

# LONGFNT: LONG-FORM SPEECH RECOGNITION WITH FACTORIZED NEURAL TRANSDUCER

Xun Gong<sup>1,2</sup>\*, Yu Wu<sup>2</sup>, Jinyu Li<sup>2</sup>, Shujie Liu<sup>2</sup>, Rui Zhao<sup>2</sup>, Xie Chen<sup>1</sup>, Yanmin Qian<sup>1</sup>

<sup>1</sup> MoE Key Lab of Artificial Intelligence, AI Institute  
X-LANCE Lab, Department of Computer Science and Engineering, Shanghai Jiao Tong University  
<sup>2</sup> Microsoft

## ABSTRACT

Traditional automatic speech recognition (ASR) systems usually focus on individual utterances, without considering long-form speech with useful historical information, which is more practical in real scenarios. Simply attending longer transcription history for a vanilla neural transducer model shows no much gain in our preliminary experiments, since the prediction network is not a pure language model. This motivates us to leverage the factorized neural transducer structure, containing a real language model, the vocabulary predictor. We propose the *LongFNT-Text* architecture, which fuses the sentence-level long-form features directly with the output of the vocabulary predictor and then embeds token-level long-form features inside the vocabulary predictor, with a pre-trained contextual encoder RoBERTa to further boost the performance. Moreover, we propose the *LongFNT* architecture by extending the long-form speech to the original speech input and achieve the best performance. The effectiveness of our LongFNT approach is validated on LibriSpeech and GigaSpeech corpora with 19% and 12% relative word error rate (WER) reduction, respectively.

**Index Terms**— long-form speech recognition, factorized neural transducer, context and speech encoder

## 1. INTRODUCTION

End-to-end (E2E) automatic speech recognition (ASR) models, such as connectionist temporal classification (CTC) [1], attention-based encoder-decoder (AED) [2, 3], and recurrent neural network transducer neural transducer (RNN-T) [4] are now dominating over traditional hybrid models [5]. A common practice is to train the E2E model with individual utterances without considering the correlation between utterances. However, real scenarios such as conversations, videos, and meetings are usually under long-form situations, which induces a significant discrepancy between training and test. To bridge the gap, long-form ASR, also known as conversational ASR, dialog-aware ASR, or large-context ASR, is proposed to capture the relationship of transcription history.

Previous approaches explored the long-form scenario in attention-based encoder-decoder models (AED) [3]. The most intuitive way is to concatenate consecutive speech or transcriptions of utterances [6, 7]. Hori et al. [8] extended their prior work to accelerate the decoding process in the streaming AED architecture. Furthermore, consecutive long-form transcriptions can be used in recurrent neural language models [9, 10] or re-scoring using BERT [11]. A more comprehensive way is to use an auxiliary encoder to model the context information in an AED manner. Masumura et al. [12, 13]

proposed a hierarchical text encoder and the distillation of large-context knowledge not limited to the current utterance, so that the long-form information can be captured while retaining the ASR model performance. Likewise, Wei et al. [14, 15, 16] utilized a latent variational module, context-aware residual attention, and pre-trained encoders to leverage acoustic and text context.

Recently, transducer-based systems are becoming more and more popular in industry [4, 17, 18, 19, 20], since it is naturally streaming with low latency, and somehow more robust than attention-based systems [21, 5]. However, there is little prior work exploring long-form neural transducer models. Narayanan et al. [22] have done primitive explorations by simulating long-form training and adaptation to obtain improvement using short utterances. Schwarz et al. [23] showed that the combination of input and context audio makes the network learn both speaker and environment adaptations. Kojima [24] explored the utilization of large context. However, how to incorporate consecutive transcription history into the neural transducer model is still not well explored.

