



# ContextNet: Improving Convolutional Neural Networks for Automatic Speech Recognition with Global Context

Wei Han\*, Zhengdong Zhang\*, Yu Zhang, Jiahui Yu, Chung-Cheng Chiu, James Qin, Anmol Gulati, Ruoming Pang, Yonghui Wu

Google Inc.

{weihan, zhangzd}@google.com

## Abstract

Convolutional neural networks (CNN) have shown promising results for end-to-end speech recognition, albeit still behind RNN/transformer based models in performance. In this paper, we study how to bridge this gap and go beyond with a novel CNN-RNN-transducer architecture, which we call ContextNet. ContextNet features a fully convolutional encoder that incorporates global context information into convolution layers by adding squeeze-and-excitation modules. In addition, we propose a simple scaling method that scales the widths of ContextNet that achieves good trade-off between computation and accuracy.

We demonstrate that on the widely used Librispeech benchmark, ContextNet achieves a word error rate (WER) of 2.1%/4.6% without external language model (LM), 1.9%/4.1% with LM and 2.9%/7.0% with only 10M parameters on the clean/noisy LibriSpeech test sets. This compares to the best previously published model of 2.0%/4.6% with LM and 3.9%/11.3% with 20M parameters. The superiority of the proposed ContextNet model is also verified on a much larger internal dataset.

**Index Terms:** speech recognition, convolutional neural networks

## 1. Introduction

Convolution Neural Network (CNN) based models for end-to-end (E2E) speech recognition is attracting an increasing amount of attention [1, 2, 3, 4]. Among them, the Jasper model [3] recently achieves close to the state-of-the-art word error rate (WER) 2.95% on LibriSpeech test-clean [5] with an external neural language model. The main feature of the Jasper model is a deep convolution based encoder with stacked layers of 1D convolutions and skip connections. Depthwise separable convolutions [6] have been utilized to further increase the speed and accuracy of CNN models [7, 4]. The key advantage of a CNN based model is its parameter efficiency; however, the WER achieved by the best CNN model, QuartzNet [4], is still behind the RNN/transformer based models [8, 9, 10, 11].

A major difference between the RNN/Transformer [9, 10, 11] based models and a CNN model is the length of the context. In a bidirectional RNN model, a cell in theory has access to the information of the whole sequence; in a Transformer model, the attention mechanism explicitly allows the nodes at two distant time stamps to attend each other. However, a naive convolution with a limited kernel size only covers a small window in the time domain; hence the context is small and the global information is not incorporated. In this paper, we argue that the lack of global context is the main cause of the gap of WER between

the CNN based ASR model and the RNN/Transformer based models.

To enhance the global context in the CNN model, we draw inspirations from the squeeze-and-excitation (SE) layer introduced in [12], and propose a novel CNN model for ASR, which we call ContextNet. An SE layer squeezes a sequence of local feature vectors into a single global context vector, broadcasts this context back to each local feature vector, and merges the two via multiplications. When we place an SE layer after a naive convolution layer, we grant the convolution output the access to global information. Empirically, we observe that adding squeeze-and-excitation layers to ContextNet introduces the most reduction in the WER on LibriSpeech test-other.

Previous works on hybrid ASR have successfully introduced the context to acoustic models by either stacking a large number of layers, or having a separately trained global vector to represent the speaker and the environment information [13, 14, 15, 16]. In [17], SE has been adopted to RNN for unsupervised adaptation. In this paper, we show that SE can also be effective for CNN encoders.

The architecture of ContextNet is also inspired by the design choices of QuartzNet [4], such as the usage of depthwise separable 1D convolution in the encoder. However, there are some key differences in the architectures in addition to the incorporation of the SE layer. For instance, we use a RNN-T decoder [18, 19, 20, 21] instead of the CTC decoder [22]. Moreover, we use the Swish activation function [23], which contributes a slight but consistent reduction in WER. Overall, ContextNet achieves the WER of 1.9%/4.1% on LibriSpeech test-clean/test-other. This is a big improvement over previous CNN based architectures such as QuartzNet[4], and it outperforms transformer and LSTM based models [11, 10, 24, 9, 25].

