



Attention Enhanced Citrinet for Speech Recognition

Xianchao Wu

NVIDIA

xianchaow@nvidia.com

Abstract

Citrinet is an end-to-end convolutional Connectionist Temporal Classification (CTC) based automatic speech recognition (ASR) model. To capture local and global contextual information, 1D time-channel separable convolutions combined with sub-word encoding and squeeze-and-excitation (SE) are used in Citrinet, making the whole architecture to be as deep as including 23 blocks with 235 convolution layers and 46 linear layers. This pure convolutional and deep architecture makes Citrinet relatively slow at convergence. In this paper, we propose to introduce multi-head attentions together with feed-forward networks in the convolution module in Citrinet blocks while keeping the SE module and residual module unchanged. For speeding up, we remove 8 convolution layers in each attention-enhanced Citrinet block and reduce 23 blocks to 13. Experiments on the Japanese CSJ-500h and Magic-1600h dataset show that the attention-enhanced Citrinet with less layers and blocks and converges faster with lower character error rates than (1) Citrinet with 80% training time and (2) Conformer with 40% training time and 29.8% model size.

Index Terms: speech recognition, human-computer interaction, computational paralinguistics

1. Introduction

Convolutions with predefined kernel sizes capturing local contexts have been successfully utilized in end-to-end ASR systems. JasperNet [1], QuartzNet [2], ContextNet [3], and Citrinet [4] follow this direction and use pure convolution networks to capture the interactive context information of input wave spectrogram. Among these models, Citrinet (Figure 1) has achieved top-level accuracy. Citrinet’s encoder combines 1D time-channel separable convolutions following QuartzNet [2] and the SE mechanism [5] which has been successfully used in ContextNet [3] for global context representation.

Alternatively, ASR systems leveraged on Transformers [6] and their variants have achieved impressively low word/character error rates (WERs/CERs) in numerous languages during recent years [7, 8, 9, 10, 11]. The Transformer architecture includes multi-head self-attention (MHSA) and cross-attention layers that can represent long-distance interactions inside and between sound-text sequences.

There are works such as Conformer [8] which introduced convolutions into Transformer architecture and achieved impressive WERs/CERs in numerous languages. In Conformer, two macaron-like [12] feed-forward networks (FFN) with half-step residual connections sandwiches the MHSA and convolution modules followed by a post layer normalization [13].

In this paper, we consider a reverse direction of introducing attention mechanisms into convolution-leading architecture. In particular, we introduce self-attention and cross-attention in a state-of-the-art deep convolutional CTC model, Citrinet, with as minimum modifications as possible. Our work is motivated

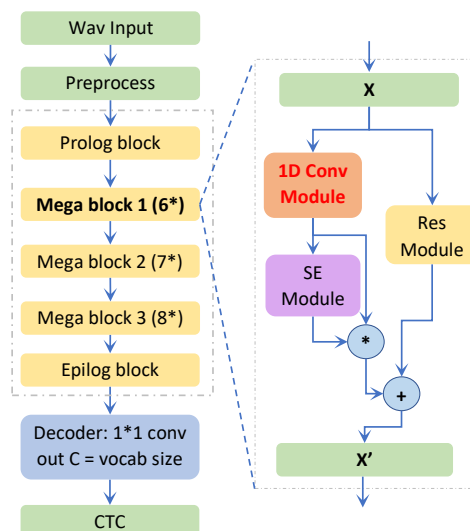


Figure 1: The end-to-end architecture of Citrinet and its major building block - “1D time-channel separable convolution and SE-attached” Jasper Block.

by the following facts and considerations. First, in order to capture long-range contextual representation, there are 235 convolution layers and 46 linear layers in a typical “thin and long” Citrinet model. As will be described in detail in Section 2, there are 10 convolutional functions stacked in one Jasper block used in Citrinet without internal residual connection. These make the training convergence to be relatively slow and the gradient back-propagation is relatively more difficult without shortcut-style residual connections [14]. Second, kernel sizes from 11 to 39 and further 41 are hard-coded in Citrinet to capture as long range information as possible. This hard setting is less friendly to long and short speech waves: a kernel size of 41 is still not enough for a long wave yet it can be a waste for short waves. Third, current Citrinet is a pure CTC model with a non-autoregressive decoder: one convolutional layer projects from the encoder’s output to the vocabulary size. We are interested in boosting Citrinet with cross-attention enhanced decoder architecture.

