



Multi-stage Progressive Compression of Conformer Transducer for On-device Speech Recognition

Jash Rathod¹, Nauman Dawalatabad², Shatrughan Singh¹, Dhananjaya Gowda¹

¹Samsung Research

²Massachusetts Institute of Technology

{jash.rathod, shatrughan.s, d.gowda}@samsung.com, nauman@csail.mit.edu

Abstract

The smaller memory bandwidth in smart devices prompts development of smaller Automatic Speech Recognition (ASR) models. To obtain a smaller model, one can employ the model compression techniques. Knowledge distillation (KD) is a popular model compression approach that has shown to achieve smaller model size with relatively lesser degradation in the model performance. In this approach, knowledge is distilled from a trained large size teacher model to a smaller size student model. Also, the transducer based models have recently shown to perform well for on-device streaming ASR task, while the conformer models are efficient in handling long term dependencies. Hence in this work we employ a streaming transducer architecture with conformer as the encoder. We propose a multi-stage progressive approach to compress the conformer transducer model using KD. We progressively update our teacher model with the distilled student model in a multi-stage setup. On standard LibriSpeech dataset, our experimental results have successfully achieved compression rates greater than 60% without significant degradation in the performance compared to the larger teacher model.

Index Terms: speech recognition, on-device, streaming ASR, model compression, knowledge distillation, conformer

1. Introduction

Automatic Speech Recognition (ASR) on smart devices traditionally involved sending data to a server, performing speech recognition, and returning the results to the device [1]. While this approach helps us leverage more computing power, it incurs increased latency. It would be better if this could be done on-device itself. This demands models to work with a smaller memory footprint [2]. Recently, there has been great adoption of End-to-End (E2E) ASR systems in smart devices. [1, 3, 4, 5, 6, 7, 8]. Also, for many use cases, it is desirable to use a streaming model as a real-time solution is expected [9]. Thus, it is crucial that our on-device ASR model adhere to these constraints and at the same time not compromise on the performance.

Recurrent Neural Network Transducer (RNN-T) is a widely used sequence-to-sequence streaming model for ASR [10]. It is a low-latency on-device solution that performs well when used with edge devices. The RNN-T model comprises a transcription network, a prediction network, and a joint network. The transcription network encodes the input audio using Long Short Term Memory (LSTM) cells [11, 12]. Recently in [13], authors show that the conformer model efficiently learns local as well as global interactions. Hence, the transcription network of the RNN-T, which originally uses unidirectional LSTM blocks, can

be replaced with conformer [13] block forming a conformer-transducer model.

Smaller models generally observe degradation in their performance as they cannot learn as effectively as their larger counterparts. This is due to the smaller model capacity. One of the methods to obtain the smaller models is by using model compression techniques [14, 15, 16]. Knowledge distillation (KD) [17, 18, 19] is a popular model compression technique that has recently been successful across various domains like natural language processing [20, 21, 22], computer vision [23, 24], and speech recognition [25, 26]. In this technique, knowledge is distilled from a larger model (teacher) to a smaller model (student). More recently, KD has also been shown to give a high compression rate with RNN-T based two-pass ASR model [26].

Using an even larger teacher model to distill a student model can be one avenue to improve the performance of a student, leveraging the rich knowledge learned by the teacher. However, our empirical results and recent works [27] suggest that a higher gap in sizes tends to impair the model performance, even though the larger teacher performs significantly better than a smaller teacher. This led us to the idea of bridging this gap by training a student model using a multi-stage approach.

In this work, we propose multi-stage progressive compression for the conformer-transducer model. While [27] discussed this idea for CNNs, our work is the first one to introduce this technique for ASR systems, to the best of our knowledge. In Stage 1, we use a large teacher model and compress it using KD such that it does not degrade significantly. In the next stage, the student model obtained in Stage 1 serves as the teacher model for the current stage and we distill a new student model for this stage. This process is repeated multiple times till a student model with the desirable size is achieved. The final student model obtained can then be used for on-device ASR. We perform detailed studies by comparing the results of the student models in every stage with models of the same configuration trained from scratch. We found that the student model always outperforms the model of the same size trained from scratch.

