



# Parameter-Efficient Conformers via Sharing Sparsely-Gated Experts for End-to-End Speech Recognition

Ye Bai, Jie Li, Wenjing Han, Hao Ni, Kaituo Xu, Zhuo Zhang, Cheng Yi, Xiaorui Wang

Kuaishou Technology Co., Ltd, Beijing, China

{baiye03, lijie03, hanwenjing, nihao, xukaituo, zhangzhuo03, chengyi03, wangxiaorui}@kuaishou.com

## Abstract

While transformers and their variant conformers show promising performance in speech recognition, the parameterized property leads to much memory cost during training and inference. Some works use cross-layer weight-sharing to reduce the parameters of the model. However, the inevitable loss of capacity harms the model performance. To address this issue, this paper proposes a parameter-efficient conformer via sharing sparsely-gated experts. Specifically, we use sparsely-gated mixture-of-experts (MoE) to extend the capacity of a conformer block without increasing computation. Then, the parameters of the grouped conformer blocks are shared so that the number of parameters is reduced. Next, to ensure the shared blocks with the flexibility of adapting representations at different levels, we design the MoE routers and normalization individually. Moreover, we use knowledge distillation to further improve the performance. Experimental results show that the proposed model achieves competitive performance with 1/3 of the parameters of the encoder, compared with the full-parameter model.

**Index Terms:** parameter-efficient, sparsely-gated mixture-of-experts, Conformer, cross-layer weight-sharing

## 1. Introduction

Nowadays, transformers and their variants have been successfully applied to end-to-end (E2E) automatic speech recognition (ASR) [1, 2, 3]. Transformers usually use stacks of self-attention and feed-forward networks (FFNs) to build an encoder and a decoder [1], and then use the attention mechanism to bridge the encoded acoustic features and the representations of text token sequences [4]. Lately, as a variant, conformers are developed to augment transformers with convolution by helping the model capture locality [3]. Combined with techniques such as relative positional representations [5, 6] and Macaron-style half-step FFNs [7], conformers further improve the performance of transformers in ASR.

Despite the promising performance, many works show the over-parameterization of transformers [8, 9], which leads the models to require much memory storage during training and inference, and hence limits the usage of the models on-device. To reduce the memory cost, some works share the parameters of one or several transformer blocks so that the total number of the parameters of the model is much reduced [9, 10, 11, 12, 13]. These models use one or a few transformer blocks to encode features in a recursive manner, thus the number of parameters is less than the original transformers with the same depth. However, because of the fewer free model parameters, the capacity of network is inevitably influenced and the performance degrades as a result.

To address this issue, we propose to share the sparsely-gated mixture-of-experts (MoE) to improve the capacity of cross-layer parameter-shared conformers, and not increase the

computation in the meanwhile. Specifically, we first design the second FFN of a conformer block to be a sparsely-gated MoE module to improve the capacity and share the grouped conformer blocks in a cross-layer manner. The sparsely-gated MoE uses a dynamic routing mechanism to activate only one or a part of experts during training and inference, which keeps comparable computation with non-MoE models and scales the capacity of the model [14, 15, 16, 17]. Then, to help the parameter-shared conformer blocks to adapt the hidden representations at different levels, we propose to make each block have its own router so that the blocks can have flexible routing paths for different level representations. We also use individual normalization layers of the blocks to make them adaptable and to ensure the statistics consistent as well [18]. Further, we use knowledge distillation [19, 20] to help parameter-shared model imitate the full-parameter model. Experimental results on the public AISHELL-1 dataset demonstrate that the proposed parameter-efficient models can achieve competitive performance with 1/3 of encoder parameters, compared with the full-parameter model.