In this paper, we propose the novel *LongFNT* to integrate long-form information into factorized neural transducer (FNT) [25] architecture to solve the above challenge. We first tried to embed long-form text into the predictor of the vanilla neural transducer, but our experiments showed that adding long-form transcriptions to the predictor part has little impact to the performance. A possible explanation is that the prediction network in the vanilla neural transducer does not act as a pure language model (LM), which limits its capability in long-form transcription modeling. It also indicates that effective methods for AED-based models [26, 27] cannot be extended to transducer-based ones, as they heavily rely on the LM characteristic of the decoder. Therefore, we try to utilize long-form transcriptions using the FNT architecture, which factorizes the blank and vocabulary prediction modules so that a standalone LM can be used for vocabulary prediction. Based on the vocabulary predictor in the FNT architecture, we propose two approaches to integrating high-level long-form features from historical transcriptions. Concretely, a context encoder is employed to yield the embedding of each token in historical transcriptions. Then we obtain a history embedding by an averaged pooling operation and fuse it to the output of the vocabulary predictor, which is called *sentence-level integration*. Meanwhile, we incorporate the token embedding sequence inside the vocabulary predictor via contextual attention, which is named as *token-level integration*. We find that the two approaches complement each other to achieve better performance, referred as *LongFNT-Text*. Moreover, we employ a pre-trained text encoder, RoBERTa [28], to further improve long-form ASR performance. Finally, we embed long-form speech into the encoder to propose the *LongFNT-Speech* model, and the combination of LongFNT-Text and -Speech is called

\*Work done during an internship at Microsoft.

*LongFNT* model. The proposed method with long-form text transcriptions achieved 17% and 9% relative word error rate (WER) reduction on LibriSpeech and GigaSpeech, respectively. By further adding long-form speech, the final improvement reached 19% and 12% relative WER reduction on these two tasks, respectively.

## 2. NEURAL TRANSDUCER

### 2.1. Conformer Transducer

Conformer [20] is a convolutional augmented transformer that is widely used in attention-based encoder-decoder and neural transducer architectures. The conformer transducer (C-T) model consists of three parts, the conformer encoder, the joint network, and the predictor network:

$$\mathbf{h}_t = \text{Encoder}(\mathbf{x}_{\leq t}), \quad (1)$$

$$\mathbf{z}_l = \text{Predictor}(\mathbf{y}_{\leq l}), \quad (2)$$

$$\mathbf{z}_{t,l} = \text{Joint}(\mathbf{h}_t, \mathbf{z}_l) \quad (3)$$

where  $t, l$  are the frame and label index, respectively. The predicted probability of the neural transducer model and loss can be computed as:

$$P_{ASR}(\hat{y}_{l+1} | \mathbf{x}_{\leq t}, y_l) = \text{softmax}(\mathbf{z}_{t,l}), \quad (4)$$

$$\mathcal{L}_{transducer} = -\log \sum_{\alpha \in \eta^{-1}(\mathbf{y})} P(\alpha | \mathbf{x}), \quad (5)$$

where  $\eta$  is a many-to-one function from all possible transducer paths to the target  $\mathbf{y}$ .

### 2.2. Modified Factorized Neural Transducer (M-FNT)

The factorized neural transducer model (FNT) [25] aims to separately predict the blank token and vocabulary tokens, so that the vocabulary predictor fully functions as an LM. The FNT model contains four main parts, the conformer encoder, the blank predictor ( $\text{Pred}^B$ ), the joint network for  $\langle \text{blank} \rangle$ , and the vocabulary predictor ( $\text{Pred}^V$ , i.e. LM manner). A main drawback of FNT is that the accuracy is slightly worse than the standard RNN-T. To tackle this issue, Zhao et al. [29] proposed some improvements for FNT, shown in Figure 1(a) by fusing the acoustic and label representations in the following way:

$$\mathbf{z}_{t,l}^B = \text{Joint}(\mathbf{h}_t, \text{Pred}^B(\mathbf{y}_{\leq l})), \quad (6)$$

$$\begin{aligned} \mathbf{z}_l^V &= \log \text{softmax}(\text{Pred}^V(\mathbf{y}_{\leq l})) \\ &= \log P_{LM}(\hat{y}_{l+1} | \mathbf{y}_{\leq l}), \end{aligned} \quad (7)$$