2. Methodology

In this section, we first describe the modules in the conformer transducer model used in this paper. We then describe the multi-stage progressive KD approach for compression of the conformer-transducer model as shown in Figure 1.

2.1. Conformer

The conformer block [13] includes a feed-forward module, followed by a convolutional network, a multi-headed self-attention [28] module, and finally another feed-forward module. The advantage of using this architecture is that the attention module

²Work done while at Samsung Research.

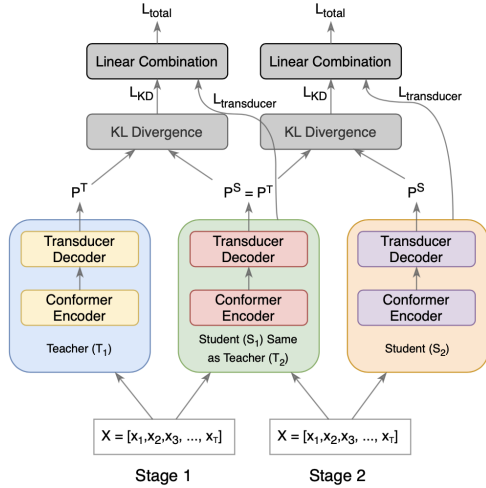


Figure 1: *Multi-stage progressive compression of conformer transducer model. The student model distilled from teacher in the first stage, is used as the teacher model for the next stage. This cascaded approach can be continued for N stages.*

helps in learning long-range global interactions well and the convolutional module exploits the local features.

Assume input given to the conformer block i is x_i and its output to be y_i . Mathematically, we can represent the conformer block as follows:

$$\begin{aligned}
 x'_i &= x_i + FF(x_i) \\
 x''_i &= x'_i + Conv(x'_i) \\
 x'''_i &= x''_i + MA(x''_i) \\
 y_i &= LayerNorm(x'''_i + FF(x'''_i))
 \end{aligned} \tag{1}$$

where FF refers to the feed-forward module, $Conv$ refers to the convolutional module, and MA refers to the multi-headed self-attention module.

The output obtained from the conformer blocks are the higher-level representations that are fed to the joint network of the transducer.

2.2. Transducer

The RNN-Transducer is a streaming sequence-to-sequence model that has an encoder-decoder architecture [10]. It includes a transcription, a prediction, and a joint network. The transcription network takes audio as the input. The prediction network, which is an autoregressive network, takes the previous output as input. They produce higher-level representations that are fed to the joint network (feed-forward network) to predict the posterior probability of the current token. In this work, we use the conformer blocks instead of LSTMs as the transcription network.

2.3. Knowledge distillation for conformer transducer

The technique of knowledge distillation aims at harnessing the rich information, that a large model learns, into a smaller model [17]. The advantage of this approach is that we obtain a smaller model that performs well and due to its size, is a better choice for being deployed to edge devices. In knowledge distillation, a teacher-student architecture is used. The output probability dis-

tribution (soft targets) represents the knowledge of the teacher model. To enable a student to make use of this, we minimize the KL-Divergence between the soft targets of the teacher and the student model. In our case, the KL-Divergence between the output of the joint network of the teacher and that of the student is minimized. If k represents the output label and (t, u) is the lattice node, the KL-Divergence loss can be represented as follows:

$$L_{KD} = \sum_{t,u} \sum_k P^T(k|t, u) \ln \frac{P^T(k|t, u)}{P^S(k|t, u)} \tag{2}$$

