

# Stream Attention-Based Multi-Array End-to-End Speech Recognition

Wang, X.; Li, R.; Mallidi, S.H.; Hori, T.; Watanabe, S.; Hermansky, H.

TR2019-005 March 29, 2019

## Abstract

Automatic Speech Recognition (ASR) using multiple microphone arrays has achieved great success in the far-field robustness. Taking advantage of all the information that each array shares and contributes is crucial in this task. Motivated by the advances of joint Connectionist Temporal Classification (CTC)/attention mechanism in the End-to-End (E2E) ASR, a stream attention-based multi-array framework is proposed in this work. Microphone arrays, acting as information streams, are activated by separate encoders and decoded under the instruction of both CTC and attention networks. In terms of attention, a hierarchical structure is adopted. On top of the regular attention networks, stream attention is introduced to steer the decoder toward the most informative encoders. Experiments have been conducted on AMI and DIRHA multi-array corpora using the encoder-decoder architecture. Compared with the best single-array results, the proposed framework has achieved relative Word Error Rates (WERS) reduction of 3.7% and 9.7% in the two datasets, respectively, which is better than conventional strategies as well.

*IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.



# STREAM ATTENTION-BASED MULTI-ARRAY END-TO-END SPEECH RECOGNITION

Xiaofei Wang<sup>1,‡</sup>, Ruizhi Li<sup>1,‡</sup>, Sri Harish Mallidi<sup>2</sup>, Takaaki Hori<sup>3</sup>, Shinji Watanabe<sup>1</sup>, Hynek Hermansky<sup>1</sup>

<sup>1</sup>The Johns Hopkins University, <sup>2</sup>Amazon, <sup>3</sup>Mitsubishi Electric Research Laboratories (MERL)  
{xiaofeiwang, ruizhili, shinjiw, hynek}@jhu.edu, mallidih@amazon.com, thori@merl.com

## ABSTRACT

Automatic Speech Recognition (ASR) using multiple microphone arrays has achieved great success in the far-field robustness. Taking advantage of all the information that each array shares and contributes is crucial in this task. Motivated by the advances of joint Connectionist Temporal Classification (CTC)/attention mechanism in the End-to-End (E2E) ASR, a stream attention-based multi-array framework is proposed in this work. Microphone arrays, acting as information streams, are activated by separate encoders and decoded under the instruction of both CTC and attention networks. In terms of attention, a hierarchical structure is adopted. On top of the regular attention networks, stream attention is introduced to steer the decoder toward the most informative encoders. Experiments have been conducted on AMI and DIRHA multi-array corpora using the encoder-decoder architecture. Compared with the best single-array results, the proposed framework has achieved relative Word Error Rates (WERs) reduction of 3.7% and 9.7% in the two datasets, respectively, which is better than conventional strategies as well.

**Index Terms**— Multiple Microphone Array, End-to-End Speech Recognition, Joint CTC/Attention, Stream Attention

## 1. INTRODUCTION

Far-field ASR using multiple microphone arrays has been a widely adopted strategy in the speech processing community. Individually, the microphone array is able to bring a substantial performance improvement with algorithms such as beamforming [1] and masking [2]. However, what kind of information can be extracted from each array and how to make multiple arrays work in cooperation are still challenging. Without any prior knowledge of speaker-array distance or video monitoring, it is difficult to figure out which array carries more reliable information or is less corrupted.

According to the reports from the CHiME-5 challenge [3], which targets the problem of multi-array conversational speech recognition in home environments, the common ways of utilizing multiple arrays in the hybrid ASR system are finding the ones with higher Signal-to-Noise/Interference Ratio (SNR/SIR) [4] or fusing the decoding results by voting for the most confident words [5], e.g. ROVER [6]. Similar to our previous work [7], combination using the classifier’s posterior probabilities followed by lattice generation has been an alternative approach [8]. The posteriors from the well-trained classifier decorrelate the input features, but reserve more distinctive speech information than the words after the full decoding stage. In terms of the combination strategy, ASR performance monitors have been designed [9], resulting in a process of stream confidence generation, guiding the linear fusion of array streams.

