



HAL
open science

Joint Beamforming and Speaker-Attributed ASR for Real Distant-Microphone Meeting Transcription

Can Cui, Imran Ahamad Sheikh, Mostafa Sadeghi, Emmanuel Vincent

► **To cite this version:**

Can Cui, Imran Ahamad Sheikh, Mostafa Sadeghi, Emmanuel Vincent. Joint Beamforming and Speaker-Attributed ASR for Real Distant-Microphone Meeting Transcription. 2023. hal-04755558

HAL Id: hal-04755558

<https://hal.science/hal-04755558v1>

Preprint submitted on 29 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

JOINT BEAMFORMING AND SPEAKER-ATTRIBUTED ASR FOR REAL DISTANT-MICROPHONE MEETING TRANSCRIPTION

Can Cui^{1,2}, Imran Sheikh², Mostafa Sadeghi¹, Emmanuel Vincent¹

¹Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

²Vivoka, Metz, France

ABSTRACT

Distant-microphone meeting transcription is a challenging task. State-of-the-art end-to-end speaker-attributed automatic speech recognition (SA-ASR) architectures lack a multichannel noise and reverberation reduction front-end, which limits their performance. In this paper, we introduce a joint beamforming and SA-ASR approach for real meeting transcription. We first describe a data alignment and augmentation method to pretrain a neural beamformer on real meeting data. We then compare fixed, hybrid, and fully neural beamformers as front-ends to the SA-ASR model. Finally, we jointly optimize the fully neural beamformer and the SA-ASR model. Experiments on the real AMI corpus show that, while state-of-the-art multi-frame cross-channel attention based channel fusion fails to improve ASR performance, fine-tuning SA-ASR on the fixed beamformer’s output and jointly fine-tuning SA-ASR with the neural beamformer reduce the word error rate by 8% and 9% relative, respectively.

Index Terms— Beamforming, delay-and-sum, FaSNet, speaker-attributed ASR, joint optimization

1. INTRODUCTION

Transcription of real distant-microphone conversational meetings or domestic data is an active research area [1–3]. It remains challenging today due to noise, reverberation, and overlapping speech. To improve performance, many studies have employed a front-end multichannel speech separation module or a series of (fixed, statistical, or neural) beamformers steered toward the speakers to extract individual speech signals from the overlapping speech mixture and subsequently feed each of them to a single-speaker ASR module [4–6]. The separation error then propagates to the ASR module. Later studies [7–9] have proposed to back-propagate the ASR training losses for all speakers to the front-end separation module using a permutation invariant training (PIT) criterion to optimize the two modules jointly. However, the system can often handle only a fixed number of speakers.

End-to-end multi-speaker ASR based on serialized output training (SOT) [10] addresses these shortcomings. The work in [11] introduced an end-to-end Transformer-based

speaker-attributed ASR (SA-ASR) system for joint recognition of speech and speaker identities from single-channel log Mel features. It was then extended to Multichannel SA-ASR (MC-SA-ASR) by fusing log Mel [12] and phase [13] features across channels with time-varying multi-frame cross-channel attention (MFCCA) weights. Such multichannel attention schemes are often believed to outperform classical beamforming and to be state-of-the-art. Yet, in contrast to a frequency-dependent complex-valued beamformer, they rely on frequency-independent real-valued weights, which achieve limited noise and reverberation reduction.

In this paper, we propose to combine SA-ASR with a beamforming-based noise and reverberation reduction front-end to improve speech and speaker recognition in far-field conditions. The beamformer fuses the mixture channels into a single-channel enhanced mixture fed to SA-ASR. While such a front-end is common in single-speaker scenarios, extending it to real multi-speaker scenarios is nontrivial. First, the beamformer must vary its spatial response over time according to the speakers’ positions and activity patterns. This is why few examples of multi-speaker beamforming front-ends are found in the literature. BeamformIt [14] was used as a front-end to PIT-based multi-speaker ASR [7] as well as to single-speaker ASR baselines for multi-speaker ASR tasks [1, 15, 16], while minimum variance distortion-less response (MVDR) beamforming was used as a front-end to SA-ASR in [17] without comparison to MFCCA. To the best of our knowledge, no full-neural beamforming front-end has been used for SA-ASR so far. Second, the pretraining of a neural beamformer on real meeting data is challenging due to the absence of ground truth noiseless dry mixture signals as pretraining targets.

The contributions of this paper are as follows. First, we introduce data alignment and augmentation techniques to pretrain a multi-speaker neural beamformer on a real meeting corpus containing both distant microphone and headset recordings. Note that the beamformer is employed to reduce noise and reverberation, but not to extract the individual speakers. Second, we propose a pipeline integrating beamforming with SA-ASR, aiming to improve both speech and speaker recognition. Third, we evaluate the differences in performance between statistical, hybrid, and neural beamformers. Finally, we jointly optimize the neural beamformer

and the SA-ASR model. Our experiments on the AMI corpus [18] reveal that, while MFCCA-based channel fusion does not improve ASR performance, fine-tuning SA-ASR on the fixed beamformer’s output and jointly fine-tuning SA-ASR with the neural beamformer reduces the WER by 8% and 9% relative, respectively.