Let α be the distillation loss weight. Then, the total loss for training a conformer transducer student model from a conformer transducer teacher model is given as follows:

$$L_{total} = (1 - \alpha)L_{transducer} + \alpha L_{KD} \tag{3}$$

2.4. Multi-stage progressive knowledge distillation

As discussed earlier, the use of a larger and better-performing teacher model should be beneficial. Yet, if the compression rate increases beyond a limit, existing works [27] and even our empirical results (discussed in Section 4.3) indicate a sharp degradation. It is likely that the complex relations learned by the larger teacher model are too difficult for a smaller student model to learn with low model capacity. As a model compressed using KD outperforms a model of the same size trained from scratch, it is desirable to use a KD-compressed model. Hence, we compress a large streaming teacher model using knowledge distillation in multiple stages to attain a better-performing streaming student model. Figure 1 gives a schematic representation of the proposed solution. In the first stage, we compress the teacher model at a compression rate that does not hurt its performance beyond a tolerable bandwidth. We obtain an intermediate high performing student model which has the potential to be compressed further. In the next stage, we compress the student model from the previous stage (which serves as the teacher for the next stage) to attain a smaller student model. This process can be performed progressively for a number of stages till the required model size is achieved. The intuition behind this is to not lose the complex interactions learned by the larger teacher by adopting a progressive compression approach, as direct compression does not allow the student to learn these interactions.

3. Experimental Setup

3.1. Data and features

The dataset used for this work is the standard LibriSpeech dataset [29] which contains 960 hours of training data. Models are evaluated on the ‘test-clean’ set (5.4 hours; contains 2620 utterances) and ‘test-other’ set (5.1 hours; contains 2939 utterances). The features used are 40-dimensional Mel Frequency Cepstral Coefficients (MFCCs) [30] which are calculated from a 25 msec window for every 10 msec shift and are passed as an input to the transcription network. The target vocab consists of 10026 Byte Pair Encoding (BPE) [31] units. We use SpecAugment [32] technique for data augmentation during training. The beam size of 8 is used for decoding.

3.2. Transducer model and KD specifications

The specifications of each of the teacher and student models used can be found in Table 1. Other attributes remain the same across all models which are discussed next in this section. The

Table 1: Multi-stage progressive compression experiment results, including number of model parameters (in million) (rounded off to whole number), compression (Comp) percentage (rounded off to whole number), WERs and SERs for test-clean and test-other LibriSpeech datasets. In experiment 1, we compress 128M params teacher model to 46M params student model in three stages. In experiment 2, we compress 128M params teacher model to 24M params student model in two stages.

Experiment	Stage	Model	Parameters (in million)	% Comp w.r.t. teacher (%)	% Comp w.r.t. T_1 (%)	test-clean		test-other	
						WER	SER	WER	SER
Experiment 1	Stage 1	T_1	128	-	-	5.54	52.71	14.99	77.65
		B_1	80	-	38	5.98	54.81	16.12	79.35
		S_1 (Student from T_1)	80	38	38	5.47	51.76	14.96	77.17
	Stage 2	$T_2 = S_1$	80	38	38	5.47	51.76	14.96	77.17
		B_2	62	-	52	6.19	55.99	16.66	79.38
		S_2 (Student from T_1)	62	52	52	5.67	52.91	15.01	77.82
		S_3 (Student from T_2)	62	23	52	5.48	52.18	14.85	77.17
	Stage 3	$T_3 = S_3$	62	23	52	5.48	52.18	14.85	77.17
		B_3	46	-	64	6.65	58.89	17.03	79.96
		S_4 (Student from T_1)	46	64	64	6.20	56.04	16.38	78.9
		S_5 (Student from T_3)	46	25	64	6.08	55.53	16.11	78.73
Experiment 2	Stage 1	T_1	128	-	-	5.54	52.71	14.99	77.65
		B_4	57	-	55	6.32	56.3	16.72	79.72
		S_6 (Student from T_1)	57	55	55	5.76	53.36	15.44	77.95
	Stage 2	$T_4 = S_6$	57	55	55	5.76	53.36	15.44	77.95
		B_5	24	-	81	8.17	64.24	20.17	83.63
		S_7 (Student from T_1)	24	81	81	7.33	60.99	18.88	83.12
		S_8 (Student from T_4)	24	58	81	7.21	60.11	18.59	83.06