Recently, E2E ASR has attracted attention in the research field. The E2E system is developed to directly transcribe human speech

into text. It integrates disjoint modules, developed from traditional hybrid methods, into one single Deep Neural Network (DNN) which can be trained from scratch. The attention-based structure [10, 11] solves the ASR problem as a sequence mapping by using an encoder-decoder architecture. Coupled with a CTC network [12, 13], the joint model [14, 15, 16] outperforms the attention-based ASR by addressing misalignment issues. While most of the E2E ASR studies engage in single-channel task or multi-channel task from one microphone array [17, 18, 19, 20], research on multi-array scenario is still unexplored within the E2E framework.

In this work, we propose an attention-based multi-array E2E architecture – the joint CTC/Attention model with the hierarchical attention mechanism, inspired by our original work [21] done at JSALT 2018, to solve the aforementioned problem. The framework has highlights as follows:

1. The output of each microphone array is modeled by a separate encoder. Multiple encoders with the same configuration act as the acoustic models for individual arrays.
2. The hierarchical attention mechanism [22, 23, 24] was introduced to dynamically combine knowledge from parallel streams. We adopt this network in multi-array scheme, where the stream-level fusion is employed on top of the per-encoder attention mechanisms.
3. Each encoder is associated with a CTC network to guide the frame-wise alignment process for each array to potentially achieve a better performance.

The remainder of this paper is organized as follows: Section 2 reviews previous work. The proposed multi-stream framework is presented in Section 3, followed by experiments and analysis in Section 4. In the end, the conclusion is given Section 5.

## 2. PRIOR WORK

### 2.1. Conventional Multi-Array ASR

In our previous work, we proposed a stream attention framework to improve the far-field performance in the hybrid approach, using distributed microphone array(s) [7]. Specifically, we generated more reliable Hidden Markov Model (HMM) state posterior probabilities by linearly combining the posteriors from each array stream, under the supervision of the ASR performance monitors.

In general, the posterior combination strategy outperformed conventional methods, such as signal-level fusion and the word-level technique ROVER [6], in the prescribed multi-array configuration. Accordingly, stream attention weights estimated from the de-correlated intermediate features should be more reliable. We adopt this assumption in the following context.

<sup>‡</sup> Both authors contributed equally to this work.

## 2.2. Joint CTC/Attention Architecture for End-to-End ASR

The joint CTC/Attention model for the E2E ASR outperforms ordinary attention-based ones by solving the misalignment issues between the speech and labels [14, 15, 16]. It digests the advantages of both CTC and attention-based model through a multi-task learning mechanism and joint decoding. Accordingly, the E2E model maps  $T$ -length acoustic features  $X = \{\mathbf{x}_t \in \mathbb{R}^D | t = 1, 2, \dots, T\}$  in  $D$  dimensional space to an  $L$ -length letter sequence  $C = \{c_l \in \mathcal{U} | l = 1, 2, \dots, L\}$  where  $\mathcal{U}$  is a set of distinct letters.

The encoder is shared by both attention and CTC networks. Typical Bidirectional Long Short-Term Recurrent (BLSTM) layers are utilized to model the temporal dependencies of the input sequence. The frame-wise hidden vector  $\mathbf{h}_t$  at frame  $t$  is derived by encoding the full input sequence  $X$ :

$$\mathbf{h}_t = \text{Encoder}(X) \quad (1)$$

For the attention-based encoder-decoder model, the letter-wise context vector  $\mathbf{r}_l$  is formed as a weighted summation of frame-wise hidden vectors  $\mathbf{h}_t$  using a content-based attention mechanism:

$$\mathbf{r}_l = \sum_{t=1}^T a_{lt} \mathbf{h}_t, \quad a_{lt} = \text{ContentAttention}(\mathbf{q}_{l-1}, \mathbf{h}_t) \quad (2)$$

where  $\mathbf{q}_{l-1}$  is the previous decoder state, and  $a_{lt}$  is the attention weight, a soft-alignment of  $\mathbf{h}_t$  for  $c_l$ . An LSTM-based decoder network predicts the next letter based on  $\mathbf{r}_l$  and the previous prediction.

