



Insertion-Based Modeling for End-to-End Automatic Speech Recognition

Yuya Fujita¹, Shinji Watanabe², Motoi Omachi¹, Xuankai Chang²

¹Yahoo Japan Corporation, Tokyo, Japan

²Center for Language and Speech Processing, Johns Hopkins University, Baltimore, MD, USA

yuyfujit@yahoo-corp.jp

Abstract

End-to-end (E2E) models have gained attention in the research field of automatic speech recognition (ASR). Many E2E models proposed so far assume left-to-right autoregressive generation of an output token sequence except for connectionist temporal classification (CTC) and its variants. However, left-to-right decoding cannot consider the future output context, and it is not always optimal for ASR. One of the non-left-to-right models is known as non-autoregressive Transformer (NAT) and has been intensively investigated in the area of neural machine translation (NMT) research. One NAT model, mask-predict, has been applied to ASR but the model needs some heuristics or additional component to estimate the length of the output token sequence. This paper proposes to apply another type of NAT called insertion-based models, that were originally proposed for NMT, to ASR tasks. Insertion-based models solve the above mask-predict issues and can generate an arbitrary generation order of an output sequence. In addition, we introduce a new formulation of joint training of the insertion-based models and CTC. This formulation reinforces CTC by making it dependent on insertion-based token generation in a non-autoregressive manner. We conducted experiments on three public benchmarks and achieved competitive performance to strong autoregressive Transformer with a similar decoding condition.

Index Terms: Transformer, speech recognition, end-to-end, non-autoregressive

1. Introduction

End-to-end (E2E) models have become mainstream in the research field of automatic speech recognition (ASR). One advantage of the E2E models is the simplicity of the model structure. A single neural network receives an acoustic feature sequence and directly generates an output token sequence. It does not need separate models such as acoustic, language and lexicon models commonly used in the conventional ASR system. There has been a lot of work that aims to improve E2E models [1–6]. Several reported that the E2E model achieved comparable or better performance to the conventional ASR system in the product system [7, 8] and publicly available corpora [9–11] such as Librispeech [12], Switchboard [13], and Corpus of Spontaneous Japanese (CSJ) [14]. One of the state-of-the-art E2E models is Transformer [15], which significantly outperformed the RNN-based E2E models [11, 16].

Many of the above E2E models assume left-to-right autoregressive generation of an output token sequence. In the speech production context, this assumption is reasonable because speech is produced in a left-to-right order given linguistic content. But in the speech perception context, it is unclear that left-to-right decoding is always the best way. For example, when we listen to speech and encounter a word whose pronunci-

ation is unclear, we leave it as uncertain and re-estimate it using the future context. Mimicking this perceptual process in ASR is scientifically quite important and also has some potential to improve the performance of left-to-right decoding.

One of the non-left-to-right E2E models is known as non-autoregressive Transformer (NAT). It is heavily investigated in the area of neural machine translation (NMT) [17–20]. Mask-predict, one of the NAT models, has been applied to speech recognition [21]. It realizes non-autoregressive output token generation by introducing a special token which masks part of an output token sequence. When training the model, some tokens in the output sequence are randomly masked and the model is trained to estimate the masked tokens with non-causal masking in the self-attention of the decoder. During decoding, the output token sequence is generated by estimating masked tokens. There is no causal masking hence non-left-to-right non-autoregressive output sequence generation is realized. It achieved competitive performance to an autoregressive model with faster decoding time on AISHELL [22].

However, mask-predict needs some additional component or heuristics to estimate the output token sequence length. To overcome this disadvantage, the insertion-based model is proposed [18]. In theory this model can generate an output token sequence with an arbitrary order without any additional component or heuristics to estimate the output token sequence length. In NMT, performance competitive to autoregressive Transformer is reported with fewer iterations in decoding [20].

This paper proposes using the insertion-based models for E2E ASR with an in-depth investigation of three insertion-based models originally proposed in NMT. In addition, we introduce a new formulation of joint modeling of connectionist temporal classification (CTC) [23] and insertion-based models. This formulation can be viewed as modeling the joint distribution between the CTC probability and the insertion-based sequence generation probability. Hence the CTC probability depends on the insertion-based output token sequence generation. With this new formulation, the monotonic alignment property of CTC is reinforced by insertion based token generation. It achieves performance competitive with an autoregressive left-to-right model decoded with a similar decoding condition in non-left-to-right and non-autoregressive manner.

