



Empowering Whisper as a Joint Multi-Talker and Target-Talker Speech Recognition System

Lingwei Meng, Jiawen Kang, Yuejiao Wang, Zengrui Jin, Xixin Wu, Xunying Liu, Helen Meng

The Chinese University of Hong Kong, Hong Kong SAR, China

{lmeng, jwkang, wangy, zrjin, wuxx, xyliu, hmmeng}@se.cuhk.edu.hk

Abstract

Multi-talker speech recognition and target-talker speech recognition, both involve transcription in multi-talker contexts, remain significant challenges. However, existing methods rarely attempt to simultaneously address both tasks. In this study, we propose a pioneering approach to empower Whisper, which is a speech foundation model, to tackle joint multi-talker and target-talker speech recognition tasks. Specifically, (i) we freeze Whisper and plug a Sidecar separator into its encoder to separate mixed embedding for multiple talkers; (ii) a Target Talker Identifier is introduced to identify the embedding flow of the target talker on the fly, requiring only three-second enrollment speech as a cue; (iii) soft prompt tuning for decoder is explored for better task adaptation. Our method outperforms previous methods on two- and three-talker LibriMix and LibriSpeechMix datasets for both tasks, and delivers acceptable zero-shot performance on multi-talker ASR on AishellMix Mandarin dataset.¹

Index Terms: multi-talker speech recognition, target-talker speech recognition, prompt tuning, domain adaptation

1. Introduction

Driven by the rapid development of deep learning along with the availability of large-scale data, automatic speech recognition (ASR) has achieved significant progress in recent years [1]. However, speech recognition in the multi-talker scenarios, where overlapping may exist, remains challenging and has attracted much attention.

To tackle the multi-talker speech recognition problem, various approaches have been explored. Conventional cascaded systems employ a speech separation module as a front-end to separate mixed speech signals, which are then fed into a single-talker ASR system for transcription [2, 3]. However, these systems usually show limited performance due to the mismatch of their optimization objectives, and may need further joint training [2]. Recently, end-to-end models have garnered interest owing to their outstanding performance. One primary challenge in training an end-to-end multi-talker ASR system is to associate prediction with the corresponding target labels for loss calculation [4]. Consequently, techniques such as Permutation Invariant Training (PIT) [5, 6, 7, 8, 9], Heuristic Error Assignment Training (HEAT) [10, 11, 12], and Serialized Output Training (SOT) [13, 14, 15] have emerged. Although these methods have yielded impressive results, they often necessitate training from scratch or performing full fine-tuning on pre-trained models, which does not fully capitalize on the existing advancements developed for standard single-talker ASR. Enlightened by findings that the ASR encoder captures more acoustic information

in its lower layers and more linguistic information in the upper layers [16, 17, 18], a recent study advocates for the use of a Conv-TasNet-like [19] Sidecar separator to tackle multi-talker speech recognition, without distorting the parameters of a well-trained single-talker ASR model [20, 21].

Target-talker ASR, which aims to efficiently recognize speech of a target talker under a multi-talker scenario, also holds substantial practical value. End-to-end approaches have been investigated and achieved substantial progress [22, 23]. However, these methods typically necessitate an external [9, 24, 25] or internal [26] module to derive the speaker embedding from the enrollment speech of target-talker, consequently increasing the model's computational burden. Moreover, they typically only output the transcripts of an assigned target talker, neglecting the speech of other talkers. This limitation hinders their applicability in situations where users may also be interested in obtaining the transcriptions of non-target talkers. Although speaker-attributed ASR can transcribe multiple speakers in a speaker-aware manner, it typically necessitates the speaker embeddings of all involved individuals [27]. As far as we know, [25] is the only study attempting to address joint multi-talker and target-talker ASR; however, it still requires an external speaker embedding extractor.

Nowadays, speech foundation models have emerged as a versatile solution for diverse speech tasks [28, 29, 18]. As an representative in this domain, Whisper [1] has demonstrated its potential across various tasks beyond ASR [30, 31] which motivated us to further extend Whisper's capabilities in tackling multi-talker and target-talker speech recognition challenges.