The objective function to be maximized is as follows:

$$\mathcal{L} = \lambda \log p_{ctc}(C|X) + (1 - \lambda) \log p_{att}^\dagger(C|X) \quad (3)$$

where the joint objective is a logarithmic linear combination of the CTC and attention training objectives, i.e.,  $p_{ctc}(C|X)$  and  $p_{att}^\dagger(C|X)$ , respectively. The attention  $p_{att}(C|X)$  is approximated during training as  $p_{att}^\dagger(C|X)$ , where the probability of a prediction is conditioned on previous true labels.  $0 \leq \lambda \leq 1$  is a trade-off parameter satisfying. In the decoding phase, the joint CTC/Attention model performs a label-synchronous beam search which jointly predicts the next character. The most probable letter sequence  $\hat{C}$  given the speech input  $X$  is computed as:

$$\hat{C} = \arg \max_{C \in \mathcal{U}^*} \{ \lambda \log p_{ctc}(C|X) + (1 - \lambda) \log p_{att}^\dagger(C|X) + \gamma \log p_{lm}(C) \} \quad (4)$$

where an external Recurrent Neural Network Language Model (RNNLM) probability  $\log p_{lm}(C)$  is added with a scaling factor  $\gamma$ .

## 3. MULTI-ARRAY END-TO-END MODEL

In this section, we present the stream attention based E2E framework [21] for the multi-array ASR task. A hierarchical attention scheme is introduced within the CTC/Attention joint training and decoding mechanism. For simplicity to understand the framework, we focus on the two-array architecture, which is shown in Fig.1.

### 3.1. Multi-Array Architecture with Stream Attention

The proposed architecture has two encoders, with each mapping the speech features of a single array to higher level representations  $\mathbf{h}_t^i$ , where we denote  $i \in \{1, 2\}$  as the index for  $\text{Encoder}_i$  corresponding to array  $i$ . Note that  $\text{Encoder}_1$  and  $\text{Encoder}_2$  have the same configurations receiving parallel speech data collected from multiple microphone arrays. Convolution Neural Networks (CNN) are often




Fig. 1. Multi-Stream Architecture Using Two Microphone-Arrays.

used together with BLSTM layers on top to extract frame-wise hidden vectors. We explore two types of encoder structures: BLSTM (RNN-based) and VGGBLSTM (CNN-RNN-based) [25]:

$$\mathbf{h}_t^i = \text{Encoder}_i(X), \quad i \in \{1, 2\} \quad (5)$$

$$\text{Encoder}_i(\cdot) = \text{BLSTM}(\cdot) \quad \text{or} \quad \text{VGGBLSTM}(\cdot) \quad (6)$$

Note that the BLSTM encoders are equipped with an additional projection layer after each BLSTM layer. In both encoder architectures, subsampling factor  $s = 4$  is applied to decrease the computational cost. Specially, the convolution layers of the VGGBLSTM encoder downsamples the input features by a factor of 4 so that there is no subsampling in the recurrent layers.

In the multi-stream setting, one inherent problem is that the contribution of each stream (array) changes dynamically. Specially, when one of the streams takes corrupted audio, the network should be able to pay more attention to other streams for the purpose of robustness. Inspired by the advances of linear posterior combination [7] and a hierarchical attention fusion [22, 23, 24], a stream-level fusion on the letter-wise context vector is introduced in this work to achieve the goal of encoder selectivity. The letter-wise context vectors,  $\mathbf{r}_l^1$  and  $\mathbf{r}_l^2$ , from individual encoders are computed similar to Eq. (2):

$$\mathbf{r}_l^i = \sum_{t=1}^{T/4} a_{lt}^i \mathbf{h}_t^i, \quad i \in \{1, 2\} \quad (7)$$

where the summation is performed from 1 to  $T/4$  due to subsampling. The fusion context vector  $\mathbf{r}_l$  is obtained as a combination of  $\mathbf{r}_l^1$  and  $\mathbf{r}_l^2$  as illustrated:

$$\mathbf{r}_l = \beta_{l1} \mathbf{r}_l^1 + \beta_{l2} \mathbf{r}_l^2 \quad (8)$$

$$\beta_{li} = \text{ContentAttention}(\mathbf{q}_{l-1}, \mathbf{r}_l^i), \quad i = 1, 2 \quad (9)$$

The stream-level attention weights  $\beta_{l1}$  and  $\beta_{l2}$  are estimated according to the feedback from the previous decoder state,  $\mathbf{q}_{l-1}$ , and context vectors,  $\mathbf{r}_l^1$  and  $\mathbf{r}_l^2$ , from individual encoders. The fusion context vector is then fed into the decoder to predict the next letter.