The paper is organized as follows. Section 2 presents the considered beamformers and SA-ASR model. Section 3 introduces our joint system and the AMI data alignment and augmentation pipeline. Section 4 describes our experimental setup and results, and we conclude in Section 5.

2. BACKGROUND

2.1. Delay-and-sum beamforming

The delay-and-sum (DAS) beamformer [19] is a fixed beamformer, which depends only on the delays between the microphone signals and a reference microphone. It was employed as a front-end in many single-speaker ASR studies, e.g., [20]. It involves computing the delays using a time difference of arrival (TDOA) estimator such as the generalized cross-correlation with phase transform (GCC-PHAT) [21], shifting the phase of the microphone signals accordingly in the complex short-time Fourier transform (STFT) domain, and summing them. Given the microphone signals $x_1(f, t), \dots, x_M(f, t)$ at time t and frequency f , the output signal $y(f, t)$ is expressed as

$$y(f, t) = \sum_{m=1}^M x_i(f, t) e^{2j\pi d_m f} \quad (1)$$

where d_m is the delay on the m -th channel.

2.2. Deep neural network-based Minimum Variance Distortionless Response Beamforming

Deep Neural Network (DNN)-based Minimum Variance Distortionless Response (MVDR) beamforming [22, 23] combines neural networks with traditional beamforming methods. The DNN is trained to estimate masks in the time-frequency domain that enhance desired signals and suppress interference. This information is then used to compute the MVDR beamformer weights, which minimize output power while preserving signals from the target direction. This method can be seen as a transition between traditional mathematical beamforming and fully neural network-based approaches.

2.3. Filter-and-Sum Network with Transform-average-concatenate

The Filter-and-Sum Network (FaSNet) system [24] aims to directly estimate time-domain beamforming filters. It employs a two-stage design: the first stage estimates filters for

a reference microphone, and the second stage estimates filters for the remaining microphones based on pairwise cross-channel features between the pre-separated output and each microphone. The FaSNet architecture utilizes dual-path recurrent neural networks (DPRNNs) [25] to extract information from both the channel and frame levels.

The Transform-average-concatenate (TAC) [26] design paradigm addresses channel permutation and is capable of handling various numbers of microphones.

2.4. Dereverberation with Weighted Prediction Error

Multichannel dereverberation using Weighted Prediction Error (WPE) [27] reduces reverberation by modeling and subtracting late reverberant components from the observed audio signal using linear prediction. This technique optimizes prediction coefficients and error weights to enhance speech clarity and intelligibility in reverberant environments. Multiple studies [28, 29] have shown that using WPE loss in conjunction with beamforming can enhance audio quality, thereby improving the performance of speech recognition or speaker identification.

2.5. Speaker-attributed ASR

A Transformer-based end-to-end speaker-attributed ASR (SA-ASR) system was proposed in [11]. Following the SOT principle [10], the output is the concatenation of all speakers’ sentences in first-in-first-out order, where each token is associated with one speaker ID and distinct speakers are separated by an <sc> token. As shown in Fig. 1, the inputs to the model consist of an acoustic (log Mel filterbank) feature sequence and a set of reference speaker embeddings. A Conformer-based ASR Encoder first encodes acoustic information, along with a Speaker Encoder to encode speaker information. Then, the Transformer-based ASR Decoder and Speaker Decoder modules decode text and speaker information, respectively. The Speaker Decoder generates a speaker representation corresponding to each token in the ASR Decoder’s output token sequence. This representation is used to assign speakers by computing a dot product with the reference speaker embeddings.




Fig. 1: Diagram of SA-ASR [11].

3. PROPOSED METHODS

3.1. Real meeting data alignment and augmentation for neural beamformer training

Neural beamforming on real-world far-field data is challenging due to the lack of ground truth enhanced signals for training. Real meeting corpora as AMI include headset recordings for each speaker, but these can't be used directly as ground truth because of variable delays caused by the speakers' positions. To address this, we generate aligned array and headset signals in three steps (see Fig. 2): (a) extract all non-overlapping speech segments for each speaker based on dataset annotations; (b) align each non-overlapping headset segment with the corresponding array segment using matched filters, then cut them into fixed-length clips; (c) randomly sample and mix array clips from different speakers to create far-field mixtures, and mix the corresponding aligned headset clips to obtain the ground-truth enhanced mixtures.

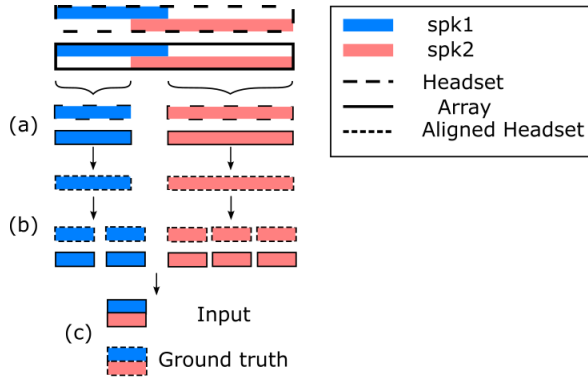


Fig. 2: Mixture generation from real meeting data.

