

SIMULATING REALISTIC SPEECH OVERLAPS IMPROVES MULTI-TALKER ASR

Muqiao Yang^{1†}, Naoyuki Kanda², Xiaofei Wang², Jian Wu², Sunit Sivasankaran², Zhuo Chen², Jinyu Li², Takuya Yoshioka²

¹ Carnegie Mellon University, Pittsburgh, PA, USA

² Microsoft, Redmond, WA, USA

ABSTRACT

Multi-talker automatic speech recognition (ASR) has been studied to generate transcriptions of natural conversation including overlapping speech of multiple speakers. Due to the difficulty in acquiring real conversation data with high-quality human transcriptions, a naïve simulation of multi-talker speech by randomly mixing multiple utterances was conventionally used for model training. In this work, we propose an improved technique to simulate multi-talker overlapping speech with realistic speech overlaps, where an arbitrary pattern of speech overlaps is represented by a sequence of discrete tokens. With this representation, speech overlapping patterns can be learned from real conversations based on a statistical language model, such as N-gram, which can be then used to generate multi-talker speech for training. In our experiments, multi-talker ASR models trained with the proposed method show consistent improvement on the word error rates across multiple datasets.

Index Terms— Multi-talker automatic speech recognition, conversation analysis, data simulation

1. INTRODUCTION

Automatic speech recognition (ASR) plays a crucial role in human-machine interactions and human conversation analyses. Even though there has been a significant progress of ASR technology in recent decades [1–3], it is still challenging to transcribe natural conversation because of the complicated acoustic and linguistic properties. Especially, natural conversation contains a considerable amount of speech overlaps [4], which significantly hurt the accuracy of conventional ASR systems designed for single-talker speech [5, 6]. To overcome the limitation, multi-talker ASR that generates transcriptions of multiple speakers has been studied [7–16]. To train such multi-talker ASR models in high quality, it is crucial to feed the model a sufficient amount of multi-talker speech samples with accurate transcriptions.

Most prior works used the simulated multi-talker speech generated from single-talker ASR training data [7–16] since it is expensive to collect a large number of real conversations with high quality transcriptions. For example, Seki et al. [8] mixed two single-talker audio samples where a shorter sample is mixed with a longer one with a random delay such that the shorter sample is fully overlapped with the longer sample. Kanda et al. [12] mixed S single-talker samples with random delays to simulate partially overlapping speech, where S was uniformly sampled from one to five. While there were minor differences in each work, all prior works naively mixed single-talker speech samples with random delays, which incurs unnatural speech overlapping pattern in the training data. On the other hand, it was suggested that a large portion of the degradation of word error rates

(WER) was caused by the insufficient training of speech overlapping patterns [17]. Recently, a few trials were made to simulate more realistic multi-talker audio in the context of speaker diarization. Landini et al. [18] used statistics about frequencies and lengths of pauses and overlaps. Yamashita et al. [19] used the transition probabilities between different overlap patterns. Although these works demonstrated superior performance for speaker diarization accuracy, such an approach has not been studied for multi-talker ASR. In addition, these methods are not directly applicable because they are assuming speaker information available for every speech sample while the ASR training data is often anonymized by excluding speaker identifiable information [17].

In this work, we propose a novel multi-talker speech data simulation method for multi-talker ASR training. We first represent an arbitrary pattern of speech overlaps by a sequence of discrete tokens. With this representation, speech overlapping patterns can be learned from real conversations based on a statistical language model (SLM), such as N-gram. The learned overlapping pattern can be then used to simulate multi-talker speech with realistic speech overlaps. We demonstrate the accuracy of multi-talker ASR can be consistently improved based on the proposed multi-talker simulation data across multiple real meeting evaluation sets. We also show that the accuracy is further improved by combining different types of simulation data as well as a small amount of real conversations.

2. RELATED WORKS

2.1. Token-level serialized output training for multi-talker ASR

The token-level serialized output training (t-SOT) was proposed to generate transcriptions of multi-talker overlapping speech in a streaming fashion [15]. In the t-SOT framework, the maximum number of active speakers at the same time is assumed to be M . In our work, we used the t-SOT with $M = 2$ for its simplicity, thus explaining the t-SOT with this condition. With t-SOT, the transcriptions of multiple speakers are serialized into a single sequence of recognition tokens (e.g., words, subwords, etc.) by sorting them in chronological order. A special token $\langle cc \rangle$, indicating a change of virtual output channels, is inserted between two adjacent words spoken by different speakers. A streaming end-to-end ASR model [20] is trained based on pairs of such serialized transcriptions and the corresponding audio samples. During inference, an output sequence including $\langle cc \rangle$ is generated by the ASR model in a streaming fashion, which is then converted to separate transcriptions based on the recognition of the $\langle cc \rangle$. The t-SOT model showed superior performance compared to prior multi-talker ASR models even with streaming processing. See [15] for more details.