This paper also studies how to reduce the computation cost of ContextNet for faster training and inference. First, we adopt a progressive downsampling scheme that is commonly used in vision models. Specifically, we progressively reduce the length of the encoded sequence eight times, significantly lower the computation while maintaining the encoder's representation power and the overall model accuracy. As a benefit, this downsampling scheme allows us to reduce the kernel size of all the convolution layers to five without significantly reducing the effective receptive field of an encoder output node.

We can scale ContextNet by globally changing the number of channels in convolutional filters. Figure 1 illustrates the trade-off of ContextNet between model size and WER, as well as its comparison against other methods. Clearly, our scaled model achieves the best trade-offs among all.

In summary, the main contributions of this paper are: (1) an improved CNN architecture with global context for ASR, (2) a progressive downsampling and model scaling scheme to

\*Equal contribution.

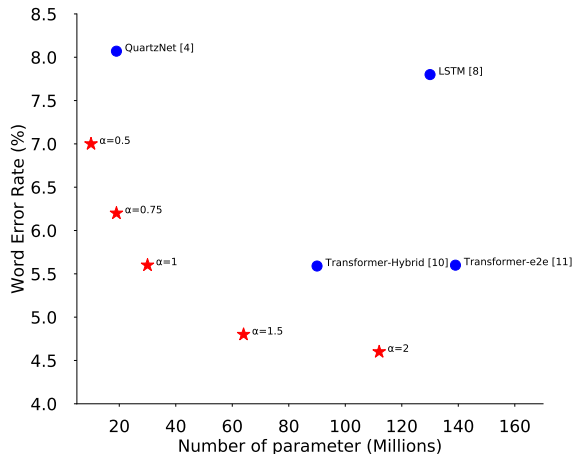


Figure 1: LibriSpeech test-other WER vs. model size. All numbers for E2E models are without external LM. For the Transformer-Hybrid model we use the encoder size. ContextNet numbers are highlighted in red, and  $\alpha$  is the model scaling parameter discussed in Section 2.2.6. Detailed results are in Table 2.

achieve superior accuracy and model size trade-off.

## 2. Model

This section introduces the architecture details of ContextNet. Section 2.1 discusses the high-level design of ContextNet. Then Section 2.2 introduces our convolutional encoder, and discusses how we progressively reduce the temporal length of the input utterance in the network to reduce the computation while maintaining the accuracy of the model.

### 2.1. End-to-end Network: CNN-RNN-Transducer

Our network is based on the RNN-Transducer framework [18, 19, 20]. The network contains three components: audio encoder on the input utterance, label encoder on the input label, and a joint network to combine the two and decode. We directly use the LSTM based label encoder and the joint network from [20], but propose a new CNN based audio encoder.

### 2.2. Encoder Design

Let the input sequence be  $\mathbf{x} = (x_1, \dots, x_T)$ . The encoder transforms the original signal  $\mathbf{x}$  into a high level representation  $\mathbf{h} = (h_1, \dots, h_{T'})$ , where  $T' \leq T$ . Our convolution based AudioEncoder( $\cdot$ ) is defined as:

$$\mathbf{h} = \text{AudioEncoder}(\mathbf{x}) = C_K(C_{K-1}(\dots C_1(\mathbf{x}))), \quad (1)$$

where each  $C_k(\cdot)$  defines a convolution block. It contains a few layers of convolutions, each followed by batch normalization [26] and an activation function. It also includes the squeeze-and-excitation component [12] and skip connections [27].

Before presenting the details of  $C(\cdot)$ , we first elaborate the important modules in  $C(\cdot)$ .

