

ON MODULAR TRAINING OF NEURAL ACOUSTICS-TO-WORD MODEL FOR LVCSR

Zhehuai Chen, Qi Liu, Hao Li and Kai Yu

Key Lab. of Shanghai Education Commission for Intelligent Interaction and Cognitive Engineering
SpeechLab, Department of Computer Science and Engineering, Shanghai Jiao Tong University, China

ABSTRACT

End-to-end (E2E) automatic speech recognition (ASR) systems directly map acoustics to words using a unified model. Previous works mostly focus on E2E training a single model which integrates acoustic and language model into a whole. Although E2E training benefits from sequence modeling and simplified decoding pipelines, large amount of transcribed acoustic data is usually required, and traditional acoustic and language modelling techniques cannot be utilized. In this paper, a novel modular training framework of E2E ASR is proposed to separately train neural acoustic and language models during training stage, while still performing end-to-end inference in decoding stage. Here, an acoustics-to-phoneme model (A2P) and a phoneme-to-word model (P2W) are trained using acoustic data and text data respectively. A phone synchronous decoding (PSD) module is inserted between A2P and P2W to reduce sequence lengths without precision loss. Finally, modules are integrated into an acoustics-to-word model (A2W) and jointly optimized using acoustic data to retain the advantage of sequence modeling. Experiments on a 300-hour Switchboard task show significant improvement over the direct A2W model. The efficiency in both training and decoding also benefits from the proposed method.

Index Terms— E2E, Connectionist Temporal Classification, Attention-based Encoder-decoder, Phone Synchronous Decoding

1. INTRODUCTION

Acoustic modeling with deep neural network (DNN) commonly operates in the hybrid framework: DNNs, as discriminative models, estimate the posterior probabilities of Hidden Markov Model (HMM) states, while DNN and HMM are optimized separately. Recent advances in more powerful neural networks enable stronger modeling effects in the context and history of sequence modeling [1, 2]. More labeled data further alleviates the sparseness and generalization problem in the modeling. Thus in automatic speech recognition (ASR), it is promising to model the acoustics in sequence level without frame decomposition and separate optimization. End-to-end (E2E) automatic speech recognition (ASR) systems directly map acoustics to words using a unified model, which integrates acoustic and language model into a whole. E2E usually utilizes larger inference granularities, which simplifies the ASR pipeline and reduces the decoding complexity, e.g. removing the lexicon design, and inferring word or grapheme sequences directly. E2E can be divided into two branches as below:

Connectionist temporal classification (CTC) is proposed to model the posterior probability of the label sequence given the acoustic feature sequence. The `blank` label unit is introduced to predict the

label sequence on unsegmented data at any time step in the input sequence. Researches have been conducted on different model granularities from tri-phone to word [3, 4]. Recently, novel structures and respective criteria [5, 6, 7] are proposed to improve the conditional independent assumption (CIA) in CTC. Attention based encoder-decoder (S2S) [8, 9] is another end-to-end model without CIA. It predicts the posterior probability of the label sequence given the acoustic feature sequence and previous inference labels. An attention mechanism weights the hidden vectors of the feature sequence to use the most related hidden vectors for prediction.

Both E2E optimization and decoding are operated in these models. Although E2E ASR benefits from sequence modeling, large amount of acoustic data is required to optimize the model jointly, e.g., a word level CTC is trained on 125K hours labeled data and outperforms hybrid models [4]. Meanwhile, the end-to-end decoding reduces the weighted finite state transducer (WFST) based search complexity [10], which is one of the key advantages.

