



InterFormer: Interactive Local and Global Features Fusion for Automatic Speech Recognition

Zhi-Hao Lai^{1,2}, Tian-Hao Zhang^{1,2}, Qi Liu^{1,2}, Xinyuan Qian¹, Li-Fang Wei^{1,2},
Song-Lu Chen^{1,2}, Feng Chen^{2,3}, Xu-Cheng Yin^{1,2,*}

¹School of Computer and Communication Engineering, University of Science and Technology Beijing, Beijing 100083, China

²USTB-EEasyTech Joint Lab of Artificial Intelligence, University of Science and Technology Beijing, Beijing 100083, China

³EEasy Technology Company Ltd., Zhuhai 519000, China

{zhihaolai, tianhaozhang}@xs.ustb.edu.cn, {xinyuanqian, songluchen}@xs.ustb.edu.cn

Abstract

The local and global features are both essential for automatic speech recognition (ASR). Many recent methods have verified that simply combining local and global features can further promote ASR performance. However, these methods pay less attention to the interaction of local and global features, and their series architectures are rigid to reflect local and global relationships. To address these issues, this paper proposes InterFormer for interactive local and global features fusion to learn a better representation for ASR. Specifically, we combine the convolution block with the transformer block in a parallel design. Besides, we propose a bidirectional feature interaction module (BFIM) and a selective fusion module (SFM) to implement the interaction and fusion of local and global features, respectively. Extensive experiments on public ASR datasets demonstrate the effectiveness of our proposed InterFormer and its superior performance over the baseline models.

Index Terms: automatic speech recognition, bidirectional interaction, parallel structure, selective feature fusion

1. Introduction

The local and global features are both essential for ASR. It is well known that convolutional neural network (CNN) can effectively exploit local features while Transformer-based models are highly effective for capturing long-range feature relationships that constitute global representations due to the self-attention mechanism. Therefore, the CNN and transformer models have been the building blocks of ASR [1, 2, 3, 4, 5, 6]. Unfortunately, global features are not well captured by CNN due to its limited receptive field [7, 8]. Whereas the details of local features are ignored by Transformer models [9, 10, 11].

To alleviate the above limitations, recent works [7, 9, 12, 13, 14, 15] have tried to combine local and global features. ContextNet [7] simply adds squeeze-and-excitation [16] module in each residual block to capture long-distance feature dependencies. Conformer [12] proposes to add a convolution module that captures local features on top of multi-head self-attention to augment the Transformer. These methods have shown that models combining local and global features outperform the models using local and global features separately in end-to-end ASR.

However, these structured serial methods rarely focus on the interaction of local and global features, which causes the following problems. First, their serial architecture makes it hard to model various ranged dependencies and reflects local and global relationships in different layers, leading to latent information loss in the fusion process. Second, the lack of effective interaction hinders further advances in feature representation ability.

* Corresponding author.

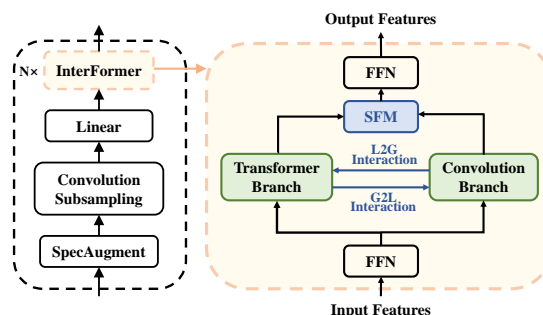


Figure 1: *The overall architecture of InterFormer.* It consists of two macaron-like feed-forward layers (FFN) and dual parallel branches. The blue arrows marked with L2G Interaction and G2L Interaction are then fused effectively in SFM. N : number of identical InterFormer blocks

To address the aforementioned issues, in this paper, we propose a parallel feature fusion network, namely InterFormer, to interactively fuse local and global features. As shown in Fig. 1, our parallel structure consists of two branches. Convolution branch employs a convolution block to capture feature neighboring information, while Transformer Branch utilizes the multi-head self-attention (MHSA) with relative positional encoding to capture global dependencies. Moreover, for the interaction between local and global features, BFIM is designed as a bridge to provide complementary information for the transformer and the convolution branches, respectively, enhancing the feature representation learning ability in local and global aspects. Taking into account the effective fusion of local and global features, SFM is proposed to dynamically choose valuable characteristics based on the weight of the two branches. Such an interactive feature fusion method has a better representation capability to improve the ASR performance.

