

TS-SEP: Joint Diarization and Separation Conditioned on Estimated Speaker Embeddings

Boeddeker, Christoph; Subramanian, Aswin Shanmugam; Wichern, Gordon; Haeb-Umbach,
Reinhold; Le Roux, Jonathan

TR2024-006 February 16, 2024

Abstract

Since diarization and source separation of meeting data are closely related tasks, we here propose an approach to perform the two objectives jointly. It builds upon the target- speaker voice activity detection (TS-VAD) diarization approach, which assumes that initial speaker embeddings are available. We replace the final combined speaker activity estimation network of TS-VAD with a network that produces speaker activity estimates at a time-frequency resolution. Those act as masks for source extraction, either via masking or via beamforming. The technique can be applied both for single-channel and multi-channel input and, in both cases, achieves a new state-of-the-art word error rate (WER) on the LibriCSS meeting data recognition task. We further compute speaker-aware and speaker-agnostic WERs to isolate the contribution of diarization errors to the overall WER performance.

IEEE/ACM Transactions on Audio, Speech, and Language Processing 2024

© 2024 ACM. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org or Publications Dept., ACM, Inc., fax +1 (212) 869-0481.

Mitsubishi Electric Research Laboratories, Inc.
201 Broadway, Cambridge, Massachusetts 02139

TS-SEP: Joint Diarization and Separation Conditioned on Estimated Speaker Embeddings

Christoph Boeddeker, Aswin Shanmugam Subramanian, Gordon Wichern,
Reinhold Haeb-Umbach, and Jonathan Le Roux

Abstract—Since diarization and source separation of meeting data are closely related tasks, we here propose an approach to perform the two objectives jointly. It builds upon the target-speaker voice activity detection (TS-VAD) diarization approach, which assumes that initial speaker embeddings are available. We replace the final combined speaker activity estimation network of TS-VAD with a network that produces speaker activity estimates at a time-frequency resolution. Those act as masks for source extraction, either via masking or via beamforming. The technique can be applied both for single-channel and multi-channel input and, in both cases, achieves a new state-of-the-art word error rate (WER) on the LibriCSS meeting data recognition task. We further compute speaker-aware and speaker-agnostic WERs to isolate the contribution of diarization errors to the overall WER performance.

Index Terms—Meeting recognition, meeting separation, speaker diarization

I. INTRODUCTION

CURRENT research in meeting transcription is not only concerned with transcribing the audio recordings of meetings into machine-readable text, but also with enriching the transcription with diarization information about “who spoke when”. Much of the difficulty of these two tasks can be attributed to the interaction dynamics of multi-talker conversational speech. Speakers articulate themselves in an intermittent manner with alternating segments of speech inactivity, single-talker speech, and multi-talker speech. In particular, overlapping speech, where two or more people are talking at the same time, is known to pose a significant challenge not only to ASR but also to diarization [1].

There is no obvious choice whether to start the processing with a diarization or with a separation/enhancement module [2]. An argument in favor of first separating the overlapped speech segments, and at the same time removing other acoustic distortions, and then carrying out diarization, is the fact that the diarization task is much easier when performed on speech without overlap. This processing order was advocated in [2]. Also, the continuous speech separation (CSS) pipeline starts with a separation module [3]. Early diarization systems, such as the clustering-based approaches that were predominant in

C. Boeddeker and R. Haeb-Umbach are with Paderborn University, Paderborn, 33098, Germany (e-mail: {boeddeker,haeb}@nt.upb.de).

A. S. Subramanian was with Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA 02139, USA, and is now with Microsoft Corporation, One Microsoft Way, Redmond, WA 98052, USA (e-mail: asubra13@alumni.jh.edu).

G. Wichern and J. Le Roux are with Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA 02139, USA (e-mail: {wichern,leroux}@merl.com).

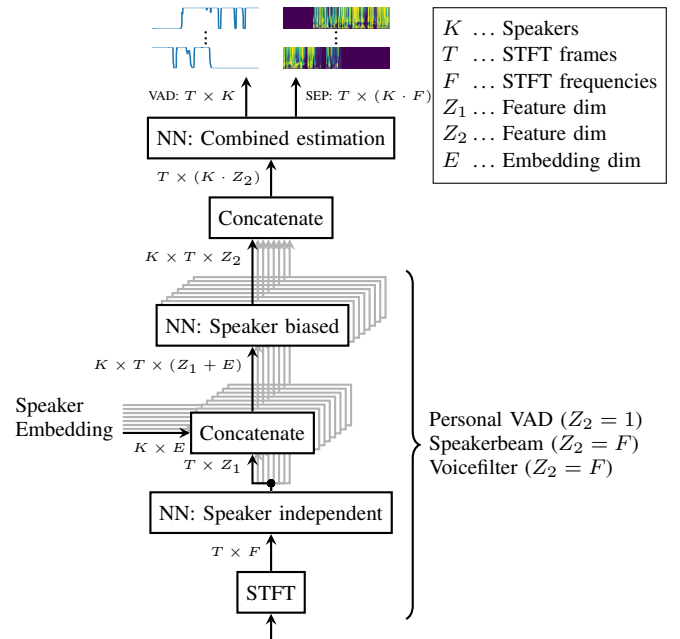


Fig. 1. TS-VAD and TS-SEP structure for one microphone. Speakers indicated as “shading”.

the first editions of the DiHARD diarization challenge, were generally unable to properly cope with overlapping speech, and thus performed poorly on data with a significant amount of concurrent speech [4].

On the other hand, starting the processing with the diarization component can be advantageous, because signal extraction is eased if information about the segment boundaries of speech absence, single-, and multi-talker speech is given. This order of processing has for example been adopted in the CHiME-6 baseline system [1]. However, this requires the diarization component to be able to cope with overlapped speech.

These different points of view highlight the interdependence of the two tasks, and they are a clear indication that a joint treatment of diarization and separation/source extraction can be beneficial. Indeed, subtasks to be solved in either of the two are similar: while diarization is tasked with determining the speakers that are active in each time frame, mask-based source extraction in essence identifies the active speakers for each time-frequency (TF) bin, the difference thus being only the resolution, time vs. time-frequency.

Early examples of joint diarization and source separation systems are spatial mixture models (SMMs) using time-

varying mixture weights [5], [6]. There, the estimates of the prior probabilities of the mixture components, after appropriate quantization, give the diarization information about who speaks when, while the posterior probabilities have TF-resolution and can be employed to extract each source present in a mixture, either by masking or by beamforming. A crucial issue is the initialization of the mixture weights or posterior probabilities. In the guided source separation (GSS) framework [7], which was part of the baseline system of the CHiME-6 challenge [1], manual annotation of the segment boundaries, and later estimates thereof [8], were used to initialize the time-varying mixture weights. An initialization scheme exploiting the specifics of meeting data, in which most of the time only a single speaker is active, was introduced in [6].

However, this approach to joint diarization and separation depends on the availability of multi-channel input, and, due to the iterative nature of the expectation maximization (EM) algorithm, it is computationally demanding. Furthermore, the EM algorithm is, at least in its original formulation, an offline algorithm that is not well suited for processing arbitrarily long meetings. There exist recursive online variants of the EM algorithm, which are nevertheless computationally intense and need careful parameter tuning, in particular in dynamic scenarios [9].

In light of their recent success, both in diarization and source separation, neural networks appear to be promising for developing an integrated solution to the problem of joint diarization and separation/extraction. Well-known methods for extracting a target speaker from a mixture include SpeakerBeam [10] and VoiceFilter [11]. Both require an enrollment utterance of the target speaker to be extracted. In the SpeakerBeam approach, the target speaker embedding and the extraction network are trained jointly, while VoiceFilter employs a separately trained speaker embedding computation network. However, their reliance on the availability of an enrollment sentence somewhat limits their usefulness for some applications. Furthermore, they have been designed for a typical source separation/extraction scenario, where it is a priori known that the target speaker is active most of the time. Only recently has an extension of the SpeakerBeam system been proposed, to deal with meeting data [12].

In the field of diarization, end-to-end neural diarization (EEND) has recently made significant inroads, because it can cope with overlapping speech by casting diarization as a multi-label classification problem [13]–[16]. EEND processes the input speech in blocks of a few seconds. This introduces a block permutation problem, which can for example be solved by subsequent clustering [17].

The target-speaker voice activity detection (TS-VAD) diarization system [8], [18], which emerged from the earlier personal VAD approach [19], has demonstrated excellent performance on the very challenging CHiME-6 data set [1]. Assuming that the total number of speakers in a meeting is known, and that embedding vectors representing the speakers are available, TS-VAD estimates the activity of all speakers simultaneously, as illustrated in Fig. 1. This combined estimation was shown to be key to high diarization performance. Rather than relying on enrollment utterances, it estimates

speaker profiles from estimated single-speaker regions of the recording to be diarized. It was later shown that the exact knowledge of the number of speakers is unnecessary, as long as a maximum number of speakers potentially present can be given [20], and the attention approach of [21] could do away even with this requirement.

However, TS-VAD is unable to deliver TF-resolution activity, which is highly desirable to be able to extract the individual speakers' signals from a mixture of overlapping speech. We here propose a joint approach to diarization and separation that is an extension of the TS-VAD architecture: its last layer is replicated F times, where F is the number of frequency bins, see Fig. 1. With this modification, the system is now able to produce time-frequency masks. We call this approach target-speaker separation (TS-SEP).

While this may appear to be a small change to the original TS-VAD architecture, the impact is profound: it results in a joint diarization and separation system that is trained under a common objective function.

Such a joint diarization and separation with neural networks (NNs), where all speakers are simultaneously considered, has so far not been done. Existing works either do not do the combined estimation of all speakers for separation [12], which was a key for the success of TS-VAD, or do experiments on short recordings [22]–[24] where diarization is not necessary and hence it is not clear how the approaches would work on long data with many speakers.