## 2. Background: Conformer-based Seq2Seq Models for ASR

As attention-based encoder-decoder (AED) models, transformers [1] use an encoder to capture the high-level representations from the acoustic features and a decoder to predict text sequences token-by-token with the attention mechanism. Formally, given an acoustic feature sequence  $x = [x_0, \dots, x_{T-1}]$  with length  $T$  and a text token sequence  $y = [y_0, \dots, y_S]$  with length  $S + 1$ , where  $y_0$  and  $y_S$  are the start-of-sentence symbol  $\langle \text{sos} \rangle$  and the end-of-sentence symbol  $\langle \text{eos} \rangle$ , the model  $Trfm$  predicts the probability of the text token:

$$P(y_s | y_{<s}, x) = Trfm(y_{<s}, x), \quad (1)$$

where  $y_{<s}$  is the prefix of  $y_s$  in the text sequence,  $1 \leq s \leq S$ . The model is trained with maximum likelihood criterion:

$$L_{\text{ml}}(\theta) = -\frac{1}{S} \sum_{s=1}^S \log P(y_s | y_{<s}, x), \quad (2)$$

where  $\theta$  is the parameters of the model  $Trfm$ . The beam-search algorithm is used to find the most likely text token sequence during inference. The overall structure is shown in Fig. 1.

Conformers [3] insert convolution layers into a transformer block to help the model capture locality of a sequence. With carefully designed fine-grained structures, including pre-norm [21], GLU [22] and Swish [23] activation functions, relative positional encodings [5, 6], conformers further improve the performance and stabilize the training. Our model chooses conformers as the basic structure of the encoder. The structure of the conformer block is shown in the middle part of Fig. 1. The details of each module in a conformer block are referred to [3].

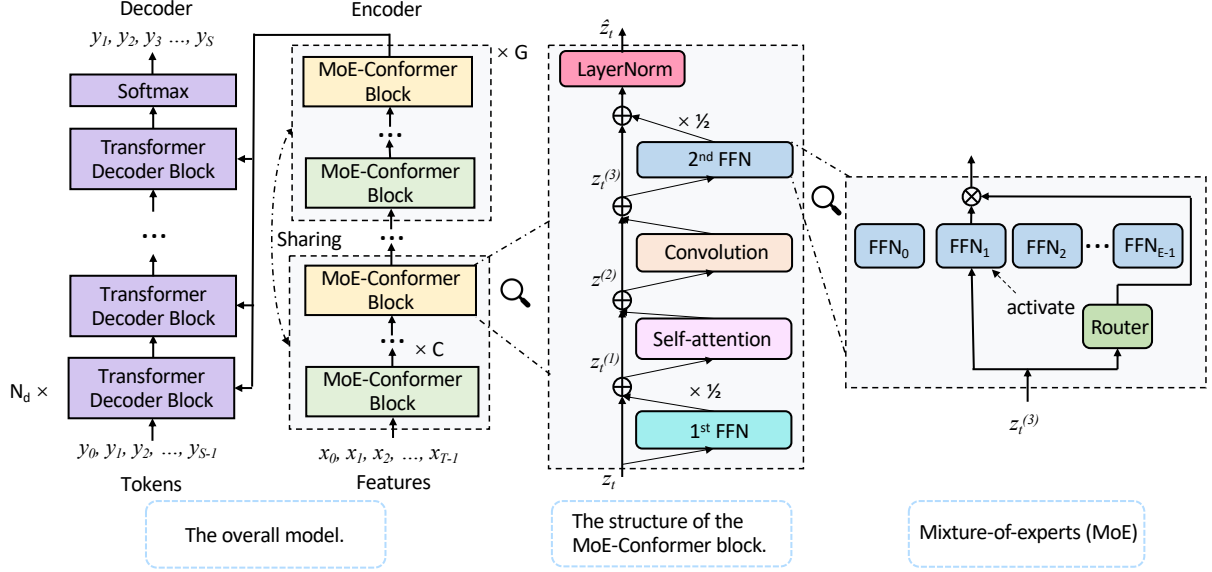


Figure 1: (Left) The overall architecture of the ASR system. The encoder consists of  $G$  groups of  $C$  consecutive MoE-conformer blocks. The parameters of the MoE-conformer blocks with the same color are shared among different groups, except the normalization modules and the routers of MoE module. The decoder consists of  $N_d$  transformer decoder blocks. (Mid.) The structure of the MoE-conformer block consists of two feed-forward network (FFN) modules, a convolution module, and a multi-head self-attention (MHSA) module. The details of each module are referred to [3]. We novelly extend the second FFN module to the mixture-of-experts (MoE) recipe. (Right) The structure of the MoE module. The MoE module consists of several parallel FFN modules and a router. During forward propagation, the input is fed into one of the FFN modules which is activated by the router.