In our attention-enhanced Citrinet, we insert a FFN module and a MHSA module in the “1D conv module” of one Jasper block and then reduce the repeating time from 5 to 1 to remove 80% of convolution layers and select 13 from 23 blocks to obtain a fat and short pipeline. We also add a bidirectional Transformer [6] decoder to ensure Citrinet’s training in a non-autoregressive way using a linear combination of CTC loss and attention losses, and inferencing in the same non-autoregressive way. We perform our experiments using the 500-hour Japanese CSJ dataset and the 1600-hour Japanese Magic dataset.

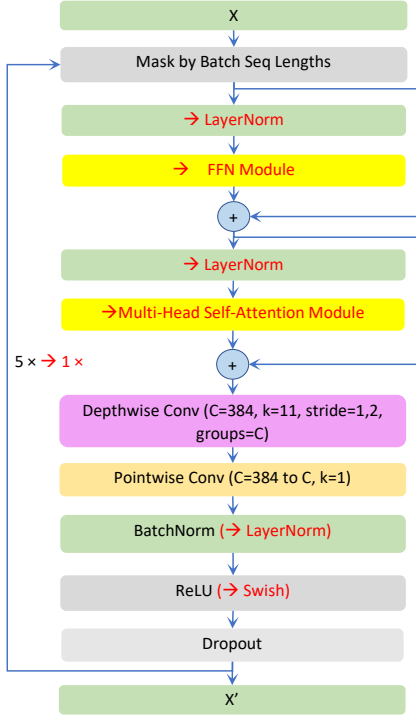


Figure 2: 1D time-channel separable convolution module enhanced by FFN and MHSA modules.

2. Attention Enhanced Citrinet

2.1. Citrinet

Figure 1 depicts the “wave to textual sequence” pipeline of Citrinet and its major building block. The preprocessing is to extract the 80-dimension FBanks and then perform spectrogram augmentation [15] on time/frequency dimensions. One (updated) Jasper block [1, 2], in the right-hand-side of the figure, contains three modules: (1) one 1D time-channel separable convolution module (Figure 2), (2) one SE module [5] and in parallel (3) one residual (Res) connection module (Figure 3). This Jasper block can be expressed by the following equations:

$$\tilde{\mathbf{X}} = \text{1D_Conv}(\mathbf{X}) \quad (1)$$

$$\mathbf{X}' = \text{Res}(\mathbf{X}) + \tilde{\mathbf{X}} \times \text{SE}(\tilde{\mathbf{X}}). \quad (2)$$

The “1D conv module” repeatedly extracts time-channel contextual representations guided by increased kernel sizes. There are two linear layers in the SE module that project from C to $C/8$ and back to C to capture the whole sequence contextual information of each channel. The contextual information is then projected into the $[0, 1]$ range by a sigmoid function to act as *gates* controlling the traffic to the following layers. Thus, the local and global information are extracted by 1D conv and SE modules, respectively. The Res module takes the original tensor \mathbf{X} as input and performs a simple 1×1 convolution followed by a batch normalization. The stride for this 1×1 convolution can take values of 1 or 2 to align with the distilling of the three mega blocks to be described hereafter.

The encoder contains 23 Jasper blocks which are assigned to 5 block sets: one prolog block, three mega blocks and finally one epilog block. Prolog block and epilog block both adapt the Jasper block architecture except that the residual module (“Res

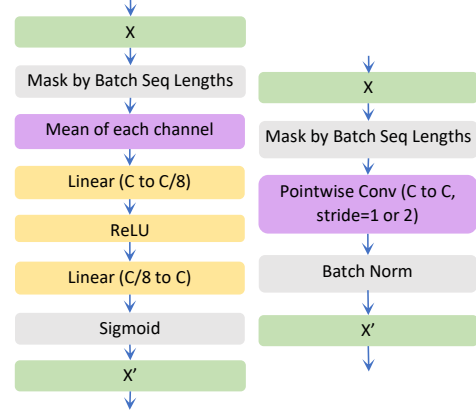


Figure 3: The SE (left) and Res (right) modules used in Citrinet’s building block.