In the following, we first introduce the signal model and notations in Section II. Then, in Section III, we describe the TS-VAD architecture, from which we derive the proposed TS-SEP in Section IV. Section V presents a detailed experimental evaluation. Notably, we achieve a new state-of-the-art WER on the LibriCSS data set.

II. SIGNAL MODEL

We assume that the observation \mathbf{y}_ℓ is the sum of K speaker signals $\mathbf{x}_{\ell,k}$ and a noise signal \mathbf{n}_ℓ :

$$\mathbf{y}_\ell = \sum_k \mathbf{x}_{\ell,k} + \mathbf{n}_\ell \in \mathbb{R}^D, \quad (1)$$

where ℓ is the sample index, and \mathbf{y}_ℓ , $\mathbf{x}_{\ell,k}$, and \mathbf{n}_ℓ are vectors of dimension D , the total number of microphones. We use d as microphone index, e.g., $y_{\ell,d}$ is the d 's entry of \mathbf{y}_ℓ . Note that we will also cover the single-channel case, where $D = 1$. Since we consider a long audio recording scenario, it is natural to assume that speaker signals $\mathbf{x}_{\ell,k}$ have active and inactive time intervals. We make this explicit by introducing a speaker activity variable $a_{\ell,k} \in \{0, 1\}$:

$$\mathbf{y}_\ell = \sum_k a_{\ell,k} \mathbf{x}_{\ell,k} + \mathbf{n}_\ell. \quad (2)$$

Note that $a_{\ell,k}$ has no effect on the equation, because $\mathbf{x}_{\ell,k}$ is zero when $a_{\ell,k}$ is zero. Going to the short-time Fourier transform (STFT) domain, we have

$$\mathbf{Y}_{t,f} = \sum_k A_{t,k} \mathbf{X}_{t,f,k} + \mathbf{N}_{t,f}, \quad (3)$$

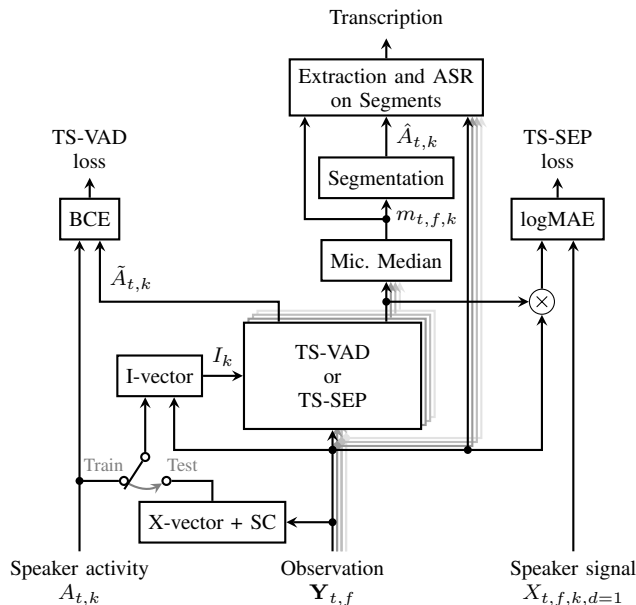


Fig. 2. System overview: Pretraining with TS-VAD loss, training with TS-SEP loss and predicting transcriptions at inference. Microphones indicated as “shading”.

where $t \in \{1, \dots, T\}$ and $f \in \{1, \dots, F\}$ are the time frame and frequency bin indices, and T and F are the total number of time frames and frequency bins, respectively, while $\mathbf{Y}_{t,f}$, $\mathbf{X}_{t,f,k}$, and $\mathbf{N}_{t,f}$ are the STFTs of \mathbf{y}_t , $\mathbf{x}_{t,k}$, and \mathbf{n}_t , respectively. Further, $A_{t,k} \in \{0, 1\}$ is the activity of the k -th speaker at frame t , which is derived from $a_{\ell,k}$ by temporally quantizing to time frame resolution.

III. TS-VAD

In the CHiME-6 challenge, target-speaker voice activity detection (TS-VAD) [8], [18] achieved impressive results for the diarization task, outperforming other approaches by a large margin. The basic idea is similar to personal VAD [19]: an NN is trained to predict the speech activity $A_{t,k}$ of a target speaker k from a mixture $\mathbf{Y}_{t,f}$, given an embedding I_k representing the speaker, as illustrated in Fig. 1.

TS-VAD has two main differences to prior work: first, the speaker embeddings are estimated from the recording instead of utilizing enrollment utterances, as illustrated in Fig. 2, and second, a combination layer estimates the activities of all speakers in a segment simultaneously, given an embedding for each speaker, as illustrated in Fig. 1.

A. Speaker embedding estimation

TS-VAD relies on the availability of initial diarization information, such as from manual annotation or from an embedding extraction system (e.g., X-vector [25]) followed by a clustering approach (e.g., spectral clustering [26]). Speaker embedding vectors (e.g., I-vectors [27]) are then computed from those segments where only a single speaker is active:

$$I_k = \text{Emb} \left\{ y_{\ell,d=1}, \forall \ell \text{ s.t. } \hat{a}_{\ell,k} = 1 \ \& \ \sum_{\bar{k} \neq k} \hat{a}_{\ell,\bar{k}} = 0 \right\} \in \mathbb{R}^E. \quad (4)$$

Here, $\text{Emb}\{\cdot\}$ symbolizes the computation of the embedding vector from segments of speech from the first microphone where a single speaker is active. Furthermore, $\hat{a}_{\ell,k}$ is an estimate for $a_{\ell,k}$.

While this approach does not require an enrollment utterance, it makes TS-VAD dependent on another, initial diarization system. Thus, TS-VAD can be viewed as a refinement system: given the estimate of a first diarization, TS-VAD is applied to estimate a better diarization. In [18], it was reported that I-vector embeddings are superior to X-vectors for TS-VAD, which is in contrast to other publications that suggest a superiority of X-vectors over I-vectors for diarization tasks [25]. We follow [18] and use I-vectors.

B. TS-VAD architecture

The TS-VAD network consists of three components, with stacking operations between them, as shown in Fig. 1. First, the logarithmic spectrogram and the logarithmic mel filterbank features of one microphone of the mixture are stacked and encoded by a few speaker-independent layers, which can be viewed as feature extraction layers, into matrices with dimension $\mathbb{R}^{T \times Z_1}$, where Z_1 is the size of a single framewise representation.

Next, for each speaker k , the framewise representation is concatenated with the speaker embedding of that speaker, resulting in an input to the second network component of dimension $\mathbb{R}^{T \times (Z_1 + E)}$. This second network component processes each speaker independently, nevertheless the NN layer parameters are shared between the speakers. As a consequence, a discrimination of the speakers can only be achieved through the speaker embeddings, not through the NN layer parameters. For each speaker, the output of the second network component has a dimension of $\mathbb{R}^{T \times Z_2}$, with Z_2 denoting the size per frame. Until this point, the NN design matches that of personal VAD [19].

Before the last NN layers, the hidden features of all speakers are concatenated to obtain one large feature matrix of dimension $\mathbb{R}^{T \times (K \cdot Z_2)}$. The final layers produce an output of dimension $\mathbb{R}^{T \times K}$, which, after thresholding and smoothing (see Section IV-C), gives the activity estimate $\hat{A}_{t,k}$ for all speakers k and frames t simultaneously. This combined estimation makes it easier to distinguish between similar speakers. In fact, it was shown in [18] that the combined estimation layers at the end were instrumental in obtaining good performance.

From this description, it is obvious that TS-VAD assumes knowledge of the total number K of speakers in the meeting to be diarized, because K defines the dimensionality of the network output. This constraint can be relaxed by incorporating an attention mechanism as was shown in [21], for the case of fully overlapped speech separation. Nevertheless, we keep the original TS-VAD stacking, only increasing the number of speakers from 4 in [18] to 8. We do so first because an upper bound of 8 speakers is already large for many meetings, and second to allow for a better comparison with [20], where this stacking is also used.

C. From single-channel to multi-channel

When multiple channels are available, there are different approaches to utilize them, e.g., via spatial features like interchannel phase differences (IPD) [28], for intermediate signal reduction such as in an attention layer [18], and output signal reduction [29]. Each of them has different advantages and disadvantages. We use here an output signal reduction: this means the NN is applied independently to each channel and the output is reduced to a single estimate with a median operation. Incorporating multiple channels in this way allowed us to train the NN with a single channel and use all channels at test time. Besides this, this makes our method agnostic to the array geometry and number of microphones, and the presence of moving speakers, as in CHiME-6 [1], becomes less critical. Finally, this approach was shown to be robust to a mismatch between simulated training data and real test data [29].

For ease of notation, we ignore microphone indices for the NN input and output in the following section. At training time, we use only a reference microphone, and at test time, the NN is applied independently to each microphone and then the median is used.

IV. FROM TS-VAD TO TS-SEP

We are now going to describe our modifications to the TS-VAD architecture in order to do joint diarization and separation. The key difference between diarization and (mask-based) separation is that diarization output has time resolution, while separation requires time-frequency resolution.

A. From Frame to Time-Frequency Resolution

Starting from the TS-VAD system, the output size of the last layer is changed from K speakers to K speakers times F frequency bins: $\mathbb{R}^{T \times K} \rightarrow \mathbb{R}^{T \times (F \cdot K)}$, where \rightarrow denotes the “repeat” operation, which is done by repeating the weights of the last layer. Rearranging the output from $\mathbb{R}^{T \times (F \cdot K)}$ to $\mathbb{R}^{T \times F \times K}$, we obtain a $(T \times F)$ -dimensional spectro-temporal output for each speaker k . With a sigmoid nonlinearity, which ensures that values are in $[0, 1]$, this can be interpreted as a mask $m_{t,f,k}, \forall t, f, k$, that can be used for source extraction, for example via masking:

$$\hat{X}_{t,f,k} = m_{t,f,k} Y_{t,f,d=1}, \quad (5)$$

where $d = 1$ is the reference microphone. We denote this system as target-speaker separation (TS-SEP).