The matched filters $f_{ij}(t)$ in step (b) are calculated in the least squares sense by solving

$$\min_{f_{ij}} \sum_t (f_{ij} \star h_j(t) - x_i(t))^2 \quad (2)$$

where $h_j(t)$ and $x_i(t)$ stand for the headset signal of speaker j and the array signal at microphone i , respectively, and \star denotes time-domain convolution. The solution is classically obtained as the finite impulse response (FIR) Wiener filter,¹ which is commonly employed for filter estimation [30, 31].

3.2. Joint multichannel beamforming and SA-ASR

We propose a joint system integrating beamforming and SA-ASR for multichannel, distant-microphone meeting transcription. As illustrated in Fig. 3, the multichannel audio is first processed by a beamformer to generate enhanced single-channel audio. The output audio is then fed to SA-ASR to obtain speech and speaker recognition results. We compare the performance of the fixed DAS beamformer, the hybrid

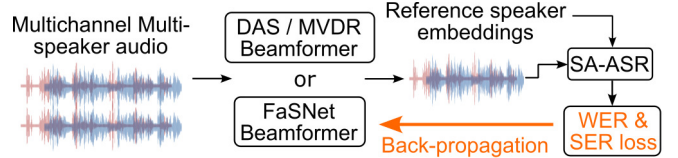


Fig. 3: Proposed joint system of Beamformer and SA-ASR

DNN-MVDR (noted as MVDR) beamformer, and the fully neural FaSNet beamformer, when fine-tuning the SA-ASR model on training data enhanced with the respective beamformer. In addition, we backpropagate the loss from SA-ASR to FaSNet, in order to fine-tune the neural beamformer according to the SA-ASR training objective.

4. EXPERIMENTAL EVALUATION

4.1. Datasets

Mixed AMI — To train the MVDR and the FaSNet beamformer, we apply the method described in Section 3.1 to the AMI meeting corpus. This method creates mixtures of real single-speaker AMI segments and their corresponding ground truths. We name this dataset Mixed AMI. We only use one-quarter of all meetings and fix the clip length to 4 s. The mixtures are generated by overlapping randomly selected clips from 1 to 4 speakers. The training, development, and test sets contain 150 h, 17 h, and 16 h of speech, respectively.

Multi-speaker LibriSpeech — To optimize performance on AMI, the SA-ASR model requires pretraining on a larger simulated distant-microphone multi-speaker dataset [32]. We created a 960 h training set and a 20 h development set from the LibriSpeech train-960 and dev-clean sets [33]. Room impulse responses are generated with gpuRIR [34] for a linear 2-, 4- or 8-microphone array with 10 cm aperture. We adopted the room simulation settings described in [13]. Each mixture contains 1 to 3 speakers. Each speaker's utterance is shifted relative to the previous speaker's by a random delay > 0.5 s. Each mixture signal is linked to 8 reference speaker embeddings obtained from two random enrollment utterances, including the actual speakers and additional random LibriSpeech speakers.

Real AMI — After it has been pretrained on Multi-speaker LibriSpeech, the SA-ASR model is fine-tuned and evaluated on real AMI multiple distant microphones (MDM) data. We utilize the segmentation method in [13] to partition the MDM data into 5 s chunks and adjust the chunk start/end times to non-overlapped word boundaries. This eliminates the need for a Voice Activity Detection model during inference. The resulting Real AMI dataset contains respectively 165 h, 19 h, and 19 h for training, development, and test. For both Mixed AMI and Real AMI, we consider 2-, 4- and 8-channel set-

¹https://en.wikipedia.org/wiki/Wiener_filter

tings. For Real AMI, the channels are taken from Array1,² namely channels 1 and 5 in the 2-channel setting, channels 1, 3, 5, and 7 in the 4-channel setting, and all 8 channels in the 8-channel setting.

4.2. Metrics

We evaluate the beamforming performance using the scale-invariant signal-to-distortion ratio (SI-SDR) and its improvement (SI-SDR_i), implemented in Asteroid toolkit [35], measured in dB on the Mixed AMI test set. We calculate the baseline SI-SDR for SA-ASR by defining the array mixture signal as the estimated signal, ensuring that the SI-SDR_i is 0 dB without beamforming. For all beamforming methods, we calculate the SI-SDR by defining the beamformed signal as the estimated signal and compute SI-SDR_i by subtracting the corresponding baseline SI-SDR. The performance of SA-ASR is evaluated by the word error rate (WER) and the sentence-level speaker error rate (SER) [36] in % on the Real AMI test set.

4.3. Model description

We utilize the implementation of DAS from the SpeechBrain toolkit [37]. For the MVDR model, we utilize the implementation in TorchAudio [22, 38], which employs the ConvTasNet [39] mask generator as the DNN module. The number of filterbank output channels and the number of bins in the estimated masks are both 513. The implementation of FaSNet with TAC is from the Asteroid toolkit. The frame size and context size of FaSNet are set to 4 and 16 ms, respectively. The encoder dimension and feature dimension are 64. The dual path blocks consist of a 4-layer dual model.