In this paper, a modular training strategy is proposed to utilize external sources to train each building block, while the end-to-end decoding is retained. The ideas we apply to E2E systems are: i) utilizing acoustic and text data in end-to-end ASR modeling by modular training strategy. A CTC acoustics-to-phoneme model (A2P) is trained using acoustic data. Text data is utilized to train a phoneme-to-word model (P2W) of CTC or S2S. ii) combining modules into an acoustics-to-word model (A2W) by phone synchronous decoding (PSD) [10] and joint optimization. Namely, a PSD module is proposed on CTC inference sequences to skip `blank`-dominated frames, which reduces the sequence length without precision loss. Finally, the A2P, PSD and P2W modules are stacked and jointly fine-tuned using acoustic data to retain the advantage of sequence modeling. The advantages include: i) easier and faster model convergence due to modularization and initialization. ii) easy to utilize traditional AM and language model (LM) technologies using text and acoustic data respectively.

The rest of the paper is organized as follows. In Section 2, end-to-end training methods are briefly reviewed. In Section 3.2, the modular training framework is proposed and compared with prior works in Section 4. In Section 5, experiments are conducted on Switchboard corpus in the context of both CTC and S2S. Finally we present our conclusions in Section 6.

2. END-TO-END SPEECH RECOGNITION

2.1. Connectionist Temporal Classification

CTC [11] directly calculates the posterior probability $P(\mathbf{l}|\mathbf{x})$ of the sequence \mathbf{l} given the feature sequence \mathbf{x} . It is proposed to label unsegmented data by introduction of the `blank` label unit to predict the label sequence at any time step in the input sequence.

$$P(\mathbf{l}|\mathbf{x}) = \sum_{\pi \in \mathcal{B}^{-1}(\mathbf{l})} P(\pi|\mathbf{x}) = \sum_{\pi} \prod_{t=1}^T P(\pi_t|\mathbf{x}) \quad (1)$$

The corresponding author is Kai Yu. This work has been supported by the National Key Research and Development Program of China under Grant No.2017YFB1002102. Experiments have been carried out on the PI supercomputer at Shanghai Jiao Tong University.

where \mathcal{B} is a many-to-one mapping defined as below.

$$\mathcal{B} : L \cup \{\text{blank}\} \mapsto L \quad (2)$$

\mathcal{B} determines the label sequence \mathbf{l} and its corresponding set of CTC paths π . The mapping is by inserting an optional and self-loop blank between each label unit l in \mathbf{l} . $P(\pi_t|\mathbf{x})$ is estimated from the neural network taking the feature sequence \mathbf{x} as the input, e.g. long short term memory (LSTM) [12]. With the conditional independent assumption (CIA), $P(\pi|\mathbf{x})$ can be decomposed into a product of posterior $P(\pi_t|\mathbf{x})$ in each frame t .

To improve CIA in CTC, RNN transducer [5], recurrent neural aligner (RNA) [6] and neural segmental model [7] are proposed. e.g. in RNA, when predicting the label of the current time step, the predicted label at last time step is used as an additional input to the recurrent model. An approximate dynamic programming method to optimize the negative log likelihood, and a sampling-based sequence discriminative training technique are designed for this model and achieve competitive performance. With these more powerful inference structures, the label context dependency can be better modeled and the model works well without external LM.

Another trend is to use larger model granularities to simplify the training framework and decoding complexity. A word level deep learning based acoustic model, A2W CTC [4], is trained on 125K hours labeled data and outperforms models with smaller granularities. To reduce the amount of acoustic data required to train this model, [13] proposes: i) initializing lower layers by monophone CTC [13] to improve acoustics extraction. ii) initializing upper dense layers by GloVe [14] word embeddings to capture word co-occurrence information.

2.2. Attention-based Encoder-decoder

Attention based encoder-decoder (S2S) [8, 9, 15, 16] is another branch of end-to-end model. It does not take CIA in CTC and predicts the posterior probability of the label sequence given both the feature sequence \mathbf{x} and the previous inference labels $\mathbf{l}_{1:i-1}$.

$$P(\mathbf{l}|\mathbf{x}) = \prod_i P(l_i|\mathbf{x}, \mathbf{l}_{1:i-1}) \quad (3)$$

An attention mechanism weights the hidden vectors of the feature sequence so that the most related hidden vectors are used for the prediction.

$$\begin{aligned} \mathbf{h} &= \text{Encoder}(\mathbf{x}) \\ P(l_i|\mathbf{x}, \mathbf{l}_{1:i-1}) &= \text{AttentionDecoder}(\mathbf{h}, \mathbf{l}_{1:i-1}) \end{aligned} \quad (4)$$

where $\text{Encoder}(\cdot)$ can be LSTM or bidirectional LSTM (BLSTM) and $\text{AttentionDecoder}(\cdot)$ can be LSTM or gated recurrent unit (GRU).

