arXiv:2204.11232v1 [eess.AS] 24 Apr 2022

Improving the Naturalness of Simulated Conversations for End-to-End Neural Diarization

Natsuo Yamashita[1*] *Shota Horiguchi*[2] *Takeshi Homma*[2]

[1]The University of Tokyo
[2]Hitachi, Ltd.
natsuo-yamashita@g.ecc.u-tokyo.ac.jp, {shota.horiguchi.wk,takeshi.homma.ps}@hitachi.com

Abstract

This paper investigates a method for simulating natural conversation in the model training of end-to-end neural diarization (EEND). Due to the lack of any annotated real conversational dataset, EEND is usually pretrained on a large-scale simulated conversational dataset first and then adapted to the target real dataset. Simulated datasets play an essential role in the training of EEND, but as yet there has been insufficient investigation into an optimal simulation method. We thus propose a method to simulate natural conversational speech. In contrast to conventional methods, which simply combine the speech of multiple speakers, our method takes turn-taking into account. We define four types of speaker transition and sequentially arrange them to simulate natural conversations. The dataset simulated using our method was found to be statistically similar to the real dataset in terms of the silence and overlap ratios. The experimental results on two-speaker diarization using the CALLHOME and CSJ datasets showed that the simulated dataset contributes to improving the performance of EEND.

Index Terms— speaker diarization, simulated conversation, conversation analysis, turn-taking

1. Introduction

Speaker diarization is the task of identifying speech segments and their speakers from audio or video recordings; in other words, a task to identify "who spoke when" [1]. It is widely utilized in a variety of applications such as meeting transcription [2, 3], conversational interaction analysis [4], content-based audio indexing [5], and conversational AI [6]. It also helps improve the accuracy of automatic speech recognition (ASR) in multi-speaker conversations [7].

A typical approach for speaker diarization [8, 9, 10, 11, 12, 13, 14] is a cascade of the following steps: speech activity detection, speaker embedding extraction, and clustering. In general, speech activity detectors and speaker embeddings extractors are constructed using neural networks, which require a large amount of training data to obtain a good performance. Since these modules can be trained from only single-speaker recordings, they can directly leverage existing large-scale datasets such as VoxCeleb [15] and SITW [16].

Alternatively, the end-to-end approach for speaker diarization is gaining attention due to its simple architecture and promising results compared to the conventional cascaded systems [17, 18, 19, 20]. In this approach, diarization models are designed to estimate each speaker's speech activities from an input multi-speaker conversational recording. Thus, they require large-scale labeled conversational recordings for training, but unfortunately, the amount of such labeled conversational data is limited—at least compared to the single-speaker datasets [21, 22]. To deal with the problem of limited real data, past studies on end-to-end neural diarization (EEND) [23, 17, 18, 24, 19] have utilized models that are first pretrained using simulated conversational data created from single-speaker datasets and then adapted with real data. Such simulated conversational data can be generated with infinite variation given single-speaker recordings, and reports have shown that using the simulated data for training can improve performance compared to training with only a small amount of real data [23, 17].

Recent studies have shown that using both simulated and real recordings during pretraining improves the diarization performance [20], especially when Conformer encoders [25] are used as a backbone architecture. It is important that the turn-taking property be natural in this case because Conformer, unlike Transformer [26] [1], captures temporal context, and thus using real conversational data during pretraining helps. Here, a research question arises as to whether improving the protocol to simulate more natural conversational data will help improve the performance of EEND.

In this paper, we propose a method to simulate natural conversational mixtures to improve the performance of EEND models. We consider the turn-taking relationship between multi-speaker utterances by introducing four utterance transition types and mixing the utterances in a sequential manner. As our experimental results demonstrate, the proposed simulation method improves the conversational similarity between the simulated and real data, and significantly decreases the diarization error rates (DERs) of the Transformer- and Conformer-based EEND models.

2. Related work

2.1. End-to-end neural diarization

End-to-end neural diarization (EEND) is a framework to estimate multi-speaker speech activities from an input recording simultaneously. Given T-length D-dimensional frame-wise acoustic features $\left(\boldsymbol{x}_t \mid \boldsymbol{x}_t \in \mathbb{R}^D\right)_{t=1}^T$, the EEND model f_{EEND} directly estimates the joint speech activities of all S speakers for each frame in an end-to-end fashion as

$$\boldsymbol{p}_1, \ldots, \boldsymbol{p}_T = f_{\mathsf{EEND}}\left(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T\right), \tag{1}$$

where $\boldsymbol{p}_t \coloneqq [p_{t,1}, \ldots, p_{t,S}]^\mathsf{T} \in (0,1)^S$ is the posterior probabilities of S speakers' speech activities at t. The backbone ar-

*This work has been done during internship at Hitachi, Ltd.

[1]Transformer encoders without positional encoding are usually used for EEND.

chitecture of f_{EEND} can be bidirectional long short-term memory [23], Transformer [17], Conformer [20], or a time-dilated neural network [19]. During training, the posterior probabilities are optimized to minimize the following permutation-free loss:

$$\mathcal{L} = \frac{1}{TS} \min_{\phi \in \Phi(S)} \sum_{t=1}^{T} H\left(\boldsymbol{y}_t, P_\phi \boldsymbol{p}_t\right), \tag{2}$$

where ϕ where $\Phi(S)$ is the set of all possible permutation of S elements, $P_\phi \in \{0,1\}^{S \times S}$ is the permutation matrix corresponding to the permutation ϕ, $\boldsymbol{y}_t \coloneqq [y_{t,1}, \ldots, y_{t,S}]^\mathsf{T} \in \{0,1\}^S$ is the ground-truth speaker activities, and $H(\cdot,\cdot)$ is the binary cross entropy, defined as

$$H\left(\boldsymbol{y}_t, \boldsymbol{p}_t\right) \coloneqq \sum_{s=1}^{S} \left\{-y_{t,s} \log p_{t,s} - (1 - y_{t,s}) \log\left(1 - p_{t,s}\right)\right\}. \tag{3}$$

Since the training objective in (2) includes the ground-truth labels $\boldsymbol{y}_t$, the training of EEND requires data with frame-wise annotations, with the exception of one study that utilized unsupervised domain adaptation of EEND [27]. The amount of real data in the target domain is sometimes limited, so most of the previous studies on EEND have utilized simulated data [23, 17, 18, 24, 19], where the EEND models are first pretrained using large-scale simulated multi-speaker datasets and then adapted to the real data in the target domain. Most of the studies used the simulation method proposed in the initial EEND paper [23, 17] (see Section 3 for the detailed protocol), but there has been no investigation into the validity of the simulation method. A recent study [20] has shown that using both simulated and real data during pretraining helps improve the diarization performance. This improvement is greater when Conformer encoders are used as the backbone architecture, as they can capture temporal contexts more effectively than Transformer encoders[2]. This result suggests that the temporal context, or *turn-taking* behavior, of the data used for pretraining should be natural, especially when Conformer encoders are used, as was discussed in [20]. In our work, we aim to provide a novel method for simulating conversational speech with natural turn-taking.

