

Target Speaker ASR with Whisper

Alexander Polok^{*†}, Dominik Klement^{*†§}, Matthew Wiesner[§], Sanjeev Khudanpur[§], Jan Černocký[†], Lukáš Burget[†]
[†]Brno University of Technology, Brno, Czech Republic
[§]Johns Hopkins University, Baltimore, United States of America

Abstract—We propose a novel approach to enable the use of large, single-speaker ASR models, such as Whisper, for target speaker ASR. The key claim of this method is that it is much easier to model *relative* differences among speakers by learning to condition on frame-level diarization outputs than to learn the space of all speaker embeddings. We find that adding even a single bias term per diarization output type before the first transformer block can transform single-speaker ASR models into target-speaker ASR models. Our approach also supports speaker-attributed ASR by sequentially generating transcripts for each speaker in a diarization output. This simplified method outperforms baseline speech separation and diarization cascade by 12.9 % absolute ORC-WER on the NOTSOFAR-1 dataset.

Index Terms—target-speaker ASR, diarization conditioning, multi-speaker ASR, Whisper

I. INTRODUCTION

Self-supervised models [1]–[3], LLMs [4], [5], and Whisper-style supervised models [6], [7] have demonstrated that scaling up models by using more parameters and extremely large amounts of data can enable the development of accurate automatic speech recognition (ASR) systems, even in relatively challenging environments. However, these models have primarily been used in single-speaker, single-channel ASR systems, whereas most conversations involve multiple speakers and are often recorded by one or more microphones.

Approaches to handling this scenario generally integrate multiple blocks performing source separation, speaker segmentation, overlapped speech detection, post-hoc speaker clustering, and ASR in order to produce speaker-attributed conversation transcripts. Alternatively, there are end-to-end systems transcribing multi-speaker speech directly using special tokens or multiple heads [8]–[11]. Another approach is the use of a semi-end-to-end system, known as target-speaker ASR (TS-ASR) [12]–[17] processing the original input mixture and transcribing each speaker individually.

The conventional TS-ASR framework extracts speaker embeddings corresponding to the target speakers and uses them as an auxiliary input to the ASR system [16], [18]. While pre-trained speaker embedding extractors [19]–[21] can guide the ASR system by highlighting relevant information in the input and filtering out irrelevant content, they inherently require the system to learn how to map speaker embeddings to ASR speech embeddings. More recent methods include the use of adaptation layers and soft prompts to modify existing ASR models to incorporate speaker embeddings [14] or speaker

enrollment audios [17]. However, since these models are often trained on simulated datasets—due to the scarcity of multi-speaker ASR datasets and the need for a large number of speaker identities—they tend to experience significant performance degradation when applied to real-world multi-speaker scenarios [22]–[24].