2.2. Multi-talker speech simulation based on random mixtures

Prior multi-talker ASR studies used a simple simulation technique to generate training data by mixing single-talker speech signals with

[†]Work performed during internship at Microsoft.

random delays [7–16]. In this work, as a baseline technique, we adopt the procedure used in [15] where multiple single-talker audio samples are mixed with random delay while limiting the maximum number of active speakers at the same time to up to two. Let S_0 denote a single-talker ASR training set. For each training data simulation, the number of audio signals \tilde{K} to be mixed is first uniformly sampled from 1 to K , where K is the maximum number of utterances to be mixed. Then, \tilde{K} utterances are randomly sampled from S_0 . Finally, 2nd to \tilde{K} -th utterances are iteratively mixed with the 1st utterance after prepending random silence signals. The duration of the silence prepended for k -th utterance is sampled from $\mathcal{U}(\text{end2}(a_{k-1}), \text{len}(a_{k-1}))$, where $\mathcal{U}(\alpha, \beta)$ is the uniform distribution function in the value range of $[\alpha, \beta]$, a_{k-1} is the mixed audio after mixing $(k-1)$ -th utterance, $\text{len}(a)$ is a function to return the duration of audio a , and $\text{end2}(a)$ is a function to return the end time of the penultimate utterance in a .

3. PROPOSED MULTI-TALKER SPEECH SIMULATION

3.1. Overview

In this work, we propose to simulate multi-talker speech based on the statistics of speech overlaps in real conversation. The overall workflow of our proposed method is as follows. (1) Convert time- and speaker-annotated transcription of real conversation into a sequence of discrete tokens, with each representing the status of speaker overlap within a short time- or word-based unit. (2) Train SLM \mathcal{M} using the discrete token sequence. (3) Simulate multi-talker speech based on single-talker ASR training data S_0 and the SLM \mathcal{M} . (4) Train multi-talker ASR by using the simulated multi-talker speech.

The core of our proposed method lies in converting the time- and speaker-annotated transcription into a discrete token sequence that represents the speaker overlapping pattern. In this paper, we propose two algorithms for such conversion, named time-based discretization and word-based discretization, which are explained in Section 3.2 and Section 3.3, respectively. Note that our algorithms use the notion of “virtual channel” proposed in t-SOT [15]. In the following explanation, we assume two virtual channels (i.e. $M = 2$) for simplicity and consistency with the t-SOT. Our algorithm can be easily extended for $M > 2$ by assuming more virtual channels.

3.2. Time-based discretization

In this method, the speech overlapping pattern is represented by a sequence of discrete tokens, where each token indicates the speech overlapping status of a d -sec time window. The algorithm is depicted in line 1–17, 26 of Algorithm 1. Suppose we have a time- and speaker-annotated transcription \mathcal{T} of a small amount of real conversation data. We first sort \mathcal{T} based on the end time of each word. We then assign each word a “virtual” channel index $c \in \{0, 1\}$. Here, $c = 0$ is always assigned to the first (i.e 0-th) word (line 4, 5). For words thereafter, the value of c is changed when the adjacent words are spoken by different speakers, which is the sign of either speech overlap or speaker turn (lines 7–9). Once the list \mathcal{C} of the channel indices is prepared, the transcription is converted to a sequence \mathcal{X} of discrete tokens $x \in \{0, 1, 2, 3\}$, which indicates the speech activity of each virtual channel for a short d -sec time window (lines 10–17). More specifically, $x = 0$ (or $q = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$) indicates no speech activity is observed during the d -sec time window. On the other hand, $x = 1, 2, 3$ (or $q = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix}$) indicates that the speech activity is observed in 0-th channel, 1-st channel, or both channels during the d -sec time window.