The source code will be publicly available in the open source E2E modeling toolkit ESPnet [24].

2. Related work

To the best of our knowledge, this is the first work to apply insertion-based models to ASR tasks.

As mentioned in Section 1, this work is another type of NAT application to ASR tasks compared with [21]. Our work does not use a special mask token or need to estimate an output sequence length in advance. Furthermore, our work can handle

autoregressive and non-autoregressive models in a single formulation and also introduces a new formulation of joint modeling of CTC and insertion-based models.

Non-autoregressive E2E ASR using a CTC-like model is proposed in Imputer [25]. It assumes that the alignment at the i -th generation step depends on the past $(i - 1)$ -th alignment. The alignment is estimated similarly to mask-predict in a non-autoregressive manner. Our work is different in using insertion-based models and joint modeling of CTC and insertion-based token sequence generation.

3. Insertion-based end-to-end models

Let $X = \{\mathbf{x}_t \in \mathbb{R}^d | t = 1, \dots, T\}$ be a d -dimensional acoustic feature sequence and $C = \{c_n \in \mathcal{V} | n = 1, \dots, N\}$ be an output token sequence. T is the input length, \mathcal{V} is a set of distinct tokens, and N is the output length. Then, decoding of the E2E model is performed to maximize the posterior probability $p^{e2e}(C|X)$:

$$\hat{C} = \arg \max_C p^{e2e}(C|X). \quad (1)$$

Training of the E2E model is also based on this criterion. The difference between various E2E models is how to define the posterior $p^{e2e}(C|X)$ in Eq. (1).

In the insertion-based models, it is assumed to be marginalized over all possible insertion orders (permutation). Let $Z = \{z_n \in \mathbb{N}^+ | n = 1, \dots, N\}$ be an insertion order. For example, suppose $C = \{\text{this, is, a, pen}\}$, Z is all the permutations of ordering 4 tokens, i.e.,

$$Z \in \{\{4, 1, 2, 3\}, \{4, 2, 1, 3\}, \dots, \{1, 2, 3, 4\}\}, \quad (2)$$

$$\begin{cases} C^{Z=\{4,1,2,3\}} = \{\text{pen, this, is, a}\} \\ C^{Z=\{4,2,1,3\}} = \{\text{pen, is, this, a}\} \\ \vdots \\ C^{Z=\{1,2,3,4\}} = \{\text{this, is, a, pen}\} \end{cases}, \quad (3)$$

where $C^Z = \{c_{z_n} \in \mathcal{V} | n = 1, \dots, N\}$ is the permuted output token sequence with an insertion order Z . The number of all permutations $|Z|$ is ${}_4P_4$. Then, the posterior in Eq. (1) is factorized with the sum and product rules as:

$$p^{e2e}(C|X) = \sum_Z p(C, Z|X) = \sum_Z p(C^Z|X)p(Z). \quad (4)$$

We assume that insertion order Z does not depend on input feature hence $P(Z|X) = P(Z)$.

Definition of $p(C^Z|X)$ in Eq. (4) is different between insertion-based models and explained in the next subsection. Note that the left-to-right autoregressive model can be interpreted as a special case where $p(Z)$ is fixed to be the left-to-right order $p(Z = \{1, 2, \dots, N\}) = 1$ and $p(C^Z|X)$ in Eq. (4) is $p(C^Z|X) = \prod_{i=1}^N p(c_n | c_{1:n-1}, X)$.

When training the model of Eq. (4), a lower bound of log-likelihood $\mathcal{L}(\theta)$ where θ is the parameters of the model is maximized under a predefined prior distribution over insertion order Z .

$$\log p^{e2e}(C|X) \geq \sum_Z p(Z) \log p(C^Z|X) =: \mathcal{L}(\theta) \quad (5)$$

From the next subsection, three existing insertion-based models are explained.