#### 2.2.1. Squeeze-and-excitation

As illustrated in Figure 2, the Squeeze-and-excitation [12] function,  $\text{SE}(\cdot)$ , performs global average pooling on the input  $x$ , transforms it into a global channelwise weight  $\theta(x)$  and

element-wise multiplies each frame by this weight. We adopt the idea to the 1D case,

$$\bar{x} = \frac{1}{T} \sum_t x_t, \theta(x) = \text{Sigmoid}(W_2(\text{Act}(W_1\bar{x} + b_1)) + b_2),$$

$$\text{SE}(x) = \theta(x) \circ x,$$

where  $\circ$  represents element-wise multiplication,  $W_1, W_2$  are weight matrices, and  $b_1, b_2$  are bias vectors.

#### 2.2.2. Depthwise separable convolution

Let  $\text{conv}(\cdot)$  represent the convolution function used in the encoder. In this paper, we choose depthwise separable convolution as  $\text{conv}(\cdot)$ , because such a design has been previously shown in various applications [6, 4, 28] to achieve better parameter efficiency without impacting accuracy.

For simplicity, we use the same kernel size on all depthwise convolution layers in the network.

#### 2.2.3. Swish activation function

Let  $\text{Act}(\cdot)$  represent the activation function in the encoder. To choose  $\text{Act}(\cdot)$ , we've experimented with both ReLU and the swish function [23] defined as:

$$\text{Act}(x) = x \cdot \sigma(\beta x) = \frac{x}{1 + \exp(-\beta x)}, \quad (2)$$

where  $\beta = 1$  for all our experiments. We've observed that the swish function works consistently better than ReLU.

#### 2.2.4. Convolution block

With all the individual modules introduced, we now present the convolution block  $C(\cdot)$  from Equation (1). Figure 3 illustrate a high-level architecture of  $C(\cdot)$ . A block  $C(\cdot)$  can contain a few  $\text{Conv}(\cdot)$  functions; let  $m$  be the number of  $\text{Conv}(\cdot)$  functions. Let  $\text{BN}(\cdot)$  be the batch normalization [29]. We define each layer as  $f(x) = \text{Act}(\text{BN}(\text{Conv}(x)))$ . Therefore,

$$C(x) = \text{Act}(\text{SE}(f^m(x)) + P(x))$$

where  $f^m$  means stacking  $m$  layers of the function  $f(\cdot)$  on the input and  $P(\cdot)$  represents a pointwise projection function on the residual. By a slight abuse of notation, we allow the first layer and the last layer to be different from the other  $m - 2$  layers: if the block needs to downsample the input sequence by two times, the last layer has a stride of two while all the rest  $m - 1$  layers has a stride of one; otherwise all  $m$  layers have a stride of one. Additionally, if the block has an input number of channels  $D_{in}$  and output number of channels  $D_{out}$ , the first layer  $f'$  turns  $D_{in}$  channels into  $D_{out}$  channels while the rest  $m - 1$  layers maintain the number of channels as  $D_{out}$ . Following the convention, the projection function  $P$  has the same number of stride as the first layer.

#### 2.2.5. Progressive downsampling

We use strided convolution for temporal downsampling. More downsampling layers reduces computation cost, but excessive downsampling in the encoder may negatively impact the decoder. Empirically, we find that a progressive  $8 \times$  downsampling scheme achieves a good trade-off between speed and accuracy. These trade-offs are discussed in Section 3.3.

#### 2.2.6. Configuration details of ContextNet

ContextNet has 23 convolution blocks  $C_0, \dots, C_{22}$ . All convolution blocks have five layers of convolution, except  $C_0$  and

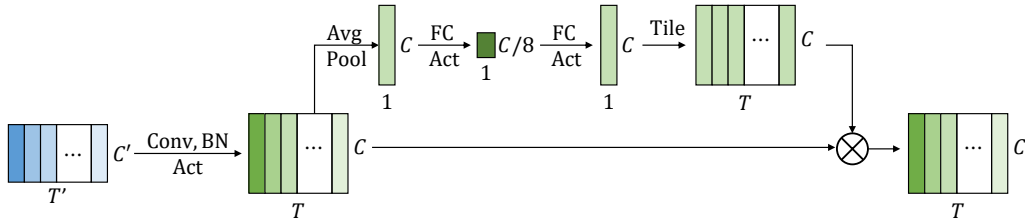


Figure 2: *1D Squeeze-and-excitation module. The input first goes through a convolution layer followed by batch normalization and activation. Then average pooling is applied to condense the conv result into a 1D vector, which is then processed by a bottleneck structure formed by two fully connected (FC) layers with activation functions. The output goes through a Sigmoid function to be mapped to (0, 1), and then tiled and applied on the conv output using pointwise multiplications.*