The sketches at the top of Fig. 1 illustrate the different outputs of TS-VAD and TS-SEP and Fig. 2 gives an overview of TS-SEP, whose components are described in the following.

B. Training Schedule and Objective

As TS-SEP emerged from TS-VAD, an obvious choice for training is a two-stage training schedule: in the first, pretraining stage, an ordinary TS-VAD system is trained, until it starts to distinguish between speakers; then, the last layer is copied F times to obtain the desired time-frequency resolution; the second training stage is then the training of the




Fig. 3. Estimated speaker activity $\hat{A}_{t,k}$ for 2 of 8 speakers from the LibriCSS DEV OV40 dataset, shown together with the corresponding clean speaker time signals $x_{l,k}$. Best viewed in color.

TS-SEP system, initialized with the parameters of the TS-VAD system.

For the TS-VAD pretraining, we used the training loss proposed in the original publication [18], which is the sum of the binary cross entropy (BCE) losses between estimated and ground-truth activities for all speakers.

There are several choices for the training objective of the TS-SEP system. We opted for a time-domain reconstruction loss, since it implicitly accounts for phase information (see [30], [31] for a discussion of time- vs. frequency-domain reconstruction losses). To be specific, a time-domain signal reconstruction loss is computed by applying an inverse STFT to $\hat{X}_{t,f,k}$ to obtain $\hat{x}_{l,k}$ and measuring the logarithmic mean absolute error (LogMAE) from the ground truth $x_{l,k}$:

$$\mathcal{L} = \log_{10} \frac{1}{L} \sum_k \sum_\ell |\hat{x}_{l,k} - x_{l,k}|. \quad (6)$$

Clearly, other loss functions, such as the MAE, mean squared error (MSE), or signal-to-distortion ratio (SDR) can also be employed. However, it is important to note that for reconstruction losses that contain a log operation, the sum across the speakers should be performed before applying the log operation. Otherwise, the loss is undefined and the training can become unstable if a speaker is completely silent [32].

We also experimented with training the TS-SEP system from scratch. However, this turned out to be far more sensitive to the choice of training loss. We found that only a modified binary cross entropy loss, where we artificially gave a higher weight to loss contributions corresponding to “target present”, was able to learn from scratch, while other choices did not converge at all¹. But even with this loss, the final performance was worse than with the two-stage training². Here, we therefore only report experiments with the described two-stage training.

C. Activity estimation / Segmentation

TS-SEP outputs a mask with TF resolution. To get an estimate like TS-VAD with only time resolution, we take the

¹The networks learned first to predict only zeros and were not able to leave these local optima.

²In some preliminary experiments, we obtained 8.72% concatenated minimum-permutation word error rate (cpWER) for two-stage training, while training from scratch led to 100% cpWER with the LogMAE loss (because of zero estimates) and 12.0% cpWER with the binary cross entropy-based loss.




Fig. 4. Thresholding and “closing” morphological operation (dilation followed by erosion) on $\hat{A}_{t,k}$ to obtain $\tilde{A}_{t,k}$, for 2 of 8 speakers from the LibriCSS DEV OV40 dataset. Best viewed in color.

mean across the frequencies:

$$\tilde{A}_{t,k} = \frac{1}{F} \sum_f m_{t,f,k}. \quad (7)$$

Typically, this estimate has spikes for each word and is close to zero between words, see Fig. 3. To fill those gaps, we followed [6] and used thresholding and the “closing” morphological operation (see Fig. 4), known from image processing [33]: in a first dilation step, a sliding window is moved over the signal with a shift of one, and the maximum value inside the window is taken as the new value for its center sample; in a subsequent erosion step, we do the same, however with the minimum operation, leading to the final smoothed activity estimate

$$\hat{A}_{t,k} = \text{Erosion}(\text{Dilation}(\delta(\tilde{A}_{t,k} \geq \tau))). \quad (8)$$

Here, τ is the threshold and $\delta(x)$ denotes the Kronecker delta which evaluates to one if x is true and zero otherwise. When the window for the dilation is larger than the window for erosion, the speech activity is overestimated, i.e., the difference divided by two is added to the beginning and end of the segment, as illustrated in Fig. 4. While an overestimation will likely worsen the diarization performance, this may not be the case for automatic speech recognition (ASR): it is indeed common to train ASR systems with recordings that contain a single utterance plus some silence (or background noise) at the beginning and end of the recording. We investigate the overestimation later in Section V-H.

At activity change points of $\hat{A}_{t,k}$, the speakers’ signals are cut, leading to segments of constant speaker activity.

This thresholding and “closing” procedure departs from the median-based smoothing typically used for TS-VAD. When the diarization system is trained to fill gaps of short inactivity (e.g., pauses between words), median-based smoothing is fine. But in TS-SEP, the system is trained to predict a time-frequency reconstruction mask for each speaker, so typically the mask values are smaller than the VAD values, short inactivity produces zeros, and values are smaller in overlap regions than in non-overlap regions. The thresholding and closing morphological operations help to fill the gaps.

D. Extraction

At evaluation time, an estimate of the signal of speaker k is obtained by multiplying the mask with the STFT of the input speech, see Eq. (5). If multi-channel input is available, an alternative to mask multiplication for source extraction is to utilize the estimated masks to compute beamformer coefficients [29]. For beamforming, the spatial covariance matrices of the desired signal

$$\Phi_{xx,f,b} = \frac{1}{|\mathcal{T}_b|} \sum_{t \in \mathcal{T}_b} m_{t,f,k_b} \mathbf{Y}_{t,f} \mathbf{Y}_{t,f}^H \quad (9)$$

and of the distortion

$$\Phi_{dd,f,b} = \frac{1}{|\mathcal{T}_b|} \sum_{t \in \mathcal{T}_b} \max \left(\varepsilon, \sum_{\bar{k} \neq k_b} m_{t,f,\bar{k}} \right) \mathbf{Y}_{t,f} \mathbf{Y}_{t,f}^H \quad (10)$$

are first computed. Here, b denotes the segment index, \mathcal{T}_b is the set of frame indices that belong to segment b , k_b is the index of the speaker active in segment b who is to be extracted³, $\varepsilon = 0.0001$ is a small value introduced for stability, and $\sum_{\bar{k} \neq k_b}$ denotes the summation over all speaker indices, except k_b .

With these covariance matrices, the beamformer coefficients can be computed for example using a minimum variance distortionless response (MVDR) beamformer in the formulation of [34]:

$$\mathbf{w}_{f,b} = \text{MVDR}(\Phi_{xx,f,b}, \Phi_{dd,f,b}). \quad (11)$$

Finally, source extraction is performed by applying the beamformer to the input speech:

$$\hat{X}_{t,f,k_b} = \mathbf{w}_{f,b}^H \mathbf{Y}_{t,f}, \quad \forall t \in \mathcal{T}_b. \quad (12)$$

In our experiments, we also combined beamforming with mask multiplication [35], which led to somewhat better suppression of competing speakers:

$$\hat{X}_{t,f,k_b} = \mathbf{w}_{f,b}^H \mathbf{Y}_{t,f} \max(m_{t,f,k_b}, \xi), \quad \forall t \in \mathcal{T}_b, \quad (13)$$

where $\xi \in [0, 1]$ is a lower bound/threshold for the mask.

The lower bound can be motivated by the observations from [36]: pure NN based speech enhancement can degrade the ASR performance compared to doing nothing; when adding a scaled version of the observation to the enhanced signal, the ASR performance can improve compared to doing nothing.

E. Guided Source Separation (GSS)

An alternative to the immediate use of the mask for extraction is to fine-tune the mask first. If multi-channel input is available, this can be achieved by an SMM, which utilizes spatial information. Here, we employ the GSS proposed in [7], where the guiding information is given by the diarization information $\hat{A}_{t,k}$: the posterior probability that the k -th mixture component is active in a frame t can only be non-zero if $\hat{A}_{t,k}$ is one. The SMM is applied to segments b whose boundaries

³The segments are obtained from $\hat{A}_{t,k}$, such that each segment b is associated with a single target speaker k_b ; when multiple speakers are active at the same time, each speaker gets a different segment index b .




Fig. 5. Distribution of the number of speakers active in 1-minute chunks of the training data, depending on the probability ζ to use superposition/mixup ($\zeta = 0$ corresponds to the original dataset without superposition).

are determined by $\hat{A}_{t,k}$, plus some context. The EM iterations using the guiding information are complemented with a non-guided EM step. The finally estimated class posterior probabilities form the speaker-specific time-frequency masks that are used for source extraction by beamforming. For details on GSS, see [7].

While earlier works initialized the SMM by broadcasting voice activity detection (VAD) information $\hat{A}_{t,k}$ over the frequency axis that was available either from manual annotation (CHiME-5 [37]) or determined automatically (CHiME-6 Track 2 [1]), we here propose to initialize the posterior of the SMM with $m_{t,f,k}$, which is only available for TS-SEP, while voice activity information $\hat{A}_{t,k}$ is available for TS-VAD and TS-SEP.

V. EXPERIMENTS

A. Training and Test Datasets

Diarization and recognition experiments are carried out on the LibriCSS data set, which is a widely used data set for evaluating meeting recognition systems [3]. It contains re-recordings of loudspeaker playback of LibriSpeech sentences mixed in a way to reflect a typical meeting scenario. The dataset consists of 10 one-hour-long sessions, where each session is subdivided into six 10-minute-long mini sessions that have different overlap ratios, ranging from 0 to 40%. In each session, a total of eight speakers are active. The recording device was a seven-channel circular microphone array. Following [2], we used the first session as development (DEV) set and the other 9 as evaluation (EVAL) set.

As training data, we employed simulated meeting data created with scripts⁴ from the LibriCSS authors. While each training meeting contains 8 speakers, we trained only on chunks of 1 minute. Since within a chunk usually fewer speakers are active, as shown in Fig. 5, we superposed segments of speech to artificially increase the number of active speakers, see Section V-D for a discussion on this.