In this study, we empower Whisper as a joint multi-talker and target-talker system in a parameter-efficient style. Specifically, we freeze the weights of Whisper and incorporate a Sidecar separator into its encoder to endow it with multi-talker speech recognition capabilities. A Target Talker Identifier (TTI) module is introduced to distinguish the target speaker's embedding branch on the fly, requiring only three seconds of the target talker's enrollment speech as a cue. Moreover, soft prompt tuning [32] for Whisper decoder is adopted to further adapt to the tasks. Our major contributions are threefold:

- We propose a pioneering framework to jointly transcribing multi-talker speech while highlighting the target talker's speech, without employing any speaker embedding extractor.
- Leveraging the frozen Whisper as the foundation model, our framework only involves limited trainable parameters, making it a parameter-efficient and loosely-coupled system.
- Extensive experiments reveal that the proposed approach achieves leading performance on two- and three-talker LibriMix and LibriSpeechMix datasets (English) on both tasks, and attains satisfactory zero-shot multi-talker ASR performance on AishellMix (Mandarin).

¹The code is available at <https://github.com/LingweiMeng/Whisper-Sidecar>

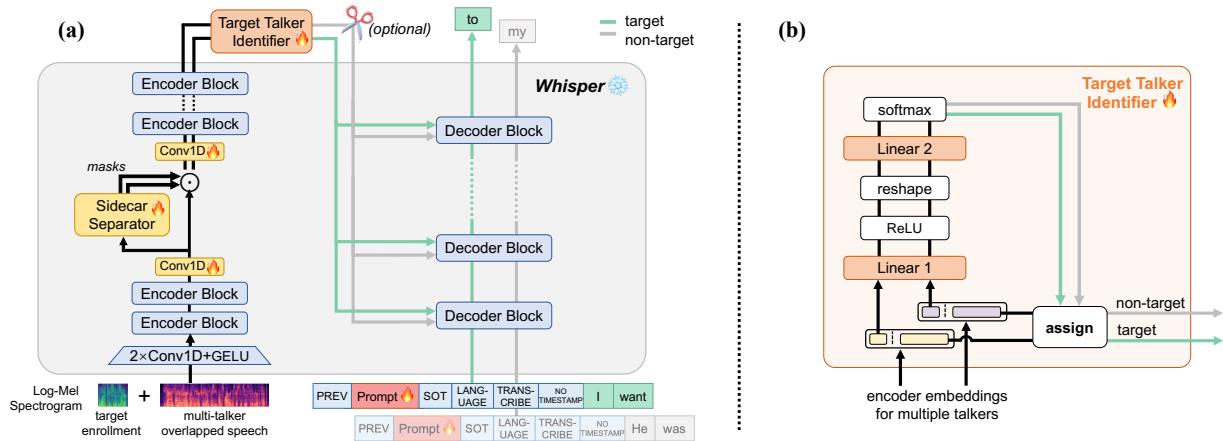


Figure 1: Take two-talker scenario as an example, the proposed system (a) take the concatenation of target enrollment speech and multi-talker speech as input. The embedding is separated by Sidecar separator. Target Talker Identifier (b) processes the prefix segments of encoder embeddings to identify the target talker branch. Optionally, non-target branches can be discarded to accelerate inference.

2. Methods

The proposed method consists of four main components — Whisper serving as the foundation model, a Sidecar separator to separate mixed embedding for multiple talkers, a Target Talker Identifier to identify the embedding flow of the target talker, and a soft prompt embedding to facilitate task adaptation.

2.1. Whisper as the Speech Foundation Model

Whisper is a speech recognition model featuring an attention-based encoder-decoder structure, which has been trained on massive amounts of web-scale labeled speech data [1]. Nowadays, it is increasingly being utilized as a speech foundation model even beyond speech recognition [24, 30, 31]. In this study, we are inspired to extend its capability to handling joint multi-talker and target-talker ASR tasks.

Whisper takes log-Mel spectrogram as input, followed by transformer encoder and decoder modules to decode the output tokens in an auto-regressive manner. Different from other ASR models, Whisper adopts several special tokens as the prefix of input sequences for decoder to specify tasks and condition information. By default, the prefix tokens are " $\langle | \text{PREV} | \rangle$ ", *text prompt*, " $\langle | \text{SOT} | \rangle$ ", " $\langle | \text{LANGUAGE} | \rangle$ ", " $\langle | \text{TRANSCRIBE} | \rangle$ ", " $\langle | \text{NO_TIMESTAMP} | \rangle$ ", where " $\langle | \text{PREV} | \rangle$ " and *text prompt* are optional.

