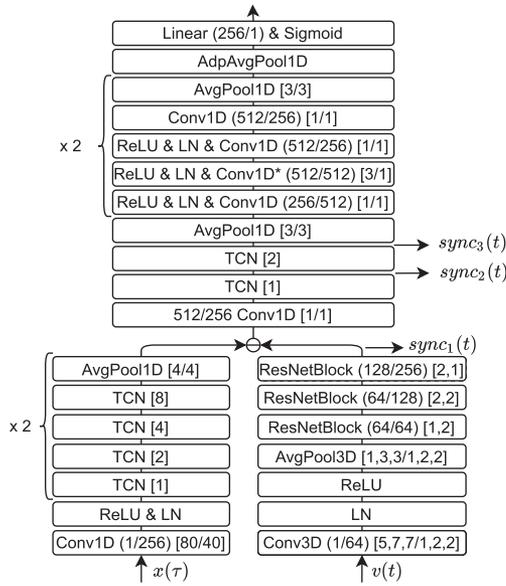


Fig. 3. The configuration of the SLSyn network. Throughout the paper, $Conv1D$, $Conv2D$, and $Conv3D$ are 1, 2, and 3-dimensional convolutional layers respectively, and their parameters are represented as $(in_channels/out_channels)[kernel_size/stride]$, $Conv1D^*$ is group convolution, with group size $in_channels$, LN is layer normalization, $AvgPool1D$ and $AvgPool3D$ are 1 and 3-dimensional average pooling layers respectively, and their parameters are represented as $[kernel_size/stride]$, $AdpAvgPool1D$ is the adaptive 1-dimensional average pooling layer, and $PReLU$ is parametric $ReLU$. The configuration of the TCN and the ResNetBlock are detailed in Fig. 4(a) and (b).

VoxCeleb2-sync dataset, we create a smaller version of the dataset named the VoxCeleb2-sync-s, which contains 10% of the video clips from the VoxCeleb2-sync dataset; and a clean dataset named the VoxCeleb2-sync-c, which is of the same size as the VoxCeleb2-sync, but all are clean speech utterances from the VoxCeleb2 dataset.

B. Model Configuration

The SLSyn network as shown in Fig. 2 consists of an audio front-end, a visual front-end, and an audio-visual back-end. The model configuration is detailed in Fig. 3. The last linear layer of the SLSyn network has an input channel size of 256 and an output channel size of 1, a $Sigmoid$ activation is used at the end for binary classification.

The reentry model as shown in Fig. 1 consists of a speech encoder, an attractor encoder, a speaker extractor, $R - 1$ speaker encoders, and a speech decoder. The attractor encoder has the same configuration as the SLSyn network as detailed in Fig. 3. The attractor encoder is followed by an adaptation layer of a stack of 5 TCN [1], with the TCN configuration detailed in Fig. 4(a). The $R - 1$ speaker encoders share the same architecture, and the configuration is detailed in Fig. 4(c). Each speaker encoder is preceded by a speech decoder and a speech encoder that have the shared weights with the speech encoder and speech decoder depicted in Fig. 1.

We use layer normalization instead of batch normalization for all models, such that a small batch size in the fine-tuning stage will not affect the normalization performance. The network

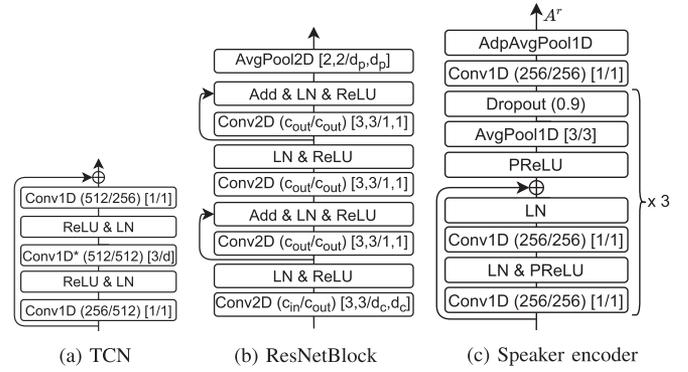


Fig. 4. (a) The configuration of a temporal convolutional network, TCN [d]. d is the stride in the $Conv1D^*$ layer, also referred to as the dilation factor. The symbol \oplus represents element-wise addition. (b) The configuration of a ResNetBlock $(c_{in}/c_{out})[d_c/d_p]$. Add refers to the element-wise addition. (c) The configuration of a speaker encoder.

parameters for the speech encoder, speaker encoders, speaker extractor, and speech decoder are selected following [47], [57]. In the speech encoder and speech decoder, the L and N are set to 40 and 256 respectively. In the speaker encoders, the $Dropout$ layer has a dropout probability of 0.9. The speaker extractor has the R and B set to 4 and 8 respectively.

C. Training of the Reentry Model

Both the SLSyn network and *reentry* model are implemented with PyTorch², and optimized by the Adam optimizer [97], with an initial learning rate of 0.001.

During the SLSyn network training, the learning rate decreases by 4% every epoch, and the training stops when the best validation loss does not improve for 4 consecutive epochs. During the *reentry* model training, the learning rate is halved when the best validation loss does not improve for 6 epochs consecutively, and the training stops when the best validation loss does not improve for 10 consecutive epochs. Both the SLSyn network and the *reentry* model are trained on 2 Tesla 32 GB V100 GPUs. The utterances are truncated to 6 seconds to fit into the GPU memory during training, while the full utterance is evaluated at run-time inference.

The overall training is carried out in three stages. In the first stage, we pre-train the SLSyn network to form the attractor encoder. In the second stage, we train the *reentry* model, which uses the pre-trained attractor encoder without updating its weights. In the last stage, we fine-tune the attractor encoder together with the *reentry* model using the re-initialized optimizer. We will justify the 3-stage strategy through an ablation experiment.

D. Baselines

Time-domain speaker extraction algorithms usually outperform their frequency-domain counterparts. MuSE [57] is a time-domain technique that reports the state-of-the-art performance

²The data generation and the *reentry* model codes are available at <https://github.com/zexupan/reentry>.

