arXiv:2309.12963v1 [eess.AS] 22 Sep 2023

Massive End-to-end Models for Short Search Queries

Weiran Wang Rohit Prabhavalkar Dongseong Hwang Qiujia Li
Khe Chai Sim Bo Li James Qin Xingyu Cai Adam Stooke
Zhong Meng CJ Zheng Yanzhang He Tara Sainath
Pedro Moreno Mengibar
Google LLC
{weiranwang,prabhavalkar,dongseong}@google.com

September 29, 2023

Abstract

In this work, we investigate two popular end-to-end automatic speech recognition (ASR) models, namely Connectionist Temporal Classification (CTC) and RNN-Transducer (RNN-T), for offline recognition of voice search queries, with up to 2B model parameters. The encoders of our models use the neural architecture of Google's universal speech model (USM), with additional funnel pooling layers to significantly reduce the frame rate and speed up training and inference. We perform extensive studies on vocabulary size, time reduction strategy, and its generalization performance on long-form test sets. Despite the speculation that, as the model size increases, CTC can be as good as RNN-T which builds label dependency into the prediction, we observe that a 900M RNN-T clearly outperforms a 1.8B CTC and is more tolerant to severe time reduction, although the WER gap can be largely removed by LM shallow fusion.

1 Introduction

There has been a recent focus on scaling up end-to-end ASR models, such as connectionist temporal classification (CTC) [1] and neural transducers (RNN-T, [2]), to extremely large sizes [3, 4, 5, 6, 7] to explore the benefits of such models at scale. Many of these models are motivated by the goal of training a single multilingual ASR model which can perform well across a range of languages. For example, the Universal Speech Model (USM, [6]) is a 2B parameter full-context CTC model trained on YouTube data from 300+ languages; Pratap et al. [7] train ASR systems on 1000+ languages. The present work follows in the footsteps of these previous works. As large language models (LLMs) have rapidly risen in popularity in the field of natural language processing [8, 9, 10], it is natural to ask how ASR models operate at large sizes. Specifically, how do the two main model classes – CTC and RNN-T – compare at large model sizes? In particular, how do these models compare on a task (voice search, in this work) where paired training data is plentiful, for a single language (English) instead of in a multilingual setup.

A related, and important practical question relates to the difficulty of training massive end-to-end models as model size increases. Since the cost of computing the outputs of the encoders in

the model (most of the model parameters are typically in the encoder) scales poorly in attention-based encoders such as transformers [11] and conformers [12], which can be prohibitive in terms of computation and memory. The RNN-T model, in particular, requires additional computation and memory for the prediction and joint network, which further compounds the problem.

In this work, we investigate the aforementioned issues and find that the two issues – how CTC and RNN-T compare, and how one can train them efficiently – are not as independent as might seem initially: we find that encoder output frame rate reduction can be applied repeatedly at multiple layers in the encoder to obtain a large output frame rate reduction; this in turn is critical for training models efficiently. While past research incorporated time reduction to lower the final frame rate to 60ms for voice search [13, 14], we find that we can increase time reduction to 6x with large CTC, and all the way up to 16x with large RNN-T, of the 40ms base frame rate with funnel pooling [15], without sacrificing much accuracy.

Additionally, we compare CTC and RNN-T as the backbone ASR system. CTC produces label probabilities that are independent of previous time frames, and is thus much faster at inference compared to RNN-T. This is one reason why CTC was preferred for USM in [6]. However, we find that RNN-T is actually more accurate than CTC, and specifically a 900M RNN-T clearly outperforms a 1.8B CTC, even though the CTC model benefits significantly from increased model size compared with a 340M CTC. Moreover, we find that RNN-T can tolerate a much larger time reduction factor. In order to address any concerns with RNN-T being slower than CTC, we run both models (including beam search) on a TPU to achieve the best possible latency.

Finally, with the much smaller number of output frames and on-TPU beam search, shallow fusion [16, 17] with a neural LM becomes a viable candidate for enhancing model accuracy. We fuse both CTC and RNN-T with external LMs trained with large amounts of text, and observe that the word error rate (WER) gap between them can be largely removed by LM fusion. Our study on shallow fusion is timely, as the community starts to become interested in the use of LLMs for ASR. We show that shallow fusion remains an effective technique for LM integration, even for large models trained with hundreds of thousands of hours of audio, and should be considered as a baseline for more advanced techniques.

2 End-to-end ASR Models

In this section, we briefly describe the end-to-end models employed in this work – connectionist temporal classification (CTC) [1], and neural transducers [2] – particularly, the hybrid autoregressive transducer (HAT) [18]. Since end-to-end models are now part of the mainstream in automatic speech recognition (ASR), we opt for brevity; interested readers can find more information in recent overview articles [19, 20].

