

AG-LSEC: Audio Grounded Lexical Speaker Error Correction

Rohit Paturi, Xiang Li, Sundararajan Srinivasan

AWS AI Labs

paturi@amazon.com, xiangzai@amazon.com, sundarsr@amazon.com

Abstract

Speaker Diarization (SD) systems are typically audio-based and operate independently of the ASR system in traditional speech transcription pipelines and can have speaker errors due to SD and/or ASR reconciliation, especially around speaker turns and regions of speech overlap. To reduce these errors, a Lexical Speaker Error Correction (LSEC), in which an external language model provides lexical information to correct the speaker errors, was recently proposed. Though the approach achieves good Word Diarization error rate (WDER) improvements, it does not use any additional acoustic information and is prone to miscorrections. In this paper, we propose to enhance and acoustically ground the LSEC system with speaker scores directly derived from the existing SD pipeline. This approach achieves significant relative WDER reductions in the range of 25-40% over the audio-based SD, ASR system and beats the LSEC system by 15-25% relative on RT03-CTS, Callhome American English and Fisher datasets.

Index Terms: Speaker Diarization, Language Models, Automatic Speech Recognition, Error Correction

1. Introduction

Multi-talker speech transcription systems answer the question "Who spoke what and when?" in an audio recording. Though single-talker speech transcription systems have advanced significantly in the recent years, there are still challenges in transcribing and diarizing multi-talker speech due to the regular speaker changes and overlapping speech in natural conversations. Multi-talker transcription systems can be broadly classified in three categories of systems that aim to transcribe and diarize multi-talker speech: 1) Speech/Speaker Separation followed by ASR [1-4], 2) Speaker attributed ASR [5-9] and 3) Modular ASR [10-13], SD systems [14-18].

Speech Separation systems separate speech from multiple speakers into different channels while performing ASR individually on these channels, thus assigning a speaker label to each transcribed word. These systems are usually trained with Permutation Invariant Training (PIT) which constrains the maximum number of speakers it can handle [1, 2], and is also prone to duplicated artifacts across channels among other challenges to handle long-form audio [3]. Speaker attributed ASR (SA-ASR) systems either follow PIT based training [8, 9] which have the same challenges as above or the more recent Serialized Output Training (SOT) [6, 7] systems which outputs transcripts of each speaker sequentially. Main drawbacks with SOT systems are the lack of timestamps for the speaker turns needed for downstream applications, and the difficulty of speaker reconciliation across transcribed blocks for long-form audios. Training both speech separation and SA-ASR systems is data intensive and needs large

amounts of multi-channel audio data. Such data is difficult to source, especially for domains outside of telephony, and could also increase cost significantly. Finally, independent modular ASR and SD systems is one of the most successful and practical way to transcribe multi-talker audios. The ASR and SD systems operate independently and their results are reconciled to assign speaker labels to each transcribed word [19]. Such systems can naturally transcribe and diarize long-form audio by operating on independent speech segments without significantly constraining the maximum number of speakers, though these approaches are error-prone around speaker change or overlapping speech regions due to SD or reconciliation errors.

Most SD systems in literature rely only on acoustic information [15, 16] to either extract speaker embeddings [20, 21] followed by clustering the embeddings or directly perform diarization with End-to-End Neural Diarization (EEND) [16, 17]. Reliance on only acoustic information can lead to speaker errors, mainly around the speaker turns and regions of speech overlaps [22, 23]. In addition to the SD errors, speakers can be attributed to the wrong words in the SD-ASR reconciliation phase also due to errors solely in ASR word timings [24]. The reconciliation errors will be higher in regions of speech overlap as a traditional ASR system can identify words corresponding to one of the overlapping speakers which can be wrongly attributed to the other overlapping speaker.

Lexical information often contains complementary information which can be useful in accurately assigning speaker labels to the transcribed words [19, 22, 23]. For instance, analyzing only the written transcript of a conversation, "how are you i am good", enables us to infer that there is likely a speaker change between the utterances "how are you" and "i am good". There have been few works [22, 23] which leverage lexical information to perform standalone SD. Though these reduce the SD errors, these are still prone to ASR reconciliation errors around speaker turns and overlapped speech similar to the other acoustic only approaches. Another work [19] models ASR and SD jointly while [25] employs a large language model (LLM) to assign speaker roles to a non-speaker-attributed transcript, but both these works are limited to two speakers with distinct roles.