We implemented SA-ASR and the MFCCA-based MC-SA-ASR system in [13] as a baseline using SpeechBrain. In SA-ASR and MC-SA-ASR, the Conformer-based encoder, the Transformer-based decoder and the speaker decoder have 12, 6 and 2 layers, respectively. All multi-head attention mechanisms have 4 heads, the model dimension is 256, and the size of the feedforward layer is 2,048. MC-SA-ASR has fewer parameters than SA-ASR due to replacing the first (feedforward) layer of the original Conformer with a smaller-sized MFCCA. The speaker embedding model is a pretrained³ ECAPA-TDNN (Emphasized Channel Attention, Propagation, and Aggregation in Time-Delay Neural Network) [40], yielding 192 dimensional embeddings.

Additionally, we test the performance obtained with WPE, implemented by [41], during the evaluation of the MC-SA-ASR model. We also compare the performance of beamforming with and without WPE during the fine-tuning of SA-ASR.

²A few meetings have two arrays, each consisting of 8 microphones.

³Available at <https://huggingface.co/speechbrain/spkrec-ecapa-voxceleb>

4.4. Training setup

The MVDR and FaSNet beamformer are trained on Mixed AMI for 200 epochs with early stopping, using the Adam optimizer with a learning rate of 10^{-3} .

The ASR modules in SA-ASR and MC-SA-ASR are pretrained on Multi-speaker LibriSpeech for 80 epochs using Adam with a learning rate of 5×10^{-4} . The ASR and speaker modules are then further pretrained on Multi-speaker LibriSpeech for 60 epochs with a learning rate of 2.5×10^{-4} .

After this, SA-ASR and MC-SA-ASR are fine-tuned on either unprocessed (baseline) or beamformed Real AMI data for 15 epochs, using Adam with a learning rate of 3×10^{-4} . We fine-tune the ASR module for the first 8 epochs and jointly fine-tune the ASR and speaker modules in the last 7 epochs.

When jointly optimizing FaSNet and SA-ASR, the FaSNet model is not trained to converge. Instead, to assess the impact of different pretraining levels, we pretrain it for 0, 5, 10, or 50 epochs. After this, the FaSNet and SA-ASR models are jointly fine-tuned on Real AMI for 15 epochs.

4.5. Evaluation results

4.5.1. Fine-tuning SA-ASR with DAS vs. with frozen MVDR and FaSNet

We initially evaluated an SA-ASR model fine-tuned on the first channel of Real AMI. The resulting WER and SER on mixtures of 1, 2, 3, or 4 speakers were 44.54% and 34.73%, respectively. However, when testing the same model on FaSNet beamformed 2-channel Real AMI, the WER, and SER increased to 64.32% and 46.30%. This discrepancy can be attributed to the distinct acoustic characteristics between far-field and beamformed scenarios. Models trained on far-field data may lack robustness to the cleaner, less reverberant conditions of beamformed audio. Therefore, in all following experiments, we fine-tune the SA-ASR model on real AMI training data enhanced using the same beamformer as the test data. This adaptation is essential to align the model with the specific conditions of the test set.

Table 1 shows the test results of the baseline models (SA-ASR and MC-SA-ASR) and the combination of SA-ASR with three beamformers, where the parameters of MVDR and FaSNet are frozen during fine-tuning. Without WPE, the WER comparison between SA-ASR (44.54%) and MC-SA-ASR (44.99%) demonstrates that, while MFCCA had achieved a 13% relative WER reduction on simulated data in [13], it is inefficient on real meeting data. In general, fine-tuning the SA-ASR model on beamformed audio improves the ASR performance. Particularly in the 8-channel setting, using the DAS beamformer leads to a 6% relative reduction in WER without WPE (41.71%) and 8% with WPE (40.96%) compared to SA-ASR. It is also interesting to note that, despite FaSNet’s superior denoising and dereverberation performance in terms of SI-SDR_i, the SA-ASR model trained on speech beamformed by FaSNet performs less effectively

Table 1: Results for models fine-tuned and tested on unprocessed (SA-ASR and MC-SA-ASR) or beamformed (DAS-SA-ASR, MVDR-SA-ASR, FaSNet-SA-ASR) data. For convenience, we denote SI-SDR and SI-SDRi as SDR and SDRi, respectively.