Recently, several model structure improvements are proposed. To help the model convergence, a pyramid structure is used in the encoder network to sub-sample the feature sequence [9]. The attention training can be further combined with CTC criterion in a multi-task learning way, which greatly improves its convergence [17]. To enable streaming recognition, online encoder-decoder structures are proposed, e.g. window based attention [18], Gaussian prediction based attention [19] and reinforcement learning based method [20].

As speech recognition is both an acoustic and linguistic pattern recognition problem, several works try to utilize text data in the initialization stage before E2E joint optimization. The decoder part of S2S can be initialized by a LM trained from text data [21, 22], which helps the model convergence and domain adaptation.

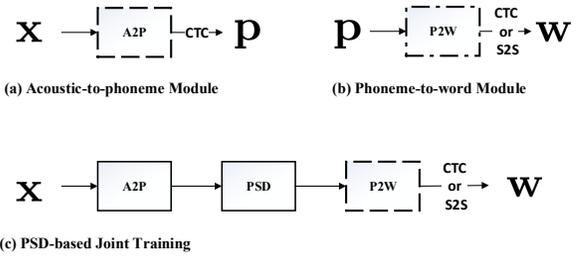


Fig. 1. Framework of Modular Training of Neural Acoustics-to-Word Model. The solid line box denotes the layers whose parameters are fixed. The dash line and dash-dot line boxes denote that models are trained from acoustic data and text data respectively.

2.3. Drawbacks of Current Frameworks

Firstly, the acoustic and text data are not utilized systematically, which results in a large acoustic data requirement. Current trials merely use text data for initialization, but not utilize both text and acoustic data in a systematical framework. Nevertheless, the traditional neural network based hidden markov model (NN-HMM) framework takes Bayesian theorem and combines acoustic model, transition model, lexicon and LM in a modular way.

Secondly, both AM and LM utilize the same model unit, which hurts the generalization and performance. In human speech perception, phoneme perception is the foundation of word recognition [23]. Grapheme and word are not designed for acoustics and do not have explicit relationship with their pronunciation. Such trials ignore the prior phoneme knowledge in human language. Meanwhile, taking character or grapheme as the model unit for LMs also deteriorates the performance [24].

Overall speaking, E2E systems form a different modeling framework compared with NN-HMM hybrid systems. Hence, prior arts and different knowledge sources of both AM and LM are hard to be applied in E2E ASR.

3. MODULAR-TRAINED NEURAL A2W MODEL

3.1. Training and Decoding Framework

Previous works of E2E ASR focus on integrating all components into a whole with both joint optimization and E2E decoding. In this work, a modular training strategy is proposed to improve the performance by utilizing external sources to train each building block, while end-to-end decoding is retained to keep the efficient decoding advantage of E2E ASR. Figure 1 shows the framework of the proposed method.

The E2E word sequence recognition is modularized as below.

$$P(\mathbf{w}|\mathbf{x}) \approx \max_{\mathbf{p}} [P(\mathbf{w}|\mathbf{p}) \cdot P(\mathbf{p}|\mathbf{x})] \quad (5)$$

where \mathbf{w} , \mathbf{p} and \mathbf{x} are word sequence, phoneme sequence and acoustic feature sequence respectively. An acoustics-to-phoneme model (A2P) is trained by CTC criterion using acoustic data. Meanwhile, a phoneme-to-word model (P2W) is trained by CTC or S2S using text.