3.1. Insertion-based decoding

Insertion-based decoding (InDIGO) [26] is an insertion-based model using Transformer with relative position representation. Let $\mathbf{R}_n^Z \in \{-1, 0, +1\}^{n \times n}$ be a relative position representation at generation step n under an insertion order Z . Element r_{ij}^Z is defined as:

$$r_{ij}^Z = \begin{cases} -1 & z_j > z_i \\ 0 & z_j = z_i \\ 1 & z_j < z_i \end{cases}. \quad (6)$$

$p(C^Z|X)$ in Eq. (4) of InDIGO is defined as:

$$p(C^Z|X) =: p(C^Z, \mathbf{R}_N^Z|X) \quad (7)$$

$$= \prod_{n=1}^N p(c_n^Z, \mathbf{r}_n^Z | c_{1:n-1}^Z, \mathbf{R}_{n-1}^Z, X), \quad (8)$$

where \mathbf{r}_n^Z is the n -th column vector of \mathbf{R}_n^Z . The factorized form $p(c_n^Z, \mathbf{r}_n^Z | c_{1:n-1}^Z, \mathbf{R}_{1:n-1}^Z, X)$ in Eq. (8) is modeled by Transformer. Let $\mathbf{H}_{\text{dec}} \in \mathbb{R}^{b \times N}$ be the final output of the decoder layer of Transformer where b is the dimension of self-attention,

$$p(c_n^Z, \mathbf{r}_n^Z | c_{1:n-1}^Z, \mathbf{R}_{1:n-1}^Z, X) =: \underbrace{p(c_n^Z | \mathbf{H}_{\text{dec}})}_{\text{Word prediction}} \underbrace{p(\mathbf{r}_n^Z | c_n^Z, \mathbf{H}_{\text{dec}})}_{\text{Position prediction}}. \quad (9)$$

For the word prediction term, a linear transform and softmax operation are applied to \mathbf{H}_{dec} , and for the position prediction term a pointer network [27] is used.

During decoding, the next token to be inserted is estimated by the word prediction term then its position is estimated by the position prediction term in Eq. (9). Because of this sequential operation, only a single token is generated per iteration during decoding. Therefore, InDIGO can be non-left-to-right but requires N iterations.

3.2. Insertion Transformer and KERMIT

Another type of insertion-based model is Insertion Transformer [18] and KERMIT (Kontextuell Encoder Representations Made by Insertion Transformations) [20]. The basic formulation of these two models is the same. Let c_n^Z be a token to be inserted and l_n^Z be a position where the token is inserted at the n -th generation step under an insertion order Z . $p(C^Z|X)$ in Eq. (4) is defined as:

$$\begin{aligned} p(C^Z|X) &=: \prod_{n=1}^N p\left(\left(c_n^Z, l_n^Z\right) \mid \left(c_1^Z, l_1^Z\right), \dots, \left(c_{n-1}^Z, l_{n-1}^Z\right), X\right) \\ &= \prod_{n=1}^N p\left(\left(c_n^Z, l_n^Z\right) \mid c_{1:n-1}^Z, X\right), \end{aligned} \quad (10)$$

where $c_{1:n-1}^Z$ is the sorted token sequence at the n -th generation step. For example, in case of $C = \{\text{this, is, a, pen}\}$ and $Z = \{3, 1, 4, 2\}$, i.e., $C^Z = \{\text{a, this, pen, is}\}$,

$$\begin{aligned} n = 1 : (c_1^Z, l_1^Z) &= (\text{a}, 1), & c_{1:1}^Z &= \{\text{a}\}, \\ n = 2 : (c_2^Z, l_2^Z) &= (\text{this}, 1), & c_{1:2}^Z &= \{\text{this, a}\}, \\ n = 3 : (c_3^Z, l_3^Z) &= (\text{pen}, 3), & c_{1:3}^Z &= \{\text{this, a, pen}\}, \\ n = 4 : (c_4^Z, l_4^Z) &= (\text{is}, 2), & c_{1:4}^Z &= \{\text{this, is, a, pen}\}. \end{aligned}$$

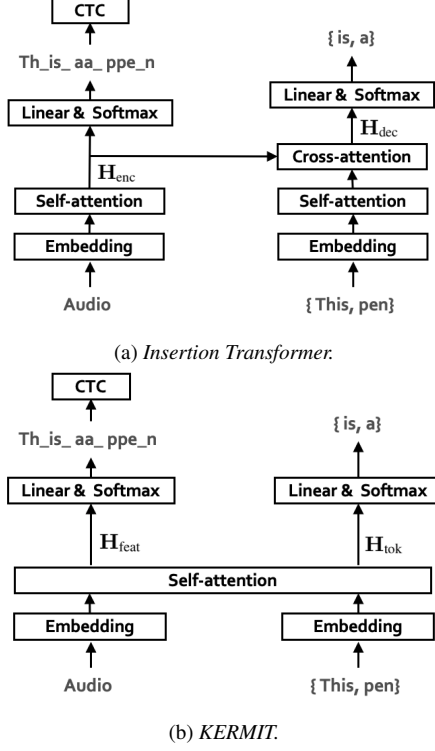


Figure 1: Schematic diagram of hybrid training with CTC. In the case of KERMIT, CTC is dependent on both audio and output token.