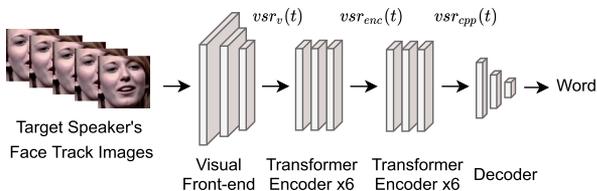


Fig. 5. Three features at different abstraction level derived from a visual speech recognition model, namely $v_{sr_v}(t)$, $v_{sr_{enc}}(t)$, and $v_{sr_{cpp}}(t)$.

in audio-visual speaker extraction, and is built on TDSE [47]. We implement the two models as the reference baselines.

1) *TDSE*: TDSE [47] has a network component and architecture similar to the *reentry* model, except that the former has neither the speaker encoders for top-down auditory attention nor the speaker classification task, which is designed to train the speaker encoders. In addition, TDSE uses a pre-trained VSR network as the attractor encoder. The implementation of TDSE-O follows the original architecture in the paper. We also implement TDSE-I as an improved version for a fair comparison with the *reentry* model. The improvements include, a) the TDSE-I replaces the batch normalization in TDSE-O with layer normalization, b) it concatenates the visual attractor to every stack of TCNs, just like what the *reentry* model does, instead of the second stack only, and c) the attractor encoder is fine-tuned similar to the *reentry* model.

2) *MuSE*: Like TDSE, MuSE [57] uses a VSR pre-trained network as the attractor encoder. We implement the original architecture, denoted as MuSE-O, and an improved version MuSE-I, with layer normalization and attractor encoder fine-tuning. MuSE-O takes the intermediate estimated speech embeddings $\hat{S}^r(t)$ and visual cue $V(t)$ as input to the speaker encoders, and applies a speaker encoder to the start of every TCN stack, while MuSE-I and *reentry* model only take $\hat{S}^r(t)$ as input to the speaker encoders starting from the second TCN stack similarly to the *reentry* model.

3) *Visual Speech Recognition*: TDSE and MuSE both use an attractor encoder pre-trained on the visual speech recognition task³. The attractor encoder generates $v_{sr_v}(t)$, which is the output of the visual front-end of the VSR-LSTM model [85]. The visual front-end consists of a convolutional 3D and a ResNet 18 layer. In a visual speech recognition study [98], this visual front-end is followed by 12 transformer layers to form a VSR-TRAN model as illustrated in Fig. 5, where $v_{sr_{enc}}(t)$ and $v_{sr_{cpp}}(t)$ are the outputs of the 6th and last layer of the transformer. $v_{sr_{cpp}}(t)$ represents the character posterior probability at the frame level. We study the three features, namely $v_{sr_v}(t)$, $v_{sr_{enc}}(t)$ and $v_{sr_{cpp}}(t)$, which represent different level of abstraction of the visual cues in our experiments.

E. Evaluation Metrics

We use the scale-invariant signal-to-noise ratio improvement (SI-SDRi); signal-to-noise ratio improvement (SDRi); perceptual evaluation of speech quality improvement (PESQi) [99]

³The visual speech recognition pre-trained models are taken from https://github.com/lordmartian/deep_avsr [69]

and Short Term Objective Intelligibility improvement (STOIi) to evaluate the extracted speech (the higher the better for all metrics). The SI-SDRi and SDRi represent the improvements in the signal quality of the extracted speech qualities compared to the mixture speech. The PESQi and STOIi represent the improvements in the perceptual quality and intelligibility of the extracted speech compared to the mixture speech.

V. EXPERIMENTS

A. Visual Speech Recognition vs. Speech-Lip Synchronization

We first compare the two pre-training models, namely the visual speech recognition (VSR) network [69] and our speech-lip synchronization (SLSyn) network, as reported in Table I. As VSR is a challenging task with limited training data, the VSR-TRAN network has a high word error rate of 61.8% on the LRS2 dataset.

For speech-lip synchronization detection, the accuracy is 86.6% with 270 hours of training data (VoxCeleb2-sync-s). If 2,700 hours of data are used (VoxCeleb2-sync), the accuracy goes up to 94.9%. If the dataset consists of clean speech only (VoxCeleb2-sync-c), the accuracy goes up to 96.8%. The number (no.) of network parameters for the SLSyn network is less than the VSR network due to the simplicity of the task. In the VoxCeleb2-sync dataset, negative samples are target-shifted speech. To ensure the SLSyn network is not learning the shift instead of the synchronization, we conduct an additional negative sample test to classify the talking face track with only the interference speaker's speech. The SLSyn network achieves the classification accuracy of 96.0%, showing that the SLSyn network is learning the speech-lip synchronization information.

We evaluate the *reentry* model on the VoxCeleb2-2mix dataset with different pre-trained attractor encoders in Table I. In Section V, we omit the subscript (t) for the attention attractors *Att* for brevity. With the VSR features, the lower level of abstraction is further away from the decision layer, and represents more the visual content than the phonetic information. We observe that the lower the feature abstraction level, the better the extracted signal quality, with v_{sr_v} having better results than $v_{sr_{enc}}$ and $v_{sr_{cpp}}$. The results suggest that the viseme-phoneme mapping objective, that optimizes $v_{sr_{cpp}}$, is not the best for speaker extraction.

With the SLSyn features, the higher level of abstraction is near the decision layer, and represents more the speech-lip synchronization/non-synchronization decision than the audio-visual content. We observe that the higher the abstraction level, the better the signal extraction quality produced, with $sync_3$ having better results than $sync_2$ and $sync_1$. When $sync_3$ is trained from the VoxCeleb2-sync-s dataset with a similar amount of data as VSR, it outperforms the best VSR feature v_{sr_v} by 0.55 dB in terms of SI-SDRi. The best synchronization feature $sync_3$ outperforms the best VSR feature v_{sr_v} by 0.69 dB in terms of SI-SDRi, showing the advantage of the speech-lip synchronization feature. Unless mentioned otherwise, the *reentry* model reported in this paper uses $sync_3$ pre-trained from the VoxCeleb2-sync dataset for the speech-lip synchronization feature representation.

