

# INTERACTIVE FEATURE FUSION FOR END-TO-END NOISE-ROBUST SPEECH RECOGNITION

Yuchen Hu, Nana Hou, Chen Chen, Eng Siong Chng

School of Computer Science and Engineering, Nanyang Technological University, Singapore  
yuchen005@e.ntu.edu.sg

## ABSTRACT

Speech enhancement (SE) aims to suppress the additive noise from noisy speech signals to improve the speech’s perceptual quality and intelligibility. However, the over-suppression phenomenon in the enhanced speech might degrade the performance of downstream automatic speech recognition (ASR) task due to the missing latent information. To alleviate such problem, we propose an interactive feature fusion network (IFF-Net) for noise-robust speech recognition to learn complementary information from the enhanced feature and original noisy feature. Experimental results show that the proposed method achieves absolute word error rate (WER) reduction of 4.1% over the best baseline on RATS Channel-A corpus. Our further analysis indicates that the proposed IFF-Net can complement some missing information in the over-suppressed enhanced feature.

**Index Terms**— Interactive feature fusion, noise-robust speech recognition, speech enhancement, joint training approach, over-suppression phenomenon

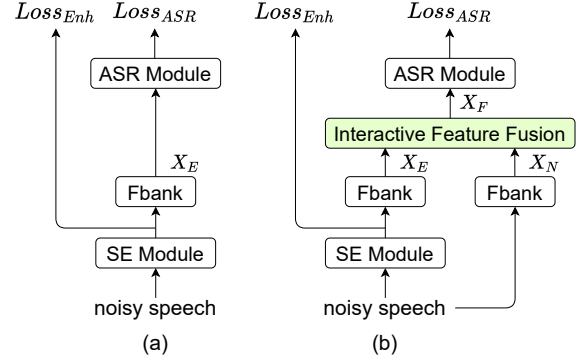
## 1. INTRODUCTION

Speech enhancement (SE) [1, 2, 3, 4, 5, 6] aims to reduce additive noise from the noisy speech to improve the speech quality for human and machine listening. It usually serves as a pre-processing module in many real-world applications, including hearing aids design [7], speaker recognition [8] and automatic speech recognition (ASR) [9].

However, recent work [10, 11, 12] observed that the enhanced speech from SE processing might not always yield good recognition accuracy for the downstream ASR task. One reason is that some important latent information in the original noisy speech are reduced by the SE processing together with the additive noise. Such over-suppression is usually undetected at the enhancement stage, but could be harmful to the downstream ASR task.

Prior work [13] proposed a cascaded framework to optimize the SE and ASR modules using ASR training objective only. Then, some studies [14, 9, 15] believed that the correlations between the SE and ASR modules can benefit each other. Therefore, they proposed a joint training approach to optimize the SE and ASR modules together via multi-task learning strategy, as shown in Figure 1(a). However, the over-suppression phenomenon still exists since the input information of ASR task only comes from the enhanced speech. Work [16] further proposed a gated recurrent fusion (GRF) method to combine enhanced and noisy speech features for ASR task. The prior studies are the source inspiration of this work.

In this paper, we propose an interactive feature fusion network (IFF-Net) for the end-to-end ASR system to improve its robustness



**Fig. 1.** Block diagrams of (a) joint training approach, (b) joint training approach with our proposed IFF-Net.

in face of noisy data. We learn a fused representation from the enhanced speech and noisy speech, which act as the input for the ASR task to complement the missing information in the enhanced speech. Specifically, the IFF-Net consists of two branches to exchange the information between enhanced and noisy features. Then, a merge module is proposed to generate a weight mask indicating which parts of enhanced feature and noisy feature could be maintained. The mask would then be used to weight sum these two features, so that we can learn clean speech information from the enhanced feature and the complementary information from the noisy feature.