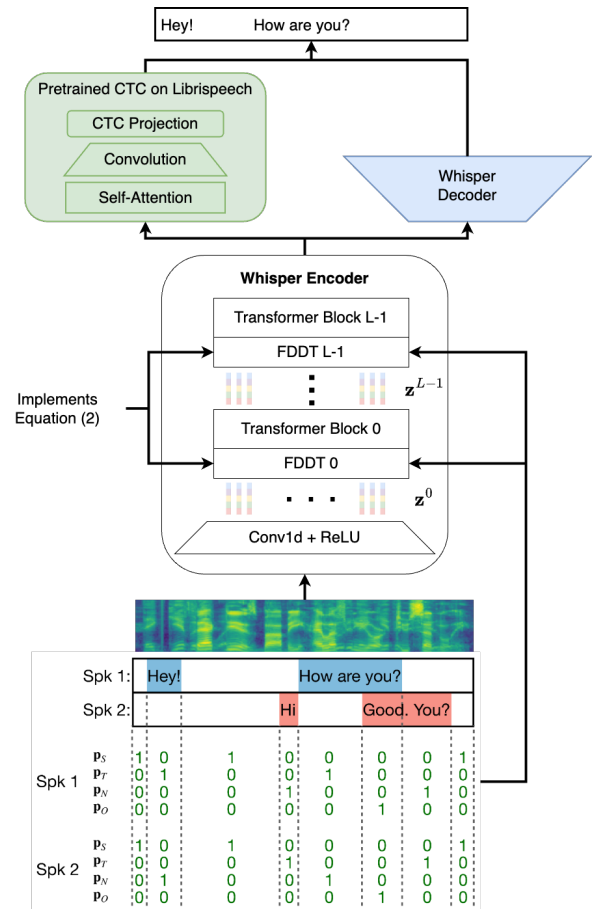


Fig. 1. Proposed Diarization-Conditioned Whisper model: An input audio segment with multiple speakers is conditioned by frame-level diarization outputs $[p_S^t \ p_T^t \ p_N^t \ p_O^t]^T$ for each STNO class at every frame t . Affine transformations, applied to intermediate input representations $z_{l:T}^t$, generate new embeddings, where l is the layer index. The final frame-level embedding is a convex combination of these embeddings for each frame.

In this paper, we propose a semi-end-to-end approach to TS-ASR that uses Whisper in a new way. Unlike previous TS-ASR methods, our system does not rely on speaker embeddings, but instead conditions directly on frame-level diarization

*These two authors contributed equally.

outputs. We believe that, compared to the aforementioned embedding-based approaches, only "relative" differentiation between speakers is needed; the TS-ASR system does not need to adapt to an existing subspace of speaker embeddings. Training our model on labeled examples of both target and non-target speech may also improve speaker discrimination and improve robustness to diarization errors.

To validate our approach, we fine-tune Whisper models on the NOTSOFAR-1 [22], AMI [23], and Libri2Mix [25] datasets using ground truth speaker segmentation. Unless stated otherwise, ground truth segmentation is also used during inference. All experiments follow the NOTSOFAR-1 Challenge guidelines¹.

II. DIARIZATION-CONDITIONED WHISPER

This section presents the Diarization-Conditioned Whisper, a model built upon the Whisper architecture, designed to perform TS-ASR by conditioning on frame-level diarization outputs. An overview of the proposed model is shown in Fig. 1. We adapt Whisper for TS-ASR by adding Frame-Level Diarization Dependent Transformations (FDDT) modules, described in Section II-C. These modules transform the model's internal representations in order to differentiate between the target- and non-target speakers in the audio.

A. Silence, Target, Non-Target, and Overlap Masks

Let $\mathbf{D} \in [0, 1]^{S \times T}$, where S is the number of speakers in the recording, and T is the number of frames, represent the diarization output, with $d(s, t)$ denoting the probability that speaker s is active at time frame t .

The dependency on the number of speakers in \mathbf{D} can be a limiting factor for easily incorporating this mask into the model. To address this, let s_k represent the target speaker. We define a distribution over the following mutually exclusive events for a frame at time t :

- \mathcal{S} : Silence at time t .
- \mathcal{T} : Only the target speaker s_k is active at t .
- \mathcal{N} : Non-target speaker(s) ($s \neq s_k$) active; target s_k inactive at t .
- \mathcal{O} : Target s_k and at least one non-target ($s \neq s_k$) active, causing overlap at t .

The probabilities of these events occurring at time frame t can be calculated as:

$$\begin{aligned} p_S^t &= \prod_{s=1}^S (1 - d(s, t)), & p_T^t &= d(s_k, t) \prod_{\substack{s=1 \\ s \neq s_k}}^S (1 - d(s, t)) \\ p_N^t &= (1 - p_S^t) - d(s_k, t), & p_O^t &= d(s_k, t) - p_T^t \end{aligned} \quad (1)$$

This definition allows us to use a fixed-sized target-speaker-dependent STNO (Silence, Target, Non-target, Overlap) mask $[p_S^t \ p_T^t \ p_N^t \ p_O^t]^\top$.

¹<https://www.chimechallenge.org/challenges/chime8/task2/>

B. Input Masking

Having the STNO mask, a straightforward way to perform target speaker ASR is to mask the signal by multiplying each frame by the probability that it is target speech or that it involves overlap with the target speaker. Hence, if the target speaker is not active, the audio signal is set to 0 (i.e., silence). We add p_T^t and p_O^t to ensure that both target speech and overlapping speech are preserved in the masked signal. However, similar to source separation approaches, this method has limitations. It can introduce artifacts because we are creating a modified version of the input signal, and errors in diarization can propagate through the system, potentially affecting the model's performance.