Module” in the Figure) is not included. Using the “1D Conv Module”, prolog block’s major task is to transform the input 80-dimension FBanks into C (e.g., 384, 1024) channels and epilog block’s major task is to capture global information again with kernel size 41 and then transform the input C channels into the dimension (e.g., 640) similar to the decoder’s input.

There are 6, 7, and 8 Jasper blocks in the three mega blocks, respectively. In each first Jasper block, a stride of 2 is used and consequently the “time-dimension” length is halved. Generally, this can be recognized as a *distilling* process in the encoder: output length is $1/2^3$ of input length. In order to capture global context information, the kernel sizes of the mega blocks increase linearly from 11 to 39. In particular, the first mega block’s kernel sizes range over 11, 13, ..., to 21; the second mega block takes values from 13, 15, ..., to 25; and the third mega block starts from kernel size of 25 and ends at 39.

The decoder is a one-layer convolutional network that projects from the encoder’s output to the vocabulary size. Finally, CTC loss [16] is used during training the architecture.

2.2. Enhance 1D Time-channel Convolution by Attention

Figure 2 shows the original 1D time-channel separable convolution module and our modifications in red font after arrows. This depth-wise separable network includes (1) a depth-wise CNN with kernel size from 11 to 39 that tries to capture relatively a long range of context and (2) a pointwise CNN that acts like a linear projection with kernel size of 1. Then, batch normalization, ReLU activation and dropout are performed. In the original Jasper block, these repeat 5 times resulting in 10 convolutional functions, 5 batch normalization and 4 ReLU/dropout functions (the final repeat does not include ReLU/dropout). These 10 convolutional layers are stacked together without any residual connections inside them. As shown in Figure 1, there is only one residual module outside the 1D conv module. These plain nets that simply stack CNN layers (without shortcut-style residual connections in between) exhibit higher training error when the depth increases as illustrated in ResNet [14].

Motivated by this observation and the successful combination of CNNs with attentions in Conformer [8], we modify this 1D conv module by (1) setting repeat time from 5 to 1, (2) appending a FFN module and a MHSA module right before the first convolutional layer, (3) changing the batch nor-

malization into layer normalization [13] which has been successfully used in Transformer [6] and (4) replacing ReLU by the Swish [17] activation function which is a self-gated function $f(x) = x \cdot \text{sigmoid}(x)$ and has been successfully used in Conformer [8].

We reduce 8 convolution functions in one Jasper block. The appending of the FFN and MHSA modules bring 6 linear projection layers, 2 layer normalization and 2 residual connections, making the resulting Citrinet architecture to be fatter and shorter. We count on these FFN+MHSA modules for capturing global interactive information among elements of the input sequence. This works together with the SE module to perform *self-attending* and *self-gating* of global information.

2.3. Enhance Decoder by Attention

We append a 3-block bidirectional Transformer decoder which encodes the textual sequences in left-to-right (*l2r*) and right-to-left (*r2l*) directions and predicts the whole target sequence by one forward step. The input and reference output to the *l2r* decoder are “ s, y_1, \dots, y_N ” and “ $y_1, \dots, y_N, /s$ ” which are taken from a reference label sequence y_1, \dots, y_N , respectively. Here, s and $/s$ stand for start/end of sequence symbols, respectively. The *r2l* decoder takes “ s, y_N, \dots, y_1 ” as input and predicts “ $y_1, \dots, y_N, /s$ ”.

The objective is to minimize \mathcal{L} , a linear combination of the CTC loss [16] and attention losses (ATT) computed by point-wise KL-divergence [18]:

$$\mathcal{L} = \lambda_1 * \text{CTC}_{loss} + (1 - \lambda_1) * \{ \lambda_2 * \text{ATT}_{l2r} + (1 - \lambda_2) * \text{ATT}_{r2l} \} \quad (3)$$

When we set $\lambda_1 = 1$, it degenerates to the original Citrinet’s loss function. In our attention-enhanced Citrinet variants, we set $\lambda_1 = 0.3$ and $\lambda_2 = 0.7$. Label smoothing with $\delta = 0.1$ is applied to the attention objective so that the references are discounted by $(1 - \delta)$.

