

EXPLORING PRE-TRAINING WITH ALIGNMENTS FOR RNN TRANSDUCER BASED END-TO-END SPEECH RECOGNITION

Hu Hu^{1,2*}, Rui Zhao¹, Jinyu Li¹, Liang Lu¹, Yifan Gong¹

¹Microsoft Speech and Language Group, Redmond, WA, USA

²Georgia Institute of Technology, Atlanta, GA, USA

ABSTRACT

Recently, the recurrent neural network transducer (RNN-T) architecture has become an emerging trend in end-to-end automatic speech recognition research due to its advantages of being capable for on-line streaming speech recognition. However, RNN-T training is made difficult by the huge memory requirements, and complicated neural structure. A common solution to ease the RNN-T training is to employ connectionist temporal classification (CTC) model along with RNN language model (RNNLM) to initialize the RNN-T parameters. In this work, we conversely leverage external alignments to seed the RNN-T model. Two different pre-training solutions are explored, referred to as encoder pre-training, and whole-network pre-training respectively. Evaluated on Microsoft 65,000 hours anonymized production data with personally identifiable information removed, our proposed methods can obtain significant improvement. In particular, the encoder pre-training solution achieved a 10% and a 8% relative word error rate reduction when compared with random initialization and the widely used CTC+RNNLM initialization strategy, respectively. Our solutions also significantly reduce the RNN-T model latency from the baseline.

Index Terms— RNN transducer, end-to-end, alignments, speech recognition, pre-training.

1. INTRODUCTION

In recent years, we have witnessed significant progress in automatic speech recognition (ASR) mainly due to the use of deep learning algorithms [1, 2]. Deep model based ASR systems mainly focus on the hybrid framework and consist of many components, including acoustic model (AM), pronunciation model, language model (LM). Those components are trained separately using different objective functions, and extra expert linguistic knowledge may be needed. Recently, an emerging trend in the ASR community is to rectify this disjoint training issue by replacing hybrid systems with end-to-end (E2E) systems [3, 4, 5, 6, 7, 8, 9, 10, 11, 12]. The three major E2E approaches are built on the Connectionist Temporal Classification (CTC) [13, 14, 15], Attention-based Encoder-Decoder (AED) [16, 17, 18, 19, 5], and recurrent neural network transducer (RNN-T) [20, 21, 22]. Different from training conventional hybrid models, token alignment information between input acoustic frames and output token sequence is not required when training the E2E models.

CTC maps the input speech frames to target label sequence by marginalizing all the possible alignments. A dynamic programming based forward-backward algorithm is usually used to train the model. An advantage of the CTC approach is that it does frame-level decoding as the conventional hybrid model, and hence can be

applied for online speech recognition. However, one disadvantage is the conditional independence assumption given the input acoustic frames. AED, on the other hand, does not have such an assumption, and is presumably more powerful than CTC for the speech recognition task. However, one drawback of the AED model is that the entire input sequence is required to start the decoding process due to the global attention mechanism, which makes it challenging for real-time streaming ASR, despite some recent attempts along this direction [23, 24].

RNN-T is an extension to CTC, which consists of three components: an encoder, a prediction network, and a joint network which integrates the outputs of encoder and prediction networks together to predict the target labels. RNN-T overcomes the conditional independence assumption of CTC with the prediction network; moreover, it allows streaming ASR because it still preforms frame-level monotonic decoding. Hence, there has been a significant research effort in promoting this approach in the ASR community [22, 21, 25, 26, 27], and RNN-T has recently been successfully deployed in embedding devices [28].