$$\begin{aligned} \mathbf{z}_t^V &= \log \text{softmax}(\text{Proj.}(\mathbf{h}_t)), \\ &= \log P_{CTC}(\pi_t | \mathbf{x}_{\leq t}), \end{aligned} \quad (8)$$

$$\mathbf{z}_{t,l}^V = \mathbf{z}_t^V + \beta \mathbf{z}_l^V \quad (9)$$

where  $P_{LM}$  is the predicted probability of the LM,  $P_{CTC}$  is the posterior of the CTC sequence,  $\beta$  is a learnable parameter. And we can compute the posterior of the transducer model:

$$P_{ASR}(\hat{y}_{l+1} | \mathbf{x}_{\leq t}, y_l) = \text{softmax}([\mathbf{z}_{t,l}^B, \mathbf{z}_{t,l}^V]) \quad (10)$$

Finally, the modified factorized neural transducer loss is computed as

$$\mathcal{L} = \mathcal{L}_{transducer} + \lambda_{LM} \mathcal{L}_{LM} + \lambda_{CTC} \mathcal{L}_{CTC}, \quad (11)$$

where  $\lambda_{LM}, \lambda_{CTC}$  are hyper-parameters.

## 3. LONGFNT: LONG-FORM FACTORIZED NEURAL TRANSDUCER

### 3.1. Long-form Context Encoder

The long-form context encoder converts consecutive historical sentences  $\{\dots, \mathbf{Y}^{p-2}, \mathbf{Y}^{p-1}\}$  into long-form contextual embeddings  $\mathbf{C}$ :

$$\mathbf{C} = \text{Context-Encoder}(\dots, \mathbf{Y}^{p-2}, \mathbf{Y}^{p-1}), \quad (12)$$

where the current sentences index is  $p$ . The basic context encoder is jointly trained with long-form FNT in a transformer manner. To extract stronger features, we directly use a pre-trained RoBERTa [28] model as the context encoder. The RoBERTa<sup>1</sup> is pre-trained using 160GB training text and is frozen during ASR training.

### 3.2. LongFNT-Text: Long-form Context Integration

Here we explore different integration designs for FNT. As shown before, FNT factorizes out the vocabulary predictor part by jointly training the speech-text pair data, therefore the historical transcriptions can be injected inside the vocabulary predictor  $\text{Pred}^V$  or after it.

**Sentence-level integration:** To integrate it in a statistical way, we first do mean and standard variance (mean+std) of contextual embedding sequence  $\mathbf{C}$  to get sentence-level  $\tilde{\mathbf{c}}$ , with  $\text{Pred}^V$ 's output  $\mathbf{z}_l^V$  shown in the yellow box of Figure 1 (b), which needs an auxiliary linear layer to match the dimension:

$$\mathbf{z}_l^{\prime V} = \mathbf{z}_l^V + \text{Linear}(\tilde{\mathbf{c}}). \quad (13)$$

Moreover, the contextual embedding  $\tilde{\mathbf{c}}$  can be further utilized at the linear layer  $\text{Linear}^V$  of the vocabulary predictor  $\text{Pred}^V$ :

$$\mathbf{p} = \text{Pred-Encoder}^V(\mathbf{y}_{\leq l}), \quad (14)$$

$$\mathbf{P} = \text{Linear}^V \cdot \text{ReLU}(\mathbf{p}), \quad (15)$$

$$\mathbf{P}' = \text{Linear}^V \cdot \text{ReLU}(\mathbf{p} \oplus \text{Projection}(\tilde{\mathbf{c}})), \quad (16)$$

where  $\mathbf{z}_l^{\prime V} = \log \text{softmax}(\mathbf{P} \text{ or } \mathbf{P}')$  to match Equation 7, and  $\oplus$  is implemented as concatenation.

**Token-level integration:** The second method is to add an auxiliary cross-attention layer inside  $\text{Pred}^V$  transformer blocks:

$$\mathbf{p}^i = \text{MHA}(\mathbf{p}^i, \mathbf{p}^i, \mathbf{p}^i), \quad (17)$$

$$\mathbf{p}^i = \text{MHA}(\mathbf{p}^i, \mathbf{C}, \mathbf{C}), \quad (18)$$

$$\mathbf{p}^i = \text{FFN}(\mathbf{p}^i), \quad (19)$$

where  $i$  is the layer index, FFN is the feed-forward layer inside the transformer encoder layer and MHA is the multi-head attention layer, respectively.