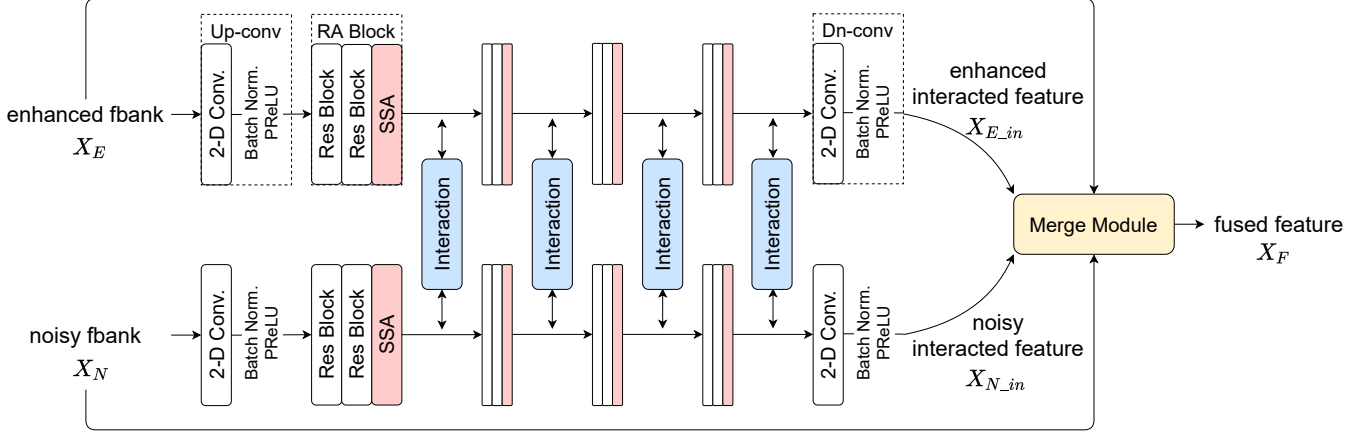
The paper is organized as follows. Section 2 introduces our proposed method. In Section 3, experimental settings and results are presented. Section 4 draws the conclusions.

## 2. PROPOSED METHOD

### 2.1. Overall Architecture

In this work, we build a joint system with speech enhancement (SE) module and automatic speech recognition (ASR) module. To further learn the complementary information from the both of enhanced feature and original noisy feature, we propose an interactive feature fusion network (IFF-Net) as shown in Figure 1(b). Specifically, we first extract the log-mel filter bank (Fbank) feature from the enhanced waveform produced by the bidirectional long short term memory (BLSTM) mask based SE module. Then the enhanced Fbank feature  $X_E$  and noisy Fbank feature  $X_N$  are fed into the proposed IFF-Net to generate a fused representation  $X_F$  as the input for the subsequent ASR module. The multi-task learning strategy is utilized in the training stage. The BLSTM-mask based SE module and the end-to-end ASR module are in same structure with the prior work [15] and [17], respectively.

This research is supported by National Research Foundation Singapore under its AI Singapore Programme (Award Number: AISG-100E-2018-006).



**Fig. 2.** Block diagram of the interactive feature fusion network (IFF-Net). “Up-conv” is the upsample operation with convolutional layers and “Dn-Conv” is the downsample operation with convolutional layers. “RA” denotes the residual attention block, “SSA” denotes the separable self-attention block, and “Interaction” denotes the interaction module.

## 2.2. Interactive Feature Fusion Network (IFF-Net)


The proposed interactive feature fusion network (IFF-Net) are in two branches (i.e., the enhanced branch and the noisy branch), which consist of two upsample convolution (Up-conv) blocks, several residual attention (RA) blocks, several interactive modules, two downsample convolution (Dn-conv) blocks and a merge module, as shown in Figure 2.

### 2.2.1. Upsample Convolution and Downsample Convolution

To learn richer and deeper information from the inputs, we firstly feed the enhanced Fbank feature  $X_E$  and the noisy Fbank feature  $X_N$  into the Up-conv block, which consists of a 2-D convolutional layer, followed by a batch normalization (BN) layer [18] and a parametric ReLU (PReLU) activation function [19]. The 2-D convolutional layer has a kernel size of (1,1) with a stride of (1,1), and the filter number is set to 64. Likewise, a corresponding Dn-conv block with same structure as Up-conv block is introduced to recover the channel dimension of the interacted features  $X_{E\_in}$  and  $X_{N\_in}$ , which keep the same with original inputs  $X_E$  and  $X_N$ . The Dn-conv block shares the same kernel size and stride with Up-conv block, and the filter number of convolutional layer is set to 1.