We conduct extensive experiments on three ASR datasets and compare the results with the state-of-the-art Transformer [17] and Conformer [17]. Experimental results show that our InterFormer model achieves the best performance. Furthermore, extensive ablation studies demonstrate the effectiveness of our proposed BFIM and SFM.

2. InterFormer

In this section, given an overview of InterFormer in Fig. 1, we elaborate each internal block, as illustrated in Fig. 2. Specifically, our model takes the raw audio waveform with a convolutional subsampling module to downsample the feature se-

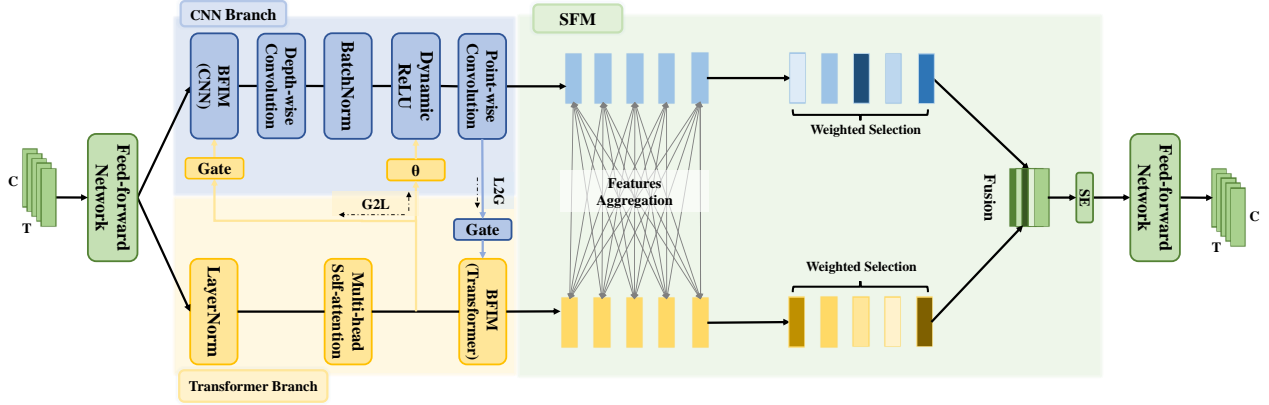


Figure 2: **The overall architecture of InterFormer.** Implementation details of the convolution branch, transformer branch, BFIM, and SFM. The convolution branch starts with a BFIM, followed by a single 1-D depthwise convolution layer, batchnorm, dynamic ReLU, and a pointwise convolution. The composition of the transformer branch is layernorm, a multi-head self-attention module and a BFIM. (SE: Squeeze-and-Excitation module; C: Channel dimension; T: Time dimension.)

quence, followed by the positional embedding addition. To take advantage of local and global representations for feature interaction, we employ a stack of N identical InterFormer blocks.

As shown in Fig. 1, InterFormer replaces the serial connection in Conformer with a parallel connection. It consists of two branches where we extract local or global features using convolution operations or attention in each branch. Considering the lack of interaction between two branches, we propose BFIM to interact local features in the convolution branch with global representations in the transformer branch. To effectively fuse local and global information extracted from the two branches, SFM is proposed to weigh the importance of local and global features and dynamically select the most crucial characteristics.

2.1. Parallel Branches for Extracting Local and Global Features

Our InterFormer adopts dual parallel branches based on the standard self-attention and convolution elaborated as follows.

Convolution branch. we employ the convolution branch to extract local feature details. As shown in Fig. 2, the composition of the convolution branch is mainly similar to the convolution module of Conformer. It starts with a BFIM, which is a BFIM used to implement the interaction of local and global information. This is followed by a single 1-D depthwise convolution layer, batchnorm, dynamic ReLU, and a pointwise convolution.

Transformer branch. The transformer branch in InterFormer aims to capture the long-distance feature dependencies that constitute global representations. As shown in Fig. 2, this branch consists of a MHSA module while adopting the relative positional encoding scheme [18] and a BFIM block.