2.2. Simulated mixtures in speech processing

The simulation of large-scale multi-speaker datasets is an important step in training data-hungry neural-network-based speech processing models, not just speaker diarization models. In the field of speech separation, WSJ0-2mix [28] is the first and still the most commonly used benchmark dataset. Each mixture in WSJ0-2mix consists of near-fully overlapped utterances of two speakers, each of which is from the Wall Street Journal (WSJ0) corpus [29]. The WHAM! [30] and WHAMR! [31] datasets are extensions of WSJ0-2mix to the noisy and reverberant conditions, respectively. Libri2Mix [32] has a similar overlap ratio to WSJ0-2mix, but the LibriSpeech dataset is used as the source dataset instead of the WSJ0 corpus. To evaluate speech separation performance under various overlap ratio conditions, LibriCSS [33] and SparseLibri2mix [32] have also been proposed, both of which are created from the LibriSpeech dataset. LibriCSS has an overlap ratio of 0–40 %, and SparseLibri2mix has an overlap ratio of 0–100 %.

[2]Most EEND studies have used Transformer encoders without positional encodings; thus, temporal context is completely ignored.

Simulated mixtures have also been used in ASR studies. The most commonly used type of simulated data consists of mixtures that each contain one utterance per speaker [34, 35, 36]. Such data is created by adding utterances of multiple speakers with random delays. One study evaluated the mixtures with various turn-taking configurations, rather than simply adding utterances with delays [37]. Other studies have used long-form simulated mixtures in which each speaker utters multiple times for training. In the research on diarization, while cascaded approaches only require single-speaker recordings to train a speaker embedding extractor, end-to-end approaches require multi-speaker recordings for the model training [38, 39, 40].

As described above, simulated mixtures have been widely utilized in speech processing research. However, the quality of the simulation, *i.e.*, how realistic the simulated multi-speaker conversation is, has rarely been investigated. In the case of speech separation and ASR, the evaluation data is also often simulation data, so the realism of the simulation data may not be as important. On the other hand, in the case of speaker diarization, where evaluation with long-form real recordings is common, it is important to know how realistic the simulation data is. In this paper, we clarify the difference in the performance of EEND depending on the simulation data used for training.

2.3. Conversation analysis

For better understanding and to establish a simulation protocol for speech diarization, we referred to research on conversation analysis. In the field of psychology, turn-taking is regarded as one of the fundamental mechanisms in conversation because conversation is a type of social interaction [41, 42]. Prior studies have analyzed the temporal patterning of turn-taking and categorized the relations between multi-speaker utterances into patterns such as gap, pause, and overlap [43] or continuation, interrupt, and turn change [44]. Inspired by their work, we take human turn-taking behaviors into account to create training data for EEND. On the basis of these earlier studies on categorization, we introduce four transition types between utterances, as described in Section 4.1 .

3. Conventional simulation method: Concat-and-sum approach

The mixture simulation protocol used in the previous EEND studies [23, 17] can be referred to as the *concat-and-sum* approach. In order to create an S-speaker mixture, a long recording $\boldsymbol{x}^{(s)}$ for each speaker $s \in \{1, \ldots, S\}$ is first prepared, where $\boldsymbol{x}^{(s)} = (x_1^{(s)}, \ldots, x_{T_s}^{(s)}) \in \mathbb{R}^{T_s}$ is the T_s-length discrete waveform of s-th speaker's speech, *i.e.*, $\left|\boldsymbol{x}^{(s)}\right| = T_s$. Given N_s utterances $\left\{\boldsymbol{u}_n^{(s)}\right\}_{n=1}^{N_s}$ of the s-th speaker, each of which is also the discrete waveform, they are concatenated with a silence between them as

$$\boldsymbol{x}^{(s)} = \boldsymbol{u}_1^{(s)} \oplus \mathbf{0}_1^{(s)} \oplus \boldsymbol{u}_2^{(s)} \oplus \mathbf{0}_2^{(s)} \oplus \cdots \oplus \mathbf{0}_{N_{s-1}}^{(s)} \oplus \boldsymbol{u}_{N_s}^{(s)}, \tag{4}$$

where $\oplus$ is the tuple concatenation operator and $\mathbf{0}_k^{(s)}$ for $1 \leq k \leq N_{s-1}$ is the all-zero tuple, whose length δ is sampled from the exponential distribution of the expected value β. Then, a random room impulse response $\boldsymbol{i}^{(s)}$ is convolved to each

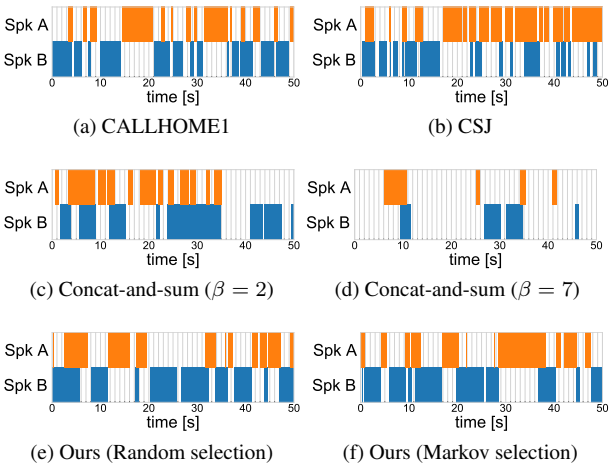

(a) CALLHOME1 (b) CSJ

(c) Concat-and-sum ($\beta = 2$) (d) Concat-and-sum ($\beta = 7$)

(e) Ours (Random selection) (f) Ours (Markov selection)

Figure 1: *Visualization of two-speaker conversations of real datasets (top), simulated datasets by the conventional method (middle), and simulated datasets by the proposed method (bottom).*

speaker's long recording, as

$$\boldsymbol{x}^{(s)} \leftarrow \boldsymbol{x}^{(s)} * \boldsymbol{i}^{(s)}. \tag{5}$$