3. Experiments

3.1. Setup

Our experiments are performed under an open-source platform, NVIDIA NeMo¹ which is a conversational AI toolkit built for researchers working on ASR, natural language processing, and text-to-speech synthesis with rich pretrained models and complete implementation of Citrinet. We investigate Citrinet and attention-enhanced Citrinet’s performances under channel sizes from 256 to 1,048. We release the code and pretrained models through NeMo. During training, we select the Novograd optimizer [19] ($\beta = [0.8, 0.25]$) with a start learning rate of 0.05. The cosine annealing [20] learning rate scheduler with 10K warm-up steps is used.

We select another open-source ASR platform WeNet² [21] as baseline. WeNet’s implementation of the Conformer baseline model consists of a 12-block Conformer encoder ($d_{\text{FFN}}=2048$, $h=8$, $d=512$, $\text{CNN}_{\text{kernel}}=31$) and a 3-block bidirectional Transformer decoder ($d_{\text{FFN}}=2048$, $h=8$, $d=512$) which also encodes the textual sequences in *l2r* and *r2l* directions. The models are trained with static batching skills. For decoding, we follow the CTC beam search + Attention-rescoring strategy. All our experiments were performed under NVIDIA DGX-A100 with 8*A100-80GB GPUs.

¹<https://github.com/NVIDIA/NeMo>

²<https://github.com/wenet-e2e/wenet>

Table 1: CERs (%) of 3 baselines and 5 Citrinet (C) and attention-enhanced Citrinet (Att-C) variants with different channel sizes. All models are trained under CSJ-500h dataset.

Model	1.73h test1	1.82h test2	1.23h test3	Model size (M)
Transformer [24]	7.6	6.1	6.3	36
Espnet+LM	6.5	4.6	5.1	NA
Conformer [8]	6.97	4.65	5.29	135.1
C-256	8.14	5.79	6.54	12.2
C-384	7.28	4.81	5.44	22.7
C-512	7.54	5.45	6.01	38.0
C-768	7.37	5.33	5.71	85.0
C-1024	7.58	5.76	5.87	141.6
Att-C-256	7.10	4.75	5.79	19.9
Att-C-384	6.27	4.16	5.01	40.3
Att-C-512	6.99	4.94	5.46	68.6
Att-C-768	7.23	4.87	5.18	123.5
Att-C-1024	7.23	5.49	5.55	201.4

3.2. Data

We evaluate attention-enhanced Citrinet variants on the Japanese “Corpus of Spontaneous Japanese” (CSJ) dataset³ and the Japanese Magic dataset⁴, which respectively consist of 500 and 1,600 hours of labeled waves. The CSJ dataset contains research presentations that are recorded from conference rooms. The Magic dataset is relatively noisier of spoken Japanese with informal Japanese grammar. Both datasets are Tokyo-region pronunciations. We use sentencepiece-bpe [22] to segment the text sequences and select a vocabulary size of 4,096. Both datasets have 3,000 around single characters and 92% characters are overlapping. During data preparation, we generate 80-dimension FBank feature vectors with a 25ms window, a 10ms frame stride and dither=1.0. SpecAugment [15] is adapted with 2 frequency masks ($F=10$) and 2 time masks ($T=50$). Global Cepstral Mean and Variance Normalization [23] technique is employed to normalize the 80-dimension FBank vectors.

3.3. Major Results on the CSJ Test Sets

Table 1 lists the CERs (%) of 5 Citrinet and attention-enhanced Citrinet variants and 3 baselines. The first baseline is a Transformer with enhanced information from speakers and speech [24]. Parameters were shared among layers of encoder and decoder blocks in Transformer. Speaker information (gender, age, education-level, speaker ID) and speech utterance properties (duration of short and long, topic of the lecture) were used to augment the training data. The second baseline is Espnet’s implementation of deep VGG-BLSTM [25] with CTC joint decoding and LM rescoring⁵. Note that only this baseline uses an external language model among all the 13 models. The third baseline is Conformer-Large [8] with 135.1M parameters.