B. Performance Measures

Since this work focuses on meeting transcription enriched with information about who spoke when, we measure both diarization and word error rate performance.

As the diarization performance measure, we employ the widely used diarization error rate (DER), which is measured with the “md-eval-22.pl” script from NIST [38]. The DER is

the sum of false alarms, missed hits, and speaker confusion errors divided by the total speaker time.

For the recognition of partially overlapping speech of multiple speakers, there are several word error rate (WER) definitions [39]. Since TS-VAD groups utterances of the same speaker, we mainly report the concatenated minimum-permutation word error rate (cpWER), where all transcriptions of one speaker are concatenated and the word error rate is computed on the mapping between estimated transcriptions and the ground-truth word sequences of the different speakers that results in the lowest word error rate.

Note that the cpWER also counts recognition errors caused by wrong diarization. For example, if the diarization swaps the speaker identity of an utterance, this adds to the deletion error count of the correct speaker’s transcription, and, furthermore, adds to the insertion error count of the wrongly identified speaker, regardless of whether the recognized words were correct or not.

Since we are interested in knowing the contribution of diarization errors, we also computed the assignment k_b of recognized segment transcriptions to the ground-truth transcription streams such that the cpWER is as small as possible. In other words, we use as the speaker label for each segment b the label that minimizes the cpWER, instead of the estimated speaker label \hat{k}_b that was obtained from $\hat{A}_{t,k}$, while keeping the estimated segment start and end times. This is similar to the MIMO-WER defined in [39], but ground-truth and estimate are swapped. The error rate computed in this way ignores contributions from diarization errors, hence we call it diarization invariant concatenated minimum-permutation word error rate (DI-cpWER). The difference between the cpWER and DI-cpWER can then be interpreted as the errors caused by diarization mistakes.

In the literature, the asclite tool is often used to compute the WER of multi-speaker transcriptions [40]. It ignores diarization errors as well, but is limited to shorter segments and fewer estimated “speakers”. It has been used on LibriCSS in the past, however with so-called continuous speech separation (CSS) systems that separate the data into two streams where each stream contains the signals of several speakers who do not overlap [3]. In our setup, we produce eight streams, one for each speaker in the meeting, and because the computation and memory demands of asclite increase exponentially with the number of streams, it is computationally infeasible to use asclite for our system.

Both asclite and DI-cpWER calculate a WER where diarization errors are ignored. The numbers obtained are thus roughly comparable, and we shall use speaker agnostic WER (SAg-WER) as a general name for those WERs that ignore diarization errors.

We used two different pretrained ASR systems from the ESPnet framework [41] to estimate the transcription. Both are used “out of the box” without any fine-tuning. The first, which we refer to as “base”⁵ [42], has a transformer architecture. It is designed and trained on LibriSpeech and achieves a WER of 2.7% on the clean test set of LibriSpeech. The second, which

⁴https://github.com/jsalt2020-asrdiar/jsalt2020_simulate

⁵<https://zenodo.org/record/3966501>

we refer to as “WavLM”⁶ [43], is a conformer based ASR system that utilizes WavLM [44] features. It is designed for CHiME-4 [45] and uses Libri-Light [46], GigaSpeech [47], VoxPopuli [48], CHiME-4 [45], and WSJ0/1 [49] for training. On the clean test set of LibriSpeech, it achieves a WER of 1.9 %.

The results in the following tables have been obtained using all seven channels of the LibriCSS dataset, unless otherwise stated. Note that masks are estimated for each channel independently, and the median is then applied across the channels to obtain the final mask, see Fig. 2.

C. Preprocessing

The speech data is transformed to the STFT domain with a 64 ms window and 16 ms shift. In the STFT domain, the logarithmic spectrogram and mel frequency cepstral coefficients (MFCC) are stacked as input features for the NN.

For the computation of the speaker embeddings, we followed [18]: at training time, we used oracle annotations to identify single-speaker regions, from which I-vectors are estimated; at test time, we employed an X-vector extractor and spectral clustering (SC) from the ESPnet LibriCSS recipe [41] to obtain estimates of the annotations, to be used for I-vector extraction⁷. Note that this initial estimate is unable to identify overlapping speech, and thus overlapping speech cannot be excluded for the embedding calculation.

In Kaldi [50], and hence also in the original implementation of TS-VAD [18], it is common to compute a mean and variance vector for each data set and normalize the input feature vectors with them. We initially skipped this normalization, but we observed a serious performance degradation without this domain adaptation, with an increase of the WER from 7.7 % to 21.5 %. Therefore, for all experiments reported here, we employed moment (mean and variance) matching between the simulated validation dataset and the rerecorded DEV and EVAL LibriCSS datasets for the input features and the I-vectors.

As an additional enhancement, weighted prediction error (WPE) based dereverberation [51], [52] was carried out prior to the NN and the extraction at test time.

D. Training Details

While developing this system, we found several aspects that had a higher impact on the convergence properties or the final performance than we expected, for example reducing the time until the network starts to learn from weeks to hours. We will not provide detailed experiments for each of the modifications we found useful, but we want to present our findings for the benefit of the community.

We used one-minute-long chunks of data for the training. To reduce the dependency on the artificial speaker ordering, we followed [18] in randomly permuting the speaker embeddings at the input and reversing the permutation after the NN.

⁶https://huggingface.co/espnet/simpleoier_librispeech_asr_train_asr_conformer7_wavlm_large_raw_en_bpe5000_sp

⁷https://kaldi-asr.org/models/13/0013_librispeech_v1_extractor.tar.gz

TABLE I
COMPARISON OF OUR TS-VAD IMPLEMENTATION WITH THAT OF [20], USING VARIOUS DIARIZATION INFORMATION TO SEGMENT THE OBSERVATION AND APPLYING VARIOUS ASR SYSTEMS, WITHOUT INTERMEDIATE SEPARATION OF OVERLAPPING SPEECH. THE FIRST MICROPHONE OF LIBRICSS IS USED.

ID	Ref.	Diarization	ASR	DER		cpWER	
				DEV	EVAL	DEV	EVAL
(1)	[20]	Oracle	HMM-DNN	–	–	–	23.1
(2)	[20]	TS-VAD	HMM-DNN	–	7.6	–	25.8
(3)	[44]	Oracle	WavLM	–	–	–	6.0
(4)		Oracle	Base	0.0	0.0	19.17	19.77
(5)		SC	Base	18.93	17.93	29.07	27.77
(6)		TS-VAD	Base	6.30	5.65	21.55	21.27
(7)		TS-VAD	WavLM	6.30	5.65	10.39	9.13
(8)		Oracle	WavLM	0.0	0.0	6.79	6.01

Furthermore, we observed a better training convergence when executing the NN two times with different permutations and averaging the outputs, after inverting the permutations, before the last nonlinearity. This idea was mentioned in [18] to be used for the evaluation of some preliminary experiments.

For separation systems, we observed better convergence properties when more overlap was present in the beginning of the training. We also observed this behavior for TS-VAD based diarization, hence we used the superposition idea from [53], which is similar to mixup [54] and took, with $\zeta = 50\%$ probability, the sum of a minute of data with the following minute as training data instead of the original data. To compute the TS-VAD targets for this superposition, we used the logical OR operator. Although this superposition introduced a mismatch between training and test (e.g., self overlap and more simultaneous and total active speakers), we observed better convergence and a better final performance. We show the distribution of the number of active speakers in the training data with and without superposition in Fig. 5.

E. Postprocessing / Segmentation

For the pretrained ASR system named “base”, we observed a strange behavior for long segments⁸, hence we split long activities at silence positions, according to $\tilde{A}_{t,k}$, such that no segment is longer than 12 s. As minimum segment length, we used 40 frames (0.64 s), but this was only necessary when no overestimation was used.

F. Baselines and reference results

The public source code of TS-VAD [18] is written in Kaldi for the CHiME-6 data, and the modifications for LibriCSS [20] are not publicly available. We reimplemented TS-VAD in PyTorch [55] and tried to follow the original implementation as closely as possible. For example, for the NN layers, we used one BLSTMP⁹ [57], two BLSTMP, and one BLSTMP followed by a feedforward layer for the three NN blocks in

⁸e.g., a transcription suddenly interrupted by a stream of single-letter words consisting of the letters E, O, and T.

⁹A BLSTMP is a bidirectional long short-term memory [56] recurrent neural network with a feedforward projection layer.

TABLE II

RESULTS OF TS-SEP WITH VARIOUS ENHANCEMENT STRATEGIES, USING SEGMENTATION HYPERPARAMETERS $\text{Thr.}\tau = 0.6$, $\text{Dil.} = 161$, AND $\text{Ero.} = 81$.

ID	Extraction			DER		cpWER	
	WPE	BF	Masking (ξ)	DEV	EVAL	DEV	EVAL
(1)	–	–	–	15.47	13.89	26.53	24.42
(2)	✓	–	–	15.35	13.68	26.26	24.08
(3)	✓	–	✓(0.0)	15.35	13.68	12.42	10.49
(4)	✓	–	✓(0.5)	15.35	13.68	22.99	21.07
(5)	✓	✓	–	15.35	13.68	10.51	8.82
(6)	✓	✓	✓(0.0)	15.35	13.68	11.08	9.29
(7)	✓	✓	✓(0.01)	15.35	13.68	11.12	9.26
(8)	✓	✓	✓(0.05)	15.35	13.68	10.93	9.00
(9)	✓	✓	✓(0.2)	15.35	13.68	10.43	8.48
(10)	✓	✓	✓(0.5)	15.35	13.68	10.16	8.42
(11)	–	✓	✓(0.5)	15.47	13.89	11.31	9.46

Fig. 1. Nevertheless, they may differ in some details. We therefore first verify if our implementation is competitive. In Table I, we see that the DER and WER of our implementation is similar to the original implementation (lines (6) vs. (2)). The TS-VAD is able to improve over SC (lines (6) vs. (5)) and, with the WavLM-based ASR system, we achieved our best WER without doing any enhancement (line (7)).