Finally, all the long recordings and the L-length noise signal $\boldsymbol{x}_{\text{noise}}$ are concatenated to be an S-speaker mixture $\boldsymbol{x}$, as

$$\boldsymbol{x} = \boldsymbol{x}_{\text{noise}} + \sum_{s=1}^{S} \boldsymbol{x}^{(s)} \oplus \underbrace{(0, \ldots, 0)}_{L-T_s}, \tag{6}$$

$$L = \max_{s \in \{1,\ldots,S\}} T_s. \tag{7}$$

In the simulation protocol above, relations between utterances of different speakers are not considered. As a result, speakers appearing in the simulated mixtures do not speak alternately (as in a real conversation shown in Figs. 1(a) and 1(b)), and there are unnatural silence intervals and overlapping speech, as shown in Figs. 1(c) and 1(d).

Indeed, it is possible to control the statistics of the simulated mixtures, *e.g.*, the overlap ratio, by varying the value of β in the exponential distribution. However, reducing the overlap ratio by increasing the value of β also makes the duration of silence longer; thus, it is difficult to make the simulated mixture closer to a natural conversation by simply adjusting the value of β. Moreover, concatenating a zero vector in (6) to align the lengths of the long recordings causes unnaturalness of the conversation; various speakers speak at the beginning of the mixtures while only some of them speak at the end. In contrast, the proposed simulation protocol can generate more natural conversations by ordering each speaker's utterances to follow the statistics calculated from real conversational data.

4. Proposed simulation method

In the proposed simulation method, utterances of multiple speakers are sequentially arranged as in a real conversation. This makes it possible to avoid unnatural conversational patterns, silence, and overlapping speech, as shown in Fig. 1(e) and Fig. 1(f) . The following subsections show how the utterance transition types are defined and arranged to simulate a conversation.

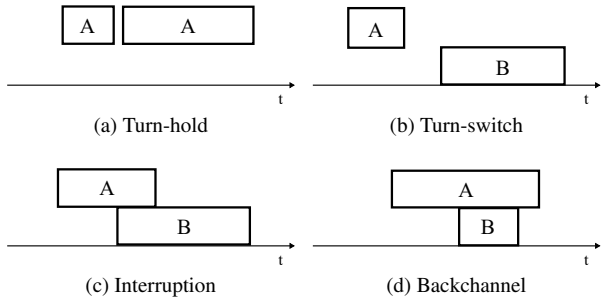

(a) Turn-hold (b) Turn-switch

(c) Interruption (d) Backchannel

Figure 2: *Four utterance transition cases in two-speaker scenario.*

4.1. Transition types between utterances

In the proposed simulation protocol, utterances are arranged in an autoregressive manner. One possible approach to this is to decide the next speaker based on the current speaker, as in VBx [14]. However, VBx only considers which speaker to transition to next, not how to transition. In other words, it does not take silence or overlap between utterances into account. Our method, in contrast, also takes into account how to transition to the next speaker by introducing variations determined between a pair of successive utterances.

Assuming that at most two speakers speak at the same time, there are four possible utterance transition cases. Let $\boldsymbol{u}_{\text{prev}}$ be the utterance with the latest end time among the utterances arranged so far, and $\boldsymbol{u}_{\text{next}}$ be the next utterance to be arranged. The four transition patterns of $\boldsymbol{u}_{\text{prev}}$ and $\boldsymbol{u}_{\text{next}}$ are shown in Fig. 2 and can be described as follows.

- **Turn-hold (TH)** (Fig. 2(a)): Two utterances from the same speaker with inter-turn silence (pause) between them. We assume that the length of pause follows the exponential distribution with the expected value β_{TH}.
- **Turn-switch (TS)** (Fig. 2(b)): Two utterances from different speakers with intra-turn silence (gap) between them. We assume that the length of pause follows the exponential distribution with the expected value β_{TS}.
- **Interruption (IR)** (Fig. 2(c)): Two utterances by different speakers that partially overlap each other.
- **Backchannel (BC)** (Fig. 2(d)): Two utterances by different speakers, one of which is fully overlapped with the other. This is also known as *aizuchi*.

For turn-hold and turn-switch, we assume that the duration of each silence δ between $\boldsymbol{u}_{\text{prev}}$ and $\boldsymbol{u}_{\text{next}}$ follows an exponential distribution, *i.e.*,

$$\delta \sim f(\delta; \beta) := \frac{1}{\beta} \exp\left(-\frac{\delta}{\beta}\right). \tag{8}$$

We use the different expected values for turn-hold and turn-switch, each of which are denoted as β_{TH} and β_{TS}, respectively.

For interruption and backchannel, on the other hand, the duration of overlap cannot be determined in the same way, since the randomly determined overlap duration may be larger than $|\boldsymbol{u}_{\text{prev}}|$ and $|\boldsymbol{u}_{\text{next}}|$. Therefore, we instead determine the ratio of the overlap ρ. Let $\boldsymbol{u}'_{\text{prev}}$ be the last part of $\boldsymbol{u}_{\text{prev}}$, which is not overlapped with all the utterances before $\boldsymbol{u}_{\text{next}}$; then, the overlap duration is calculated as

$$\delta = \rho \cdot \min\left(\left|\boldsymbol{u}'_{\text{prev}}\right|, \left|\boldsymbol{u}_{\text{next}}\right|\right), \tag{9}$$

where ρ is sampled from the truncated exponential distribution whose probability density function is defined as

$$\rho \sim g(\rho;\beta) := \begin{cases} \frac{f(\rho;\beta)}{\int_{\epsilon}^{1-\epsilon} f(\rho';\beta)d\rho'} & (\epsilon \leq \rho \leq 1-\epsilon) \\ 0 & \text{(otherwise)} \end{cases}, \quad (10)$$

where $\epsilon \ll 1$ is set to 0.03 in this paper. In this case too, we use the different values β for interruption and backchannel, each of which are denoted as β_{IR} and β_{BC}, respectively. Note that the starting time of $\boldsymbol{u}_{\text{next}}$ is uniformly determined in the case of backchannel.

After the arrangement of $\boldsymbol{u}_{\text{next}}$, $\boldsymbol{u}_{\text{prev}}$ is updated for the next step. In the case of turn-hold, turn-switch, and interruption, $\boldsymbol{u}_{\text{prev}}$ is updated as

$$\boldsymbol{u}_{\text{prev}} \leftarrow \boldsymbol{u}_{\text{next}}, \quad (11)$$

while it is not updated in the case of backchannel.