TABLE I

WE REPORT THE EVALUATION OF THE PRE-TRAINED (PT) NETWORK VSR-TRAN IN TERMS OF WORD ERROR RATE (WER), AND THAT OF THE SLSYN NETWORK IN TERMS OF BINARY CLASSIFICATION ACCURACY (ACC.). WE ALSO EVALUATE THE *REENTRY* MODEL ON THE VOXCELEB2-2MIX DATASET WITH DIFFERENT TOP-DOWN ATTENTION ATTRACTORS (*Att*) DERIVED FROM THE PRE-TRAINED NETWORKS. THE VALIDATION SET (Val.) AND TEST SET SI-SDRi ARE REPORTED WITH (w/) AND WITHOUT (w/o) THE ATTRACTOR ENCODER FINE-TUNING. THE NO. OF NETWORK PARAMETERS (PARAMS) IS REPORTED IN MILLION (M)

PT network	PT Dataset	Hours	WER (%)	Acc. (%)	<i>Att</i>	Val. SI-SDRi (dB) w/o	Test SI-SDRi (dB) w/o	Val. SI-SDRi (dB) w/	Test SI-SDRi (dB) w/	Params (m)
VSR-LSTM	LRW [70]	160	-	-	<i>vsr_v</i>	12.29	11.88	12.30	11.91	24.5
VSR-TRAN	LRS2 [98]	200	61.8	-	<i>vsr_enc</i>	12.07	11.61	12.12	11.73	75.1
					<i>vsr_cpp</i>	11.25	10.91	11.51	11.12	130.1
SLSyn	VoxCeleb2-sync	2700	-	94.9	<i>sync1</i>	12.17	11.70	12.27	11.80	16.0
					<i>sync2</i>	12.43	12.14	12.70	12.34	18.5
					<i>sync3</i>	12.58	12.34	13.01	12.60	18.8
					a face image	3.19	2.92	3.29	3.00	18.8
SLSyn	VoxCeleb2-sync-s	270	-	86.6	<i>sync3</i>	12.27	12.00	12.87	12.46	18.8
	VoxCeleb2-sync-c	2700				12.33	11.91	12.66	12.26	18.8

TABLE II

WE STUDY THE TRAINING STRATEGY FOR THE *REENTRY* MODEL THAT INVOLVES THE ATTRACTOR ENCODER IN FOUR EXPERIMENTS. IN ALL EXPERIMENTS, WE TRAIN THE *REENTRY* MODEL OF THE SAME ARCHITECTURE WITH DIFFERENT VARIANTS OF THE PROPOSED TRAINING STRATEGY ON THE VOXCELEB2-2MIX DATASET, AND REPORT THE SI-SDRi FOR THE VALIDATION AND TEST SETS. *INIT* DENOTES THE USE OF PRE-TRAINED ATTRACTOR ENCODER; *FIX* DENOTES FREEZING THE PRE-TRAINED ATTRACTOR ENCODER; *FT* DENOTES FINE-TUNING OF THE ATTRACTOR ENCODER

Expt	<i>Att</i>	<i>Init</i>	<i>Fix</i>	<i>FT</i>	Val. SI-SDRi (dB)	Test SI-SDRi (dB)
1	<i>sync3</i>	✗	N/A	N/A	12.27	11.92
2		✓	✗	N/A	12.42	12.02
3		✓	✓	✗	12.58	12.34
4		✓	✓	✓	13.01	12.60

Among the synchronization features, the visual-only feature *sync1* lags behind audio-visual features *sync2* and *sync3* by about 0.5 dB in terms of SI-SDRi, showing the importance of the speech in the synchronization cue. When *sync3* is trained on the single-talker dataset (VoxCeleb2-sync-c), the SI-SDRi degrades by 0.34 dB, showing that the pre-training domain-mismatch affects the speaker extraction quality significantly.

The face-voice association can also be used as the top-down attention cue for speaker extraction [43]. We would like to investigate if the attractor encoder in the *reentry* model indeed benefits from the face-voice association or the speech-lip synchronization. We study the use of a single face image as the (*Att*) in Table I. During the *reentry* model training, we first randomly pick a face image from the video sequence and repeat the same image to form an image sequence. We then use the image sequence in place of the video sequence as the top-down attention cue, and report an SI-SDRi of 3.00 dB, which is far behind that of 12.60 dB compared with *sync3*. While the attractor encoder seems to take advantage of the face-voice association cue, it mainly benefits from the speech-lip synchronization cue, as most of SI-SDRi gain comes from the latter.

B. Study of Training Strategy

We perform an ablation study to justify the proposed 3-stage training strategy. In Table II, Expt 1 represents a training strategy

in which we train the entire *reentry* model from scratch without any pre-training; Expt 2 represents a training strategy that uses the pre-trained SLSyn network to initialize the attractor encoder, then performs a full *reentry* model training; Expt 3 represents a variant of Expt 2 that uses the pre-trained SLSyn network to initialize the attractor encoder, then performs a full *reentry* model training but freezing the attractor encoder; Expt 4 represents the proposed 3-stage training strategy, that is to initialize the attractor encoder with pre-trained SLSyn network, then fix the attractor encoder for model training, finally fine-tune the attractor encoder and the rest of the *reentry* model together. As the *sync3* feature from the SLSyn network has the best representation, we adopt *sync3* to justify our training strategy for the *reentry* model.

As shown in Table II, the proposed 3-stage training strategy (Expt 4) provides an SI-SDRi improvement of 0.68 dB over that without pre-training (Expt 1), which suggests that the *reentry* model benefits from the knowledge learned from the pre-training step. The proposed 3-stage training strategy (Expt 4) also outperforms that with a simple initialization (Expt 2) at retaining the transferred knowledge, with an SI-SDRi improvement of 0.58 dB. This could be because, during the first few epochs of speaker extractor training, the loss is high. In Expt 2, such high loss back-propagates and adversely affects the integrity of the attractor encoder. We observe in Expt 3, that by fixing the pre-trained attractor encoder, we gain 0.32 dB of SI-SDRi, and by fine-tuning the entire *reentry* model resulting from Expt 3, we gain another 0.26 dB of SI-SDRi.

C. Study of Speaker Encoders

We study the self-enrolled speaker encoders to justify their role through experiments reported in Table III.

System 10 has the same architecture as systems 1–7, except that the former has neither the speaker encoders nor the speaker classification task, which is designed to train the speaker encoders. In the speaker extractor in Fig. 1, the interlaced architecture between the speaker encoders and the TCN stacks is designed to estimate the mask $M(t)$. With system 10, we would like to examine the contribution of the speaker encoders.