SLM \mathcal{M} can be trained based on the derived sequence of discrete tokens, and then used to yield speech overlapping patterns for simulating more realistic multi-talker speech. The algorithm

Algorithm 1: Discretization of speech overlapping pattern with up to two concurrent utterances

```

1 Input: A speech sample with a time- and speaker-annotated
   transcription  $\mathcal{T} = (w_i, b_i, e_i, s_i)_{i=0}^{N-1}$  where  $w_i$  is  $i$ -th word,
    $b_i, e_i, s_i$  is the begin time, end time, and speaker index of  $w_i$ ,
   respectively. Duration of one discretization unit  $d$ .
2 Output: discretized speaker overlapping pattern  $\mathcal{X}$ 
3 Sort  $\mathcal{T}$  in ascending order of  $e_i$ .
4 Current channel index  $c \leftarrow 0$ .
5 Initialize a list of channel indices  $\mathcal{C} \leftarrow [0]$ .
6 for  $i \leftarrow 1$  to  $N - 1$  do
7   if  $s_i \neq s_{i-1}$  then
8      $c \leftarrow 1 - c$ 
9   Append  $c$  to  $\mathcal{C}$ 
10 if (Time-based discretization) then
11   for  $t \leftarrow 0$  to  $\lfloor e_{N-1}/d \rfloor$  do
12      $q \leftarrow \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ 
13     for  $i \leftarrow 0$  to  $N - 1$  do
14       if  $[b_i, e_i]$  is overlapped with  $[d \cdot t, d \cdot (t + 1)]$  then
15          $q[\mathcal{C}[i]] \leftarrow 1$ 
16        $x \leftarrow q[0] + 2 \cdot q[1]$  // Encode to discrete token
17       Append  $x$  to  $\mathcal{X}$ 
18 else // (Word-based discretization)
19   for  $i \leftarrow 0$  to  $N - 1$  do
20      $q \leftarrow \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ 
21     for  $j \leftarrow 0$  to  $N - 1$  do
22       if  $[b_j, e_j]$  is overlapped with  $[b_i, e_i]$  then
23          $q[\mathcal{C}[j]] \leftarrow 1$ 
24      $x \leftarrow q[0] + 2 \cdot q[1]$  // Encode to discrete token
25     Append  $x$  to  $\mathcal{X}$ 
26 Return  $\mathcal{X}$ 

```

of generating a multi-talker speech sample is described in lines 1–12, 19, 20 of Algorithm 2 and depicted as follows. We first sample a discretized token sequence $\tilde{\mathcal{X}}$ from SLM \mathcal{M} (line 3) based on the time-based discretization. We then apply Decode function that converts $\tilde{\mathcal{X}}$ into a sequence $\tilde{\mathcal{Q}}$, consisting of a list of binary pairs $q \in \{\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix}\}$, where 1 in c -th element of q indicates the existence of speech in the c -th channel (line 4). Next, a target mixed audio a is initialized with empty (line 5). We then call ConsecutiveOne to extract the beginning index i^b and ending index i^e of the consecutive sequence of 1 in either 0-th or 1-st channel of $\tilde{\mathcal{Q}}$, in chronological order of i^b (line 7). For example, when $\tilde{\mathcal{Q}} = [\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}]$, ConsecutiveOne returns $i^b = 0, i^e = 3$ for the first iteration, and returns $i^b = 2, i^e = 4$ for the second iteration. Based on i^b and i^e , we sample a speech segment u from source ASR training data S_0 , with the duration within or nearest to the expected interval $[d^{\min}, d^{\max}]$ (lines 9–11).¹ A duration γ of the silence prepended to u is randomly determined such that u is mixed to the audio a with the expected position from i^b (line 12, 19). These procedures are repeated to form the mixed audio a , which is returned as the training sample for multi-talker ASR (line 20).

¹To save computation, we created a pool of 10K utterances from S_0 , and used it to sample u . When we created the pool, utterances in S_0 were short-segmented at the point of silence longer than 0.5 sec. We randomly sampled u among samples whose duration is within $[d^{\min}, d^{\max}]$. If there was no such sample, we sample an utterance having closest duration to $(d^{\min}, d^{\max})/2$.