However, compared to CTC or AED, RNN-T is much more difficult to train due to the model structure, and the synchronous decoding constraint. Besides, its training is very memory demanding due to the 3-dimensional output tensor [21, 22]. In [22], an approach is proposed to reduce the memory cost, and it enables large mini-batch training. To tackle the training difficulty, initializing the encoder and prediction networks of an RNN-T with a CTC model and an RNNLM respectively is proven to be beneficial [21, 27]. In this paper, we explore other model initialization approaches to overcome the training difficulty of RNN-T models. Specifically, we propose to utilize external token alignment information to pre-train RNN-T. Two types of pre-training methods are investigated, which are referred to as encoder pre-training and whole-network pre-training respectively. Encoder pre-training refers to initializing the encoder in the RNN-T only, while the other components are trained from the random initialization. The whole-network pre-training, as its name suggests, pre-trains the whole network by an auxiliary objective function instead of the RNN-T loss. The proposed methods are evaluated on 3400 hours voice assistant data and 65,000 hours production data. The experimental results show that the accuracy of RNN-T model can be significantly improved with our proposed pre-training methods, with up-to 28% relative word error rate (WER) reduction.

The rest of this paper is organized as follows: Section 2 briefly introduces the basic RNN-T model, including the model training and decoding. The proposed two types of pre-training methods are described in Section 3 and Section 4, respectively. Next, Section 5 shows the experimental results and analysis. Section 6 gives the conclusions.

*The work was done as an intern at Microsoft.

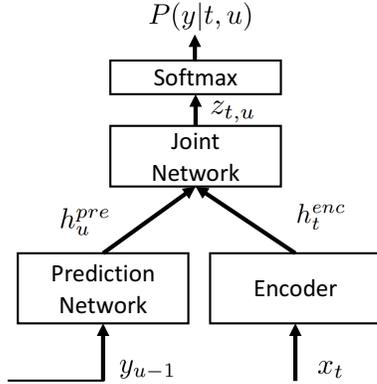


Fig. 1. RNN-Transducer model structure.

2. RNN TRANSDUCER MODEL

The RNN-T model was proposed in [20] as an extension to the CTC model. A typical RNN-T model has three components, as shown in the Figure 1, namely encoder, prediction network and joint network. Compared with CTC, RNN-T does not have the conditional independence assumption because of the prediction network, which emits an output tokens conditioned on the previous prediction results.

To be more precise, the encoder in an RNN-T model is an RNN that maps each acoustic frame x_t to a high-level feature representation h_t^{enc} , where t is the time index:

$$h_t^{enc} = f^{enc}(x_t). \quad (1)$$

The prediction network, which is also based on RNNs, converts previous non-blank output token y_{u-1} to a high-level representation h_u^{pre} , where u is the label index of each output token.

$$h_u^{pre} = f^{pre}(y_{u-1}). \quad (2)$$

Given the hidden representations of both acoustic features and labels from the encoder and prediction network, the joint network integrates the information using a feed-forward network as:

$$z_{t,u} = f^{joint}(h_t^{enc}, h_u^{pre}). \quad (3)$$

The posterior probability $P(y|t, u)$ can be obtained by taking the Softmax operation on the output of the joint network. Then a forward-backward algorithm [20] is performed on the three-dimension output to compute the total probability $P(\mathbf{y}|\mathbf{x})$ of the output sequence \mathbf{y} , conditioned on the input sequence \mathbf{x} . The negative log-loss of the target sequence is used as the objective function to train the model,

$$L_{RNN-T} = -\log P(\mathbf{y}|\mathbf{x}). \quad (4)$$

The decoding of RNN-T is operated in a frame-by-frame fashion. Starting from the first frame fed to encoder, if the current output is not *blank*, then the prediction network is updated with that output token. Otherwise, if the output is *blank*, then the encoder is updated with the next frame. The decoding terminates when the last frame of input sequence is consumed. In this way, real-time streaming is satisfied. Greedy search and beam search can be used in the decoding stage, which stores different numbers of intermediate states.

3. ENCODER PRE-TRAINING

In an RNN-T model, encoder and prediction network usually have different model structures, which make it difficult to train them well at the same time. Directly training RNN-T from random initialization may get a biased model toward one of the model components, i.e., dominated by the acoustic or the language input. Most groups adopt a initialization strategy that initializes the encoder with a CTC model and the prediction network with a RNNLM [27, 21, 29]. However, the output sequence of CTC is a series of spikes, separated by the *blank* [13]. Thus after CTC based pre-training, most encoder output h_t^{enc} leads to generate *blank*, which results in wrong inference for the RNN-T model.