Note that l_n^Z is a position relative to the hypothesis at the $(n-1)$ -th generation step. In the case of $n = 3$ explained above, $l_n^Z \in \{1, 2\}$ because there are two tokens in the previous hypothesis.

The difference between Insertion Transformer and KERMIT is the matrix $\mathbf{H} \in \mathbb{R}^{b \times N}$ used when the posterior is calculated. In the case of Insertion Transformer, the final output of the decoder layer of Transformer $\mathbf{H}_{\text{dec}} \in \mathbb{R}^{b \times N}$ is used as \mathbf{H} . On the other hand, KERMIT uses only the encoder block of Transformer. Acoustic feature sequence and token embedding are concatenated and fed into the encoder block. The final output of the encoder layer is sliced as $\mathbf{H}_{\text{tok}} \in \mathbb{R}^{b \times N}$ then used as \mathbf{H} . This difference is depicted in Figure 1.

By using the \mathbf{H} matrix, the posterior in Eq. (10) is calculated as:

$$p\left(\left(c_n^Z, l_n^Z\right) \mid c_{1:n-1}^Z, X\right) =: \underbrace{p\left(c_n^Z \mid l_n^Z, \mathbf{H}\right)}_{\text{Word prediction}} \underbrace{p\left(l_n^Z \mid \mathbf{H}\right)}_{\text{Position prediction}}. \quad (11)$$

The word and position prediction term is calculated by a softmax followed by a linear transformation of \mathbf{H} .

There are two ways of decoding. The first one is autoregressive greedy decoding directly using the posterior in Eq. (11):

$$\left(\hat{c}, \hat{l}\right) = \arg \max_{c, l} p\left(\left(c, l\right) \mid c_{1:n-1}^Z, X\right). \quad (12)$$

The second way is non-autoregressive parallel decoding using only the word prediction term in Eq. (11):

$$\hat{c} = \arg \max_c p\left(c \mid l, c_{1:n-1}^Z, X\right). \quad (13)$$

When the balanced binary insertion order proposed in [18] is used as $p(Z)$, parallel decoding finishes empirically with $\log_2(N)$ iterations. This order is to insert centermost tokens of current hypothesis. For example, suppose $C = \{c_1, \dots, c_9\}$, then the hypothesis grows like $\{c_5\} \rightarrow \{c_3, c_5, c_7\} \rightarrow \{c_2, c_3, c_4, c_5, c_6, c_7, c_8\} \rightarrow \{c_1, c_2, c_3, c_4, c_5, c_6, c_7, c_8, c_9\}$.

4. Insertion-based/CTC joint modeling

Speech is generated in a left-to-right order hence the alignment is monotonic. Therefore, an E2E model trained with a CTC objective is reported to achieve faster convergence and high accuracy [5]. It is natural to apply this technique to insertion-based models. In the case of InDIGO and Insertion Transformer, its network is composed of an encoder and decoder so it is easy to apply this technique the same way as in [5]. However, for KERMIT, because it consists of only an encoder, a new formulation must be introduced.

Let Y be an output token sequence to be modeled by CTC. Usually, Y is set as $Y = C$. Joint modeling is to extend $P(C^Z|X)$ in Eq. (4) as:

$$P(C^Z|X) =: P(C^Z, Y|X) = P(Y|X, C^Z)P(C^Z|X). \quad (14)$$

The term $P(Y|X, C^Z)$ in Eq. (14) is modeled by CTC. Let A be a sequence of tokens extended with a blank symbol, $A = \{a_t \in \mathcal{V} \cup \{\text{blank}\} \mid t = 1, \dots, T\}$. $\mathcal{F}(\cdot)$ is a mapping function which deletes repetitions and the blank symbol from a sequence A hence $\mathcal{F}(A) = Y$. The CTC probability is formulated as:

$$p(Y|X, C^Z) =: \sum_{A \in \mathcal{F}^{-1}(Y)} p(A|X, C^Z). \quad (15)$$

In the case of InDIGO and Insertion Transformer, the final output of the encoder layer $\mathbf{H}_{\text{enc}} \in \mathbb{R}^{b \times T}$ is used to calculate $p(A|X, C^Z)$ in Eq. (15) as:

$$p(A|X, C^Z) \simeq p(A|X) =: p(A|\mathbf{H}_{\text{enc}}), \quad (16)$$

where $p(A|\mathbf{H}_{\text{enc}})$ is calculated by applying a linear transformation and softmax to \mathbf{H}_{enc} .

For KERMIT, $p(A|X, C^Z)$ in Eq. (15) can not be approximated as in Eq. (16). KERMIT consists of only an encoder and acoustic feature sequence and token embedding are concatenated then fed into the encoder block. Therefore, the output of the encoder block depends on both the acoustic feature and output token sequence. There might be several ways how to calculate $p(A|X, Z, C)$ in Eq. (15). In this work, output of KERMIT encoder is sliced as $\mathbf{H}_{\text{feat}} \in \mathbb{R}^{b \times T}$ and used:

$$p(A|X, C^Z) =: p(A|\mathbf{H}_{\text{feat}}). \quad (17)$$

This process is depicted in Figure 1(b). This formulation can reinforce CTC by making it dependent not only on the acoustic feature sequence but also on the output token sequence from insertion-based generation. Note that this formulation still retains non-autoregressive characteristics.

When training the model, in order to adjust the range of the two terms in Eq. (14), the CTC weight α is introduced as:

$$\begin{aligned} & \log P(Y|X, C^Z)P(C^Z|X) \\ & \simeq \alpha \log P(Y|X, C^Z) + (1.0 - \alpha) \log P(C^Z|X). \end{aligned} \quad (18)$$

During decoding, either the CTC part $p(Y|X, C^Z)$ or the insertion part $p(C^Z|X)$ in Eq. (14) can be used.

Table 1: CER of CSJ, AISHELL and WER of TEDLIUM2.

Model	Beam	Iterations	$p(Z)$	CTC weight α		CSJ 271h			TEDLIUM2		AISHELL	
				train	decode	Eval1	Eval2	Eval3	dev	test	dev	test
AT	10	$10N$	-	0.3	0.3	7.9	5.7	13.7	10.6	9.1	6.5	7.2
	1	N				8.1	5.4	13.9	12.7	10.1	6.7	8.1
InDIGO	1	N	L2R	0.0	0.0	8.4	6.2	14.7	-	-	-	-
				0.3		7.8	5.5	13.3	13.6	9.6	6.1	6.7
Insertion Transformer	1	N	L2R	0.0	0.0	8.7	6.3	16.1	-	-	-	-
				0.3		8.3	5.4	13.9	11.2	9.6	6.8	7.6
KERMIT	1	N	L2R	0.0	0.0	11.0	8.0	18.9	-	-	-	-
				0.9		9.2	6.7	15.7	14.9	12.4	7.7	8.9
					1.0	9.5	6.7	14.8	16.1	15.4	7.8	8.8
CTC	1	1	-	1.0	1.0	8.5	6.1	13.8	16.1	16.3	6.8	7.6
Insertion Transformer	1	$\simeq \log_2(N)$	BBT	0.0	0.0	15.0	12.4	21.6	-	-	-	-
				0.3		14.1	10.8	18.0	19.1	16.3	9.6	10.6
KERMIT	1	$\simeq \log_2(N)$	BBT	0.0	0.0	12.5	9.7	18.5	-	-	-	-
				0.9 (Proposed formulation)	1.0	11.5	9.1	16.7	18.8	15.0	9.8	10.9
						7.9	5.5	13.6	12.0	10.7	6.9	7.8

5. Experiments

5.1. Setup

We used three corpora, the Corpus of Spontaneous Japanese (CSJ) [14], TEDLIUM2 [28] and AISHELL [22]. As a baseline model we chose CTC [23], which is similar to the work in [29] using the Transformer encoder layers. Another baseline is autoregressive Transformer (AT) [16].

