



# Cross-Attention End-to-End ASR for Two-Party Conversations

Suyoun Kim<sup>1</sup>, Siddharth Dalmia<sup>2</sup>, Florian Metze<sup>2</sup>

<sup>1</sup>Electrical & Computer Engineering

<sup>2</sup>Language Technologies Institute, School of Computer Science  
Carnegie Mellon University  
Pittsburgh, PA 15213; U.S.A.

{suyoung1|sdalmia|fmetze}@andrew.cmu.edu

## Abstract

We present an end-to-end speech recognition model that learns interaction between two speakers based on the turn-changing information. Unlike conventional speech recognition models, our model exploits two speakers history of conversational-context information that spans across multiple turns within an end-to-end framework. Specifically, we propose a speaker-specific cross-attention mechanism that can look at the output of the other speaker side as well as the one of the current speaker for better at recognizing long conversations. We evaluated the models on the Switchboard conversational speech corpus and show that our model outperforms standard end-to-end speech recognition models.

**Index Terms:** conversational speech recognition, end-to-end speech recognition

## 1. Introduction

Contextual information plays an important role in automatic speech recognition (ASR), especially in processing a long conversation since semantically related words, or phrases often re-occur across sentences. Typically, a long contextual information is only modeled in the language model (LM) which is trained only on text data separately [1, 2, 3, 4, 5, 6], then the language model combined with the acoustic model trained on isolated utterances in decoding phase. Such disjoint modeling process may not exploit the useful contextual information fully.

Several recent work attempted to use the contextual information with a recent progress *end-to-end* speech recognition framework, promises to integrate all available information into a single model, which is jointly optimized [7, 8, 9, 10, 11, 12]. In [7, 8, 9, 10], they proposed to use conversational-context embeddings which encodes previous utterance prediction to predict the current output token and showed promising results, however they did not consider the two-speaker interaction in a conversation. In [11, 12], they proposed to use phrase list (e.g. *song list*, *contact list* with attention mechanism and showed significant performance improvement, however, the model assumes that it has access to the phrase list at inference.

In this work, we create a cross-attention end-to-end speech recognizer capable of incorporating the two speaker conversational-context information to better process long conversations. Specifically, we first propose to use an additional attention mechanism to find more informative utterance representation among the multiple histories and focus more on it. Additionally, we propose to use LSTM with attention mechanism that specifically track the interactions between two speakers. We evaluate our model on the Switchboard conversational speech corpus [13, 14], and show that our model outperforms the standard end-to-end speech recognition model.

## 2. Related work

Several recent studies have considered to incorporate the context information within a end-to-end speech recognizer [11, 12]. In contrast with our method which uses a conversational-context information for processing a long two-speaker conversation, their methods use distinct phrases (i.e. *play a song*) with an attention mechanism in specific tasks, *contact names*, *songs names*, etc. Their model assumes there exists such a list of phrases at inference.

In study [5], they proposed contextual RNN language models that track the interactions between speakers by using additional RNN. In contrast with our work, they built a language model which is trained only on text corpus.

Several recent studies have considered to embed a longer context information within a end-to-end framework [7, 9, 10]. In contrast with our method, we consider multiple utterance histories for each speaker and use the attention mechanism with LSTM for learning the interaction between the two speakers.

## 3. Model

In this section, we first review conversational-context aware end-to-end speech recognition model [7, 9, 10]. We then present our proposed cross-attention end-to-end speech recognition model for processing two speaker conversations.

### 3.1. Encoder for history of the utterances

The key idea of the conversational-context aware end-to-end speech recognition model [7, 9, 10] is to use the history of utterance representation for predicting each output token. In order to obtain the history of utterance representation, an additional utterance encoder is used within the decoder network in the standard sequence-to-sequence framework [15, 16, 17, 18, 19]. This utterance encoder maps the variable-length output tokens from multiple preceding utterances into the fixed-length single vector, then it is fed into the decoder at each output time step.