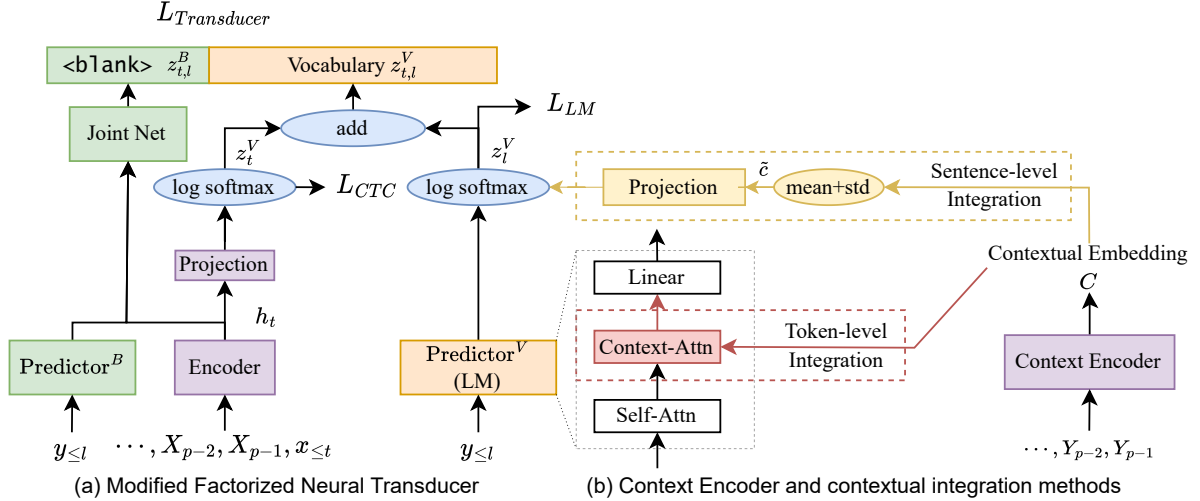
Finally, the sentence-level integration and the token-level integration can be combined to achieve better utilization of contextual embedding sequence  $\mathbf{C}$ , which is named as **LongFNT-Text**.

### 3.3. LongFNT-Speech: Long-form Speech Encoder

Motivated by [23], we can propose the LongFNT-Speech encoder with extended historical speech  $\{\dots, \mathbf{X}^{p-2}, \mathbf{X}^{p-1}\}$  to utilize the long-form speech information:

$$\dots, \mathbf{h}^{p-1}, \mathbf{h}_{\{1:t\}}^p = \text{Encoder}(\dots, \mathbf{X}^{p-1}, \mathbf{X}_{\{1:t\}}^p), \quad (20)$$

<sup>1</sup><https://huggingface.co/sentence-transformers/all-roberta-large-v1>



**Fig. 1.** Architecture of modified factorized neural transducer (M-FNT) and the long-form factorized neural transducer. Two integration methods are proposed for long-form text, and the speech encoder is extended for long-form speech.

where  $\mathbf{h} = \mathbf{h}^p$  matches Equation 1. More specifically, historical speech is only used in attention score computation, and the label representations  $z_t^V$  is calculated by  $\mathbf{h}_t^p$  as in Equation 8. Using such extension, the speech encoder receives a longer history and thus benefits training and evaluation. Combining LongFNT-Text and LongFNT-Speech, the final proposed method is called as **LongFNT**.