Then modules are integrated into an A2W model (A2W) by phone synchronous decoding (PSD) [10] and joint optimization.

$$P(\mathbf{w}|\mathbf{x}) \approx \max_{\mathbf{p}} [P(\mathbf{w}|\mathbf{p}) \cdot \text{PSD}(P(\mathbf{p}|\mathbf{x}))] \quad (6)$$

In decoding stage, the jointly optimized A2W model is used as a whole to infer word sequences directly, whose complexity is similar to that of traditional A2W systems [13]. For CTC, the maximum inference labels in each output step are concatenated as the decoding

result. Viterbi beam search algorithm is applied in S2S. The A2W models can be further combined with external language models to improve the performance. In this case, the n-gram language model is compiled to a word WFST. Thus, PSD search algorithm [10] can be applied in word level to speed up the pipeline.

3.2. Modularization

As phoneme is a prior knowledge to define all possible pronunciation the ASR system cares about, it has explicit relationship with acoustics and it is taken as the acoustic model unit. The A2P module utilizes acoustic data and predicts $P(\mathbf{p}|\mathbf{x})$ as Figure 1(a), which is the same to typical phoneme CTC models [25]. Notably, although CTC is applied in this work for better comparison purpose, other traditional acoustic models, inferring phoneme given acoustics, can also be applied in this module, e.g. RNA [6] and LF-MMI [26].

Different from the end-to-end models discussed in Section 2, the LM here takes the word as the inference unit and predicts $P(\mathbf{w}|\mathbf{p})$ given the phoneme sequence, i.e. phoneme-to-word model in Figure 1(b). Another key difference is that P2W module utilizes text data and a lexicon, nor acoustic data. Overall speaking, the P2W module is similar to traditional LMs except that: i) P2W consumes phoneme sequences and implicitly does the tokenization on them. ii) P2W infers the word sequence given the phoneme sequence. Hence different from traditional LMs, inferring the next word given the previous word sequence, P2W gets more hints from the next word. Our experiment results also show that P2W performs better than LMs in predicting word sequences. iii) P2W is trained by sequence criteria, CTC and S2S, which automatically learn the alignment between phoneme sequences and word sequences.

An extra word boundary unit \mathbf{wb} is introduced into the phoneme set to improve the tokenization discussed above. \mathbf{wb} exists in the end of the phoneme sequence belonging to each word. e.g., the lexicon “okay ow k ey” is changed to “okay ow k ey \mathbf{wb} ”. The motivation is to take \mathbf{wb} as the hint of tokenization, e.g. distinguishing short words in case that its phoneme sequence is a substring of longer words.

The advantages of the modularization include: i) both AM and LM use proper inference unit. ii) systematically combining acoustic and text data in the model training.

3.3. Phone Synchronous Decoding

With different inference units, namely phoneme and word, PSD is applied to cope with different information rates.

PSD is originally proposed to speed up the ASR decoding [10]. For frames where \mathbf{blank} symbol dominates in the CTC inference sequence, it is redundant to do searching as no phoneme information is provided. Given this observation, phone synchronous decoding is proposed by skipping the search of blank-dominated time steps during CTC decoding. Statistic in [10] shows that the remaining inference distributions are very compact, which greatly reduces the information rate without precision loss.

Here, $PSD(\cdot)$ is applied as a sub-sampling layer upon the A2P inference sequence, which reduces the input sequence length of P2W layers in the joint optimization. The advantages of PSD module include: i) ease the burden of LSTM temporal modeling. ii) speed up the joint optimization. Besides, PSD module also speeds up the decoding stage [10].

3.4. Joint Optimization

Finally, modules are stacked as Figure 1(c). The acoustic data is used to fine-tune the stacked model. Word level CTC criterion is applied similar to [4]. Meanwhile, S2S is applied in word level for the first time. During the optimization, only P2W module is fine-tuned and the reason includes: i) the A2P module, mono-phone level

CTC model, can always achieve good modeling effects [25, 3]. ii) fixing A2P and combining PSD module can greatly speed up the joint optimization. The advantage of sequence modeling in end-to-end model is retained in this stage.