entire training pipeline is developed using Tensorflow 2 [33]. A dropout [34] of 0.1 is used succeeding every layer other than the first one. Max-pooling of 2 for the first three layers (and an overall factor of 8) is used which helps in shortening the training time and latency of the model. To reduce the number of parameters and regularize our model, the same neural network serves as the joint network and the embedding layer in the prediction network. We use a learning rate represented by lr of 0.0001 and decay it using the decay function as follows:

$$\lambda = \frac{step}{warmup}$$

$$\gamma = \frac{step - warmup}{decay_steps} \quad (4)$$

$$lr = \begin{cases} lr * \lambda^3 + 10^{-7}, & \text{if } step < warmup \\ \min(lr, 5 * lr * 0.9^\gamma), & \text{otherwise} \end{cases}$$

where $step$ represents current training time step, $warmup$ is set to be 10000, and $decay_steps$ is set to be 40000.

As for knowledge distillation, our preliminary experiments show that a distillation loss weight of 0.02 gives the best results. For all our experiments, a temperature [17] of 1.0 has been used for both, the teacher and the student models.

3.3. Notations used

In this paper, we denote a teacher model with T and a student model with S . For a detailed comparison, we have also trained baselines models B with the same configuration as the student models, trained from scratch. Thus, T and B are models trained without knowledge distillation. A subscript acts as identification between different models of the same type.

Table 1 includes the experiments performed, their model sizes in number of parameters in million (M), the compression percentages (%), and the evaluation metrics - Word Error Rate (WER) and Sentence Error Rate (SER). The $T_i = S_j$ indicates that S_j now serves as T_i and both are the same models.

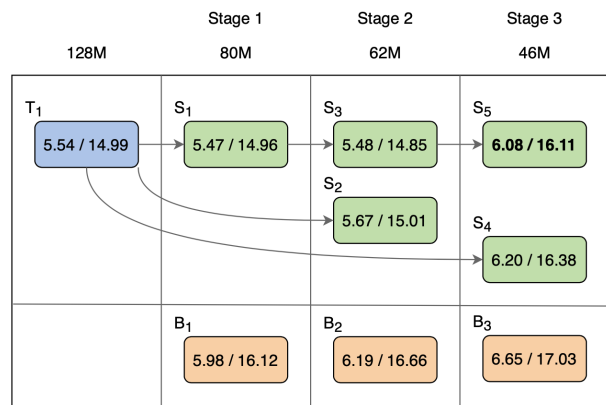


Figure 2: Results for multi-stage KD compression of conformer transducer Experiment 1 (compression rate of under 50% in each stage) on test-clean / test-other sets. T_i , S_j , and B_k represent teacher, student, and baseline models respectively.

4. Results and Discussion

In this section, we exhibit the efficacy of our proposed multi-stage progressive compression approach. We present results on varying compression rates and their performance as discussed in Section 2.3. All models in these experiments are streaming models. Each model was trained on two NVIDIA A100 GPUs for 80 epochs. Also, all the results are without the use of a Language Model (LM). For every stage, compression is achieved by reduction in encoder and decoder dimensions and/or number of layers.