System	# Prm	# Chn	Mixed AMI test set								Real AMI test set							
			1-spkr		2-spkr		3-spkr		1,2,3,4-spkr		1-spkr		2-spkr		3-spkr		1,2,3,4-spkr	
			SDR	SDRi	SDR	SDRi	SDR	SDRi	SDR	SDRi	WER	SER	WER	SER	WER	SER	WER	SER
SA-ASR	69M	1	5.41	0	5.75	0	5.79	0	5.66	0	26.76	11.92	40.23	32.66	52.31	45.11	44.54	34.73
MC-SA-ASR +WPE (test)	59M	2	5.41	0	5.75	0	5.79	0	5.66	0	26.41	11.73	40.79	32.64	52.59	43.82	44.99	34.43
		2	5.72	0.31	5.93	0.18	5.92	0.13	5.87	0.21	26.43	12.13	40.80	32.54	52.32	44.14	44.72	34.65
DAS-SA-ASR +WPE	69M	2	5.62	0.21	5.42	-0.33	5.23	-0.56	5.39	-0.27	25.59	12.82	40.36	33.87	52.04	45.55	44.03	35.56
		4	5.66	0.25	5.47	-0.28	5.25	-0.54	5.35	-0.31	24.43	12.23	39.51	33.26	50.25	42.98	42.34	34.29
		8	5.66	0.25	5.48	-0.27	5.30	-0.49	5.38	-0.28	23.51	12.13	38.41	33.12	50.43	43.44	41.71	34.40
		2	6.08	0.67	5.70	-0.05	5.40	-0.39	5.66	0.00	24.65	11.59	38.64	32.14	50.29	43.87	42.39	34.05
+WPE	69M	4	6.33	0.92	5.82	0.07	5.40	-0.39	5.67	0.01	23.50	11.26	37.22	32.00	49.34	42.34	40.96	33.33
		8	6.18	0.77	5.54	-0.21	5.04	-0.75	5.35	-0.31	23.49	11.43	37.89	32.67	50.12	43.59	41.37	33.84
MVDR-SA-ASR +WPE	74M	2	7.40	1.98	7.42	1.67	7.46	1.80	7.44	1.78	26.54	12.94	41.07	34.47	52.81	45.18	44.39	35.92
		4	7.94	2.52	7.98	2.23	7.99	2.19	7.99	2.33	27.15	12.63	41.35	34.92	52.72	44.94	44.76	35.78
		8	8.14	2.72	8.10	2.34	8.10	2.30	8.11	2.44	27.31	12.42	41.27	34.52	52.63	45.07	44.23	35.21
		2	7.75	2.34	7.48	1.73	7.48	1.69	7.55	1.89	26.40	12.78	41.09	33.66	52.00	44.89	44.29	35.54
+WPE	74M	4	8.25	2.83	8.01	2.26	7.97	2.18	8.05	2.39	26.70	13.04	41.28	34.57	52.34	44.19	44.35	35.54
		8	8.40	2.99	8.09	2.34	8.03	2.24	8.14	2.48	26.35	12.93	41.03	33.72	52.12	44.26	44.12	35.37
FaSNet-SA-ASR +WPE	72M	2	10.21	4.79	9.85	4.09	9.56	3.76	9.76	4.10	26.86	11.33	40.91	35.67	52.57	47.12	44.57	36.24
		4	10.22	4.80	9.89	4.13	9.63	3.83	9.82	4.15	26.82	12.23	40.29	34.87	51.52	45.73	43.85	35.90
		8	10.41	4.99	10.01	4.25	9.72	3.92	9.96	4.29	26.53	10.73	39.93	34.78	51.70	45.28	44.11	35.51
		2	7.35	1.93	7.18	1.42	7.04	1.24	7.16	1.50	26.75	12.40	41.07	34.75	52.38	46.16	44.48	36.01
+WPE	72M	4	6.48	1.07	6.26	0.51	6.13	0.34	6.24	0.58	25.86	12.63	39.27	33.15	51.78	44.16	43.29	34.97
		8	5.88	0.47	5.88	0.47	5.66	-0.13	5.85	0.19	26.16	10.78	39.35	34.89	51.22	45.25	43.39	35.41

Note: For all the tables, we employed the SCTK toolkit [42] to conduct significance tests, specifically the Matched Pair Sentence Segment test. For the mixture of 1,2,3 and/or 4 speakers, the best WER/SER and the results statistically equivalent to it at a 0.05 significance level are highlighted.

than the one trained on speech beamformed by DAS. In the 8-channel setting, without WPE, using DAS results in a 5% relative reduction in WER compared to using FaSNet (from 44.11% to 41.71%), and a reduction of 4% relative in SER is observed for 4-channel (from 35.90% to 34.29%). The WER relative reduction is up to 6% (from 44.23% to 41.71%) compared to the MVDR-SA-ASR system. The latter system has a similar performance to the FaSNet-SA-ASR system.

To find the reason for the difference between DAS-SA-ASR and FaSNet-SA-ASR, we visualize the spectrogram of one 8-channel Mixed AMI test chunk before and after beamforming (see Fig. 4). It can be seen that, although FaSNet exhibits effective denoising, it also removes a portion of the speech signal, as highlighted by the white columns in the figure. On the contrary, DAS can preserve a significant portion of all speech signals while providing some denoising, which results in better speech and speaker recognition results.

4.5.2. Effectiveness of adding WPE in frozen beamformers

Table 1 also shows the performance differences for each system with or without WPE for dereverberation. First, even without beamforming, using WPE only during the inference phase for MC-SA-ASR results in a 0.21 dB improvement in SI-SDR and a slight absolute WER reduction of 0.27% (from 44.99% to 44.72%). For systems using beamformed signals,




Fig. 4: Spectrogram of one 8-channel Mixed AMI test chunk. From top to bottom: 1st array channel; DAS beamformed signal; FaSNet beamformed signal; ground truth.

integrating WPE during the beamforming phase improves the SI-SDRi for DAS and MVDR but not for FaSNet. However, using WPE during beamforming to fine-tune the SA-ASR model systematically improves ASR and speaker identi-

Table 2: Results for jointly trained 2-channel FaSNet and SA-ASR models.