Since convolutions play a leading position in Citrinet, we take channel sizes of 256, 384, 512, 768, and 1024 to investigate if enlarging the model size yields lower CERs. From Table

³<https://ccd.ninjal.ac.jp/csj/en/>

⁴<https://www.magicdatatech.com/datasets?language=Japanese>

⁵<https://github.com/espnet/espnet/blob/master/egs/csj/asr1/RESULTS.md>

Table 2: CERs (%) of 3 baselines and 3 Citrinet (C) and attention-enhanced Citrinet (Att-C) variants with different channel sizes. All models are trained under the combined CSJ-500h + Magic-1600h dataset.

Model	1.73h test1	1.82h test2	1.23h test3	2.01h Magic-test
Conformer [8]	11.22	6.35	8.78	11.73
C-256	13.68	7.78	7.29	12.65
C-384	8.51	5.80	8.36	10.56
C-512	12.39	7.25	7.72	13.98
Att-C-256	12.02	7.22	7.54	10.80
Att-C-384	7.94	5.62	6.92	9.13
Att-C-512	11.53	7.21	7.60	11.98

1, we observe that C-384 performs the best and larger channel sizes hardly bring significant reduction of CERs. Similar tendency happens when we attach an attention mechanism to Citrinet. This does not strictly align with Citrinet’s performance on the famous LibriSpeech dataset. We make a detailed analysis of the length of the waves and find that 76.6% of LibriSpeech’s waves are longer than 10 seconds, yet 97.8% of CSJ’s waves are less than 5 seconds. This reflects that large channel sizes in CNN layers are less helpful to ASR corpus with short waves.

In addition, we have the following observations. First, Conformer performs well among the baselines, even with a relatively large model size. Conformer’s performance is comparable with Espnet where an external LM is employed for rescoring.

Second, under all the 5 channel sizes, attention-enhanced Citrinet (Att-C) performs better than their Citrinet counterparts, with averagely 0.94%, 0.70%, 0.54%, 0.38%, and 0.31% CER reducing, respectively. These reflect that attentions do help Citrinet capture global contextual representations. The model sizes are nearly doubled by adding FFN+MHSA modules and $l2r, r2l$ decoders. We set $h=8, d=256$ for Att-C-256, $d=384$ for Att-C-384, and $d=512$ for the other three variants, and $d_{FFN}=4d$. With these modifications, Att-C converges faster than Citrinet and only requires 80% of the total training time.

Third, the best Att-C-384 also outperformed Conformer by reducing 0.49% (relatively 8.7%) CER scores on average with 29.8% model size (40.3M vs. 135.1M). Conformer and Att-C-384 are set the same number of training epochs. Training one epoch for Conformer costs 10min and Att-C-384’s is only 4min. We thus only require 40% of Conformer’s training time in total.

3.4. Major Results on the CSJ+Magic Datasets

Considering the fact that the Japanese CSJ dataset only contains 500 hours of training data and most of its waves are less than 5 seconds, we append a 1600-hour Japanese Magic dataset of spoken Japanese contents which has 67.64% wave files that exceed 10 seconds. The CSJ dataset contains presentations of research conferences. This brings a challenge of domain adaptation. We are interested in disclosing if the Conformer baseline and Citrinet variants can keep their performance. From the 1600-hour Magic dataset, we randomly pick out 20 speaker ids and take half as the development set and the remaining half as the test set. Both sets contain 2 hours and there is no overlapping of the speaker ids in the training data anymore. We combine the two training datasets and retrain the sentencepiece-bpe tokenizer with a same vocabulary size of 4,096.

Table 3: CERs (%) of variants when deleting individual modules on the CSJ test sets.

Model	test1	test2	test3
C-384	7.28	4.81	5.44
Att-C-384	6.27	4.16	5.01
without bi-decoder	7.08	4.51	5.24
without FFN	6.95	4.65	5.25
layer-norm \rightarrow batch-norm	6.58	4.46	5.21
Swish \rightarrow ReLU	6.35	4.29	5.23

Table 2 lists the CERs of three Citrinet and attention-enhanced Citrinet variants and the Conformer baseline with the same configuration as Table 1. We choose channel sizes 256, 384, and 512. From Table 2, we observe that C-384 again performs the best and larger channel sizes do not bring reduction of CERs: C-256 and C-512’s performances are comparable. Similar tendency happens in attention-enhanced Citrinet.