Let we have  $K$ -utterances and acoustic features  $(x_1, \dots, x_T)^k$  and output word sequence,  $(w_1, \dots, w_U)$  for each  $k$ -th utterance. At each output step  $u$  in prediction of  $k$ -th utterance, decoder predicts the word distribution by conditioning on 1) acoustic embedding ( $e_a^k$ ) from Encoder, 2) previous word embedding ( $e_w^{u-1}$ ), and additionally 3) conversational-context embedding ( $e_c^k$ ):

$$e_a^k = \text{Encoder}(x^k) \quad (1)$$

$$w_u^k \sim \text{Decoder}(e_c^k, e_w^k, e_a^k) \quad (2)$$




Figure 1: Overall architecture of our proposed end-to-end speech recognition model using speaker-specific conversational-context embedding generated from the utterance history of both other speaker and current speaker. Our framework uses the utterance encoder to generate the utterance embedding, and uses additional attention mechanism for the multiple utterance embeddings from two speakers. Then, speaker-specific conversational-context embedding is forwarded to the conversational end-to-end speech recognizer. All of these model is a single network and trained in end-to-end manner.

For representing the utterance into the fixed length vector representation,  $e_c^k$ , we can use 1) mean of one-hot word vectors,  $\text{mean}(e_w^{k-1} + \dots + e_w^{k-1})$ , or 2) mean of external word embeddings, (i.e. Word2Vec [2], GloVe [20], fastText [21], etc), or 3) external sentence embedding resource (i.e. ELMo [22], BERT [23], etc). In this work, we are focusing on learning interactions of two speaker conversation rather than focusing on exploring the method of conversational-context representation, we thus use 3) BERT representation method in all of our experiments.

In order to pass the conversational-context across mini-batches during training and decoding, the dataset is serialized based on their onset times and their dialogs rather than random shuffling of data. Then, we create the mini-batches that contain only one utterance of each dialog to pass the context embedding to the next mini-batches properly.

### 3.2. Cross-Attention for two speakers' conversation

Although the previously proposed conversational-context aware end-to-end ASR model exploits the utterance of history as an additional information, there are two limitations. First, the model does not consider multiple speaker case and interaction between them, which is common and useful information in processing the conversation. Second, the model simply concatenates the multiple utterance embeddings then projects to a fixed dimensional vector or use the mean of the multiple utterance embeddings, so it does not explicitly allow to attend more on the more important utterance embedding.

Based on the two observations above, we therefore propose two methods to extend the current conversational-context aware end-to-end ASR for processing two-party conversations. The overall architecture of our proposed model is described in Figure 1. Specifically, our model works as follows.

We first represent an utterance to a fixed-length vector representation as described in Section 3.1. The utterance encoder maps the sequence of one-hot word vectors to the single, dense vector, the utterance embedding.

Next, we create a queue for each speaker to store the history of utterance embeddings as described in Figure 1. In this work, we consider two-speaker conversations and assume that the turn-changing information is known so that the utterance embeddings can be stored separately.

We then use an attention mechanism to generate speaker-specific conversational-context embedding given the history of *what other speaker said* and the history of *what current speaker said*. Note that, based on what current speaker is, we swap the queues properly. We propose two methods to generate the attended context embeddings.

#### 3.2.1. Attention over each speaker's utterance history

First method is simply using an additional attention mechanism over the utterance embeddings. Given the  $N$ -size of the utterance history for speaker A,  $e_{k-N}^{u-A}, \dots, e_{k-1}^{u-A}$ , the conversational-embedding,  $\text{att\_}e_k^A$ , is generated as follows:

$$\mathbf{G}_k^A = \tanh(\mathbf{W}e_{k-N:k-1}^A + \mathbf{b}) \quad (3)$$

$$\alpha_k^A = \text{softmax}(\mathbf{w}^T \mathbf{G}_k^A + \mathbf{b}) \quad (4)$$

$$\text{att\_}e_k^A = \sum_N \alpha_k^A \odot e_{k-N:k-1}^A \quad (5)$$

where  $\mathbf{W}$ ,  $\mathbf{b}$  are trainable parameters.  $e_k^B$  is generated in the same way.