Notations We assume that the input audio signal has been parameterized into suitable acoustic feature vectors: $\mathbf{x} = [\mathbf{x}_1, \ldots, \mathbf{x}_{T'}]$, of length T', where $\mathbf{x}_t \in \mathbb{R}^d$ (128-dimensional log-mel features, in this work). The input acoustic features are processed using an *encoder*, a suitable neural network (a Conformer [12], in this work), which transforms the input into a higher-level representation: $\mathbf{h} = [\mathbf{h}_1, \ldots, \mathbf{h}_T]$, where the length of the encoded representation is typically shorter than the original input length ($T \leq T'$). We assume that we have an input transcript corresponding to each utterance: $\mathbf{y} = [y_0 = \langle \mathrm{s} \rangle, y_1, \ldots, y_U]$, where each $y_u \in \mathcal{V}$, the set of output symbols (word-pieces [21], in this work), and $\langle \mathrm{s} \rangle$ represents a special start-of-sentence symbol.

2.1 Connectionist Temporal Classification

Connectionist Temporal Classification (CTC) was introduced by Graves et al. [1], as a way to train sequence-to-sequence models which can transduce the input sequence $\mathbf{x}$, into the output sequence $\mathbf{y}$ when the alignments between the two sequences are unknown. CTC accomplishes this by modeling the conditional distribution, $P(\mathbf{y}|\mathbf{x})$, by marginalizing over all possible alignment paths between the two sequences:

$$P(\mathbf{y}|\mathbf{x}) = P(\mathbf{y}|\mathbf{h}(\mathbf{x})) = \sum_{\mathbf{a}\in \mathrm{B}_{\mathrm{ctc}}(\mathbf{y})} \prod_{t=1}^{T} P(a_t|\mathbf{h}) \tag{1}$$

where, the $\mathrm{B}_{\mathrm{ctc}}(\mathbf{y})$ corresponds to the set of all valid alignments. Specifically, an alignment $\mathbf{a} \in \mathrm{B}_{\mathrm{ctc}}(\mathbf{y})$ is a valid alignment if it contains $|\mathbf{h}(\mathbf{x})| = T$ symbols from the set of outputs augmented with a special *blank* symbol – $\mathcal{V} \cup \{\langle \mathrm{b} \rangle\}$; and if additionally, removing consecutive repeated non-blank symbols and then removing all $\langle \mathrm{b} \rangle$ symbols produces the original label sequence $\mathbf{y}$. As can be seen in (1), CTC models make a strong conditional independence assumption – that the output labels are conditionally independent given the input acoustic encoded features. However, these models work well in practice [22, 23], especially with large encoders [6]. Finally, as can be seen in (1), the CTC model produces one output symbol (blank, or non-blank) per encoder time step.

2.2 Neural Transducers

The recurrent neural network transducer (RNN-T) was proposed by Graves et al. [2], as an improvement over the CTC model, to reduce conditional independence assumptions in the model. This is achieved by introducing a separate *prediction network*, which models label dependency (and thus makes model outputs conditionally dependent on the sequence of previous predictions).

$$P(\mathbf{y}|\mathbf{x}) = P(\mathbf{y}|\mathbf{h}(\mathbf{x})) = \sum_{\mathbf{a}\in \mathrm{B}_{\mathrm{nt}}(\mathbf{y})} \prod_{\tau=1}^{T+U} P(a_\tau|\mathbf{q}_{i_\tau}, \mathbf{h}_{\tau-i_\tau})$$

where, $\mathrm{B}_{\mathrm{nt}}(\mathbf{y})$ is the set of all valid alignments – sequences of length $T+U$, which are identical to $\mathbf{y}$ after removing all blanks (i.e., each sequence has exactly T blank symbols); i_τ represents the number of non-blank labels in the partial alignment $a_1, \ldots, a_{\tau-1}$; and, $\mathbf{q}_j$ represents the output of the prediction network after processing the first $j-1$ output labels: $\mathbf{q}_j = \mathrm{PredNetwork}(y_{j-1}, \ldots, y_0)$. While early works employed recurrent LSTM networks to model the prediction network, in this work, the prediction network is modeled as $|V|^2$ embedding network [24], whose output only depends on the last two non-blank labels y_{j-1}, and y_{j-2}. Notably, neural transducers allow the model to output multiple (non-blank) output labels at each frame; blanks correspond to transitions to the next input encoder frame.

2.2.1 Hybrid Autoregressive Transducer

The hybrid autoregressive transducer (HAT) model proposed by Variani et al. [18] improves over the basic neural transducer structure in two ways. First, it factorizes the output distribution $P(a_\tau|\mathbf{q}_{i_\tau}, \mathbf{h}_{\tau-i_\tau})$ into two separate distributions: a Bernoulli distribution to model blank symbols (i.e., if $a_\tau = \langle \mathrm{b} \rangle$), and a separate probability distribution for the non-blank labels. Second, the model proposes the notion of an internal language model (ILM), $P_{\mathrm{ILM}}(\mathbf{y})$, which represents the LM learned by the model based on the training data; the ILM can thus be subtracted out from the posterior distribution, before fusion with an external LM, $P_{\mathrm{EXT}}(\mathbf{y})$, during decoding, by adding

tunable hyperparameters α, and β which can be set based on a development set to produce the most likely hypothesis [25, 26, 27, 28]:

$$\mathbf{y}^* = \arg\max_{\mathbf{y}} \log P(\mathbf{y}|\mathbf{x}) - \alpha \log P_{\text{ILM}}(\mathbf{y}) + \beta \log P_{\text{EXT}}(\mathbf{y})$$