In our work, we propose to utilize external alignments to pre-train the encoder with the Cross Entropy (CE) criterion. The encoder is regarded as a token classification model rather than a CTC model. As shown in the right part of Figure 2, an RNN based token classification model is trained first with the CE loss. In this paper we use 'CE loss' to represent the cross entropy loss function, and 'CTC loss' to represent the CTC forward-backward algorithm based loss function, and 'RNN-T loss' to represent the RNN-T loss function.

In our experiments, we use word piece units as target tokens [30], which have been explored in the context of machine translation, and successfully applied in E2E ASR [21, 15]. With word-level alignments, we can get the boundary frame index of each word. For the word divided into more than one word piece, we equally allocate the total frames inside the word boundary to its word pieces. There will be a marginal case in which a word contains more word pieces than frames, which prevents us from generating token alignments. The total ratio of this special case is less than 0.01% of all the training utterances, so we just remove those utterance in the pre-training stage. In this way, we can obtain the hard alignments of target tokens for all the frames.

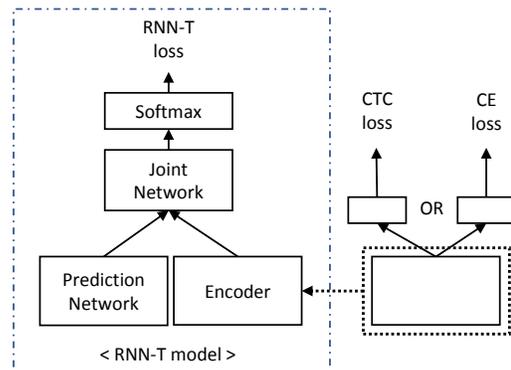


Fig. 2. Illustration of encoder pre-training for RNN-T. Dashed arrow means initializing from a pre-trained model.

Based on the encoder structure, one extra fully connected layer is added on top of the encoder, in which the output h_t^{enc} is used for token classification. The objective is

$$L_{enc} = -\sum_{k=1}^K y_{t,k} * \log(\text{softmax}(f^{fc}(h_{t,k}^{enc}))), \quad (5)$$

where f^{fc} represents a fully connected layer, k is the label index and K denotes the target dimension, which is also the dimension of $z_{t,u}$. And y_t is the word piece label for each input frame x_t . After

the encoder pre-training, each output h_t^{enc} , which is the high-level representation of input acoustic features, is expected to contain the information about the alignments.

4. WHOLE-NETWORK PRE-TRAINING

Among encoder pre-training methods, encoder is regarded to perform token mapping (CTC loss pre-training) or token aligning (CE loss pre-training). However, these pre-training methods only consider part of the RNN-T model. In this paper, we also explore the whole-network pre-training method with the use of external token alignments information. Different from other models, the output of the RNN-T is three-dimensional. Thus, the key challenge for the whole-network pre-training is the label tensor \mathbf{y} design. The optimizer reduces the CE loss between the output z of the model and a crafted three-dimensional label tensor \mathbf{y} .

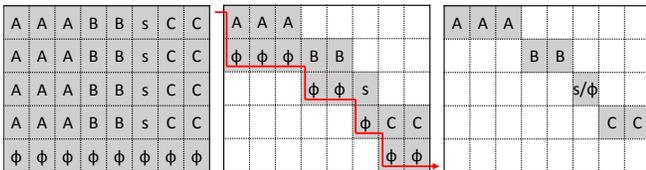


Fig. 3. Examples of three designed label tensors for whole-network pre-training. Each grid represents an one-hot vector. The example 8-frame utterance is 'A B s C', and its alignment is 'A A A B B s C C'. In each label tensor, 's' means *space* and 'Φ' means *blank*. From left to right, each sub-figure represents \mathbf{y}_1 to \mathbf{y}_3 as the order. Only gray grids are used for CE computing. The read arrow in \mathbf{y}_2 represents the decoding path when decoding.

