



A Unified Cascaded Encoder ASR Model for Dynamic Model Sizes

Shaojin Ding*, Weiran Wang*, Ding Zhao*, Tara N. Sainath, Yan Zhang He, Robert David, Rami Botros, Xin Wang, Rina Panigrahy, Qiao Liang, Dongseong Hwang, Ian McGraw, Rohit Prabhavalkar, Trevor Strohman

Google LLC, USA

{shaojinding, weiranwang, dingzhao}@google.com

Abstract

In this paper, we propose a dynamic cascaded encoder Automatic Speech Recognition (ASR) model, which unifies models for different deployment scenarios. Moreover, the model can significantly reduce model size and power consumption without loss of quality. Namely, with the dynamic cascaded encoder model, we explore three techniques to maximally boost the performance of each model size: 1) Use separate decoders for each sub-model while sharing the encoders; 2) Use funnel-pooling to improve the encoder efficiency; 3) Balance the size of causal and non-causal encoders to improve quality and fit deployment constraints. Overall, the proposed large-medium model has 30% smaller size and reduces power consumption by 33%, compared to the baseline cascaded encoder model. The triple-size model that unifies the large, medium, and small models achieves 37% total size reduction with minimal quality loss, while substantially reducing the engineering efforts of having separate models.

Index Terms: End-to-end ASR, RNN-T, Conformer

1. Introduction

End-to-end (E2E) models [1, 2, 3, 4, 5] have gained popularity over the past few years, particularly for on-device automatic speech recognition (ASR), as they can achieve similar recognition performance compared to conventional hybrid systems [6] at a fraction of the size. Over the past few years, developing an E2E model that surpasses conventional models in both quality and latency in diverse test conditions has been an active research area across many research groups [7, 8, 9, 10, 11, 12].

Recently, we presented an on-device E2E model based on a two-pass cascaded encoder which outperforms a conventional model in terms of word error rate (WER) on both search and long-tail queries, as well as endpointer latency metrics [13]. We further adapted the cascaded encoder to a small 1st-pass (50M parameters) large 2nd-pass (100M parameters) architecture to improve computational latency for both cloud and edge tensor processing units (TPUs), while maintaining quality [14].

However, on-device ASR systems often require different model sizes for deployment to a variety of edge devices with different hardware constraints, e.g. mobile phones, home speakers, or cars. Even in the same device, different model sizes might still be required for various application constraints, e.g. a large model might be used for short-form applications (like voice search) to obtain the best quality, while a medium or a small model might be required for long-running applications (like dictation or video captioning) to maintain low power consumption. It is inefficient to train these different-sized models separately with duplicate efforts and high maintenance cost,

especially for multiple languages.

To support such diversity of scenarios, we propose an approach by extending the cascaded encoder architecture in [13] to unify multiple size configurations in a single model during training. By only running a subset of the model layers at inference time, the model can be executed as different sizes with similar accuracies as the independently trained models of the corresponding sizes. This greatly reduces both the training overhead and the management complexity of deployment processes, and also allows run-time on-the-fly model size adjustment for variable resource usage. Furthermore, we apply the following novel optimizations to improve quality, memory and latency: 1) Replace the shared decoder in the cascaded encoder model with separate decoders, which we will show is more robust to smaller encoder sizes; 2) Replace the stacking layer for downsampling in the causal encoder with a funnel-pooling layer to help reduce the size of the encoder [15]; 3) Balance the size of causal and non-causal encoders to improve quality and fit deployment constraints. We conduct extensive experiments on large scale tasks including voice search and dictation. Results show that our unified large-medium model achieves the same accuracy as the cascaded encoder baselines, with only about 70% of model size, significantly reducing power consumption in the dictation task. Moreover, the unified large-medium-small model obtains minimal accuracy loss along with 37% size reductions, compared to the upper-bounded individually trained models.

Relation to prior work. Several prior studies also explored the idea of jointly training ASR models with different sizes. The closest works to ours are [16, 17], which investigated encoder and decoder weight sharing among large/medium/small models. However, all their encoder layers are non-causal, leading to significant latency increase at inference time. By contrast, our proposed model unifies both causal and non-causal layers, which makes it more efficient and flexible under different hardware constraints. More importantly, in these works, the model of each size has leveraged dedicated encoder layers that are not shared with other model sizes, which increases the overall model size. However, as we have shown in the experiments, using smaller separate decoders avoids additional model size overhead and even allows the use of smaller encoders without any performance degradation. Secondly, [16, 17, 18] had additional distillation loss terms during the joint model training. In contrast, our preliminary experiments show that it is not straightforward to perform distillation between the causal layers and non-causal layers to improve the performance of causal layers, potentially due to the different right context; this direction is left as future work. Lastly, compared with the alternative approach of model shrinking with sparsity networks [19, 20], our model is dense and requires no additional hardware support. Furthermore, it is more convenient to control the amount of right context in each size within our framework, and our training pipeline

* Equal Contribution. Listed in alphabetical order.