2. Related Work

Recent works [26, 27] have explored leveraging a LLM for performing word level SD or correcting speaker errors after ASR reconciliation. [27] models SD as a beam search at the word level, with distinct speaker paths for each word in the transcript, where the probability that a given word is emitted by a particular speaker is computed using a LLM. However, a significant drawback of this approach is the high computational cost and latency involved in prompting the LLM for each word




Figure 1: (a) AG-LSEC with Late Fusion of Speaker Scores, (b) Speaker Scores Extraction from the SD pipeline, (c) AG-LSEC with Early Fusion of Speaker Scores

in the transcript, as well as the need to tune the LLM to emit the response in a specific format, which otherwise can be susceptible to the non-deterministic nature of the LLM.

The authors of [26] utilize a LLM to correct the outputs of a speaker-attributed transcript using a transcript preserving mechanism to prevent LLM hallucinations. They demonstrate that a zero or one-shot approach using a frozen LLM degrades diarization performance, and fine-tuning the LLM provides good WDER improvements. However, fine-tuning LLMs [26, 27] can be data-intensive and computationally expensive and can result in higher inference latency. In addition, methods involving LLMs would warrant special mechanisms to prevent LLM hallucinations [26].

An alternative method [24] proposes a simple yet effective way to perform speaker error correction by applying a smaller encoder language model to the transcriptions and fusing these with speaker labels from SD to correct speaker and reconciliation errors. One limitation of [24–26] is their reliance solely on lexical information, which can lead to both under-corrections and over-corrections, as conversational speech often contains artifacts such as speech overlaps, disfluencies, incomplete speaker turns, etc may not be captured in text-based language models.

In this work, we propose to enhance the LSEC model in [24] by grounding the model with acoustic information extracted from the speaker diarization module.

3. Audio Grounded Lexical Error Corrector

In this section, we introduce our proposed approach for leveraging both acoustic and lexical information from the speech signals to perform the speaker error correction. The proposed AG-LSEC consists of three main components: an acoustic speaker score extractor, backbone language model (LM) and a Transformer Encoder Front-end to predict the speaker labels. We first introduce the extraction of speaker scores from the EEND system followed by two fusion techniques to incorporate the speaker scores into the lexical speaker error correction model.

3.1. Acoustic Speaker Scores

3.1.1. End-to-end Neural Diarization

We choose the End-to-end neural diarization (EEND) for our SD system due to its ability to handle overlapping speech efficiently compared to traditional embedding based approaches. In addition, EEND is a natural choice also because it directly provides soft speaker scores that can be used to ground lexical error correc-

tion. Given frame-level acoustic features $\{\mathbf{x}_t\}_{t=1}^T$, $\mathbf{x}_i \in \mathbb{R}^{1 \times F}$, where $t \in \{1, \dots, T\}$ is the frame index and F is the feature dimension, EEND estimates the corresponding speaker label sequence $\{\mathbf{y}_t\}_{t=1}^T$ where $\mathbf{y}_t = \{y_{s,t}\}_{s=1}^S$ denotes speech activities of S speakers at frame t and is defined as

$$y_{s,t} = \begin{cases} 0 & (\text{Speaker } s \text{ is inactive at } t) \\ 1 & (\text{Speaker } s \text{ is active at } t) \end{cases} \quad (1)$$

Given the sequence of frame-level acoustic features $\{\mathbf{x}_t\}_{t=1}^T$, EEND models speaker diarization as a multi-label classification problem using a neural network f_{EEND} as

$$(\mathbf{p}_1, \dots, \mathbf{p}_T) = f_{\text{EEND}}(\mathbf{x}_1, \dots, \mathbf{x}_T), \quad (2)$$

where $\mathbf{p}_t = \{p_{s,t}\}_{s=1}^S$ are the frame level posterior probabilities of S speakers at frame index t . The EEND system is trained to minimize the permutation-free loss between the output posteriors \mathbf{p}_t and the reference speaker label \mathbf{y}_t , as follows:

$$\mathcal{L}_{\text{diar}} = \frac{1}{TS} \min_{\phi \in \text{perm}(S)} \sum_t \text{BCE}(\mathbf{y}_t^\phi, \mathbf{p}_t), \quad (3)$$

where $\text{perm}(S)$ is the set of all permutations of the sequence $(1, \dots, S)$, $\phi = (\phi_1, \dots, \phi_S)$ is the permuted sequence, $\mathbf{y}_t^\phi = [y_{\phi_1,t}, \dots, y_{\phi_S,t}] \in \{0, 1\}^S$ is the permuted ground-truth labels using ϕ , and $\text{BCE}(\cdot, \cdot)$ is the binary cross entropy loss.

3.1.2. Posterior post-processing

The EEND frame posteriors tend to be noisy and can produce unreasonably short speaker turns. So, a median filter is applied to the posteriors to produce the smoothed posteriors $\hat{\mathbf{p}}_t = \{\hat{p}_{s,t}\}_{s=1}^S$ as

$$\hat{p}_{s,t} = \text{median}(p_{s,t-M/2}, p_{s,t-M/2+1}, \dots, p_{s,t+M/2}) \quad (4)$$

where M is the number of frames over which median is calculated. The ASR system produces a sequence of words $\{\mathbf{W}_i\}_{i=1}^N$ along with a set of word timings which can be converted to EEND frame indices $\{t_{i,\text{start}}\}_{i=1}^N$, $\{t_{i,\text{end}}\}_{i=1}^N$ which denote the start and end frames corresponding to each word respectively, where N is the number of transcribed words in the speech segment. The frame-level smoothed posteriors $\hat{p}_{s,t}$ are pooled at the word-level to provide word-level aggregated posteriors $\{\mathbf{a}_i\}_{i=1}^N$, $\mathbf{a}_i = \{a_{s,i}\}_{s=1}^S$ as

$$a_{s,i} = \frac{1}{t_{i,\text{end}} - t_{i,\text{start}}} \sum_{t=t_{i,\text{start}}}^{t_{i,\text{end}}} \hat{p}_{s,t} \quad (5)$$

Table 1: WDER of different models on Fisher test, RT03-CTS and CHAE sets.

Model Type	Fisher Test	RT03-CTS		CHAE		CH109
		Validation	Test	Validation	Test	
Baseline (No Correction)	2.56	2.37	2.64	3.98	3.45	4.06
LSEC [24]	2.03	1.93	2.10	3.5	2.89	3.56
AG-LSEC Late Fusion	1.80	1.70	1.80	3.26	2.72	3.32
AG-LSEC Early Fusion	1.56	1.58	1.56	2.86	2.48	3.01

The word-level posteriors are normalized in the range (0,1) to produce the word-level acoustic speaker scores $\{\hat{a}_i\}_{i=1}^N$, $\hat{a}_i = \{\hat{a}_{s,i}\}_{s=1}^S$ where

$$\hat{a}_{s,i} = \frac{a_{s,i}}{\sum_{s=1}^S a_{s,i}} \quad (6)$$

3.2. Acoustic Grounding

3.2.1. Early Fusion

The words W_i are tokenized into sub-words and passed to the backbone LM to obtain contextual sub-word embeddings $\{\mathbf{E}_i\}_{i=1}^{N_1}$, $\mathbf{E}_j \in \mathbb{R}^{1 \times W}$ where N_1 is the number of sub-words in the word sequence and W is the word embedding dimension.

$$\{\mathbf{E}_i\}_{i=1}^{N_1} = \text{EncoderLM}(\{\mathbf{W}_i\}_{i=1}^N) \quad (8)$$

The word level speaker scores \hat{a}_i are mapped to the first sub-word token of the word if the word has more than 1 token and a special “don’t care” vector is assigned to any of the subsequent word tokens. The corrected word-level speaker posteriors $\{\mathbf{z}_i\}_{i=1}^N$, $\mathbf{z}_i \in \mathbb{R}^{1 \times M}$ are then extracted as

$$\{\mathbf{c}_i\}_{i=1}^{N_1} = \text{concat}(\{\mathbf{E}_i\}_{i=1}^{N_1}, \{\hat{a}_i\}_{i=1}^{N_1}) \quad (9)$$

$$\{\mathbf{z}_i\}_{i=1}^N = \text{Softmax}(\text{Encoder}(\{\mathbf{c}_i\}_{i=1}^{N_1})) \quad (10)$$

where posteriors corresponding to the first sub-word tokens are selected to represent the word level speaker posteriors. These posteriors are optimized for Cross Entropy (CE) loss on the ground-truth speaker labels. This early fusion process is shown in Figure 1c.