4.1. Experiment 1 - compression rate of under 50% in each stage

Table 1 and Figure 2 highlight results of multi-stage progressive compression approach where we perform two sets of ex-

	Stage 1	Stage 2
	128M	57M
T_1	5.54 / 14.99	7.21 / 18.59
S_1	5.76 / 15.44	7.33 / 18.88
B_1	6.32 / 16.72	8.17 / 20.17

Figure 3: Results for multi-stage KD compression of conformer transducer Experiment 2 (compression rate of over 50% in each stage) on test-clean / test-other sets. T_i , S_j , and B_k represent teacher, student, and baseline models respectively.

periments. In experiment 1, we compress a 128M parameters teacher model to a 46M parameters student model in three stages progressively. The 128M teacher model has 16 encoder layers with 512 dimensions, 4 attention heads, and 2 decoder layers with 1024 dimensions. In each stage, we have kept the compression rate below 50% compared to the teacher model size in that stage. It can be observed that in Stage 1, we obtained the student model S_1 (distilled from T_1) with better WER than the teacher and 38% smaller in size. For comparison, we also trained a baseline model B_1 (same configuration model but trained from scratch without KD). As expected, B_1 is found to be worse than the student model S_1 .

In Stage 2 of multi-stage progressive KD, the S_1 model is used as the teacher model T_2 . Even here, we observe that our student model S_3 has a WER almost the same as that of the teacher while being 23% smaller than T_2 . Again, clearly student model S_3 outperforms B_2 and S_2 (distilled from T_1). It is worth noticing that, compared to the T_1 , our student model S_3 trained through the proposed progressive KD has been compressed by 52% with no loss in WER performance.

For Stage 3, S_3 model now acts as the teacher model T_3 . We went a step ahead and compressed the model by 25% which brings the model size to 46M parameters. This time we observe approximately 11% and 8.5% relative degradation in WER on test-clean and test-other respectively. This minor degradation in WER is reasonable given a high compression rate of 64% with respect to original teacher model T_1 . Also, when compared to the baseline B_3 model, the student S_4 and S_5 performs consistently better again. S_5 performing better than S_4 shows that the student model obtained using our multi-stage progressive compression technique produces the most superior results. Compared to original teacher model T_1 , the minor relative degradations of 10% and 7.5% in WERs on test-clean and test-other datasets are observed. At the same time, a significantly high compression rate of 64% compared to the teacher model is achieved.

4.2. Experiment 2 - compression rate of over 50% in each stage

In order to experiment with still higher compression rates (beyond 55%), we performed experiment 2 which uses two stages. The results of this set of experiments are shown in bottom row of Table 1 and in Figure 3. For Stage 1, we use the same 128M

Table 2: Higher compression rates on 128M param teacher model using single stage knowledge distillation

Model	M Params (% Comp)	test-clean		test-other	
		WER	SER	WER	SER
T_1	128 (-)	5.54	52.71	14.99	77.65
S_9	41 (68%)	6.21	56.34	16.53	79.41
S_{10}	37 (71%)	6.25	56.26	16.63	79.72
S_{11}	30 (76%)	7.09	60.08	18.05	81.59
S_{12}	27 (79%)	6.96	59.25	18.28	82.44
S_{13}	29 (77%)	7.81	62.75	20.07	84.38
S_{14}	19 (85%)	9.32	68.21	22.55	87.24

parameters teacher T_1 . We distill a student model S_6 from T_1 with a compression rate of 55% and observe minor degradation of approximately 4% and 3% (relative) in WER on the test-clean and test-other datasets respectively. Also, we observe that the performance of the baseline B_4 is significantly worse.

For Stage 2, we compress the teacher model T_4 (i.e. S_6) by 58% to obtain a model with only 24M parameters. While S_7 performs better than the B_5 , it is marginally worse than S_8 . We get a degradation in WER of about 24% (relative) on both test-clean and test-other. This degradation in WER is expected given a high compression rate of 81% with respect to T_1 .

4.3. Ablation study on single stage KD

In this section, we demonstrate the effect of compression rates higher than 65% using only a single stage. From results in Table 2, we can observe that generally as the compression increase, the WERs tend to degrade significantly. S_9 and S_{10} are identical except for a higher decoder compression rate used in the latter. S_{11} and S_{12} are also identical except for a higher decoder compression rate used in the latter and both have higher encoder compression rate than S_9 and S_{10} . Considering sets (S_9 , S_{13}) and (S_{12} , S_{14}), each models in their respective sets are identical except for higher encoder compression rates used in the latter. Based on our studies, we have been able to draw some insights as in [35]. Firstly, compression in decoder doesn't affect the performance of the models as much. And secondly, even a slightly higher compression in the encoder parameters (number of layers and input sizes) affects the model performance.