TABLE III

A COMPARISON AMONG THE *REENTRY* MODEL, AND ITS VARIANTS ON THE VOXCELEB2-2MIX DATASET. γ IS THE SCALING FACTOR IN THE LOSS FUNCTION. R IS THE NO. OF TCN STACKS IN THE SPEAKER EXTRACTOR PIPELINE. ‘SHARED WEIGHTS’ DENOTES A NETWORK ARCHITECTURE WHERE THE WEIGHTS IN $R - 1$ SPEAKER ENCODERS ARE SHARED. ‘WITH A^r ’ DENOTES THE USE OF THE SELF-ENROLLED SPEAKER EMBEDDING A^r AS PART OF THE TOP-DOWN ATTENTION CUE

System	γ	R	Shared Weights	With A^r	Val. SI-SDRi (dB)	Test SI-SDRi (dB)
1	0.5	4			12.51	12.03
2	0.1	4			12.92	12.51
3	0.05	4			12.89	12.45
4	0.01	4	✗	✓	12.99	12.52
5	0.005	4			13.01	12.60
6	0.001	4			12.90	12.51
7	0	4			12.68	12.27
8	0.005	3	✗	✓	11.87	11.90
9	0.005	2			11.07	11.13
10	N/A		N/A	N/A	12.64	12.30
11	0.005	4	✗	✗	12.77	12.39
12	0.005		✓	✓	12.83	12.46
13	0.1				13.40	13.01
14	0.005	4	✗	✓	13.38	13.02
15	0				13.34	12.95
16	N/A	4	N/A	N/A	13.15	12.72

For systems 1–7, we vary the scaling factor γ , which weights the contributions between speaker classification and speaker extraction during the multi-task learning. With $\gamma = 0$ (system 7), we ignore the speaker classification task during the training of speaker encoders and obtain an SI-SDRi of 12.27 dB, which is similar to that of system 10. The results suggest that the $R - 1$ speaker encoders are contributing only when they are trained under speaker classification supervision in multi-task learning. When $\gamma = 0.005$ (system 5), the *reentry* model achieves the best SI-SDRi of 12.60 dB, which represents 0.33 dB improvement over that for $\gamma = 0$ (system 7). The results clearly suggest that the speaker encoders benefit from multi-task learning.

We also present systems 13–15 as an upper-bound model, where the speaker encoders take the target speaker’s clean speech as the input instead of the intermediate speech, during both training and evaluation. System 14 delivers the best SI-SDRi of 13.02 dB. The proposed *reentry* model (system 5) lags only 0.42 dB behind the upper-bound (system 14).

Furthermore, we build system 16, which uses a speaker embedding extracted from the target speaker’s clean speech with a pre-trained speaker verification model. During both training and evaluation, the speaker extractor takes such speaker embedding as input instead of speaker embedding estimated on the fly. Despite having clean speech as input, system 16 only attains an SI-SDRi of 12.72 dB. This suggests that the independently trained speaker encoders don’t work the best for speaker extraction, thus once again justifies the proposed joint training between the speaker encoders and the speaker extraction network.

D. Study of Speaker Extractor

We consider that the speaker encoders and speaker classification task may jointly contribute to the speaker extraction in two

ways. First, A^r provides the top-down attention cue for speaker extraction. Second, the speaker classification task supervises the training of speaker encoders and TCN stacks to ensure the output speech carries the target speaker identity. The question is in what way the speaker encoders with speaker classification task contribute to the extracted speech quality.

To answer the question, we implement system 11 as shown in Table III, in which we alter the standard *reentry* model (system 5) by keeping the speaker encoders during multi-task learning, but removing the top-down attention cue A^r from the interlace. During training, the gradients of the speaker classification loss back-propagate through the speaker encoders to update the TCN stacks, therefore, ensuring the speaker identity of the estimated speech. However, at run-time, only the TCN stacks pipeline participates in the mask estimation without A^r .

The difference between system 11 and system 10 is that the former involves the speaker encoders, as well the speaker classification loss, during training, but the latter does not. In Table III, we observe that system 11 has improved SI-SDRi over system 10 from 12.30 dB to 12.39 dB, which shows the speaker classification task, without explicitly using A^r , offers a small gain of 0.09 dB. However, when A^r is used as the attractor along with $V(t)$ in the *reentry* model (system 5), we obtain a higher SI-SDRi of 12.60 dB. The results suggest that the use of self-enrolled A^r as a top-down attention cue dominates the contribution from the speaker encoders and speaker classification task.

The speaker extractor in the *reentry* model consists of R TCN stacks that interlace with $R - 1$ speaker encoders. We conduct further experiments to show the interlaced architecture supports a progressively refined speaker self-enrollment process, and report in Table III. We vary the number of TCN stacks R from 2 to 4 in systems 9, 8, and 5, thus the speaker encoders $R - 1$ from 1 to 3, in the speaker extractor. It is seen that, as R increases, the SI-SDRi of *reentry* increases from 11.13 dB to 12.60 dB at $\gamma = 0.005$, showing that the progressive refinement architecture outperforms a single-stage extraction.

We finally analyze the weight sharing among the speaker encoders. System 12 shares the weights across $R - 1$ speaker encoders, while system 5 has individual weights for each speaker encoder. The results suggest that the individual weights scheme outperforms the shared weights counterpart by 0.14 dB in terms of SI-SDRi, which confirms our intuition to adopt individual weights in the *reentry* model.

E. Benchmark Against Baselines

We now compare the *reentry* model with two recent models TDSE and MuSE on mixture speech of two-speaker (VoxCeleb2-2mix) and three-speaker (VoxCeleb2-3mix) datasets as summarized in Table IV. We observe that the improved versions of TDSE and MuSE, namely TDSE-I and MuSE-I outperform their original model, i.e., TDSE-O [47] and MuSE-O [57], thus serving as competitive baselines. Since MuSE-I employs the same speaker encoders architecture and multi-task learning framework as the *reentry* model, we used the $\gamma = 0.005$ for MuSE-I training as it proves to learn the best speaker representation according to Table III. It is seen