Algorithm 2: Multi-talker speech generation.

```

1 Input: Single-talker ASR training data set  $S_0$ . SLM learned from
  the discretized speaker overlapping pattern  $\mathcal{M}$ . Duration of one
  discretization unit  $d$ .
2 Output: Simulated multi-talker speech sample  $a$ 
3 Sample a discretized token sequence  $\tilde{\mathcal{X}}$  from  $\mathcal{M}$ 
4  $\tilde{\mathcal{Q}} \leftarrow \text{Decode}(\tilde{\mathcal{X}})$  //  $0 \rightarrow \begin{bmatrix} 0 \\ 0 \end{bmatrix}, 1 \rightarrow \begin{bmatrix} 1 \\ 0 \end{bmatrix}, 2 \rightarrow \begin{bmatrix} 0 \\ 1 \end{bmatrix}, 3 \rightarrow \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ 
5 Initialize mixed audio  $a \leftarrow \emptyset$ 
6 End times of each word  $D \leftarrow \emptyset$  // Only for word-based disc.
7 for  $(i^b, i^e) \leftarrow \text{ConsecutiveOne}(\tilde{\mathcal{Q}})$  do
8   if (Time-based discretization) then
9      $d^{\min} \leftarrow d \cdot (i^e - i^b)$ 
10     $d^{\max} \leftarrow d \cdot (i^e - i^b + 1)$ 
11     $u \leftarrow \text{DurationNearestSample}(d^{\min}, d^{\max}, S_0)$ 
12     $\gamma \leftarrow d \cdot i^b + \mathcal{U}(0, d^{\max} - \text{len}(u))$  // Silence duration
13  else // (Word-based discretization)
14     $\mathcal{I} = [i_0, \dots, i_{M-1}] \leftarrow \text{WordIndices}(i^b, i^e, \tilde{\mathcal{Q}})$ 
15     $u \leftarrow \text{WordCountNearestSample}(\text{len}(\mathcal{I}), S_0)$ 
16     $\gamma \leftarrow D[i_0 - 2] + \mathcal{U}(0, D[i_0 - 1] - D[i_0 - 2])$ 
17    for  $j \leftarrow 0$  to  $M - 1$  do
18       $D[i_j] \leftarrow \text{end time of } j\text{-th word in } u$ 
19   $a \leftarrow a + [\text{sil}(\gamma), u]$ 
20 Return  $a$ 

```

3.3. Word-based discretization

The word-based discretization algorithm differs from the time-based one in a sense that the discretized unit representing the overlapping status is based on “word” rather than a d -sec time window. The algorithm of converting \mathcal{T} into the discretized speaker overlapping pattern \mathcal{X} is depicted in lines 1–9, 18–26 of Algorithm 1. The word-based discretization shares the same procedure with the time-based discretization for the first 9 steps, and then performs discretization for each word in \mathcal{T} , where x represents the overlapping status of speech overlaps of each token (lines 18–25). Note that x cannot be 0 (or q cannot be $\begin{bmatrix} 0 \\ 0 \end{bmatrix}$) in this algorithm. Fundamentally, word-based discretization is assumed to be applied to a short segment that does not include long-silence signals. This can be achieved by segmenting the transcription based on the existence of long silence.

The multi-speaker speech generation algorithm with word-based discretization is introduced in Algorithm 2. The first 7 steps are the same as the time-based discretization algorithm except that we create a buffer D to keep the end time of each word of utterances being mixed (line 6). The buffer D is necessary to determine the duration γ of silence when we mix u to a with targeted overlaps (lines 16–18).² In line 14, we first apply WordIndices function, which returns a list \mathcal{I} of indices indicating where each word should be allocated. This function is necessary because around half of the consecutive sequence of $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ represents the words in the other channel. For example, given $\tilde{\mathcal{Q}} = [\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}]$, ConsecutiveOne first returns $i^b = 0, i^e = 3$, and WordIndices returns $[0, 1, 2]$. For the second iteration, ConsecutiveOne returns $i^b = 2, i^e = 4$, and WordIndices returns $[3, 4]$. We then sample u whose number of words are closest to the length of \mathcal{I} (line 15).³ We mix u to a with a silence signal whose duration γ is determined such that the word in u is overlapping with a as represented by $\tilde{\mathcal{Q}}$ (lines 16, 19).

²We assume D will return 0 for an undefined index.

³We used a pool of 10K utterances from S_0 , and randomly sampled u among samples whose number of words is $\text{len}(\mathcal{I})$. If there was no such sample, we sample an utterance having the closest word count to $\text{len}(\mathcal{I})$.

Table 1. Datasets used in the experiments.

	AMI			ICSI			MS ^{mtg}	MS ^{asr}
	train	dev	eval	train	dev	eval		
Size (hr)	80.2	9.7	9.1	66.6	2.3	2.8	15.5	75K
# meetings	135	18	16	70	2	3	33	-
# speakers / meeting	3–5	4	3–4	3–10	6–7	7	2–19	-
Overlap SLM training	✓	-	-	-	-	-	-	-
t-SOT pre-training	-	-	-	-	-	-	-	✓
t-SOT fine-tuning	✓	-	-	✓	-	-	-	✓
Evaluation	-	✓	✓	-	✓	✓	✓	-

Table 2. WERs (%) of t-SOT TT18 (0.16-sec latency) trained with different algorithms for multi-talker audio simulation. 10K steps of fine-tuning was performed from the pre-trained t-SOT model.

