arXiv:2211.00490v1 [eess.AS] 31 Oct 2022

DELAY-PENALIZED TRANSDUCER FOR LOW-LATENCY STREAMING ASR

Wei Kang[*,1], *Zengwei Yao*[*,1], *Fangjun Kuang*[1], *Liyong Guo*[1], *Xiaoyu Yang*[1], *Long lin*[1], *Piotr Żelasko*[2], *Daniel Povey*[1]

[1] Xiaomi Corp., Beijing, China [2] Meaning.Team Inc, USA

{kangwei1, yaozengwei, dpovey}@xiaomi.com, pzelasko@meaning.team

ABSTRACT

In streaming automatic speech recognition (ASR), it is desirable to reduce latency as much as possible while having minimum impact on recognition accuracy. Although a few existing methods are able to achieve this goal, they are difficult to implement due to their dependency on external alignments. In this paper, we propose a simple way to penalize symbol delay in transducer model, so that we can balance the trade-off between symbol delay and accuracy for streaming models without external alignments. Specifically, our method adds a small constant times (T/2 - t), where T is the number of frames and t is the current frame, to all the non-blank log-probabilities (after normalization) that are fed into the two dimensional transducer recursion. For both streaming Conformer models and unidirectional long short-term memory (LSTM) models, experimental results show that it can significantly reduce the symbol delay with an acceptable performance degradation. Our method achieves similar delay-accuracy trade-off to the previously published FastEmit, but we believe our method is preferable because it has a better justification: it is equivalent to penalizing the average symbol delay. Our work is open-sourced and publicly available[1].

Index Terms— speech recognition, delay-penalized, transducer, streaming, low latency

1. INTRODUCTION

End-to-end models have achieved remarkable success in Automatic Speech Recognition (ASR). As a prominent example, transducer [1–3] has gained more and more popularity for real-time ASR system development, because it is naturally streaming and demonstrates superior performance. However, one limitation of transducer is that it focuses on maximizing the total log-probability over all alignments but ignores their specific symbol delays. We hypothesize that the streaming model would augment those alignments emitting symbols later to access more contexts for better performance, leading to higher emission latency in practical ASR application.

There are several classical methods [4–8] to reduce the model latency by constraining the alignments between the frames and transcriptions based on the alignment references generated from external models. Whilst this type of methods achieve good trade-offs between accuracy and latency, it suffers from two limitations: 1) the model performance heavily depends on the precision of the reference alignments; 2) it defeats the advantage of end-to-end model training since it requires an extra frame-level token-time alignments.

To address these limitations, another line of research [9–11] tends to regularize the objective function in a sequence-level manner. A prominent example is FastEmit [9], which encourages the model to emit symbols earlier by scaling up the derivatives of emitting non-blank tokens in backpropagation. Another work named Self alignment [11] proposes to boost the log-probability of the alignment that is one frame to the left of the Viterbi forced-alignment, which requires an extra recursion with a time complexity of $\mathcal{O}(T \times U)$ to obtain the Viterbi forced-alignment, where T and U are the lengths of frame sequence and token sequence respectively. Our method, like FastEmit [9], is simple to implement, but we are able to provide a more detailed demonstration explaining why our method would cause alignment times to change.

In this paper, we propose a novel method of delay penalization for transducer which is able to balance the trade-off between symbol delay and accuracy for streaming models in a simple and efficient way. Different from FastEmit [9] that directly changes the derivatives, we modify the log-probabilities of emitting symbols by adding a small constant λ times the frame offsets relative to middle frame. We mathematically prove that it is approximately equivalent to adding a regularization term that aims to decrease the averaged symbol delay on the regular transducer objective function.

The main contributions of this paper are:

- We propose the delay-penalized transducer, which penalizes the symbol delay without extra token-time alignment.
- We provide a detailed proof why it can encourage the low-delay alignments and penalize the high-delay alignments.
- We show that a tunable trade-off between latency and accuracy can be achieved by adjusting the hyperparameter λ.