2.2. Bidirectional Feature Interaction Module (BFIM)

Considering the lack of information interaction and the feature misalignment between convolution and transformer features, BFIM is designed as the bridge for information exchange. The parallel branches contains the local-to-global (L2G) and global-to-local (G2L) interactions, respectively. As shown in Fig. 2, the information in the convolution branch flows to the other branch through the L2G interaction, strengthening the global features modeling ability. Meanwhile, the G2L Interaction enables global features to flow from the transformer branch to the convolution branch, thus enhancing the local representa-

tion ability. As a result, the proposed bidirectional interactions provide complementary clues for each other and eliminate the feature misalignment between convolution and transformer features.

The L2G interaction. The input in the Transformer branch is passed through the attention mechanism to obtain the global feature G , followed by a layernorm and pointwise convolution. Then we concatenate the local feature L and global feature G where L comes from the convolution branch, and the local information L acts as gating information to influence global features. The interaction process is depicted in Fig. 3. We compute the hidden features $h = \{h_0, h_1, \dots, h_t\}$ in the convolution branch as:

$$h_g(x) = (PW \text{ Conv}(x) + b) \odot \sigma(L) \quad (1)$$

where $PW \text{ Conv}$ is the pointwise convolution, $x \in R^{T \times d}$, T is the sequence length, d is the feature dimension, $b \in R^T$, σ is the sigmoid function, and \odot is the element-wise product between matrices.

The G2L interaction. For the G2L direction, we apply the same operation with the L2G direction to realize the interaction of local and global features:

$$h_l(x) = (PW \text{ Conv}(x) + b) \odot \sigma(G) \quad (2)$$

To further exploit the global information for better information interaction, the following adjustments are applied to the activation function for the convolution branch. Because the interaction between local and global information is a dynamic adjustment process, using the previous static activation function, such as swish and ReLU, may degrade the performance. Therefore, we replace the previous static activation function with the dynamic ReLU activation function to achieve dynamic interaction of local and global features. Specifically, the acquisition of θ is different from the original dynamic ReLU, which uses the original input x to obtain. In our method, the output of θ is obtained through the global feature G from the transformer branch. It applies two MLP layers (with a ReLU between them) and then uses a function to map the global information between -1 and 1. The process of dynamic ReLU is formulated as:

$$\theta = 2\sigma(W_2\gamma(W_1G)) - 1 \quad (3)$$

where γ is the ReLU activation, W_1 and W_2 denotes the MLP layer. The above two methods achieve feature interaction.

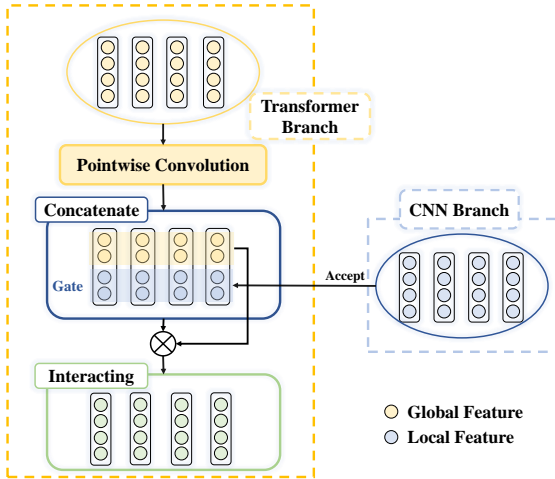


Figure 3: **Detailed design of the Interactions.** The Transformer branch accepts information from the CNN branch for interaction. Local features from CNN branch act as gating information to influence global features. The information flow is shown by the blue arrow in Fig. 2.

2.3. Fusion of Local and Global Features

Both local and global features are essential to improve the representation capability of the model. To perform the fusion of local and global features, most previous methods use direct addition [9, 12] or concatenation [19, 20], which are too simple and effective to selectively highlight the varying significance of the local and global features. To weigh the importance of each branch in different layers and efficiently combine information from the dual branches, we propose SFM to fuse local and global representations in a dynamic way.