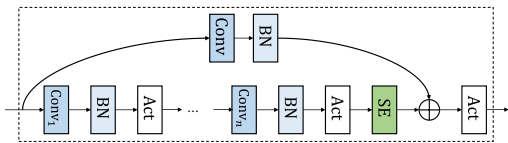


Figure 3: *A convolution block  $C_i$  contains a number of convolutions, each followed by batch normalization and activation. A squeeze-and-excitation (SE) block operates on the output of the last convolution layer. A skip connection with projection is applied on the output of the squeeze-and-excitation block.*

Table 1: *Configuration of the ContextNet encoder.  $\alpha$  controls the number of output channels, and thus the scaling of our model. The kernel size is for the window size in the temporal domain; the convolutions are across frequency. If the stride of a convolution block is 2, its last conv layer has a stride of two while the rest of the conv layers has a stride of one, as discussed in Section 2.2.*

Block ID	#Conv layers	#Output channels	Kernel size	Other
$C_0$	1	$256 \times \alpha$	5	No residual
$C_1$ - $C_2$	5	$256 \times \alpha$	5	stride is 2
$C_3$	5	$256 \times \alpha$	5	
$C_4$ - $C_6$	5	$256 \times \alpha$	5	stride is 2
$C_7$	5	$256 \times \alpha$	5	
$C_8$ - $C_{10}$	5	$256 \times \alpha$	5	stride is 2
$C_{11}$ - $C_{13}$	5	$512 \times \alpha$	5	
$C_{14}$	5	$512 \times \alpha$	5	
$C_{15}$ - $C_{21}$	5	$512 \times \alpha$	5	No residual
$C_{22}$	1	$640 \times \alpha$	5	

$C_{22}$ , which only have one layer of convolution each. Table 1 summarizes the architecture details. Note that a global parameter  $\alpha$  controls the scaling of our model. Increasing  $\alpha$  when  $\alpha > 1$  increases the number of channels of the convolutions, giving the model more representation power with a larger model size.

### 3. Experiments

We conduct experiments on the Librispeech [5] dataset which consists of 970 hours of labeled speech and an additional text only corpus for building language model. We extract 80 dimensional filterbanks features using a 25ms window with a stride of

10ms.

We use the Adam optimizer [30] and a transformer learning rate schedule [31] with 15k warm-up steps and a peak learning rate of 0.0025. An  $\ell_2$  regularization with  $10^{-6}$  weight is also added to all the trainable weights in the network. We use a single layer LSTM as decoder with input dimension of 640. Variational noise is introduced to the decoder as a regularization.

We use SpecAugment [25, 8] with mask parameter ( $F = 27$ ), and ten time masks with maximum time-mask ratio ( $p_S = 0.05$ ), where the maximum size of the time mask is set to  $p_S$  times the length of the utterance. Time warping is not used. We use a 3-layer LSTM LM with width 4096 trained on the LibriSpeech language model corpus with the LibriSpeech960h transcripts added, tokenized with the 1k WPM built from LibriSpeech 960h. The LM has word-level perplexity 63.9 on the dev-set transcripts. The LM weight  $\lambda$  for shallow fusion is tuned on the dev-set via grid search. All models are implemented with Lingvo toolkit [32].

#### 3.1. Results on LibriSpeech

We evaluate three different configurations of ContextNet on LibriSpeech. The models are all based on Table 1, but differ in the network width,  $\alpha$ ; hence, they differ in model size. Specifically, we choose  $\alpha$  in  $\{0.5, 1, 2\}$  for the small, medium and large ContextNet. We also build our own LSTM baseline as a reference.