2.2. Empowering Whisper as a Multi-Talker ASR System

Recently, the Sidecar separator (SS) has been introduced as a parameter-efficient module to convert a well-trained single-talker ASR model into a multi-talker one [20]. In this work, we incorporated the Sidecar separator with Whisper to harness its capability acquired from extensive training data.

The Sidecar separator is a temporal convolutional network inserted between the early layers of the ASR encoders. It consists of stacked 1-D dilated convolutional blocks inspired by Conv-TasNet [19]. As the shallower layers of the ASR encoder are believed to encode more acoustic information rather than the linguistic ones [17, 20], the Sidecar separator is able to separate the mixed representation with talker-related masks, producing disentangled representation of different speakers.

As depicted in Figure 1, the Sidecar separator accompanied by two 1-D convolutional layers is positioned after the second encoder block. Talker-dependent masks are generated,

which are element-wisely multiplied with the mixed embedding, yielding separated embeddings of each talker. The subsequent encoder blocks and decoder process these branches, ultimately transcribing the corresponding text for each talker.

2.3. Target Talker Identifier

We introduce the Target Talker Identifier (TTI) module, which equips the system with the capability for target-talker ASR.

During the forward process, as illustrated in Figure 1 (b), the encoder-output embeddings corresponding to different talkers will be segmented into two distinct segments: a prefix segment aligning with the length of three-second enrollment speech, and a main segment that corresponds to the duration of the multi-talker speech. The prefix segments are then fed into the TTI module, which determines the branch associated with the target talker, while only the main embedding segments are sent to the Whisper decoder for transcription.

Specifically, the prefix segments hold a tensor shape of $(B \times S, 150, C)$, where B denotes batch size, S denotes the number of talkers, C denotes the number of channels, and 150 denotes the number of time frames. Given that each time frame spans a duration of 20 ms, 150 coincides with the three-second duration of the enrollment speech. As shown in Figure 1, within the TTI module, the prefix segments traverse a linear layer followed by the ReLU activation function, yielding a tensor with a shape of $(B \times S, 150, 1)$. Upon squeezing and reshaping, the tensor proceeds through another linear layer and the softmax function to produce the probability (B, S) of each branch being the target talker.

Consequently, the target talker branch is efficiently determined on the fly, introducing minimal computational overhead. Underpinned by the superior performance of the separation module, TTI can be considered as performing target-talker activity detection, which is a more economical task compared with methods necessitating speaker embedding extraction.

2.4. Soft Prompt Tuning

The original Whisper model allows for the inclusion of text prompt tokens as prefix to the decoder's input sequences, conditioned on which the model yields improved ASR performance on ambiguous audios [1]. In this context, by exploiting this inherent characteristic of Whisper combined with soft prompt

tuning technique [32], we aim to adapt the model more efficiently to multi-talker and target-talker ASR tasks.

Specifically, as shown in Figure 1 (a), we insert a learnable embedding as soft prompt between $\langle | \text{PREV} | \rangle$ and $\langle | \text{SOT} | \rangle$ tokens where hard prompt tokens were originally specified. Note that we mask the position of the soft prompt when calculating the training loss, since the model does not require learning to generate them. The soft prompt embedding will be updated as the model learns to transcribe the multi-talker speech.

2.5. Training objectives

At each training step, there’s an 80% probability of undertaking multi-talker ASR training, while a 20% probability for appending a three-second enrollment speech for joint multi-talker and target-talker ASR training. Both ASR loss and TTI’s cross-entropy loss necessitate a permutation assignment for speaker order to address the label ambiguity issue [33]. In this study, the permutation is determined by Permutation Invariant Training (PIT) based on ASR loss, and is then assigned for TTI’s cross-entropy loss calculation. The permutation is derived as:

$$\hat{\pi} = \arg \min_{\pi \in \mathcal{P}} \sum_{s=1}^S \text{Loss}_{\text{ASR}}(Y^s, R^{\pi(s)}) \quad (1)$$

where \mathcal{P} denotes the set of all permutations on $= \{1, \dots, S\}$, $\pi(s)$ denotes the s -th element in a permutation π , Y^s is the predicted token sequences of the s -th branch, and R is the reference labels for S talkers. At last, the final objective function is the sum of PIT-ASR loss and corresponding TTI loss multiplied by a coefficient λ . Therefore we have,

$$\mathcal{L} = \mathcal{L}_{\text{ASR}} + \lambda \mathcal{L}_{\text{TTI}} \quad (2)$$

$$\mathcal{L}_{\text{ASR}} = \sum_s \text{Loss}_{\text{ASR}}(Y^s, R^{\hat{\pi}(s)}) \quad (3)$$

$$\mathcal{L}_{\text{TTI}} = \text{Loss}_{\text{CE}}(Z, D^{\hat{\pi}}) \quad (4)$$

where Z is the probability of each branch is of the target talker, and $D^{\hat{\pi}}$ is the ground truth after permutation.