2.3.1. Direct Addition and Concatenation

Let Y_{conv} and Y_{trans} be the output sequences of the convolution branch and transformer branch, respectively. In the direct addition method, we add the features of the two branches along the time dimension directly:

$$Y_{fused} = Y_{conv} + Y_{trans} \quad (4)$$

In the method of concatenation, we first concatenate the features of the two branches along the feature dimension and then use a linear transformation to make the result of the original dimension:

$$Y_{fused} = Concat(Y_{conv}, Y_{trans})W_{fuse} \in R^{2d \times d} \quad (5)$$

2.3.2. Selective Fusion Module (SFM)

In the SFM, it weighs the importance of local and global features and selects the corresponding features based on the obtained attention weights. From a macro perspective, the information obtained from the transformer branch is equivalent to a global receptive field acting on the transformer branch, while the information obtained from the convolution branch corresponds to acting on the convolution branch in a local receptive field. Inspired by Selective Kernel Networks [21], we selec-

tively fuse the two interacted features L and G from the convolution branch and transformer branch, respectively. The core idea is to control whether the information of the interactions is retained or forsaken, as shown in Fig. 2. To this end, we first concatenate features from two branches:

$$X = Concat(L, G) \quad (6)$$

Then, we calculate the mean value of each row of the input tensor in time dimension, followed by an MLP layer with a ReLU activation to reduce the channel dimension for better efficiency:

$$I = F_m = \frac{1}{T} \sum_{i=1}^T x_i \\ \hat{X} = F_f = \gamma(W_f(I)) \quad (7)$$

where γ is the ReLU activation, $W_f \in R^{c \times d}$ denotes the MLP layer. Further, as shown in Eq.(8), we use two MLP layers to act on \hat{X} respectively to obtain two features that are restored to the original dimension and then concatenate them for information aggregation. These operations can generate a feature Z , guiding accurate and adaptive selections.

$$Z = Concat(X_1 = W_{u1}(\hat{X}), X_2 = W_{u2}(\hat{X})) \quad (8)$$

where $W_{u1} \in R^{d \times c}$, $W_{u2} \in R^{d \times c}$. It is well-known that the softmax function can introduce competition between attention weights, which is significant for obtaining selective attention. Thus, a softmax function is used to learn a selective factor α , and it can select appropriate features from local and global aggregated features. The output fusing convolution and transformer branch X_f is generated by Eq.(9):

$$X_f = Softmax(Z) \odot X \quad (9)$$

Finally, to exploit contextual information, we use the Squeeze-and-Excitation module [16] to weigh the sum of the output by channel attention mechanism.

3. Experiments

3.1. Experimental Setting

In this paper, we train and evaluate our proposed models on three public ASR datasets: Aishell-1 [22], Tedlium3 [23] and Librispeech [24]. Aishell-1 corpus has 170 hours of labeled speech collected from 400 speakers. Tedlium3 corpus was made from audio talks and contains 456 hours of 16kHz English TED talks. Librispeech consists of about 960 hours of 16kHz English read audiobooks. We choose them for experiments since they are well-known benchmarks for ASR tasks.

All our experiments are conducted on ESPNet toolkit [17], which uses PyTorch [25] as a deep learning engine. For feature extraction, we use kaldi to extract input spectrograms of 80-dimensional mel-scale log filter banks computed from a 25ms window with a stride of 10ms. During acoustic front-end processing, we used SpecAugment [26] and Speed Perturbation [27] on Aishell-1, for Tedlium3 and LibriSpeech dataset, we only use SpecAugment for data augmentation. The 2D-CNN front-end is comprised of 2 convolution layers with kernel size 3x3, filter 256, and stride 2 on the time dimension for 4x down-sampling. For different datasets, we adjust the attention head, attention dimension, hidden size and the number of encoder and decode layers to get better results. The Adam [28] optimizer and warm up with $\beta_1 = 0.9$, $\beta_2 = 0.98$, and $\epsilon = 10^{-9}$. we employ label smoothing [29] of value 0.1 and dropout [30] with

Table 1: Experimental results present the Character Error Rate (in %) of our proposed model and other previous models on the Aishell-1 Mandarin ASR task.

Method	Dev	Test	Params(M)
Citrinet[8]	5.2	5.71	142
ESPNet[17]			
Transformer	6.0	6.7	29.7
Conformer	4.6	5.1	46.2
InterFormer(ours)	4.4	4.9	46.8

Table 2: Word Error Rates (in %) comparison among ASR models with Transformer, Conformer and InterFormer encoders.

Datasets	Evaluation Sets	Transformer	Conformer	InterFormer
LibriSpeech	test/test-other	4.2/10.5	4.1/9.2	3.7/9.0
Tedlium3	test	11.6	10.5	9.5

probability 0.1 to prevent overfitting. Our experiments are conducted using 2 Nvidia 2080ti GPUs and use CER and WER as evaluation metrics, with the Chinese dataset using CER as an evaluation metric and the English dataset using WER as an evaluation metric.