Table 2 summarizes the evaluation results as well as the comparisons with a few previously published systems. The results suggest improvements of ContextNet over previously published systems. Our medium model, ContextNet(M), only has 31M parameters and achieves similar WER compared with much larger systems [11, 33]. The large model, ContextNet(L), outperforms the previous SOTA by 13% relatively on test-clean and 18% relatively on test-other. Our scaled-down model, ContextNet(S), also shows an improvement to previous systems of similar size [4], with or without a language model.

#### 3.2. Effect of Context Size

To validate the effectiveness of adding global context to the CNN model for ASR, we perform an ablation study on how the squeeze-and-excitation module affects the WER on LibriSpeech test-clean/test-other. ContextNet in Table 1 with all squeeze-and-excitation modules removed and  $\alpha = 1.25$  serves as the baseline of zero context.

The vanilla squeeze-and-excitation module uses the whole utterance as context. To investigate the effect of different context sizes, we replace the global average pooling operator of the

Table 2: WER on Librispeech. Compared to previous models, ContextNet achieves superior performance both with and without language models

Method	#Params (M)	Without LM		With LM	
		testclean	testother	testclean	testother
<b>Hybrid</b>					
Transformer [10]	-	-	-	2.26	4.85
<b>CTC</b>					
QuartzNet (CNN) [4]	19	3.90	11.28	2.69	7.25
<b>LAS</b>					
Transformer [9]	-	-	-	2.6	5.7
Transformer [33]	270	2.89	6.98	2.33	5.17
LSTM	360	2.6	6.0	2.2	5.2
<b>Transducer</b>					
Transformer [11]	139	2.4	5.6	2.0	4.6
<b>This Work</b>					
ContextNet(S)	10.8	<b>2.9</b>	<b>7.0</b>	<b>2.3</b>	<b>5.5</b>
ContextNet(M)	31.4	<b>2.4</b>	<b>5.4</b>	<b>2.0</b>	<b>4.5</b>
ContextNet(L)	112.7	<b>2.1</b>	<b>4.6</b>	<b>1.9</b>	<b>4.1</b>

squeeze-and-excitation module by a stride-one pooling operator where the context can be controlled by the size of the pooling window. In this study, we compare the window size of 256, 512 and 1024 on all convolutional blocks.

As illustrated in Table 3, the SE module provides major improvement over the baseline. In addition, the benefit becomes greater as the length of the context window increases. This is consistent with the observation in a similar study of SE on image classification models [34].

Table 3: Effect of the context window size on WER. All models have  $\alpha = 1.25$ .

Context	dev clean	dev other	test clean	test other
None	2.6	7.0	2.6	6.9
256	2.1	5.4	2.3	5.5
512	2.1	5.1	2.3	5.2
1024	2.1	5.0	2.3	5.1
global	<b>2.0</b>	<b>4.9</b>	<b>2.3</b>	<b>4.9</b>

### 3.3. Depth, Width, Kernel Size and Downsampling

*Depth:* We perform a sweeping on the number of convolutional blocks and our best configuration is in Table 1. We find that with this configuration, we can train a model in a day with stable convergence.

*Width:* We globally scale the width of the network (i.e., the number of channels) on all encoder layers and study how it impacts the model performance. Specifically, we take the ContextNet model from Table 1, sweep  $\alpha$ , and report the model size

Table 4: The effect of temporal reduction and convolution kernel size on FLOPS and model accuracy.

Reduction	Kernel size	GFLOPS <sup>1</sup>	testclean	testother
2x	3	2.131	2.7	6.3
	5	2.137	2.6	5.8
	11	2.156	2.4	5.4
	23	2.194	2.3	5.0
8x	3	1.036	2.3	5.1
	5	1.040	2.3	5.0
	11	1.050	2.3	5.0
	23	1.071	2.3	5.2

Table 5: Effect of Model scaling by network width on WER.

$\alpha$	#Params(M)	dev clean	dev other	test clean	test other
0.5	10.8	2.7	7.0	2.9	7.0
1	31.4	2.2	5.1	2.4	5.4
1.5	65.4	2.0	4.7	2.2	4.8
2	112.7	2.0	4.6	2.1	4.6

and the WER on LibriSpeech. Table 5 summarizes the result; it demonstrates the good trade-off between model size and WER of ContextNet.