Considering that the original Whisper was not trained using CTC loss, we refrain from employing an additional CTC loss for early permutation assignment as done in [5, 6].

3. Experimental Setup

3.1. Datasets

The experiments are conducted on three multi-talker public datasets, namely LibriMix [34] and LibriSpeechMix [13] in English, and Aishell1Mix² in Mandarin. Audio exceeding Whisper’s maximum handling duration of 30 seconds are time-stretched to conform to this limit. For target-talker ASR on LibriMix and LibriSpeechMix, we randomly trim three-second clips from LibriSpeech as enrollment speech for each talker.

LibriMix. The dataset simulates audio mixtures in a left-aligned manner, involving two or three speakers from the LibriSpeech-clean corpus. Thus, the shorter source speech is entirely overlapped with the longer one from the start, presenting significant challenge in separating overlaps. We focus on its two-speaker-mixed and three-speaker-mixed clean subset, denoted as Libri2Mix and Libri3Mix in the following.

LibriSpeechMix. The utterances are simulated from LibriSpeech, comprising mixtures from two or three speakers. Unlike LibriMix, the delay time for each speaker is randomly sampled, resulting in partially overlapped mixtures. Since only official dev and test sets are released, we created our training set

²<https://github.com/huangzj421/Aishell1Mix>

from the 960-hour LibriSpeech following the same protocol as in [13], except that the mixtures are kept under 30 seconds.

Aishell1Mix. It is a Mandarin multi-talker speech dataset, source from Aishell1 corpus. It simulate mixtures with a same protocol of the LibriMix. We focus on its two-speaker-clean subset for analysis, denoted as AishellMix in the following.

3.2. Model Settings and Evaluation Metrics

Throughout this study, we employ Whisper-small, -medium, and -large-v3 as the foundation models, respectively. We freeze these models and only train the Sidecar separator, Target Talker Identifier, and soft prompt embedding. The number of trainable parameters for systems using different foundation models and for various numbers of talkers are listed in Table 1.

The Conv-TasNet-like Sidecar separator [20] comprises a series of K temporal convolutional blocks with dilation rates ranging from 1 to 2^{K-1} , with each block repeats up to R times. Consistent with the protocol in [20, 21], we use $K = 8$ and $R = 3$ and plug it between the second and third encoder blocks. The length for soft prompt embeddings are investigated through ablation experiments (Section 4.4). As a result, we establish a length of 4, which gives the best performance.

For systems with the TTI module, at each training step, there’s an 80% probability of undertaking multi-talker ASR training, while a 20% probability for joint multi-talker and target-talker ASR training. We set the coefficient of TTI loss λ to 0.01. The systems are trained and evaluated on two- and three-talker subsets of LibriMix and LibriSpeechMix, respectively. Each training session lasts for a maximum of 200k steps on 8 NVIDIA V100 GPUs with a total batch size of 16, employing AdamW optimizer with an initial learning rate of 2e-4 that decreases linearly to 1e-4.

Permutations with minimum errors are used to compute word error rate (WER) or character error rate (CER) for multi-talker ASR as in prior studies [14, 20]. For target-talker ASR, we use standard WER for evaluation. Both the model’s predictions and the references are normalized following [1].