G. Extraction and enhancement

In Table II, we analyze the performance of our proposed TS-SEP, which computes a mask to extract the source signals of the speakers from the observation, under various extraction strategies. We first use no enhancement (line (1)) to see the diarization performance. Note that we used a dilation of 161 frames in Table II instead of the 81 frames used in Table I. The higher DER is caused by the resulting overestimation, and the cpWER gets slightly worse compared to TS-VAD (we explain why in the next section). Including WPE-based dereverberation slightly improved the performance (line (2)), and using masking for enhancement gave a big jump down to 10.49% cpWER (line (3)). With beamforming (line (5)), we were able to further improve the cpWER to 8.82%. Doing masking after beamforming (line (6)) yielded worse results, probably due to the artifacts introduced by masking. But by selecting a proper lower threshold $\xi = 0.5$ for the mask, we were able to obtain a slight improvement and achieved 8.42% cpWER (line (10)). We also include line (11) to help illustrate the effect of WPE.

The effect of the lower threshold is in line with the observations made in [36] for speech enhancement¹⁰: for ASR, the artifacts produced by masking have a negative effect, but using a carefully chosen lower threshold for the mask improves the ASR performance.

¹⁰In [36], a scaled version of the observation is added to the estimate, which in our system corresponds to adding a constant to the mask. Note that we do not add the observation, since it contains the cross-talker and beamformed signals are known to work well for ASR [29].

TABLE III

RESULTS OF TS-SEP WITH VARIOUS SEGMENTATION HYPERPARAMETERS, USING WPE+BF+MASKING ($\xi = 0.5$) AS ENHANCEMENT.

ID	Segmentation				DER		cpWER	
	Thr.	τ	Dil.	Ero.	DEV	EVAL	DEV	EVAL
(1)	0.3	201	81	24.91	21.94	10.41	8.26	
(2)	0.3	161	81	17.23	14.69	10.06	7.72	
(3)	0.3	121	81	9.89	7.89	9.99	7.70	
(4)	0.3	81	81	7.61	6.49	12.98	11.16	
(5)	0.1	161	81	19.55	16.46	10.36	7.98	
(6)	0.2	161	81	18.22	15.38	10.27	7.79	
(8)	0.3	161	81	17.23	14.69	10.06	7.72	
(9)	0.4	161	81	16.44	14.16	9.95	7.72	
(10)	0.6	161	81	15.35	13.68	10.16	8.42	
(11)	0.3	141	61	17.14	14.61	10.01	7.72	
(12)	0.3	161	81	17.23	14.69	10.06	7.72	
(13)	0.3	181	101	17.41	14.85	10.06	7.75	

H. Segmentation

In preliminary experiments not reported here, we optimized the segmentation parameters to minimize the DER when using a TS-VAD system. With the TS-SEP system, two things changed. First, an overestimation of the interval length of activity might be non-critical for source extraction, because potentially active cross-talkers at the borders can be suppressed by the enhancement. Further, an overestimation yields utterance boundaries that better match the typical training data of ASR systems, as mentioned in Section IV-C. Second, TS-SEP is trained to predict a reconstruction mask, so the activity estimates $\tilde{A}_{t,k}$ are often smaller than with TS-VAD, especially in TF bins where the speaker is inactive or less active, as shown in Fig. 3. The TS-VAD system, on the other hand, is trained to yield values close to one for the full duration of an utterance. Because of this training mismatch, we expected TS-SEP to prefer a smaller threshold τ in Eq. (8).

We verified both hypotheses in the experiments of Table III. When optimizing the segmentation hyperparameters, the DER and WER showed opposite behavior. For the DER, as expected, an overestimation has a negative effect (lines (4) vs. (1) to (3)). For the WER, on the contrary, some overestimation is helpful, as argued in the previous paragraph. While in preliminary experiments a threshold of $\tau = 0.6$ was a good choice for TS-VAD, smaller values in the range of 0.3 to 0.4 led to best performance for TS-SEP (lines (5) to (10)). Simultaneously expanding or reducing the dilation and erosion windows had only a minor effect (lines (11) to (13)).

I. Single-channel system

In Table IV, we replaced components of our system with single-channel counterparts to end up with a single-channel system: beamforming, masking, dereverberation, and domain adaptation (DA). Using beamforming obviously requires multiple channels, so the single-channel alternative is simply not to use any beamforming. By contrast, a mask for source extraction can be computed either from a single (reference) channel or in a multi-channel fashion by applying the mask estimation NN independently to each channel and reducing the

TABLE IV

TOWARDS A SINGLE-CHANNEL SYSTEM. A ✓ INDICATES THAT A COMPONENT IS SINGLE-CHANNEL AND A – INDICATES THAT THE COMPONENT USES MULTIPLE CHANNELS. DA MEANS DOMAIN ADAPTATION, WHERE WE DISTINGUISH WHETHER IT IS DONE FOR ALL CHANNELS SIMULTANEOUSLY (–) OR PER CHANNEL (✓).

ID	WPE	Mask Reduction	Extraction	cpWER				DI-cpWER			
				DEV		EVAL		DEV		EVAL	
				DA: –	DA: ✓	DA: –	DA: ✓	DA: –	DA: ✓	DA: –	DA: ✓
(1)	–	– (median)	– (BF+Masking ($\xi = 0.5$))	10.06	9.83	7.72	7.70	5.41	5.12	5.92	5.85
(2)	–	✓ (reference)	– (BF+Masking ($\xi = 0.5$))	8.96	8.45	9.05	7.80	5.94	5.37	6.59	5.98
(3)	–	– (median)	✓ (Masking ($\xi = 0.0$))	12.14	11.90	10.11	10.00	7.55	7.27	8.34	8.20
(4)	–	✓ (reference)	✓ (Masking ($\xi = 0.0$))	11.83	11.07	12.04	10.30	8.80	8.07	9.70	8.51
(5)	✓	✓ (reference)	✓ (Masking ($\xi = 0.0$))	13.40	12.12	13.27	11.55	10.45	9.01	10.83	9.74

TABLE V

TS-VAD VS TS-SEP WITH DIFFERENT EXTRACTIONS AND ASR SYSTEMS, WHERE DIL. = 161, ERO. = 81 AND THR. IS SET TO $\tau = 0.6$ (TS-VAD) OR $\tau = 0.3$ (TS-SEP).

ID	Dia./Sep.	Extraction	ASR	DER		cpWER		DI-cpWER	
				DEV	EVAL	DEV	EVAL	DEV	EVAL
(1)	TS-VAD	BF+Masking ($\xi = 0.5$)	Base	18.72	16.28	17.26	15.98	13.22	13.87
(2)	TS-VAD	GSS T-init	Base	18.72	16.28	9.09	6.43	4.74	4.32
(3)	TS-SEP	BF+Masking ($\xi = 0.5$)	Base	17.23	14.69	10.06	7.72	5.41	5.92
(4)	TS-SEP	GSS T-init	Base	17.23	14.69	8.93	6.32	4.23	4.54
(5)	TS-SEP	GSS TF-init	Base	17.23	14.69	8.80	6.15	4.07	4.37
(6)	TS-VAD	BF+Masking ($\xi = 0.5$)	WavLM	18.72	16.28	12.13	9.50	7.95	7.35
(7)	TS-VAD	GSS T-init	WavLM	18.72	16.28	8.44	5.47	4.20	3.32
(8)	TS-SEP	BF+Masking ($\xi = 0.5$)	WavLM	17.23	14.69	8.55	5.76	3.85	3.92
(9)	TS-SEP	GSS T-init	WavLM	17.23	14.69	8.12	5.10	3.40	3.29
(10)	TS-SEP	GSS TF-init	WavLM	17.23	14.69	8.11	5.06	3.32	3.26

resulting masks to a single mask with the median operation. Similarly, WPE-based dereverberation is usually done jointly over all available channels, but can also be applied to each channel independently. For the domain adaptation, we used by default one channel of the simulated validation set to compute the training time statistics, and all channels from the rerecorded test set to compute the test time statistics. To make it single-channel, we compute the test time statistics independently on each channel. For the DEV set, we observed some outliers for the cpWER¹¹, hence we report here also the DI-cpWER, which is more robust.

Switching one component after each other to the single-channel variant increased the cpWER, except for the domain adaptation, which had a positive effect. We investigated this surprising result with more experiments. It turned out that on the LibriCSS data, which are true recordings with a circular microphone array including a microphone in the center of the circle, the reference microphone, which was the center microphone, had notably different statistics compared to the other microphones. This is different from the simulated data, which employed omnidirectional microphone characteristics, where no relevant differences in the statistics between the microphones were observable. With a completely single-channel system, we finally arrive at a cpWER of 11.55% and a DI-cpWER of 9.74%.

¹¹The DEV set was chosen for diarization-agnostic metrics, where 6 recordings are enough. On diarization-aware metrics, one wrong assignment can have a large impact because of the set's small size.




Fig. 6. TS-SEP with GSS for different numbers of guided EM steps and different amounts of temporal context, where VAD (T-init) or mask initialization (TF-init) is used for GSS.

J. TS-VAD vs. TS-SEP and mask refinement with GSS

GSS can be considered to be an alternative to TS-SEP to expand voice activity information from time to time-frequency resolution. However, it can also be used as a mask refinement operation to be applied to the TS-SEP output. In Table V, we explore these two different setups.

To start with, we compared the performance of TS-VAD with TS-SEP without postprocessing their outputs with GSS (lines (1) and (3)). We see a clear advantage of using a mask with time-frequency resolution for the cpWER (7.72% vs. 15.98%). Next, only temporal activity information (T-init) instead of a time-frequency mask are used from TS-SEP

(same for TS-VAD, naturally). Expanded by GSS to a time-frequency mask, the mask is used for source extraction (lines (2) and (4)). Both configurations benefit from GSS, however TS-VAD benefits considerably more than TS-SEP, such that the margin between them is drastically reduced to 6.43% vs. 6.32% cpWER. In contrast to TS-VAD, TS-SEP can utilize GSS as mask refinement (TF-init), which slightly improves the cpWER to 6.15%.