## 4. EXPERIMENTS

### 4.1. Experimental Setup

We conduct experiments with two datasets, LibriSpeech [30] and GigaSpeech Middle (abbr. as GigaSpeech) [31]. LibriSpeech has around 960 hours of audiobook speech, while GigaSpeech has around 1,000 hours of audiobook, podcasting, and YouTube audio. The word error rate (WER) averaged over each test set is reported. For acoustic feature extraction, 80-dimensional mel filterbank (Fbank) features are extracted with global level cepstral mean and variance normalization. Frame length and frame shift are 25ms and 10ms respectively. Standard SpecAugment [32] is used for both datasets, respectively. Each utterance has two frequency masks with parameter ( $F = 27$ ) and ten time masks with maximum time-mask ratio ( $pS = 0.05$ ). 5,000 sentence pieces [33] are trained using LibriSpeech and GigaSpeech datasets separately. The baseline follows the settings of FNT [25, 29]. The subsampling layer is a VGG2L-like network, which contains four convolution layers with the down-sampling rate of 4. The encoder has 18 conformer layers, in which the inner size of the feed-forward layer is 1,024, and the attention dimension is 512 with 8 heads. The  $\text{Pred}^B$  has two LSTM layers with 1,024 hidden size and the joint dimension is set to 512. The  $\text{Pred}^V$  use vanilla 8 transformer layers, which has 256 attention dimension with 8 heads. The hyper-parameter weights are fixed as  $\lambda_{CTC} = 0.1$ ,  $\lambda_{LM} = 0.5$ . The context encoder has the same shape as  $\text{Pred}^V$  if training from scratch, and using the frozen RoBERTa model else. As for the FNT's large LM adaptation, we use 16 transformer layers with 512 attention dimension with 8 heads. For LibriSpeech, we use the extra text corpus, and for GigaSpeech, we use the 10,000 hours of training text data.

### 4.2. Evaluation Results

Model	$\mathcal{M}$	Libri-test		Giga	
		clean	other	dev	test
C-T	-	3.1	6.6	16.1	15.7
+ long text	gt	3.1	6.5	15.8	15.5
+ long text	hyp	3.1	6.7	16.1	15.9
+ sentence-level integ.	gt	3.1	6.5	16.0	15.5
+ sentence-level integ.	hyp	3.2	6.9	16.4	16.2
M-FNT	-	3.2	6.4	16.8	16.3
+ long text	gt	3.2	6.3	16.5	16.2
+ long text	hyp	3.2	6.4	16.7	16.4
+ large LM	-	3.0	6.1	16.4	16.0
+ large LM and long text	hyp	3.0	6.1	16.4	16.1
LongFNT-Text	hyp	2.5	5.5	15.2	14.9
LongFNT-Speech	hyp	2.8	6.0	15.9	15.7
LongFNT	hyp	<b>2.4</b>	<b>5.4</b>	<b>14.8</b>	<b>14.3</b>

**Table 1.** LongFNT performance (WER) (%) on LibriSpeech test sets and GigaSpeech dev/test sets.  $\mathcal{M}$  has two modes, ‘gt’ denotes the historical text is obtained from ground truth, while ‘hyp’ is from the decoded hypotheses.

We first explore whether the vanilla C-T model can be improved by the long transcription history, as shown in the first block of Table 1. It indicates that the C-T model reaches little gain with extended long-form text or the sentence-level integration method. If we use it in real cases (i.e. the previous text is decoded instead of ground truth), the performance even gets worse. Our explanation is the predictor network is not a pure language model, so it cannot utilize longer history.

In terms of M-FNT, shown in the second block of Table 1, enlarging the text input (+‘long text’) does a little help to the GigaSpeech corpus, but no more gain to LibriSpeech. When the predictor is trained with unpair text data (+‘large LM’), the performance of all models is increased by at least 2%. It demonstrates that

the M-FNT architecture can be benefited from a powerful language model. However, the long-form text does not benefit from large LM either, which means it is non-trivial to explore how to better leverage long-form information in the context of FNT.

Our proposed LongFNT model achieves 19/12% rel. WERR compared to the M-FNT baseline and 20/9% rel. WERR compared to the C-T baseline. It demonstrates that LongFNT can better utilize long-form history information. Furthermore, we can observe both LongFNT-Text and LongFNT-Speech are better than baselines, indicating two types of historical information are helpful for long-form speech recognition. The LongFNT-Text is slightly better than the LongFNT-Speech model, showing the transcription history is more valuable to our model. Compared to the limited gain from the long text in C-T and M-FNT, the WER reduction from long text is relatively larger, around 10% (LongFNT v.s. LongFNT-Speech), verifying the effectiveness of the proposed method.