Figure 1: *Dynamic cascaded encoder model structure.*

is much simpler, without the need for warm-starting a sparse model with a trained dense model.

2. Method

In this section, we first introduce the proposed dynamic cascaded encoder model architecture, followed by the detailed descriptions of each of our novel designs. Finally, we present two specific dynamic cascaded encoder model architectures for practical applications.

2.1. Dynamic cascaded encoder model

The baseline Conformer-based [21] cascaded encoder model [14] is comprised of a causal conformer encoder with N layers, followed by a non-causal conformer encoder [22] with M layers and an embedding RNN-T decoder [23]. To improve the flexibility in unifying different models, we reformulate the cascaded model architecture to allow easy extractions of models with different sizes, as shown in Figure 1. In our model, each causal layer can be connected to the decoder or the first non-causal layer. We also allow connections from any non-causal layer to the decoder. From the super-net, we extract K sub-models, each containing the first n_k ($0 \leq n_k \leq N$) causal layers, and the first m_k ($0 \leq m_k \leq M$) non-causal layers, which can be used under different model size and latency restrictions:

$$\mathbf{y}_k = \text{Dec}(\text{Enc}_k^{nc}(\text{Enc}_k^c(\mathbf{x}))) \quad (1)$$

where \mathbf{x} and \mathbf{y}_k denote the input and output of the k -th sub-model (all the sub-models have the same input). Enc_k^c is the causal encoder containing n_k causal layers, Enc_k^{nc} is the non-causal encoder containing m_k non-causal layers, and Dec is the shared decoder. Note that each of our sub-models does not have any dedicated encoder layer during training to minimize the total memory and storage cost in practice.

2.2. Separate decoders

The original cascaded encoder model [22] uses a shared RNN-T decoder. The decoder works with a causal encoder in the first pass to provide streaming recognition results, and works with an additional non-causal encoder that sits on top of the causal encoder to provide more accurate final results, leveraging audio right context extracted by the noncausal encoder. Therefore, the same decoder has to deal with features of different context, and we observe tension between the performance of the passes as we try to reduce the model size, i.e., as we assign more loss weights for the causal pass to satisfy WER target, the accuracy of the non-causal pass degrades.

In this work, we propose to use smaller separate decoders in each sub-model, to better cope with the different context, and this significantly alleviates the tension between different sub-models:

$$\mathbf{y}_k = \text{Dec}_k(\text{Enc}_k^{nc}(\text{Enc}_k^c(\mathbf{x}))) \quad (2)$$




Figure 2: *Sub-model with separate decoders in a Dynamic cascaded encoder model.*

Figure 2 shows an example of a sub-model with separate decoders: solid arrows are the connections used by this sub-model, and dotted arrows are connections used by other sub-models. As we will show in the experiments, empirically we can keep increasing the loss weight of the causal pass for better streaming results, without sacrificing performance of the non-causal pass. This allows us to use smaller separate decoders to replace the shared decoder, thus saving total memory cost and improving the inference speed of each sub-model.

2.3. Funnel-pooling layers

To reduce the overall computational cost, prior models usually use a stacking layer in the causal encoder to down-sample the input frame rate. The stacking layer concatenates features of two consecutive frames, and thus doubling the dimension of its output, which is used as input to the next attention layer and results in large amount of weight parameters in that layer. However, it is extremely parameter-inefficient. To address the issue, we explore alternative down-sampling techniques. The most straight-forward substitution could be average pooling. However, using average pooling at the bottom layers usually introduce inevitable performance regressions [15]. Observing this, we propose to use funnel pooling [15] to down-sample the input frame rate, which has been shown to be able to preserve the model performance while reducing the frame rate in the middle of a sequential model.