4.2. Sequential determination of transition types

In our mixture simulation protocol, multi-speaker utterances are mixed by sequentially determining the next speaker and transition type. We consider two types of transition type generation: *random selection* and *Markov selection*.

Random selection: The random selection assumes that the transition type at each transition is independent of the others, *i.e.*, the transition type is randomly selected according to the probability P_{ind}:

$$P_{\text{ind}} = [p_{\text{TH}}, p_{\text{TS}}, p_{\text{IR}}, p_{\text{BC}}], \quad (12)$$

where p_{X} for $X \in \{\text{TH}, \text{TS}, \text{IR}, \text{BC}\}$ is the probability of occurrence of each transition type.

Markov selection: The Markov selection assumes that the transition types follow a simple Markov process, *i.e.*, the probability of each transition type depends on the previous transition type. The transition matrix $P \in [0,1]^{4\times 4}$ is denoted as

$$P_{\text{Markov}} = \begin{bmatrix} p_{\text{TH}\to\text{TH}} & p_{\text{TS}\to\text{TH}} & p_{\text{IR}\to\text{TH}} & p_{\text{BC}\to\text{TH}} \\ p_{\text{TH}\to\text{TS}} & p_{\text{TS}\to\text{TS}} & p_{\text{IR}\to\text{TS}} & p_{\text{BC}\to\text{TS}} \\ p_{\text{TH}\to\text{IR}} & p_{\text{TS}\to\text{IR}} & p_{\text{IR}\to\text{IR}} & p_{\text{BC}\to\text{IR}} \\ p_{\text{TH}\to\text{BC}} & p_{\text{TS}\to\text{BC}} & p_{\text{IR}\to\text{BC}} & p_{\text{BC}\to\text{BC}} \end{bmatrix}, \quad (13)$$

where $p_{X\to Y}$ for $X, Y \in \{\text{TH}, \text{TS}, \text{IR}, \text{BC}\}$ is the probability that the next state will be Y given the current state X.

The parameters P_{ind} and P_{Markov} decide the naturalness of the simulated mixtures. In this paper, we use these parameters extracted from the real conversational data.

4.3. Overall simulation protocol

The algorithm of the proposed simulation protocol is presented in Algorithm 1 . As inputs, set of speakers $\mathcal{U}$, noises $\mathcal{N}$, room impulse responses (RIRs) $\mathcal{I}$, signal-to-noise ratios $\mathcal{R}$, and utterances for each speaker $\mathcal{U}_1, \ldots, \mathcal{U}_{|\mathcal{S}|}$ are given. The numbers of speakers N_{spk} and utterances N_{utt} within a mixture are also given as inputs.

First, the set of N_{spk} speakers and corresponding RIRs are sampled (L1–L2). The speakers and their utterances are determined sequentially and arranged on the basis of the utterance transition types described in Section 4.1 and Section 4.2 to construct a long-form audio (L3–L29). Finally, a noise is added with a randomly determined SNR to obtain a final simulated mixture (L30–34).

Algorithm 1: Mixture simulation.

Input: $\mathcal{S}, \mathcal{N}, \mathcal{I}, \mathcal{R}$ // Set of speakers, noises, RIRs, and SNRs
$\mathcal{U}_1, \ldots, \mathcal{U}_{|\mathcal{S}|}$ // Set of utterances for each speaker
$N_{\text{spk}}, N_{\text{utt}}$ // No. of speakers / no. of utterances per mixture
Output: $\boldsymbol{x}$ // Mixture

1 Sample a set of N_{spk} speakers $\mathcal{S}'$ from $\mathcal{S}$
2 Sample $\boldsymbol{i}^{(s)}$ for $s \in \mathcal{S}'$ from $\mathcal{I}$ // RIR
3 Sample the first speaker s_1 from $\mathcal{S}'$
4 Initialize $\boldsymbol{x}$ with $\boldsymbol{u} * \boldsymbol{i}^{(s_1)}$, where $\boldsymbol{u} \sim U_{s_1}$
5 **for** $t \leftarrow 2$ **to** N_{utt} **do**
6 Decide the next utterance transition type Z_t
7 **if** $Z_t = \text{TH}$ **then** // Turn-hold
8 $s_t = s_{t-1}$
9 Sample $\boldsymbol{u}$ from $\mathcal{U}_{s_t}$
10 Prepare a silence signal $\mathbf{0}_t$ using (8)
11 $\boldsymbol{x} \leftarrow \boldsymbol{x} \oplus \mathbf{0}_t \oplus \boldsymbol{u} * \boldsymbol{i}^{(s_t)}$
12 **else if** $Z_t = \text{TS}$ **then** // Turn-switch
13 Sample s_t from $\mathcal{S}' \setminus \{s_{t-1}\}$
14 Sample $\boldsymbol{u}$ from $\mathcal{U}_{s_t}$
15 Prepare a silence signal $\mathbf{0}_t$ using (8)
16 $\boldsymbol{x} \leftarrow \boldsymbol{x} \oplus \mathbf{0}_t \oplus \boldsymbol{u} * \boldsymbol{i}^{(s_t)}$
17 **else if** $Z_t = \text{IR}$ **then** // Interruption
18 Sample s_t from $\mathcal{S}' \setminus \{s_{t-1}\}$
19 Sample $\boldsymbol{u}$ from $\mathcal{U}_{s_t}$
20 $\boldsymbol{u} \leftarrow \boldsymbol{u} * \boldsymbol{i}^{(s_t)}$
21 Decide the overlap duration δ using (9) (10)
22 $\boldsymbol{x} \leftarrow \boldsymbol{x}_{1:|\boldsymbol{x}|-\delta} \oplus \left(\boldsymbol{x}_{|\boldsymbol{x}|-\delta:|\boldsymbol{x}|} + \boldsymbol{u}_{1:\delta}\right) \oplus \boldsymbol{u}_{\delta:|\boldsymbol{u}|}$
23 **else if** $Z_t = \text{BC}$ **then** // Backchannel
24 Sample s_t from $\mathcal{S}' \setminus \{s_{t-1}\}$
25 Sample $\boldsymbol{u}$ from $\mathcal{U}_{s_t}$
26 $\boldsymbol{u} \leftarrow \boldsymbol{u} * \boldsymbol{i}^{(s_t)}$
27 Decide the overlap duration δ using (9) (10)
28 Decide the start time τ of $\boldsymbol{u}$
29 $\boldsymbol{x} \leftarrow \boldsymbol{x}_{1:\tau} \oplus \left(\boldsymbol{x}_{\tau+|\boldsymbol{u}|} + \boldsymbol{u}\right) \oplus \boldsymbol{x}_{\tau+|\boldsymbol{u}|:|\boldsymbol{x}|}$
30 Sample $\boldsymbol{n}$ from $\mathcal{N}$ // Background noise
31 Sample r from $\mathcal{R}$ // SNR
32 Determine a mixing scale p from $r, \boldsymbol{x}$, and $\boldsymbol{n}$
33 $\boldsymbol{n}' \leftarrow$ repeat $\boldsymbol{n}$ until reaching the length of $\boldsymbol{x}$
34 $\boldsymbol{x} \leftarrow \boldsymbol{x} + p \cdot \boldsymbol{n}'$