3 Model architecture

3.1 USM Conformer

The USM uses the convolution-augmented Transformer [12], or Conformer architecture. Each Conformer block consists of a feed-forward module (FFN), a multi-head self-attention module (MHSA), a convolution module (CONV), and a second feed-forward module. If the input to l-th Conformer block is $\mathbf{x}^{(l)}$, then its output (and, the input to the next block), $\mathbf{x}^{(l+1)}$, is computed as:

$$\tilde{\mathbf{x}}^{(l)} = \mathbf{x}^{(l)} + \frac{1}{2}\text{FFN}_1\big(\mathbf{x}^{(l)}\big)$$

$$\begin{aligned}
\mathbf{x}'^{(l)} &= \tilde{\mathbf{x}}^{(l)} + \text{MHSA}\big(W_q\tilde{\mathbf{x}}^{(l)}, W_k\tilde{\mathbf{x}}^{(l)}, W_v\tilde{\mathbf{x}}^{(l)}\big) \\
\mathbf{x}''^{(l)} &= \mathbf{x}'^{(l)} + \textsc{Conv}\big(\mathbf{x}'^{(l)}\big) \\
\mathbf{x}^{(l+1)} &= \textsc{LayerNorm}\big(\mathbf{x}''^{(l)} + \frac{1}{2}\text{FFN}_2(\mathbf{x}''^{(l)})\big)
\end{aligned} \tag{2}$$

For the standard conformer layer, the input length and output length are the same.

3.2 Funnel Pooling

We improve training and inference speed in our models by reducing the encoder embedding sequence length relative to the audio length. At selected Conformer blocks within the encoder, we employ the pooling technique in the MHSA module introduced in Funnel-Transformers [15]. At a funnel self-attention layer, the entire input is used to produce the key and value sequences as usual, but strided-average pooling is applied to the time dimension of the input used in producing the query vectors. That is, we replace the MHSA module in (2) by (3) below.

$$\begin{aligned}
\hat{\mathbf{x}}^{(l)} &= \textsc{StridedPooling}(\tilde{\mathbf{x}}^{(l)}), \\
\mathbf{x}'^{(l)} &= \hat{\mathbf{x}}^{(l)} + \text{MHSA}\big(W_q\hat{\mathbf{x}}^{(l)}, W_k\tilde{\mathbf{x}}^{(l)}, W_v\tilde{\mathbf{x}}^{(l)}\big).
\end{aligned} \tag{3}$$

The stride of pooling in (3) is the time reduction factor in this layer. In a language model, this pool-query-only method was shown to provide a slight advantage over simply pooling the entire hidden embedding sequence between layers [15]; likely, the higher granularity in the key and value sequences allows the network to learn a less lossy compression. Computational savings accrue in all subsequent Conformer layers according to the $O(T^2D)$ complexity for self-attention, and linearly for feed-forward networks.

Related work on time reduction Time reduction has been a useful technique for achieving a balance between input and output lengths for ASR, and has evolved over time with the choice of neural architecture and modeling units. In [29, 30], the lower frame rates were achieved by concatenating input frames and striding, for DNN or LSTM models predicting context-dependent

HMM states. [31, 32, 33] proposed to use hierarchical/pyramidal RNNs where outputs of consecutive steps are combined before feeding to the next layer. After the community switched to end-to-end systems and word-piece type modeling units, a popular frame rate is 40ms as achieved by convolutional subsampling, and is adopted by widely-used open-source libraries [34]. With the Conformer architecture [12], previous work on voice search mostly had a final frame rate of 60ms, achieved by a 30ms input frame rate and a 2x time reduction layer early in the encoder, with stacking [13] or funnel pooling [14]. With regards to the loss function, recent work of [35] proposed to use CTC to (irregularly) select encoder output frames for RNN-T modeling, effectively reducing the frame rate for the decoder. However, such a CTC-based encoder output selection does not save any computations in the encoder, but only reduces the computation in downstream modules. Our goal in this work is to aggressively reduce the sequence length early in the architecture for computational savings.

4 Experiments

4.1 Datasets

Our experiments focus on short voice search queries. For a majority of the experiments, we use 520M utterances of voice search queries for training; the total amount of audio is 490K hours and the average duration per utterance is 3.4 seconds. A small percentage of the training data is human transcribed while the rest is pseudo-labeled by a 600M bidirectional RNN-T teacher [36]. We tokenize training transcripts with word-piece models [21].

We use both real audio and TTS-generated data for evaluation. The real audio utterances are representative of typical voice search traffic, with an average duration of 3.9 seconds. Our development set consists of 9K real audio utterances (denoted as VS-dev); we use a separate held-out test set consisting of 5K utterances for testing (denoted as VS-test). The TTS sets contain rare proper nouns (RPN) which appear fewer than 5 times in the training set, and they are good testbeds for external LM integration. Each TTS set contains 10K utterances and covers one of five domains: Maps (denoted as RPNM), News (RPNN), Play (RPNP), Search query logs (RPNS), and Youtube (RPNY); they have average durations of 5.9, 10.1, 5.3, 5.4, and 5.8 seconds respectively. We use VS-dev, RPNM and RPNN for tuning the model architecture and other hyperparameters and report final WERs on the rest sets.

