

# END-TO-END MONAURAL MULTI-SPEAKER ASR SYSTEM WITHOUT PRETRAINING

Xuankai Chang<sup>1,2</sup>, Yanmin Qian<sup>1†</sup>, Kai Yu<sup>1</sup>, Shinji Watanabe<sup>2†</sup>

<sup>1</sup>SpeechLab, Department of Computer Science and Engineering, Shanghai Jiao Tong University, China

<sup>2</sup>Center for Language and Speech Processing, Johns Hopkins University, U.S.A

xuank@sjtu.edu.cn, yanminqian@sjtu.edu.cn, kai.yu@sjtu.edu.cn, shinjiw@jhu.edu

## ABSTRACT

Recently, end-to-end models have become a popular approach as an alternative to traditional hybrid models in automatic speech recognition (ASR). The multi-speaker speech separation and recognition task is a central task in cocktail party problem. In this paper, we present a state-of-the-art monaural multi-speaker end-to-end automatic speech recognition model. In contrast to previous studies on the monaural multi-speaker speech recognition, this end-to-end framework is trained to recognize multiple label sequences completely from scratch. The system only requires the speech mixture and corresponding label sequences, without needing any indeterminate supervisions obtained from non-mixture speech or corresponding labels/alignments. Moreover, we exploited using the individual attention module for each separated speaker and the scheduled sampling to further improve the performance. Finally, we evaluate the proposed model on the 2-speaker mixed speech generated from the WSJ corpus and the wsj0-2mix dataset, which is a speech separation and recognition benchmark. The experiments demonstrate that the proposed methods can improve the performance of the end-to-end model in separating the overlapping speech and recognizing the separated streams. From the results, the proposed model leads to  $\sim 10.0\%$  relative performance gains in terms of CER and WER respectively.

**Index Terms**— Cocktail party problem, multi-speaker speech recognition, end-to-end speech recognition, CTC, attention mechanism

## 1. INTRODUCTION

In the deep learning era, single-speaker automatic speech recognition systems have achieved a lot of progress. Deep neural networks (DNN) and hidden markov model (HMM) based hybrid systems have attained surprisingly good performance [1, 2, 3, 4]. Recently, there has been a growing interest in developing end-to-end models for speech recognition [5, 6, 7], in which the various modules of the hybrid systems, such as the acoustic model (AM) and language model (LM), are folded into a single neural network model. Two major approaches of end-to-end speech recognition systems are connectionist temporal classification [8, 9] and attention-based encoder-decoder [10, 11]. The performance of deep learning based conventional speech recognition systems has been reported to be comparable with, or even surpassing, human performance [3]. However, it is still extremely difficult to solve the cocktail party problem

[12, 13, 14, 15], which refers to the task of separating and recognizing the speech from a specific speaker when it is interfered by noise and speech from other speakers.

To address the monaural multi-speaker speech separation and recognition problem, there has been a lot of research in single-channel multi-speaker speech separation and recognition, which aims to separate the overlapping speech and recognize the resulting separated speech individually, given a single-channel multiple-speaker mixed speech. In [16, 17], a method called deep clustering (DPCL) was proposed for speech separation. DPCL separates the mixed speech by training a neural network to project each time-frequency (T-F) unit into a high-dimensional embedding space, in which pairs of T-F units are close to each other if they have the same dominating speaker and farther away otherwise. In addition to segmentation using k-means clustering, a permutation-free mask objective was proposed to refine the output [17]. In [18, 19], a speech separation method called permutation invariant training (PIT) was proposed to train a compact deep neural network with the objective that minimizes the average minimum square error of the best output-target assignment at the utterance level. PIT was later extended to train a speech recognition model for multi-speaker speech mixture by directly optimizing with the ASR objective [20, 21, 22, 23]. In [24, 25], a joint CTC/attention-based encoder-decoder network for end-to-end speech recognition [5, 6] was applied to multi-speaker speech recognition. First, an encoder separates the mixed speech into hidden vector sequences for every speaker. Then an attention-based decoder is used to generate the label sequence for each speaker. To avoid label permutation problem, a CTC objective is used in permutation-free manner right after the encoder to determine the order of the label sequences. However, the model needs to first be pre-trained on single-speaker speech so that decent performance can be achieved.