*Downsampling and kernel size:* Table 4 summarizes the FLOPS and WER on LibriSpeech with various choices of downsampling and filter size. We use the same model with only one downsampling layer added to  $C_3$  as the baseline; hence the baseline only does  $2\times$  temporal reduction. We sweep the kernel size in  $\{3, 5, 11, 21\}$ , each kernel size is applied to all the depthwise convolution layers. The results suggest that progressive downsampling introduces significant saving in the number of FLOPS. Moreover, it actually benefits the accuracy of the model slightly. In addition, with progressive downsampling, increasing the kernel size decreases the WER of the model.

### 3.4. Large Scale Experiments

Finally, we show that the proposed architecture is also effective on large scale datasets. We use an experiment setup similar to [35], where the training set has public Youtube videos with semi-supervised transcripts generated by the approach in [36]. We evaluate on 117 videos with a total duration of 24.12 hours. This test set has diverse and challenging acoustic environments<sup>2</sup>. Table 6 summarizes the result. We can see that ContextNet outperforms the previous best architecture from [35], which is a combination of convolution and bidirectional LSTM, by 12% relatively with fewer parameters and FLOPS.

Table 6: Comparing ContextNet with previous best results on Youtube test sets

Model	#Params (M)	GFLOPS	Youtube WER
TDNN [35]	192	3.834	9.3
ContextNet	112	2.647	8.2

## 4. Conclusion

In this work, we proposed and evaluated a CNN based architecture for end-to-end speech recognition. A couple of modeling choices are discussed and compared. This model achieves a better accuracy on the LibriSpeech benchmark with much fewer parameters compared to previously published CNN models. The proposed architecture can easily be used to search for small ASR models by limiting the width of the network. Initial study on a much larger and more challenging dataset also confirms our findings.

<sup>1</sup>We report the average encoder FLOPS for processing one second of audio.

<sup>2</sup>Reproduced results. The train and eval set has been changed recently so the numbers in Table 6 are different from reported in [35].