2. TRANSDUCER

Let $\mathbf{x} = \{x_t\}_0^{T-1}$ be a sequence of T parameterized input feature frames. Let $\mathbf{y} = \{0 \leq y_u < V\}_0^{U-1}$ be a sequence of

* stands for equal contribution

[1] https://github.com/k2-fsa/k2

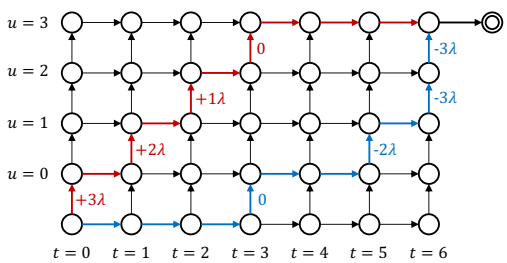

Fig. 1. Delay penalized transducer lattice.

U transcript tokens, where V is the vocabulary size containing the blank token $\varnothing$. As shown in Figure 1, transducer [1] learns alignments between these two sequences $\mathbf{x}$ and $\mathbf{y}$ with different lengths. The vertical transition leaving node (t, u) represents emitting non-blank token y_{u+1} with the log-probility $y(t, u)$, while the horizontal transition represents emitting blank token $\varnothing$ with the log-probility $\varnothing(t, u)$.

The objective function of transducer is to maximize the total log-probability $\mathcal{L}$ over all alignment paths:

$$\mathcal{L} = \log \sum_i \exp(s_i), \tag{1}$$

where s_i is the log-probability of path i summing over all contained transitions. The forward-backward algorithm [1] is usually employed to calculate $\mathcal{L}$ in an efficient manner. Let $\alpha(t, u)$ be the log-probability at node (t, u), which represents emitting tokens $y_{0\ldots u}$ after seeing features $x_{0\ldots t}$. Then $\alpha(t, u)$ could be calculated recursively as:

$$\begin{aligned}\alpha(t, u) = \text{LogAdd}(&\alpha(t, u-1) + y(t, u-1), \\ &\alpha(t-1, u) + \varnothing(t-1, u)),\end{aligned} \tag{2}$$

where LogAdd is defined as:

$$\text{LogAdd}(a, b) = \log\left(e^a + e^b\right). \tag{3}$$

Herein, $\alpha(0, 0)$ is initialized as 0. The total log-probability of over all alignments path $\mathcal{L}$ is:

$$\mathcal{L} = \alpha(T-1, U) + \varnothing(T-1, U). \tag{4}$$

One limitation of transducer is that it is optimized to maximize the total log-probability $\mathcal{L}$ over all alignments, regardless of their respective symbol delays. As shown in Figure 1, the blue alignment emitting symbols later has a higher delay compared with the red alignment. Unlike non-streaming model that could access full contexts in an utterance, the streaming model tends to concentrate on those alignments emitting symbols later, such as the blue alignment in Figure 1, thus to access more future contexts for a better recognition performance. The blue line in Figure 2 presents the mean alignment delay of the streaming Conformer, which constantly increase as the training goes on.

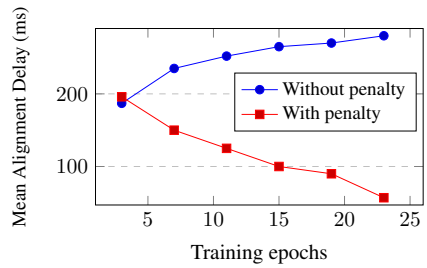

Fig. 2. Mean alignment delay of streaming model during training.

3. DELAY-PENALIZED TRANSDUCER

To penalize symbol delay, we (conceptually) add an extra term in the loss function:

$$\mathcal{L}_{\text{aug}} = \mathcal{L} + \mathcal{L}_{\text{delay}}. \tag{5}$$

Herein, $\mathcal{L}_{\text{delay}}$ represents the scaled weighted average delay score over all alignments, which is formulated as:

$$\mathcal{L}_{\text{delay}} = \lambda \sum_i d_i w_i, \tag{6}$$

where d_i is the delay score of alignment i, λ is a scaling hyper-parameter, and w_i is the path weight:

$$w_i = \frac{\partial \mathcal{L}}{\partial s_i} = \frac{\exp(s_i)}{\sum_i \exp(s_i)}. \tag{7}$$

Herein, the sum of weight w_i over all alignments is 1. We can get the derivatives of $\mathcal{L}_{\text{aug}}$ with respect to s_i as:

$$\frac{\partial \mathcal{L}_{\text{aug}}}{\partial s_i} = \frac{\partial \mathcal{L}}{\partial s_i} + \frac{\partial \mathcal{L}_{\text{delay}}}{\partial s_i} \tag{8}$$