5. Experiments

5.1. Data

5.1.1. Real datasets

We utilized the two-speaker portion of the CALLHOME dataset [45] for the experiment. Following the setting in previous studies [17, 18], we used the split provided in the Kaldi recipe[3] to create an adaptation set consisting of 155 recordings (CALLHOME1) and a test set consisting of 148 recordings (CALLHOME2). We also used 54 dialogue recordings in the Corpus of Spontaneous Japanese (CSJ) for evaluation.

[3] https://github.com/kaldi-asr/kaldi/tree/master/egs/callhome_diarization/v2

Table 1: Parameters extracted from CALLHOME1.

Parameter	Value
$[\beta_{\text{TH}}, \beta_{\text{TS}}, \beta_{\text{IR}}, \beta_{\text{BC}}]$	$[0.57, 0.40, 0.10, 0.44]$
P_{ind}	$[0.15, 0.31, 0.44, 0.10]$
P_{Markov}	$\begin{bmatrix} 0.26 & 0.11 & 0.09 & 0.31 \\ 0.23 & 0.38 & 0.29 & 0.29 \\ 0.27 & 0.45 & 0.53 & 0.31 \\ 0.24 & 0.06 & 0.09 & 0.09 \end{bmatrix}$

5.1.2. Simulated mixtures

The sets of utterances used for the simulation were extracted from the Switchboard-2 (Phase I, II, III), Switchboard Cellular (Part1, Part2), and NIST Speaker Recognition Evaluation (2004, 2005, 2006, 2008) datasets. All recordings are telephone speech sampled at 8 kHz. There are 6,381 speakers in total, and we used 5,743 speakers among them to create the training sets following previous studies [23, 17, 18, 24, 20]. We extracted utterances from these datasets by using the time-delay-neural-network-based speech activity detector[4]. While most previous studies filtered out detected utterances of less than 1.5 s, one study reported that utilizing these short utterances helps improve diarization performance. Therefore, in this paper we examined using utterances only 1.5 s or longer and using all the detected utterances.

With the set of utterances, we created two-speaker simulated conversational datasets. For the conventional simulation method, we used $\beta = 2$ following previous studies [23, 17, 18, 24]. We also used $\beta = 7$ to obtain a similar overlap ratio to CALLHOME1. For the proposed simulation method, we determine the parameters based on those extracted from CALLHOME1, as shown in Table 1 . Each simulated dataset has 100,000 mixtures and the number of utterances within each mixture was adjusted so that the total duration of each dataset was roughly the same (see Table 2).

The set of background noises consists of 37 recordings annotated as background noises in the MUSAN corpus [46]. The Simulated Room Impulse Response Database [47] was used as the set of RIRs for all the simulation methods. The SNR values were sampled from $\{5, 10, 15, 20\}$ dB.

5.2. Models

We used attractor-based EEND (EEND-EDA) [18, 24] for the experiments. As the backbone architecture, we used four-stacked Transformer encoders [26] or Conformer encoders [25]. For both models, the number of attention units was set to 256, and that of attention heads to four. The dimensionality of each position-wise feed-forward layer in the Transformer and Conformer encoders was set to 1,024 and 256, respectively. Note that the purpose of this research is to investigate the optimal mixture simulation method; thus, we do not utilize updates introduced in the original Conformer-based EEND paper [20] such as SpecAugment [48] or convolutional subsampling.

Each model was pretrained using simulated data for 100 epochs using the Adam optimizer [49] with a Noam scheduler [26] and then adapted to CALLHOME1 for another 100 epochs using Adam with a fixed learning rate of 1×10^{-5}. Diarization performance was evaluated using the CALLHOME2 and CSJ datasets.

[4] https://github.com/kaldi-asr/kaldi/tree/master/egs/aspire/s5

5.3. Metrics

We used the standard silence and overlap ratio as metrics to investigate the statistics of simulated and real conversation data, each of which are defined as

$$\text{SilenceRatio} = \frac{T_{\#Spk=0}}{T_{\#Spk=0} + T_{\#Spk\geq 1}}, \tag{14}$$

$$\text{OverlapRatio} = \frac{T_{\#Spk\geq 2}}{T_{\#Spk\geq 1}}, \tag{15}$$

where $T_{\#Spk=0}$, $T_{\#Spk\geq 1}$, and $T_{\#Spk\geq 2}$ denote the total duration where no speaker is active, at least one speaker is active, and multiple speakers are active, respectively.

We also utilized the similarity metric based on the earth movers' distance (EMD) [20] to evaluate the distribution of the duration of overlap and silence. Let U and V be the cumulative distribution functions of the overlap (or silence) duration of the datasets to be compared. The EMD between U and V is determined as

$$\text{EMD}\left(U, V\right) = \frac{\sum_{i=1}^{m} \sum_{j=1}^{n} f_{ij} d_{ij}}{\sum_{i=1}^{m} f_{ij}}, \tag{16}$$

where f_{ij} is the optimal flow that minimize the overall moving cost, d_{ij} is the L_1 distance between the i-th element of U and the j-th element of V, and m and n are the number of elements in U and V, respectively. The silence similarity or overlap similarity is determined as

$$\text{sim}\left(U, V\right) = \exp\left(-\gamma \text{EMD}\left(U, V\right)\right), \tag{17}$$

where γ is the scaling parameter, which is set to 0.001 in this paper.