### 3. Sharing Sparsely-Gated Experts

The core idea of the proposed parameter-efficient model is to reuse the conformer encoder blocks recursively to make the most of them. Crucially, the sparsely-gated MoE modules are used to improve the capacity of the modules without increasing computation. What's equally important, the routers and the normalization layers are further designed in an individual way, so that they can be used as adapters to help the reused blocks adapt representations at different levels. The method can also be used in the other network structures, such as transformer decoder blocks and convolutional neural networks, as well as in other E2E ASR models, such as transducers [24] and CTC [25].

#### 3.1. Parameter-Sharing for Conformers

The structure of a conformer, as shown in the middle part of Fig. 1, consists of two FFN modules, a multi-head attention module, and a convolution module. All the modules use pre-norm style combination of the residual connection and the layer normalization. Besides, the FFN modules use Swish activation functions. The multi-head self-attention module uses relative positional encodings. The convolution module is a time-depth separable style convolutional block with GLU and Swish activation functions. More details of each module are referred to [3]. Here, formally, for the input representation  $z_t$  at  $t$  time-step, the computation of a conformer block is as follows:

$$\begin{aligned}
 z_t^{(1)} &= z_t + \frac{1}{2}FFN(z_t), \\
 z_t^{(2)} &= z_t^{(1)} + MHSA(z_t^{(1)}), \\
 z_t^{(3)} &= z_t^{(2)} + Conv(z_t^{(2)}), \\
 \hat{z}_t &= LayerNorm(z_t^{(3)} + \frac{1}{2}FFN^{(MoE)}(z_t^{(3)})),
 \end{aligned} \tag{3}$$

where  $FFN$ ,  $MHSA$ ,  $Conv$  and  $FFN^{(MoE)}$  denote the first FFN module, the multi-head self-attention module, the convolution

module, and the second FFN enhanced with MoE, respectively.  $LayerNorm$  denotes layer normalization [26]. The details of  $FFN^{(MoE)}$  are described in Section 3.2

As shown in the left part of Fig. 1, we share the parameters of different blocks. Specifically,  $C$  consecutive conformer blocks are grouped and  $G$  groups are stacked. For the conformer block at the same position in different groups, the parameters of each module are shared. It can be viewed as one group of conformer blocks are reused  $G$  times, and the computation is implemented in a recursive iteration manner. Thus the model makes the most use of the parameters.

#### 3.2. Dynamic Routing for Mixture of Experts

By parameter-sharing, the parameters of the encoder are much reduced. However, the capacity of the model is also reduced which influences the performance of the model negatively. So, to improve the model capacity but not increase the computation, we introduce sparsely-gated MoE [14, 17] to the second FFN module, as shown in the right part of Fig. 1.

The sparsely-gated MoE mechanism consists of  $E$  parallel experts and a router. The input  $z_t^{(3)}$  is first fed into the router to select one of the experts<sup>1</sup> and then is computed by the activated expert. The formal computation is as follows:

$$\begin{aligned}
 g &= [g_0, \dots, g_{E-1}] = softmax(router(z_t^{(3)})), \\
 i^* &= \arg \max_{0 \leq i \leq E-1} g_i,
 \end{aligned} \tag{4}$$

$$FFN^{(MoE)}(z_t^{(3)}) = g_{i^*} FFN_{i^*}(z_t^{(3)}),$$

where  $FFN_i$  denotes the  $i$ -th expert,  $g_i$  denotes the gating value regarding to the  $i$ -th expert, and  $i^*$  is the index of the selected expert. One may notice that, the procedure of MoE is actually similar to the attention mechanism: the input  $z_t^{(3)}$  can be

<sup>1</sup>We use top-1 MoE to keep the number of the activated parameters the same with the non-MoE model in this paper.

viewed as the query vector in the attention mechanism, and the gating scores  $g$  can be viewed as the attention coefficients [1]. However, the attention procedure is in a “hard” way, namely, the non-maximum coefficients are all set to zero.

