



## RWTH ASR Systems for LibriSpeech: Hybrid vs Attention

Christoph Lüscher<sup>1</sup>, Eugen Beck<sup>1,2</sup>, Kazuki Irie<sup>1</sup>, Markus Kitza<sup>1</sup>, Wilfried Michel<sup>1,2</sup>, Albert Zeyer<sup>1,2</sup>,  
Ralf Schlüter<sup>1</sup>, Hermann Ney<sup>1,2</sup>

<sup>1</sup>Human Language Technology and Pattern Recognition, Computer Science Department,  
RWTH Aachen University, 52074 Aachen, Germany

<sup>2</sup>AppTek GmbH, 52062 Aachen, Germany

{luescher, beck, irie, kitza, michel, zeyer, schlueter, ney}@cs.rwth-aachen.de

### Abstract

We present state-of-the-art automatic speech recognition (ASR) systems employing a standard hybrid DNN/HMM architecture compared to an attention-based encoder-decoder design for the LibriSpeech task. Detailed descriptions of the system development, including model design, pretraining schemes, training schedules, and optimization approaches are provided for both system architectures. Both hybrid DNN/HMM and attention-based systems employ bi-directional LSTMs for acoustic modeling/encoding. For language modeling, we employ both LSTM and Transformer based architectures. All our systems are built using RWTH's open-source toolkits RASR and RETURNN. To the best knowledge of the authors, the results obtained when training on the full LibriSpeech training set, are the best published currently, both for the hybrid DNN/HMM and the attention-based systems. Our single hybrid system even outperforms previous results obtained from combining eight single systems. Our comparison shows that on the LibriSpeech 960h task, the hybrid DNN/HMM system outperforms the attention-based system by 15% relative on the clean and 40% relative on the other test sets in terms of word error rate. Moreover, experiments on a reduced 100h-subset of the LibriSpeech training corpus even show a more pronounced margin between the hybrid DNN/HMM and attention-based architectures.

**Index Terms:** speech recognition, hybrid BLSTM/HMM, attention, LibriSpeech

### 1. Introduction

Over the last years, automatic speech recognition (ASR) systems have improved significantly. Especially the rise of deep neural networks (NNs) has accelerated this development immensely [1]. Convolutional NNs and recurrent NNs are the state-of-the-art architectures for most ASR tasks. State-of-the-art systems are largely based on the hybrid deep neural network (DNN) based standard architectures. However, the general progress in deep learning/machine learning also triggered a diversification of ASR architectures into a series of so-called end-to-end approaches. Most notably, this includes the attention-based encoder-decoder architecture, for which good performance has been reported on a number of tasks, including the LibriSpeech task.

The LibriSpeech task comprises English read speech data based on the LibriVox project [2]. Previous results on LibriSpeech using hybrid models are presented in [3, 4]. While [3] uses a Gaussian mixture model/hidden Markov model (GMM/HMM) as the basis for their system, further training is conducted with a hybrid DNN/HMM with a densely connected topology. The densely connected NNs in [3] are composed of different types of NN layers: convolutional NNs, time

delay neural network and bi-directional long short-term memory (LSTMs). Lattice-free maximum mutual information (LFMMI) is applied during training. A recurrent NN language model (LM) is used for rescaling. The final best result in [3] is achieved with a system combination of eight systems. In [4], a lattice-free state-level minimum bayes risk (sMBR) training method is used.

End-to-end results on LibriSpeech were presented in [5–9]. The end-to-end approach in [8] uses the raw waveform and a convolutional NN acoustic model with gated linear units. An end-to-end attention-based encoder-decoder approach with a pretraining scheme is presented in [5, 6]. In [7] a training procedure based on edit distance for sequence to sequence model optimization is presented. An exploration of target units (phoneme, grapheme and word-piece) in relation to training size was performed in [9]. A data augmentation method called SpecAugment was presented in [10]. So far, while end-to-end approaches show competitive performance, they are outperformed by hybrid approaches. We compare the conventional hybrid DNN/HMM approach on phone level to the encoder-decoder-attention model which directly operates on the word or sub-word level and is thus often referred to as an end-to-end model. In addition, we use word-level and subword-level neural language models to further improve the performance of both systems. We describe the development of our hybrid system and show which factors were especially important for the performance of the system. To the best of our knowledge, the results obtained on the LibriSpeech task reflect state-of-the-art performance for both hybrid and attention-based modeling, with a clear margin still for hybrid DNN/HMM modeling when no data augmentation scheme is applied.