### 2.2.2. Residual Attention (RA) Block

To capture both local and global dependencies in the representations, we introduce the residual attention (RA) block [20] after the Up-conv block, which consists of two residual blocks and a separable self-attention (SSA) module, as shown in Figure 3. Each residual block includes two 2-D convolutional layers with a kernel size of (3,3), a stride of (1,1) and a filter number of 64 to extract deep local features from inputs. The output of two residual blocks  $X^{Res}$  is then fed into the temporal and frequency-wise self-attention blocks in parallel to capture global dependencies along both temporal and frequency dimensions. Here, we take the temporal-wise self-attention as an example to formulate the temporal and frequency-wise self-attention mechanism:



**Fig. 3.** Block diagram of the residual attention (RA) block. “Res Block” is short for the residual block.


$$\begin{aligned}
 X_t^i &= Reshape^t(X^{Res}), \quad i \in \{q, k, v\}, \\
 SA^t &= Softmax(X_t^q \cdot (X_t^k)^T / \sqrt{C \times F}) \cdot X_t^v, \quad (1) \\
 X^{Temp} &= X^{Res} + Reshape^{t,inv}(SA^t),
 \end{aligned}$$

where  $X^{Res} \in \mathbb{R}^{C \times T \times F}$ ,  $X_t^i \in \mathbb{R}^{T \times (C \times F)}$ ,  $SA^t \in \mathbb{R}^{T \times (C \times F)}$ , and  $X^{Temp} \in \mathbb{R}^{C \times T \times F}$ .  $C$  is the filter number,  $T$  is the frame number and  $F$  is the frequency-bin number.  $Reshape^t$  refers to the tensor reshape from  $\mathbb{R}^{C \times T \times F}$  to  $\mathbb{R}^{T \times (C \times F)}$  along the  $T$  dimension and  $Reshape^{t,inv}$  is the inverse operation. Likewise, there is a similar reshape operation  $Reshape^f$  and its inverse operation  $Reshape^{f,inv}$  in the frequency-wise self-attention, which are conducted along the  $F$  dimension.


Finally, the two generated deep features  $X^{Temp}$  and  $X^{Freq}$ , together with  $X^{Res}$ , are concatenated and fed into a 2-D convolutional layer to obtain the block output  $X^{RA}$ .

### 2.2.3. Interaction Module

To learn the complementary information from the deep features of the enhanced branch and the noisy branch, we propose an interaction module to exchange information between them as shown in Figure 4. The module consists of two interaction directions, noisy-to-enhanced (n2e) and enhanced-to-noisy (e2n). Taking the n2e direction for example, we first concatenate the input enhanced and noisy features  $X_E^{RA}$  and  $X_N^{RA}$  and then input them into a 2-D convolutional layer, followed by a BN layer and a Sigmoid activate function to generate a multiplicative mask  $M_N$ , which predicts the suppressed and maintained information of  $X_N^{RA}$ . Then a residual feature



**Fig. 4.** Block diagram of the interaction module. Here we take the n2e interaction direction as an example for illustration.  $\otimes$  denotes element-wise multiplication, and  $\oplus$  is residual connection.



**Fig. 5.** Block diagram of the merge module.  $\otimes$  denotes element-wise multiplication, and  $\oplus$  is residual connection.

$R_{N2E}$  is generated by multiplying  $M_N$  and  $X_N^{RA}$  elementally. Finally, the module adds  $X_E^{RA}$  and  $R_{N2E}$  to get a “filtered” version of the enhanced feature,  $X_E^{IM}$ . Likewise, we conduct the same process for the e2n interaction direction, where only the inputs “Feature-1” and “Feature-2” are exchanged.

#### 2.2.4. Merge Module

To fuse the two interacted features  $X_{E,in}$  and  $X_{N,in}$ , we propose a merge module as shown in Figure 5. Such interacted features  $X_{E,in}$  and  $X_{N,in}$  are stacked with the input Fbank features  $X_E$  and  $X_N$  and fed into the merge module. The merge module consists of a 2-D convolutional layer, a temporal self-attention block from SSA to capture global temporal dependencies and another convolutional layer to learn an element-wise mask  $M$ . It acts as a gate to control whether the information of interacted features is kept or discarded. The kernel size of both convolutional layers is (3, 3) and the filter number is set to 4 and 1, respectively. Finally, the output fused feature  $X_F$  is generated by:

$$X_F = X_{E,in} * M + X_{N,in} * (1 - M). \quad (2)$$

### 3. EXPERIMENTS AND RESULTS

#### 3.1. Dataset

We conduct experiments on Robust Automatic Transcription of Speech (RATS) dataset. As RATS corpus is a chargeable dataset under LDC, we only release its Fbank features and several listening samples on Github for reference<sup>1</sup>. RATS corpus consists of eight channels, and in this work we only focus on the Channel-A data which belongs to the ultra high frequency (UHF) data category. The Channel-A data in RATS corpus includes 44.3 hours training data, 4.9 hours valid data and 8.2 hours test data. One can refer to [21] for more information.