To evaluate the diarization performance, we used the diarization error rate (DER). Following the previous studies [23, 17, 18, 24], we allowed 0.25 s of collar at the boundary of each speech segment.

6. Results

6.1. Statistics of real and simulated data

The statistics of the real and simulated conversation data are shown in Table 2 . We report the silence and overlap ratios in addition to the silence and overlap similarities of each dataset against the CALLHOME1, CALLHOME2, and CSJ datasets. First, the silence similarity and overlap similarity between the three real datasets are both high (> 0.9), while the silence and overlap ratios are bit far apart between CALLHOME1/2 and CSJ. This suggests that these similarities are more important in evaluating the naturalness of conversational datasets than the silence or overlap ratios themselves.

With the conventional concat-and-sum approach, the silence and overlap similarities are low (< 0.5) when $\beta = 2$. $\beta = 7$ brought the overlap ratio closer to the real datasets, but resulted in only a slight improvement in overlap similarity with severe degradation in silence similarity. Using the utterances of less than 1.5 s improves the overlap similarities to about 0.6, but they are still not as good as those between the real datasets.

With the proposed simulation method using the parameters extracted from the CALLHOME1 dataset, we obtained about 0.95 silence similarity and 0.86 overlap similarity with CALLHOME1. Using the utterances of less than 1.5 s did not improve

Table 2: *Data statistics of silence and overlap ratio, and silence/overlap similarities between real and simulated datasets. Parameters used for the proposed simulation method were extracted from CALLHOME1.*

					CALLHOME1		CALLHOME2		CSJ	
Dataset	Utterance length	Total duration	Silence ratio	Overlap ratio	Silence similarity	Overlap similarity	Silence similarity	Overlap similarity	Silence similarity	Overlap similarity
CALLHOME1	>0 s	3 h	0.090	0.141	1.000	1.000	0.968	0.962	0.956	0.930
CALLHOME2	>0 s	3 h	0.098	0.131	0.968	0.962	1.000	1.000	0.951	0.939
CSJ	>0 s	12 h	0.168	0.201	0.956	0.930	0.951	0.939	1.000	1.000
Concat-and-sum ($\beta = 2$)	>1.5 s	2461 h	0.212	0.341	0.425	0.363	0.416	0.350	0.428	0.355
Concat-and-sum ($\beta = 7$)	>1.5 s	2428 h	0.550	0.137	0.018	0.369	0.018	0.357	0.018	0.362
Ours (Random selection)	>1.5 s	2459 h	0.075	0.141	0.954	0.859	0.964	0.828	0.936	0.822
Ours (Markov selection)	>1.5 s	2473 h	0.074	0.144	0.954	0.861	0.966	0.829	0.937	0.824
Concat-and-sum ($\beta = 2$)	>0 s	2401 h	0.305	0.262	0.437	0.611	0.428	0.590	0.440	0.599
Concat-and-sum ($\beta = 7$)	>0 s	2393 h	0.661	0.093	0.019	0.630	0.018	0.609	0.019	0.618
Ours (Random selection)	>0 s	2452 h	0.107	0.119	0.954	0.862	0.965	0.891	0.936	0.865
Ours (Markov selection)	>0 s	2458 h	0.106	0.121	0.954	0.861	0.966	0.890	0.937	0.864

the similarities with CALLHOME1, but it did improve the overlap similarities with CALLHOME2 and CSJ. Note that the type of utterance transition generation (random selection or Markov selection) did not affect the similarities very much. The reason is that silence or overlap similarity only consider the distribution of their duration, not their order.

6.2. Diarization performance

Table 3 shows the DERs of the Transformer-based and Conformer-based EEND-EDA on the real datasets: CALLHOME2 and CSJ. Each model was evaluated with and without domain adaptation on CALLHOME1. We observed that Conformer-based EEND-EDA trained on the dataset simulated with >0 s utterances using the proposed method with Markov selection performed best, regardless of the adaptation.

In terms of the utterance length, Transformer-based EEND-EDA performed better than Conformer-based EEND-EDA when only the utterances longer than 1.5 s were used for simulation. This indicates that Conformer, which takes temporal context into account, could only recognize speech segments lasting more than 1.5 s if only the utterances longer than 1.5 s were used for simulation. In contrast, Conformer-based EEND-EDA was better than Transformer-based EEND-EDA when all the utterances were used for simulation, and in most cases the DERs were better than the cases where the utterances longer than 1.5 s were used. From here, we focus on the results of Conformer-based EEND-EDA trained using the simulation datasets created with all the utterances (>0 s).

In terms of the simulation methods, the datasets that were simulated with the proposed method contributed to improving the DERs. Comparing the proposed method using random selection and using Markov selection, Markov selection was always better due to the improvement of the overlap similarity. These results clearly show that the EEND performance can be improved by making the simulated mixtures more like real conversations, without using real data for pretraining as in [20].

Finally, we discuss the differences in EEND-EDA with and without domain adaptation. Adjusting the value of β of the concat-and-sum approach to obtain an overlap ratio similar to the real datasets helped improve DERs without adaptation ($\beta = 2$ vs. $\beta = 7$). However, the DERs after adaptation were rather worse in the case of $\beta = 7$ than in the case of $\beta = 2$. We presume this depends on whether the model overfits to the adaptation data (CALLHOME1) or generalizes well to the real data. In Table 2 , we can see that concat-and-sum with $\beta = 7$ is closer to the real datasets than that with $\beta = 2$ and thus was overfitted to CALLHOME1 during the adaptation phase; therefore, the DER on CALLHOME2, which has the same domain as CALLHOME1, was improved from $10.70\,\%$ to $8.93\,\%$, but the DER on CSJ, which has a different domain, was degraded from $19.37\,\%$ to $21.15\,\%$. These results indicate that the improvement of the similarities between simulated and real datasets is the key to improving the DER on CSJ. The datasets that are simulated using the proposed method, which are more similar to the real datasets, indeed contributed to improving the DERs, especially on CSJ: from $17.52\,\%$ to $14.50\,\%$ with random selection and from $17.09\,\%$ to $12.41\,\%$ with Markov selection.

7. Conclusion

In this paper, we proposed a method for simulating conversations that contain natural turn-taking. We introduced four utterance transition types and sequentially arranged utterances in accordance with these types to simulate a long-form conversation. The datasets that were simulated using the proposed method showed higher similarity with the real dataset than those simulated using a conventional method. The experimental results showed that EEND-EDA pretrained using the datasets simulated by the proposed method contributed to improving DERs on both CALLHOME and CSJ.