WPE usage	Mixed AMI test set					Real AMI test set							
	Pretrained		Fine-tuned			1-spkr		2-spkr mix		3-spkr		1,2,3,4-spkr	
	# Epo	SDR	SDRi	SDR	SDRi	WER	SER	WER	SER	WER	SER	WER	SER
None	0	5.66	0	-16.21	-21.87	25.31	11.37	38.04	32.28	48.63	43.07	41.71	33.68
	5	9.27	3.61	5.13	-0.53	24.91	13.06	36.84	33.54	47.49	45.00	40.60	34.87
	50	9.69	4.02	7.05	1.39	24.54	13.51	36.93	33.36	47.71	43.87	40.52	34.43
Fine-tune	0	5.87	0.21	-14.37	-20.03	25.51	11.75	38.58	32.03	50.23	42.53	43.41	33.68
	5	7.63	1.97	4.37	-1.29	25.04	12.59	37.46	31.16	48.05	41.59	41.00	32.84
	50	7.29	1.63	5.82	0.16	24.75	12.59	37.87	30.89	48.01	42.23	40.92	33.14
Test	50	7.29	1.63	7.29	1.63	25.11	11.90	38.03	31.73	48.61	42.52	41.46	33.33

fication performance. Specifically, the DAS-SA-ASR system achieves a 3% relative reduction in WER (from 42.39% to 40.96%) and SER (from 34.29% to 33.33%) with WPE in the 4-channel setting. This demonstrates that WPE aids in reverberation reduction for fixed beamformers. However, the improvement in recognition performance for neural beamformers (MVDR and FaSNet) is less pronounced, likely because these beamformers have already learned to reduce reverberation during their training process.

4.5.3. Joint optimization of FaSNet and SA-ASR

Since fine-tuning SA-ASR with a frozen MVDR and FaSNet does not significantly improve SA-ASR performance, we conduct the joint optimization of SA-ASR and FaSNet, since FaSNet has relatively better performance and fewer parameters than MVDR. More specifically, we pretrain FaSNet for 0, 5, 10, or 50 epochs and subsequently fine-tune it for 15 epochs jointly with SA-ASR by backpropagating the SA-ASR loss.

The results in Table 2 show that, without WPE, joint optimization of FaSNet and SA-ASR (40.52%) reduces the WER by 9% relative to the frozen FaSNet (44.57%) and to SA-ASR (44.54%). We also observed a 7% relatively lower SER (33.68%) than using the frozen FaSNet (36.24%). However, the fine-tuned FaSNet exhibits a smaller SI-SDRi than the pretrained one. This indicates that joint training optimizes FaSNet for ASR performance rather than maximum noise and reverberation reduction at the cost of greater speech distortion. Furthermore, while the number of FaSNet pretraining epochs significantly impacts the SI-SDRi, it does not significantly impact the result of joint optimization, provided it's nonzero. This demonstrates the insensitivity of the joint optimization to the pretraining level of FaSNet.

Moreover, we tested different uses of WPE during the joint optimization phase. First, incorporating WPE during the fine-tuning optimization of beamforming does not result in better recognition performance. On the contrary, using WPE led to a 4% relative increase in WER (from 41.71% to 43.41%) when training FaSNet from scratch with the training objectives of SA-ASR. Second, applying WPE to FaSNet after the joint optimization is completed also does not bring any

benefits (41.46%). This indicates that the joint optimization of FaSNet and SA-ASR does not benefit from the dereverberation provided by WPE.

5. CONCLUSION

This paper explored the integration of beamforming with SA-ASR for joint speech and speaker recognition of far-field meeting audio. We evaluated the impact of fine-tuning SA-ASR on the outputs of DAS, MVDR, or FaSNet beamformer and jointly fine-tuning SA-ASR with the last, and compared it with state-of-the-art MFCCA-based channel fusion. Experiments reveal that, in contrast to previously published results on simulated data, MFCCA is inefficient on real AMI data. This highlights the importance of systematically evaluating SA-ASR on real meeting data. Utilizing a DAS beamformer and jointly optimizing SA-ASR with the FaSNet beamformer lead to relative WER reductions of 8% and 9%, respectively. The use of WPE in the DAS-SA-ASR system can bring a 3% relative reduction in both WER and SER.

6. ACKNOWLEDGMENTS

Experiments presented in this paper were carried out using (a) the Grid'5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>), and (b) HPC resources from GENCI-IDRIS (Grant 2023-[AD011013881]).

7. REFERENCES

- [1] Shinji Watanabe, Michael Mandel, Jon Barker, Emmanuel Vincent, Ashish Arora, Xuankai Chang, Sanjeev Khudanpur, Vimal Manohar, Daniel Povey, et al., "CHiME-6 Challenge: Tackling multispeaker speech recognition for unsegmented recordings," in *CHiME*, 2020, pp. 1–7.
- [2] Fan Yu, Shiliang Zhang, Yihui Fu, Lei Xie, Siqi Zheng, Zhihao Du, Weilong Huang, Pengcheng Guo, Zhi-