Three different types of label tensor \mathbf{y} are explored in this study, and examples are given in Figure 3. In each label tensor, the horizontal axis represents the time dimension, and the vertical axis represents the output token dimension. We represent *blank* as an additional class, and it is shown as one-hot vector in the label tensor. At the first label tensor \mathbf{y}_1 , we set all the output target grids of each frame in \mathbf{y}_1 to the one-hot vector corresponding to its alignment label. The last row of the label tensor is set to all *blank*, which indicates the end of the utterance. Thus, after pre-training, encoder output h_t^{enc} is supposed to contain the alignment information. However, \mathbf{y}_1 only considers the frame-by-frame alignment, but ignores the output token level information. If we directly perform the RNN-T decoding on \mathbf{y}_1 , we can not obtain the correct inference sequence.

Thus, taken the decoding process into consideration, we design another label tensor \mathbf{y}_2 . Each frame is assigned to its token alignment. Target token position is determined by its sequence order. When perform pre-training, we only compute the CE from the non-empty part of the label tensor. *blank* token is inserted under each target token to ensure the correct decoding results. If we directly perform the RNN-T decoding algorithm on the label tensor \mathbf{y}_2 , correct results should be obtained. The decoding path is illustrated by the red arrow on the label tensor \mathbf{y}_2 in Figure 3. Thus, by directly performing the decoding on \mathbf{y}_2 of the given example, the inference result is 'A Φ Φ Φ B Φ Φ s Φ C Φ Φ'. After removing *blank* tokens, the final result is 'A B s C', which is the same with the alignment of this utterance.

However, in \mathbf{y}_2 , almost half of the valid part is *blank*, so that *blank* tokens dominate the pre-training process. Therefore, we design the label tensor \mathbf{y}_3 , which only keeps the non-blank part of

\mathbf{y}_2 . The label tensor only remains one grid with its corresponding alignment for each frame. In order to provide the blank information during the pre-training stage, we set short-pause (space token less than 3 frames) of each utterance to *blank*. That means some *space* in the valid part of the label tensor will become *blank*. After the pre-training is done, we replace the CE loss with the RNN-T loss, and proceed to the standard RNN-T training.

5. EXPERIMENTS AND ANALYSIS

5.1. Experimental setup

The proposed methods are evaluated on the 3400-hour Cortana voice assistant data, and 65,000-hour Microsoft production data. For the Cortana data, the training and test sets consist of approximately 3400 hours and 6 hours of English audio, respectively. The 65,000 hours production data are transcribed data from all kinds of Microsoft products. The test sets cover 13 application scenarios such as Cortana and far-field speaker, with totally 1.8 million (M) words. Training and test material is anonymized, with personally identifiable information removed. In this work, we evaluate the methods on Cortana data at first, and then evaluate the selected best method on very large scale 65,000-hour production data.

The input feature is a vector of 80-dimension log Mel filter bank for every 10 milliseconds (ms) of speech. Eight vectors are stacked together to form an input frame to the encoder, and the frame shift is 30 ms. All RNN-T models adopt the configuration recommended in [22, 28]. All encoders (Enc.) have 6 hidden-layer LSTMs, and all prediction networks (Pred.) have 2 hidden-layer LSTMs. The joint network has two linear layers without any activation functions. Layer normalization is used in all LSTM layers, and the hidden dimension is 1280 with the projection dimension equal to 640. The output layer models 4000 word piece units together with *blank* token. The word piece units are generated by running byte pair encoding [31] on the acoustic training texts.

5.2. Evaluation on Cortana data

Experimental results of whole-network pre-training are shown in the Table 1. The RNN-T baseline is trained from random initialization. For pre-trained models, the whole network is pre-trained with CE loss, then it is trained with RNN-T loss. Using the pre-trained network as a seed model, the final word error rate (WER) can be significantly reduced. All designed label tensors can improve the RNN-T training, achieving 10% to 12% relative WER reduction.

Table 1. WER comparison of different whole-network pre-training methods on 3400 hours Cortana data. Pre-train (all align) uses \mathbf{y}_1 , Pre-train (correct decoding) uses \mathbf{y}_2 , and Pre-train (align path - sp blank) uses \mathbf{y}_3 as the target label sensors, respectively.