We have the following observations. First, CSJ’s three test sets’ CERs become worse when different domain dataset is merged into the relatively clean CSJ dataset. Second, the original Citrinet performs better than the Conformer, with 1.21% (relatively 12.7%) CER decreasing on average of the four test sets. This reflects that a deeper convolutional model trained using large-scale datasets is still robust. Finally, attention-enhanced Citrinet outperforms Citrinet of all the four test sets with a 0.91% (relatively 10.9%) CER reduction on average.

3.5. Impact of Individual Modules

We perform an ablation study of the individual ideas in the attention-enhanced Citrinet. Table 3 shows the CER results when independently removing bi-decoder, dropping FFN modules, replacing layer normalization with batch normalization and changing Swish activation function back to ReLU. Bi-decoder and FFN are relatively important among the four ideas.

3.6. Ablation Study of Jasper Block Numbers

We finally investigate the number of total Jasper blocks with a value set of {8, 9, 10, 11, 12, 13, 15, 17, 20}. We averagely prune ending Jasper blocks in the 3 mega block set (Figure 2). The variants achieve stable comparable CERs with a standard deviation of 0.3% except 9 and 12 with +2% CERs. We thus select the first stable block number of 13.

4. Conclusions

In this paper, we introduce attention mechanisms into a convolution-leading architecture, Citrinet, for end-to-end non-autoregressive Japanese ASR. We keep Citrinet’s architecture of using 1D time-channel separable convolutions, SE and residual modules and append FFN and MHSA before each 1D time-channel convolution. We reduce 8 convolution layers in each attention-enhanced Citrinet block and only utilize 13 such blocks. Experiments on the Japanese CSJ-500h and Magic-1600h datasets show that the attention-Citrinet requires less layers of blocks and converges faster with lower character error rates than Citrinet with 80% training time and Conformer [8] with 40% training time and 29.8% model size. Future work includes using linear attentions [26, 27] and adaptive span attentions [28] for building more efficient Citrinet.