Table 1: *The amount of trainable parameters, with numbers in parentheses indicating their proportion in the total parameters.*

Foundation Model	2-speaker	3-speaker
Whisper-small	8.69 M (3.47%)	8.79 M (3.51%)
Whisper-medium	13.16 M (1.69%)	13.29 M (1.71%)
Whisper-large	18.41 M (1.18%)	18.58 M (1.19%)

4. Results and Discussions

4.1. Multi-Talker ASR Results

We compared the performance of various systems for multi-talker ASR on the two- and three-talker LibriMix and LibriSpeechMix test sets, as shown in Table 2. Empowered by Sidecar separator (SS), TTI, and soft prompt, our systems (f)-(k) consistently illustrated improved performance across both the two- and three-speaker subsets. Even with Whisper-small-SS-TTI (g), owing to Whisper’s extensive pre-training, our method has already surpassed the original Sidecar scheme[21]. As the size of the Whisper model increases, we observed a steady improvement in performance, which aligns with our expectations. Ultimately, our systems outperformed previous approaches across all datasets except for LibriSpeechMix-3spk, demonstrating the superiority of the proposed approach.

Interestingly, we find that systems with the TTI module (g) (i) (k) outperform their counterparts without TTI (f) (h) (j) in

multi-talker ASR task, even though the TTI is specifically designed to support target-talker ASR. This suggests that the training objective of learning to distinguish the target talker can benefit the Sidecar’s ability to separate embeddings, thereby facilitating the task of multi-talker ASR.

4.2. Target-Talker ASR Results

We evaluate target-talker ASR performance on two- and three-speaker subsets of LibriMix and LibriSpeechMix, as illustrated in Table 3.³ Our systems outperform previous state-of-the-art method on LibriMix dataset by a large margin, and we are the first to perform target-talker ASR task on Libri3Mix. To guarantee a fair comparison, we further trained three additional systems with limited training data to ensure consistency with the data used in [9, 24]. These systems are denoted as ”-limited”. The results demonstrate that our method still outperforms [9, 24] though under this restriction.

For LibriSpeechMix, the speech signals are partially overlapped, which means the target talker’s speech can incur a considerable delay before it commences, resulting in a substantial time interval away between it and the enrollment speech. Nevertheless, despite the existence of delays, our systems still demonstrate good performance on the LibriSpeechMix dataset, validating the effectiveness of the proposed method.

4.3. Zero-Shot Multi-Lingual Evaluation

We investigated whether the multi-lingual characteristics of Whisper are retained after fine-tuning it on an English multi-talker dataset. Specifically, we conducted evaluations for systems (g) (i) (k) listed in Table 1 using the two-speaker AishellMix Mandarin dataset, which is the first time to be used on multi-talker ASR task. The evaluations are performed using two schemes: zero-shot and one-batch-tuning. Zero-shot refers to directly evaluating the system on the AishellMix, while one-batch-tuning implies conducting an additional training epoch on the AishellMix training set prior to the evaluation.

As Table 5 illustrates, the medium and large models demonstrated acceptable CER performance even under zero-shot conditions. With just one batch tuning, these models exhibited satisfactory results. This suggests that our method largely maintains the inherent multilingual capabilities of Whisper.

4.4. Ablation Study

We investigated the optimal prompt length by examining the multi-talker ASR performance on the Libri2Mix dataset with Whisper-medium and Whisper-large models. As shown in Table 5, a soft prompt of length 4 yields the best performance. However, as the soft prompt length increases to 16, the systems see a decline in performance. This may be due to overly long sequence sequences, making the model difficult to optimize, given the original Whisper model is frozen.

4.5. Limitations and Future Work

This study has several limitations. Firstly, our method relies on PIT which requires pre-defining of the maximum number of speakers. Future efforts will integrate SOT [13] or HEAT [11] to address this issue and reduce training costs. Secondly, when the target talker’s speech undergoes excessive delay, there could be potential degradation in the target-talker ASR’s performance. We anticipate future work to enable the TTI module synthesize information across the entire utterance duration rather than only the three-second enrollment speech.

³We did not include results reported in [26], which delivers better performance but undergoes about ten times training efforts as ours.

Table 2: Multi-talker ASR on the test sets of LibriMix and LibriSpeechMix. Evaluated by WER (%). “SS” denotes “Sidecar Separator”, “TTI” denotes “Target Talker Identifier”.