Model	WER(%)
RNN-T	15.11
+ Pre-train (all align)	13.53
+ Pre-train (correct decoding)	13.66
+ Pre-train (align path - sp blank)	13.23

In Table 2, we evaluate the encoder pre-training method on 3400 hours Cortana data. Using a pre-trained CTC model to initialize the encoder does not improve the accuracy. This is because the output of CTC is a sequence of spikes, in which there are lots of *blank*

tokens without any meaning. Hence, if we use the pre-trained CTC as the seed for the encoder of RNN-T, most encoder output h_t^{enc} will generate *blank*, which does not help the RNN-T training. When we use CE loss pre-trained encoder to initialize the encoder of RNN-T, it achieves significant improvement compared with training from the random initialization. It obtains 28% relative WER reduction from the RNN-T baseline and CTC based encoder pre-training.

Table 2. WER comparison of different encoder pre-training methods on 3400 hours Cortana data. Greedy search is used. ‘CTC’ means pre-training of CTC loss with target sequence, ‘CE’ means pre-training of CE loss with target alignment. ‘no’ means training from the random initialization.

Model	Enc. Pre-train	WER (%)
RNN-T baseline	no	15.11
+ Pre-train	CTC	15.07
+ Pre-train	CE	10.83

Among all the encoder pre-training experiments in Table 2, prediction network and joint network are all trained from the random initialization. The only difference is the parameters seed of encoder. When comparing CTC loss and CE loss based encoder pre-training methods, there is a huge WER gap between these two approaches. Initializing the encoder as a token aligning model rather than a sequence mapping model results in the much better accuracy. This is because the RNN-T encoder performs the frame-to-token aligning, which extracts the high-level features of each input frame.

5.3. Evaluation on very large scale data

From our experiments, both the encoder pre-training and the whole-network pre-training can improve the performance of RNN-T model. In order to get more convincing results, we evaluate our proposed methods on very large scale data, where we use the 65,000 hours Microsoft production data set. The results are shown in Table 3. Due to the very large resource requirement and computation cost, we only evaluate CE-based encoder pre-training method, which obtained the best accuracy in Cortana experiments. All the results are obtained using beam search, and the beam width is 5.

Besides our proposed methods, we evaluate the widely used CTC-RNNLM pre-training strategy [21, 29, 27] as comparison. It used a well trained CTC model to initialize the encoder, and a well trained RNNLM to initialize the prediction network. This CTC+RNNLM initialization approach reduced the average WER from 12.63 to 12.29 in 13 test scenarios with 1.8 M words. In contrast, our proposed approach, which pre-trains the encoder with alignments using the CE loss, outperforms the other methods significantly, achieving a 11.41 WER on the average. Compared with training from the random initialization, our proposed method can obtain 10% relative WER reduction in the very large scale task.

Table 3. WER comparison of different encoder pre-training methods on the task with 65,000 hours production training data.

Model	WER (%)
RNN-T baseline	12.63
+ Pre-train (Enc. CTC, Pred. LM) [21]	12.29
+ Pre-train (Enc. CE)	11.34

5.4. Output time delay comparison

Although RNN-T is a natural streaming model, it still has latency compared to hybrid models [22]. With the help of alignments for model initialization, we hope to reduce the latency of RNN-T. To better understand the advantages of our proposed pre-training methods, we compare the gap between the ground truth word alignment and the word alignment generated by greedy decoding from different RNN-T models. The visualization is performed on the test set of Cortana data. As shown in the Figure 4, the central axis represents the ground truth word alignment. The output alignment distributions are normalized to the normal distribution. The horizontal axis represents the number of frames away from the ground truth word alignment, and the vertical axis represents the ratio of words.

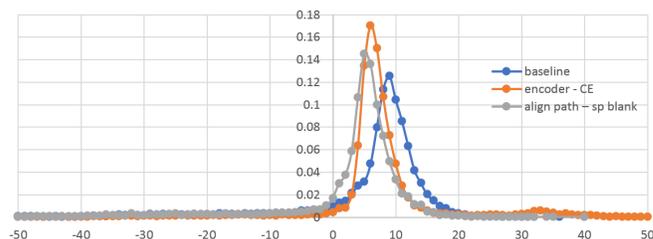


Fig. 4. Frame delay difference between ground truth word alignment and the word alignment generated from different RNN-T models.