3.2.2. Late Fusion

The word level speaker labels for the LSEC model are obtained by assigning the speaker with the highest word-level acoustic scores from (6) for each word. For the acoustic grounding, we perform lexical correction with the LSEC model and fuse the word level acoustic speaker scores, $\{\hat{a}_i\}_{i=1}^N$ with the corresponding word level lexical posteriors from the LSEC model. A simple Neural network (FusionNet) takes in the acoustic and lexical scores and predicts the correct speaker label as shown in Figure 1a.

3.3. Training Methodology

The AG-LSEC model is trained by initializing it with the LSEC model (impact of the initialization is shown in Figure 2) since the speaker scores used in the AG-LSEC correspond to the soft scores of the speaker labels used to train the LSEC model. It is trained with paired audio-text data, where the audio is diarized with the SD system, transcribed by an ASR system and reconciled to get the words and the word level acoustic scores as in Figure 1.

Similar to the AG-LSEC training data, word-level speaker scores are extracted using the paired audio-text data and the

lexical posteriors are extracted using the LSEC model which are fed to the Fusion Net to predict the corrected speaker label.

The AG-LSEC, LSEC and Fusion Net are all trained over local windows with two hypothesized speakers.

3.4. Inference Setup

During inference, we perform error correction on overlapping sliding windows with a fixed number of ASR transcribed words. Similar to LSEC [24], AG-LSEC can also handle cases where more than two speakers are detected globally across the audio by only correcting sliding windows comprising of one or two speakers and bypassing the remaining windows.

4. Experiments

4.1. Data and Metrics

We use the Fisher dataset [28, 29] to train the AG-LSEC system. We create a single channel version of Fisher by merging both the channels and leverage the transcripts for each channel along with the timestamps as the ground-truths similar to [24, 26] and we use the same held-out Fisher test split for evaluation as used in [24, 26, 30]. For evaluation, in addition to the held out Fisher test split, we use 2 other telephony datasets, namely CALLHOME American English (CHAE) [31] and RT03-CTS [32].

The metric we use to evaluate our model is the Word Diarization Error Rate (WDER) proposed in [19] as it aptly captures both ASR and SD errors at the word level and accounts for errors in the overlapping speech regions as well. In order to align the multi-speaker references to multi-speaker hypothesis including regions of speech overlap, we use aslrite tool [33].

4.2. Baseline System

The baseline EEND system follows [16] and consists of 6 stacked self-attention-based Transformer layers, 8 attention heads with a hidden size of 256 and 1024 internal units in the position-wise feed-forward layer. It is trained on 3-speaker simulated mixtures similar to [34] and adapted on the CALLHOME (CH) dataset [35]. During inference, a median filter is applied over 11 frames in the post-processing of EEND posteriors and the entire audio is diarized in a single pass without the use of any segmentation. This baseline EEND system achieves a DER of 7.36 on the CH test set compared to a DER of 10.76 reported in [16].

We use an end-to-end ASR system [11, 36] with a Conformer Acoustic model [37] and a n-gram Language model trained on several tens of thousands of audio and text data. The ASR system achieves a WER of 12.2 on the single channel Fisher test set compared to 15.48 reported in [26]. The combined SD-ASR system achieves a WDER of 2.56 on Fisher test set which is a much stronger baseline compared to the WDER of 5.32 reported in [26].




Figure 2: Ablation of the AG-LSEC models with different amounts of training data and initializations.

4.3. AG-LSEC System

We follow the same architecture and training procedure for the LSEC model as outlined in [24] which uses the Roberta-base model [38] as the backbone LM and a 1 layer Transformer Encoder of size 128 hidden states for the Front-end model. We use this same model for our late fusion experiments. The backbone and front-end model architectures are the same for the AG-LSEC model as well. The AG-LSEC model is trained with Adam Optimizer with a batch size of 32 and an average sequence length of 30 words per batch. We use a learning rate of $1e-4$ and train the model for 20 epochs on a machine with 8 GPUs.