<sup>1</sup><https://github.com/YUCHEN005/RATS-Channel-A-Speech-Data>

#### 3.2. Experimental Setup

##### 3.2.1. Network Configurations

The joint SE-ASR system with the proposed IFF-Net consists of the three modules: the SE module, the proposed IFF-Net and the ASR module. The SE module aims to predict the mask for the noisy magnitude feature by 3 layer of 896-unit bidirectional long short-term memory (BLSTM) [22] and a 257-unit linear layer followed by the ReLU [23] activation function. We adopt the state-of-the-art Conformer [17] for downstream ASR task, where the encoder consists of 12 Conformer layers, and the decoder has 6 transformer [24] layers, both with 256 hidden units. We use 994 byte-pair-encoding (BPE) [25] tokens as ASR output.

During the training stage, the noisy waveform is cut into some segments with a duration of several seconds each for batch training. The network is optimized by the Adam algorithm [26], where the learning rate first warms up linearly to 0.002 in 25,000 steps and then decreases proportional to the inverse square root of the step number. The weight for enhancement loss in multi-task learning is set to 0.3 and the batch size is set to 64. For fair comparison, the training epoch is set to 50 for all experiments.

##### 3.2.2. Reference Baselines

To evaluate the effectiveness of the proposed system, we build four competitive baselines from prior work for comparison. For fair evaluation, we adopt the same architecture for all SE modules and ASR modules included, respectively.

1. **E2E ASR System** [17]: an End-to-End ASR system based on Conformer. It has achieved the state-of-the-art performance on ASR, but may not work well for the noise-robust speech recognition task.
2. **Cascaded SE and ASR System** [13]: a cascaded system consisting of a front-end BLSTM-based SE module and a back-end Conformer-based ASR module. The system is optimized using ASR training objective only.
3. **Joint Training Approach** [15]: a same structure as previous cascaded system, but adopt the multi-task learning strategy for joint training. It builds correlations between SE and ASR, thus leads to better recognition results.
4. **GRF Network** [16]: a gated recurrent fusion (GRF) method to combine enhanced speech and noisy speech for subsequent ASR training. Based on the joint training approach, it further improves the performance of noise-robust ASR.

**Table 1.** WER% results in an ablation study of the proposed IFF-Net. “# blocks” denotes number of RA blocks in each branch of IFF-Net, “# filters” denotes filter number of the convolutional layer in the RA blocks, and “# params.” denotes number of parameters of IFF-Net. Different configurations have been explored to maximize our GPU memory usage.

Method	# blocks	# filters	# params.(M)	WER(%)
IFF-Net	2	32	0.19	49.1
	2	64	0.74	47.9
	4	32	0.37	47.7
	4	64	1.49	<b>46.2</b>

**Table 2.** WER% results in a comparative study of the proposed IFF-Net. “IFF-Net w/o noisy branch” means without the noisy branch, interaction modules and the merge module.

Method	WER(%)
IFF-Net w/o noisy branch	49.5
IFF-Net w/o SSA	49.2
IFF-Net w/o both Interaction Modules	48.0
IFF-Net w/o n2e Interaction Module	47.8
IFF-Net w/o e2n Interaction Module	47.4
IFF-Net	<b>46.2</b>

**Table 3.** WER% results of the proposed IFF-Net and competitive baselines.

Method	WER(%)
E2E ASR System [17]	54.3
Cascaded SE and ASR System [13]	53.1
Joint Training Approach [15]	51.8
GRF Network [16]	50.3
IFF-Net	<b>46.2</b>

### 3.3. Results

#### 3.3.1. Effect of the number of RA blocks and filters in IFF-Net


We first analyse and summarize the performance with different number of RA blocks and filters in the proposed IFF-Net. As shown in Table 1, the second and third column denote the number of RA blocks in each branch of IFF-Net and the filter number of the convolutional layer in the RA blocks. We observe that the performance improves as the number of RA blocks and filters increase. We obtain the best performance with 4 RA blocks and 64 filters in convolutional layers, only at cost of 1.49M parameters. We adopt this setting for the proposed IFF-Net hereafter.