The three insertion-based models described in Section 3 are compared to the baseline. For parameters with these models, we followed the Transformer recipe of ESPnet [24] based on [16]. The numbers of layers for the encoder and decoder were 12 and 6, respectively. We increased the number of encoder layers to 18 for CTC and KERMIT because they are composed of only an encoder. We focused on two types of priors, left-to-right (L2R) and balanced binary tree (BBT) for $p(Z)$ in Eq. (4) to simplify the comparison. L2R is evaluated in order to see if there is a performance difference to AT from explicit modeling of the insertion position of tokens. BBT is chosen because it can decode an N length sequence with empirically $\log_2(N)$ iterations. In training with the BBT prior, we increased the number of epochs from 50 to 300 because only a single step of the output token sequence generation is trained in one minibatch while with the L2R prior we can train the whole sequence generation.

Since our insertion-based models do not have a beam search algorithm, we mainly compare the L2R-prior insertion-based models with AT (beam=1) and the BBT-prior insertion-based models with CTC (beam=1).

5.2. Results

First, the results of AT and insertion-based models with the L2R prior are shown in the upper part of Table 1. When compared to AT without beam search, the performance of insertion-based models trained with the CTC objective is mostly better except for the eval2 set of CSJ.

Next, the models with the BBT prior are compared in the lower part of Table 1. These are non-autoregressive models hence performance is first compared to CTC. Unfortunately, Insertion Transformer with the BBT prior cannot compete with

CTC even with hybrid training with CTC. KERMIT with joint training with CTC, which is a new formulation introduced in Section 4, achieved better performance than CTC on CSJ and TEDLIUM2. Notably, it achieved better performance in a non-autoregressive manner than AT without beam search on several test sets highlighted by underlined numbers in Table 1.

5.3. Discussion

One of the remarks we got from the experiments of the L2R prior is that explicit modeling of the position of an output token and hybrid training with the CTC objective worked complementarily and the quality of a hypothesis in decoding was improved.

Another remark is that when the BBT prior is used, the new formulation of joint training with CTC introduced in this paper seems to make use of both benefits of left-to-right and non-left-to-right generation orders at least on CSJ and TEDLIUM2. However, the performance of the BBT prior on AISHELL was not as good as CTC. In this case, the performance of CTC is close to that of AT unlike the other tasks, and the effectiveness of the generation order may depend on the language or task.

6. Conclusions

This paper proposes applying three insertion-based models, originally proposed for NMT, to ASR tasks. In addition, we introduce a new formulation for joint training of the insertion-based model and CTC. Our experiments show that InDIGO and Insertion Transformer trained with the L2R prior achieved comparable or better performance than autoregressive Transformer without beam search. Models trained with the BBT prior and the proposed formulation, which retains non-autoregressive characteristics, achieved better performance than CTC and competitive with autoregressive Transformer without beam search on CSJ and TEDLIUM2 but was not effective on AISHELL. Therefore, we will investigate more solid use of insertion-based models including an extension of decoding algorithm with beam search.