Suppose we have a feature map $\mathbf{h} \in \mathbb{R}^{T \times D}$ as the input to a self-attention layer, where T and D denote the original sequence length and feature dimensions, respectively. We first create a down-sampled sequence of $\mathbf{h}' \in \mathbb{R}^{T' \times D}$ through average pooling:

$$\mathbf{h}' = \text{AvgPool}(\mathbf{h}) \quad (3)$$

where $T' = T/2$ in our case (down-sampled by a factor of 2). Instead of simply feeding \mathbf{h}' to the self-attention, we only use \mathbf{h}' as the query vector \mathbf{q} in the self-attention layer. The key \mathbf{k} and value vectors \mathbf{v} are still based on the original input feature map \mathbf{h} :

$$\mathbf{h}'' = \text{SelfAttention}(\mathbf{q} = \mathbf{h}', \mathbf{k} \mathbf{v} = \mathbf{h}) \quad (4)$$

where $\mathbf{h}'' \in \mathbb{R}^{T' \times D}$ is the output feature maps.

2.4. Sub-model joint training

We perform standard two-stage training as done in previous work. During maximum likelihood estimation training, we forward a minibatch through all sub-models and compute the loss of each sub-model:

$$L_k = \text{LOSS}_{\text{RNNT}}(\mathbf{y}_k) \quad (5)$$

and the losses for all sub-models are combined linearly,

$$L = \sum_{k=0}^K \lambda_k \cdot L_k \quad (6)$$




Figure 3: *Triple-sized large-medium-small model.*

where λ_k is the weight of the k -th sub-model, and all the weights sum to 1. After that, we continue fine-tuning the model with discriminative training using the MWER criteria [24]. For each step of MWER training, we randomly sample each sub-model with a probability equal to its loss weight, and use the sampled decoder to perform beam search on the minibatch to generate the top-4 hypotheses. The (full-sum) negative log-likelihood are computed for the hypotheses using the same sampled pass, and re-normalized in the top-4 space (so that the conditional "probabilities" sum to 1) to approximate the expected word error loss for minimization.

2.5. Dynamic cascaded encoder model in practice

With the flexibility of the dynamic cascaded encoder model, we establish a large-medium super-net and a large-medium-small super-net that work for most of the practical use cases. The large-medium super-net has a 46.8M causal encoder for the medium sub-model and an additional 60M non-causal encoder for the large pass, each having a 4.4M separate decoder. With the balanced size of causal and non-causal encoders, we show that it improves quality and fits deployment constraints better in Section 4.3. Our large-medium model only has around 70% of model size, compared to the previous models in [13, 14]. Similarly, the large-medium-small super-net is comprised of a 20M causal encoder for the small sub-model, an additional 26.8M causal encoder for the medium sub-model, and a final 60M non-causal encoder for the large sub-model, as shown in Figure 3. The non-causal layer is only added to the large sub-model, because it requires fast hardware to catch up delays introduced by the right context, although it gives considerable quality gain. Each of the separate decoders also has 4.4M parameters.

3. Experimental setup

3.1. Dataset

Similar to [25, 26], all models are trained with $\sim 400k$ hours English audio-text pairs from multiple domains, such as YouTube and anonymized voice search traffic. YouTube data is transcribed in a semi-supervised fashion [27]. All other domains are anonymized and hand-transcribed. Our data handling abides by *Google AI Principles* [28]. We use a mixed-case word-piece vocabulary for all our experiments for on-device ASR to avoid a separate capitalization normalizer after decoding. This is different from previous studies [26, 13, 14] that are conducted using lowercase wordpieces for cloud-based E2E models. To avoid domain overfitting and increase data diversity, we apply two data augmentation techniques, including "multi-style training" (MTR) [29] and Spec-Augmentation [30].

During testing, we use the Voice Search (VS) test set and the Gboard Dictation Donation (Dictation) test set to evaluate the system performance. Voice Search contains around 12k voice search utterances, each having an average length of 5.5 seconds. Gboard Dictation Donation has 15k utterances

and is collected as part of a voluntary program where users may choose to donate snippets of dictation speech to help improve speech models. Both search and dictation utterances are anonymized and hand-transcribed.

3.2. Implementation details

In our large-medium super-net, the causal encoder for the medium sub-model has seven 512-dimensional conformer layers (first three layers have no self-attention) with 23-frame left context per layer, and no right context to strictly prevent the model from using future inputs. The additional non-causal encoder for large pass has six 640-dimensional conformer layers, with additional 30-frame right context across six layers that processes 900ms speech from the future. All the self-attention layers have eight heads. Each separate RNN-T decoder is comprised of an 320-dimensional embedding prediction network and a 384-dimensional fully-connected joint network. We jointly train the super-net as described in Sec 2.2, and we experimented with the weights in Section 4.1. The large-medium-small super-net, has six 256-dimensional conformer layers for small sub-model, an additional six 512-dimensional causal conformer layers for the medium sub-model, and another six 640-dimensional non-causal layers for the large sub-model. The loss weights during joint model training are set to [0.80, 0.15, 0.05] for small, medium, and large sub-models, respectively.