In this paper, we explore several new methods to refine the end-to-end speech recognition model for multi-speaker speech. Firstly, we revise the model in [25] so that pretraining on single-speaker speech is not required without loss of performance. Secondly, we propose to use speaker parallel attention modules. In previous work, the separated speech streams were treated equally in the decoder, regardless of the energy and speaker characteristics. We bring in multiple attention modules [26] for each speaker to enhance the speaker tracing ability and to alleviate the burden of the encoder similar to [22]. Another method is to use scheduled sampling [27] to randomly choose the token from either the ground truth or the model prediction as the history information, which reduces the gap between training and inference in the sequence prediction tasks. This would be extremely helpful in our setup, since the separation is not always perfect and we often observe mixed label results. Schedule sampling can help to recover such errors during inference.

The rest of the paper is organized as follows: In Section 2, the

This work was done while Xuankai Chang was an intern at the Johns Hopkins University.

<sup>†</sup>Yanmin Qian and Shinji Watanabe are the corresponding authors

end-to-end monaural multi-speaker ASR model and the proposed new methods are described. In Section 3, we evaluate the proposed approach on the 2-speaker mixing WSJ data set, and the experiments and analysis are given. Finally the paper is concluded in Section 4.

## 2. END-TO-END MULTI-SPEAKER JOINT CTC/ATTENTION-BASED ENCODER-DECODER

In this section, we first describe the end-to-end ASR system for multi-speaker speech that has been used in [25]. Then we introduce two techniques to improve the training process and performance of the end-to-end ASR multi-speaker system, namely the speaker parallel attention and scheduled sampling [27].

### 2.1. End-to-End Multi-speaker ASR

In [5, 6, 28], an end-to-end speech recognition model was proposed to take advantage of both the Connectionist Temporal Classification (CTC) and attention-based encoder-decoder, in aim of using the CTC to enhance the alignment ability of the model. An end-to-end model for multi-speaker speech recognition was brought up in [25], extending the joint CTC/attention-based encoder-decoder network to be applied on multi-speaker speech mixtures and to allow the permutation-free training in the objective function to address the permutation problem. The model is shown in Fig.1, in which the modules *Attention 1* and *Attention 2* share parameters. The input speech mixture is first explicitly separated into multiple sequences of vectors in the encoder, each representing a speaker source. These sequences are fed into the decoder to compute the conditional probabilities.

The encoder of the model can be divided into three stages, namely the  $\text{Encoder}_{\text{Mix}}$ ,  $\text{Encoder}_{\text{SD}}$  and  $\text{Encoder}_{\text{Rec}}$ . Let  $\mathbf{O}$  denote an input speech mixture from  $S$  speakers. The first stage,  $\text{Encoder}_{\text{Mix}}$ , is the mixture encoder, which encodes the input speech mixture  $\mathbf{O}$  as an intermediate representation  $\mathbf{H}$ . Then, the representation  $\mathbf{H}$  is processed by  $S$  speaker-different (SD) encoders,  $\text{Encoder}_{\text{SD}}$ , with the outputs being referred to as feature sequences  $\mathbf{H}^s, s = 1, \dots, S$ .  $\text{Encoder}_{\text{Rec}}$ , the last stage, transforms the features sequences to high-level representations  $\mathbf{G}^s, s = 1, \dots, S$ . The encoder is computed as

$$\mathbf{H} = \text{Encoder}_{\text{Mix}}(\mathbf{O}) \quad (1)$$

$$\mathbf{H}^s = \text{Encoder}_{\text{SD}}^s(\mathbf{H}), s = 1, \dots, S \quad (2)$$

$$\mathbf{G}^s = \text{Encoder}_{\text{Rec}}(\mathbf{H}^s), s = 1, \dots, S \quad (3)$$