C. Frame-Level Diarization Dependent Transformations

To overcome issues of Input Masking, we designed a soft version called Frame-Level Diarization Dependent Transformations (FDDT). This approach modifies the frame-by-frame model inputs based on the diarization outputs.

Let $\mathbf{Z}^l \in \mathbb{R}^{d_m \times T}$ represent the frame-by-frame inputs to the l -th (Transformer) layer. We transform these hidden representations by applying four affine STNO layer- and class-specific transformations: $\mathbf{W}_S^l, \mathbf{W}_T^l, \mathbf{W}_N^l, \mathbf{W}_O^l \in \mathbb{R}^{d_m \times d_m}$ together with biases $\mathbf{b}_S^l, \mathbf{b}_T^l, \mathbf{b}_N^l, \mathbf{b}_O^l \in \mathbb{R}^{d_m}$ to obtain new speaker-specific hidden representations $\hat{\mathbf{Z}}^l = [\hat{\mathbf{z}}_1^l, \dots, \hat{\mathbf{z}}_T^l]$ as

$$\begin{aligned} \hat{\mathbf{z}}_t^l &= (\mathbf{W}_S^l \mathbf{z}_t^l + \mathbf{b}_S^l) p_S^t + (\mathbf{W}_T^l \mathbf{z}_t^l + \mathbf{b}_T^l) p_T^t \\ &\quad + (\mathbf{W}_N^l \mathbf{z}_t^l + \mathbf{b}_N^l) p_N^t + (\mathbf{W}_O^l \mathbf{z}_t^l + \mathbf{b}_O^l) p_O^t, \end{aligned} \quad (2)$$

These transformations generate four distinct representations of the frame-by-frame inputs, each highlighting one of the STNO classes. A target-speaker-specific representation is formed by a convex combination of these terms, with weights derived from the STNO mask. The same transformation is applied to all frames with identical STNO masks.

Fine-tuning the model using randomly initialized FDDT matrices could easily disrupt the internal representations of the model. Therefore, we propose initialization strategies to mitigate this risk:

- *Identity Initialization* (Non-Disturbing Init): Here, biases are initialized with zero vectors, and weights are initialized as identity matrices. This method ensures that the model's internal representations are not altered.
- *Suppressive Initialization*: To bias the model toward masking other speakers, we initialize the $\mathbf{W}_S^l, \mathbf{W}_N^l$ weights as diagonal matrices with values close to zero, e.g., 0.1. This approach helps the model to distinguish between different types of speech, reinforcing the separation between the STNO classes.

III. EXPERIMENTS

We primarily conducted our experiments on the new NOTSOFAR-1 dataset [22], which includes 280 meetings, each averaging 6 minutes, capturing diverse real-world acoustic conditions and conversational dynamics. To assess generalization and competitiveness, we also evaluated our best

TABLE I

COMPARISON OF THE PROPOSED SYSTEM BUILT WITH UPDATED WHISPER-LARGE-V3-TURBO [28] ALONGSIDE VARIOUS MULTI-TALKER ASR SYSTEMS. THE TOP SECTION INCLUDES SYSTEMS WHERE NO ADDITIONAL GROUND TRUTH INFORMATION ABOUT SPEAKER IDENTITY OR SEGMENTATION IS PROVIDED. THE BOTTOM SECTION FEATURES MODELS THAT DIRECTLY OR INDIRECTLY UTILIZE GROUND TRUTH SEGMENTATION INFORMATION. RESULTS MARKED WITH † ARE EVALUATED ON UTTERANCE GROUPS. PROPOSED ORC WER RESULTS MARKED WITH * WERE APPROXIMATED BY INCREASING THE COLLAR FOR tcORC WER.

	AMI-sdm test ORC WER	NOTSOFAR-1 eval-small tcORC WER	Libri2Mix test-both ORC WER
Raj et al. [8]	44.6 [†]	60.9 [†]	
Vinnikov et al. [22]		35.5	
Niu et al. [29]		17.7	
Proposed	18.0*	22.6	14.9
Input masking	79.1	76.6	56.7
Ma et al. [14]			26.4
Zhang et al. [15]			23.5
Proposed	16.5*	19.1	10.9

models on the synthetic Libri2Mix dataset [25] and on real-world meeting dataset AMI [23]. Experiments are divided into two parts: In Section III-C, we analyze the FDDT behavior under different weight structure constraints, initializations, numbers of additional parameters, and provided information. In Section III-D, we evaluate the framework’s performance as more parameters and data are added.