We use the 128-dimensional log Mel-filterbank energies (extracted from 32ms window and 10ms shift) as the frontend feature, and then we stack the contiguous 4 frames, sub-sampled by a factor of 3, and append a 16-dimensional one-hot domain-ID vector [26]. All our evaluations are running on an on-device inference pipeline, where we first convert the TensorFlow graphs to TensorFlow Lite format, and leverage the 8-bit post training quantization to further reduce the model file size. Additionally, we did not use any language model in our experiments, as this is orthogonal to the end-to-end model improvements. The dictation power consumption is measured for recognizing a 14-minute continuous speech recording on a Pixel 6 mobile phone with the edge TPU on the Google Tensor chip.

4. Results

We conduct four sets of experiments to evaluate our proposed approach. First, we conduct two ablation studies verifying the impact of separate decoders and funnel pooling in the proposed dynamic cascaded encoder model, based on our large-medium model. Following this, we compare our best-performing large-medium model and large-medium-small model to the corresponding baseline methods, respectively, to show the effectiveness of our proposed approach.

4.1. Impact of separate decoders

We first examine the impact of the newly proposed separate decoders, by comparing with the previously used shared decoder approach [14]. We provide the WERs on the VS testset in Table 3. MWER training tends to reduce the WERs by similar amounts for both type of models, as shown in E4.

As we skew the loss weight towards the small sub-model, shared decoder models do get improved accuracy for the small sub-model, and the WER reduces from 9.0% to 8.2% when the its weight increase from 0.6 to 0.95. However, this comes at the cost of a worse second pass, whose WER increase from 6.5% to 6.9%. In comparison, for models with separate decoders, as the medium sub-model WER decrease from 9.0% to 8.5%, the large sub-model WER only degraded by 0.1% from 6.1% to

Table 1: Comparing proposed large-medium models to cascaded encoder baselines. The numbers in grey indicate that the model cannot be used in practical application due to their high power consumption.

Exp	Model	VS WER		Dictation WER		Dictation Power (mW)		Size (MB)	
		medium	large	medium	large	medium	large	medium	large
B0	Conf. cascaded encoder [13]	6.9	5.8	5.8	5.3	272	410	120	152
B1	Small 1st/Large 2nd [14]	8.6	5.9	7.0	5.3	259	418	56	155
E6	Proposed large-medium model	7.9	5.8	6.6	5.6	190	273	44	108

Table 2: Comparing large-medium-small model with separately trained large/medium/small models.

Exp	Model	VS WER			Size (MB)
		Small	Medium	Large	
B2	Separately trained models	10.0	7.3	5.7	180
E6	Proposed large-medium-small model	10.6	7.7	6.1	115

Table 3: Evaluations on the effectiveness of separate decoders. We provide VS WERs (%) for all models (E0 to E3 without MWER; E4 with MWER).

Exp	medium/large weights	Shared dec.		Separate decs.	
		medium	large	medium	large
E0	0.6/0.4	9.0	6.5	9.0	6.1
E1	0.8/0.2	8.7	6.5	8.7	6.2
E2	0.9/0.1	8.4	6.6	8.5	6.2
E3	0.95/0.05	8.2	6.9	8.5	6.2
E4	0.9/0.1 w/ MWER	7.8	6.2	7.9	5.8

6.2%. Therefore, we stick to the separate decoders setup with 0.9 vs 0.1 loss weights.

4.2. Impact of funnel pooling

To evaluate the effectiveness of funnel pooling, we compare it against two variants, i.e., using stacking and using average pooling for down-sampling. Results are shown in Table 4. As we expect, the model with funnel pooling can achieve the same WERs as the model based on stacking. Additionally, comparing funnel pooling and average pooling, we do see a 0.2 WER regression in the model based on average pooling for both medium and large sub-models, further demonstrating the necessity of funnel pooling.

Table 4: Evaluations on the effectiveness of funnel pooling.