We demonstrated the effectiveness of the proposed method under the two-speaker condition, but since Algorithm 1 is described in a way that does not limit the number of speakers, our future work will include evaluations in cases of more than two speakers. In addition, since the diarization performance was improved by using Markov selection instead of random selection, further improvement can be expected by determining the utterance transition type according to a higher-order Markov chain. Finally, the similarity metric used in this study does not capture the difference between the datasets created using random selection and Markov selection. Constructing a new similarity metric that can express this difference is thus a possible avenue for future work.

Table 3: *DERs (%) on CALLHOME2 and CSJ.*

		CALLHOME2				CSJ			
		Transformer		Conformer		Transformer		Conformer	
Simulation method	Utt. length	w/o adapt	w/ adapt	w/o adapt	w/ adapt	w/o adapt	w/ adapt	w/o adapt	w/ adapt
Concat-and-sum ($\beta = 2$)	>1.5 s	14.46	7.90	18.03	8.36	19.30	16.98	20.42	15.95
Concat-and-sum ($\beta = 7$)	>1.5 s	10.64	8.66	12.21	10.02	18.95	17.09	19.79	22.79
Ours (Random selection)	>1.5 s	10.35	7.63	12.40	7.94	18.85	18.58	21.01	17.43
Ours (Markov selection)	>1.5 s	10.32	7.74	11.97	7.28	**16.99**	**14.36**	20.30	16.23
Concat-and-sum ($\beta = 2$)	>0 s	12.26	7.80	11.25	7.53	19.41	17.00	18.86	16.55
Concat-and-sum ($\beta = 7$)	>0 s	10.11	8.61	10.70	8.93	19.51	17.92	19.37	21.15
Ours (Random selection)	>0 s	10.39	7.83	9.74	7.32	17.96	17.71	17.52	14.50
Ours (Markov selection)	>0 s	**9.65**	**7.53**	**9.65**	**7.18**	17.85	15.26	**17.09**	**12.41**

8. References

[1] Tae Jin Park, Naoyuki Kanda, Dimitrios Dimitriadis, Kyu J Han, Shinji Watanabe, and Shrikanth Narayanan, "A review of speaker diarization: Recent advances with deep learning," *Computer Speech & Language*, vol. 72, no. 7, pp. 101317, 2022.

[2] Takuya Yoshioka, Dimitrios Dimitriadis, Andreas Stolcke, William Hinthorn, Zhuo Chen, Michael Zeng, and Xuedong Huang, "Meeting transcription using asynchronous distant microphones," in *INTERSPEECH*, 2019, pp. 2968–2972.

[3] Takuya Yoshioka, Igor Abramovski, Cem Aksoylar, Zhuo Chen, Moshe David, Dimitrios Dimitriadis, Yifan Gong, Ilya Gurvich, Xuedong Huang, Yan Huang, et al., "Advances in online audio-visual meeting transcription," in *ASRU*, 2019, pp. 276–283.

[4] Manoj Kumar, So Hyun Kim, Catherine Lord, and Shrikanth Narayanan, "Improving speaker diarization for naturalistic child-adult conversational interactions using contextual information," *The Journal of the Acoustical Society of America*, vol. 147, no. 2, pp. EL196–EL200, 2020.

[5] Allen Guo, Arlo Faria, and Korbinian Riedhammer, "Remeeting-deep insights to conversations.," in *INTERSPEECH*, 2016, pp. 1964–1965.

[6] Angus Addlesee, Yanchao Yu, and Arash Eshghi, "A comprehensive evaluation of incremental speech recognition and diarization for conversational AI," in *COLING*, 2020, pp. 3492–3503.

[7] Naoyuki Kanda, Shota Horiguchi, Yusuke Fujita, Yawen Xue, Kenji Nagamatsu, and Shinji Watanabe, "Simultaneous speech recognition and speaker diarization for monaural dialogue recordings with target-speaker acoustic models," in *ASRU*, 2019, pp. 31–38.

[8] Stephen H Shum, Najim Dehak, Réda Dehak, and James R Glass, "Unsupervised methods for speaker diarization: An integrated and iterative approach," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 10, pp. 2015–2028, 2013.

[9] Gregory Sell and Daniel Garcia-Romero, "Speaker diarization with plda i-vector scoring and unsupervised calibration," in *SLT*, 2014, pp. 413–417.

[10] Daniel Garcia-Romero, David Snyder, Gregory Sell, Daniel Povey, and Alan McCree, "Speaker diarization using deep neural network embeddings," in *ICASSP*, 2017, pp. 4930–4934.

[11] Aonan Zhang, Quan Wang, Zhenyao Zhu, John Paisley, and Chong Wang, "Fully supervised speaker diarization," in *ICASSP*, 2019, pp. 6301–6305.

[12] Tae Jin Park, Kyu J. Han, Manoj Kumar, and Shrikanth Narayanan, "Auto-tuning spectral clustering for speaker diarization using normalized maximum eigengap," *IEEE Signal Processing Letters*, vol. 27, pp. 381–385, 2020.

[13] Qiujia Li, Florian L. Kreyssig, Chao Zhang, and Philip C. Woodland, "Discriminative neural clustering for speaker diarisation," in *SLT*, 2021.

[14] Federico Landini, Ján Profant, Mireia Diez, and Lukáš Burget, "Bayesian HMM clustering of x-vector sequences (VBx) in speaker diarization: theory, implementation and analysis on standard tasks," *Computer Speech & Language*, vol. 71, pp. 101254, 2022.

[15] Arsha Nagrani, Joon Son Chung, Weidi Xie, and Andrew Zisserman, "VoxCeleb: Large-scale speaker verification in the wild," *Computer Speech & Language*, vol. 60, pp. 101027, 2020.

[16] Mitchell McLaren, Luciana Ferrer, Diego Castan, and Aaron Lawson, "The 2016 Speakers in the Wild Speaker Recognition Evaluation," in *INTERSPEECH*, 2016, pp. 823–827.

[17] Yusuke Fujita, Naoyuki Kanda, Shota Horiguchi, Yawen Xue, Kenji Nagamatsu, and Shinji Watanabe, "End-to-end neural speaker diarization with self-attention," in *ASRU*, 2019, pp. 296–303.

[18] Shota Horiguchi, Yusuke Fujita, Shinji Watanabe, Yawen Xue, and Kenji Nagamatsu, "End-to-end speaker diarization for an unknown number of speakers with encoder-decoder based attractors," in *INTERSPEECH*, 2020, pp. 269–273.