### 4.3. Ablation Study

Model	$\mathcal{M}$	Libri-test		Giga	
		clean	other	dev	test
LongFNT-Text - <i>RoBERTa</i>	hyp	2.5	5.5	15.2	14.9
	hyp	2.5	5.6	15.6	15.2
- token-level integ. - - <i>RoBERTa</i>	hyp	2.8	5.8	15.9	15.3
	hyp	2.9	5.9	16.0	15.5
- sentence-level integ. - - <i>RoBERTa</i>	hyp	2.6	5.6	15.4	15.0
	hyp	2.6	5.7	15.8	15.4
LongFNT-Text (gt or hyp) w/o large LM and <i>RoBERTa</i>					
- - token-level integ.	gt	2.9	6.1	16.0	15.8
	hyp	3.0	6.3	16.5	16.2
- - sentence-level integ.	gt	2.9	6.0	15.7	15.4
	hyp	3.0	6.1	16.0	15.7

**Table 2.** The ablation study of different components in LongFNT-Text.

In this subsection, we validate the effectiveness of the proposed LongFNT model to explore which part is the most important. Shown in Table 2, we explore the efficiency of different integration methods for LongFNT-Text based on M-FNT or LongFNT-Speech architecture. The token-level integration is more important than sentence-level integration, as after removing token-level integration (i.e. only sentence-level integration), the system degrades by 12/5% relative WER reduction in LibriSpeech and 5/3% relative WER reduction in GigaSpeech while removing sentence-level integration only degrades  $\sim 4\%$  and  $\sim 1\%$  in LibriSpeech/GigaSpeech corpus.

Then, we also explore the importance of contextual encoder replacement from the train-from-scratch one to the pre-trained RoBERTa. Results show that the importance is different from LibriSpeech and GigaSpeech datasets. For LibriSpeech, we can see that the performance only degrades by 0.1 absolute WER reduction, and even no degradation can be found in the test-clean set for the LongFNT-Text model. As for GigaSpeech, we can observe a consistent drop by at least 0.3 $\sim$ 0.5 absolute WER reduction. This interesting phenomenon is that the RoBERTa model has little impact on the LibriSpeech which may be because the influence of transcriptions is relatively smaller than the GigaSpeech dataset, and the contextual encoder has a similar ability to model the long-form textual information.

We also evaluate our models when the large pre-trained predictor network is absent. Detailed training data and model description can be found in the setup Section 4.1. After replacing the larger in-domain LM with a trained-from-scratch  $\text{Pred}^V$ , the performance of all models decreases by at least 2%, in the third block of Table 2. Furthermore, we find that long-form transcriptions are further utilized with large LM. Strong large LM also outperforms the original one in both integration methods.

As mentioned previously in Section 4.2, in the real scenario, ground truth (gt) text can not be accessed, and we evaluate the above methods using decoded transcriptions (hyp) to get real performance and explore the importance of those two types in different LongFNT-Text modes. Shown in the 3rd block of Table 2, models that is decoded using ‘(hyp)’ drop the performance by 1 $\sim$ 5% compared to the ‘(gt)’ one. Results show the degradation is much higher on GigaSpeech compared with on LibriSpeech. One possible reason is that the basic error rate influences the performance of hypo mode, a system with low WER is necessary for the long-form speech recognition improvement. Moreover, we find that the sentence-level integration (i.e. -- token-level integ.) is more sensitive to the long-form transcriptions compared to the token-level one (i.e. -- sentence-level integ.), as it drops more 0.1 $\sim$ 0.2 absolute WER. And this indicates the shortcoming of statistical averaging pooling for sentence-level integration.