Fine-tuning data		AMI			ICSI		MS ^{mtg}
Algorithm	Filt.	dev (1-spkr / m-spkr)	eval	dev	eval		
-	-	35.5 (24.4 / 42.5)	39.2	32.2	33.3	27.3	
Rand ($K = 2$)	-	35.9 (25.6 / 42.4)	40.4	33.8	34.0	27.9	
Rand ($K = 2$)	✓	31.6 (24.5 / 36.0)	37.3	29.9	27.8	25.3	
Rand ($K = 5$)	✓	31.2 (25.2 / 34.9)	36.6	29.9	27.3	24.8	
Rand ($K = 8$)	✓	31.1 (25.0 / 34.8)	36.8	29.9	27.1	25.0	
Word	✓	30.6 (25.5 / 33.7)	36.1	29.6	26.5	25.0	
Time ($d = 0.10$)	✓	31.4 (26.5 / 34.4)	36.0	30.2	27.4	24.7	
Time ($d = 0.25$)	✓	30.8 (26.0 / 33.8)	35.5	29.8	26.7	24.5	
Time ($d = 0.50$)	✓	32.6 (26.6 / 36.4)	37.3	30.8	28.4	25.6	

4. EXPERIMENTS**4.1. Evaluation data**

Table 1 shows the dataset used in our experiments. As public meeting evaluation sets, we used the AMI meeting corpus [21] and the ICSI meeting corpus [22]. For the AMI corpus, we used the first channel of the microphone array signals, also known as the single distant microphone (SDM) signal. For the ICSI corpus, we used the SDM signal from the D2 microphone. Both corpora were pre-processed and split into training, development, and evaluation sets by using the Kaldi toolkit [23]. In addition to these public corpora, we also used 33 internal meeting recordings based on an SDM, denoted as MS^{mtg}. For all datasets, we applied causal logarithmic-loop-based automatic gain control to normalize the significant volume differences among different recordings. Evaluation was conducted based on the utterance-group segmentation [17], and speaker-agnostic WER was calculated by using the multiple dimension Levenshtein edit distance [16, 24].

For our proposed multi-talker audio simulation, the AMI training data with official word-level time stamps was used to train the SLM, where we trained N-gram with $N = 30$ based on the NLTK toolkit [25]. For the source data (S_0) of the multi-talker audio simulation, we used 75 thousand (K) hours of 64 million (M) anonymized and transcribed single-talker English utterances, denoted as MS^{asr} [17]. MS^{asr} consists of audio signals from various domains, such as dictation and voice commands, and each audio was supposed to contain single-talker speech. However, we found it occasionally contained untranscribed background human speech, which broke the assumption for the multi-talker audio simulation. Therefore, we filtered out the audio sample that potentially contains more than two speaker audio. To detect such audio samples, we applied serialized output training-based multi-talker ASR pre-trained by MS^{asr} and fine-tuned by AMI [17] for all audio samples in MS^{asr} with a beam size of 1. With this procedure, transcriptions of more than one speaker were generated for 14M utterances out of 64M utterances, which were excluded in the multi-talker audio simulation. The effect of this data filtering is examined in Section 4.3.1.

Table 3. WERs (%) of t-SOT TT18 (0.16-sec latency) with different combinations of multi-talker simulation algorithms. 20K steps of fine-tuning was performed from the pre-trained t-SOT model.