### 2. Hybrid system

#### 2.1. Acoustic model

##### 2.1.1. GMM/HMM system

We use 16-dim Mel frequency cepstral coefficients (MFCC) adding first and second order derivatives, and additionally energy features as input for the GMM/HMM system. The transition probabilities are set manually and applied to all HMMs.

The first step is linear time alignment where the features are uniformly distributed over the audio. We iterate the repetition of the parameter estimation based on the linear time alignment five times. Afterwards, we perform a non-linear time alignment to improve the alignment. Afterwards we perform parameter estimation. Initially, this process was iterated 10 times. Increasing the number of iterations showed constant improvement. Therefore we continued adding training iterations until word error rate (WER) convergence. Training of a state-

tied triphone GMM/HMM model is the following step. The states are tied using a phonetic classification and regression tree (CART). We experimented with different numbers of CART labels ranging from  $9k - 20k$  plus silence. All state-tied triphones use three HMM states. We switch the input features from 16-dim. MFCC with derivatives to a context window of MFCC features resulting in a 144-dim. feature vector on which linear discriminant analysis (LDA) is performed. The LDA output has a dimension of 48. After training the state-tied triphone GMM/HMM model, vocal tract length normalization (VTLN) is applied, followed by speaker adaptive training (SAT) with constrained maximum likelihood linear regression to adapt the Gaussian mixture model parameters to a speaker. After adapting the parameters, a realignment is performed.

### 2.1.2. Hybrid DNN/HMM system

The NN acoustic model architecture is a bi-directional LSTM [11, 12]. This architecture achieves good performance in acoustic modelling [12, 13]. For the hybrid DNN/HMM system we extract several different features: 16-dim MFCC with derivatives and energy, 48-dim. features from the triphone system and Gammatone filters [14] with 25, 50 or 100-dim. The extracted 50-dim. Gammatone filters had the best performance. All features are used as input into the bi-directional LSTM along with the generated alignments from the GMM/HMM system. We continue to use CART labels for the state-tied phones. The same range of CART labels was used for experimentation. The network topology consists of six bi-directional LSTM layers with 1000 units for backward and forward direction each. We experimented with smaller bi-directional LSTM sizes (number of layers and number of units per layer) but found them to be worse in performance. The output layer is comprised of a softmax layer with output units corresponding to the number of the CART labels. Frame-wise cross-entropy loss criterion and Adam optimization with Nesterov momentum (Nadam) are used for the mini-batch training of the network [15, 16]. Newbob learning rate scheduling [13] is applied to control the learning rate reduction with a learning rate decay rate of 0.9.  $L_2$  regularization was used to prevent overfitting. The  $L_2$  hyperparameter was set to 0.01. Further regularization is done with dropout [17]. We experimented with dropout in the range of 5% - 30% and found a dropout of 10% to work best for us. Gradient noise [18] with a variance of 0.1 was employed. We experimented with different learning rates and batch sizes in various combinations. So far a batch size of  $20k$  and a learning rate of 0.008 have shown the best performance. Additionally, learning rate warm up proved to be helpful. We start with a learning rate of 0.003 and increase the learning rate to 0.008 over the first ten subepochs. A subepoch is 1/40th of the training data. The training data is seen 12.5 times. During decoding the LM scale is an important hyperparameter which will effect the WER directly. We found a scale between 11-13 worked best for us.

### 2.1.3. Sequence discriminative training

Sequence discriminative training is performed using a lattice-based version of the state-level minimum Bayes risk (sMBR) criterion [19].

The hybrid DNN/HMM model is used to generate lattices for all of the training data. The training is then continued from the hybrid DNN/HMM model with a lower learning rate.

We use cross-entropy smoothing with a smoothing factor of 0.1 and early stopping to prevent overfitting.