In the single-speaker joint CTC/attention-based encoder-decoder network, the CTC objective function is used to train the attention model encoder as an auxiliary task right after the encoder [5, 6, 28]. While in the multi-speaker framework, the CTC objective function is also used to perform the permutation-free training as in Eq.4, which is referred to as permutation invariant training in [15, 18, 20, 21, 22, 23, 29, 30].

$$\hat{\pi} = \arg \min_{\pi \in \mathcal{P}} \sum_s \text{Loss}_{\text{ctc}}(\mathbf{Y}^s, \mathbf{R}^{\pi(s)}), \quad (4)$$

where  $\mathbf{Y}^s$  is the output sequence variable computed from the encoder output  $\mathbf{G}^s$ ,  $\pi(s)$  is the  $s$ -th element in a permutation  $\pi$  of  $\{1, \dots, S\}$ , and  $\mathbf{R}$  is the reference labels for  $S$  speakers. Later, the permutation  $\hat{\pi}$  with minimum CTC loss is used for the reference labels in the attention-based decoder in order to reduce the computational cost.

After obtaining the representations  $\mathbf{G}^s, s = 1, \dots, S$  from the encoder, an attention-based decoder network is used to decode these streams and output label sequence  $\mathbf{Y}^s$  for each representation stream according to the permutation determined by the CTC objective function. For each pair of representation and reference label index  $(s, \hat{\pi}(s))$ , the decoding process is described as the following equations:

$$p_{\text{att}}(\mathbf{Y}^{s, \hat{\pi}(s)} | \mathbf{O}) = \prod_n p_{\text{att}}(y_n^{s, \hat{\pi}(s)} | \mathbf{O}, y_{1:n-1}^{s, \hat{\pi}(s)}) \quad (5)$$

$$c_n^{s, \hat{\pi}(s)} = \text{Attention}(a_{n-1}^{s, \hat{\pi}(s)}, e_{n-1}^{s, \hat{\pi}(s)}, \mathbf{G}^s) \quad (6)$$

$$e_n^{s, \hat{\pi}(s)} = \text{Update}(e_{n-1}^{s, \hat{\pi}(s)}, c_{n-1}^{s, \hat{\pi}(s)}, y_{n-1}^{s, \hat{\pi}(s)}) \quad (7)$$

$$y_n^{s, \hat{\pi}(s)} \sim \text{Decoder}(c_n^{s, \hat{\pi}(s)}, y_{n-1}^{s, \hat{\pi}(s)}) \quad (8)$$

where  $c_n^{s, \hat{\pi}(s)}$  denotes the context vector,  $e_n^{s, \hat{\pi}(s)}$  is the hidden state of the decoder, and  $r_n^{\hat{\pi}(s)}$  is the  $n$ -th element in the reference label sequence. During training, the reference label  $r_{n-1}^{\hat{\pi}(s)}$  in  $\mathbf{R}$  is used as a history in the manner of teacher-forcing, instead of  $y_{n-1}^{\hat{\pi}(s)}$  in Eq.7 and Eq.8. And, Eq.5 means the probability of the target label sequence  $\mathbf{Y} = \{y_1, \dots, y_N\}$  that the attention-based encoder-decoder predicted, in which the probability of  $y_n$  at  $n$ -th time step is dependent on the previous sequence  $y_{1:n-1}$ .