TABLE IV
PERFORMANCE OF THE *REENTRY* MODEL AND BASELINES ON THE TEST SET OF SIMULATED VOXCELEB2 DATASET MIXTURES

Dataset Mixtures	Task	Model	<i>Att</i>	SI-SDRi (dB)	SDRi (dB)	PESQi	STOIi	Params (m)
VoxCeleb2-2mix	Speech separation	Conv-TasNet [16]	N/A	10.94	11.27	0.835	0.230	8.7
		DPRNN [18]	N/A	11.57	11.88	0.902	0.235	2.6
	Speaker extraction	TDSE-O [47]	<i>vsrv</i>	10.64	10.93	0.823	0.226	20.2
		TDSE-I		11.62	11.96	0.958	0.240	20.1
		MuSE-O [57]		11.67	12.04	0.969	0.241	25.0
		MuSE-I		11.91	12.22	1.007	0.245	24.5
	<i>reentry</i>	<i>sync3</i>	12.60	12.92	1.103	0.256	18.8	
VoxCeleb2-3mix	Speech separation	Conv-TasNet [16]	N/A	9.31	9.74	0.264	0.230	8.8
		DPRNN [18]	N/A	9.83	10.23	0.283	0.239	2.7
	Speaker extraction	TDSE-O [47]	<i>vsrv</i>	9.78	10.21	0.360	0.262	20.2
		TDSE-I		11.54	12.04	0.500	0.303	20.1
		MuSE-O [57]		11.64	12.18	0.513	0.306	25.0
		MuSE-I		12.21	12.68	0.559	0.318	24.5
	<i>reentry</i>	<i>sync3</i>	12.63	13.08	0.613	0.327	18.8	

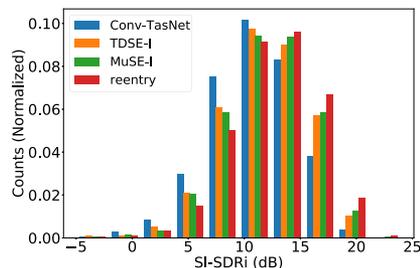


Fig. 6. The histogram of SI-SDRi by Conv-TasNet, TDSE-I, MuSE-I and *reentry* model on VoxCeleb2-2mix.

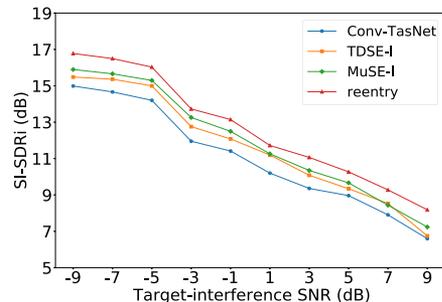


Fig. 7. The average SI-SDRi by Conv-TasNet, TDSE-I, MuSE-I and *reentry* model on VoxCeleb2-2mix as a function of the target-interference SNR.

that the *reentry* model outperforms all the baselines in terms of SI-SDRi, SDRi, PESQi, and STOIi on the VoxCeleb2-2mix and VoxCeleb2-3mix datasets.

We also re-implement two speech separation models, namely Conv-TasNet [16] and DPRNN [18] following the original paper, and report their results on the VoxCeleb2-2mix and VoxCeleb2-3mix datasets for comparison. Since Conv-TasNet and DPRNN are audio-only speech separation algorithms, they don't use the visual information from the VoxCeleb2-2mix dataset mixture. Hence we expect that the audio-visual speaker extraction outperforms the audio-only models.

In Fig. 6, we show the histogram of SI-SDRi for the VoxCeleb2-2mix test set samples by the Conv-TasNet, TDSE-I, MuSE-I, and the *reentry* model. It shows that the *reentry* model has more extracted speech samples with higher SI-SDRi gain than the baselines. In Fig. 7, we show the distribution of SI-SDRi gain for mixtures with various target-interference SNR. The target-interference SNR is defined as the energy contrast between the target speaker and the interference speaker in terms of SNR. It is seen that the *reentry* model consistently achieves a higher SI-SDRi than the baselines. As the input mixture becomes less noisy (high SNR), the SI-SDRi gain becomes smaller.

F. Speaker Extraction in Face of Visual Occlusion

As the proposed *reentry* model relies on the speech-lip synchronization cue, any visual occlusion will adversely impact

TABLE V
THE RESULTS ON THE VOXCELEB2-2MIX-OCCL TEST SET FOR MODELS THAT TRAINED WITH AND WITHOUT OCCLUSION DATA

Model	Train data	SI-SDRi	SDRi	PESQi	STOIi
TDSE-I	VoxCeleb2-2mix	4.98	6.18	0.671	0.091
MuSE-I		4.02	5.40	0.625	0.066
<i>reentry</i>		6.68	7.70	0.801	0.125
TDSE-I	VoxCeleb2-2mix-occl	9.21	9.82	0.856	0.188
MuSE-I		9.29	9.87	0.875	0.188
<i>reentry</i>		10.27	10.74	0.909	0.209

the performance. We evaluate three systems, namely the *reentry* model, TDSE-I, and MuSE-I to observe how systems perform in face of visual occlusions.

We create a two-speaker visual occlusion dataset, named the VoxCeleb2-2mix-occl, based on the VoxCeleb2-2mix dataset, but blackout the visual signals randomly and for a random duration while keeping the audio signals intact. This simulates the scenarios where the face-tracking algorithm fails to detect the presence of the target speaker.

In Table V, it is observed that none of the models generalizes well if not trained on occlusion data. Nonetheless, the *reentry* model outperforms the best baselines for all evaluation metrics. When trained on occlusion data, all models improve, with the *reentry* model still showing the best performance.

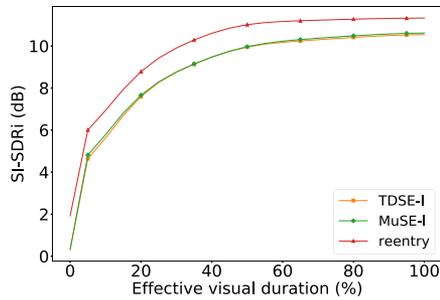


Fig. 8. The average SI-SDRi by TDSE-I, MuSE-I and *reentry* model on VoxCeleb2-2mix-occl as a function of effective visual duration.