### 3.2.2. Cross-attention between two speakers' utterance history

Second method is using LSTM with an attention mechanism. Inspired by the `matchLSTM` model which has been widely used in question answering tasks and natural language inference (NLI) [24, 25], we consider to track the interaction between two speakers sequentially, by attending *what other speaker said* at each utterance timestamp. The idea of the `matchLSTM` is to attempt to take the question (premise) and the passage (hypothesis) along with an answer pointer [26] pointing to the start and the end of the answer to make predictions. The `matchLSTM` tries to obtain a question-aware representation of the passage, by attending over the representations of the question tokens for each token in passage.

The key difference in our work is that the question (premise) is a sequence of utterance-embedding from other speaker (*what other speaker said*), and passage (hypothesis) is a sequence of utterance-embedding from current speaker (*what current speaker said*). The embedding of *what current speaker said* takes into consideration the alignment between the *what current speaker said* and *what other speaker said*.

By using `matchLSTM` over the first simple attention method, there are two benefits -

- First, the model is able to handle a longer utterance-history
- Second, the model can learn the interaction between the two speakers, as the `matchLSTM` can potentially track the flow of the conversations.

Specifically, the attended conversational embedding at  $i$ -th utterance-history step is generated as follows:

$$\mathbf{G}_{k_i}^A = \tanh(\mathbf{W}e_{k-N:k-1}^B + \mathbf{V}e_i^A + \mathbf{U}\mathbf{h}_{i-1}^A + \mathbf{b})$$

$$\alpha_{k_i}^A = \text{softmax}(\mathbf{w}^T \mathbf{G}_{k_i}^A + b)$$

where  $\mathbf{W}, \mathbf{V}, \mathbf{U}, \mathbf{b} \in \mathbb{R}^{h \times h}$ ,  $\mathbf{b} \in \mathbb{R}^h$ ,  $b \in \mathbb{R}$  are trainable pa-

rameters. Each hidden state  $\mathbf{h}_{i-1} \in \mathbb{R}^h$  comes from the output of the `matchLSTM` that is fed the following  $\mathbf{z}_i$  as input.

$$\mathbf{z}_i^A = [e_i^A, e_{k-N:k-1}^B \odot \alpha_{k_i}^A]$$

$$\mathbf{h}_i^A = \text{matchLSTM}(\mathbf{z}_i^A, \mathbf{h}_{i-1}^A)$$

Using a LSTM, there are  $N$  such  $h$ -dimensional hidden states, and we take the final hidden states for our attended conversational-context embedding:

$$\text{att\_e}_k^A = \mathbf{h}_{k-1}^A$$

Finally, the decoder network takes the above attended conversational-context embedding,  $\text{att\_e}_k^A$ , generated from the either method, in addition to the usual inputs, acoustic embedding from encoder network and the previous word embedding in current utterance. In this work, we used the same, 100 dimension for the conversational embedding for all of our experiments.

## 4. Experiments

### 4.1. Datasets

We trained our model on 300 hours of two-party conversational speech corpus, the Switchboard LDC corpus (97S62), and tested on the HUB5 Eval 2000 LDC corpora (LDC2002S09, LDC2002T43). In Section 5, we show separate results for the CallHome English (CH) and Switchboard (SWB) sets. We denote the number of utterances, the number of dialogs, the average number of utterances per dialog, and the number of speakers for each training, validation, evaluation sets in Table 1.

Table 1: *Experimental dataset description. We used 300 hours of Switchboard conversational corpus. Note that any pronunciation lexicon or Fisher transcription was not used.*

numbers	training	validation	SWB	CH
utterances	192,656	4,000	1,831	2,627
dialogs	2402	34	20	20
utterances / dialog	80	118	92	131
speakers	4804	68	40	40

The input feature for each frame which was sampled audio data at 16kHz is represented by a 83-dimensional feature vector, consisting of 80-dimensional log-mel filterbank coefficients and 3-dimensional pitch features. We used same output units in previous work [10], consisting of 10,038 the word units and the single character units.