The final loss function is defined as

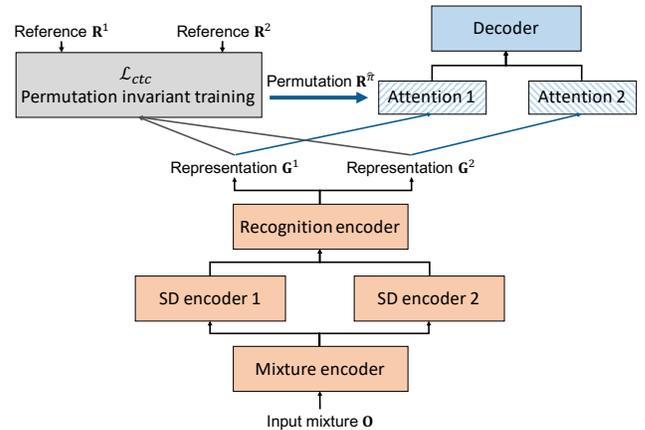
$$\mathcal{L}_{\text{mtl}} = \lambda \mathcal{L}_{\text{ctc}} + (1 - \lambda) \mathcal{L}_{\text{att}}, \quad (9)$$

$$\mathcal{L}_{\text{ctc}} = \sum_s \text{Loss}_{\text{ctc}}(\mathbf{Y}^s, \mathbf{R}^{\hat{\pi}(s)}), \quad (10)$$

$$\mathcal{L}_{\text{att}} = \sum_s \text{Loss}_{\text{att}}(\mathbf{Y}^{s, \hat{\pi}(s)}, \mathbf{R}^{\hat{\pi}(s)}), \quad (11)$$

where  $\lambda$  is the interpolation factor, and  $0 \leq \lambda \leq 1$ .

### 2.2. Speaker parallel attention modules



**Fig. 1.** End-to-End Multi-speaker Speech Recognition Model in the 2-Speaker Case

Due to the differences in the characteristics of speakers and energy, the encoder usually has to compensate for those differences while separating the speech. The motivation of speaker parallel attention module that we proposed is to alleviate the burden for the encoder and to make the attention-decoder learn to filter the separated

speech as well while keeping the model compact. In light of [22], we proposed to use independent attention modules called speaker parallel attention. Fig.1 illustrates the architecture of the model, in which *Attention 1* and *Attention 2* are not sharing. The computation process in Eq.6 should be rewritten in a stream-specific way, in particular for the  $s$ -th stream, as:

$$c_n^{s,\hat{\pi}(s)}, a_n^{s,\hat{\pi}(s)} = \text{Attention}^s(a_{n-1}^{s,\hat{\pi}(s)}, c_{n-1}^{s,\hat{\pi}(s)}, \mathbf{G}^s) \quad (12)$$

### 2.3. Scheduled sampling

We generally trained the decoder network in a teacher-forcing fashion, which means the reference label token  $r_n$ , not the predicted token  $y_n$ , is used to predict the next token in the sequence during training. However, during inference, we are only accessible to the predicted token  $y_n$  from the model itself. This difference may lead to performance degradation, especially in the multi-speaker speech recognition task susceptible to the label permutation problem. We alleviate this problem by using the scheduled sampling technique [27]. During training, whether the history information is chosen from the ground truth label or the prediction is done randomly with a probability of  $p$  from the prediction and  $(1 - p)$  from ground truth. Thus Eq.7 and Eq.8 should be changed as:

$$e_n^{s,\hat{\pi}(s)} = \text{Update}(e_{n-1}^{s,\hat{\pi}(s)}, c_{n-1}^{s,\hat{\pi}(s)}, h), \quad (13)$$

$$y_n^{s,\hat{\pi}(s)} \sim \text{Decoder}(c_n^{s,\hat{\pi}(s)}, h), \quad (14)$$

where

$$b \sim \text{Bernoulli}(p), \quad (15)$$

$$h = \begin{cases} r_{n-1}^{\hat{\pi}(s)}, & \text{if } b = 0 \\ y_{n-1}^{\hat{\pi}(s)}, & \text{if } b = 1 \end{cases} \quad (16)$$

## 3. EXPERIMENT

### 3.1. Experimental setup

To evaluate our method, we used the artificially generated single-channel two-speaker mixed signals from the Wall Street Journal (WSJ) speech corpus according to [25], using the tool released by MERL<sup>1</sup>. We used the WSJ SI284 to generate the training data, Dev93 for development and Eval92 for evaluation. The durations for the training, development and evaluation sets of the mixed data are 98.5 hr, 1.3 hr, and 0.8 hr respectively. In section 3.4, we also compared our model with previous works on the wsj0-2mix dataset, which is a standard speech separation and recognition benchmark [16, 17, 24].