In addition, to encourage all the experts to be used in balance, the load balancing loss [17] is used as follows:

$$L_{\text{balance}} = E \sum_{i=0}^{E-1} f_i \bar{g}_i, \quad (5)$$

where  $f_i$  is the active frequency of the  $i$ -th expert in a batch, and  $\bar{g}_i$  is the mean of the gating values computed for the  $i$ -th expert. Otherwise, Gaussian noises are added to the routers to make the expert selection various during training.

With MoE, the parameters are extended so that the model capacity is increased. However, since only one FFN is activated actually, the computation is not increased.

### 3.3. Individual Routers and Normalization

To further improve the ability of the reused MoE modules, we propose to make each MoE module have its own router. The underneath thinking is to help the routing path achieve more flexibility. With this, the MoE modules in different MoE-conformer blocks can thus be adapted to different levels of representations. Furthermore, all normalization layers (including layer normalization and batch normalization) are built individually, thus, to maintain the statistics of the normalization layers corresponding to representations at different levels is consistent. And the scale and offset parameters in the normalization layers can be seen as parameter-efficient bias adapters [27].

### 3.4. Distilling Knowledge from Hidden Embedding

We use knowledge distillation [19, 20] to transfer the knowledge from a full-parameter model to further improve the performance of the shared-parameter model. Specifically, we minimize the  $L_2$  distance between the outputs of the shared-parameter encoder (student) and the full-parameter encoder (teacher):

$$L_{kd} = \frac{1}{T} \sum_{t=0}^{T-1} \|h_t - h'_t\|, \quad (6)$$

where  $h_t$  denotes the output of the shared-parameter encoder and  $h'_t$  denotes the output of the full-parameter encoder.

### 3.5. Learning

The model is learned by minimizing the overall loss:

$$L = L_{\text{nil}} + \alpha \frac{1}{C} \sum L_{\text{balance}} + \beta L_{kd}, \quad (7)$$

where  $C$  is the number of MoE module (see Fig. 1),  $\alpha$  and  $\beta$  are hyperparameters to balance the values of the losses.

## 4. Relation to Prior Work

**Conditional computation of mixture-of-experts.** MoE has been shown as an effective way to scale the capacity of neural networks without increasing computation [14, 15, 16, 17]. However, previous works aim to scale the model sizes to billions or trillions, which needs extremely much resources and model parallelization during training and inference. Different from these works, we aim to use MoE in a parameter-efficient way. We reuse the MoE modules to make the most use of them. **Cross-layer weight sharing.** Cross-layer weight sharing is first used in transformer with adaptive computation time [10]. [9] uses this technique to reduce the parameters of BERT.

Table 1: *The overall character error rates on AISHELL-1.  $N_{pe}$  denotes the total number of the parameters of the encoder. Dev. and Test denote the character error rates (CERs) on the development set and the test set, respectively.*

Model	$N_{pe}$	Dev.	Test
C12	21.58M	4.46	4.93
C2	3.74M	5.86	6.50
C2-MoE4	6.89M	5.77	6.22
C2-G6	3.74M	5.18	5.62
C2-MoE4-G6	6.95M	4.67	5.08
C2-MoE4-G6-KD	6.95M	<b>4.65</b>	<b>5.03</b>

[12, 13, 28] use the similar techniques for ASR. However, directly sharing parameters may influence the capacity of the model negatively. To address this issue, we propose to use the MoE mechanism to improve the model capacity without increasing computation. Recently, [18] shares the MoE module for ALBERT and ViT and applies the models to NLP and CV tasks. However, their work uses two experts, which increases computation cost. Otherwise, sharing the routers limits the capacity of the models. Different from their work, this paper focuses on more efficient architecture of Conformers [3] in ASR tasks. We use individual routers to help the model to have diverse routing paths at different levels. And we use group strategy to improve the model capacity in depth.

## 5. Experiments

### 5.1. Experimental Setup

We conduct experiments on a publicly available Chinese Mandarin AISHELL-1<sup>2</sup> dataset [29], which includes about 150 hours of speech for training, about 18 hours of speech for development, and about 10 hours speech for test.