Exp	Model	VS WER		Size (MB)
		medium	large	
E4	Stacking	7.9	5.8	115
E5	Average pooling	8.1	6.0	108
E6	Funnel pooling	7.9	5.8	108

4.3. Comparisons between the large-medium model and baseline cascaded encoder models

After validating the use of separate decoders and funnel pooling, we discuss the performance of the large-medium model. We consider two conformer cascaded encoder baselines: (B0) The original conformer cascaded encoder model in [13], and (B1) the small 1st/large 2nd conformer cascaded encoder model [14] that is optimized for cloud TPU.

Results are shown in Table 1. Comparing between the two baselines, we confirm the medium sub-model degradation issue of model B1 (6.9 vs. 8.6), which is one of the motivations of this study. Our proposed model (E6) can significantly mitigate the degradation and improve the first pass WER from 8.6 to 7.9. More importantly, E6 has a much smaller total model size (108MB) compared to the baselines (~30% relative reduction), while retaining the large sub-model VS WER. Besides quality-wise improvements, the proposed model also benefits in terms of the power consumption. When using B0 or B1 in recognizing continuous speech, although large sub-model has a better WER, we still rely on only the medium sub-model, since

running the large sub-model leads to much higher power consumption (e.g., B0: 270mW vs. 410mW). By contrast, with the reduced model size, the large sub-model of E6 achieves similar power consumption to that of the baselines so that it can be used for long-running applications, while obtaining 0.2 and 1.4 absolute dictation WER reduction compared to the medium sub-models of B0 and B1 respectively.

4.4. Comparisons between the large-medium-small model and the separately trained models

Finally, we illustrate the capability of our triple-size model that unifies the large, medium, and small model production models. We compare it against a baseline (B2) of separately trained large, medium, and small models. B2 can be treated as an upper-bound to the proposed model, as there is no weight sharing and each size has a dedicated optimized model. Table 2 shows the results of the two models. Compared to separately trained models, our unified model reduces 37% model size with only a minimal WER regression, and the 6.1 WER on the large sub-model has already surpassed the quality of the server conventional model [26]. The unified model allows us to use smaller sub-models to reduce model loading or computational latency during model cold-start or bursty audio situations, while switching to larger sub-models afterwards for better quality without increasing much memory, similar to [17]. Also, it reduces the engineering efforts in model tuning and runtime optimizations, which is beneficial to large scale productionizations.

5. Conclusions

We have proposed a dynamic cascaded encoder ASR model based on separate decoders, which generalizes well to different model sizes, unifying the large, medium, and small models for different deployment scenarios. Moreover, the model can significantly reduce model size and power consumption compared to prior methods. Our experimental results confirmed that the separate decoders obtained a more promising performance compared to the shared decoder. In addition, with separate decoders, we showed that the efficiency of the encoders can be further improved via funnel pooling and deliberately designing between causal/non-causal encoder sizes, resulting in a 30% smaller model size without any performance loss. Compared to baseline models, the proposed model reduces dictation power consumption on large sub-model by 33%, which makes it possible to run inference with large sub-model for dictation with improved quality. Compared to separately trained large, medium, and small models, the proposed architecture achieves 37% total size reduction, with slight performance degradations.