- jie Yan, et al., “M2Met: The ICASSP 2022 Multi-Channel Multi-Party Meeting Transcription Challenge,” in *ICASSP*, 2022, pp. 6167–6171.
- [3] Samuele Cornell, Matthew S. Wiesner, Shinji Watanabe, Desh Raj, Xuankai Chang, Paola Garcia, Yoshiki Masuyam, Zhong-Qiu Wang, Stefano Squartini, and Sanjeev Khudanpur, “The CHiME-7 DASR Challenge: Distant meeting transcription with multiple devices in diverse scenarios,” in *CHiME*, 2023, pp. 1–6.
- [4] Takuya Yoshioka, Hakan Erdogan, Zhuo Chen, Xiong Xiao, and Fil Allewa, “Recognizing overlapped speech in meetings: A multichannel separation approach using neural networks,” in *Interspeech*, 2018, pp. 3038–3042.
- [5] Hangting Chen, Pengyuan Zhang, and Yonghong Yan, “Multi-talker MVDR beamforming based on extended complex Gaussian mixture model,” *arXiv preprint arXiv:1910.07753*, 2019.
- [6] Naoyuki Kanda, Christoph Boeddeker, Jens Heitkaemper, Yusuke Fujita, Shota Horiguchi, Kenji Nagamatsu, and Reinhold Haeb-Umbach, “Guided source separation meets a strong ASR backend: Hitachi/Paderborn University joint investigation for dinner party ASR,” in *Interspeech*, 2019, pp. 1248–1252.
- [7] Xuankai Chang, Wangyou Zhang, Yanmin Qian, Jonathan Le Roux, and Shinji Watanabe, “MIMO-Speech: End-to-end multi-channel multi-speaker speech recognition,” in *ASRU*, 2019, pp. 237–244.
- [8] Jian Wu, Zhuo Chen, Sanyuan Chen, Yu Wu, Takuya Yoshioka, Naoyuki Kanda, Shujie Liu, and Jinyu Li, “Investigation of practical aspects of single channel speech separation for ASR,” in *Interspeech*, 2021, pp. 3066–3070.
- [9] Jing Shi, Xuankai Chang, Shinji Watanabe, and Bo Xu, “Train from scratch: Single-stage joint training of speech separation and recognition,” *Computer Speech & Language*, vol. 76, pp. 101387, 2022.
- [10] Naoyuki Kanda, Yashesh Gaur, Xiaofei Wang, Zhong Meng, and Takuya Yoshioka, “Serialized output training for end-to-end overlapped speech recognition,” in *Interspeech*, 2020, pp. 2797–2801.
- [11] Naoyuki Kanda, Guoli Ye, Yashesh Gaur, Xiaofei Wang, Zhong Meng, Zhuo Chen, and Takuya Yoshioka, “End-to-end speaker-attributed ASR with Transformer,” in *Interspeech*, 2021, pp. 4413–4417.
- [12] Fan Yu, Shiliang Zhang, Pengcheng Guo, Yuhao Liang, Zhihao Du, Yuxiao Lin, and Lei Xie, “MFCCA: Multi-frame cross-channel attention for multi-speaker ASR in multi-party meeting scenario,” in *SLT*, 2023, pp. 144–151.
- [13] Can Cui, Imran Sheikh, Mostafa Sadeghi, and Emmanuel Vincent, “End-to-end multichannel speaker-attributed ASR: Speaker guided decoder and input feature analysis,” in *ASRU*, 2023.
- [14] Xavier Anguera, Chuck Wooters, and Javier Hernando, “Acoustic beamforming for speaker diarization of meetings,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 7, pp. 2011–2022, 2007.
- [15] Tsubasa Ochiai, Shinji Watanabe, Takaaki Hori, John R Hershey, and Xiong Xiao, “Unified architecture for multichannel end-to-end speech recognition with neural beamforming,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1274–1288, 2017.
- [16] Hung-Shin Lee, Yu-Huai Peng, Pin-Tuan Huang, Ying-Chun Tseng, Chia-Hua Wu, Yu Tsao, and Hsin-Min Wang, “The Academia Sinica systems of speech recognition and speaker diarization for the CHiME-6 Challenge,” in *CHiME*, 2020, pp. 29–32.
- [17] Naoyuki Kanda, Jian Wu, Xiaofei Wang, Zhuo Chen, Jinyu Li, and Takuya Yoshioka, “Vararray meets T-Sot: Advancing the state of the art of streaming distant conversational speech recognition,” in *ICASSP*, 2023, pp. 1–5.
- [18] Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Mael Guillemot, Thomas Hain, Jaroslav Kadlec, Vasilis Karaiskos, Wessel Kraaij, et al., “The AMI meeting corpus: A pre-announcement,” in *International Workshop on Machine Learning for Multimodal Interaction*, 2005, pp. 28–39.
- [19] Don H. Johnson and Dan E. Dudgeon, *Array signal processing: concepts and techniques*, Prentice Hall, 1993.
- [20] Xiong Xiao, Chenglin Xu, Zhaofeng Zhang, Shengkui Zhao, Sining Sun, Shinji Watanabe, Longbiao Wang, Lei Xie, Douglas L Jones, et al., “A study of learning based beamforming methods for speech recognition,” in *CHiME*, 2016, pp. 26–31.
- [21] C. Knapp and G. Carter, “The generalized correlation method for estimation of time delay,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 24, no. 4, pp. 320–327, 1976.
- [22] Yen-Ju Lu, Samuele Cornell, Xuankai Chang, Wangyou Zhang, Chenda Li, Zhaoheng Ni, Zhong-Qiu Wang, and Shinji Watanabe, “Towards low-distortion multi-channel speech enhancement: The espnet-se submission to the I3das22 challenge,” in *ICASSP*, 2022, pp. 9201–9205.