4. RELATION TO PRIOR WORK

End-to-end ASR models always utilize large amount of acoustic data to optimize the sequence level model jointly. e.g., in [4], a word level CTC acoustic model is trained on 125K hours labeled data and outperforms hybrid models with smaller granularity. As speech recognition is both an acoustic and linguistic pattern recognition problem, several works try to utilize text data in the initialization stage before end-to-end joint optimization. In CTC, [13] proposes to initialize the word output dense layers with GloVe [14] word embeddings to capture word co-occurrence information. In attention-based encoder-decoder, the decoder part can be initialized by a LM trained from text data [21, 22], which helps the model convergence and domain adaptation. This work can be applied in both CTC and S2S. The key differences include: i) modularizing the end-to-end speech recognition by Bayesian theorem. ii) utilizing respective inference units for acoustic and language modeling. iii) the LM generalizes word sequences and lexicons jointly. Recently, novel structures [5, 6, 7] are proposed to improve the CIA in CTC by jointly training a LM with the vanilla CTC. A key difference is that the proposed AM and LM infer phonemes and words respectively, which utilizes the linguistic knowledge as discussed in 3.2.

5. EXPERIMENTS

5.1. Experimental Setup

Experiments were conducted on Switchboard corpus [27], which contains about 300 hours of speech. 36-dimensional filterbank over 25 ms frames every 10 ms from the input speech signal was extracted. Neural networks were trained by Torch [28] and Kaldi [29, 25]. The model unit of phoneme CTC was 45 mono-phones and a \mathbf{blank} . The baseline phoneme CTC was with 5-layer LSTMs, each with 1024 memory cells and 256 nodes projection layer [1]. The baseline hybrid system was trained by cross entropy (CE) criterion and with the same structure except the last layer, which was with 8K clustered tri-phone states. The CTC model was initialized by the baseline hybrid system above [3] and the training procedure was similar to [25]. In modular training, the baseline phoneme CTC was taken as the A2P module. The CTC P2W module was with 4-layer LSTMs, each with 700 memory cells and 256 nodes projection layer. The S2S P2W module used 5-layer LSTMs with 700 nodes for encoder and 1-layer LSTM with 700 nodes for decoder. The vocabulary size was 30K as the standard evaluation setup in this corpus and it was taken as the output layer of P2W module¹. As a baseline of the end-to-end system without modular training, the direct A2W CTC, similar to [13], was with the same structure as the phoneme CTC except the last layer with 30K words. It was initialized by the phoneme CTC.

Evaluation was carried out on the Switchboard (swbd) and Callhome (callhm) subset of the NIST 2000 CTS test set. A 30k-vocabulary tri-gram LM trained from the transcription of Switchboard corpus, without Fisher corpus interpolation, was used for the decoding of baseline phoneme CTC. The P2W module was also trained from the Switchboard transcription to make them comparable with each other. The decoding procedures of CD-phone CE and CI-phone CTC were the same to [25]. The decoding algorithm of A2W systems was discussed in Section 3.1. S2S

¹Specific filtering on the vocabulary may improve the performance [4]. Nevertheless, we didn't do that to make it comparable with other traditional systems in this corpus and also simplify the pipeline.

utilized a beam of 20. Word error rate (WER) and phoneme error rate (PER) were taken as the metric.

5.2. Modularization

Table 1 shows the performance of each module in the validation set (CV). Systems with bold-font are used in the later experiments.

Table 1. Performance of Each Module.

Module	Model	Inf. Label	Word bound.	PER/WER CV (%)
A2P	CTC	phoneme	×	13.0
			✓	12.0
P2W	CTC	word	×	16.0
			✓	4.3
	S2S	word	×	13.9
			✓	2.8

In A2P, the baseline phoneme recognition performance is parallel with [11]. *wb* discussed in Section 3.2 does not hurt the performance. The slight improvement stems from the statistic of PER includes *wb*. Further statistic shows *wb* prediction error rate is 4%.