For all the experiments, we use 80-dimension Mel-filter bank features (FBANK) as the input, which are extracted every 10 ms with 25 ms window. We use global cmvn as feature normalization. Speed perturbation with factors of 0.9, 1.0 and 1.1 is used as audio augmentation [30]. All the feature processing is employed with Kaldi toolkit [31]. We use 4235 Chinese characters as the vocabulary, including <eos> and <eos>.

We use a 2-layer CNN as a subsampling module. Each layer is a  $3 \times 3$  convolutional layer with 32 output channels, and the stride is 2. Thus, the frame rate is subsampled to 25 Hz. For the encoder, we set the dimension of an MoE-Conformer module to 256, the number of heads of MHSA to 4, the kernel size of Conv to 15. The intermediate dimension of an FFN module is 1024. We use 4 experts for the second FFN in an MoE-Conformer module. We compare effects of the different number of MoE-Conformer modules and groups, i.e.,  $C$  and  $G$  in Fig. 1. For the decoder, we use the transformer structure. To control experimental variables, we fix the number of the decoder blocks to 4. The dimension of the decoder module is also 256 and the intermediate dimension of the FFN in the decoder module is 1024. We set dropout rate to 0.1 and use SpecAugmentation [32] and time stretch [33] to avoid overfitting. The values of  $\alpha$  and  $\beta$  in Eq.(7) are set to 0.01 and 0.005, respectively. The standard deviation of Gaussian noise for the MoE gate is set to 0.1 in training. CTC loss is also used to improve the alignment with weight 0.2. The learning rate schedule is inverse square root with 4000 warm-up steps. All models are trained for 80 epochs with 8 GPU cards. One batch includes 32000 frames. We use PyTorch [34] and FastMoE [35] for implementation.

<sup>2</sup><https://www.openslr.org/33/>

Table 2: Ablation studies on AISHELL-1. “n.” denotes normalization modules, and “r.” denotes routers. “indiv.” means that the corresponding modules are not shared.

(a) w/ MoE vs. w/o MoE.

Model	$N_{pe}$	Dev.	Test
C1	1.95M	7.53	8.41
C1-MoE4	3.53M	7.26	8.05
C2	3.74M	5.86	6.50
C2-MoE4	6.89M	5.77	6.22

(b) w/ parameter-sharing vs. w/o parameter-sharing.

Model	$N_{pe}$	Dev.	Test
C1	1.95M	7.53	8.41
C1-G12	1.95M	5.65	6.07
C1-MoE4-G12	3.59M	5.01	5.40
C2	3.74M	5.86	6.50
C2-G6	3.74M	5.18	5.62
C2-MoE4-G6	6.95M	4.67	5.08

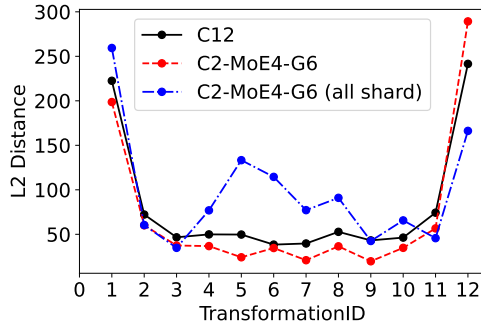


Figure 2: L2 distances between the input and output for each transformation.

## 5.2. Results and Analysis

**Overall.** Table 1 shows the overall performance of our model. C denotes the number of conformer blocks and G denotes the number of groups (see Fig. 1). MoE4 and KD mean whether to use MoE and knowledge distillation, respectively. We can see that with the proposed methods, C2-MoE4-G6-KD achieves competitive performance with 1/3 parameters of the encoder, compared with the full-parameter model C12. Directly reducing the number of blocks to 2 hurts the performance of the model (C2). MoE improves the performance of the models and does not increase the activated parameters.

**w/ MoE vs. w/o MoE.** We compare the shallow encoders with MoE and the ones without MoE in Table 2a. We can see that MoE improves the capacity so that the performance of C1-MoE4 is better than C1 and the performance of C2-MoE4 is better than C2. And with more conformer blocks, the model can perform better but with more parameters (C2 vs. C1 and C2-MoE4 vs. C1-MoE4).