TABLE VI
CROSS-DATASETS EVALUATIONS FOR MODELS THAT ARE TRAINED ON THE VOXCELEB2-2MIX DATASET AND TESTED ON OTHER DATASETS

Domain	Dataset	Model	SI-SDRi	SDRi	PESQi	STOIi
Wild	LRS2	TDSE-I	11.61	12.05	0.874	0.234
		MuSE-I	12.21	12.60	0.937	0.245
		<i>reentry</i>	12.66	13.08	1.036	0.250
	LRS3	TDSE-I	13.39	13.71	1.102	0.241
		MuSE-I	13.69	13.98	1.142	0.245
		<i>reentry</i>	14.24	14.54	1.244	0.251
Studio	Grid	TDSE-I	9.33	11.51	0.822	0.077
		MuSE-I	10.18	12.06	0.931	0.090
		<i>reentry</i>	9.83	12.09	0.893	0.098
	TCD-TIMIT	TDSE-I	13.70	14.22	0.989	0.174
		MuSE-I	14.36	14.87	1.067	0.183
		<i>reentry</i>	14.96	15.55	1.118	0.188

We analyze the performance with different effective visual duration after the visual occlusion in Fig. 8. The effective visual duration is defined as the percentage of the duration of the unoccluded visual reference to that of the entire utterance. We plot the SI-SDRi of the *reentry* model and baselines with different effective visual duration. It is seen that when the visual information is present for about 10%, the *reentry* model is able to achieve near 7 dB SI-SDRi. When the visual information is present for about 50%, the performance of the *reentry* model starts to saturate, the model is able to extract the speech segments where the visual information is occluded according to the prosody of the extracted speech when the visual information is present.

G. Cross-Dataset Evaluation

We are interested in how the *reentry* model trained on the VoxCeleb2 dataset performs on other datasets. We evaluate the *reentry* model with two competing models on Grid [100], TCD-TIMIT [101], LRS2 [98], and LRS3 [102] datasets. Grid and TCD-TIMIT are ‘studio’ videos while LRS2 and LRS3 are ‘wild’ videos from BBC and TED. These datasets are pre-processed with face detection and tracking algorithms same as the VoxCeleb2 dataset to minimize the visual mismatch. We generate 3000 utterances to form a test set for each of the above datasets, following the VoxCeleb2-2mix protocol.

As shown in Table VI, on the LRS2 and LRS3 datasets, the *reentry* model shows higher SI-SDRi, SDRi, PESQi, and STOIi, which are consistent with the improvement on the VoxCeleb2

dataset. On the Grid and TCD-TIMIT datasets, which belong to another domain, the *reentry* model still shows a relative improvement over the competing models, with the exception that MuSE-I outperforms the *reentry* model on the Grid dataset for SI-SDRi and PESQi.

VI. CONCLUSION

We propose an end-to-end audio-visual speaker extraction network to emulate human selective auditory attention in the cocktail party. The study is motivated by the *reentry* theory by Edelman which shows the cross-modal temporal integration in human auditory attention. The proposed *reentry* model is particularly useful in situations where a pre-enrolled face or speech reference is not available. In summary, the *reentry* model presents a step forward in solving the cocktail party problem using a computational model. It will potentially enable applications, such as intelligent hearing aids, or downstream tasks such as speech recognition and speaker verification.

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] A. W. Bronkhorst, “The cocktail party phenomenon: A review of research on speech intelligibility in multiple-talker conditions,” *Acta Acustica United Acustica*, vol. 86, no. 1, pp. 117–128, 2000.
- [3] F. Katsuki and C. Constantinidis, “Bottom-up and top-down attention: Different processes and overlapping neural systems,” *Neuroscientist*, vol. 20, no. 5, pp. 509–521, 2014.
- [4] Y. Pinto, A. R. van der Leij, I. G. Sligte, V. A. Lamme, and H. S. Scholte, “Bottom-up and top-down attention are independent,” *J. Vis.*, vol. 13, no. 3, pp. 16–16, 2013.
- [5] R. Tao, R. K. Das, and H. Li, “Audio-visual speaker recognition with a cross-modal discriminative network,” in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2020, pp. 2242–2246.
- [6] W. Rao, C. Xu, E. S. Chng, and H. Li, “Target speaker extraction for multi-talker speaker verification,” in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2019, pp. 1273–1277.
- [7] C. Xu, W. Rao, J. Wu, and H. Li, “Target speaker verification with selective auditory attention for single and multi-talker speech,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 2696–2709, Jul. 2021.
- [8] T. Liu, R. K. Das, K. A. Lee, and H. Li, “Neural acoustic-phonetic approach for speaker verification with phonetic attention mask,” *IEEE Signal Process. Lett.*, vol. 29, pp. 782–786, 2022.
- [9] Y. Ma, K. A. Lee, V. Hautamaki, and H. Li, “PL-EESR: Perceptual loss based end-to-end robust speaker representation extraction,” in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2021, pp. 106–113.
- [10] 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.
- [11] 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.
- [12] Y. Zhou, X. Tian, and H. Li, “Multi-task waveRNN with an integrated architecture for cross-lingual voice conversion,” *IEEE Signal Process. Lett.*, vol. 27, pp. 1310–1314, 2020.
- [13] Y. Zhou, X. Tian, and H. Li, “Language agnostic speaker embedding for cross-lingual personalized speech generation,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 3427–3439, Nov. 2021.
- [14] K. Zhou, B. Sisman, R. Liu, and H. Li, “Seen and unseen emotional style transfer for voice conversion with a new emotional speech dataset,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2021, pp. 920–924.
- [15] 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.