The input feature is 80-dimensional log Mel filterbank coefficients with pitch features and their delta and delta delta features extracted using the Kaldi [31]. Zero mean and unit variance are used to normalize the input features. All the joint CTC/attention-based encoder-decoder networks for end-to-end speech recognition were built based on the ESPnet [7] framework. The networks were initialized randomly from uniform distribution in the range  $-0.1$  to  $0.1$ . We used the AdaDelta algorithm with  $\rho = 0.95$  and  $\epsilon = 1e - 8$ . During training, we set the interpolation factor  $\lambda$  in Eq.9 to be 0.2. We revise the deep neural network, replacing the original encoder layers with shallower but wider layers [32], so that the performance can be good enough without pre-training on single-speaker speech.

<sup>1</sup><http://www.merl.com/demos/deep-clustering/create-speaker-mixtures.zip>

To make the model comparable, we set all the neural network models to have the same depth and similar size. We use the VGG-motivated CNN layers and bidirectional long-short term memory recurrent neural networks with projection (BLSTMP) as the encoder. The total depth of the encoder is 5, namely two CNN blocks and three layer BLSTMP layers. For all models, the decoder network has 1 layer of unidirectional long-short term memory network (LSTM) with 300 cells.

During decoding, we combined both the joint CTC/attention score and the pretrained word-level recurrent neural network language model (RNNLM) score, which had 1-layer LSTM with 1000 cells and was trained on the transcriptions from WSJ SI284, in a shallow fusion manner. We set the beam width to be 30. The interpolation factor  $\lambda$  we used during decoding was 0.3, and the weight for RNNLM was 1.0.

### 3.2. Performance of baseline systems

In this section, we describe the performance of the baseline E2E ASR systems on multi-speaker mixed speech. The first baseline system is the joint CTC/attention-based encoder-decoder network for single-speaker speech trained on WSJ corpus, whose performance is 0.9% in terms of CER and 1.9% in terms of WER on the eval92\_5k test set with the closed vocabulary. In the encoder, there are 3 layers of BLSTMP following the CNN and each BLSTMP layer has 1024 memory cells in each direction. The second baseline system is the joint CTC/attention-based encoder-decoder network for multi-speaker speech. The 2-layer CNN is used as the  $\text{Encoder}_{\text{Mix}}$ . The depth of the following BLSTMP layers is also 3 including 1 layer of BLSTMP as the  $\text{Encoder}_{\text{SD}}$  and 2 layers of BLSTMP as the  $\text{Encoder}_{\text{Rec}}$ . The attention-decoder in the multi-speaker system is shared among representations  $\mathbf{G}^s$ , which is of the same architecture with single-speaker system. The results are shown in Table 1.

Model	dev CER	eval CER
single-speaker	79.13	76.52
multi-speaker [25]	n/a	13.7
multi-speaker	15.14	12.20
Model	dev WER	eval WER
single-speaker	113.47	112.21
multi-speaker	24.90	20.43

**Table 1.** Performance (Avg. CER & WER) (%) on 2-speaker mixed WSJ corpus. Comparison between End-to-End single-speaker and multi-speaker joint CTC/attention-based encoder-decoder systems

In the case of single-speaker, the CER and WER is measured by comparing the output against the reference labels of both speakers. From the table, we can see that the speech recognition system designed for multi-speaker can improve the performance for the overlapped speech significantly, leading to more than 80.0% relative reduction on both average CER and WER. As a comparison, we also include the CER result from [25] in the table, and it shows that the newly constructed end-to-end multi-speaker system without pretraining in this work can achieve better performance.