## 5. References

- [1] Y. Zhang, W. Chan, and N. Jaitly, "Very deep convolutional networks for end-to-end speech recognition," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 4845–4849.
- [2] N. Zeghidour, Q. Xu, V. Liptchinsky, N. Usunier, G. Synnaeve, and R. Collobert, "Fully convolutional speech recognition," *arXiv preprint arXiv:1812.06864*, 2018.
- [3] 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 preprint arXiv:1904.03288*, 2019.
- [4] 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," *arXiv preprint arXiv:1910.10261*, 2019.
- [5] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [6] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1251–1258.
- [7] A. Hannun, A. Lee, Q. Xu, and R. Collobert, "Sequence-to-sequence speech recognition with time-depth separable convolutions," *arXiv preprint arXiv:1904.02619*, 2019.
- [8] D. S. Park, Y. Zhang, C.-C. Chiu, Y. Chen, B. Li, W. Chan, Q. V. Le, and Y. Wu, "SpecAugment on large scale datasets," in *ICASSP*, 2020.
- [9] S. Karita, N. Chen, T. Hayashi, T. Hori, H. Inaguma, Z. Jiang, M. Someki, N. E. Y. Soplin, R. Yamamoto, X. Wang *et al.*, "A comparative study on transformer vs rnn in speech applications," *arXiv preprint arXiv:1909.06317*, 2019.
- [10] 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," *arXiv preprint arXiv:1910.09799*, 2019.
- [11] Q. Zhang, H. Lu, H. Sak, A. Tripathi, E. McDermott, S. Koo, and S. Kumar, "Transformer transducer: A streamable speech recognition model with transformer encoders and rnn-t loss," *arXiv preprint arXiv:2002.02562*, 2020.
- [12] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [13] V. Peddinti, G. Chen, V. Manohar, T. Ko, D. Povey, and S. Khudanpur, "Jhu aspire system: Robust lvcsr with tdnns, ivector adaptation and rnn-lms," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE, 2015, pp. 539–546.
- [14] S. Xue, O. Abdel-Hamid, H. Jiang, L. Dai, and Q. Liu, "Fast adaptation of deep neural network based on discriminant codes for speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 12, pp. 1713–1725, 2014.
- [15] M. Karafiát, L. Burget, P. Matějka, O. Glembek, and J. Černocký, "ivector-based discriminative adaptation for automatic speech recognition," in *2011 IEEE Workshop on Automatic Speech Recognition & Understanding*. IEEE, 2011, pp. 152–157.
- [16] G. Saon, H. Soltau, D. Nahamoo, and M. Picheny, "Speaker adaptation of neural network acoustic models using i-vectors," in *2013 IEEE Workshop on Automatic Speech Recognition and Understanding*. IEEE, 2013, pp. 55–59.
- [17] H. B. Sailor, S. Deena, M. A. Jalal, R. Lileikyte, and T. Hain, "Un-supervised adaptation of acoustic models for asr using utterance-level embeddings from squeeze and excitation networks," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 980–987.
- [18] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [19] K. Rao, H. Sak, and R. Prabhavalkar, "Exploring architectures, data and units for streaming end-to-end speech recognition with rnn-transducer," in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 193–199.
- [20] Y. He, T. N. Sainath, R. Prabhavalkar, I. McGraw, R. Alvarez, D. Zhao, D. Rybach, A. Kannan, Y. Wu, R. Pang, Q. Liang, D. Bhatia, Y. Shangquan, B. Li, G. Pundak, K. C. Sim, T. Bagby, S. yin Chang, K. Rao, and A. Gruenstein, "Streaming end-to-end speech recognition for mobile devices," in *ICASSP*, 2019.
- [21] T. N. Sainath, Y. He, B. Li, A. Narayanan, R. Pang, A. Bruguier, S.-y. Chang, W. Li, R. Alvarez, Z. Chen, and *et al.*, "A streaming on-device end-to-end model surpassing server-side conventional model quality and latency," in *ICASSP*, 2020.
- [22] 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.
- [23] P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for activation functions," *arXiv preprint arXiv:1710.05941*, 2017.
- [24] A. Zeyer, P. Bahar, K. Irie, R. Schlüter, and H. Ney, "A comparison of transformer and lstm encoder decoder models for asr," in *IEEE Automatic Speech Recognition and Understanding Workshop, Sentosa, Singapore*, 2019.
- [25] 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," in *Interspeech*, 2019.
- [26] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT press, 2016.
- [27] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [28] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4510–4520.
- [29] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *arXiv preprint arXiv:1502.03167*, 2015.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [31] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [32] J. Shen, P. Nguyen, Y. Wu, Z. Chen, and *et al.*, "Lingvo: a modular and scalable framework for sequence-to-sequence modeling," 2019.
- [33] G. Synnaeve, Q. Xu, J. Kahn, T. Likhomanenko, E. Grave, V. Pratap, A. Sriram, V. Liptchinsky, and R. Collobert, "End-to-end asr: from supervised to semi-supervised learning with modern architectures," 2019.
- [34] J. Hu, L. Shen, S. Albanie, G. Sun, and A. Vedaldi, "Gather-excite: Exploiting feature context in convolutional neural networks," in *Advances in Neural Information Processing Systems*, 2018, pp. 9401–9411.
- [35] C.-C. Chiu, A. Narayanan, W. Han, R. Prabhavalkar, Y. Zhang, N. Jaitly, R. Pang, T. N. Sainath, P. Nguyen, L. Cao, and Y. Wu, "Rnn-t models fail to generalize to out-of-domain audio: Causes and solutions," *arXiv preprint arXiv:2005.03271*, 2020.
- [36] H. Liao, E. McDermott, and A. Senior, "Large scale deep neural network acoustic modeling with semi-supervised training data for youtube video transcription," in *2013 IEEE Workshop on Automatic Speech Recognition and Understanding*. IEEE, 2013, pp. 368–373.