**Recursive iteration.** Table 2b shows that more recursive iterations make better performance with the same number of parameters. Specifically, for C1-G12, the 12 groups of blocks are shared, which can be seen as the group is computed recursively, and the performance is much better than the non-shared model C1. Similarly, C2-G6 performs better than C2 with the same number of parameters. C2-G6, which has more blocks in one group than C1-G12, performs better than C1 with the same computation iteration. MoE improves capacity and performance of C1-G12 and C2-G6.

**Individual routers and normalization.** Table 2c compares the

(c) Individual routers and normalization.

Model	$N_{pe}$	Dev.	Test
C1-MoE4-G12 (all shared)	3.53M	6.39	6.90
C1-MoE4-G12 (indiv. n.)	3.58M	5.19	5.57
C1-MoE4-G12 (indiv. n. & r.)	3.59M	5.01	5.40
C2-MoE4-G6 (all shared)	6.89M	5.60	6.00
C2-MoE4-G6 (indiv. n.)	6.94M	4.72	5.21
C2-MoE4-G6 (indiv. n. & r.)	6.95M	4.67	5.08

(d) Knowledge distillation from hidden embeddings.

Model	$N_{pe}$	Dev.	Test
C12 (teacher)	21.58M	4.46	4.93
C1-MoE4-G12	3.53M	5.01	5.40
C1-MoE4-G12-KD	3.53M	4.99	5.43
C2-MoE4-G6	6.95M	4.67	5.08
C2-MoE4-G6-KD	6.95M	4.65	5.03

effect of the individual routers and normalization. We can see that if the routers and the normalization modules are all shared, the performance of the parameter-shared model is hurt heavily. Keeping each conformer block having its own normalization module so that the performance of the parameter-shared model is better. And the individual routers make the MoE module be able to select proper experts at different levels. Thus, the model with individual routers and normalization can achieve better performance.

**Knowledge distillation from hidden embeddings.** We further use knowledge distillation to make the parameter-sharing model imitate the full-parameter model (Table 2d). We can see that with knowledge distillation, the performance of the parameter-shared model is further improved for C2-MoE4-G6 model. However, the improvement is not very significant for C1-MoE4-G12. This is probably because, C1 model has far less conformer blocks when comparing to C12 model (1 vs. 12), then such a big divergence influences the effect of knowledge distillation [36] between them.

**L2 distances between the input and the output.** Fig. 2 shows the L2 distances of the input and output of each transformation for an example utterance. We can see that C2-MoE4-G6 shows a similar behavior with the full-parameter model C12. Whereas the curve of the all-shared model is oscillating. This shows that the individual routers and normalization have an effect on stabilizing network parameters.

## 6. Conclusions and Future Works

This paper explores sharing the sparsely-gated mixture-of-experts (MoE) to build a parameter-efficient conformer model for speech recognition. Specifically, we first use MoE to extend the capacity of a conformer block. Then, we share the parameters of the grouped conformer blocks so that the parameters are much reduced compared with the full-parameter model. To ensure the representations adapt at different levels, we make the routers of MoE and normalization modules individual. Moreover, we use knowledge distillation to further improve the performance. The experimental results demonstrate that the proposed model can achieve competitive performance with about 1/3 of the parameters of the encoder, compared with the full-parameter model. In the future, we will extend the proposed method to more large-scale datasets and other ASR models, such as transducers and CTC.