From (6) and (7), we can get:

$$\frac{\partial \mathcal{L}_{\text{delay}}}{\partial s_i} = \lambda \left(\frac{d_i \exp(s_i)}{\sum_i \exp(s_i)} - \frac{d_i (\exp(s_i))^2}{(\sum_i \exp(s_i))^2} \right), \tag{9}$$

which can be rearranged as:

$$\frac{\partial \mathcal{L}_{\text{delay}}}{\partial s_i} = \lambda \frac{(d_i - d_{\text{avg}}) \exp(s_i)}{\sum_i \exp(s_i)}, \tag{10}$$

where d_{avg} is:

$$d_{avg} = \sum_i d_i w_i. \tag{11}$$

From (7), (8) and (10) we can get:

$$\frac{\partial \mathcal{L}_{\text{aug}}}{\partial s_i} = \frac{(1 + \lambda(d_i - d_{\text{avg}})) \exp(s_i)}{\sum_i \exp(s_i)}. \tag{12}$$

For a small λ, $1 + \lambda(d_i - d_{\text{avg}})$ is close to $\exp(\lambda(d_i - d_{\text{avg}}))$, we can approximate (12) as:

$$\frac{\partial L_{\text{aug}}}{\partial s_i} \approx \frac{\exp(\lambda(d_i - d_{\text{avg}}) + s_i)}{\sum_i \exp(s_i)}. \tag{13}$$

According to (10) and (11), the sum of derivative $\frac{\partial \mathcal{L}_{\text{delay}}}{\partial s_i}$ over all alignments is 0. By plugging in $\frac{\partial \mathcal{L}}{\partial s_i}$ from (7), we can get:

$$\sum_i \frac{\partial \mathcal{L}_{\text{aug}}}{\partial s_i} = \sum_i \frac{\partial \mathcal{L}}{\partial s_i} + \sum_i \frac{\partial \mathcal{L}_{\text{delay}}}{\partial s_i} = 1. \tag{14}$$

Then we can equivalently normalize (13) as:

$$\frac{\partial \mathcal{L}_{\text{aug}}}{\partial s_i} \approx \frac{\exp\left(\lambda(d_i - d_{\text{avg}}) + s_i\right)}{\sum_i \exp\left(\lambda(d_i - d_{\text{avg}}) + s_i\right)}. \tag{15}$$

without changing its numerical value (for small λ). There is no difference between (13) and (15) for a small λ; in any case the change is equivalent to multiplying the loss function by a constant that is very close to 1. As softmax is invariant under translation, $\mathcal{L}_{\text{delay}}$ actually makes no difference to the expression as it cancels, so (15) can be written as:

$$\frac{\partial \mathcal{L}_{\text{aug}}}{\partial s_i} \approx \frac{\exp\left(\lambda d_i + s_i\right)}{\sum_i \exp\left(\lambda d_i + s_i\right)}. \tag{16}$$

Therefore, we can get these path derivatives of the augmented objective function $\mathcal{L}_{\text{aug}}$, by simply computing the regular transducer loss (1) with the modified inputs:

$$s'_i = \lambda d_i + s_i. \tag{17}$$

Let $\pi = \{\pi_u\}_0^{U-1}$ be the frame indexes that emit tokens $\mathrm{y}_{0\ldots U-1}$. As we want the alignments with a lower delay to have a larger delay score, we define d_i as the sum of offsets relative to the middle frame in each utterance:

$$d_i = \sum_u \left(\frac{T-1}{2} - \pi_u\right). \tag{18}$$

Adding the middle-frame offset will make no difference to the derivatives; it is done to prevent the delay-penalty from changing the numerical value of the loss function too much, which would make diagnostics harder to interpret. As shown in Figure 1, we can equivalently implement (17) by adding the offsets on the log-probabilities of emitting non-blank tokens $y(t, u-1)$ according to the specific frame indexes $0 \leq t < T$:

$$y'(t,u) = y(t,u) + \lambda \times \left(\frac{T-1}{2} - t\right). \tag{19}$$

Therefore, by replacing $y(t,u)$ with $y'(t,u)$ in (2), it would encourage low-delay alignments while maximizing the total log-probability $\mathcal{L}$, to prevent the transducer from avidly enhancing the high-delay alignments to access more future contexts [2]. As shown in the red line in Figure 2, by applying the delay penalty on transducer, we can gradually achieve a lower symbol delay for the streaming Conformer.