### 4.2. Training and decoding

For the standard-end-to-end speech recognition model, we used joint CTC/Attention framework [19, 33] which is based on the sequence-to-sequence framework [15, 16, 17, 18] with using CTC [34] as an auxiliary objective function. We followed the same network architecture as the prior study in [35, 36]. We used CNN-BLSTM encoder and LSTM decoder. The CNN-BLSTM encoder consists of 6-layer CNN and 6-layer BLSTM with 320 cells. The LSTM decoder was 2-layer LSTM with 300 cells. Our proposed model requires approximately 2M trainable parameters for the speaker-specific attention mechanisms compared to our baseline. In Table 2 shows the total number of trainable parameters.

For optimization, we used AdaDelta [37] with gradient clipping [38]. We bootstrapped the training our proposed two-speaker conversational end-to-end models from the vanilla conversational models and the baseline end-to-end models.

For decoding, we used left-right beam search algorithm [39] with the beam size 10. We adjusted the score by adding a length penalty since the model has a small bias for shorter utterances. The final score is normalized with a length penalty 0.1.

The models are implemented using the PyTorch deep learning library [40], and ESPnet toolkit [19, 33, 41].

## 5. Results

The Table 2 summarizes the results of our baseline and proposed model, Attention and `matchLSTM`, with various number of utterance history. The Baseline is our baseline which is trained on isolated utterances without using

Table 2: Comparison of word error rates (WER) on Switchboard 300h with standard end-to-end speech recognition models and our proposed end-to-end speech recognition models with two-speaker conversational context. Note that our baselines did not use the external language model, Fisher text data, CD phones information, or the layer-wise pre-training technique [27]. (The \* mark denotes our estimate for the number of parameters used in the previous work.

Model	Output Units	Number of trainable parameters	Number of utterance-history	External LM	SWB (WER%)	CH (WER%)
<b>Prior Models</b>						
LF-MMI [28]	CD phones	N/A		✓	9.6	19.3
CTC [29]	Char	53M		✓	19.8	32.1
CTC [30]	Char, BPE-300,1k,10k	26M		✓	12.5	23.7
CTC [31]	Word (Phone init.)	N/A		✓	14.6	23.6
Seq2Seq [27]	BPE-10k	150M*		✗	13.5	27.1
Seq2Seq [32]	Word-10k	N/A		✗	23.0	37.2
<b>Our Baselines</b>						
Baseline	Word-10k	32M	✗	✗	17.9	30.6
<b>Our Models</b>						
Attention	Word-10k	34M	6	✗	16.7	30.0
matchLSTM	Word-10k	34M	6	✗	16.4	29.9
Attention	Word-10k	34M	10	✗	16.6	30.0
matchLSTM	Word-10k	34M	10	✗	16.4	29.9
Attention	Word-10k	34M	20	✗	16.6	30.0
matchLSTM	Word-10k	34M	20	✗	16.4	29.8

conversational-context information. Our proposed models were bootstrapped the training from the vanilla conversational model, and the vanilla model was also bootstrapped the training from the Baseline. We found that this pre-training procedure is necessary since we need to learn the additional parameters in decoder network.

We also shows the results of prior models from other literature. Note that our baselines used smaller number of model parameters, did not use the external language model, Fisher text data, CD phones information, or the layer-wise pre-training technique [27].

We first observed that our conversational models with both matchLSTM method and Attention method achieved better performance than the baseline. Specifically, the proposed model with matchLSTM using 20 utterance-history performed best, showing 16.4% WER and 29.8% WER in Switchboard (SWB) and CallHome (CH) evaluation sets, respectively. Figure 2 visualizes the attention weights over the utterance-history. It shows that the model focuses on the long, informative utterances, rather than the short, meaningless utterances, i.e. *oh*, *heck yeah*, etc.