5. Conclusion

In this work, we propose a multi-stage progressive approach to compress conformer transducer model using knowledge distillation. The proposed approach begins with a 128M parameters streaming teacher model and distills a smaller streaming model without any degradation in performance of the model. In the next stage, this student model is used as the teacher model to obtain a smaller model for KD in this stage. Using this cascaded approach, we obtain two models: a 46M parameter model with a compression of 64% in three stages, and a 24M parameter model with a compression of 81% in two stages. The former shows a minor degradation, yielding 10% and 7.5% relative in WERs on LibriSpeech's test-clean and test-other datasets respectively. These compression rates are good enough for medium-end mobile devices. Further, we also demonstrate the effects of single stage compression of a large teacher and share our insights. In future, we aim to focus specifically on the encoder compression as it makes up a sizable portion of the model size and has a significant impact on performance.

6. References

- [1] Y. He, T. N. Sainath, R. Prabhavalkar, I. McGraw, R. Alvarez, D. Zhao, D. Rybach, A. Kannan, Y. Wu, R. Pang *et al.*, “Streaming end-to-end speech recognition for mobile devices,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6381–6385.
- [2] T. N. Sainath, Y. He, B. Li, A. Narayanan, R. Pang, A. Bruguier, S.-y. Chang, W. Li, R. Alvarez, Z. Chen *et al.*, “A streaming on-device end-to-end model surpassing server-side conventional model quality and latency,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6059–6063.
- [3] T. N. Sainath, R. Pang, D. Rybach, Y. He, R. Prabhavalkar, W. Li, M. Visontai, Q. Liang, T. Strohmaier, Y. Wu, I. McGraw, and C.-C. Chiu, “Two-Pass End-to-End Speech Recognition,” in *InterSpeech*, 2019, pp. 2773–2777.
- [4] G. Venkatesh, A. Valliappan, J. Mahadeokar, Y. Shangguan, C. Fuegen, M. L. Seltzer, and V. Chandra, “Memory-Efficient Speech Recognition on Smart Devices,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 8368–8372.
- [5] K. Kim, K. Lee, D. Gowda, J. Park, S. Kim, S. Jin, Y.-Y. Lee, J. Yeo, D. Kim, S. Jung *et al.*, “Attention based on-device streaming speech recognition with large speech corpus,” in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019, pp. 956–963.
- [6] A. Garg, D. Gowda, A. Kumar, K. Kim, M. Kumar, and C. Kim, “Improved Multi-Stage Training of Online Attention-based Encoder-Decoder Models,” in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019, pp. 70–77.
- [7] C. Kim, S. Kim, K. Kim, M. Kumar, J. Kim, K. Lee, C. Han, A. Garg, E. Kim, M. Shin, S. Singh, L. Heck, and D. Gowda, “End-to-end training of a large vocabulary end-to-end speech recognition system,” in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019, pp. 562–569.
- [8] A. Garg, G. Vadiseti, D. Gowda, S. Jin, A. Jayasimha, Y. Han, J. Kim, J. Park, K. Kim, S. Kim, Y. Lee, K. Min, and C. Kim, “Streaming on-device end-to-end ASR system for privacy-sensitive voicetyping,” in *InterSpeech*, 2020, pp. 3371–3375.
- [9] J. Yu, W. Han, A. Gulati, C.-C. Chiu, B. Li, T. N. Sainath, Y. Wu, and R. Pang, “Dual-mode asr: Unify and improve streaming asr with full-context modeling,” *arXiv preprint arXiv:2010.06030*, 2020.
- [10] A. Graves, “Sequence Transduction with Recurrent Neural Networks,” *arXiv preprint arXiv:1211.3711*, 2012.
- [11] H. Sak, A. W. Senior, and F. Beaufays, “Long short-term memory recurrent neural network architectures for large scale acoustic modeling,” in *InterSpeech*, 2014, pp. 338–342.