For GSS, we used a context length of ± 15 s, such that each segment to which the EM algorithm is applied is expanded by 30 s. While it appears that the advantage of TS-SEP over TS-VAD is almost lost if GSS is used as a postprocessing step, this conclusion should be reconsidered in light of the high computational complexity of GSS¹². In Fig. 6, we considered different initialization based on the output of TS-SEP and different context lengths, and used up to 20 guided and one non-guided EM steps, which are the default. It can be seen that mask-based initialization (TF-init) of GSS is always better than temporal activity based initialization (T-init). Further, using the mask allows one to reduce the number of EM steps without a relevant effect on the WER. That fewer EM steps are necessary is expected, since the mask is already a solid starting point.

Further, the context by which the segments are expanded can be significantly reduced with mask initialization. This can also be well explained by the ambiguity incurred if two speakers overlap for the whole duration of a segment: while a mask-based TF-level initialization can then still distinguish between the speakers, the temporal activity information is sometimes insufficient to do so, and left and right contexts need to be added to have a greater chance that the activity patterns between speakers are sufficiently different. See [7] for a detailed discussion on the virtue of context expansion for GSS.

Table V further contains results with the WavLM model, from which similar conclusions can be drawn. Using the stronger WavLM-based ASR system, we obtained our best cpWER of 5.06% for TS-SEP.

Note that the table also includes the DI-cpWER results, which is speaker agnostic and thus does not count the contribution of diarization errors to the WER. Here, the best configuration achieves a WER of 3.26%.

K. Comparison with the literature

LibriCSS is a publicly available dataset and thus allows for a comparison of our system with the performance achieved by others. We compiled some relevant results in Table VI. For completeness, we furthermore report the performance of our systems as a function of the speech overlap condition in Table VII.

In our experiments until here, we followed the convention proposed in [2] that the first session of LibriCSS is used for DEV and the remaining 9 for EVAL. For some of the results from the literature that we report here, it is not clear from the publication if this convention is followed, or if the results are

¹²For example in the CHiME-6 baseline [1], GSS used 12 microphones instead of 24 microphones, because the runtime was too large.

TABLE VI
LITERATURE COMPARISON. THE STYLE COLUMN INDICATES HOW THE SYSTEM IS DESIGNED, WHERE D, S, AND A SIGNIFY DIARIZATION, SEPARATION, AND ASR, RESPECTIVELY. AN \rightarrow INDICATES A SEQUENCE AND A + INDICATES A COUPLING BETWEEN COMPONENTS. THE \dagger INDICATES THAT THE CPWER IS CALCULATED ON ROUGHLY 1 MINUTE CHUNKS OF THE 10 MINUTE DATA.

System (other)	Style	Single Ch.	cpWER	SAg-WER asclite
CSS with DOA Dia [58]	S \rightarrow D \rightarrow A	–	12.98	–
CSS with DOA Dia [58]	S \rightarrow D \rightarrow A	–	12.40 \dagger	–
SMM [6]	D \rightarrow S \rightarrow A	–	5.9 \dagger	–
CSS \rightarrow SC [2]	S \rightarrow D \rightarrow A	–	12.7	–
t-SOT TT [59]	E2E	✓	–	7.6
CSS [60]	S \rightarrow A	✓	–	10.15
CSS [60]	S \rightarrow A	–	–	5.85
Transcribe-to-Diarize [61]	E2E	✓	11.6	–
TS-VAD + Speakerbeam [12]	D \rightarrow S \rightarrow A	✓	18.8	–
TS-VAD + GSS [12]	D \rightarrow S \rightarrow A	–	11.2	–
SC + GSS [62]	D \rightarrow S \rightarrow A	–	12.12	–
System (mixed)	Style	Single Ch.	cpWER	SAg-WER DI-cpWER
SC [62] \rightarrow GSS \rightarrow WavLM	D \rightarrow S \rightarrow A	–	13.85	8.82
System (our)	Style	Single Ch.	cpWER	SAg-WER DI-cpWER
TS-VAD \rightarrow GSS \rightarrow Base	D \rightarrow S \rightarrow A	–	6.70	4.36
TS-VAD \rightarrow WavLM	D \rightarrow A	✓	9.26	7.16
TS-VAD \rightarrow GSS \rightarrow WavLM	D \rightarrow S \rightarrow A	–	5.77	3.40
TS-SEP \rightarrow Mask. \rightarrow Base	D + S \rightarrow A	✓	11.61	9.66
TS-SEP \rightarrow GSS \rightarrow Base	D + S \rightarrow A	–	6.42	4.34
TS-SEP \rightarrow Mask. \rightarrow WavLM	D + S \rightarrow A	✓	7.81	5.80
TS-SEP \rightarrow GSS \rightarrow WavLM	D + S \rightarrow A	–	5.36	3.27

TABLE VII
BREAKDOWN OF CPWER FOR DIFFERENT OVERLAP SUBSETS, FROM NO OVERLAP WITH LONG (OL) OR SHORT (OS) PAUSES TO OVERLAP RATIOS BETWEEN 10 AND 40 PERCENT.

System	Overlap ratio in %						avg
	0L	0S	10	20	30	40	
TS-VAD \rightarrow GSS \rightarrow Base	5.76	4.34	6.73	8.12	8.04	6.53	6.70
TS-VAD \rightarrow WavLM	5.05	5.04	7.89	10.78	12.27	11.90	9.26
TS-VAD \rightarrow GSS \rightarrow WavLM	4.66	3.51	5.96	7.59	7.02	5.24	5.77
TS-SEP \rightarrow Mask. \rightarrow Base	7.91	5.57	9.34	13.27	15.09	15.58	11.61
TS-SEP \rightarrow GSS \rightarrow Base	6.15	3.71	6.32	7.90	7.12	6.80	6.42
TS-SEP \rightarrow Mask. \rightarrow WavLM	5.25	3.63	6.63	9.53	9.84	10.10	7.81
TS-SEP \rightarrow GSS \rightarrow WavLM	5.02	2.75	5.35	7.14	5.91	5.55	5.36

for DEV+EVAL. For the sake of fairness, we thus report in this table our results for DEV+EVAL, as they are in our case slightly worse than for EVAL alone.

The best WERs for LibriCSS that we found are 3.34% from Raj et al. [62] and 3.97% from Wang et al. [63]. But both used oracle diarization to obtain those numbers. Besides those experiments, Wang et al. [63] also reported, as far as we know, the best asclite-based WER on LibriCSS of 5.85% by using multiple channels. They used a CSS pipeline, where a system produced two streams from the observation, each stream containing no overlap. When they restricted the system to single-channel data, they obtained 10.15% WER. In a

side experiment of [59]¹³, Kanda et al. reported an asclite-based WER of 7.6%, only utilizing a single channel for their serialized output training based system, which yields directly the transcriptions, instead of producing an intermediate overlap-free audio signal. Later, Kanda et al. [61] proposed Transcribe-to-Diarize, which includes diarization information, and reported a cpWER of 11.6%, which is the best single-channel cpWER that we found.

In [12], Delcroix et al. reported a combination of TS-VAD and Speakerbeam for single-channel processing and obtained a cpWER of 18.8%. To interpret this WER, they also tried TS-VAD with GSS and obtained a cpWER of 11.2%, which is the best cpWER (without single-channel constraint) that we found for LibriCSS.

In [6], an SMM was applied to LibriCSS, but it had issues with the memory consumption and was therefore only applied to the subsegmented data with an average duration of around 1 min, which is typically used for CSS pipelines with asclite-based WER evaluations. The SMM obtained a cpWER of 5.9% on the subsegmented data. In [58], a direction of arrival (DOA) approach is used to do diarization after a CSS-based separation. This is one of the few publications that report cpWER on both the 1 min data and the full 10 minute long mini sessions, and thus can be used to estimate the impact of the session length on the WER performance. They observed an improvement of 0.58% when using the shorter data. An improvement is reasonable, since shorter data can compensate some diarization errors in cpWER.

By using the publicly available pretrained “base” ASR system from mid 2020, which was available for all mentioned references and used in [6], we are able to outperform or be similar to all systems, except [59] for the asclite-based WER under the single-channel constraint. By using the recently released “WavLM” based ASR, we are able to outperform all systems. Note that the combination of our reimplemented TS-VAD and the pretrained “WavLM”, without any enhancements, is a single-channel system and has a better cpWER on the 10-minute data than all references, regardless of whether they use single- or multi-channel data.

The diarization estimates of [62] are publicly available¹⁴. We applied GSS and the “WavLM” based ASR to them and observed a degradation of the cpWER, which highlights the strength of our diarization performance.

Interestingly, we also obtained some WERs that are better than those from the literature while solving a more difficult task. Without using oracle diarization, we were able to obtain a speaker agnostic WER of 3.27%, while the so far best WER is 3.34% [62], where oracle diarization was used. While [63] uses a speaker agnostic WER to obtain 5.85%, our proposed system is able to achieve a WER of 5.36% while additionally counting errors caused by diarization.

VI. CONCLUSION AND OUTLOOK

By expanding the output from time to time-frequency resolution, we were able to extend the TS-VAD diarization system

to a joint diarization and separation method. The approach, which we named TS-SEP, is general in the sense that it can be used both for single- and multi-channel input data. While TS-SEP can be viewed as an alternative to GSS, it can also be combined with it for further performance improvement on multi-channel data. We discussed various training details of TS-VAD and TS-SEP, and presented an extensive experimental evaluation on the LibriCSS meeting data set. Performance was measured in terms of diarization error rate and in terms of both diarization dependent and independent word error rates, to assess the impact of diarization errors on ASR. Obtaining 11.61% and 6.42% for the single- and multi-channel cases with a “base” ASR model, and 7.81% and 5.36% with WavLM features for ASR, we were able to achieve a new state of the art in diarization dependent WER. The best result translates into a WER of 3.27% for diarization error independent evaluation.