Source codes and recipes² are built on top of the transformers library [26]. All models are evaluated with the Time-Constrained Optimal Reference Combination Word Error Rate (tcORC WER) [27], referred to as WER throughout the text.

A. Training details

To enhance Whisper’s performance, we incorporated an additional CTC (Connectionist Temporal Classification) head, following the hybrid CTC-attention-based training scheme proposed in [30]. Given Whisper’s large 50k vocabulary size and fixed sequence length, adding an extra projection layer poses memory challenges. To overcome this, we added two convolutional layers, each with a subsampling factor of two, along with an additional self-attention layer. Both the CTC head and the decoder are trained with timestamp tokens, and the CTC loss weight is set to 0.3.

For the ablation experiments, we used Whisper-medium.en, while the final model was trained with Whisper-large-v3-turbo. All models are trained with an overall batch size of 64 samples using bf16 precision and the AdamW optimizer [31]. The learning rate is set to 2×10^{-6} , with a weight decay of 1×10^{-6} , a linear decay scheduler, and 2k warm-up steps. The new parameters introduced by FDDT are trained with a learning rate of 2×10^{-4} . By default, FDDT modules are inserted before all layers of the encoder with the diagonal constraint, meaning that only the diagonal values of the weight

TABLE II

ANALYSIS OF DIFFERENT CONSTRAINTS APPLIED TO THE FDDT PARAMETERS AND METHODS USED TO INITIALIZE THEM EVALUATED WITH WHISPER-MEDIUM.EN ON NOTSOFAR-1 EVAL-SMALL. THE FDDT PARAMETERS COLUMN SPECIFIES WHICH PARAMETERS ARE USED TO CONDITION THE MODEL. $\mathbf{W}_{diag} = diag(\mathbf{w})$, WHERE $diag(\mathbf{w})$ IS A DIAGONAL MATRIX WITH ELEMENTS OF \mathbf{w} IN THE DIAGONAL.

FDDT parameters	Initialization Method		
	Random	Identity	Suppressive
b	28.4	28.0	28.0
$\mathbf{W}_{diag}, \mathbf{b}$	129.4	27.3	26.7
\mathbf{W}, \mathbf{b}	129.0	46.1	44.6

matrices are updated during training and can be non-zero. Unless otherwise stated, the CTC head undergoes an initial “CTC preheating” phase, where it is trained on LibriSpeech for 10k steps, with the rest of the model being frozen. Afterwards, FDDT and CTC parameters are trained for a single epoch (FDDT preheating). Finally, the full model is trained for up to 50k steps, with early stopping set to a patience of 5 epochs. Most of the models typically converge within ten epochs. For the final evaluation, we always select the best-performing checkpoint based on the development set WER.

B. Comparison to Baselines

Table I presents a comparison of the proposed method with various end-to-end and modular systems. The top section includes systems where no additional ground truth information about speaker identity or segmentation is provided, and the system must infer this information. For the proposed system, we utilized DiariZen [32]³ to condition our model. The bottom section features models that directly or indirectly utilize ground truth segmentation information. It can be seen that our approach significantly outperforms the naive input masking baseline across all three datasets, largely due to its fine-tuning capabilities and robust handling of overlapped speech. Although our method does not achieve the top performance on the NOTSOFAR dataset, it surpasses baseline models, including the NOTSOFAR baseline [22] and the fine-tuned SURT model [8]. Furthermore, our approach delivers strong results on the AMI and Libri2Mix datasets, underscoring its effectiveness across diverse scenarios.

C. Frame-Level Diarization Dependent Transformation

To assess the impact of FDDT, we evaluated whether reducing the number of parameters in the additional modules affects the performance and examined the importance of proper initialization. As shown in Table II, using biases alone performs comparably to the combination of diagonal transformation matrices and biases, indicating that biasing frame-by-frame representations is sufficient to effectively focus the model on frames belonging to the same STNO class. The table also highlights that randomly initializing FDDT parameters is suboptimal and can significantly harm model performance,

²<https://github.com/BUTSpeechFIT/TS-ASR-Whisper>

³<https://github.com/BUTSpeechFIT/DiariZen>

TABLE III
EFFECT OF INCREASING THE NUMBER OF ENCODER LAYERS, WHERE FDDTs ARE APPLIED ON NOTSOFAR-1 EVAL-SMALL WITH WHISPER-MEDIUM.EN (24 ENCODER LAYERS).