In the P2W module, both CTC and S2S are examined. Without *wb* inserted in the phoneme sequence, both CTC and S2S obtain large WERs. As discussed in Section 3.2, *wb* gives hints of phoneme sequence tokenization. Hence both CTC and S2S with *wb* improves significantly. S2S consistently shows better performance versus CTC, which benefits from removal of the CIA in CTC [9]. Different from traditional LMs, the perplexity (PPL) is not reported because the phoneme sequence and the word sequence have different lengths, while the alignment between them is uncertain and it is automatically learned by sequence criteria.

5.3. Joint Optimization

After modular initialization, models are jointly optimized in Table 2. To better support the results, we compare them with the recent A2W system in this corpus [13]. The different setups include: i) i-vector based adaptation ii) utilizing BLSTM iii) the LM is interpolated with Fisher corpus. Hence the gaps between baselines of this work and those of [13] are always relative 20-30%.

Table 2. Performance Comparison with or w/o Modular Training.

Name	E2E Opt.	Modularization		WER (%)	
		A2P	P2W	swbd	callhm
CD-phone CE	×	HMM	WFST	14.9	27.6
CI-phone CTC	×	CTC	WFST	19.4	33.5
Word CTC	✓	n/a	n/a	29.6	41.7
Mod. CTC	✓	CTC	CTC	24.9	36.5
	✓	CTC	+WFST	23.0	35.1
Mod. S2S	✓	CTC	S2S	31.2	40.5

The baseline hybrid system (CD-phone CE) and phoneme CTC (CI-phone CTC) are in line 1 and 2 respectively. Both of them are decoded with a WFST generated from the 30K lexicon and n-gram LM. The performance of CI-phone CTC is worse than that of CD-phone CE², whose gap is similar to that in [13]³. The direct A2W CTC (Word CTC) is in line 3. It is with phoneme initialization but without GloVe initialization [13]. The performance is significantly worse than CI-phone CTC⁴. This setup is taken as the naive A2W

²CTC performs better than hybrid systems with more data [3]. Our previous work [10] also shares similar finding in a large dataset.

³In the same corpus, the phoneme CTC baseline WER in [13] is 14.5% and 25.1% for *swbd* and *callhm* respectively.

⁴The WER in [13] is 24.4% and 34.1%, which shows similar trend.

baseline, where Glove is not implemented as we believe the modular training is a better way to capture the linguistic information.

The proposed modular-trained A2W CTC (Mod. CTC) is in line 4. The PSD based joint optimization is applied here, whose effects will be examined in Table 3 later. Mod. CTC outperforms the naive A2W in line 3 significantly. The modular training framework benefits from: i) easier and faster model convergence due to modularization and initialization. ii) easy to utilize standard AM and LM technologies using text and acoustic data respectively.

Table 3 shows modeling effects and training speeds with or without PSD. All results are reported on a single Titan GPU. “fr./s.” denotes the number of acoustics frames processed per second. The training speedup stems from two folds: i) PSD reduces the sequence length to be processed by P2W in each sequence. ii) As the sequence length is reduced, more sequences can be loaded into GPU memory for parallel training. Meanwhile, the performance is significantly improved. We believe it also results from sequence length reduction. Although LSTM is used, the model is still hard to remember a very long input sequence. Nevertheless, for A2W modeling, the history to be remembered before inferring each word is much longer than that of the traditional CI-phone CTC or hybrid systems. Similar problem is reported in [9] and solved by a pyramid model structure. The PSD framework shows another choice to cope with this problem.

Table 3. Performance and Speed Comparison with or w/o PSD.