One noticeable thing is that the matchLSTM method slightly outperformed the Attention method in various number of utterance-history. It is possible that the matchLSTM method can track the flow of the conversation because it generates the better speaker-specific conversational context representation by conditioning on what others and current speaker said at each step. This point needs to be verified with additional experiments in future work.

## 6. Conclusions

We have introduced an end-to-end speech recognizer with speaker-specific cross-attention mechanism for two-party conversations. Unlike conventional speech recognition models, our model generates output tokens conditioning on two speakers'

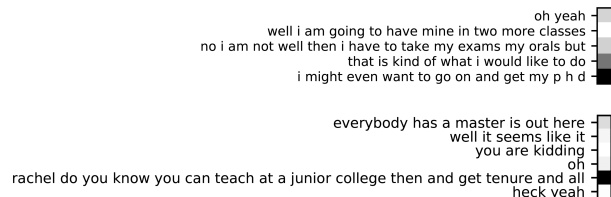


Figure 2: The attention weights over utterance-history of the speaker A (top) and the speaker B (bottom) when the model predicts the utterance (come out here to California) in evaluation set. The dark color represents higher attention weight.

conversational history, and consequently improves recognition accuracy of a long conversation. Our proposed speaker-specific cross-attention mechanism can look at what other speaker said in addition to what current speaker said. We evaluated the models on the Switchboard conversational speech corpus and show that our proposed model using cross-attention achieves WER improvements over the baseline end-to-end model for Switchboard and CallHome.

A future direction would be to explore the variant of the current cross-attention model, such as taking the word-level history rather than the utterance-level history. We also plan to analyze the effect of the attention mechanism to get better understanding.

## 7. Acknowledgments

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research. This work also used the Bridges system, which is supported by NSF award number ACI-1445606, at the Pittsburgh Supercomputing Center (PSC).