### 3.3. Performance of speaker parallel attention with scheduled sampling

In this section we report the results of the evaluation of our proposed methods. The first method is the speaker parallel attention, introducing independent attention modules for each speaker source

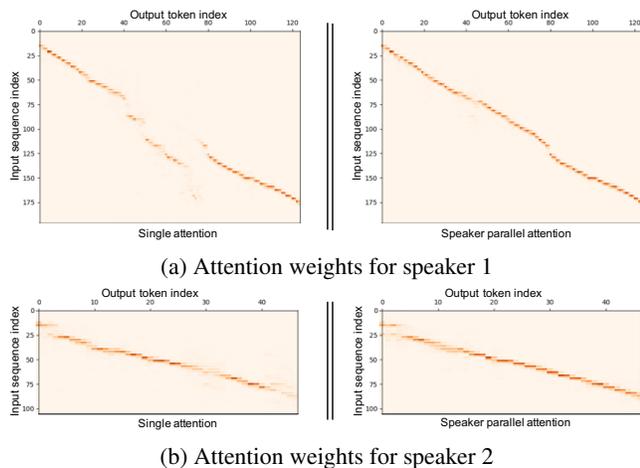
instead of using a shared attention module. The rest of the network is kept the same as the baseline multi-speaker model, containing a 2-layer CNN Encoder<sub>Mix</sub>, 1-layer BLSTMP Encoder<sub>SD</sub>, a 2-layer BLSTMP Encoder<sub>Rec</sub>, and a shared 1-layer LSTM as the decoder network. The performance is illustrated in the Table 2. The speaker parallel attention module reduces the average CER by 9% and average WER by 8% relatively. From the results we can tell that the CER is high, so the gap is large between the training and inference using the teacher-forcing fashion. Thus we adopted the scheduled sampling method with probability  $p = 0.2$  in Eq. 15, which lead to a further improvement in performance. Finally, the system using both speaker parallel attention and scheduled sampling can obtain relative  $\sim 10.0\%$  reduction on both CER and WER on the evaluation set.

Model	dev CER	eval CER
multi-speaker (baseline)	15.14	12.20
+ speaker parallel attention	14.80	11.11
++ scheduled sampling	<b>14.78</b>	<b>10.93</b>

Model	dev WER	eval WER
multi-speaker (baseline)	24.90	20.43
+ speaker parallel attention	24.88	18.76
++ scheduled sampling	<b>24.52</b>	<b>18.44</b>

**Table 2.** Performance (Avg. CER & WER) (%) on 2-speaker mixed WSJ corpus. Comparison between End-to-End multi-speaker joint CTC/attention-based encoder-decoder systems



**Fig. 2.** Visualization of the attention weights sequences for two overlapped speakers. The left part is from the previous single-attention multi-speaker end-to-end model and the right part is from the proposed speaker-parallel-attention multi-speaker end-to-end model.

We show the visualization of the attention weights sequences for two overlapped speakers, generated by the baseline single-attention multi-speaker end-to-end model and the proposed speaker-parallel-attention multi-speaker end-to-end model individually. The horizontal axis represents the output token sequence and the vertical axis represents the input sequence to the attention module. The left parts of Figures.2 (a) and (b) show the attention weights for speaker 1 and speaker 2 generated by the previous single-attention model. The right parts show the attention weights generated by the proposed speaker-parallel-attention model. We can observe that the right parts are more smooth and clear, and the attention weights are more concentrated. This observation conforms with the characteristics of alignments between output sequence and input sequence for

speech recognition, and further shows the superiority of the proposed speaker parallel attentions.