4. EXPERIMENTS

4.1. Latency metrics

We measure the latency of streaming models with two types of delay metrics described below: (1) Mean Alignment Delay (MAD) and (2) Mean End Delay (MED). The ground-truth word-time alignments are obtained by performing forced alignment with the Montreal Forced Aligner tool [3]. For simplicity, we only consider the correctly recognized words for both metrics. Specifically, MAD is the mean of word time difference between the predicted alignments and ground truth, which is defined as:

$$\text{MAD} := \frac{1}{\sum_{n=0}^{N-1} S_n} \sum_{n=0}^{N-1} \sum_{s=0}^{S_n-1} \left(\hat{t}_s^n - t_s^n\right) \tag{20}$$

Herein, $\hat{t}_s^n$ and t_s^n are the timestamps of the s-th word in prediction and ground truth respectively. N is the number of utterances. S_n is the number of matched words between prediction and reference in the n-th utterance. MED only considers the emitting time of the last word in an utterance, which is calculated as:

$$\text{MED} := \frac{1}{N} \sum_{n=0}^{N-1} \left(\hat{t}_{end}^n - t_{end}^n\right), \tag{21}$$

where $\hat{t}_{end}^n$ and t_{end}^n are the timestamps of the last word in prediction and ground truth respectively.

4.2. Experimental Setup

Our experiments are conducted on the popularly used open-source dataset LibriSpeech [12], containing 1000 hours of English reading speech. We employ Lhotse [13] for data preparation. The acoustic features are 80-dimension Mel filterbank with a frame length of 25 ms and frame shift of 10 ms. SpecAugment [14] and noise augmentation based on MUSAN [15] are applied during training to improve generalization capability. Furthermore, speed perturbation [16] with factors 0.9 and 1.1 are used to triple the training set. The transcripts are tokenized into 500-class word pieces with Byte Pair Encoding (BPE) [17].

To evaluate the effectiveness and robustness of the proposed method, we adopt streaming Conformer [18] and unidirectional LSTM as encoders respectively. Both of the Conformer and LSTM consist of 12 layers, where a convolutional downsampling layer with a factor of 4 is first used to obtain the 512-dimension feature embedding. Similar to [19], we train the Conformer with block-triangular masks to limit the future context within a dynamic chunk size and infer it with a fixed chunk size of 640 ms. For each of Conformer encoder layer, the attention dimension and the feed-forward dimension are 512 and 2048, respectively. The LSTM layers adopt similar residual connection structure as in Conformer [18], each of which is composed of a unidirectional LSTM layer with 1024 hidden units and a feed-forward layer with a hidden dimension of 2048. We use a stateless decoder [20], which consists of an embedding layer followed by a 1-D convolutional layer with a kernel size of 2. Pruned transducer

[2] An alternative way to implement (17) is to apply the delay penalty on log-probability of emitting blank tokens $\varnothing(t,u)$ in opposite direction.

[3] https://github.com/MontrealCorpusTools/Montreal-Forced-Aligner

loss [3] is adopted for low memory usage and efficient computation. The delay penalty is applied on both the simple loss and pruned loss. We conduct experiments with 5 values of λ in (19), including 0.0015, 0.0030, 0.0060, 0.0075, 0.0100.

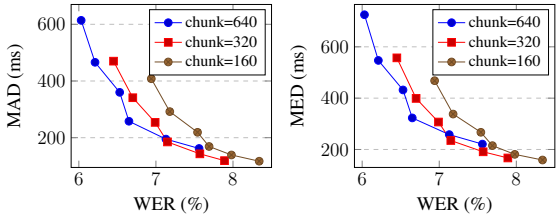

(a) MAD-WER trade-off (b) MED-WER trade-off

Fig. 3. Delay-accuracy trade-off comparison using different decoding chunk sizes (ms) for streaming Conformer.