In comparison to fusion on frame-wise hidden vectors  $\mathbf{h}_t^i$ , stream-level fusion can deal with temporal misalignment from multiple arrays at the stream level. Furthermore, adding an extra microphone array  $j$  could be simply implemented with an additional term  $\beta_{lj} \mathbf{r}_t^j$  in Eq.(8).

### 3.2. Training and Decoding with Per-encoder CTC

We assign each encoder with a separate CTC network. During multi-task training and joint decoding, we follow the similar formulas depicted by Eq. (3) and Eq. (4). The only difference is that we have per-encoder CTC objective to compute the loss:

$$\log p_{ctc}(C|X) = \frac{1}{2} \lambda (\log p_{ctc_1}(C|X) + \log p_{ctc_2}(C|X)), \quad (10)$$

where the equal weight is assigned to each CTC network.

## 4. EXPERIMENT AND DISCUSSION

### 4.1. Dataset (AMI and DIRHA) Description

The AMI Meeting Corpus consists of 100 hours of far-field recordings from 3 meeting rooms (Edinburgh, Idiap and TNO Room) [26]. The recordings use a range of signals synchronized to a common time line. There are two arrays placed in each meeting room to record the sentences, with one 10 cm radius circular array between the speakers consisting of 8 omni-directional microphones. The setups of the second microphone array are different among the rooms, detailed by Table 1.

The DIRHA dataset was collected in a real apartment setting with typical domestic background noise and reverberation [27]. In the configuration, a total of 32 microphones were placed in the living-room (26 microphones) and in the kitchen (6 microphones). The microphone network consists of 2 circular arrays of 6 microphones (located on the ceiling of the living-room and the kitchen), a linear array of 11 sensors (located in the living-room) and 9 microphones distributed on the living-room walls. During the recording, the speaker was asked to move to a different position and take a different orientation after reading several sentences.

In both datasets, we chose two microphone arrays as parallel streams (noted by Str1 and Str2) to train the proposed E2E system, which is also shown by Table 1. For each microphone array, all the simulations or recordings were synthesized into the single channel using delay-and-sum (DS) beamforming with the BeamformIT Toolkit [28]. The AMI training set consists of 81 hours of speech. The development (Dev) and evaluation (Eval) set respectively contain 9 hours of meeting recordings. We used Dev set for cross validation and Eval set for testing. Contaminated version of the original WSJ (Wall Street Journal) corpus is used for DIRHA training. Two streams were generated using the WSJ0 and WSJ1 clean utterances convolved by the circular array impulse responses and the linear ones, respectively. Recorded noises were added as well. We used the DIRHA Simulation set (generated via the same way as training data) for cross validation and DIRHA Real set for testing, which consisted of 3 Male and 3 Female native US speakers uttering 409 WSJ sentences.

All the experiments were implemented by ESPnet, an end-to-end speech processing toolkit [29] with the configuration as described in Table 2:

**Table 1.** Description of the array configuration in the two-stream E2E experiments.

Dataset	Str1 (Stream 1)	Str2 (Stream 2)
AMI	8-mic Circular Array	Edinburgh: 8-mic Circular Array Idiap: 4-mic Circular Array TNO: 10-mic Linear Array
DIRHA	6-mic Circular Array	11-mic Linear Array

**Table 2.** Experimental configuration

Feature	
Single Stream	80-dim fbank + 3-dim pitch
Multi Stream	Array <sub>1</sub> :80+3; Array <sub>2</sub> :80+3
Model	
Encoder type	BLSTM or VGGBLSTM
Encoder layers	BLSTM:4; VGGBLSTM[25]:6(CNN)+4
Encoder units	320 cells (BLSTM layers)
(Stream) Attention	Content-based
Decoder type	1-layer 300-cell LSTM
CTC weight $\lambda$ (train)	AMI:0.5; DIRHA:0.2
CTC weight $\lambda$ (decode)	AMI:0.3; DIRHA:0.3
RNN-LM	
Type	Look-ahead Word-level RNNLM [30]
Train data	AMI:AMI; DIRHA:WSJ0-1+extra WSJ text data
LM weight $\gamma$	AMI:0.5; DIRHA:1.0