System	LibriMix		LibriSpeechMix	
	2spk	3spk	2spk	3spk
(a) WavLM Base+ PIT [9]	18.45	-	-	-
(b) C-HuBERT-Large [35]	7.80	-	-	-
(c) SURT [11]	-	-	7.20	-
(d) SOT-Conformer [27]	-	-	4.90 [†]	6.20[†]
(e) D2V-Sidecar-DB [21]	9.69	33.91	7.49	11.94
(f) Whisper-small-SS	10.04	29.20	5.27	9.85
(g) Whisper-small-SS-TTI	9.39	26.76	5.18	8.61
(h) Whisper-medium-SS	6.95	22.58	4.32	7.80
(i) Whisper-medium-SS-TTI	6.56	21.47	4.01	7.50
(j) Whisper-large-SS	4.98	17.55	3.81	7.13
(k) Whisper-large-SS-TTI	4.66	16.79	3.43	6.80

[†] with extremely heavier training efforts.

Table 3: Target-talker ASR on LibriMix and LibriSpeechMix. Evaluated by WER (%). ”-limited” denotes using the same training data as in [24].

System	LibriMix		LibriSpeechMix	
	2spk	3spk	2spk	3spk
WavLM-Base ⁺ -TSE [9]	12.32	-	-	-
Whisper-TS-ASR [24]	11.98	-	-	-
Whisper-small-SS-TTI-limited	15.75	-	-	-
Whisper-medium-SS-TTI-limited	11.39	-	-	-
Whisper-large-SS-TTI-limited	10.79	-	-	-
Whisper-small-SS-TTI	11.81	30.52	8.89	15.85
Whisper-medium-SS-TTI	9.14	25.75	7.58	12.4
Whisper-large-SS-TTI	7.97	21.97	6.99	11.4

Table 4: Zero-shot and one-batch-tuning multi-talker ASR on Aishell1Mix Mandarin dataset. Evaluated by CER (%).

System	zero-shot	one-batch-tuning
Whisper-small-SS-TTI	55.87	28.95
Whisper-medium-SS-TTI	36.28	19.83
Whisper-large-SS-TTI	28.94	17.81

Table 5: Ablation study on soft prompt, evaluated by WER (%).

System	Soft Prompt Length				
	0	2	4	8	16
Whisper-medium-SS-TTI	7.21	6.82	6.56	6.84	7.5
Whisper-large-SS-TTI	5.27	4.98	4.66	4.74	5.43

5. Conclusions

In this study, we introduce a novel methodology that harnesses Whisper, a speech foundation model, to jointly transcribe multi-talker speech meanwhile highlighting the target talker’s speech, without employing any speaker embedding extractor. Specifically, we freeze whisper and insert a Sidecar separator into its encoder to separate mixed embedding for multiple talkers. Subsequently, a Target Talker Identifier module is introduced to identify the embedding flow of the target talker on the fly, requiring only three-second enrollment speech as a cue. The soft prompt tuning is further utilized to facilitate task adaptation.

Extensive experiments reveal that our approach outperforms previous methods on LibriMix and LibriSpeechMix on both tasks. Moreover, it achieves acceptable zero-shot performance on multi-talker ASR on AishellMix Mandarin dataset.

6. Acknowledgements

This research is partially supported by the HKSARG Research Grants Council's Theme-based Research Grant Scheme (Project No. T45-407/19N) and by the CUHK Stanley Ho Big Data Decision Analytics Research Centre.