Name	PSD	Training Speed		WER (%)	
		Seq./GPU	fr./s.	swbd	callhm
Mod. CTC	×	5	1027	32.0	42.5
	✓	30	5851	24.9	36.5

To alleviate the deterioration of the word sequence modeling effect because of the CIA in CTC, two methods are further investigated⁵. Firstly, a WFST generated from the previous n-gram LM is used to decode with the proposed system. The result is shown in Table 2 line 5 with moderate improvement. Hence, the performance gap between line 2 and line 5, CI-phone CTC versus the A2W system, is reduced to relative 15%⁶. Another method is replacing CTC with S2S discussed in Section 2.2. The proposed modular-trained A2W S2S (Mod. S2S) is in Table 2 the last line. Different from the observation in Table 1, the S2S based system does not achieve improvement. Analysis in the decoding result shows that S2S is prone to the phoneme recognition errors from the A2P module. After joint optimization, S2S can not restore the errors, which is similar to the observation in [16]. Besides, grapheme based systems are not included in this work, although they have been investigated in most of the S2S works, whereas the grapheme is not suitable for language modeling as discussed in Section 2.3.

6. CONCLUSION

In this paper, modular training strategy of E2E ASR is proposed, while end-to-end decoding is retained. Experiments on 300 hours Switchboard corpus show significant improvement on the naive word CTC (A2W). Compared with the recent A2W system [13] in this corpus, our performance gap between CI-phone CTC and the A2W model is reduced⁷, which also confirms the efficiency of this work. To further reduce the gap, future works in the modular training framework include: i) utilizing external knowledge sources. ii) applying novel models [5, 6]. iii) improving S2S-based P2W.

⁵We notice that several recent advances of CTC discussed in Section 2.1 also cope with this problem, which can be our future works.

⁶Our recent trials on more potent P2W structure finally eliminates the gap.

⁷For *swbd*, 19.4% to 23.0% in our setup versus 14.5% to 21.7% in [13]. Meanwhile, other performance gaps show similar trends, e.g. CD-phone CE versus CI-phone CTC, and word CTC versus naive word CTC.