Going forward, we plan to apply TS-SEP to real recordings such as CHiME-6. Since TS-SEP can also be trained with the TS-VAD objective, a hybrid training is possible, where we simultaneously train with the VAD objective on real data and with a separation objective on simulated data, in order to enable training on real data for separation without a separation target.

ACKNOWLEDGMENTS

Part of this work was done while Christoph Boeddeker was an intern at Mitsubishi Electric Research Laboratories (MERL). He was later funded by Deutsche Forschungsgemeinschaft (DFG), project no. 448568305. Computational resources were provided by the Paderborn Center for Parallel Computing.

REFERENCES

- [1] S. Watanabe, M. Mandel, J. Barker, E. Vincent, A. Arora, X. Chang, S. Khudanpur, V. Manohar, D. Povey, D. Raj, D. Snyder, A. S. Subramanian, J. Trmal, B. B. Yair, C. Boeddeker, Z. Ni, Y. Fujita, S. Horiguchi, N. Kanda, T. Yoshioka, and N. Ryant, “CHiME-6 challenge: Tackling multispeaker speech recognition for unsegmented recordings,” in *Proc. CHiME-6 Workshop*, 2020.
- [2] D. Raj, P. Denisov, Z. Chen, H. Erdogan, Z. Huang, M. He, S. Watanabe, J. Du, T. Yoshioka, Y. Luo, N. Kanda, J. Li, S. Wisdom, and J. R. Hershey, “Integration of speech separation, diarization, and recognition for multi-speaker meetings: System description, comparison, and analysis,” in *Proc. IEEE SLT*, 2021, pp. 897–904.
- [3] Z. Chen, T. Yoshioka, L. Lu, T. Zhou, Z. Meng, Y. Luo, J. Wu, X. Xiao, and J. Li, “Continuous speech separation: Dataset and analysis,” in *Proc. IEEE ICASSP*, 2020, pp. 7284–7288.
- [4] N. Ryant, K. Church, C. Cieri, A. Cristia, J. Du, S. Ganapathy, and M. Liberman, “The second DIHARD diarization challenge: Dataset, task, and baselines,” in *Proc. ISCA Interspeech*, 2019, pp. 978–982.
- [5] N. Ito, S. Araki, and T. Nakatani, “Permutation-free convolutive blind source separation via full-band clustering based on frequency-independent source presence priors,” in *Proc. IEEE ICASSP*, 2013, pp. 3238–3242.
- [6] C. Boeddeker, T. Cord-Landwehr, T. von Neumann, and R. Haeb-Umbach, “An initialization scheme for meeting separation with spatial mixture models,” in *Proc. ISCA Interspeech*, 2022, pp. 271–275.
- [7] C. Boeddeker, J. Heitkaemper, J. Schmalenstroer, L. Drude, J. Heymann, and R. Haeb-Umbach, “Front-end processing for the CHiME-5 dinner party scenario,” in *Proc. CHiME-5 Workshop*, 2018.
- [8] I. Medennikov, M. Korenevsky, T. Prisyach, Y. Khokhlov, M. Korenevskaya, I. Sorokin, T. Timofeeva, A. Mitrofanov, A. Andrusenko, I. Podluzhny, A. Laptev, and A. Romanenko, “The STC system for the CHiME-6 challenge,” in *Proc. CHiME-6 Workshop*, 2020, pp. 36–41.

¹³Visible in appendix of arXiv version at <https://arxiv.org/abs/2202.00842>

¹⁴<https://github.com/desh2608/diarizer>

- [9] A. Eisenberg, B. Schwartz, and S. Gannot, "Online blind audio source separation using recursive expectation-maximization," in *Proc. ISCA Interspeech*, 2021, pp. 3480–3484.
- [10] K. Žmolíková, M. Delcroix, K. Kinoshita, T. Higuchi, A. Ogawa, and T. Nakatani, "Speaker-aware neural network based beamformer for speaker extraction in speech mixtures," in *Proc. ISCA Interspeech*, 2017, pp. 2655–2659.
- [11] Q. Wang, H. Muckenhirn, K. Wilson, P. Sridhar, Z. Wu, J. R. Hershey, R. A. Saurous, R. J. Weiss, Y. Jia, and I. L. Moreno, "Voicefilter: Targeted voice separation by speaker-conditioned spectrogram masking," in *Proc. ISCA Interspeech*, 2019, pp. 2728–2732.
- [12] M. Delcroix, K. Zmolikova, T. Ochiai, K. Kinoshita, and T. Nakatani, "Speaker activity driven neural speech extraction," in *Proc. IEEE ICASSP*, 2021, pp. 6099–6103.
- [13] Y. Fujita, N. Kanda, S. Horiguchi, K. Nagamatsu, and S. Watanabe, "End-to-end neural speaker diarization with permutation-free objectives," in *Proc. ISCA Interspeech*, 2019, pp. 4300–4304.
- [14] S. Horiguchi, Y. Fujita, S. Watanabe, Y. Xue, and K. Nagamatsu, "End-to-end speaker diarization for an unknown number of speakers with encoder-decoder based attractors," in *Proc. ISCA Interspeech*, 2020, pp. 269–273.
- [15] S. Horiguchi, S. Watanabe, P. García, Y. Xue, Y. Takashima, and Y. Kawaguchi, "Towards neural diarization for unlimited numbers of speakers using global and local attractors," in *Proc. IEEE ASRU*, 2021, pp. 98–105.
- [16] K. Kinoshita, T. von Neumann, M. Delcroix, C. Boeddeker, and R. Haeb-Umbach, "Utterance-by-utterance overlap-aware neural diarization with Graph-PIT," in *Proc. ISCA Interspeech*, 2022, pp. 1486–1490.
- [17] K. Kinoshita, M. Delcroix, and N. Tawara, "Integrating end-to-end neural and clustering-based diarization: Getting the best of both worlds," in *Proc. IEEE ICASSP*, 2021, pp. 7198–7202.
- [18] I. Medennikov, M. Korenevsky, T. Prisyach, Y. Khokhlov, M. Korenevskaya, I. Sorokin, T. Timofeeva, A. Mitrofanov, A. Andrusenko, I. Podluzhny, A. Laptev, and A. Romanenko, "Target-speaker voice activity detection: A novel approach for multi-speaker diarization in a dinner party scenario," in *Proc. ISCA Interspeech*, 2020, pp. 274–278.
- [19] S. Ding, Q. Wang, S.-Y. Chang, L. Wan, and I.-L. Moreno, "Personal VAD: Speaker-conditioned voice activity detection," in *Proc. Odyssey The Speaker and Language Recognition Workshop*, 2020, pp. 433–439.
- [20] M. He, D. Raj, Z. Huang, J. Du, Z. Chen, and S. Watanabe, "Target-speaker voice activity detection with improved i-vector estimation for unknown number of speaker," in *Proc. ISCA Interspeech*, 2021, pp. 3555–3559.
- [21] D. Wang, X. Xiao, N. Kanda, T. Yoshioka, and J. Wu, "Target speaker voice activity detection with transformers and its integration with end-to-end neural diarization," in *Proc. IEEE ICASSP*. IEEE, 2023.
- [22] X. Xiao, Z. Chen, T. Yoshioka, H. Erdogan, C. Liu, D. Dimitriadis, J. Droppo, and Y. Gong, "Single-channel speech extraction using speaker inventory and attention network," in *Proc. IEEE ICASSP*, 2019, pp. 86–90.
- [23] 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.
- [24] S. Maiti, Y. Ueda, S. Watanabe, C. Zhang, M. Yu, S.-X. Zhang, and Y. Xu, "EEND-SS: Joint end-to-end neural speaker diarization and speech separation for flexible number of speakers," in *Proc. IEEE SLT*, 2023, pp. 480–487.
- [25] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *Proc. IEEE ICASSP*, 2018, pp. 5329–5333.
- [26] T. J. Park, K. J. Han, M. Kumar, and S. Narayanan, "Auto-tuning spectral clustering for speaker diarization using normalized maximum eigengap," *IEEE Signal Process. Lett.*, vol. 27, pp. 381–385, 2019.
- [27] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 19, no. 4, pp. 788–798, 2010.
- [28] Z.-Q. Wang, J. Le Roux, and J. R. Hershey, "Multi-channel deep clustering: Discriminative spectral and spatial embeddings for speaker-independent speech separation," in *Proc. IEEE ICASSP*, 2018, pp. 1–5.
- [29] J. Heymann, L. Drude, and R. Haeb-Umbach, "Neural network based spectral mask estimation for acoustic beamforming," in *Proc. IEEE ICASSP*, 2016, pp. 196–200.
- [30] J. Heitkaemper, D. Jakobeit, C. Boeddeker, L. Drude, and R. Haeb-Umbach, "Demystifying TasNet: A dissecting approach," in *Proc. IEEE ICASSP*, 2020, pp. 6359–6363.
- [31] T. Cord-Landwehr, C. Boeddeker, T. Von Neumann, C. Zorilä, R. Dodipatla, and R. Haeb-Umbach, "Monaural source separation: From anechoic to reverberant environments," in *Proc. IEEE IWAENC*, 2022.
- [32] 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 ICASSP*, 2022, pp. 6022–6026.
- [33] R. M. Haralick, S. R. Sternberg, and X. Zhuang, "Image analysis using mathematical morphology," *IEEE Trans. Pattern Anal. Mach. Intell.*, no. 4, pp. 532–550, 1987.
- [34] M. Souden, J. Benesty, and S. Affes, "On optimal frequency-domain multichannel linear filtering for noise reduction," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 18, no. 2, pp. 260–276, 2009.
- [35] A. S. Subramanian, C. Weng, M. Yu, S.-X. Zhang, Y. Xu, S. Watanabe, and D. Yu, "Far-field location guided target speech extraction using end-to-end speech recognition objectives," in *Proc. IEEE ICASSP*, 2020, pp. 7299–7303.
- [36] K. Iwamoto, T. Ochiai, M. Delcroix, R. Ikeshita, H. Sato, S. Araki, and S. Katagiri, "How bad are artifacts?: Analyzing the impact of speech enhancement errors on ASR," in *Proc. Interspeech 2022*, 2022, pp. 5418–5422.
- [37] J. Barker, S. Watanabe, E. Vincent, and J. Trmal, "The fifth 'CHiME' speech separation and recognition challenge: dataset, task and baselines," in *Proc. ISCA Interspeech*, 2018.
- [38] O. Sadjadi, C. Greenberg, E. Singer, L. Mason, and D. Reynolds, "NIST 2021 speaker recognition evaluation plan," 2021.
- [39] T. von Neumann, C. Boeddeker, K. Kinoshita, M. Delcroix, and R. Haeb-Umbach, "On word error rate definitions and their efficient computation for multi-speaker speech recognition systems," in *Proc. IEEE ICASSP*, 2023.
- [40] J. G. Fiscus, J. Ajot, N. Radde, and C. Laprun, "Multiple dimension Levenshtein edit distance calculations for evaluating automatic speech recognition systems during simultaneous speech," in *Proc. LREC*, 2006, pp. 803–808.
- [41] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. Enrique Yalta Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, "ESPnet: End-to-end speech processing toolkit," in *Proc. ISCA Interspeech*, 2018, pp. 2207–2211.
- [42] S. Watanabe, "ESPnet2 pretrained model, Shinji Watanabe/librispeech_asr_train_asr_transformer_e18_raw_bpe_sp_valid_acc.best, fs=16k, lang=en," 2020. [Online]. Available: <https://doi.org/10.5281/zenodo.3966501>
- [43] X. Chang, T. Maekaku, Y. Fujita, and S. Watanabe, "End-to-end integration of speech recognition, speech enhancement, and self-supervised learning representation," in *Proc. ISCA Interspeech*, 2022, pp. 3819–3823.
- [44] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, X. Yu, and F. Wei, "WavLM: Large-scale self-supervised pre-training for full stack speech processing," *IEEE J. Sel. Top. Signal Process.*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [45] E. Vincent, S. Watanabe, A. A. Nugraha, J. Barker, and R. Marxer, "An analysis of environment, microphone and data simulation mismatches in robust speech recognition," *Comput. Speech Lang.*, vol. 46, pp. 535–557, 2017.
- [46] J. Kahn, M. Riviere, W. Zheng, E. Kharitonov, Q. Xu, P.-E. Mazaré, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A.-r. Mohamed, and E. Dupoux, "Libri-Light: A benchmark for ASR with limited or no supervision," in *Proc. IEEE ICASSP*, 2020, pp. 7669–7673.
- [47] G. Chen, S. Chai, G.-B. Wang, J. Du, W.-Q. Zhang, C. Weng, D. Su, D. Povey, J. Trmal, J. Zhang, M. Jin, S. Khudanpur, S. Watanabe, S. Zhao, W. Zou, X. Li, X. Yao, Y. Wang, Z. You, and Z. Yan, "GigaSpeech: An evolving, multi-domain ASR corpus with 10,000 hours of transcribed audio," in *Proc. ISCA Interspeech*, 2021, pp. 3670–3674.
- [48] C. Wang, M. Riviere, A. Lee, A. Wu, C. Talnikar, D. Haziza, M. Williamson, J. Pino, and E. Dupoux, "VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation," in *Proc. ACL*. Association for Computational Linguistics, Aug. 2021, pp. 993–1003.
- [49] D. B. Paul and J. Baker, "The design for the Wall Street Journal-based CSR corpus," in *Proc. DARPA Speech and Natural Language Workshop*, 1992.
- [50] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Vesely, "The Kaldi speech recognition toolkit," in *Proc. IEEE ASRU*, 2011.