### 4.2. Results

We define two kinds of E2E architectures in these results discussions: single-stream architecture, which has only one encoder without stream attention and multi-stream architecture, which has several encoders with each corresponding to one microphone array and has stream attention mechanism as well.

#### 4.2.1. Single-array results

First of all, we explore the ASR performance for the individual array (single stream). As illustrated in Table 3, the single stream system with the VGGBLSTM based encoder outperforms the one with BLSTM encoder, both in Character Error Rate (CER) and WER. Joint training of CTC and attention based model helps since CTC can enforce the monotonic behavior of attention alignments, rather than merely estimating the desired alignment for long sequence. With the RNNLM, we can see a dramatical decrease of the WERs on both datasets. The Str1 WERs of AMI Eval and DIRHA Real are 56.9% and 35.1%, respectively. For simplicity, we only keep the CTC/Attention based single-stream results with RNNLM for Str2 since the same trend can be found and only the WER will be compared in the following results.

#### 4.2.2. Multi-array results

As shown in Table 4, the proposed stream attention framework achieves 3.7% (56.9 to 54.9) and 9.7% (35.1 to 31.7) relative WERs reduction on AMI and DIRHA datasets, respectively. Hierarchical attention plays a role that emphasizing the more reliable stream.

**Table 3.** Exploration of best encoder and decoding strategy for single-stream E2E model.

Model (Single Stream)	AMI Eval		DIRHA Real	
	CER	WER	CER	WER
<i>BLSTM</i> (Str1)				
Attention	45.1	60.9	42.7	68.7
+ CTC	41.7	63.0	38.5	74.8
+ Word RNNLM	41.7	59.1	29.4	47.4
<i>VGGBLSTM</i> (Str1)				
Attention	43.2	59.7	39.5	71.4
+ CTC	40.2	62.0	30.1	61.8
+ Word RNNLM	<b>39.6</b>	<b>56.9</b>	<b>21.2</b>	<b>35.1</b>
<i>VGGBLSTM</i> (Str2)	<b>45.6</b>	<b>64.0</b>	<b>22.5</b>	<b>38.4</b>

In addition, we compare the multi-stream framework with conventional strategies using single-stream system trained by the Fbank and pitch features, either concatenated by the Str1 and Str2 features or extracted from the speech audio through alignment and average between the streams. The multi-stream framework outperforms the others. To explain the improvement is not from the boost of the number of model parameters, we doubled the BLSTM layers (4 to 8) in the VGGBLSTM encoder and train the single-stream CTC/Attention system with a comparable amount of parameters (33.7M vs 31.6M). Our system still has strong competitiveness.


**Table 4.** WER(%) Comparison between the proposed multi-stream approach and alternative single-stream strategies.

Encoder <i>VGGBLSTM</i> (Att + CTC + RNNLM)	#Param	AMI Eval	DIRHA Real
Single-stream model			
Concatenating Str1&Str2	23.3M	56.7	33.5
WAV alignment and average	26.2M	56.7	43.5
+ model parameter extension	33.7M	56.9	39.6
Multi-stream model			
Proposed framework	31.6M	<b>54.9</b>	<b>31.7</b>

During the inference stage of the multi-stream model, we examine how the stream attention weights change once one of the streams is corrupted by noise. Fig.2 shows an example in the DIRHA Real set that whether the input features of Str1 is affected by an additive Gaussian noise with zero mean and unit variance. After the corruption, the alignment between characters and acoustic frames of Str1 becomes blurred (Fig.2(c)), indicating that the information from Str1 should be less trusted. Therefore, as expected, a positive shift of the attention weights for Str2 can be observed (upper line in Fig.2(e)).