## 2.2. Language model

We report performance of hybrid systems using both a 4-gram count based language model [20] and an LSTM language model [21] in the *first pass decoding* [22]. We use the 4-gram count model officially distributed with the LibriSpeech dataset [2]. For the LSTM language model, we train our own model using our toolkit RETURNN [23]. Two training datasets are available for language modeling: 800M-word text only data and 960h of audio transcriptions which corresponds to 10M-word text data. These two sets are merged to form one training dataset for language model training. Our LSTM language model has two recurrent layers with 4096 LSTM nodes in each layer, an input projection layer of size 128, and a output softmax layer over the full 200k vocabulary. We train the model using the stochastic gradient descent with gradient norm clipping and Newbob learning rate scheduling.

In addition, we carry out *rescoring of lattices* generated by the LSTM language model using a Transformer [24] language model. Our Transformer model has 96 layers with the self-attention total dimension of 512 using 8 heads and the inner feed-forward dimension of 2048 in each layer, which gave the best development perplexity in our preliminary experiments [25]. We use push-forward algorithm [26] with recombination pruning of order 9. We linearly interpolate the two models with interpolation weights optimized on the development perplexity. We found 0.71 to be the optimal weight on the Transformer model which gave the development perplexity of 52.3, while the LSTM and Transformer models have the individual development perplexity of 60.2 and 53.7 respectively.

## 3. Encoder-Decoder-Attention system

The *encoder-decoder framework with attention* has initially been introduced for machine translation where it dominates the field now [27-29]. Recent investigations have shown promising results by applying the same approach for speech recognition [5, 7, 30-33]. Among end-to-end approaches for ASR, the attention model seems to perform best [6]. Our model operates on sub-word units via byte-pair encoding [34]. As input 40-dim MFCC feature vectors are used. Our presented results outperform the best LibriSpeech attention system presented in [6]. Compared to the system in [6] we use an extended pretraining variant where we not only grow the encoder depth but also grow the hidden dimension of the LSTMs. Specifically, we start with 2 layers in the encoder of dimension 512 and increase to 6 layers with dimension 1024. Additionally, we train the first pre-train construction step first without dropout. We improved upon that model by tuning the curriculum learning schedule slightly, i.e. we have these 4 steps with different portions of the dataset:

1. from 25% of the whole data, take only train-clean, and filter randomly such that the max mean number of characters in the transcriptions of each sequence is 50,
2. from the next 25% of the whole data, take only train-clean, and filter randomly such that the max mean number of characters is 75,
3. from the next 50% of the whole data, take only train-clean,
4. from now on, take everything.

Also, in the pretraining, we repeat the first step once more, with 2 layers of dimension 512, without dropout. The next improvement came from just training longer, i.e. we trained with our learning rate scheduling until it converged, then took the best

Table 1: GMM/HMM and hybrid DNN/HMM results on LibriSpeech with 12k CART labels and evaluated with the official 4-gram LM.

phoneme context	acoustic model	VTLN	SAT	sMBR	WER [%]			
					dev		test	
					clean	other	clean	other
mono	GMM	no	no	no	24.3	52.6	24.1	56.1
tri		yes			12.1	34.5	12.9	36.9
		no			12.0	35.1	11.2	36.4
		yes	yes		8.0	21.9	8.6	22.9
			no		7.6	22.0	8.4	23.1
LSTM		yes	yes		4.0	9.6	4.4	10.0
	no		3.4	8.3	3.8	8.8		

model, and continued training with a reset learning rate scheduling. We repeated this twice. In the first iteration, we went over the whole data 12.5 times, then another 6.6 times and finally another 8.3 times, i.e. in total 27.4 times.

To further enhance end-to-end system’s performance, we train byte-pair encoding (BPE)-level language models and apply them to the system by shallow fusion [35,36]. We report the performance of LSTM based and Transformer based language models separately. Our LSTM model has 4 recurrent layers with 2048 LSTM nodes. We use a 24-layer Transformer model with 8-head self-attention and feed-forward dimensionality of 1024 and 4096 respectively, which we obtained in [25]. We select the language model checkpoints for the recognition experiments based on the development perplexity. For shallow fusion, we apply a single weight on the language model score (the weight on the score of the attention model is 1) and we use a beam size of 64 as well as an end-of-sentence penalty [37]. We optimize the weights separately on the dev-clean and dev-other sets, then respectively apply them to the test-clean and test-other sets. We found optimal weights to be similar for both models; 0.5 and 0.56 for the LSTM language model, and 0.52 and 0.54 for the Transformer model, respectively on the clean and other sets.