7. References

- [1] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," in *Proc. Advances in Neural Information Processing Systems (NIPS)* 28, 2015, pp. 577–585. [Online]. Available: <http://papers.nips.cc/paper/5847-attention-based-models-for-speech-recognition.pdf>
- [2] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *Proc. 2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2016, pp. 4960–4964.
- [3] D. Amodei *et al.*, "Deep speech 2: End-to-end speech recognition in english and mandarin," in *Proc. of the 33rd International Conference on International Conference on Machine Learning (ICML)*, 2016, pp. 173–182.
- [4] 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*, 2017, pp. 939–943. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2017-233>
- [5] 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, Dec 2017.
- [6] A. Zeyer, K. Irie, R. Schlüter, and H. Ney, "Improved training of end-to-end attention models for speech recognition," in *Proc. Interspeech 2018*, 2018, pp. 7–11. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1616>
- [7] C. Chiu *et al.*, "State-of-the-art speech recognition with sequence-to-sequence models," in *Proc. 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, April 2018, pp. 4774–4778.
- [8] T. N. Sainath *et al.*, "A streaming on-device end-to-end model surpassing server-side conventional model quality and latency," in *Proc. 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6059–6063.
- [9] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition," in *Proc. Interspeech 2019*, 2019, pp. 2613–2617. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2019-2680>
- [10] C. Lüscher, E. Beck, K. Irie, M. Kitza, W. Michel, A. Zeyer, R. Schlüter, and H. Ney, "RWTH ASR Systems for LibriSpeech: Hybrid vs Attention," in *Proc. Interspeech 2019*, 2019, pp. 231–235. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2019-1780>
- [11] S. Karita *et al.*, "A comparative study on transformer vs RNN in speech applications," *arXiv preprint arXiv:1909.06317*, 2019.
- [12] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *Proc. 2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, April 2015, pp. 5206–5210.
- [13] J. J. Godfrey, E. C. Holliman, and J. McDaniel, "Switchboard: telephone speech corpus for research and development," in *Proc. 1992 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, vol. 1, March 1992, pp. 517–520 vol.1.
- [14] K. Maekawa, H. Koiso, S. Furui, and H. Isahara, "Spontaneous speech corpus of Japanese," in *Proc. the Second International Conference on Language Resources and Evaluation (LREC'00)*, 2000.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. Advances in Neural Information Processing Systems (NIPS)* 30, 2017, pp. 5998–6008. [Online]. Available: <http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>
- [16] S. Karita, N. E. Y. Soplin, S. Watanabe, M. Delcroix, A. Ogawa, and T. Nakatani, "Improving Transformer-Based End-to-End Speech Recognition with Connectionist Temporal Classification and Language Model Integration," in *Proc. Interspeech 2019*, 2019, pp. 1408–1412. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2019-1938>
- [17] J. Gu, J. Bradbury, C. Xiong, V. O. Li, and R. Socher, "Non-autoregressive neural machine translation," *arXiv preprint arXiv:1711.02281*, 2017.
- [18] M. Stern, W. Chan, J. Kiros, and J. Uszkoreit, "Insertion transformer: Flexible sequence generation via insertion operations," in *Proc. of International Conference on Machine Learning (ICML)*, 2019, pp. 5976–5985.
- [19] J. Gu, C. Wang, and J. Zhao, "Levenshtein transformer," in *Proc. Advances in Neural Information Processing Systems (NIPS)* 32, 2019, pp. 11 181–11 191. [Online]. Available: <http://papers.nips.cc/paper/9297-levenshtein-transformer.pdf>
- [20] W. Chan, N. Kitaev, K. Guu, M. Stern, and J. Uszkoreit, "Kermit: Generative insertion-based modeling for sequences," *arXiv preprint arXiv:1906.01604*, 2019.
- [21] N. Chen, S. Watanabe, J. Villalba, and N. Dehak, "Listen and fill in the missing letters: Non-autoregressive transformer for speech recognition," *arXiv preprint arXiv:1911.04908*, 2020.
- [22] X. N. B. W. H. Z. Hui Bu, Jiayu Du, "Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline," in *Proc. Oriental COCODSA 2017*, 2017.
- [23] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *Proc. of the 23rd International Conference on Machine Learning (ICML)*, 2006, pp. 369–376.
- [24] S. Watanabe *et al.*, "Espnet: End-to-end speech processing toolkit," in *Proc. Interspeech 2018*, 2018, pp. 2207–2211. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1456>
- [25] W. Chan, C. Saharia, G. Hinton, M. Norouzi, and N. Jaitly, "Imputer: Sequence modelling via imputation and dynamic programming," *arXiv preprint arXiv:2002.08926*, 2020.
- [26] J. Gu, Q. Liu, and K. Cho, "Insertion-based decoding with automatically inferred generation order," *Transactions of the Association for Computational Linguistics*, vol. 7, pp. 661–676, 2019.
- [27] O. Vinyals, M. Fortunato, and N. Jaitly, "Pointer networks," in *Proc. Advances in Neural Information Processing Systems (NIPS)* 28, 2015, pp. 2692–2700. [Online]. Available: <http://papers.nips.cc/paper/5866-pointer-networks.pdf>
- [28] A. Rousseau, P. Deléglise, and Y. Estève, "Enhancing the TED-LIUM corpus with selected data for language modeling and more TED talks," in *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, May 2014, pp. 3935–3939.
- [29] J. Salazar, K. Kirchhoff, and Z. Huang, "Self-attention networks for connectionist temporal classification in speech recognition," in *Proc. 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 7115–7119.