FDDT parameters	# layers	Initilization Method		
		Random	Identity	Suppressive
b	1	28.7	30.9	29.3
	12	28.7	27.6	27.6
	24	28.4	28.0	28.0
W_{diag}, \mathbf{b}	1	117.7	27.8	27.0
	12	118.9	27.4	27.1
	24	129.4	27.3	26.7

suggesting that suppressive initialization is preferable. Interestingly, using non-restricted weights, meaning full matrices without diagonal constraints leads to noticeable performance degradation.

Table III demonstrates the effect of increasing the number of encoder layers where FDDTs are applied. The number of layers refers to how many layers, starting from the first encoder layer, are modified by FDDT parameters. For example, when only one layer is used, only the input to the Whisper encoder is modified. The results indicate that even a single layer of bias-only parameters can achieve performance comparable to the best diagonal setup. However, using a random initialization with diagonal matrices, even for just the first layer, severely impacts the model's performance. Notably, the use of a diagonal setup with both W_{diag} and \mathbf{b} across 24 layers yields the best performance, with the suppressive initialization method achieving the lowest WER of 26.7.

Table IV presents the performance of the FDDT method under different reductions of diarization information. It demonstrates that while the FDDT model still significantly outperforms input masking, some performance degradation is evident when comparing the STNO and T configurations. Interestingly, using three classes (TNO) results in worse performance than using (TN) classes only, even though the STNO and TNO configurations provide the same amount of information.

D. Scaling System With More Data And Parameters

Table V explores the impact of incorporating additional training data on system performance. The results show that adding the AMI dataset improves the performance on the NOTSOFAR-1, highlighting the benefits of more real-world data. Furthermore, incorporating Libri2Mix data provides ad-

TABLE IV
THE PERFORMANCE OF FDDT GIVEN A REDUCTION OF INFORMATION PROVIDED FROM DIARIZATION OUTPUT ON NOTSOFAR-1 EVAL-SMALL WITH WHISPER-MEDIUM.EN. WHEN EMPLOYING THE STNO MASK, ALL FRAMES ARE TRANSFORMED, WHILE WITH TNO, FRAMES CORRESPONDING TO THE SILENCE ARE LEFT UNCHANGED.

STNO	TNO	TN	T
26.7	30.0	28.7	34.8

TABLE V
DIFFERENT SIZES OF TRAINING CORPORA AFFECTING THE PERFORMANCE OF WHISPER-MEDIUM.EN. TESTED ON NOTSOFAR-1 EVAL SMALL.

NOTSOFAR-1	+ AMI	+ Libri2Mix
26.7	25.6	24.8

ditional gains, raising the question of whether pretraining on synthetic data could offer even greater performance gains.

Finally, Table VI provides a performance analysis across different model sizes, highlighting the improvements from employing an additional CTC head. Notably, incorporating a randomly initialized CTC head does not improve performance and leads to degradation. However, preheating the CTC head provides some improvements, especially for the small and medium models, even without joint decoding. Additionally, an FDDT preheating phase further enhances performance, with the best results observed in the large-v3 model.

TABLE VI
INFLUENCE OF CTC HEAD AND SIZE OF THE MODEL EVALUATED ON NOTSOFAR-1 EVAL-SMALL.

	small.en	medium.en	large-v3
without CTC	30.3	28.1	24.6
with CTC	31.0	28.9	25.8
+ CTC Preheating	29.2	26.7	25.2
+ FDDT Preheating	30.3	27.4	24.5

IV. CONCLUSIONS AND LIMITATIONS

In this study, we introduced FDDT into the Whisper model, making it a target-speaker ASR (TS-ASR) system conditioned on diarization output. We analyzed factors such as model size, FDDT placement, dataset size, and parameter initialization, evaluating the system on both real and synthetic datasets with consistently strong results. We also examined the impact of incorporating a CTC head, noting some improvements when it was preheated.

While the FDDT-based TS-ASR method performed effectively across diverse datasets, further validation on varied datasets, conditions, and languages is needed. We identified strengths and limitations, particularly regarding robustness against diarization errors, which need further attention for better real-world performance. Finally, the approach is extendable to other pre-trained ASR models, and a comparative analysis across ASR architectures could offer deeper insights into its adaptability and benefits.