6. References

- [1] D. Wang, X. Wang, and S. Lv, "An overview of end-to-end automatic speech recognition," *Symmetry*, vol. 11, no. 8, p. 1018, 2019.
- [2] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Sathesh, S. Sengupta, A. Coates *et al.*, "Deep speech: Scaling up end-to-end speech recognition," *arXiv preprint arXiv:1412.5567*, 2014.
- [3] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [4] J. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," in *ICONIP*, 2015, pp. 577–585.
- [5] L. Dong, S. Xu, and B. Xu, "Speech-transformer: a no-recurrence sequence-to-sequence model for speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5884–5888.
- [6] G. Pundak and T. N. Sainath, "Lower frame rate neural network acoustic models," in *Proc. Interspeech*, 2016.
- [7] J. Li, Y. Wu, Y. Gaur *et al.*, "On the Comparison of Popular End-to-End Models for Large Scale Speech Recognition," in *Proc. Interspeech*, 2020.
- [8] Y. He, T. N. Sainath, R. Prabhavalkar *et al.*, "Streaming End-to-end Speech Recognition For Mobile Devices," in *Proc. ICASSP*, 2019.
- [9] C.-C. Chiu, T. N. Sainath, Y. Wu *et al.*, "State-of-the-art Speech Recognition With Sequence-to-Sequence Models," in *Proc. ICASSP*, 2018.
- [10] S. Kim, T. Hori, and S. Watanabe, "Joint CTC-attention based end-to-end speech recognition using multi-task learning," in *Proc. ICASSP*, 2017.
- [11] J. Li, R. Zhao, H. Hu, and Y. Gong, "Improving RNN transducer modeling for end-to-end speech recognition," in *Proc. ASRU*, 2019.
- [12] A. Zeyer, A. Merboldt, R. Schlüter, and H. Ney, "A new training pipeline for an improved neural transducer," in *Proc. Interspeech*, 2020.
- [13] T. N. Sainath, Y. He, A. Narayanan, R. Botros *et al.*, "An Efficient Streaming Non-Recurrent On-Device End-to-End Model with Improvements to Rare-Word Modeling," in *Proc. of Interspeech*, 2021.
- [14] T. N. Sainath, Y. He, A. Narayanan *et al.*, "Improving the Latency and Quality of Cascaded Encoders," in *ICASSP*, 2022.
- [15] Z. Dai, G. Lai, Y. Yang, and Q. Le, "Funnel-transformer: Filtering out sequential redundancy for efficient language processing," *Advances in neural information processing systems*, vol. 33, pp. 4271–4282, 2020.
- [16] V. Nagaraja, Y. Shi, G. Venkatesh, O. Kalinli, M. L. Seltzer, and V. Chandra, "Collaborative training of acoustic encoders for speech recognition," *arXiv preprint arXiv:2106.08960*, 2021.
- [17] Y. Shi, V. Nagaraja, C. Wu, J. Mahadeokar, D. Le, R. Prabhavalkar, A. Xiao, C.-F. Yeh, J. Chan, C. Fuegen *et al.*, "Dynamic encoder transducer: A flexible solution for trading off accuracy for latency," *arXiv preprint arXiv:2104.02176*, 2021.
- [18] R. V. Swaminathan, B. King, G. P. Strimel, J. Droppo, and A. Mouchtaris, "Codert: Distilling encoder representations with co-learning for transducer-based speech recognition," *arXiv preprint arXiv:2106.07734*, 2021.
- [19] Z. Wu, D. Zhao, Q. Liang, J. Yu, A. Gulati, and R. Pang, "Dynamic sparsity neural networks for automatic speech recognition," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6014–6018.
- [20] H. Yang, Y. Shangguan, D. Wang, M. Li, P. Chuang, X. Zhang, G. Venkatesh, O. Kalinli, and V. Chandra, "Omni-sparsity dnn: Fast sparsity optimization for on-device streaming e2e asr via supernet," *arXiv preprint arXiv:2110.08352*, 2021.
- [21] 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," *Proc. Interspeech 2020*, pp. 5036–5040, 2020.
- [22] A. Narayanan, T. N. Sainath, R. Pang *et al.*, "Cascaded encoders for unifying streaming and non-streaming ASR," in *Proc. ICASSP*, 2021.
- [23] R. Botros, T. Sainath, R. David, E. Guzman, W. Li, and Y. He, "Tied & reduced rnn-t decoder," in *Proc. Interspeech*, 2021.
- [24] R. Prabhavalkar, T. N. Sainath, Y. Wu, P. Nguyen, Z. Chen, C.-C. Chiu, and A. Kannan, "Minimum word error rate training for attention-based sequence-to-sequence models," in *ICASSP*, 2018.
- [25] A. Narayanan, R. Prabhavalkar, C.-C. Chiu, D. Rybach, T. N. Sainath, and T. Strohmaier, "Recognizing long-form speech using streaming end-to-end models," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 920–927.
- [26] 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 *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6059–6063.
- [27] H. Liao, E. McDermott, and A. Senior, "Large Scale Deep Neural Network Acoustic Modeling with Semi-supervised Training Data for YouTube Video Transcription," in *Proc. ASRU*, 2013.
- [28] Google, "Artificial Intelligence at Google: Our Principles." [Online]. Available: <https://ai.google/principles/>
- [29] C. Kim, A. Misra, K. Chin *et al.*, "Generation of Large-Scale Simulated Utterances in Virtual Rooms to Train Deep-Neural Networks for Far-Field Speech Recognition in Google Home," in *Proc. Interspeech*, 2017.
- [30] D. S. Park, W. Chan, Y. Zhang, C. Chiu, B. Zoph, E. Cubuk, and Q. Le, "SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition," in *Proc. Interspeech*, 2019.