## 7. References

- [1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [2] L. Dong, S. Xu, and B. Xu, "Speech-transformer: a no-recurrence sequence-to-sequence model for speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5884–5888.
- [3] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu *et al.*, "Conformer: Convolution-augmented transformer for speech recognition," *arXiv preprint arXiv:2005.08100*, 2020.
- [4] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *2016 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2016, pp. 4960–4964.
- [5] P. Shaw, J. Uszkoreit, and A. Vaswani, "Self-attention with relative position representations," *arXiv preprint arXiv:1803.02155*, 2018.
- [6] Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Q. V. Le, and R. Salakhutdinov, "Transformer-xl: Attentive language models beyond a fixed-length context," *arXiv preprint arXiv:1901.02860*, 2019.
- [7] 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.
- [8] A. Fan, E. Grave, and A. Joulin, "Reducing transformer depth on demand with structured dropout," *arXiv preprint arXiv:1909.11556*, 2019.
- [9] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, "Albert: A lite bert for self-supervised learning of language representations," *arXiv preprint arXiv:1909.11942*, 2019.
- [10] M. Dehghani, S. Gouws, O. Vinyals, J. Uszkoreit, and Ł. Kaiser, "Universal transformers," *arXiv preprint arXiv:1807.03819*, 2018.
- [11] S. Li, C. Ding, X. Lu, P. Shen, T. Kawahara, and H. Kawai, "End-to-end articulatory attribute modeling for low-resource multilingual speech recognition," in *INTERSPEECH*, 2019, pp. 2145–2149.
- [12] Y. Zhao, C. Ni, C.-C. Leung, S. R. Joty, E. S. Chng, and B. Ma, "Universal speech transformer," in *INTERSPEECH*, 2020, pp. 5021–5025.
- [13] T. Komatsu, "Non-autoregressive asr with self-conditioned folded encoders," *arXiv preprint arXiv:2202.08474*, 2022.
- [14] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean, "Outrageously large neural networks: The sparsely-gated mixture-of-experts layer," *arXiv preprint arXiv:1701.06538*, 2017.
- [15] C. Riquelme, J. Puigcerver, B. Mustafa, M. Neumann, R. Jenatton, A. Susano Pinto, D. Keysers, and N. Houlsby, "Scaling vision with sparse mixture of experts," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [16] Z. You, S. Feng, D. Su, and D. Yu, "Speechmoe: Scaling to large acoustic models with dynamic routing mixture of experts," *arXiv preprint arXiv:2105.03036*, 2021.
- [17] W. Fedus, B. Zoph, and N. Shazeer, "Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity," *arXiv preprint arXiv:2101.03961*, 2021.
- [18] F. Xue, Z. Shi, F. Wei, Y. Lou, Y. Liu, and Y. You, "Go wider instead of deeper," *arXiv preprint arXiv:2107.11817*, 2021.
- [19] G. Hinton, O. Vinyals, J. Dean *et al.*, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, vol. 2, no. 7, 2015.
- [20] J. Li, R. Zhao, J.-T. Huang, and Y. Gong, "Learning small-size dnn with output-distribution-based criteria," in *Fifteenth annual conference of the international speech communication association*, 2014.
- [21] R. Xiong, Y. Yang, D. He, K. Zheng, S. Zheng, C. Xing, H. Zhang, Y. Lan, L. Wang, and T. Liu, "On layer normalization in the transformer architecture," in *International Conference on Machine Learning*. PMLR, 2020, pp. 10524–10533.
- [22] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, "Language modeling with gated convolutional networks," in *International conference on machine learning*. PMLR, 2017, pp. 933–941.
- [23] P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for activation functions," *arXiv preprint arXiv:1710.05941*, 2017.
- [24] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [25] 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*, 2006, pp. 369–376.
- [26] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [27] E. B. Zaken, S. Ravfogel, and Y. Goldberg, "Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models," *arXiv preprint arXiv:2106.10199*, 2021.
- [28] 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 *InterSpeech*, 2019, pp. 4400–4404.
- [29] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline," in *2017 20th Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment (O-COCOSDA)*. IEEE, 2017, pp. 1–5.
- [30] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, "Audio augmentation for speech recognition," in *Sixteenth annual conference of the international speech communication association*, 2015.
- [31] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The kaldı speech recognition toolkit," in *IEEE 2011 workshop on automatic speech recognition and understanding*, no. CONF. IEEE Signal Processing Society, 2011.
- [32] 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," *arXiv preprint arXiv:1904.08779*, 2019.
- [33] T.-S. Nguyen, S. Stueker, J. Niehues, and A. Waibel, "Improving sequence-to-sequence speech recognition training with on-the-fly data augmentation," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7689–7693.
- [34] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," *Advances in neural information processing systems*, vol. 32, 2019.
- [35] J. He, J. Qiu, A. Zeng, Z. Yang, J. Zhai, and J. Tang, "Fastmoe: A fast mixture-of-expert training system," *arXiv preprint arXiv:2103.13262*, 2021.
- [36] J. H. Cho and B. Hariharan, "On the efficacy of knowledge distillation," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 4794–4802.