### 3.4. Comparison with previous work

We then compared our work with other related work. We trained and tested our model on wsj0-2mix dataset that was first used in [16]. Table 3 shows the WER results of hybrid systems including PIT-ASR [23], DPCL-based speech separation with Kaldi-based ASR [17], and the end-to-end systems constructed in [25] and ours in this paper. These were evaluated under the same evaluation data and metric as in [17] based on the wsj0-2mix. Noted that the model in [25] was trained on a different, larger training dataset than that used in other experiments. From Table. 3, we can observe that our new system constructed by the proposed methods in this paper is significantly better than the others.

Model	Avg. WER
DPCL+ASR [17]	30.8
PIT-ASR [23]	28.2
End-to-end ASR (Char/Word-LM) [25]	28.2
Proposed End-to-end ASR with SPA (Word LM)	<b>25.4</b>

**Table 3.** WER (%) on 2-speaker mixed **WSJ0** corpus. The comparison is done between our proposed end-to-end ASR with speaker parallel attention (SPA) and previous works including DPCL+ASR, PIT-ASR and end-to-end ASR systems.

## 4. CONCLUSION

In this paper, we have introduced a state-of-the-art end-to-end multi-speaker speech recognition system under the joint CTC/attention-based encoder-decoder framework. More specifically, a new neural network architecture enabled us to train the model from random initialization. And we adopted the speaker parallel attention module and scheduled sampling to improve performance over the previous end-to-end multi-speaker speech recognition system. The experiments on the 2-speaker mixed speech recognition show that the proposed new strategy can obtain a relative  $\sim 10.0\%$  improvement on CER and WER reduction.

## 5. ACKNOWLEDGEMENT

Xuankai Chang, Yanmin Qian and Kai Yu were supported by the China NSFC project (No. 61603252), the China NSFC project (No. U1736202) and the Shanghai Sailing Program No. 16YF1405300.

We are also grateful to Matthew Maciejewski and Tian Tan for their comments on an earlier version of the manuscript.

## 6. REFERENCES

- [1] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman 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, pp. 82–97, 2012.
- [2] Tara N Sainath, Abdel-rahman Mohamed, Brian Kingsbury, and Bhuvana Ramabhadran, “Deep convolutional neural networks for LVCSR,” in *IEEE (ICASSP)*, 2013, pp. 8614–8618.