7. References

- [1] A. Radford, J. W. Kim *et al.*, "Robust speech recognition via large-scale weak supervision," in *Proceedings of the 40th International Conference on Machine Learning*, vol. 202, 2023, pp. 28 492–28 518.
- [2] S. Settle, J. Le Roux *et al.*, "End-to-end multi-speaker speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [3] S. Li, B. Ouyang *et al.*, "Real-time end-to-end monaural multi-speaker speech recognition," in *Proc. Interspeech 2021*, 2021, pp. 3750–3754.
- [4] J. Kang, L. Meng *et al.*, "Cross-speaker encoding network for multi-talker speech recognition," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 1–5.
- [5] H. Seki, T. Hori *et al.*, "A purely end-to-end system for multi-speaker speech recognition," in *ACL*, 2018.
- [6] W. Zhang, X. Chang *et al.*, "Improving end-to-end single-channel multi-talker speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1385–1394, 2020.
- [7] X. Chang, Y. Qian *et al.*, "End-to-end monaural multi-speaker ASR system without pretraining," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [8] X. Chang, W. Zhang *et al.*, "End-to-end multi-speaker speech recognition with transformer," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [9] Z. Huang, D. Raj *et al.*, "Adapting self-supervised models to multi-talker speech recognition using speaker embeddings," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [10] A. Tripathi, H. Lu *et al.*, "End-to-end multi-talker overlapping speech recognition," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6129–6133.
- [11] L. Lu, N. Kanda *et al.*, "Streaming end-to-end multi-talker speech recognition," *IEEE Signal Processing Letters*, vol. 28, pp. 803–807, 2021.
- [12] D. Raj, D. Povey *et al.*, "Surt 2.0: Advances in transducer-based multi-talker speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 3800–3813, 2023.
- [13] N. Kanda, Y. Gaur *et al.*, "Serialized output training for end-to-end overlapped speech recognition," in *Proc. Interspeech 2020*, 2020, pp. 2797–2801.
- [14] N. Kanda, J. Wu *et al.*, "Streaming Multi-Talker ASR with Token-Level Serialized Output Training," in *Proc. Interspeech 2022*, 2022, pp. 3774–3778.
- [15] C. Li, Y. Qian *et al.*, "Adapting Multi-Lingual ASR Models for Handling Multiple Talkers," in *Proc. INTERSPEECH 2023*, 2023, pp. 1314–1318.
- [16] A. Pasad, J.-C. Chou *et al.*, "Layer-wise analysis of a self-supervised speech representation model," in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2021, pp. 914–921.
- [17] K. Shim, J. Choi *et al.*, "Understanding the role of self attention for efficient speech recognition," in *International Conference on Learning Representations*, 2022.
- [18] S. Chen, C. Wang *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [19] Y. Luo and N. 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.
- [20] L. Meng, J. Kang *et al.*, "A sidecar separator can convert a single-talker speech recognition system to a multi-talker one," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [21] L. Meng, J. Kang *et al.*, "Unified Modeling of Multi-Talker Overlapped Speech Recognition and Diarization with a Sidecar Separator," in *Proc. INTERSPEECH 2023*, 2023, pp. 3467–3471.
- [22] W. Zhang and Y. Qian, "Weakly-Supervised Speech Pre-training: A Case Study on Target Speech Recognition," in *Proc. INTERSPEECH 2023*, 2023, pp. 3517–3521.
- [23] T. Moriya, H. Sato *et al.*, "Knowledge Distillation for Neural Transducer-based Target-Speaker ASR: Exploiting Parallel Mixture/Single-Talker Speech Data," in *Proc. INTERSPEECH 2023*, 2023, pp. 899–903.
- [24] H. Ma, Z. Peng *et al.*, "Extending whisper with prompt tuning to target-speaker ASR," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 1–5.
- [25] R. Masumura, N. Makishima *et al.*, "End-to-End Joint Target and Non-Target Speakers ASR," in *Proc. INTERSPEECH 2023*, 2023, pp. 2903–2907.
- [26] Y. Zhang, K. C. Puvvada *et al.*, "Conformer-based target-speaker automatic speech recognition for single-channel audio," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [27] N. Kanda, G. Ye *et al.*, "End-to-End Speaker-Attributed ASR with Transformer," in *Proc. Interspeech 2021*, 2021, pp. 4413–4417.
- [28] A. Baevski, Y. Zhou *et al.*, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Advances in neural information processing systems*, vol. 33, 2020, pp. 12 449–12 460.
- [29] W.-N. Hsu, B. Bolte *et al.*, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [30] Y. Gong, S. Khurana *et al.*, "Whisper-AT: Noise-Robust Automatic Speech Recognizers are Also Strong General Audio Event Taggers," in *Proc. INTERSPEECH 2023*, 2023, pp. 2798–2802.
- [31] P. Peng, B. Yan *et al.*, "Prompting the Hidden Talent of Web-Scale Speech Models for Zero-Shot Task Generalization," in *Proc. INTERSPEECH 2023*, 2023, pp. 396–400.
- [32] B. Lester, R. Al-Rfou *et al.*, "The power of scale for parameter-efficient prompt tuning," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, Nov. 2021, pp. 3045–3059.
- [33] D. Yu, M. Kolbæk *et al.*, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 241–245.
- [34] J. Cosentino, M. Pariente *et al.*, "LibriMix: An open-source dataset for generalizable speech separation," *arXiv preprint arXiv:2005.11262*, 2020.
- [35] M. Fazel-Zarandi and W.-N. Hsu, "Cocktail hubert: Generalized self-supervised pre-training for mixture and single-source speech," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.