Fine-tuning data			AMI		ICSI		MS ^{mtg}
rand	word	time	dev (1-spkr / m-spkr)	eval	dev	eval	
1.0	-	-	31.0 (25.1 / 34.6)	36.4	29.6	26.9	24.8
-	1.0	-	31.0 (26.3 / 34.0)	36.2	29.6	26.4	25.0
-	-	1.0	31.1 (26.6 / 33.9)	35.9	29.8	26.7	24.6
0.8	0.2	-	30.5 (24.9 / 33.9)	35.9	29.1	26.1	24.5
0.5	0.5	-	30.3 (24.8 / 33.7)	35.7	29.3	25.8	24.3
0.2	0.8	-	30.7 (25.5 / 33.9)	35.7	29.5	26.1	24.8
0.8	-	0.2	30.5 (25.0 / 34.0)	35.8	29.2	26.5	24.3
0.5	-	0.5	30.5 (25.0 / 33.9)	36.0	29.3	26.2	24.4
0.2	-	0.8	30.2 (25.0 / 33.4)	35.4	28.8	25.9	24.0
0.1	0.1	0.8	30.3 (25.4 / 33.3)	35.3	29.0	25.9	24.1
0.2	0.2	0.6	30.0 (25.2 / 33.1)	35.3	28.9	25.5	24.0
0.3	0.3	0.4	30.0 (25.0 / 33.2)	35.3	28.7	25.4	24.0
0.4	0.4	0.2	30.3 (24.8 / 33.8)	35.7	28.8	25.8	24.2

4.2. Multi-talker ASR configuration

As an instance of multi-talker ASR, we trained t-SOT based transformer transducer (TT) [26] with chunk-wise look-ahead [27], where the algorithmic latency of the model can be controlled based on the chunk size of the attention mask. We trained TT models with 18 and 36 layers of transformer encoders, which were denoted by TT18 and TT36, respectively. Each transformer block consisted of a 512-dim multi-head attention with 8 heads and a 2048-dim point-wise feed-forward layer. The prediction network consisted of two layers of 1024-dim long short-term memory. 4,000 word pieces plus blank and ⟨cc⟩ tokens were used as the recognition units. We used 80-dim log mel-filterbank extracted for every 10 msec.

All models were first pre-trained by using the multi-talker simulation data based on MS^{asr} with the random simulation algorithm with $K = 2$. We performed 425K training steps with 32 GPUs, with each GPU processing a mini-batch of 24K frames. A linear decay learning rate schedule with a peak learning rate of $1.5e-3$ after 25K warm-up iterations were used. After the pre-training, the model was further fine-tuned by using the simulation data based on MS^{asr} and/or AMI & ICSI training data. We performed various fine-tuning configurations, which will be presented in the next section.

4.3. Experimental results and analysis

4.3.1. Comparison of simulation algorithms

The WERs with different simulation algorithms are summarized in Table 2. In this experiment, we conducted 10K steps of fine-tuning with 8 GPUs where each GPU node consumed 24K frames of training samples for each training step. A linear decay learning rate schedule starting from a learning rate of $1.5e-4$ was used. The 1st and 2nd rows are the results of random simulation with $K = 2$ for both pre-training and fine-tuning. Because the data and algorithm used for the pre-training and fine-tuning are the same, we didn’t observe any improvement by additional fine-tuning. Then, we applied the data filtering explained in Section 4.1 (3rd row). We observed significant WER improvements for all evaluation sets, confirming the effectiveness of the data filtering. We also evaluated the effect of different K as shown in the 3rd to 5th rows, where we observed $K = 5$ and $K = 8$ provides better results than $K = 2$.

Then, we evaluated the proposed simulation algorithms, whose results are shown in the bottom four rows. We observed that the best results for AMI-dev, ICSI-dev and ICSI-eval were obtained by the proposed method with word-based discretization while the best results for AMI-eval and MS^{asr} were obtained by the proposed method

Table 4. WERs (%) of t-SOT TT18 (0.16-sec latency) and t-SOT TT36 (2.56-sec latency) with different combinations of real training data (AMI and ICSI) and proposed simulation training data.

Model	Fine-tuning data		AMI		ICSI		MS ^{mtg}
	AMI+ICSI	MS ^{asr} -sim	dev	eval	dev	eval	
t-SOT TT18	1.0	-	21.9	25.7	20.3	17.9	24.9
	0.5	0.5	21.6	25.4	19.5	17.1	24.6
	0.2	0.8	21.8	25.5	19.9	17.3	23.4
	0.1	0.9	22.4	26.1	21.0	18.2	23.1
	-	1.0	30.0	35.3	28.7	25.4	24.0
t-SOT TT36	1.0	-	16.9	19.7	15.6	14.0	19.9
	0.2	0.8	16.8	19.7	15.3	13.8	18.9

with time-based discretization with $d = 0.25$. We also noticed that both of the proposed methods improved the WER of multi-speaker (m-spkr) regions, while had a degradation of the performance for the single-speaker (1-spkr) regions. We speculate this was caused because we short-segmented the sample in S_0 when we created the pool in the simulation process (see footnote 1), which introduced an unnatural onset / offset of the audio.