For training LMs for fusion, each minibatch is sampled 50/50 from the transcripts of acoustic training data, and text-only data which contains 50B utterances. The text-only data contains textual search queries from the domains of Maps, Textual search query logs, News, Play, and Youtube, and a frequency-based pruning strategy, designed to improve rare word modeling, is implemented to adjust the probability of selecting each query [37]. We train transformer LMs of 128M and 1B parameters (not counting parameters in the final softmax) with wordpieces compatible with those of E2E models.

4.2 Model Architectures

We use the 128-dimensional log Mel-filterbank energies (extracted from 32ms window and 10ms shift) as the frontend features. After two 2D-convolution layers, both with strides (2,2), the resulting feature sequence has a frame rate of 40ms and becomes the input to our Conformer architecture. This architecture mimics that of Google's universal speech model (USM, [6]). The number of attention heads used in Conformer blocks is 8, and the intermediate dimension of the FFNs is 4

times the model dimension. The Conformer blocks use local self-attention with a large attention span, and the encoder output has a large enough receptive field to cover the entire utterance. As observed by [3], unsupervised pre-training does not help with ASR accuracy when using a large amount of supervised audio data, and thus we skip the pre-training step in this work (except for including a baseline CTC model with pretraining in Table 1).

We explore two different encoder configurations for CTC, with different sizes: the smaller encoder configuration consists of 24 Conformer layers of dimension 768, leading to a total of 340M parameters; the larger encoder has 32 Conformer layers of dimension 1536, leading to a total of 1.8B parameters. For RNN-T, the encoder consists of 16 Conformer layers of dimension 1536, resulting in a total of 870M parameters. We use a $|V|^2$ embedding decoder [24], i.e., the prediction network computes LM features based on two previous non-blank tokens, which was shown to work well on voice search data. The model output uses the HAT factorization [18] which was shown to benefit external LM integration.

Deviating from the USM, we perform significant time reduction in the encoder architecture. Following [14], we initially start funnel pooling at layer index 4 (zero-based), and apply pooling in subsequent layers to achieve the desired factor of time reduction. As an example, if we perform pooling at layers with (zero-based) indices 4 and 5, each with the reduction factor 2, we achieve a total reduction factor 4 for layers 6 and onwards, i.e., the sequence length after layer 6 is $1/4$ of the original input length and the frame rate at the encoder output is 160ms, which is also the frame rate at which the decoder operates. Note that [14] used funnel pooling for on-device modeling with stringent latency requirements, whereas here we use funnel pooling for the full-context model to speed up training and inference. We observe that the ASR accuracy turns out to be more tolerant of time reduction. In Sec 4.5, we investigate the effect of starting funnel pooling at earlier layers to further reduce inference costs.

Each model is trained with the Adafactor optimizer [38] with a batch size of 4096 utterances. We train CTC models to 500K steps and RNN-T models to 300K steps, by which point their WERs on development sets have stabilized. In a pilot study, we have conducted discriminative training with the minimum word error rate objective [39]. We achieved no WER gain on VS sets and 3% relative improvement on the RPN sets. We only present results with maximum likelihood training in this work, and it is future work to carefully study discriminative training for large models.

For CTC, we perform greedy decoding when not using external LM, in which case beam search does not provide additional WER gain. We do perform prefix beam search [40] with a beam size of 8 when fusing with external LM. Two types of prior probabilities were used for CTC shallow fusion: the uniform prior over non-blanks as implemented by the blank probability downscaling technique [41], versus the uni-gram prior based on the model posteriors on training set [42]. We perform label synchronous beam search (with path merging) for RNN-T with a beam size of 8, and perform internal LM score subtraction in the case of shallow fusion.

4.3 CTC Results

Intuitively, with a larger vocabulary, the label sequences are shorter and we can afford more time reduction. The hard constraint for CTC is that, since it emits only one token (blank or non-blank) at each encoder output frame, the encoder output sequence length must remain longer than the label sequence length (RNN-T however is not subject to this constraint). During CTC training, we discard utterances that violate this constraint, although such cases are uncommon in the VS sets. For example, for the 16K vocabulary size, at the final frame rate of 320ms (8x reduction), roughly 1% utterances violate this constraint as estimated on VS-dev.

Table 1: WERs (%) on dev sets by CTC, with different architectures and vocabulary sizes. Funnel pooling is employed at layers 4 and 5, with factors 2,2 for 160ms frame rate and 3,2 for a 240 frame rate.