#### 3.3.2. Effect of each component in IFF-Net


We further report the effect of the each component including the noisy branch, SSA module and interactive module in Table 2. We observe that the proposed system degrades 3.3% absolute WER without the entire noisy branch. Similar phenomenon can be observed when we remove the SSA modules. Then we find that the interactive module could contribute 1.8% of WER improvement and the two interaction directions of n2e/e2n (noisy-to-enhanced or enhanced-to-noisy) are both important for the proposed IFF-Net.

#### 3.3.3. IFF-Net vs. Other Competitive Methods

Table 3 summarizes the comparison between the proposed IFF-Net and other competitive methods. Specifically, E2E ASR System denotes the best result of the Conformer-based ASR system, which indicates the difficulty of noise-robust ASR. Cascaded SE and ASR System could slightly improve the performance with help of the enhanced speech. Joint Training Approach achieves 2.5% absolute improvement over E2E ASR System by jointly optimizing the SE module and ASR module. GRF Network can further lower the WER by combining the information of enhanced speech and noisy speech. We observe that the proposed IFF-Net obtains the best result with 4.1% absolute WER reduction over the best baseline.



**Fig. 6.** Spectrums of (a) clean fbank, (b) noisy fbank; and ASR input of (c) Cascaded SE and ASR System, (d) Joint Training Approach, (e) GRF Network, (f) IFF-Net. The colorbar is for all the spectrograms.



**Fig. 7.** Spectrums of (a) enhanced fbank, (b) enhanced interacted feature, (c) noisy interacted feature, (d) weight mask, (e) fused feature in IFF-Net. The colorbar is for the weight mask.

#### 3.3.4. Visualization of Spectrums

To further show the contribution of the proposed IFF-Net approach, we illustrate the spectrums of the ASR input of different methods using an example (fe\_03\_1781-02235-B-024485-024688-A.wav) from the RATS testset as shown in Figure 6. First, from (a) and (b) we can observe a lot of noise in the noisy fbank feature. Then comparing with the methods (c-e), we observe that the proposed IFF-Net approach could significantly reduce more background noise and meanwhile maintain richer clean information.

To analyse the reasons why the proposed IFF-Net could yield better performance, we further illustrate the spectrums of enhanced fbank, two interacted features and the fused feature in IFF-Net, as shown in Figure 7. We observe from (a) that some speech content in enhanced spectrum are over-suppressed and distorted. The enhanced interacted feature (b) could recover some information with help from the noisy branch. Meanwhile, the noisy interacted feature (c) reduces noise and extracts some coarse-grained structure information. Then, to further address the remaining distortions in (b), the merge module generates a weight mask (d) where the light spots indicate speech content and the dark spots indicate noise distortions. Therefore, by weighting sum the interacted features (b) and (c) with this mask, we can finally obtain the output fused feature (e), which contains richer speech information with less distortions, thus benefit ASR and yield better performance.

## 4. CONCLUSION

In this paper, we propose a joint SE and ASR system with an interaction feature fusion network (IFF-Net) for noise-robust speech recognition. Specifically, we exchange and fuse the enhanced and noisy speech features to complement the missing information for the downstream ASR task. Experimental results on RATS Channel-A corpus show that our proposed IFF-Net is more effective for noise-robust ASR than other competitive approaches.