4.3.2. Combination of multiple simulation algorithms

Given the observation about the trade-off between 1-spkr and m-spkr performance, we investigated the combination of the data from different simulation algorithms. The results are shown in Table 3. In this experiment, we increased the fine-tuning steps from 10K to 20K to make sure sufficient data were generated from each multi-talker audio simulation algorithm. We used $K = 5$ for the random simulation, and $d = 0.25$ for the time-based discretization. From the table, we can observe that the trade-off between 1-spkr and m-spkr regions were effectively resolved by mixing different types of multi-talker simulation data sets. The best results for all evaluation sets were obtained when we combined all three multi-talker audio simulation algorithms with the mixture ratio of 0.3, 0.3, 0.4 for the random simulation, the proposed simulation with word-based discretization, and the proposed simulation with time-based discretization, respectively. Finally, the paired t-test between the random simulation and the proposed method with the best mixture ratio was performed, and the difference of WERs was judged to be significant in 99% confidence.

4.3.3. Combination of real and simulated multi-talker audio

Finally, we also evaluated the combination of real meeting data (AMI and ICSI training data) and the proposed simulation data. For the simulation data, we used the best combination of three simulation algorithms as in the previous section, and we performed 20K steps of fine-tuning. Note that we performed only 2.5K steps of fine-tuning with 8 GPUs (12K frames of mini-batch per GPU) when the fine-tuning data does not include the simulation data to avoid over-fitting phenomena. As shown in the table, we observed the best results for all evaluation sets when we mixed the real and simulation data with a 0.2 / 0.8 ratio. We also performed the paired t-test between the model fine-tuned by AMI+ICSI and that trained by the proposed mixed data, and the WER differences were judged to be significant in 99% confidence for both TT18 and TT36.

5. CONCLUSION

This paper presented improved multi-talker audio simulation techniques for multi-talker ASR modeling. We proposed two algorithms to represent the speech overlap patterns based on an SLM, which was then used to simulate multi-talker audio with realistic speech overlaps. In our experiments using multiple meeting evaluation sets, we demonstrated that multi-talker ASR models trained with the proposed method consistently showed improved WERs across datasets.