Model size (frame rate)	VS-dev	RPNM	RPNN
baselines: no pooling, vocab size=4096			
340M (40ms)	4.7	15.1	12.4
1.8B (40ms), w. pretraining	4.2	14.2	**10.4**
with funnel pooling, vocab size=16384			
340M (160ms)	4.5	15.3	16.6
340M (240ms)	4.5	14.9	21.6
1.8B (160ms)	4.3	**13.7**	12.7
1.8B (240ms)	**4.2**	13.8	17.3
1.8B (320ms)	5.0	14.0	26.4
1.8B (160ms) + 128M LM fusion			
blank downscaling [41]	3.8	10.7	11.0
model-based prior [42]	3.8	10.5	10.8
1.8B (160ms) + 1B LM fusion			
blank downscaling [41]	3.8	10.1	10.4
model-based prior [42]	**3.8**	**9.8**	**10.1**

We report WERs of a selection of CTC models in Table 1. As baselines, we report the WERs with 4K vocabulary at a 40ms frame rate, a setup closely following that of the USM architecture [6]. We match these baselines on VS with the 16K vocabulary size and much lower frame rates of 160ms and 240ms: for 340M CTC, we obtain an improvement from 4.7% to 4.5% on VS-dev, while for 1.8B CTC, we match the 4.2% without pretraining. For supervised training with 1.8B CTC, we observe a 4x speedup in training time with a 160ms frame rate, and a 4.5x speedup with a 240ms frame rate compared to the 40ms model (despite the use of a more costly softmax operation due to larger vocabulary). However, we do observe that with a 320ms frame rate, the WER on VS-dev significantly degrades to 5.0%, and this could not be alleviated by increasing the vocabulary size to 32K. This suggests that a too coarse time resolution does not work well with the CTC loss and its underlying independence assumptions.

The trend of WER on RPNM is similar to that of VS-dev. However, we do observe worse degradation on RPNN as we apply heavier time reduction. As mentioned in Sec 4.1, the transcriptions of RPNN come from the News domain which has different linguistic characteristics from voice search, and the audio length is quite longer (10 secs on average) than the VS set (4 secs on average). We hypothesize that models with lower frame rates may not generalize well to unseen audio length, and further investigate this issue in Sec 4.6.

Given that the recent work on large models is based on CTC [5, 6], one may speculate that as the bidirectional CTC model gets larger, and with large amounts of training data, the underlying modeling assumption of CTC holds approximately and it can achieve state-of-the-art accuracy by itself. We challenge this speculation by performing LM shallow fusion to the 1.8B CTC model with a 160ms frame rate. We observe a significant WER reduction from 4.2% to $3.7 - 3.8\%$ with a smaller 128M LM already on the in-domain VS-dev set. The results we present in Table 1 uses LM weights that achieve a good balance between VS-dev and RPN sets. Had we focused only on VS,

Table 2: Dev set WERs (%) by RNN-T with various vocabulary sizes and frame rates. The encoder contains 870M parameters. Funnel pooling starts at layer 4 and applies to adjacent layers. We use the notation 3x2 to indicate that layer 4 has a reduction factor of 3, and layer 5 has a reduction factor of 2, which yield the final frame rate of 240ms.

Frame rate (reduction factors)	VS-dev	RPNM	RPNN
vocab size=4096, decoder size=10M			
240ms (3x2)	3.8	12.6	**12.1**
320ms (2x2x2)	3.8	12.9	14.2
vocab size=16384, decoder size=33M			
160ms (2x2)	**3.7**	**12.3**	12.8
240ms (3x2)	3.8	12.5	12.2
320ms (2x2x2)	3.8	12.5	13.6
400ms (5x2)	3.8	12.7	15.6
480ms (3x2x2)	3.9	12.6	19.5
640ms (2x2x2x2)	3.9	12.9	28.7
vocab size=32768, decoder size=65M			
320ms (2x2x2)	3.8	12.9	14.8
640ms (2x2x2x2)	3.9	12.7	25.9
vocab size=16384, 128M LM fusion			
160ms (2x2)	3.8	**11.0**	32.6

the best we could achieve on VS-dev is 3.7%, with clearly worse WERs on RPN sets. Among the two prior estimation methods, model-based prior [42] consistently outperforms blank probability downscaling [41] over different WER operating points. When increasing the external LM size to 1B, we could not further improve on VS-dev but achieved sizeable gains on rare word sets.

4.4 RNN-T Results

We conduct a similar set of time reduction experiments with RNN-T and the results are shown in Table 2. Overall RNN-T is quite robust to different vocabulary sizes. With the 16K vocabulary, we studied the most time reduction settings, and we observed only small WER degradations from 160ms frame rate all the way to 640ms frame rate on VS-dev and RPNM. Like CTC, the performance on the "out-of-domain" dataset RPNN degrades as the frame rate reduces. Comparing the models at the same frame rate of 160ms and vocabulary size 16K, RNN-T has a WER of 3.7% on VS-dev, outperforming the 4.2% by CTC by a large margin, and achieving parity with CTC + shallow fusion on this set. This demonstrates the benefit of having a learnable LM feature encoder for modeling label dependency in end-to-end ASR.