## 8. References

- [1] T. Mikolov, M. Karafiát, L. Burget, J. Černocký, and S. Khudanpur, "Recurrent neural network based language model," in *Eleventh Annual Conference of the International Speech Communication Association*, 2010.
- [2] T. Mikolov and G. Zweig, "Context dependent recurrent neural network language model." *SLT*, vol. 12, pp. 234–239, 2012.
- [3] T. Wang and K. Cho, "Larger-context language modelling," *arXiv preprint arXiv:1511.03729*, 2015.
- [4] Y. Ji, T. Cohn, L. Kong, C. Dyer, and J. Eisenstein, "Document context language models," *arXiv preprint arXiv:1511.03962*, 2015.
- [5] B. Liu and I. Lane, "Dialog context language modeling with recurrent neural networks," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5715–5719.
- [6] W. Xiong, L. Wu, J. Zhang, and A. Stolcke, "Session-level language modeling for conversational speech," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018, pp. 2764–2768.
- [7] S. Kim and F. Metze, "Dialog-context aware end-to-end speech recognition," *SLT*, 2018.
- [8] S. Kim, S. Dalmia, and F. Metze, "Situation informed end-to-end asr for chime-5 challenge," *ChIME5 workshop*, 2018.
- [9] S. Kim and F. Metze, "Acoustic-to-word models with conversational context information," *NAACL*, 2019.
- [10] S. Kim, S. Dalmia, and F. Metze, "Gated embeddings in end-to-end speech recognition for conversational-context fusion," *ACL*, 2019.
- [11] G. Pundak, T. N. Sainath, R. Prabhavalkar, A. Kannan, and D. Zhao, "Deep context: end-to-end contextual speech recognition," *arXiv preprint arXiv:1808.02480*, 2018.
- [12] U. Alon, G. Pundak, and T. N. Sainath, "Contextual speech recognition with difficult negative training examples," *arXiv preprint arXiv:1810.12170*, 2018.
- [13] J. Godfrey and E. Holliman, "Switchboard-1 release 2 ldc97s62," *Linguistic Data Consortium, Philadelphia*, vol. LDC97S62, 1993.
- [14] J. J. Godfrey, E. C. Holliman, and J. McDaniel, "Switchboard: Telephone speech corpus for research and development," in *Acoustics, Speech, and Signal Processing, 1992. ICASSP-92., 1992 IEEE International Conference on*, vol. 1. IEEE, 1992, pp. 517–520.
- [15] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [16] J. Chorowski, D. Bahdanau, K. Cho, and Y. Bengio, "End-to-end continuous speech recognition using attention-based recurrent NN: First results," *arXiv preprint arXiv:1412.1602*, 2014.
- [17] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," in *Advances in Neural Information Processing Systems*, 2015, pp. 577–585.
- [18] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, "Listen, attend and spell," *arXiv preprint arXiv:1508.01211*, 2015.
- [19] S. Kim, T. Hori, and S. Watanabe, "Joint ctc-attention based end-to-end speech recognition using multi-task learning," in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 4835–4839.
- [20] J. Pennington, R. Socher, and C. Manning, "Glove: Global vectors for word representation," in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.
- [21] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, "Enriching word vectors with subword information," *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 135–146, 2017.
- [22] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, "Deep contextualized word representations," in *Proc. of NAACL*, 2018.
- [23] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [24] S. Wang and J. Jiang, "Learning natural language inference with lstm," *arXiv preprint arXiv:1512.08849*, 2015.
- [25] —, "Machine comprehension using match-lstm and answer pointer," *arXiv preprint arXiv:1608.07905*, 2016.
- [26] O. Vinyals, M. Fortunato, and N. Jaitly, "Pointer networks," in *Advances in Neural Information Processing Systems*, 2015, pp. 2692–2700.
- [27] A. Zeyer, K. Irie, R. Schlüter, and H. Ney, "Improved training of end-to-end attention models for speech recognition," *arXiv preprint arXiv:1805.03294*, 2018.
- [28] D. Povey, V. Peddinti, D. Galvez, P. Ghahremani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely sequence-trained neural networks for asr based on lattice-free mmi," in *Interspeech*, 2016, pp. 2751–2755.
- [29] G. Zweig, C. Yu, J. Droppo, and A. Stolcke, "Advances in all-neural speech recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 4805–4809.
- [30] R. Sanabria and F. Metze, "Hierarchical multi task learning with ctc," *arXiv preprint arXiv:1807.07104*, 2018.
- [31] K. Audhkhasi, B. Kingsbury, B. Ramabhadran, G. Saon, and M. Picheny, "Building competitive direct acoustics-to-word models for english conversational speech recognition," *arXiv preprint arXiv:1712.03133*, 2017.
- [32] S. Palaskar and F. Metze, "Acoustic-to-word recognition with sequence-to-sequence models," *arXiv preprint arXiv:1807.09597*, 2018.
- [33] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. 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.
- [34] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd international conference on Machine learning*. ACM, 2006, pp. 369–376.
- [35] Y. Zhang, W. Chan, and N. Jaitly, "Very deep convolutional networks for end-to-end speech recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 4845–4849.
- [36] T. Hori, S. Watanabe, Y. Zhang, and W. Chan, "Advances in joint ctc-attention based end-to-end speech recognition with a deep cnn encoder and rnn-lm," *arXiv preprint arXiv:1706.02737*, 2017.
- [37] M. D. Zeiler, "Adadelata: an adaptive learning rate method," *arXiv preprint arXiv:1212.5701*, 2012.
- [38] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *International Conference on Machine Learning*, 2013, pp. 1310–1318.
- [39] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Advances in neural information processing systems*, 2014, pp. 3104–3112.
- [40] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," in *NIPS-W*, 2017.
- [41] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen *et al.*, "Espnet: End-to-end speech processing toolkit," *arXiv preprint arXiv:1804.00015*, 2018.