### 4.4. The Number of Long-form Sentences

#Previous	1	2	3
+ sentence-level integ.	16.6/16.2	16.4/16.1	16.3/16.1
++ speech	16.0/15.7	15.5/15.1	15.2/15.9

**Table 3.** Performance (WER) (%) comparison of successive sentence counts. ‘++ speech’ here means use (LongFNT-Speech with sentence-level integration). The baseline performance is 16.8/16.3%, which means # Previous=0 (no historical sentence).

From Table 3, we can find the correlation between the number of historical sentences and the recognition error rate. As #Previous grows, the WER of the current sentence is reduced. However, after #Previous > 2, the improvement is limited but training resources are consumed, where the most important information can be achieved in the previous two sentences. Therefore, we select two previous consecutive sentences to learn appropriate long-form information for all experiments.

## 5. CONCLUSION

In this paper, we propose a novel long-form architecture **LongFNT** based on factorized neural transducer (FNT) architecture. Long-form sentence-/token-level transcription integration methods are proposed with pre-trained RoBERTa to acquire the LongFNT-Text model. LongFNT-Text achieve 17% relative WER reduction on LibriSpeech test sets and 9% relative WER reduction on GigaSpeech corpus. Then LongFNT-Speech is proposed to integrate long-form speech. And LongFNT is obtained combining LongFNT-Text and LongFNT-Speech. After utilizing history speech, the final system archives 19% relative WER reduction on LibriSpeech test sets and 12% relative WER reduction on GigaSpeech corpus. In future work, we will explore streaming mode FNT, and more efficient training and testing methods to improve long-form speech.