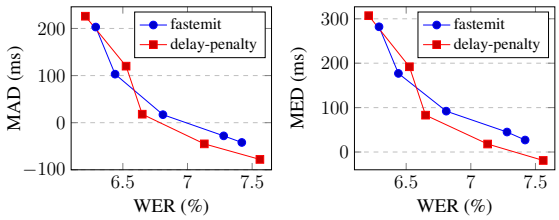

(a) MAD-WER trade-off (b) MED-WER trade-off

Fig. 4. Delay-accuracy trade-off comparison on Conformer.

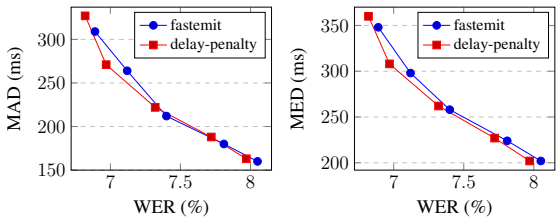

(a) MAD-WER trade-off (b) MED-WER trade-off

Fig. 5. Delay-accuracy trade-off comparison on LSTM.

4.3. Delay and accuracy trade-off

Table 1 presents experimental results using different λ in (19). For both Conformer and LSTM models, a larger λ consistently leads to a lower symbol delay as well as a higher WER. It manifests that we can balance the trade-off between symbol delay and accuracy for both streaming Conformer and LSTM models in a simple and effective way by tuning λ. Note that for the Conformer that can access a chunk of future context, the MAD and MED is further reduced to below zero, which indicates that the model is regularized to emit symbols before they are spoken.

We also investigate the effect of the decoding chunk size of the streaming Conformer model. Figure 3 shows the delay-accuracy trade-offs with decoding chunk size of 640 ms, 320 ms, and 160 ms, respectively, where the presented results for each experiment are averaged over *test-clean* and *test-other*. Note that the MAD and MED here are the total delays including the latency introduced by the chunk-wise decoding, which equals half of the chunk length (i.e., 320 ms, 160 ms, and 80 ms). The results manifest that it is preferable to decode with a larger chunk size while employing the delay penalty, which yields a better trade-off between symbol delay and accuracy.

Table 1. ASR results on LibriSpeech using Conformer and LSTM as streaming encoder respectively.

Method	λ	test-clean			test-other		
		WER (%)	MAD (ms)	MED (ms)	WER (%)	MAD (ms)	MED (ms)
Conformer	0	3.4	373	485	8.66	374	484
	0.0015	3.42	213	295	9.01	238	319
	0.0030	3.70	102	176	9.35	137	208
	0.0060	3.74	-3	62	9.56	39	104
	0.0075	4.13	-62	-1	10.13	-27	37
	0.0100	4.67	-93	-37	10.44	-62	-1
LSTM	0	3.78	418	437	9.55	419	425
	0.0015	3.82	316	353	9.82	337	366
	0.0030	3.86	257	299	10.08	284	317
	0.0060	4.11	206	250	10.53	237	273
	0.0075	4.52	172	214	10.91	203	240
	0.0100	4.53	148	189	11.40	178	214

4.4. Comparison with FastEmit

We also conduct experiments to compare our proposed delay penalization method with FastEmit [9]. For FastEmit, λ is set to 0.0030, 0.0060, 0.0100, 0.0150, and 0.0200, respectively. The FastEmit mechanism is also applied on both of the simple loss and pruned loss in pruned transducer [3]. Figure 4 and Figure 5 present the delay-accuracy trade-offs of applying FastEmit [9] and delay penalty as latency regularization on both Conformer model and LSTM model, respectively. For an overall comparison, the presented results for each experiment are averaged over *test-clean* and *test-other*. It shows that our method achieves similar delay-accuracy trade-offs to FastEmit [9], while our method provides a more detailed demonstration explaining why it is able to cause alignment time to change.

5. CONCLUSION

We propose a method of delay penalty on transducer, which is able to penalize the symbol delay in a simple and efficient way without any extra token-time alignments. We provide a detailed proof explaining why our method is able to cause alignment time to change, so as to reduce the symbol delay. We verify the proposed method on both streaming Conformer and LSTM models. The experimental results show that we can get a promising trade-off between symbol delay and accuracy by tuning the hyper parameter λ.

6. REFERENCES

[1] Alex Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.