#### 4. Experimental setup

The two systems, a hybrid-DNN and an attention-based encoder-decoder are both trained on the 960h training data from the LibriSpeech corpus. For comparison, also a 100h subset is used. Unless specified otherwise, the training was performed using the full training set of 960h. The data is in English but the content ranges from different time periods and different English speaking countries. Having the consequence of different English styles being within the corpus.

The hybrid model was trained and decoded with RASR [38] and RETURNN [23, 39]. The monophone and triphone system to generate the alignments was built in RASR while the NN model was trained in RETURNN. The decoding process was setup in RASR. Our encoder-decoder-attention model was trained and decoded using RETURNN [23]. Both toolkits are open-source. All the config files used for training and recognition of all our results are publicly available online [40].

We evaluate the models on the dev and test sets provided with the LibriSpeech corpus: *dev-clean*, *dev-other*, *test-clean* and *test-other*. The difference between *clean* and *other* is the quality of the audio and its corresponding transcription. The *clean* quality is higher than the *other*.

#### 5. Experimental results

The development stages of our acoustic model are shown in Table 1. We start the training of the GMM/HMM model from scratch using linear alignments. Afterwards we utilize non-linear alignments. To further improve the GMM/HMM model we introduce triphones. Adding VTLN on top of the triphone system only shows improvements on clean but degradation on other. However adding SAT to the triphone system improves the WER. Combining VTLN and SAT gives mixed WER: *clean* improves, *other* degrades. Introducing an hybrid DNN/HMM improves the system WER results. Continuing with sequence discriminative training improves the performance even further.

Table 2: Hybrid DNN/HMM results on LibriSpeech with different numbers of CART labels. For all systems the official 4-gram word LM is used.

# of CART labels	WER [%]			
	dev		test	
	clean	other	clean	other
9001	6.2	14.9	5.8	15.9
12001	4.0	9.6	4.4	10.0
20001	4.9	11.3	5.4	12.3

We evaluated the influence of the number of CART labels with the hybrid DNN/HMM model and the official 4-gram LM (Table 2). 9k CART labels show the worst performance. In contrast, 20k CART labels shows improved performance. But the best performance was shown by 12k CART labels.

Table 3: Comparison between hybrid DNN/HMM and encoder-decoder-attention model results on LibriSpeech with different training corpus sizes. *train-clean-100* is a official subset of the training corpus. *train-960* is the complete training corpus. (Clustered) context-dependent phones (CDp) are utilized for the hybrid model, and sub-word units for the attention model.

training set	model	LM	WER [%]			
			dev		test	
			clean	other	clean	other
train-clean-100	hybrid	4-gram	5.0	19.5	5.8	18.6
	attention	none	14.7	38.5	14.7	40.8
train-960	hybrid	4-gram	4.0	9.6	4.4	10.0
	attention	none	4.7	14.3	4.8	15.4

We compare the hybrid model with the encoder-decoder-attention model. We trained both models on the *train-clean-100* training subset and on the *train-960* complete training set.

Table 4: The WER results from our most interesting models and important results from other papers on LibriSpeech 960 h. CDp are (clustered) context-dependent phones. BPE are sub-word units. 4-gr LM is the official 4-gram word LM. GCNN are gated convolutional NN. RNN are recurrent NN.

paper	model	label unit		LM	WER [%]			
		AM	LM		dev		test	
					clean	other	clean	other
Han et al. [3]	hybrid, seq. disc., single	CDp	word	RNN	3.0	8.8	3.6	8.7
	hybrid, seq. disc., ensemble				2.6	7.6	3.2	7.6
Zeghidour et al. [8]	end-to-end GCNN	chars	words	GCNN	3.2	10.1	3.4	11.2
Irie et al. [9]	end-to-end attention	Word Piece Model		LSTM	3.3	10.3	3.6	10.3
Zeyer et al. [5]		BPE			LSTM	3.5	11.5	3.8
this work				None	4.3	12.9	4.4	13.5
				LSTM	2.9	8.9	3.2	9.9
		Transformer	<b>2.6</b>	<b>8.4</b>	<b>2.8</b>	<b>9.3</b>		
this work	hybrid	CDp	word	4-gr	4.0	9.6	4.4	10.0
	hybrid, seq. disc.			4-gr	3.4	8.3	3.8	8.8
				+ LSTM	<b>2.2</b>	<b>5.1</b>	<b>2.6</b>	<b>5.5</b>
				Transformer resc.	<b>1.9</b>	<b>4.5</b>	<b>2.3</b>	<b>5.0</b>
Park et. al. [10]	end-to-end attention/SpecAugment	Word Piece Model		LSTM	-	-	2.5	5.8