[19] Soumi Maiti, Hakan Erdogan, Kevin Wilson, Scott Wisdom, Shinji Watanabe, and John R. Hershey, "End-to-end diarization for variable number of speakers with local-global networks and discriminative speaker embeddings," in *ICASSP*, 2021, pp. 7183–7187.

[20] Yi Chieh Liu, Eunjung Han, Chul Lee, and Andreas Stolcke, "End-to-end neural diarization: From transformer to conformer," in *INTERSPEECH*, 2021, pp. 3081–3085.

[21] Jean Carletta, "Unleashing the killer corpus: experiences in creating the multi-everything AMI Meeting Corpus," *Language Resources and Evaluation*, vol. 41, no. 2, pp. 181–190, 2007.

[22] Jon Barker, Shinji Watanabe, Emmanuel Vincent, and Jan Trmal, "The fifth 'CHiME' Speech Separation and Recognition Challenge: dataset, task and baselines," in *INTERSPEECH*, 2018, pp. 1561–1565.

[23] Yusuke Fujita, Naoyuki Kanda, Shota Horiguchi, Kenji Nagamatsu, and Shinji Watanabe, "End-to-end neural speaker diarization with permutation-free objectives," in *INTERSPEECH*, 2019, pp. 4300–4304.

[24] Shota Horiguchi, Yusuke Fujita, Shinji Watanabe, Yawen Xue, and Paola Garcia, "Encoder-decoder based attractors for end-to-end neural diarization," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2022.

[25] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, et al., "Conformer: Convolution-augmented transformer for speech recognition," in *INTERSPEECH*, 2020, pp. 5036–5040.

[26] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin, "Attention is all you need," in *NeurIPS*, 2017, pp. 5998–6008.

[27] Yuki Takashima, Yusuke Fujita, Shota Horiguchi, Shinji Watanabe, Paola Garcia, and Kenji Nagamatsu, "Semi-supervised training with pseudo-labeling for end-to-end neural diarization," in *INTERSPEECH*, 2021, pp. 3096–3110.

[28] John R Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *ICASSP*, 2016, pp. 31–35.

[29] Douglas B Paul and Janet Baker, "The design for the Wall Street Journal-based CSR corpus," in *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York, February 23-26, 1992*, 1992.

[30] Gordon Wichern, Joe Antognini, Michael Flynn, Licheng Richard Zhu, Emmett McQuinn, Dwight Crow, Ethan Manilow, and Jonathan Le Roux, "WHAM!: Extending speech separation to noisy environments," in *INTERSPEECH*, 2019, pp. 1368–1372.

[31] Matthew Maciejewski, Gordon Wichern, and Jonathan Le Roux, "WHAMR!: Noisy and reverberant single-channel speech separation," in *ICASSP*, 2020, pp. 696–700.

[32] Joris Cosentino, Manuel Pariente, Samuele Cornell, Antoine Deleforge, and Emmanuel Vincent, "LibriMix: An open-source dataset for generalizable speech separation," arXiv:2005.11262, 2020.

[33] Zhuo Chen, Takuya Yoshioka, Liang Lu, Tianyan Zhou, Zhong Meng, Yi Luo, Jian Wu, Xiong Xiao, and Jinyu Li, "Continuous speech separation: Dataset and analysis," in *ICASSP*, 2020, pp. 7284–7288.

[34] Hiroshi Seki, Takaaki Hori, Shinji Watanabe, Jonathan Le Roux, and John R Hershey, "A purely end-to-end system for multi-speaker speech recognition," in *ACL*, 2018, pp. 2620–2630.

[35] 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.

[36] Anshuman Tripathi, Han Lu, and Hasim Sak, "End-to-end multi-talker overlapping speech recognition," in *ICASSP*, 2020, pp. 6129–6133.

[37] Takuya Yoshioka, Hakan Erdogan, Zhuo Chen, and Fil Alleva, "Multi-microphone neural speech separation for far-field multi-talker speech recognition," in *ICASSP*, 2018, pp. 5739–5743.

[38] Thilo von Neumann, Keisuke Kinoshita, Marc Delcroix, Shoko Araki, Tomohiro Nakatani, and Reinhold Haeb-Umback, "All-neural online source separation, counting, and diarization for meeting analysis," in *ICASSP*, 2019, pp. 91–95.

[39] Keisuke Kinoshita, Marc Delcroix, Shoko Araki, and Tomohiro Nakatani, "Tackling real noisy reverberant meetings with all-neural source separation, counting, and diarization system," in *ICASSP*, 2020, pp. 381–385.

[40] Xuankai Chang, Naoyuki Kanda, Yashesh Gaur, Xiaofei Wang, Zhong Meng, and Takuya Yoshioka, "Hypothesis stitcher for end-to-end speaker-attributed asr on long-form multi-talker recordings," in *ICASSP*, 2021, pp. 6763–6767.

[41] Stephen C Levinson and Francisco Torreira, "Timing in turn-taking and its implications for processing models of language," *Frontiers in psychology*, vol. 6, pp. 731, 2015.

[42] Harvey Sacks, Emanuel A. Schegloff, and Gail Jefferson, "A simplest systematics for the organization of turn-taking for conversation," *Language*, vol. 50, no. 4, pp. 696–735, 1974.

[43] Mattias Heldner and Jens Edlund, "Pauses, gaps and overlaps in conversations," *Journal of Phonetics*, vol. 38, no. 4, pp. 555–568, 2010.

[44] Louis Ten Bosch, Nelleke Oostdijk, and Jan Peter De Ruiter, "Turn-taking in social talk dialogues: temporal, formal and functional aspects," in *SPECOM*, 2004, pp. 454–461.

[45] "2000 NIST Speaker Recognition Evaluation," `https://catalog.ldc.upenn.edu/LDC2001S97`.

[46] David Snyder, Guoguo Chen, and Daniel Povey, "MUSAN: A music, speech, and noise corpus," arXiv:1510.08484, 2015.

[47] Tom Ko, Vijayaditya Peddinti, Daniel Povey, Michael L Seltzer, and Sanjeev Khudanpur, "A study on data augmentation of reverberant speech for robust speech recognition," in *ICASSP*, 2017, pp. 5220–5224.

[48] Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," in *INTERSPEECH*, 2019, pp. 2613–2617.

[49] Diederik P. Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2015.