- [16] 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.
- [17] 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.
- [18] 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.
- [19] 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.
- [20] 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, Jul. 2021.
- [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. Kolbæk, D. Yu, Z.-H. Tan, and J. Jensen, "Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks," *IEEE/ACM 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] 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, Apr. 2020.
- [25] M. Ge, C. Xu, L. Wang, E. S. Chng, J. Dang, and H. Li, "SpEx : A complete time domain speaker extraction network," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2020, pp. 1406–1410.
- [26] Q. Wang *et al.*, "VoiceFilter: Targeted voice separation by speaker-conditioned spectrogram masking," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2019, pp. 2728–2732.
- [27] 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.
- [28] 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.
- [29] K. Žmolíková *et al.*, "SpeakerBeam: Speaker aware neural network for target speaker extraction in speech mixtures," *IEEE J. Sel. Top. Signal Process.*, vol. 13, no. 4, pp. 800–814, Aug. 2019.
- [30] 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.
- [31] J. Shi, J. Xu, Y. Fujita, S. Watanabe, and B. Xu, "Speaker-conditional chain model for speech separation and extraction," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2020, pp. 2707–2711.
- [32] 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.
- [33] 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.
- [34] T. Ochiai, M. Delcroix, K. Kinoshita, A. Ogawa, and T. Nakatani, "Multimodal SpeakerBeam: Single channel target speech extraction with audio-visual speaker clues," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2019, pp. 2718–2722.
- [35] 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. Annu. Conf. Int. Speech Commun. Assoc.*, 2021, pp. 1469–1473.
- [36] W. Wang, C. Xu, M. Ge, and H. Li, "Neural speaker extraction with speaker-speech cross-attention network," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2021, pp. 3535–3539.
- [37] 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.
- [38] G. M. Edelman, *Neural Darwinism: The Theory of Neuronal Group Selection*. New York, NY, USA: Basic Books, 1987.
- [39] J. T. Fleming, A. L. Noyce, and B. G. Shinn-Cunningham, "Audio-visual spatial alignment improves integration in the presence of a competing audio-visual stimulus," *Neuropsychologia*, vol. 146, 2020, Art. no. 107530.
- [40] 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, Mar. 2021.
- [41] E. Z. Golombic, 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.
- [42] 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.
- [43] S.-W. Chung, S. Choe, J. S. Chung, and H.-G. Kang, "FaceFilter: Audio-visual speech separation using still images," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2020, pp. 3481–3485.
- [44] L. Qu, C. Weber, and S. Wermter, "Multimodal target speech separation with voice and face references," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2020, pp. 1416–1420.
- [45] R. Gao and K. Grauman, "VisualVoice: Audio-visual speech separation with cross-modal consistency," in *Proc. IEEE/CVF Conf. Comput. Vis. Patt. Recognit.*, 2021, pp. 15 490–15 500.
- [46] T. Afouras, J. S. Chung, and A. Zisserman, "The conversation: Deep audio-visual speech enhancement," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2018, pp. 3244–3248.
- [47] J. Wu *et al.*, "Time domain audio visual speech separation," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2019, pp. 667–673.
- [48] V.-N. Nguyen, M. Sadeghi, E. Ricci, and X. Alameda-Pineda, "Deep variational generative models for audio-visual speech separation," in *Proc. IEEE 31st Int. Workshop Mach. Learn. Signal Process.*, 2021, pp. 1–6.
- [49] C. Li and Y. Qian, "Deep audio-visual speech separation with attention mechanism," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 7314–7318.
- [50] Z. Pan, M. Ge, and H. Li, "USEV: Universal speaker extraction with visual cue for sparsely overlapped multi-talker speech," 2021, *arXiv:2109.14831*.
- [51] R. Lu, Z. Duan, and C. Zhang, "Audio-visual deep clustering for speech separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 11, pp. 1697–1712, Nov. 2019.
- [52] A. Gabbay, A. Shamir, and S. Peleg, "Visual speech enhancement," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2018, pp. 1170–1174.
- [53] G. Morrone, S. Bergamaschi, L. Pasa, L. Fadiga, V. Tikhonoff, and L. Badino, "Face landmark-based speaker-independent audio-visual speech enhancement in multi-talker environments," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2019, pp. 6900–6904.
- [54] R. Lu, Z. Duan, and C. Zhang, "Listen and look: Audio-visual matching assisted speech source separation," *IEEE Signal Process. Lett.*, vol. 25, no. 9, pp. 1315–1319, Sep. 2018.
- [55] D. Soderoy, L. Girin, C. Jutten, and J.-L. Schwartz, "Developing an audio-visual speech source separation algorithm," *Speech Commun.*, vol. 44, no. 1–4, pp. 113–125, 2004.
- [56] 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.
- [57] 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.
- [58] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2018, pp. 1086–1090.
- [59] K. T. Hill and L. M. Miller, "Auditory attentional control and selection during cocktail party listening," *Cereb. Cortex*, vol. 20, no. 3, pp. 583–590, 2010.
- [60] 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.
- [61] 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.