When fusing the best RNN-T model with a 128M LM, we observe interestingly that it tends to significantly degrade RPNN (with heavy deletion errors) and the WERs are quite sensitive to internal and external LM weights. We list one set of results in Table 2 (bottom panel) which achieves a good balance between VS-dev and RPNM, yet it degrades RPNN from 12.8% to 32.6%. We hypothesize this is due to the bias of the internal language model of HAT learned purely on short utterances. In Sec 4.6, we provide further evidence for this, by demonstrating that training on length-diverse data improves the WER and robustness to shallow fusion on RPNN.

Table 3: Dev set WERs (%) by 900M RNN-T with a vocab size of 16K and a 240ms frame rate (3x2 reduction), with funnel pooling started in earlier Conformer layers. The first row is taken from Table 2.

Start layer index	VS-dev	RPNM	RPNN
4	3.8	12.5	12.2
3	3.7	12.6	12.0
2	3.7	**12.4**	**11.7**
1	3.7	12.6	12.3

Table 4: Dev set WERs (%) by 900M RNN-T with 16K vocabulary and two frame rates 240ms and 640ms trained on multi-domain data. Funnel pooling starts from layer 4. First two rows are taken from Table 2.

Frame rate	VS-dev	RPNM	RPNN
Voice search training			
240ms	3.8	12.5	12.2
640ms	3.9	12.9	28.7
Multi-domain training			
240ms	**3.6**	13.6	**6.9**
640ms	3.8	12.6	9.5
+ 128M LM fusion			
240ms	3.7	**11.5**	7.1
Multi-domain 600M teacher, vocab size=4096			
60ms	4.0	**11.5**	7.1

4.5 Location of Funnel Pooling

So far, we have started funnel pooling from layer 4, following prior work [14] that worked on smaller models with hard latency constraints. In this section, we start pooling earlier in the architecture, which leads to more efficient inference.

We use the 240ms frame rate RNN-T model from Table 2 as the baseline, with time reduction factors of 3 and 2 in two consecutive layers, and change the pooling start layer index to 3, 2, and 1. Training of the model with pooling starting at layer 0 diverged with the same learning parameters and we do not report its performance. The results of these models on development sets are reported in Table 3, which shows that on top of the 40ms base frame rate, the model is quite robust to the location of pooling layers, and in fact we obtain small WER gains on RPN sets by starting pooling at layer 2. We plan to replace the 2D convolutional subsampling layers before the Conformer layer with funnel pooling layers to explore more possibilities for frame reduction.

4.6 Adding Long-Form Training Data

To verify that the poor performance of end-to-end models on RPNN was due to the lack of longer training audio, we repeat several RNN-T experiments with additional multi-domain training data [43]. Most notably, we include segmented YT audio data containing 520M utterances with an average duration of 9.8 seconds, giving us a total of 600K hours of longer-form training data.

The comparisons between RNN-T models trained on voice search data and multi-domain data

Table 5: Test set WERs (%) by CTC and RNN-T, with 16K vocabulary. Funnel pooling starts from layer 4. By default, we use VS training data.

VS-test	RPNP	RPNS	RPNY
1.8B CTC, 160ms frame rate			
4.9	39.8	23.1	26.0
+ 128M LM fusion with model prior			
4.5	34.1	17.0	20.8
900M RNN-T, 240ms frame rate			
4.5	37.8	20.6	23.3
+ Multi-domain training data			
4.4	36.4	19.9	22.3
+ 128M LM fusion with ILM			
4.4	**33.9**	**16.9**	**20.2**
Multi-domain 120M RNN-T [14] 60ms frame rate, 0.9s look-ahead			
5.0	35.9	19.2	23.2

are presented in Table 4. For two vastly different frame rates 240ms and 640ms, the additional long-form training data improves the WERs on both VS-dev and RPNN by a lot. With only voice search training data, the WER gap on RPNN between the two frame rates used to be very large (12.2% vs 28.7%), whereas with multi-domain training, the gap is significantly reduced (6.9% vs 9.5%) with much improved absolute WERs. When performing shallow fusion for the multi-domain 240ms frame rate model with the same 128M LM used in Sec 4.4, we achieve a better balance between in-domain and out-of-domain test sets, improving RPNM from 13.6% to 11.5% and without affecting RPNN much. .

It is interesting to add the bidirectional RNN-T teacher [36] into comparison, as shown in the bottom panel of Table 4. The 600M teacher is trained on an earlier version of multi-domain data, where the voice search portion is smaller but fully hand-transcribed, and is used to generate the pseudo-labels of voice search data used in this paper. Note that we are surpassing or on par with the teacher's performance on VS and RPNN, probably due to a larger amount of training data and more consistent labeling. The teacher does outperform new models on RPNM, probably because, as an end-to-end model, it predicts even fewer rare words in its pseudo-labels than human transcriptions, for the student to imitate.