4.4. Results

We benchmark WDER on the test sets with our proposed acoustic grounding techniques and compare these to the baseline ASR, SD system as well as the LSEC model in Table 1. It can be observed that the proposed AG-LSEC model with early fusion achieves significant relative WDER improvements of 25% to 40% across all the test sets and also improves the WDER by more than 15% relative over the LSEC model. The late fusion model also improves the WDER over the LSEC model, though produces smaller relative gains in the order of of 5-10%.

In order to analyze the paired data requirements for training the models, we evaluate the WDER of the models on Fisher test set when trained with limited amount of speaker tagged, paired audio-text data. For this analysis specifically, we train the LSEC model only on text data with synthetic simulated errors [24] and do not include any paired audio-text data. From Figure 2, it can be observed that the late fusion model is able to achieve the same consistent improvement of 7% over the text only trained LSEC model with as little as 20 minutes of paired data. It can also be observed that the LSEC initialization for the AG-LSEC model surpasses the late fusion model around 10hrs of data while also much better than the random initialized AG-LSEC model till almost 100hrs of training data, while still being marginally better beyond 100hrs of data. This shows the importance of the LSEC initialization in both cutting down the data requirements as well to achieve better WDER improvements.

We further analyze the behaviour of the AG-LSEC models in correcting errors whilst also introducing new errors over the baseline SD-ASR system. For this, we benchmark the LSEC model, AG-LSEC model with Early and Late Fusion. These results are documented in Table 2. It can be seen that the AG-LSEC Early Fusion introduces lesser errors than the LSEC model, thus

Table 2: Percentage of errors corrected and new errors introduced with LSEC and variants of AG-LSEC on Fisher test set.

Model	Errors	
	Corrected (%) \uparrow	Introduced (%) \downarrow
LSEC [24]	29.2	8.4
AG-LSEC Late Fusion	38.18	9.2
AG-LSEC Early Fusion	44.53	6.6

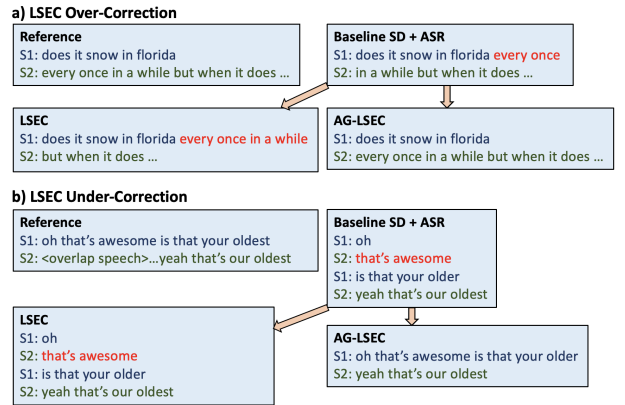


Figure 3: Qualitative examples of over-correction and under-correction with LSEC model rectified by the AG-LSEC Early Fusion model. The words highlighted in red are speaker errors.

preventing over-corrections. At the same time, both Early and Late Fusion correct more speaker errors than the LSEC model, which also suffers from under-corrections due to its excessive reliance on lexical cues.

Some qualitative examples of such over-corrections and under-corrections by the LSEC model and its rectification by the AG-LSEC model are presented in Figure 3. It can be seen from 3a) that the LSEC model over-corrects to a lexically more plausible yet incorrect speaker transcript which the audio grounding prevents with the AG-LSEC model. Similarly, from 3b), it can be seen that the LSEC model fails to correct the lexically plausible baseline due to its excessive dependence on lexical characteristics. On the contrary, the AG-LSEC model has the potential to correct such errors due to the inclusion of acoustic speaker scores along with the lexical features, thus improving the correction performance significantly.

5. Conclusion and Future Work

In this work, we propose a novel acoustic fusion strategy to the recently proposed Lexical Speaker Error Corrector (LSEC) to correct word-level speaker label errors from a conventional multi-talker speech transcription system. We extract acoustic speaker scores directly from the existing EEND pipeline and outline two fusion strategies by leveraging these speaker scores, both showing improvements over the baseline LSEC system. The early fusion AG-LSEC model achieves significant relative WDER improvements of above 25% on the test datasets over the the baseline SD-ASR system and also handily beats the LSEC model with relative WDER improvements of over 15%.