7. REFERENCES

- [1] Haşim Sak, Andrew Senior, and Françoise Beaufays, “Long short-term memory recurrent neural network architectures for large scale acoustic modeling,” in *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [2] Yanmin Qian, Mengxiao Bi, Tian Tan, and Kai Yu, “Very deep convolutional neural networks for noise robust speech recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 12, pp. 2263–2276, 2016.
- [3] Haşim Sak, Andrew Senior, Kanishka Rao, and Françoise Beaufays, “Fast and accurate recurrent neural network acoustic models for speech recognition,” *arXiv preprint arXiv:1507.06947*, 2015.
- [4] Hagen Soltau, Hank Liao, and Hasim Sak, “Neural speech recognizer: Acoustic-to-word lstm model for large vocabulary speech recognition,” *arXiv preprint arXiv:1610.09975*, 2016.
- [5] Alex Graves, “Sequence transduction with recurrent neural networks,” *arXiv preprint arXiv:1211.3711*, 2012.
- [6] Haşim Sak, Matt Shannon, Kanishka Rao, and Françoise Beaufays, “Recurrent neural aligner: An encoder-decoder neural network model for sequence to sequence mapping,” *Proc. Interspeech 2017*, pp. 1298–1302, 2017.
- [7] Hao Tang, Liang Lu, Lingpeng Kong, Kevin Gimpel, Karen Livescu, Chris Dyer, Noah A Smith, and Steve Renals, “End-to-end neural segmental models for speech recognition,” *IEEE Journal of Selected Topics in Signal Processing*, 2017.
- [8] Ronan Collobert, Christian Puhrsch, and Gabriel Synnaeve, “Wav2letter: an end-to-end convnet-based speech recognition system,” *arXiv preprint arXiv:1609.03193*, 2016.
- [9] William Chan, *End-to-End Speech Recognition Models*, Ph.D. thesis, Carnegie Mellon University Pittsburgh, PA, 2016.
- [10] Z. Chen, Y. Zhuang, Y. Qian, and K. Yu, “Phone synchronous speech recognition with ctc lattices,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 1, pp. 86–97, Jan 2017.
- [11] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of the 23rd international conference on Machine learning*. ACM, 2006, pp. 369–376.
- [12] Sepp Hochreiter and Jürgen Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [13] Kartik Audhkhasi, Bhuvana Ramabhadran, George Saon, Michael Picheny, and David Nahamoo, “Direct acoustics-to-word models for english conversational speech recognition,” *arXiv preprint arXiv:1703.07754*, 2017.
- [14] Jeffrey Pennington, Richard Socher, and Christopher Manning, “Glove: Global vectors for word representation,” in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.
- [15] Ehsan Variiani, Tom Bagby, Erik McDermott, and Michiel Bacchiani, “End-to-end training of acoustic models for large vocabulary continuous speech recognition with tensorflow,” *Proc. Interspeech 2017*, pp. 1641–1645, 2017.
- [16] Rohit Prabhavalkar, Kanishka Rao, Tara N Sainath, Bo Li, Leif Johnson, and Navdeep Jaitly, “A comparison of sequence-to-sequence models for speech recognition,” *Proc. Interspeech 2017*, pp. 939–943, 2017.
- [17] Suyoun Kim, Takaaki Hori, and Shinji Watanabe, “Joint ctc-attention based end-to-end speech recognition using multi-task learning,” in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 4835–4839.
- [18] Rohit Prabhavalkar, Tara N Sainath, Bo Li, Kanishka Rao, and Navdeep Jaitly, “An analysis of attention in sequence-to-sequence models,” *Proc. Interspeech 2017*, pp. 3702–3706, 2017.
- [19] Junfeng Hou, Shiliang Zhang, and Lirong Dai, “Gaussian prediction based attention for online end-to-end speech recognition,” *Proc. Interspeech 2017*, pp. 3692–3696, 2017.
- [20] Yuping Luo, Chung-Cheng Chiu, Navdeep Jaitly, and Ilya Sutskever, “Learning online alignments with continuous rewards policy gradient,” in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 2801–2805.
- [21] Takaaki Hori, Shinji Watanabe, Yu Zhang, and William Chan, “Advances in joint ctc-attention based end-to-end speech recognition with a deep cnn encoder and rnn-lm,” *arXiv preprint arXiv:1706.02737*, 2017.
- [22] Anuroop Sriram, Heewoo Jun, Sanjeev Satheesh, and Adam Coates, “Cold fusion: Training seq2seq models together with language models,” *arXiv preprint arXiv:1708.06426*, 2017.
- [23] David B Pisoni, Howard C Nusbaum, Paul A Luce, and Louisa M Slowiaczek, “Speech perception, word recognition and the structure of the lexicon,” *Speech communication*, vol. 4, no. 1-3, pp. 75–95, 1985.
- [24] Rafal Jozefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu, “Exploring the limits of language modeling,” *arXiv preprint arXiv:1602.02410*, 2016.
- [25] Yajie Miao, Mohammad Gowayyed, and Florian Metze, “Eesen: End-to-end speech recognition using deep rnn models and wfst-based decoding,” in *Automatic Speech Recognition and Understanding (ASRU), 2015 IEEE Workshop on*. IEEE, 2015, pp. 167–174.
- [26] Daniel Povey, Vijayaditya Peddinti, Daniel Galvez, Pegah Ghahramani, Vimal Manohar, Xingyu Na, Yiming Wang, and Sanjeev Khudanpur, “Purely sequence-trained neural networks for asr based on lattice-free mmi,” *Submitted to Interspeech*, 2016.
- [27] John J Godfrey, Edward C Holliman, and Jane McDaniel, “Switchboard: Telephone speech corpus for research and development,” in *Acoustics, Speech, and Signal Processing, 1992. ICASSP-92., 1992 IEEE International Conference on*. IEEE, 1992, vol. 1, pp. 517–520.
- [28] Adam Paszke, Sam Gross, and Soumith Chintala, “Pytorch,” 2017.
- [29] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, et al., “The kald speech recognition toolkit,” in *IEEE 2011 workshop on automatic speech recognition and understanding*. IEEE Signal Processing Society, 2011, number EPFL-CONF-192584.