3.2. Main Results and Performance Analysis

In this section, we compare the proposed InterFormer model with other state-of-the-art models. For fairness, we reproduced Transformer and Conformer baselines with the same configurations, training, and decoding strategies. All models are trained using the ESPNet toolkit without using any language model.

Table 1 compares the parameter sizes and character error rates (CER) on Aishell-1 dataset. The experimental results show that our model is superior to the Conformer and Transformer models. Compared with the Conformer model, our model outperforms the Conformer by 0.2%.

Table 2 compares WERs on LibriSpeech and Tedlium3. For the LibriSpeech dataset, because we use a configuration file with fewer parameters, without speed perturbation and language model, and fewer training epochs, the reproduced baselines are worse than those reported by other open-source toolkits. When evaluated on test/test-other of the Librispeech dataset under the same configuration, the InterFormer outperforms Transformer and Conformer models with promising gains: +0.5/1.5% and +0.4/0.2%, respectively. When evaluated on the Tedlium3 dataset, the InterFormer achieve results of 9.5 on test outperforming the Transformer and Conformer.

3.3. Ablation Studies

In this section, we perform several detailed ablation studies to evaluate the contribution of each component and the effectiveness of different methods. All models in ablation experiments are trained for 50 epochs on 2 Nvidia 2080 Ti GPUs with batch of 32 per GPU.

3.3.1. Effect of BFIM

Table 3 shows the results of the proposed BFIM. Starting from a Conformer block, We parallelize its convolution module and Transformer module and then continue to add various interactive components. According to the results, we can see that both L2G and G2L interactions outperform the model with only par-

Table 3: Ablation study of BFIM on the Aishell-1 corpus. The baseline model in this table is Conformer with a successive design and has no interactions. * means to remove dynamic ReLU.

Parallel	Interaction		Aishell1	
	L2G	G2L	dev	test
			4.6	5.1
✓			4.6	5.0
✓	✓		4.5	5.0
✓		✓	4.5	4.9
✓		✓*	4.5	5.0
✓	✓	✓	4.4	4.9
		Δ	+0.2	+0.2

Table 4: Comparison of different fusion methods.

Method	Interaction	Params(M)	CER(%)	
			Dev	Test
Conformer	✗	46.2	4.6	5.1
Direct addition	✗	46.2	4.7	5.1
Direct addition	✓	46.5	4.5	5.0
Concatenation	✗	47.8	4.6	5.1
Concatenation	✓	48.4	4.6	5.0
SFM(InterFormer)	✗	46.6	4.6	5.0
SFM(InterFormer)	✓	46.8	4.4	4.9

tial interaction or no interaction. We only use a simple and light-weight design for BFIM, the gains are extraordinary, which shows that BFIM can effectively provide important complementary clues for the convolution branch and the Transformer branch.

3.3.2. Comparison of Different Fusion Operations

Our proposed SFM method and the other two fusion methods are discussed in Section 2.3. Compared to the other two fusion methods, our proposed SFM has the ability to dynamically fuse the features of the convolutional and Transformer branches. Table 4 compares their performance on Aishell-1. According to the table information, we can see that our proposed SFM is superior to the above two methods, which proves the effectiveness of our proposed fusion operations.

4. Conclusion

In this paper, we propose InterFormer as an efficient dual parallel branches encoder for ASR. Within InterFormer, we leverage the convolution branch to extract local features and the transformer branch to capture global representations. BFIM promotes the feature interaction between the convolution branch and the transformer branch and alleviates the feature mismatch between the two branches. Moreover, we design a SFM to fuse local features and global representations effectively, enhancing the modeling ability of audio representations. Experiments show that the proposed InterFormer outperforms the Conformer with comparable parameters. Furthermore, extensive ablation experiments have proved the effectiveness of BFIM and SFM.

Acknowledgements: The research is supported by National Key Research and Development Program of China (2020AAA0109700), National Natural Science Foundation of China (62076024, 62006018, U22B2055).