- [62] 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.
- [63] Z. Huang, S. Wang, and K. Yu, "Angular softmax for short-duration text-independent speaker verification," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2018, pp. 3623–3627.
- [64] C. Xu, W. Rao, E. S. Chng, and H. Li, "Time-domain speaker extraction network," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2019, pp. 327–334.
- [65] 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.
- [66] C. Xu, W. Rao, E. S. Chng, and H. Li, "Optimization of speaker extraction neural network with magnitude and temporal spectrum approximation loss," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2019, pp. 6990–6994.
- [67] T. Afouras, J. S. Chung, and A. Zisserman, "My lips are concealed: Audio-visual speech enhancement through obstructions," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2019, pp. 4295–4299.
- [68] K. Ito, M. Yamamoto, and K. Nagamatsu, "Audio-visual speech enhancement method conditioned in the lip motion and speaker-discriminative embeddings," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2021, pp. 6668–6672.
- [69] T. Afouras, J. S. Chung, and A. Zisserman, "Deep lip reading: A comparison of models and an online application," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2018, pp. 3514–3518.
- [70] J. S. Chung and A. Zisserman, "Lip reading in the wild," in *Proc. Asian Conf. Comput. Vis.*, 2016, pp. 87–103.
- [71] J. Lu, B. Sisman, R. Liu, M. Zhang, and H. Li, "VisualTTS: TTS with accurate lip-speech synchronization for automatic voice over," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 8032–8036.
- [72] E. M. Rudd, M. Günther, and T. E. Boulton, "MOON: A mixed objective optimization network for the recognition of facial attributes," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 19–35.
- [73] 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.
- [74] A. Owens and A. A. Efros, "Audio-visual scene analysis with self-supervised multisensory features," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 631–648.
- [75] T. Afouras, A. Owens, J. S. Chung, and A. Zisserman, "Self-supervised learning of audio-visual objects from video," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 208–224.
- [76] J. W. Fisher and T. Darrell, "Speaker association with signal-level audiovisual fusion," *IEEE Trans. Multimedia*, vol. 6, no. 3, pp. 406–413, Jun. 2004.
- [77] J. Hershey and J. R. Movellan, "Audio-vision: Locating sounds via audio-visual synchrony," in *Proc. Adv. Neural Inf. Process. Syst.*, 1999, pp. 813–819.
- [78] E. Marcheret, G. Potamianos, J. Vopicka, and V. Goel, "Detecting audio-visual synchrony using deep neural networks," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2015, pp. 548–552.
- [79] J. S. Chung and A. Zisserman, "Out of time: Automated lip sync in the wild," in *Proc. Asian Conf. Comput. Vis.*, 2016, pp. 251–263.
- [80] S.-W. Chung, J. S. Chung, and H.-G. Kang, "Perfect match: Improved cross-modal embeddings for audio-visual synchronisation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2019, pp. 3965–3969.
- [81] Y. Shalev and L. Wolf, "End to end lip synchronization with a temporal autoencoder," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, 2020, pp. 341–350.
- [82] 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.
- [83] A. van den Oord *et al.*, "WaveNet: A generative model for raw audio," in *Proc. 9th ISCA Speech Synth. Workshop*, 2016, pp. 125–125.
- [84] D. Rethage, J. Pons, and X. Serra, "A wavenet for speech denoising," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 5069–5073.
- [85] T. Stafylakis and G. Tzimiropoulos, "Combining residual networks with LSTMs for lipreading," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2017, pp. 3652–3656.
- [86] M. Gutmann and A. Hyvärinen, "Noise-contrastive estimation: A new estimation principle for unnormalized statistical models," in *Proc. 13th Int. Conf. Artif. Intell. Statist.*, 2010, pp. 297–304.
- [87] Z. Pan, Z. Luo, J. Yang, and H. Li, "Multi-modal attention for speech emotion recognition," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2020, pp. 364–368.
- [88] Y. Li and D. Wang, "On the optimality of ideal binary time-frequency masks," *Speech Commun.*, vol. 51, no. 3, pp. 230–239, 2009.
- [89] A. Narayanan and D. Wang, "Ideal ratio mask estimation using deep neural networks for robust speech recognition," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2013, pp. 7092–7096.
- [90] Y. Wang, A. Narayanan, and D. Wang, "On training targets for supervised speech separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 22, no. 12, pp. 1849–1858, Dec. 2014.
- [91] H. Erdogan, J. R. Hershey, S. Watanabe, and J. Le Roux, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2015, pp. 708–712.
- [92] A. Oppenheim and R. Schaffer, *Theory and Application of Digital Signal Processing*. Englewood Cliffs, NJ, USA: Pearson Educ., 1978.
- [93] 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.
- [94] 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.
- [95] W. Liu *et al.*, "SSD: Single shot multibox detector," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 21–37.
- [96] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A large-scale speaker identification dataset," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2017, pp. 2616–2620.
- [97] D. P. Kingma and J. Ba, "Adam, A method for stochastic optimization," in *Proc. 3rd Int. Conf. Learn. Representations*, 2015.
- [98] T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman, "Deep audio-visual speech recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, to be published, doi: [10.1109/TPAMI.2018.2889052](https://doi.org/10.1109/TPAMI.2018.2889052).
- [99] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (PESQ)—A new method for speech quality assessment of telephone networks and codecs," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2001, pp. 749–752.
- [100] M. Cooke, J. Barker, S. Cunningham, and X. Shao, "An audio-visual corpus for speech perception and automatic speech recognition," *J. Acoustical Soc. Amer.*, vol. 120, no. 5, pp. 2421–2424, 2006.
- [101] N. Harte and E. Gillen, "TCD-TIMIT: An audio-visual corpus of continuous speech," *IEEE Trans. Multimedia*, vol. 17, no. 5, pp. 603–615, May 2015.
- [102] T. Afouras, J. S. Chung, and A. Zisserman, "LRS3-TED: A large-scale dataset for visual speech recognition," 2018, [arXiv:1809.00496](https://arxiv.org/abs/1809.00496).



Zexu Pan received the B.Eng. degree in electrical and electronic engineering from Nanyang Technological University, Singapore, in 2019. He is currently working toward the Ph.D. degree with the Institute of Data Science, National University of Singapore, Singapore. His research interests include speech extraction, speech enhancement, multitalker robust automatic speech recognition, and multimodality representation learning.



Ruijie Tao received the B.Eng. degree from Soochow University, Suzhou, China, in 2018, and the M.Sc. degree in 2019 from the National University of Singapore, Singapore, where he is currently working toward the Ph.D. degree. His research interests include speaker recognition, speech extraction, and audio-visual representation learning.



recognition.

Chenglin Xu (Member, IEEE) received the B.Eng. and M.Sc. degrees from Northwestern Polytechnical University, Xi'an, China, in 2012 and 2015, respectively, and the Ph.D. degree from Nanyang Technological University, Singapore, in 2020. He was a Research Fellow with HLT Laboratory, Department of Electrical and Computer Engineering, National University of Singapore. He is currently with Kuaishou Technology. His research interests include speech extraction, speech enhancement, source separation, multitalker speaker verification, and robust speech



Haizhou Li (Fellow, IEEE) received the B.Sc., M.Sc., and Ph.D. degrees in electrical and electronic engineering from the South China University of Technology, Guangzhou, China, in 1984, 1987, and 1990, respectively. From 1988 to 1990, he taught in The University of Hong Kong, Hong Kong, and the South China University of Technology from 1990 to 1994. From 1994 to 1995, he was a Visiting Professor with CRIN, France, Research Manager with Apple-ISS Research Centre from 1996 to 1998, Research Director in Lernout and Hauspie Asia Pacific from 1999 to 2001, the Vice President with InfoTalk Corp. Ltd. from 2001 to 2003, and the Principal Scientist and Department Head of Human Language Technology with the Institute for Infocomm Research, Singapore from 2003 to 2016. He is currently a 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. His research interests include automatic speech recognition, speaker and language recognition, and natural language processing. Dr. Li was the Editor-in-Chief of the IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH AND LANGUAGE PROCESSING during 2015–2018, an elected Member of IEEE Speech and Language Processing Technical Committee during 2013–2015, the President of the International Speech Communication Association during 2015–2017, President of Asia Pacific Signal and Information Processing Association during 2015–2016, and President of Asian Federation of Natural Language Processing during 2017–2018. He was the General Chair of ACL 2012, INTERSPEECH 2014, ASRU 2019 and ICASSP 2022. Since 2012, he has been a Member of the Editorial Board of Computer Speech and Language. Dr. Li is a Fellow of the ISCA and the Academy of Engineering Singapore. He was the recipient of the National Infocomm Award 2002 and President's Technology Award 2013 in Singapore. He was named one of the two Nokia Visiting Professors in 2009 by the Nokia Foundation, and U Bremen Excellence Chair Professor in 2019.