- [3] Wayne Xiong, Jasha Droppo, Xuedong Huang, Frank Seide, Mike Seltzer, Andreas Stolcke, Dong Yu, and Geoffrey Zweig, “The Microsoft 2016 conversational speech recognition system,” in *IEEE (ICASSP)*, 2017, pp. 5255–5259.
- [4] Zhehuai Chen and Jasha Droppo, “Sequence modeling in unsupervised single-channel overlapped speech recognition,” in *IEEE (ICASSP)*, 2018.
- [5] Suyoun Kim, Takaaki Hori, and Shinji Watanabe, “Joint ctc-attention based end-to-end speech recognition using multi-task learning,” in *IEEE (ICASSP)*, 2017, pp. 4835–4839.
- [6] Shinji Watanabe, Takaaki Hori, Suyoun Kim, John R. Hershey, and Tomoki Hayashi, “Hybrid ctc/attention architecture for end-to-end speech recognition,” *J. Sel. Topics Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [7] Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplín, Jahn Heymann, Matthew Wiesner, Nanxin Chen, et al., “Espnet: End-to-end speech processing toolkit,” *arXiv preprint arXiv:1804.00015*, 2018.
- [8] Alex Graves and Navdeep Jaitly, “Towards end-to-end speech recognition with recurrent neural networks,” in *ICML*, 2014, pp. 1764–1772.
- [9] Yajie Miao, Mohammad Gowayyed, and Florian Metze, “Eesen: End-to-end speech recognition using deep rnn models and wfst-based decoding,” in *IEEE Workshop on (ASRU)*, 2015, pp. 167–174.
- [10] Jan Chorowski, Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, “End-to-end continuous speech recognition using attention-based recurrent nn: first results,” *arXiv preprint arXiv:1412.1602*, 2014.
- [11] William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *IEEE (ICASSP)*, 2016, pp. 4960–4964.
- [12] 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.
- [13] Martin Cooke, John R Hershey, and Steven J Rennie, “Monaural speech separation and recognition challenge,” *Computer Speech & Language*, vol. 24, no. 1, pp. 1–15, 2010.
- [14] Jon Barker, Shinji Watanabe, Emmanuel Vincent, and Jan Trmal, “The fifth ‘chime’ speech separation and recognition challenge: Dataset, task and baselines,” *arXiv preprint arXiv:1803.10609*, 2018.
- [15] Yanmin Qian, Chao Weng, Xuankai Chang, Shuai Wang, and Dong Yu, “Past review, current progress, and challenges ahead on the cocktail party problem,” *Frontiers of Information Technology & Electronic Engineering*, vol. 19, no. 1, pp. 40–63, Jan 2018.
- [16] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *IEEE (ICASSP)*, 2016, pp. 31–35.
- [17] Yusuf Isik, Jonathan Le Roux, Zhuo Chen, Shinji Watanabe, and John R. Hershey, “Single-channel multi-speaker separation using deep clustering,” in *(INTERSPEECH)*, 2016, pp. 545–549.
- [18] Dong Yu, Morten Kolbæk, Zheng-Hua Tan, and Jesper Jensen, “Permutation invariant training of deep models for speaker-independent multi-talker speech separation,” in *IEEE (ICASSP)*, 2017, pp. 241–245.
- [19] Morten Kolbæk, Dong Yu, Zheng-Hua Tan, and Jesper Jensen, “Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks,” *IEEE/ACM (TASLP)*, vol. 25, no. 10, pp. 1901–1913, 2017.
- [20] Dong Yu, Xuankai Chang, and Yanmin Qian, “Recognizing multi-talker speech with permutation invariant training,” in *(INTERSPEECH)*, 2017, pp. 2456–2460.
- [21] Z. Chen, J. Droppo, J. Li, and W. Xiong, “Progressive joint modeling in unsupervised single-channel overlapped speech recognition,” *IEEE/ACM (TASLP)*, vol. 26, no. 1, pp. 184–196, Jan 2018.
- [22] Xuankai Chang, Yanmin Qian, and Dong Yu, “Monaural multi-talker speech recognition with attention mechanism and gated convolutional networks,” in *(INTERSPEECH)*, 2018, pp. 1586–1590.
- [23] Yanmin Qian, Xuankai Chang, and Dong Yu, “Single-channel multi-talker speech recognition with permutation invariant training,” *Speech Communication*, vol. 104, pp. 1–11, 2018.
- [24] Shane Settle, Jonathan Le Roux, Takaaki Hori, Shinji Watanabe, and John R Hershey, “End-to-end multi-speaker speech recognition,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4819–4823.
- [25] 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) (Volume 1: Long Papers)*, 2018, pp. 2620–2630.
- [26] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [27] Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer, “Scheduled sampling for sequence prediction with recurrent neural networks,” in *(NIPS) - Volume 1*, 2015, pp. 1171–1179.
- [28] Takaaki Hori, Shinji Watanabe, and John Hershey, “Joint ctc/attention decoding for end-to-end speech recognition,” in *(ACL) (Volume 1: Long Papers)*, 2017, vol. 1, pp. 518–529.
- [29] Xuankai Chang, Yanmin Qian, and Dong Yu, “Adaptive permutation invariant training with auxiliary information for monaural multi-talker speech recognition,” in *IEEE (ICASSP)*, 2018.
- [30] Tian Tan, Yanmin Qian, and Dong Yu, “Knowledge transfer in permutation invariant training for single-channel multi-talker speech recognition,” in *IEEE (ICASSP)*, 2018.
- [31] 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,” .
- [32] Albert Zeyer, Kazuki Irie, Ralf Schlüter, and Hermann Ney, “Improved training of end-to-end attention models for speech recognition,” *arXiv preprint arXiv:1805.03294*, 2018.