## 6. REFERENCES

- [1] Alex Graves, Santiago Fernández, et al., “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proc. ICML*, 2006, pp. 369–376.
- [2] Ashish Vaswani, Noam Shazeer, et al., “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, Dec. 2017.
- [3] Takaaki Hori, Shinji Watanabe, and John R Hershey, “Joint CTC/attention decoding for end-to-end speech recognition,” in *Proc. ACL*, 2017, pp. 518–529.
- [4] Alex Graves, “Sequence transduction with recurrent neural networks,” *Proc. ICML*, Nov. 2012.
- [5] Jinyu Li, “Recent advances in end-to-end automatic speech recognition,” *APSIPA Transactions on Signal and Information Processing*, vol. 11, no. 1, 2022.
- [6] Suyoun Kim and Florian Metze, “Dialog-Context Aware end-to-end Speech Recognition,” in *Proc. SLT*, Dec. 2018, pp. 434–440.
- [7] Takaaki Hori, Niko Moritz, et al., “Transformer-Based Long-Context End-to-End Speech Recognition,” in *Proc. Interspeech*, Oct. 2020, pp. 5011–5015.
- [8] Takaaki Hori, Niko Moritz, Chiori Hori, and Jonathan Le Roux, “Advanced long-context end-to-end speech recognition using context-expanded transformers,” in *Proc. Interspeech*, 2021, pp. 2097–2101.
- [9] Kazuki Irie, Albert Zeyer, et al., “Training language models for long-span cross-sentence evaluation,” in *Proc. ASRU*, Dec. 2019, pp. 419–426.
- [10] Tomas Mikolov and Geoffrey Zweig, “Context dependent recurrent neural network language model,” in *Proc. SLT*, Miami, FL, USA, Dec. 2012, pp. 234–239.
- [11] Shih-Hsuan Chiu, Tien-Hang Lo, et al., “Cross-sentence Neural Language Models for Conversational Speech Recognition,” in *Proc. IJCNN*, July 2021, pp. 1–7.
- [12] Ryo Masumura, Tomohiro Tanaka, et al., “Large Context End-to-end Automatic Speech Recognition via Extension of Hierarchical Recurrent Encoder-decoder Models,” in *Proc. ICASSP*, May 2019, pp. 5661–5665.
- [13] Ryo Masumura, Naoki Makishima, et al., “Hierarchical transformer-based large-context end-to-end asr with large-context knowledge distillation,” in *Proc. ICASSP*, 2021, pp. 5879–5883.
- [14] Kun Wei, Yike Zhang, Sining Sun, Lei Xie, and Long Ma, “Conversational speech recognition by learning conversation-level characteristics,” in *Proc. ICASSP*, 2022, pp. 6752–6756.
- [15] Kun Wei, Yike Zhang, Sining Sun, Lei Xie, and Long Ma, “Leveraging Acoustic Contextual Representation by Audio-textual Cross-modal Learning for Conversational ASR,” in *Proc. Interspeech*, 2022, pp. 1016–1020.
- [16] Kun Wei, Pengcheng Guo, and Ning Jiang, “Improving Transformer-based Conversational ASR by Inter-Sentential Attention Mechanism,” in *Proc. Interspeech*, 2022, pp. 3804–3808.
- [17] Ching-Feng Yeh, Jay Mahadeokar, Kaustubh Kalgaonkar, Yongqiang Wang, Duc Le, Mahaveer Jain, Kjell Schubert, Christian Fuegen, and Michael L Seltzer, “Transformer-transducer: End-to-end speech recognition with self-attention,” *arXiv preprint arXiv:1910.12977*, 2019.
- [18] Qian Zhang, Han Lu, et al., “Transformer transducer: A streamable speech recognition model with transformer encoders and RNN-T loss,” in *Proc. ICASSP*, 2020, pp. 7829–7833.
- [19] Xie Chen, Yu Wu, Zhenghao Wang, Shujie Liu, and Jinyu Li, “Developing real-time streaming transformer transducer for speech recognition on large-scale dataset,” in *Proc. ICASSP*. IEEE, 2021, pp. 5904–5908.
- [20] Anmol Gulati, James Qin, et al., “Conformer: Convolution-augmented transformer for speech recognition,” in *Proc. Interspeech*, May 2020, pp. 5036–5040.
- [21] Chung-Cheng Chiu, Wei Han, et al., “A Comparison of End-to-End Models for Long-Form Speech Recognition,” in *Proc. ASRU*, Dec. 2019, pp. 889–896.
- [22] Arun Narayanan, Rohit Prabhavalkar, et al., “Recognizing Long-Form Speech Using Streaming End-to-End Models,” in *Proc. ASRU*, Dec. 2019, pp. 920–927.
- [23] Andreas Schwarz, Ilya Sklyar, and Simon Wiesler, “Improving RNN-T ASR Accuracy Using Context Audio,” in *Proc. Interspeech*, 2021, pp. 1792–1796.
- [24] Atsushi Kojima, “Large-context automatic speech recognition based on rnn transducer,” in *2021 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2021, pp. 460–464.
- [25] Xie Chen, Zhong Meng, et al., “Factorized neural transducer for efficient language model adaptation,” in *Proc. ICASSP*, 2022, pp. 8132–8136.
- [26] Rami Botros, Tara N. Sainath, et al., “Tied & reduced rnn-t decoder,” in *Proc. Interspeech*, Aug. 2021, pp. 4563–4567.
- [27] Mohammadreza Ghodsi, Xiaofeng Liu, et al., “Rnn-transducer with stateless prediction network,” in *Proc. ICASSP*, May 2020, pp. 7049–7053.
- [28] Yinhan Liu, Myle Ott, et al., “Roberta: A robustly optimized bert pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [29] Rui Zhao, Jian Xue, et al., “Fast and accurate factorized neural transducer for text adaption of end-to-end speech recognition models,” in *submitted to Proc. ICASSP 2023*, 2022.
- [30] Vassil Panayotov, Guoguo Chen, et al., “Librispeech: An ASR corpus based on public domain audio books,” in *Proc. ICASSP*, 2015, pp. 5206–5210.
- [31] Guoguo Chen, Shuzhou Chai, et al., “GigaSpeech: An Evolving, Multi-Domain ASR Corpus with 10,000 Hours of Transcribed Audio,” in *Proc. Interspeech*, 2021, pp. 3670–3674.
- [32] Daniel S. Park, William Chan, et al., “SpecAugment: A simple data augmentation method for automatic speech recognition,” *Proc. Interspeech*, pp. 2613–2617, Sept. 2019.
- [33] Taku Kudo and John Richardson, “Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2018, pp. 66–71.