From the Figure 4, different RNN-T models have different time delay compared with the ground truth. That’s because the RNN-T model tends to see several future frames, which can provide more information for the token recognition. The baseline RNN-T model has around 10 frames average delay. In contrast, when performed the proposed pre-training methods, the average delay can be significantly reduced. Using CE pre-trained encoder to initialize the RNN-T model can reduce the average delay to 6 frames, and using whole-network pre-training method can reduce it to 5 frames. The reason for the time delay reduction is that pre-training provides the alignment information to the RNN-T model, which will guide the model to make decision earlier. This shows the advantage of our proposed pre-training methods in terms of time delay during the decoding stage.

6. CONCLUSION

In this work, we explore the training strategy of an RNN-T model and propose two pre-training approaches with the use of external alignment information. Two types of pre-training methods have been evaluated, referred to as encoder pre-training and whole-network pre-training. Encoder pre-training used CE loss to pre-train the encoder of RNN-T only. Whole-network pre-training pre-trains the whole RNN-T model with CE loss. Three kinds of designed label tensors are used for the whole-network pre-training. The proposed methods are evaluated on 3400 hours Cortana data and 65000 hours production data. When compared with training from the random initialization, the whole-network pre-training obtains a 12% relative WER reduction. And the encoder pre-training obtains a 28% and a 10% relative WER reduction on 3400 hours Cortana and 65,000 hours production data, respectively. Compared to the widely used CTC+RNNLM initialization strategy on very large scale data, encoder pre-training still outperforms it by a 8% relative WER reduction. Our proposed methods can also significantly reduce the time delay of RNN-T model.