ACKNOWLEDGEMENT

The work was supported by Ministry of Education, Youth and Sports of the Czech Republic (MoE) through the OP JAK project "Linguistics, Artificial Intelligence and Language and Speech Technologies: from Research to Applications" (ID:CZ.02.01.01/00/23_020/0008518) and Brno Ph.D. Talent Scholarship Programme. Computing on IT4I supercomputer was supported by MoE through the e-INFRA CZ (ID:90254).

REFERENCES

- [1] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *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.
- [2] W.-N. Hsu, Y.-H. H. Tsai, B. Bolte, R. Salakhutdinov, and A. Mohamed, “HuBERT: How much can a bad teacher benefit ASR pre-training?” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6533–6537.
- [3] V. Pratap, A. Tjandra, B. Shi, P. Tomasello, A. Babu, S. Kundu, A. Elkahky, Z. Ni, A. Vyas, M. Fazel-Zarandi *et al.*, “Scaling speech technology to 1,000+ languages,” *Journal of Machine Learning Research*, vol. 25, no. 97, pp. 1–52, 2024.
- [4] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, “GPT-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023.
- [5] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, “Llama: Open and efficient foundation language models,” *arXiv preprint arXiv:2302.13971*, 2023.
- [6] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *International conference on machine learning*. PMLR, 2023, pp. 28 492–28 518.
- [7] Y. Peng, J. Tian, B. Yan, D. Berrebbi, X. Chang, X. Li, J. Shi, S. Arora, W. Chen, R. Sharma, W. Zhang, Y. Sudo, M. Shakeel, J.-W. Jung, S. Maiti, and S. Watanabe, “Reproducing Whisper-style training using an open-source toolkit and publicly available data,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2023, pp. 1–8.
- [8] D. Raj, D. Povey, and S. Khudanpur, “Surt 2.0: Advances in transducer-based multi-talker speech recognition,” *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 31, p. 3800–3813, sep 2023. [Online]. Available: <https://doi.org/10.1109/TASLP.2023.3318398>
- [9] C. Li, Y. Qian, Z. Chen, N. Kanda, D. Wang, T. Yoshioka, Y. Qian, and M. Zeng, “Adapting multi-lingual ASR models for handling multiple talkers,” in *Interspeech 2023*, 2023, pp. 1314–1318.
- [10] N. Kanda, Y. Gaur, X. Wang, Z. Meng, and T. Yoshioka, “Serialized output training for end-to-end overlapped speech recognition,” in *Interspeech 2020*, 2020, pp. 2797–2801.
- [11] Y. Qian, X. Chang, and D. Yu, “Single-channel multi-talker speech recognition with permutation invariant training,” *Speech Communication*, vol. 104, pp. 1–11, 2018.
- [12] N. Kanda, S. Horiguchi, R. Takashima, Y. Fujita, K. Nagamatsu, and S. Watanabe, “Auxiliary interference speaker loss for target-speaker speech recognition,” in *Interspeech 2019*, 2019, pp. 236–240.
- [13] Y. Zhang, K. C. Puvvada, V. Lavrukhin, and B. Ginsburg, “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.
- [14] H. Ma, Z. Peng, M. Shao, J. Li, and J. Liu, “Extending Whisper with prompt tuning to target-speaker ASR,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 12 516–12 520.
- [15] W. Zhang and Y. Qian, “Weakly-supervised speech pre-training: A case study on target speech recognition,” in *Interspeech 2023*, 2023, pp. 3517–3521.
- [16] Z. Huang, D. Raj, P. García, and S. Khudanpur, “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.
- [17] L. Meng, J. Kang, Y. Wang, Z. Jin, X. Wu, X. Liu, and H. Meng, “Empowering Whisper as a joint multi-talker and target-talker speech recognition system,” in *Interspeech 2024*, 2024, pp. 4653–4657.
- [18] M. Karafiát, L. Burget, P. Matějka, O. Glembek, and J. Černocký, “iVector-based discriminative adaptation for automatic speech recognition,” in *2011 IEEE Workshop on Automatic Speech Recognition & Understanding*. IEEE, 2011, pp. 152–157.