In future, we will experiment with a multi-lingual AG-LSEC model to handle languages other than English and will handle more than 2 speakers locally to increase the coverage of lexical correction.

6. REFERENCES

- [1] T. von Neumann, K. Kinoshita, L. Drude, *et al.*, “End-to-end training of time domain audio separation and recognition,” in *ICASSP 2020*, pp. 7004–7008, IEEE, 2020.
- [2] S. Chen, Y. Wu, Z. Chen, *et al.*, “Continuous speech separation with conformer,” in *ICASSP 2021*, pp. 5749–5753, 2021.
- [3] R. Paturi, S. Srinivasan, *et al.*, “Directed speech separation for automatic speech recognition of long form conversational speech,” in *Proc. INTERSPEECH 2022*, pp. 5388–5392, 2022.
- [4] S. Zhao and B. Ma, “Mossformer: Pushing the performance limit of monaural speech separation using gated single-head transformer with convolution-augmented joint self-attentions,” in *ICASSP 2023*, pp. 1–5, IEEE, 2023.
- [5] X. Chang, W. Zhang, Y. Qian, *et al.*, “End-to-end multi-speaker speech recognition with transformer,” in *ICASSP 2020*, pp. 6134–6138, IEEE, 2020.
- [6] N. Kanda, J. Wu, Y. Wu, *et al.*, “Streaming Multi-Talker ASR with Token-Level Serialized Output Training,” in *Proc. Interspeech 2022*, pp. 3774–3778, 2022.
- [7] M. Shi, Z. Du, Q. Chen, *et al.*, “CASA-ASR: Context-Aware Speaker-Attributed ASR,” in *Proc. INTERSPEECH 2023*, pp. 411–415, 2023.
- [8] I. Sklyar, A. Piunova, and Y. Liu, “Streaming multi-speaker asr with rnn-t,” in *ICASSP 2021*, pp. 6903–6907, IEEE, 2021.
- [9] S. Berger, P. Vieting, C. Boeddeker, *et al.*, “Mixture Encoder for Joint Speech Separation and Recognition,” in *Proc. INTERSPEECH 2023*, pp. 3527–3531, 2023.
- [10] Y. Wang, A. Mohamed, D. Le, *et al.*, “Transformer-based acoustic modeling for hybrid speech recognition,” in *ICASSP 2020*, pp. 6874–6878, IEEE, 2020.
- [11] Y. Yang, P. Wang, and D. Wang, “A conformer based acoustic model for robust automatic speech recognition,” *arXiv preprint arXiv:2203.00725*, 2022.
- [12] D. Rekish, N. R. Koluguri, S. Krizan, *et al.*, “Fast conformer with linearly scalable attention for efficient speech recognition,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pp. 1–8, IEEE, 2023.
- [13] 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*, pp. 28492–28518, PMLR, 2023.
- [14] T. J. Park, N. Kanda, D. Dimitriadis, *et al.*, “A review of speaker diarization: Recent advances with deep learning,” *Computer Speech & Language*, vol. 72, p. 101317, 2022.
- [15] Q. Wang, C. Downey, L. Wan, P. A. Mansfield, and I. L. Moreno, “Speaker diarization with lstm,” in *ICASSP 2018*, pp. 5239–5243, IEEE, 2018.
- [16] Y. Fujita, N. Kanda, S. Horiguchi, *et al.*, “End-to-end neural speaker diarization with self-attention,” in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pp. 296–303, IEEE, 2019.
- [17] K. Kinoshita, M. Delcroix, and N. Tawara, “Integrating end-to-end neural and clustering-based diarization: Getting the best of both worlds,” in *ICASSP 2021*, pp. 7198–7202, IEEE, 2021.
- [18] H. Bredin, “pyannote.audio 2.1 speaker diarization pipeline: principle, benchmark, and recipe,” in *Proc. INTERSPEECH 2023*, pp. 1983–1987, 2023.
- [19] L. E. Shafey, H. Soltan, and I. Shafran, “Joint Speech Recognition and Speaker Diarization via Sequence Transduction,” in *Proc. Interspeech 2019*, pp. 396–400, 2019.