7. REFERENCES

- [1] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A. r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, and B. Kingsbury, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal Process. Mag.*, vol. 29, no. 6, pp. 82–97, 2012.
- [2] Frank Seide, Gang Li, and Dong Yu, "Conversational speech transcription using context-dependent deep neural networks," in *Proc. INTERSPEECH*, 2011.
- [3] H. Sak, A. Senior, K. Rao, O. Irsoy, A. Graves, F. Beaufays, and J. Schalkwyk, "Learning Acoustic Frame Labeling for Speech Recognition with Recurrent Neural Networks," in *Proc. ICASSP*. IEEE, 2015, pp. 4280–4284.
- [4] Y. Miao, M. Gowayyed, and F. Metze, "EESSEN: End-to-End Speech Recognition using Deep RNN Models and WFST-based Decoding," in *Proc. ASRU*. IEEE, 2015, pp. 167–174.
- [5] William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *Proc. ICASSP*, 2016, pp. 4960–4964.
- [6] R. Prabhavalkar, K. Rao, T. N. Sainath, B. Li, L. Johnson, and N. Jaitly, "A Comparison of Sequence-to-Sequence Models for Speech Recognition," in *Proc. Interspeech*, 2017, pp. 939–943.
- [7] Chung-Cheng Chiu, Tara N Sainath, Yonghui Wu, Rohit Prabhavalkar, Patrick Nguyen, Zhifeng Chen, Anjuli Kannan, Ron J Weiss, Kanishka Rao, Katya Gonina, et al., "State-of-the-art speech recognition with sequence-to-sequence models," in *Proc. ICASSP*, 2018.
- [8] E. Battenberg, J. Chen, R. Child, A. Coates, Y. Gaur, Y. Li, H. Liu, S. Satheesh, D. Seetapun, A. Sriram, et al., "Exploring Neural Transducers for End-to-End Speech Recognition," in *Proc. ASRU*, 2017.
- [9] Hasim Sak, Matt Shannon, Kanishka Rao, and Françoise Beaufays, "Recurrent neural aligner: An encoder-decoder neural network model for sequence to sequence mapping," in *Proc. Interspeech*, 2017.
- [10] Hossein Hadian, Hossein Sameti, Daniel Povey, and Sanjeev Khudanpur, "Towards Discriminatively-trained HMM-based End-to-end models for Automatic Speech Recognition," in *Proc. ICASSP*, 2018.
- [11] Tara N Sainath, Chung-Cheng Chiu, Rohit Prabhavalkar, Anjuli Kannan, Yonghui Wu, Patrick Nguyen, and ZhiJeng Chen, "Improving the performance of online neural transducer models," in *Proc. ICASSP*, 2018, pp. 5864–5868.
- [12] Jinyu Li, Guoli Ye, Amit Das, Rui Zhao, and Yifan Gong, "Advancing acoustic-to-word CTC model," in *Proc. ICASSP*, 2018, pp. 5794–5798.
- [13] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks," in *ICML*, 2006, pp. 369–376.
- [14] A. Graves and N. Jaitley, "Towards End-to-End Speech Recognition with Recurrent Neural Networks," in *PMLR*, 2014, pp. 1764–1772.
- [15] Amit Das, Jinyu Li, Guoli Ye, Rui Zhao, and Yifan Gong, "Advancing acoustic-to-word ctc model with attention and mixed-units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 12, pp. 1880–1892, 2019.
- [16] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation," in *EMNLP*, 2014.
- [17] D. Bahdanau, K. Cho, and Y. Bengio, "Neural Machine Translation by Jointly Learning to Align and Translate," in *ICLR*, 2015.
- [18] Dzmitry Bahdanau, Jan Chorowski, Dmitriy Serdyuk, Philemon Brakel, and Yoshua Bengio, "End-to-end attention-based large vocabulary speech recognition," in *Proc. ICASSP*, 2016, pp. 4945–4949.
- [19] J. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-Based Models for Speech Recognition," in *NIPS*, 2015.
- [20] A. Graves, "Sequence Transduction with Recurrent Neural Networks," *CoRR*, vol. abs/1211.3711, 2012.
- [21] Kanishka Rao, Haşim Sak, and Rohit Prabhavalkar, "Exploring architectures, data and units for streaming end-to-end speech recognition with rnn-transducer," in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 193–199.
- [22] Jinyu Li, Rui Zhao, Hu Hu, and Yifan Gong, "Improving RNN Transducer Modeling for End-to-End Speech Recognition," in *Proc. ASRU*, 2019.
- [23] Chung-Cheng Chiu and Colin Raffel, "Monotonic chunkwise attention," *arXiv preprint arXiv:1712.05382*, 2017.
- [24] Niko Moritz, Takaaki Hori, and Jonathan Le Roux, "Triggered attention for end-to-end speech recognition," in *Proc. ICASSP*, 2019, pp. 5666–5670.
- [25] Anjuli Kannan, Arindrima Datta, Tara N Sainath, Eugene Weinstein, Bhuvana Ramabhadran, Yonghui Wu, Ankur Bapna, Zhifeng Chen, and Seungji Lee, "Large-scale multilingual speech recognition with a streaming end-to-end model," *arXiv preprint arXiv:1909.05330*, 2019.
- [26] Eric Battenberg, Jitong Chen, Rewon Child, Adam Coates, Yashesh Gaur Yi Li, Hairong Liu, Sanjeev Satheesh, Anuroop Sriram, and Zhenyao Zhu, "Exploring neural transducers for end-to-end speech recognition," in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 206–213.
- [27] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton, "Speech recognition with deep recurrent neural networks," in *2013 IEEE international conference on acoustics, speech and signal processing*. IEEE, 2013, pp. 6645–6649.
- [28] Yanzhang He, Tara N Sainath, Rohit Prabhavalkar, Ian McGraw, Raziq Alvarez, Ding Zhao, David Rybach, Anjuli Kannan, Yonghui Wu, Ruoming Pang, et al., "Streaming end-to-end speech recognition for mobile devices," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6381–6385.
- [29] Senmao Wang, Pan Zhou, Wei Chen, Jia Jia, and Lei Xie, "Exploring rnn-transducer for chinese speech recognition," *CoRR*, vol. abs/1811.05097, 2018.
- [30] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al., "Google's neural machine translation system: Bridging the gap between human and machine translation," *arXiv preprint arXiv:1609.08144*, 2016.
- [31] Rico Sennrich, Barry Haddow, and Alexandra Birch, "Neural machine translation of rare words with subword units," *arXiv preprint arXiv:1508.07909*, 2015.