#### 4.2.3. Comparison with hybrid system

Table 5 shows the comparison between the proposed E2E framework and the conventional hybrid ASR approach. In [7], we designed three scenarios using different subsets from the 32 microphones and 2 arrays in the DIRHA dataset. Our proposed DNN posterior combination approach and ROVER technique could relatively reduce the



**Fig. 2.** Comparison of the alignments between characters (y-axis) and acoustic frames (x-axis) before ((a) Str1; (b) Str2) and after ((c) Str1; (d) Str2) noise corruption of Str1. (e) shows the attention weight shift of Str2 between two cases (x-axis is the letter sequence).

WER of the hybrid system by 7.2% and 5.8% respectively, if we average the WERs of the Real test sets among three cases. Meanwhile, a relative 9.7% WER reduction has already been achieved in the stream attention-based two-stream E2E system, even though we have less number of streams (two) than the hybrid one (six). Ignoring the WER gap between the hybrid and E2E ASR systems, we still believe that the proposed E2E approach has much potential to do better with more array streams.

**Table 5.** WER(s) Comparison between the hybrid and end-to-end system on DIRHA dataset. #Num denotes the number of streams.

System	#Num	Method	Best Stream	WER
Hybrid	6	post. comb.	29.2	27.1 (7.2%)
	6	ROVER	29.2	27.5 (5.8%)
E2E	2	proposed	35.1	31.7 (9.7%)

## 5. CONCLUSION

In this paper, we presented a multi-stream End-to-End ASR framework targeting the distributed microphone array situation. Stream attention was achieved through a hierarchical connection between the decoder and encoders, with each modeling one array into higher-level representations. Thanks to the success of joint training of per-encoder CTC and attention, substantial WER reduction was shown in both AMI and DIRHA corpora, demonstrating the potentials of the proposed architecture. For further research, an extension to more streams efficiently and exploration of schedule training of the encoders would be interesting.