- [20] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, “Generalized end-to-end loss for speaker verification,” in *ICASSP 2018-2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4879–4883, IEEE, 2018.
- [21] N. Dawalatabad, M. Ravanelli, F. Grondin, *et al.*, “ECAPA-TDNN Embeddings for Speaker Diarization,” in *Proc. Interspeech 2021*, pp. 3560–3564, 2021.
- [22] T. J. Park, K. J. Han, J. Huang, *et al.*, “Speaker Diarization with Lexical Information,” in *Proc. Interspeech 2019*, pp. 391–395, 2019.
- [23] W. Xia, H. Lu, Q. Wang, A. Tripathi, *et al.*, “Turn-to-diarize: Online speaker diarization constrained by transformer transducer speaker turn detection,” in *ICASSP 2022*, pp. 8077–8081, IEEE, 2022.
- [24] R. Paturi, S. Srinivasan, and X. Li, “Lexical Speaker Error Correction: Leveraging Language Models for Speaker Diarization Error Correction,” in *Proc. INTERSPEECH 2023*, pp. 3567–3571, 2023.
- [25] A. Adedeji, S. Joshi, and B. Doohan, “The sound of healthcare: Improving medical transcription asr accuracy with large language models,” *arXiv preprint arXiv:2402.07658*, 2024.
- [26] Q. Wang, Y. Huang, G. Zhao, E. Clark, W. Xia, and H. Liao, “Diarizationlm: Speaker diarization post-processing with large language models,” *arXiv preprint arXiv:2401.03506*, 2024.
- [27] T. J. Park, K. Dhawan, N. Koluguri, *et al.*, “Enhancing speaker diarization with large language models: A contextual beam search approach,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 10861–10865, IEEE, 2024.
- [28] C. Cieri *et al.*, “Fisher english training speech part 1 speech ldc2004s13,” *Web Download. Philadelphia: Linguistic Data Consortium*, 2004.
- [29] C. Cieri *et al.*, “Fisher english training part 2, speech ldc2005s13,” *Web Download. Philadelphia: Linguistic Data Consortium*, 2005.
- [30] Q. Wang, Y. Huang, H. Lu, G. Zhao, and I. L. Moreno, “Highly efficient real-time streaming and fully on-device speaker diarization with multi-stage clustering,” *arXiv preprint arXiv:2210.13690*, 2022.
- [31] D. G. Canavan, Alexandra and G. Zipperlen, “Callhome american english speech ldc97s42,” *Web Download. Philadelphia: Linguistic Data Consortium*, 1997.
- [32] J. G. Fiscus *et al.*, “2003 nist rich transcription evaluation data ldc2007s10,” *Web Download. Philadelphia: Linguistic Data Consortium*, 2007.
- [33] J. G. Fiscus, J. Ajot, N. Radde, C. Laprun, *et al.*, “Multiple dimension levenshtein edit distance calculations for evaluating automatic speech recognition systems during simultaneous speech,” in *LREC*, pp. 803–808, Citeseer, 2006.
- [34] Y. Fujita, N. Kanda, S. Horiguchi, K. Nagamatsu, and S. Watanabe, “End-to-End Neural Speaker Diarization with Permutation-Free Objectives,” in *Proc. Interspeech 2019*, pp. 4300–4304, 2019.
- [35] G. R. Doddington, M. A. Przybocki, A. F. Martin, and D. A. Reynolds, “The nist speaker recognition evaluation—overview, methodology, systems, results, perspective,” *Speech communication*, vol. 31, no. 2-3, pp. 225–254, 2000.
- [36] W. Zhou, W. Michel, K. Irie, *et al.*, “The rwth asr system for ted-lium release 2: Improving hybrid hmm with specaugment,” in *ICASSP 2020*, pp. 7839–7843, IEEE, 2020.
- [37] A. Gulati, J. Qin, C.-C. Chiu, *et al.*, “Conformer: Convolution-augmented Transformer for Speech Recognition,” in *Proc. Interspeech 2020*, pp. 5036–5040, 2020.
- [38] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, “Roberta: A robustly optimized bert pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.