6. REFERENCES

- [1] Frank Seide, Gang Li, and Dong Yu, “Conversational speech transcription using context-dependent deep neural networks,” in *Proc. Interspeech*, 2011, pp. 437–440.
- [2] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdelrahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al., “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *IEEE Signal processing magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [3] Yanmin Qian, Mengxiao Bi, Tian Tan, and Kai Yu, “Very deep convolutional neural networks for noise robust speech recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 12, pp. 2263–2276, 2016.
- [4] Özgür Çetin and Elizabeth Shriberg, “Analysis of overlaps in meetings by dialog factors, hot spots, speakers, and collection site: Insights for automatic speech recognition,” in *Proc. Interspeech*, 2006, pp. 293–296.
- [5] Zhuo Chen, Takuya Yoshioka, Liang Lu, Tianyan Zhou, Zhong Meng, Yi Luo, Jian Wu, and Jinyu Li, “Continuous speech separation: dataset and analysis,” in *Proc. ICASSP*, 2020, pp. 7284–7288.
- [6] Desh Raj, Pavel Denisov, Zhuo Chen, Hakan Erdogan, Zili Huang, Maokui He, Shinji Watanabe, Jun Du, Takuya Yoshioka, Yi Luo, et al., “Integration of speech separation, diarization, and recognition for multi-speaker meetings: System description, comparison, and analysis,” in *Proc. SLT*, 2021, pp. 897–904.
- [7] Dong Yu, Xuankai Chang, and Yanmin Qian, “Recognizing multi-talker speech with permutation invariant training,” *Proc. Interspeech*, pp. 2456–2460, 2017.
- [8] Hiroshi Seki, Takaaki Hori, Shinji Watanabe, Jonathan Le Roux, and John R Hershey, “A purely end-to-end system for multi-speaker speech recognition,” *Proc. ACL*, 2018.
- [9] Xuankai Chang, Yanmin Qian, Kai Yu, and Shinji Watanabe, “End-to-end monaural multi-speaker ASR system without pre-training,” in *Proc. ICASSP*, 2019, pp. 6256–6260.
- [10] Xuankai Chang, Wangyou Zhang, Yanmin Qian, Jonathan Le Roux, and Shinji Watanabe, “MIMO-SPEECH: End-to-end multi-channel multi-speaker speech recognition,” in *Proc. ASRU*, 2019, pp. 237–244.
- [11] Anshuman Tripathi, Han Lu, and Hasim Sak, “End-to-end multi-talker overlapping speech recognition,” in *Proc. ICASSP*. IEEE, 2020, pp. 6129–6133.
- [12] Naoyuki Kanda, Yashesh Gaur, Xiaofei Wang, Zhong Meng, and Takuya Yoshioka, “Serialized output training for end-to-end overlapped speech recognition,” *Proc. Interspeech*, pp. 2797–2801, 2020.
- [13] Liang Lu, Naoyuki Kanda, Jinyu Li, and Yifan Gong, “Streaming end-to-end multi-talker speech recognition,” *IEEE Signal Processing Letters*, vol. 28, pp. 803–807, 2021.
- [14] Ilya Sklyar, Anna Piunova, and Yulan Liu, “Streaming multi-speaker ASR with RNN-T,” in *Proc. ICASSP*, 2021, pp. 6903–6907.
- [15] Naoyuki Kanda, Jian Wu, Yu Wu, Xiong Xiao, Zhong Meng, Xiaofei Wang, Yashesh Gaur, Zhuo Chen, Jinyu Li, and Takuya Yoshioka, “Streaming multi-talker ASR with token-level serialized output training,” *Proc. Interspeech*, pp. 3774–3778, 2022.
- [16] Naoyuki Kanda, Jian Wu, Xiaofei Wang, Zhuo Chen, Jinyu Li, and Takuya Yoshioka, “VarArray meets t-SOT: Advancing the state of the art of streaming distant conversational speech recognition,” *arXiv preprint arXiv:2209.04974*, 2022.
- [17] Naoyuki Kanda, Guoli Ye, Yu Wu, Yashesh Gaur, Xiaofei Wang, Zhong Meng, Zhuo Chen, and Takuya Yoshioka, “Large-scale pre-training of end-to-end multi-talker ASR for meeting transcription with single distant microphone,” *Proc. Interspeech*, pp. 3430–3434, 2021.
- [18] Federico Landini, Alicia Lozano-Diez, Mireia Diez, and Lukáš Burget, “From simulated mixtures to simulated conversations as training data for end-to-end neural diarization,” *Proc. Interspeech*, 2022.
- [19] Natsuo Yamashita, Shota Horiguchi, and Takeshi Homma, “Improving the naturalness of simulated conversations for end-to-end neural diarization,” *Proc. Odyssey*, 2022.
- [20] Jinyu Li, “Recent advances in end-to-end automatic speech recognition,” *APSIPA Transactions on Signal and Information Processing*, vol. 11, no. 1, 2022.
- [21] Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Mael Guillemot, Thomas Hain, Jaroslav Kadlec, Vasilis Karaiskos, Wessel Kraaij, Melissa Kronenthal, et al., “The ami meeting corpus: A pre-announcement,” in *International workshop on machine learning for multimodal interaction*. Springer, 2005, pp. 28–39.
- [22] Adam Janin, Don Baron, Jane Edwards, Dan Ellis, David Gelbart, Nelson Morgan, Barbara Peskin, Thilo Pfau, Elizabeth Shriberg, Andreas Stolcke, et al., “The ICSI meeting corpus,” in *Proc. ICASSP*. IEEE, 2003, vol. 1, pp. I–I.
- [23] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, et al., “The Kaldi speech recognition toolkit,” in *Proc. ASRU*, 2011.
- [24] Jonathan G Fiscus, Jerome Ajot, Nicolas Radde, and Christophe Laprun, “Multiple dimension Levenshtein edit distance calculations for evaluating automatic speech recognition systems during simultaneous speech,” in *Proc. LREC*, 2006, pp. 803–808.
- [25] Steven Bird, “NLTK: the natural language toolkit,” in *Proceedings of the COLING/ACL 2006 Interactive Presentation Sessions*, 2006, pp. 69–72.
- [26] Qian Zhang, Han Lu, Hasim Sak, Anshuman Tripathi, Erik McDermott, Stephen Koo, and Shankar Kumar, “Transformer transducer: A streamable speech recognition model with transformer encoders and RNN-T loss,” in *Proc. ICASSP*, 2020, pp. 7829–7833.
- [27] Xie Chen, Yu Wu, Zhenghao Wang, Shujie Liu, and Jinyu Li, “Developing real-time streaming transformer transducer for speech recognition on large-scale dataset,” in *Proc. ICASSP*, 2021, pp. 5904–5908.