- [51] T. Nakatani, T. Yoshioka, K. Kinoshita, M. Miyoshi, and B.-H. Juang, "Speech dereverberation based on variance-normalized delayed linear prediction," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 18, no. 7, pp. 1717–1731, 2010.
- [52] L. Drude, J. Heymann, C. Boeddeker, and R. Haeb-Umbach, "NARA-WPE: A python package for weighted prediction error dereverberation in numpy and tensorflow for online and offline processing," in *Speech Communication; 13th ITG-Symposium*. VDE, 2018, pp. 1–5.
- [53] J. Ebberts and R. Haeb-Umbach, "Self-trained audio tagging and sound event detection in domestic environments," in *Proc. DCASE*, 2021.
- [54] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *Proc. ICLR*, 2018.
- [55] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimeshain, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 8024–8035.
- [56] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [57] H. Sak, A. Senior, and F. Beaufays, "Long short-term memory based recurrent neural network architectures for large vocabulary speech recognition," *arXiv preprint arXiv:1402.1128*, 2014.
- [58] Z.-Q. Wang and D. Wang, "Localization based sequential grouping for continuous speech separation," in *Proc. IEEE ICASSP*, 2022, pp. 281–285.
- [59] N. Kanda, J. Wu, Y. Wu, X. Xiao, Z. Meng, X. Wang, Y. Gaur, Z. Chen, J. Li, and T. Yoshioka, "Streaming multi-talker ASR with token-level serialized output training," in *Proc. ISCA Interspeech*, 2022, pp. 3774–3778.
- [60] Z.-Q. Wang, P. Wang, and D. Wang, "Multi-microphone complex spectral mapping for utterance-wise and continuous speech separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 2001–2014, 2021.
- [61] N. Kanda, X. Xiao, Y. Gaur, X. Wang, Z. Meng, Z. Chen, and T. Yoshioka, "Transcribe-to-diarize: Neural speaker diarization for unlimited number of speakers using end-to-end speaker-attributed ASR," in *Proc. IEEE ICASSP*, 2022, pp. 8082–8086.
- [62] D. Raj, D. Povey, and S. Khudanpur, "GPU-accelerated guided source separation for meeting transcription," in *Proc. ISCA Interspeech*, 2023, pp. 3507–3511.
- [63] Z.-Q. Wang, P. Wang, and D. Wang, "Multi-microphone complex spectral mapping for utterance-wise and continuous speech separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 2001–2014, 2021.



USA, respectively.

Christoph Boeddeker received the bachelor's and master's degrees in electrical engineering from Paderborn University, where he is currently working toward the Ph.D. degree, under the supervision of Reinhold Haeb-Umbach. His research interests range from multichannel speech separation, beamforming, and dereverberation to automatic speech recognition on meetings with a focus on combining statistical models and neural networks. In 2017 and 2022 he pursued a research internship with Microsoft Research, Redmond, USA and MERL, Cambridge,



Aswin Shanmugam Subramanian received the M.S.E. and Ph.D. degrees from Johns Hopkins University, Baltimore, MD, USA, in 2017 and 2022, respectively. From 2021 to 2022, he was a Research Scientist with Mitsubishi Electric Research Laboratories, Cambridge, MA, USA. He is currently a Sr. Applied Scientist with Microsoft Corporation, Redmond, WA, USA. His research interests include automatic speech recognition, microphone array processing, source separation, and speaker diarization.



Gordon Wichern is a Senior Principal Research Scientist at Mitsubishi Electric Research Laboratories (MERL) in Cambridge, Massachusetts. He received his B.Sc. and M.Sc. degrees from Colorado State University in electrical engineering and his Ph.D. from Arizona State University in electrical engineering with a concentration in arts, media and engineering, where he was supported by a National Science Foundation (NSF) Integrative Graduate Education and Research Traineeship (IGERT) for his work on environmental sound recognition. He was previously a member of the research team at iZotope, inc. where he focused on applying novel signal processing and machine learning techniques to music and post production software, and a member of the Technical Staff at MIT Lincoln Laboratory where he worked on radar signal processing. He is the Chair of the AES Technical Committee on Machine Learning and Artificial Intelligence (TC-MLAI), and a member of the IEEE Audio and Acoustic Signal Processing Technical Committee (AASP-TC).



Reinhold Haeb-Umbach is a professor of Communications Engineering at Paderborn University, Germany. He holds a Dr.-Ing. degree from RWTH Aachen University, and he has more than 30 years of experience in speech research, which he acquired both in an industrial and academic environment. His main research interests are in speech enhancement, acoustic beamforming and source separation, as well as automatic speech recognition and unsupervised learning from speech and audio. He has more than 300 scientific publications, and his students have

received several best student paper awards. From 2015 to 2020 he was member of the IEEE Signal Processing Society Speech and Language Technical Committee, and he is currently (2021 to 2023) member of the Audio and Acoustic Signal Processing Technical Committee. He is a fellow of the International Speech Communication Association (ISCA) and a fellow of the IEEE.



Jonathan Le Roux is a Distinguished Research Scientist and the Speech and Audio Senior Team Leader at Mitsubishi Electric Research Laboratories (MERL) in Cambridge, Massachusetts. He completed his B.Sc. and M.Sc. degrees in Mathematics at the Ecole Normale Supérieure (Paris, France), his Ph.D. degree at the University of Tokyo (Japan) and the Université Pierre et Marie Curie (Paris, France), and worked as a postdoctoral researcher at NTT's Communication Science Laboratories from 2009 to 2011. His research interests are in signal processing

and machine learning applied to speech and audio. He has contributed to more than 140 peer-reviewed papers and 30 granted patents in these fields. From 2014 to 2019, he was a member of the IEEE Signal Processing Society Audio and Acoustic Signal Processing Technical Committee (AASP-TC). He is a founder and chair of the Speech and Audio in the Northeast (SANE) series of workshops, and a Fellow of the IEEE.