5. References

- [1] S. Krivan, 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," in *ICASSP IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2020, pp. 6124–6128.
- [2] V. Pratap, A. Hannun, Q. Xu, J. Cai, J. Kahn, G. Synnaeve, V. Liptchinsky, and R. Collobert, "Wav2letter++: A fast open-source speech recognition system," in *ICASSP IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2019, pp. 6460–6464.
- [3] C.-F. Yeh, J. Mahadeokar, K. Kalgaonkar, Y. Wang, D. Le, M. Jain, K. Schubert, C. Fuegen, and M. L. Seltzer, "Transformer-transducer: End-to-end speech recognition with self-attention," *arXiv preprint arXiv:1910.12977*, 2019.
- [4] T. Hori, N. Moritz, C. Hori, and J. Le Roux, "Transformer-based long-context end-to-end speech recognition," in *Interspeech*, 2020.
- [5] Y. Wang, A. Mohamed, D. Le, C. Liu, A. Xiao, J. Mahadeokar, H. Huang, A. Tjandra, X. Zhang, F. Zhang *et al.*, "Transformer-based acoustic modeling for hybrid speech recognition," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6874–6878.
- [6] C. Liang, M. Xu, and X.-L. Zhang, "Transformer-Based End-to-End Speech Recognition with Residual Gaussian-Based Self-Attention," in *Interspeech*, 2021.
- [7] 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," *Interspeech*, 2020.
- [8] 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," *arXiv preprint arXiv:2104.01721*, 2021.
- [9] G. Zhao, X. Sun, J. Xu, Z. Zhang, and L. Luo, "Muse: Parallel multi-scale attention for sequence to sequence learning," *arXiv preprint arXiv:1911.09483*, 2019.
- [10] M. Mao, R. Zhang, H. Zheng, T. Ma, Y. Peng, E. Ding, B. Zhang, S. Han *et al.*, "Dual-stream network for visual recognition," *Advances in Neural Information Processing Systems*, vol. 34, pp. 25 346–25 358, 2021.
- [11] Z. Fan, Y. Gong, D. Liu, Z. Wei, S. Wang, J. Jiao, N. Duan, R. Zhang, and X. Huang, "Mask attention networks: Rethinking and strengthen transformer," in *North American Chapter of the Association for Computational Linguistics*, 2021.
- [12] 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," *Interspeech*, 2020.
- [13] S. Kim, A. Gholami, A. E. Shaw, N. Lee, K. Mangalam, J. Malik, M. W. Mahoney, and K. Keutzer, "Squeezeformer: An efficient transformer for automatic speech recognition," in *Advances in Neural Information Processing Systems*, 2022.
- [14] M. Radfar, R. Barnwal, R. V. Swaminathan, F.-J. Chang, G. Strimel, N. Susanj, and A. Mouchtaris, "Convrmn-t: Convolutional augmented recurrent neural network transducers for streaming speech recognition," in *Interspeech*, 2022.
- [15] K. Li, Y. Wang, G. Peng, G. Song, Y. Liu, H. Li, and Y. Qiao, "Uniformer: Unified transformer for efficient spatial-temporal representation learning," in *International Conference on Learning Representations*, 2022.
- [16] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [17] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen *et al.*, "Espnet: End-to-end speech processing toolkit," *arXiv preprint arXiv:1804.00015*, 2018.
- [18] Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Q. V. Le, and R. Salakhutdinov, "Transformer-xl: Attentive language models beyond a fixed-length context," *Association for Computational Linguistics*, pp. 2978–2988, 2019.
- [19] Z.-H. Jiang, W. Yu, D. Zhou, Y. Chen, J. Feng, and S. Yan, "Convbert: Improving bert with span-based dynamic convolution," *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 837–12 848, 2020.
- [20] Y. Peng, S. Dalmia, I. Lane, and S. Watanabe, "Branchformer: Parallel mlp-attention architectures to capture local and global context for speech recognition and understanding," in *International Conference on Machine Learning*. PMLR, 2022, pp. 17 627–17 643.
- [21] X. Li, W. Wang, X. Hu, and J. Yang, "Selective kernel networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 510–519.
- [22] 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 international coordinating committee on speech databases and speech I/O systems and assessment (O-COCOSDA)*. IEEE, 2017, pp. 1–5.
- [23] F. Hernandez, V. Nguyen, S. Ghannay, N. Tomashenko, and Y. Esteve, "Ted-lium 3: twice as much data and corpus repartition for experiments on speaker adaptation," in *International conference on speech and computer*. Springer, 2018, pp. 198–208.
- [24] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *ICASSP IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2015, pp. 5206–5210.
- [25] 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.
- [26] 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.
- [27] 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.
- [28] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *ICLR International Conference for Learning Representations*, 2014.
- [29] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [30] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.