## 6. REFERENCES

- [1] Emmanuel Vincent, Shinji Watanabe, Aditya Arie Nugraha, Jon Barker, and Ricard Marxer, “An analysis of environment, microphone and data simulation mismatches in robust speech recognition,” *Computer Speech & Language*, vol. 46, pp. 535–557, 2017.
- [2] Ziteng Wang, Xiaofei Wang, Xu Li, Qiang Fu, and Yonghong Yan, “Oracle performance investigation of the ideal masks,” in *IWAENC 2016*. IEEE, 2016, pp. 1–5.
- [3] Jon Barker, Shinji Watanabe, Emmanuel Vincent, and Jan Trmal, “The fifth ‘chime’ speech separation and recognition challenge: Dataset, task and baselines,” in *Interspeech 2018*, 2018, pp. 1561–1565.
- [4] Jun Du et al., “The ustc-ifytek systems for chime-5 challenge,” in *CHiME-5*, 2018.
- [5] Naoyuki Kanda et al., “The hitachi/jhu chime-5 system: Advances in speech recognition for everyday home environments using multiple microphone arrays,” in *CHiME-5*, 2018.
- [6] Jonathan G Fiscus, “A post-processing system to yield reduced word error rates: Recognizer output voting error reduction (rover),” in *ASRU 1997*. IEEE, 1997, pp. 347–354.
- [7] Xiaofei Wang, Ruizhi Li, and Hynek Hermansky, “Stream attention for distributed multi-microphone speech recognition,” in *Interspeech 2018*, 2018, pp. 3033–3037.
- [8] Feifei Xiong et al., “Channel selection using neural network posterior probability for speech recognition with distributed microphone arrays in everyday environments,” in *CHiME-5*, 2018.
- [9] Sri Harish Mallidi, Tetsuji Ogawa, and Hynek Hermansky, “Uncertainty estimation of dnn classifiers,” in *ASRU 2015*. IEEE, 2015, pp. 283–288.
- [10] William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *ICASSP 2016*. IEEE, 2016, pp. 4960–4964.
- [11] Jan K Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio, “Attention-based models for speech recognition,” in *NIPS 2015*, 2015, pp. 577–585.
- [12] Alex Graves and Navdeep Jaitly, “Towards end-to-end speech recognition with recurrent neural networks,” in *ICML 2014*, 2014, pp. 1764–1772.
- [13] Yajie Miao, Mohammad Gowayed, and Florian Metze, “EESSEN: End-to-end speech recognition using deep RNN models and WFST-based decoding,” in *ASRU 2015*, 2015, pp. 167–174.
- [14] Suyoun Kim, Takaaki Hori, and Shinji Watanabe, “Joint CTC-attention based end-to-end speech recognition using multi-task learning,” in *ICASSP 2017*, 2017, pp. 4835–4839.
- [15] Takaaki Hori, Shinji Watanabe, Yu Zhang, and William Chan, “Advances in joint CTC-attention based end-to-end speech recognition with a deep CNN encoder and RNN-LM,” in *Interspeech 2017*, 2017.
- [16] Shinji Watanabe, Takaaki Hori, Suyoun Kim, John R Hershey, and Tomoki Hayashi, “Hybrid ctc/attention architecture for end-to-end speech recognition,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [17] Tsubasa Ochiai, Shinji Watanabe, Takaaki Hori, John R Hershey, and Xiong Xiao, “Unified architecture for multichannel end-to-end speech recognition with neural beamforming,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1274–1288, 2017.
- [18] Stefan Braun, Daniel Neil, Jithendar Anumula, Enea Ceolini, and Shih-Chii Liu, “Multi-channel attention for end-to-end speech recognition,” in *Interspeech 2018*, 2018, pp. 17–21.
- [19] Tsubasa Ochiai, Shinji Watanabe, Takaaki Hori, and John R Hershey, “Multichannel end-to-end speech recognition,” *arXiv preprint arXiv:1703.04783*, 2017.
- [20] Suyoun Kim, Ian Lane, S. Kim, and I. Lane, “End-to-end speech recognition with auditory attention for multi-microphone distance speech recognition,” in *Interspeech 2017*, 2017, pp. 3867–3871.
- [21] Ruizhi Li, Xiaofei Wang, Sri Harish Mallidi, Takaaki Hori, Shinji Watanabe, and Hynek Hermansky, “Multi-encoder multi-resolution framework for end-to-end speech recognition,” *arXiv preprint arXiv:1811.04897*, 2018.
- [22] Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy, “Hierarchical attention networks for document classification,” in *NAACL HLT*, 2016, pp. 1480–1489.
- [23] Chiori Hori, Takaaki Hori, Teng-Yok Lee, Ziming Zhang, Bret Harsham, John R Hershey, Tim K Marks, and Kazuhiko Sumi, “Attention-based multimodal fusion for video description,” in *ICCV 2017*. IEEE, 2017, pp. 4203–4212.
- [24] Jindřich Libovický and Jindřich Helcl, “Attention strategies for multi-source sequence-to-sequence learning,” in *ACL 2017*, 2017, vol. 2, pp. 196–202.
- [25] Jaejin Cho, Murali Karthick Baskar, Ruizhi Li, Matthew Wiesner, Sri Harish Mallidi, Nelson Yalta, Martin Karafiat, Shinji Watanabe, and Takaaki Hori, “Multilingual sequence-to-sequence speech recognition: architecture, transfer learning, and language modeling,” in *SLT 2018*, 2018.
- [26] Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Mael Guillemot, Thomas Hain, Jaroslav Kadlec, Vasiliis 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.
- [27] Mirco Ravanelli, Piergiorgio Svaizer, and Maurizio Omologo, “Realistic multi-microphone data simulation for distant speech recognition,” in *Interspeech 2016*, 2016.
- [28] Xavier Anguera, Chuck Wooters, and Javier Hernando, “Acoustic beamforming for speaker diarization of meetings,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 7, pp. 2011–2022, 2007.
- [29] Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplín, Jahn Heymann, Matthew Wiesner, Nanxin Chen, Adithya Renduchintala, and Tsubasa Ochiai, “Espnet: End-to-end speech processing toolkit,” in *Interspeech 2018*, 2018, pp. 2207–2211.
- [30] Takaaki Hori, Jaejin Cho, and Shinji Watanabe, “End-to-end speech recognition with word-based RNN language models,” *arXiv preprint arXiv:1808.02608*, 2018.