- [19] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, “Front-end factor analysis for speaker verification,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2010.
- [20] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, “X-vectors: Robust DNN embeddings for speaker recognition,” in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 5329–5333.
- [21] H. Wang, C. Liang, S. Wang, Z. Chen, B. Zhang, X. Xiang, Y. Deng, and Y. Qian, “Wespeaker: A research and production oriented speaker embedding learning toolkit,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [22] A. Vinnikov, A. Ivry, A. Hurvitz, I. Abramovski, S. Koubi, I. Gurvich, S. Peer, X. Xiao, B. M. Elizalde, N. Kanda, X. Wang, S. Shaer, S. Yagev, Y. Asher, S. Sivasankaran, Y. Gong, M. Tang, H. Wang, and E. Krupka, “Notsofar-1 challenge: New datasets, baseline, and tasks for distant meeting transcription,” in *Interspeech 2024*, 2024, pp. 5003–5007.
- [23] I. Mccowan, J. Carletta, W. Kraaij, S. Ashby, S. Bourban, M. Flynn, M. Guillemot, T. Hain, J. Kadlec, V. Karaiskos, M. Kronenthal, G. Lathoud, M. Lincoln, A. Lisowska Masson, W. Post, D. Reidsma, and P. Wellner, “The AMI meeting corpus,” *Int’l. Conf. on Methods and Techniques in Behavioral Research*, 01 2005.
- [24] F. Yu, S. Zhang, Y. Fu, L. Xie, S. Zheng, Z. Du, W. Huang, P. Guo, Z. Yan, B. Ma, X. Xu, and H. Bu, “M2Met: The ICASSP 2022 multi-channel multi-party meeting transcription challenge,” *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6167–6171, 2021.
- [25] J. Cosentino, M. Pariente, S. Cornell, A. Deleforge, and E. Vincent, “LibriMix: An open-source dataset for generalizable speech separation,” 2020. [Online]. Available: <https://arxiv.org/abs/2005.11262>
- [26] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush, “Transformers: State-of-the-art natural language processing,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Online: Association for Computational Linguistics, Oct. 2020, pp. 38–45. [Online]. Available: <https://www.aclweb.org/anthology/2020.emnlp-demos.6>
- [27] T. v. Neumann, C. B. Boeddeker, M. Delcroix, and R. Haeb-Umbach, “MeetEval: A Toolkit for Computation of Word Error Rates for Meeting Transcription Systems,” in *Proceedings of the 7th International Workshop on Speech Processing in Everyday Environments (CHIEME 2023)*, 2023, pp. 27–32.
- [28] A. Polok, D. Klement, M. Kocour, J. Han, F. Landini, B. Yusuf, M. Wiesner, S. Khudanpur, J. Černocký, and L. Burget, “Dicow: Diarization-conditioned whisper for target speaker automatic speech recognition,” 2024, submitted to Computer Speech & Language (CSL) journal special issue on Multi-Speaker, Multi-Microphone, and Multi-Modal Distant Speech Recognition. [Online]. Available: <https://arxiv.org/abs/2501.00114>
- [29] S. Niu, R. Wang, J. Du, G. Yang, Y. Tu, S. Wu, S. Qian, H. Wu, H. Xu, X. Zhang, G. Zhong, X. Yu, J. Chen, M. Wang, D. Cai, T. Gao, G. Wan, F. Ma, J. Pan, and J. Gao, “The USTC-NERCSLIP Systems for the CHIEME-8 NOTSOFAR-1 Challenge,” in *8th International Workshop on Speech Processing in Everyday Environments (CHIEME 2024)*, 2024, pp. 31–36.
- [30] T. Hori, S. Watanabe, and J. Hershey, “Joint CTC/attention decoding for end-to-end speech recognition,” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, R. Barzilay and M.-Y. Kan, Eds. Vancouver, Canada: Association for Computational Linguistics, Jul. 2017, pp. 518–529. [Online]. Available: <https://aclanthology.org/P17-1048>
- [31] I. Loshchilov and F. Hutter, “Decoupled Weight Decay Regularization,” in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=Bkg6RiCqY7>
- [32] J. Han, F. Landini, J. Rohdin, A. Silnova, M. Diez, and L. Burget, “Leveraging self-supervised learning for speaker diarization,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025.