## 5. REFERENCES

- [1] Yuxuan Wang, Arun Narayanan, and DeLiang Wang, “On training targets for supervised speech separation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 12, pp. 1849–1858, 2014.
- [2] Santiago Pascual, Antonio Bonafonte, and Joan Serra, “Segan: Speech enhancement generative adversarial network,” *arXiv preprint arXiv:1703.09452*, 2017.
- [3] Nana Hou, Chenglin Xu, Eng Siong Chng, and Haizhou Li, “Domain adversarial training for speech enhancement,” in *2019 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2019, pp. 667–672.
- [4] Nana Hou, Chenglin Xu, Joey Tianyi Zhou, Eng Siong Chng, and Haizhou Li, “Multi-task learning for end-to-end noise-robust bandwidth extension,” in *INTERSPEECH*, 2020, pp. 4069–4073.
- [5] Nana Hou, Chenglin Xu, Joey Tianyi Zhou Van Tung Pham, Joey Tianyi Zhou, Eng Siong Chng, and Haizhou Li, “Speaker and phoneme-aware speech bandwidth extension with residual dual-path network,” in *INTERSPEECH*, 2020, pp. 4064–4068.
- [6] Nana Hou, Chenglin Xu, Eng Siong Chng, and Haizhou Li, “Learning disentangled feature representations for speech enhancement via adversarial training,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 666–670.
- [7] Igor Fedorov, Marko Stamenovic, Carl Jensen, Li-Chia Yang, Ari Mandell, Yiming Gan, Matthew Mattina, and Paul N Whatmough, “Tynlstms: Efficient neural speech enhancement for hearing aids,” *arXiv preprint arXiv:2005.11138*, 2020.
- [8] J. Ortega-Garcia and J. Gonzalez-Rodriguez, “Overview of speech enhancement techniques for automatic speaker recognition,” in *Proceeding of Fourth International Conference on Spoken Language Processing, ICSLP '96*. IEEE, 1996.
- [9] Ashutosh Pandey, Chunxi Liu, Yun Wang, and Yatharth Saraf, “Dual application of speech enhancement for automatic speech recognition,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021.
- [10] I. Mporas, T. Ganchev, O. Kocsis, and N. Fakotakis, “Speech enhancement for robust speech recognition in motorcycle environment,” *International Journal on Artificial Intelligence Tools*, vol. 19, pp. 159–173, 2010.
- [11] Philipos C. Loizou and Gibak Kim, “Reasons why current speech-enhancement algorithms do not improve speech intelligibility and suggested solutions,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 1, pp. 47–56, 2011.
- [12] Peidong Wang, Ke Tan, and De-Liang Wang, “Bridging the gap between monaural speech enhancement and recognition with distortion-independent acoustic modeling,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 39–48, 2020.
- [13] Aswin Shanmugam Subramanian, Xiaofei Wang, Murali Karthick Baskar, Shinji Watanabe, Toru Taniguchi, Dung Tran, and Yuya Fujita, “Speech enhancement using end-to-end speech recognition objectives,” in *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2019, pp. 234–238.
- [14] Bin Liu, Shuai Nie, Shan Liang, Wenju Liu, Meng Yu, Lianwu Chen, Shouye Peng, and Changliang Li, “Jointly adversarial enhancement training for robust end-to-end speech recognition,” in *Interspeech*, 2019, pp. 491–495.
- [15] Duo Ma, Nana Hou, Van Tung Pham, Haihua Xu, and Eng Siong Chng, “Multitask-based joint learning approach to robust asr for radio communication speech,” *arXiv preprint arXiv:2107.10701*, 2021.
- [16] Cunhang Fan, Jiangyan Yi, Jianhua Tao, Zhengkun Tian, Bin Liu, and Zhengqi Wen, “Gated recurrent fusion with joint training framework for robust end-to-end speech recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 198–209, 2021.
- [17] Anmol Gulati, James Qin, Chiu Chung-Cheng, 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 *Interspeech*, 2020, pp. 5036–5040.
- [18] Sergey Ioffe and Christian Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *Proceedings of the 32nd International Conference on Machine Learning*. 2015, vol. 37, pp. 448–456, PMLR.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [20] Chengyu Zheng, Xiulian Peng, Yeliang Zhang, Sriram Srinivasan, and Y. Lu, “Interactive speech and noise modeling for speech enhancement,” in *AAAI*, 2021.
- [21] David Graff, Kevin Walker, Stephanie M. Strassel, Xiaoyi Ma, Karen Jones, and Ann Sawyer, “The rats collection: Supporting hlt research with degraded audio data,” in *Proceedings of the Ninth International Conference on Language Resources and Evaluation*. 2014, pp. 1970–1977, European Language Resources Association (ELRA).
- [22] Sepp Hochreiter and Jürgen Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [23] Xavier Glorot, Antoine Bordes, and Yoshua Bengio, “Deep sparse rectifier neural networks,” in *Proceedings of the fourteenth international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings*, 2011, pp. 315–323.
- [24] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [25] Taku Kudo and John 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*. 2018, pp. 66–71, Association for Computational Linguistics.
- [26] Diederik P. Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.