- [12] A. Sherstinsky, “Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network,” *Physica D: Nonlinear Phenomena*, vol. 404, p. 132306, 2020.
- [13] 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,” in *InterSpeech*, 2020, pp. 5036–5040.
- [14] Y. Cheng, D. Wang, P. Zhou, and T. Zhang, “Model compression and acceleration for deep neural networks: The principles, progress, and challenges,” *IEEE Signal Processing Magazine*, vol. 35, no. 1, pp. 126–136, 2018.
- [15] J. Wu, C. Leng, Y. Wang, Q. Hu, and J. Cheng, “Quantized convolutional neural networks for mobile devices,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4820–4828.
- [16] X. Yu, T. Liu, X. Wang, and D. Tao, “On compressing deep models by low rank and sparse decomposition,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7370–7379.
- [17] G. Hinton, O. Vinyals, and J. Dean, “Distilling the Knowledge in a Neural Network,” in *NIPS Deep Learning and Representation Learning Workshop*, 2015.
- [18] Y. Chebotar and A. Waters, “Distilling Knowledge from Ensembles of Neural Networks for Speech Recognition,” in *InterSpeech*, 2016, pp. 3439–3443.
- [19] J. Gou, B. Yu, S. J. Maybank, and D. Tao, “Knowledge Distillation: A Survey,” *International Journal of Computer Vision*, vol. 129, no. 6, pp. 1789–1819, Mar 2021.
- [20] I. Turc, M.-W. Chang, K. Lee, and K. Toutanova, “Well-read students learn better: On the importance of pre-training compact models,” *arXiv preprint arXiv:1908.08962*, 2019.
- [21] Y. Kim and A. M. Rush, “Sequence-level knowledge distillation,” in *EMNLP*, 2016, pp. 1317–1327.
- [22] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, “DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter,” *arXiv preprint arXiv:1910.01108*, 2019.
- [23] Y. Liu, K. Chen, C. Liu, Z. Qin, Z. Luo, and J. Wang, “Structured knowledge distillation for semantic segmentation,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2604–2613.
- [24] G. Chen, W. Choi, X. Yu, T. Han, and M. Chandraker, “Learning efficient object detection models with knowledge distillation,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [25] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “Fastspeech 2: Fast and high-quality end-to-end text to speech,” *arXiv preprint arXiv:2006.04558*, 2020.
- [26] N. Dawalatabad, T. Vatsal, A. Gupta, S. Kim, S. Singh, D. Gowda, and C. Kim, “Two-Pass End-to-End ASR Model Compression,” in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2021, pp. 403–410.
- [27] S. I. Mirzadeh, M. Farajtabar, A. Li, N. Levine, A. Matsukawa, and H. Ghasemzadeh, “Improved knowledge distillation via teacher assistant,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020, pp. 5191–5198.
- [28] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [29] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [30] Z. Wanli and L. Guoxin, “The research of feature extraction based on MFCC for speaker recognition,” in *Proceedings International Conference on Computer Science and Network Technology*, 2013, pp. 1074–1077.
- [31] R. Sennrich, B. Haddow, and A. Birch, “Neural machine translation of rare words with subword units,” *arXiv preprint arXiv:1508.07909*, 2015.
- [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] M. Abadi, A. Agarwal, P. Barham *et al.*, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015, software available from tensorflow.org. [Online]. Available: <https://www.tensorflow.org/>
- [34] 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.
- [35] R. Botros, T. N. Sainath, R. David, E. Guzman, W. Li, and Y. He, “Tied & Reduced RNN-T Decoder,” in *InterSpeech*, 2021.