[2] Qian Zhang, Han Lu, Hasim Sak, Anshuman Tripathi, Erik McDermott, Stephen Koo, and Shankar Kumar, "Transformer transducer: A streamable speech recognition model with transformer encoders and rnn-t loss," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7829–7833.

[3] Fangjun Kuang, Liyong Guo, Wei Kang, Long Lin, Mingshuang Luo, Zengwei Yao, and Daniel Povey, "Pruned RNN-T for fast, memory-efficient ASR training," in *Proc. Interspeech*, 2022, pp. 2068–2072.

[4] Andrew Senior, Haşim Sak, Félix de Chaumont Quitry, Tara Sainath, and Kanishka Rao, "Acoustic modelling with cd-ctc-smbr lstm rnns," in *IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, 2015, pp. 604–609.

[5] Haşim Sak, Andrew Senior, Kanishka Rao, and Françoise Beaufays, "Fast and accurate recurrent neural network acoustic models for speech recognition," *arXiv preprint arXiv:1507.06947*, 2015.

[6] Tara N Sainath, Ruoming Pang, David Rybach, Basi García, and Trevor Strohman, "Emitting word timings with end-to-end models.," in *Proc. Interspeech*, 2020, pp. 3615–3619.

[7] Hirofumi Inaguma, Yashesh Gaur, Liang Lu, Jinyu Li, and Yifan Gong, "Minimum latency training strategies for streaming sequence-to-sequence asr," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6064–6068.

[8] Jay Mahadeokar, Yuan Shangguan, Duc Le, Gil Keren, Hang Su, Thong Le, Ching-Feng Yeh, Christian Fuegen, and Michael L Seltzer, "Alignment restricted streaming recurrent neural network transducer," in *IEEE Spoken Language Technology Workshop (SLT)*, 2021, pp. 52–59.

[9] Jiahui Yu, Chung-Cheng Chiu, Bo Li, Shuo-yiin Chang, Tara N Sainath, Yanzhang He, Arun Narayanan, Wei Han, Anmol Gulati, Yonghui Wu, et al., "Fastemit: Low-latency streaming asr with sequence-level emission regularization," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6004–6008.

[10] Yuan Shangguan, Rohit Prabhavalkar, Hang Su, Jay Mahadeokar, Yangyang Shi, Jiatong Zhou, Chunyang Wu, Duc Le, Ozlem Kalinli, Christian Fuegen, and Michael L. Seltzer, "Dissecting User-Perceived Latency of On-Device E2E Speech Recognition," in *Proc. Interspeech*, 2021, pp. 4553–4557.

[11] Jaeyoung Kim, Han Lu, Anshuman Tripathi, Qian Zhang, and Hasim Sak, "Reducing streaming asr model delay with self alignment," *arXiv preprint arXiv:2105.05005*, 2021.

[12] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2015, pp. 5206–5210.

[13] *Lhotse: a speech data representation library for the modern deep learning ecosystem*, NeurIPS Data-Centric AI Workshop, 2021.

[14] Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le, "SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition," in *Proc. Interspeech*, 2019, pp. 2613–2617.

[15] David Snyder, Guoguo Chen, and Daniel Povey, "Musan: A music, speech, and noise corpus," *arXiv preprint arXiv:1510.08484*, 2015.

[16] Tom Ko, Vijayaditya Peddinti, Daniel Povey, and Sanjeev Khudanpur, "Audio augmentation for speech recognition," in *Proc. Interspeech*, 2015, pp. 3586–3589.

[17] Rico Sennrich, Barry Haddow, and Alexandra Birch, "Neural machine translation of rare words with subword units," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 2016, pp. 1715–1725.

[18] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang, "Conformer: Convolution-augmented Transformer for Speech Recognition," in *Proc. Interspeech*, 2020, pp. 5036–5040.

[19] Binbin Zhang, Di Wu, Zhuoyuan Yao, Xiong Wang, Fan Yu, Chao Yang, Liyong Guo, Yaguang Hu, Lei Xie, and Xin Lei, "Unified streaming and non-streaming two-pass end-to-end model for speech recognition," *arXiv preprint arXiv:2012.05481*, 2020.

[20] Mohammadreza Ghodsi, Xiaofeng Liu, James Apfel, Rodrigo Cabrera, and Eugene Weinstein, "Rnn-transducer with stateless prediction network," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7049–7053.