- [23] Minseung Kim, Sein Cheong, and Jong Won Shin, “DNN-based Parameter Estimation for MVDR Beamforming and Post-filtering,” in *Interspeech*, 2023, pp. 3879–3883.
- [24] Yi Luo, Cong Han, Nima Mesgarani, Enea Ceolini, and Shih-Chii Liu, “FaSNet: Low-latency adaptive beamforming for multi-microphone audio processing,” in *ASRU*, 2019, pp. 260–267.
- [25] Yi Luo, Zhuo Chen, and Takuya Yoshioka, “Dual-path RNN: efficient long sequence modeling for time-domain single-channel speech separation,” in *ICASSP*, 2020, pp. 46–50.
- [26] Yi Luo, Zhuo Chen, Nima Mesgarani, and Takuya Yoshioka, “End-to-end microphone permutation and number invariant multi-channel speech separation,” in *ICASSP*, 2020, pp. 6394–6398.
- [27] Tomohiro Nakatani, Takuya Yoshioka, Keisuke Kinoshita, Masato Miyoshi, and Biing-Hwang Juang, “Speech dereverberation based on variance-normalized delayed linear prediction,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 7, pp. 1717–1731, 2010.
- [28] Joon-Young Yang and Joon-Hyuk Chang, “Virtual acoustic channel expansion based on neural networks for weighted prediction error-based speech dereverberation,” in *Interspeech*, 2020, pp. 3930–3934.
- [29] Sandipana Dowerah, Romain Serizel, Denis Jouvét, Mohammad Mohammadamini, and Driss Matrouf, “How to leverage DNN-based speech enhancement for multi-channel speaker verification?,” in *International Conference on Advances in Signal Processing and Artificial Intelligence (ASPAI' 2022)*, 2022.
- [30] Emmanuel Vincent, Rémi Gribonval, and Cédric Févotte, “Performance measurement in blind audio source separation,” *IEEE transactions on audio, speech, and language processing*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [31] Jonathan Le Roux, Scott Wisdom, Hakan Erdogan, and John R Hershey, “Sdr-half-baked or well done?,” in *ICASSP*, 2019, pp. 626–630.
- [32] Muqiao Yang, Naoyuki Kanda, Xiaofei Wang, Jian Wu, Sunit Sivasankaran, Zhuo Chen, Jinyu Li, and Takuya Yoshioka, “Simulating realistic speech overlaps improves multi-talker ASR,” in *ICASSP*, 2023, pp. 1–5.
- [33] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *ICASSP*, 2015, pp. 5206–5210.
- [34] David Diaz-Guerra, Antonio Miguel, and Jose R. Beltran, “gpuRIR: A Python library for room impulse response simulation with GPU acceleration,” *Multimedia Tools and Applications*, vol. 80, pp. 5653–5671, 2021.
- [35] Manuel Pariente, Samuele Cornell, Joris Cosentino, Sunit Sivasankaran, Efthymios Tzinis, Jens Heitkaemper, Michel Olvera, Fabian-Robert Stöter, Mathieu Hu, et al., “Asteroid: the PyTorch-based audio source separation toolkit for researchers,” in *Interspeech*, 2020.
- [36] Naoyuki Kanda, Yashesh Gaur, Xiaofei Wang, Zhong Meng, Zhuo Chen, Tianyan Zhou, and Takuya Yoshioka, “Joint speaker counting, speech recognition, and speaker identification for overlapped speech of any number of speakers,” in *Interspeech*, 2020, pp. 36–40.
- [37] Mirco Ravanelli, Titouan Parcollet, Peter Plantinga, Aku Rouhe, Samuele Cornell, Loren Lugosch, Cem Subakan, Nauman Dawalatabad, Abdelwahab Heba, et al., “SpeechBrain: A general-purpose speech toolkit,” 2021, arXiv:2106.04624.
- [38] Yao-Yuan Yang, Moto Hira, Zhaoheng Ni, Anjali Chourdia, Artyom Astafurov, Caroline Chen, Ching-Feng Yeh, Christian Puhersch, David Pollack, Dmitriy Genzel, Donny Greenberg, Edward Z. Yang, Jason Lian, Jay Mahadeokar, Jeff Hwang, Ji Chen, Peter Goldsborough, Prabhat Roy, Sean Narenthiran, Shinji Watanabe, Soumith Chintala, Vincent Quenneville-Bélaïr, and Yangyang Shi, “Torchaudio: Building blocks for audio and speech processing,” *arXiv preprint arXiv:2110.15018*, 2021.
- [39] Yi Luo and Nima Mesgarani, “Conv-tasnet: Surpassing ideal time–frequency magnitude masking for speech separation,” *IEEE/ACM transactions on audio, speech, and language processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [40] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuynck, “ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in TDNN based speaker verification,” in *Interspeech*, 2020, pp. 3830–3834.
- [41] Lukas Drude, Jahn Heymann, Christoph Boeddeker, and Reinhold Haeb-Umbach, “NARA-WPE: A python package for weighted prediction error dereverberation in Numpy and Tensorflow for online and offline processing,” in *ITG Fachtagung Sprachkommunikation*, 2018.
- [42] NIST, “SCTK,” <https://github.com/usnistgov/SCTK.git>, 2024, GitHub repository.