5. References

- [1] J. Li, V. Lavrukhin, B. Ginsburg, R. Leary, O. Kuchaiev, J. M. Cohen, H. Nguyen, and R. T. Gadde, "Jasper: An end-to-end convolutional neural acoustic model," *ArXiv*, vol. abs/1904.03288, 2019.
- [2] S. Kriman, S. Beliaev, B. Ginsburg, J. Huang, O. Kuchaiev, V. Lavrukhin, R. Leary, J. Li, and Y. Zhang, "Quartznet: Deep automatic speech recognition with 1d time-channel separable convolutions," *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6124–6128, 2020.
- [3] W. Han, Z. Zhang, Y. Zhang, J. Yu, C.-C. Chiu, J. Qin, A. Gulati, R. Pang, and Y. Wu, "Contextnet: Improving convolutional neural networks for automatic speech recognition with global context," *ArXiv*, vol. abs/2005.03191, 2020.
- [4] S. Majumdar, J. Balam, O. Hrinchuk, V. Lavrukhin, V. Noroozi, and B. Ginsburg, "Citrinet: Closing the gap between non-autoregressive and autoregressive end-to-end models for automatic speech recognition," 2021.
- [5] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," *CoRR*, vol. abs/1709.01507, 2017. [Online]. Available: <http://arxiv.org/abs/1709.01507>
- [6] 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 *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017.
- [7] H. Miao, G. Cheng, C. Gao, P. Zhang, and Y. Yan, "Transformer-based online ctc/attention end-to-end speech recognition architecture," *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6084–6088, 2020.
- [8] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented Transformer for Speech Recognition," in *Proc. Interspeech 2020*, 2020, pp. 5036–5040.
- [9] X. Chen, Y. Wu, Z. Wang, S. Liu, and J. Li, "Developing real-time streaming transformer transducer for speech recognition on large-scale dataset," *CoRR*, vol. abs/2010.11395, 2020. [Online]. Available: <https://arxiv.org/abs/2010.11395>
- [10] L. Dong, S. Xu, and B. Xu, "Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition," *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5884–5888, 2018.
- [11] N.-Q. Pham, T. S. Nguyen, J. Niehues, M. Müller, and A. H. Waibel, "Very deep self-attention networks for end-to-end speech recognition," in *INTERSPEECH*, 2019.
- [12] Y. Lu, Z. Li, D. He, Z. Sun, B. Dong, T. Qin, L. Wang, and T.-Y. Liu, "Understanding and improving transformer from a multi-particle dynamic system point of view," *arXiv preprint arXiv:1906.02762*, 2019.
- [13] J. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *ArXiv*, vol. abs/1607.06450, 2016.
- [14] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [15] 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," *Interspeech 2019*, Sep 2019. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2019-2680>
- [16] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd International Conference on Machine Learning*, ser. ICML '06. New York, NY, USA: Association for Computing Machinery, 2006, p. 369–376. [Online]. Available: <https://doi.org/10.1145/1143844.1143891>
- [17] P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for activation functions," *CoRR*, vol. abs/1710.05941, 2017. [Online]. Available: <http://arxiv.org/abs/1710.05941>
- [18] S. Kullback and R. A. Leibler, "On Information and Sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, 1951. [Online]. Available: <https://doi.org/10.1214/aoms/1177729694>
- [19] B. Ginsburg, P. Castonguay, O. Hrinchuk, O. Kuchaiev, V. Lavrukhin, R. Leary, J. Li, H. Nguyen, and J. M. Cohen, "Stochastic gradient methods with layer-wise adaptive moments for training of deep networks," *CoRR*, vol. abs/1905.11286, 2019. [Online]. Available: <http://arxiv.org/abs/1905.11286>
- [20] I. Loshchilov and F. Hutter, "SGDR: stochastic gradient descent with restarts," *CoRR*, vol. abs/1608.03983, 2016. [Online]. Available: <http://arxiv.org/abs/1608.03983>
- [21] B. Zhang, D. Wu, C. Yang, X. Chen, Z. Peng, X. Wang, Z. Yao, X. Wang, F. Yu, L. Xie, and X. Lei, "Wenet: Production first and production ready end-to-end speech recognition toolkit," *CoRR*, vol. abs/2102.01547, 2021. [Online]. Available: <https://arxiv.org/abs/2102.01547>
- [22] T. Kudo and J. 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*. Brussels, Belgium: Association for Computational Linguistics, Nov. 2018, pp. 66–71. [Online]. Available: <https://aclanthology.org/D18-2012>
- [23] O. Viikki and K. Laurila, "Cepstral domain segmental feature vector normalization for noise robust speech recognition," *Speech Commun.*, vol. 25, no. 1–3, p. 133–147, aug 1998. [Online]. Available: [https://doi.org/10.1016/S0167-6393\(98\)00033-8](https://doi.org/10.1016/S0167-6393(98)00033-8)
- [24] S. Li, D. Raj, X. Lu, P. Shen, T. Kawahara, and H. Kawai, "Improving transformer-based speech recognition systems with compressed structure and speech attributes augmentation," in *INTER-SPEECH*, 2019.
- [25] X. Ouyang, S. Kawaai, E. Goh, S. Shen, W. Ding, H. Ming, and D.-Y. Huang, "Audio-visual emotion recognition using deep transfer learning and multiple temporal models," 11 2017, pp. 577–582.
- [26] S. Zhuoran, Z. Mingyuan, Z. Haiyu, Y. Shuai, and L. Hongsheng, "Efficient attention: Attention with linear complexities," in *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2021, pp. 3530–3538.
- [27] S. Li, M. Xu, and X.-L. Zhang, "Efficient conformer-based speech recognition with linear attention," *2021 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pp. 448–453, 2021.
- [28] X. Chang, A. S. Subramanian, P. Guo, S. Watanabe, Y. Fujita, and M. Omachi, "End-to-end asr with adaptive span self-attention," in *INTERSPEECH*, 2020.