4.7 Final Evaluation

Finally, we compare the best configurations from both CTC and RNN-T on the test sets, namely VS-test, RPNP, RPNS, and RPNY. Taking into account both in-domain and out-of-domain performance, we choose the 1.8B CTC model with 160ms frame rate, and the 900M RNN-T model with 240ms frame rate. The results are shown in Table 5. Relative merits between methods are consistent with the performance on development sets. That is, with the same voice search training data, the 900M RNN-T outperforms the 1.8B CTC on all test sets, and by 8.0% relative margin for the in-domain VS-test set (4.5% vs 4.9%). For more stable shallow fusion performance on out-of-domain test sets, it is beneficial to train RNN-T on multi-domain data, and we expect the same for

CTC. As a reference, we provide results of another small RNN-T model, with a 60ms frame rate and a limited right context of 0.9s, from previous work [14].

5 Conclusions

We have compared two major end-to-end ASR models, CTC and RNN-T, on a large-scale voice search task. We have verified that RNN-T is clearly more accurate than CTC for in-domain test data even with a smaller model size, although the gap can be largely removed with LM fusion which compensates for the label independence assumption underlying CTC. We have also observed that, for large models, time reduction is effective at reducing inference and training costs without sacrificing accuracy, and is thus useful for offline transcription tasks. In the future, we will further optimize the model architecture and extend the usage to much longer audio.

References

[1] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *ICML*, 2006.

[2] Alex Graves, "Sequence transduction with recurrent neural networks," in *ICML Workshop on Representation Learning*, 2012.

[3] Yu Zhang, Daniel S Park, Wei Han, James Qin, Anmol Gulati, Joel Shor, Aren Jansen, Yuanzhong Xu, Yanping Huang, Shibo Wang, et al., "Bigssl: Exploring the frontier of large-scale semi-supervised learning for automatic speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1519–1532, 2022.

[4] Bo Li, Ruoming Pang, Yu Zhang, Tara N. Sainath, Trevor Strohman, Parisa Haghani, Yun Zhu, Brian Farris, Neeraj Gaur, and Manasa Prasad, "Massively multilingual asr: A lifelong learning solution," in *ICASSP*, 2022.

[5] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever, "Robust speech recognition via large-scale weak supervision," in *ICML*, 2023.

[6] Yu Zhang, Wei Han, James Qin, Yongqiang Wang, Ankur Bapna, Zhehuai Chen, Nanxin Chen, Bo Li, Vera Axelrod, Gary Wang, et al., "Google USM: Scaling automatic speech recognition beyond 100 languages," *arXiv preprint arXiv:2303.01037*, 2023.

[7] Vineel Pratap, Andros Tjandra, Bowen Shi, Paden Tomasello, Arun Babu, Sayani Kundu, Ali Elkahky, Zhaoheng Ni, Apoorv Vyas, Maryam Fazel-Zarandi, et al., "Scaling speech technology to 1,000+ languages," *arXiv preprint arXiv:2305.13516*, 2023.

[8] OpenAI, "GPT-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.

[9] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample, "LLaMA: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.

[10] Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, et al., "PaLM 2 technical report," *arXiv preprint arXiv:2305.10403*, 2023.

[11] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin, "Attention is all you need," in *NeurIPS*, 2017.

[12] Anmol Gulati, James Qin, Chung-Cheng Chiu, et al., "Conformer: Convolution-augmented transformer for speech recognition," in *Interspeech*, 2020.

[13] Yanzhang He, Tara N. Sainath, Rohit Prabhavalkar, Ian McGraw, Raziel Alvarez, Ding Zhao, David Rybach, Anjuli Kannan, Yonghui Wu, Ruoming Pang, Qiao Liang, Deepti Bhatia, Yuan Shangguan, Bo Li, Golan Pundak, Khe Chai Sim, Tom Bagby, Shuo-yiin Chang, Kanishka Rao, and Alexander Gruenstein, "Streaming end-to-end speech recognition for mobile devices," in *ICASSP*, 2019.

[14] Shaojin Ding, Weiran Wang, Ding Zhao, Tara N. Sainath, Yanzhang He, Robert David, Rami Botros, Xin Wang, Rina Panigrahy, Qiao Liang, Dongseong Hwang, Ian McGraw, Rohit Prabhavalkar, and Trevor Strohman, "A unified cascaded encoder ASR model for dynamic model sizes," in *Interspeech*, 2022.

[15] Zihang Dai, Guokun Lai, Yiming Yang, and Quoc Le, "Funnel-Transformer: Filtering out sequential redundancy for efficient language processing," in *NeurIPS*, 2020.

[16] Caglar Gulcehre, Orhan Firat, Kelvin Xu, Kyunghyun Cho, Loic Barrault, Huei-Chi Lin, Fethi Bougares, Holger Schwenk, and Yoshua Bengio, "On using monolingual corpora in neural machine translation," *arXiv:1503.03535*, 2015.

[17] Jan Chorowski and Navdeep Jaitly, "Towards better decoding and language model integration in sequence to sequence models," in *Interspeech*, 2017.

[18] Ehsan Variani, David Rybach, Cyril Allauzen, and Michael Riley, "Hybrid autoregressive transducer (HAT)," in *ICASSP*, 2020.

[19] Jinyu Li et al., "Recent advances in end-to-end automatic speech recognition," *APSIPA Transactions on Signal and Information Processing*, vol. 11, no. 1, 2022.

[20] Rohit Prabhavalkar, Takaaki Hori, Tara N Sainath, Ralf Schlüter, and Shinji Watanabe, "End-to-end speech recognition: A survey," *arXiv preprint arXiv:2303.03329*, 2023.

[21] Mike Schuster and Kaisuke Nakajima, "Japanese and korean voice search," in *ICASSP*, 2012.

[22] Yajie Miao, Mohammad Gowayyed, and Florian Metze, "EESEN: End-to-end speech recognition using deep RNN models and WFST-based decoding," in *ASRU*, 2015.

[23] Shigeki Karita, Nelson Enrique Yalta Soplin, Shinji Watanabe, Marc Delcroix, Atsunori Ogawa, and Tomohiro Nakatani, "Improving Transformer-based end-to-end speech recognition with connectionist temporal classification and language model integration," in *Interspeech*, 2019.

[24] Rami Botros, Tara N. Sainath, Robert David, Emmanuel Guzman, Wei Li, and Yanzhang He, "Tied & reduced RNN-T decoder," in *Interspeech*, 2021.

[25] Nelson Morgan and Hervé Bourlard, "Continuous speech recognition using multilayer perceptrons with hidden Markov models," in *ICASSP*, 1990.

[26] Ehsan Variani, Erik McDermott, and Georg Heigold, "A Gaussian mixture model layer jointly optimized with discriminative features within a deep neural network architecture," in *ICASSP*, 2015.

[27] Naoyuki Kanda, Xugang Lu, and Hisashi Kawai, "Minimum Bayes risk training of CTC acoustic models in maximum a posteriori based decoding framework," in *ICASSP*, 2017.

[28] Erik McDermott, Hasim Sak, and Ehsan Variani, "A density ratio approach to language model fusion in end-to-end automatic speech recognition," in *ASRU*, 2019.

[29] Vincent Vanhoucke, Matthieu Devin, and Georg Heigold, "Multiframe deep neural networks for acoustic modeling," in *ICASSP*, 2013.

[30] Yajie Miao, Jinyu Li, Yongqiang Wang, Shi-Xiong Zhang, and Yifan Gong, "Simplifying long short-term memory acoustic models for fast training and decoding," in *ICASSP*, 2016.

[31] Salah Hihi and Yoshua Bengio, "Hierarchical recurrent neural networks for long-term dependencies," in *NeurIPS*, 1996.

[32] Jan Koutnik, Klaus Greff, Faustino Gomez, and Juergen Schmidhuber, "A clockwork RNN," in *ICML*, 2014.

[33] William Chan, Navdeep Jaitly, Quoc V. Le, and Oriol Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *ICASSP*, 2016.

[34] Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplin, Jahn Heymann, Matthew Wiesner, Nanxin Chen, Adithya Renduchintala, and Tsubasa Ochiai, "ESPnet: End-to-end speech processing toolkit," in *Interspeech*, 2018.

[35] Yongqiang Wang, Zhehuai Chen, Chengjian Zheng, Yu Zhang, Wei Han, and Parisa Haghani, "Accelerating RNN-T training and inference using CTC guidance," in *ICASSP*, 2023.

[36] Dongseong Hwang, Khe Chai Sim, Zhouyuan Huo, and Trevor Strohman, "Pseudo label is better than human label," in *Interspeech*, 2022.

[37] W. Ronny Huang, Cal Peyser, Tara N. Sainath, Ruoming Pang, Trevor Strohman, and Shankar Kumar, "Sentence-select: Large-scale language model data selection for rare-word speech recognition," in *Interspeech*, 2022.

[38] Noam Shazeer and Mitchell Stern, "Adafactor: Adaptive learning rates with sublinear memory cost," *arXiv preprint arXiv:1804.04235*, 2018.

[39] Rohit Prabhavalkar, Tara N. Sainath, Yonghui Wu, Patrick Nguyen, Zhifeng Chen, Chung-Cheng Chiu, and Anjuli Kannan, "Minimum word error rate training for attention-based sequence-to-sequence models," in *ICASSP*, 2018.

[40] Awni Y. Hannun, Andrew L. Maas, Daniel Jurafsky, and Andrew Y. Ng, "First-pass large vocabulary continuous speech recognition using bi-directional recurrent DNNs," *arXiv preprint arXiv:1408.2873*, 2014.

[41] Hasim Sak, Andrew W. Senior, Kanishka Rao, Ozan Irsoy, Alex Graves, Françoise Beaufays, and Johan Schalkwyk, "Learning acoustic frame labeling for speech recognition with recurrent neural networks," in *ICASSP*, 2015.

[42] Qiujia Li, Chao Zhang, and Philip C. Woodland, "Integrating source-channel and attention-based sequence-to-sequence models for speech recognition," in *ASRU*, 2019.

[43] Arun Narayanan, Rohit Prabhavalkar, Chung-Cheng Chiu, David Rybach, Tara Sainath, and Trevor Strohman, "Recognizing long-form speech using streaming end-to-end models," in *ASRU*, 2019.