These are not the best models but utilize a baseline model for both approaches. The hybrid DNN/HMM model outperforms the encoder-decoder-attention model constantly. But the difference in performance shrinks substantially with the much larger training set.

Our encoder-decoder-attention model in combination with a Transformer LM gives a WER of 3.2% on *test-clean* and 9.9% on *test-other* (Table 4). Evaluating our sequence discriminatively trained acoustic model with our LSTM LM results in a WER of 2.6% on *test-clean* and 5.5% on *test-other*. Rescoring with a Transformer language model further improves the performance of our hybrid DNN/HMM system resulting in a WER of 2.3% on *test-clean* and 5.0% on *test-other*. The previous best hybrid system was presented in [3] while the best end-to-end system without data augmentation was presented in [8, 9] (Table 4). Additionally we present the best end-to-end system with data augmentation [10]. Our best encoder-decoder-attention model improves the state-of-the-art for end-to-end models without data augmentation by 17.6% relative WER on *test-clean* and by 3.9% relative WER on *test-other*. Our best hybrid DNN/HMM system without Transformer LM rescoring improves the state-of-the-art by 18.8% relative WER on *test-clean* and by 27.6% relative WER on *test-other*. If we add rescoring with a Transformer LM we improve further by 11.5% relative WER on *test-clean* and by 9.1% relative WER on *test-other*. In comparison, the hybrid DNN/HMM system still outperforms the encoder-decoder-attention system by over 15% relative WER on *test-clean* and by over 40% relative WER on *test-other*. Our best hybrid model even outperforms the end-to-end attention model with SpecAugment [10] by 8% relative WER on *test-clean* and by 13.8% relative WER on *test-other*. These results reflect the state-of-the-art performance for both hybrid and attention-based models on LibriSpeech, to the best of the authors’ knowledge.

WERs become very small, especially for *dev-clean* and *test-clean*. When analyzing the errors, it is noticeable that some of the errors would not be recognized as primary errors by a human. These can be categorized as, for example: word contractions or American vs British English spelling. Examples of such errors are: I am ↔ I’m, tyrannise ↔ tyrannize, color ↔ colour, oh ↔ o. So far we have not employed a normalization

strategy for these errors.

## 6. Conclusions

In this paper we presented two ASR systems for the LibriSpeech task. One System was a hybrid DNN/HMM system based on a GMM/HMM system, the other system was an attention-based encoder-decoder system.

We described how we built the systems and described how to incrementally improve the systems to get competitive results. For the hybrid DNN/HMM system a large NN acoustic model, the sequence discriminative training and the employment of an LSTM LM was important for the good performance. The encoder-decoder-attention approach utilized an extended pre-training variant and a tuned curriculum learning schedule. This enabled the model to achieve competitive results in comparison to other end-to-end approaches.

The presented encoder-decoder-attention system showed state-of-the-art performance on the LibriSpeech 960h task in comparison with end-to-end systems without data augmentation. But our comparison shows that on the LibriSpeech 960h task, the hybrid DNN/HMM system outperforms the attention-based system by 15% relative on the clean and 40% relative on the other test sets. Our hybrid system even outperforms previous results presented in the literature. Moreover, experiments on a reduced 100h-subset of the LibriSpeech training corpus even show a more pronounced margin between the hybrid DNN/HMM and attention-based architectures. To the best knowledge of the authors, the results obtained when training on the full LibriSpeech training set, are the best published currently, both for the hybrid DNN/HMM and the attention-based systems presented in this work.

## Acknowledgements

This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 694537, project “SEQCLAS”) and from a Google Focused Award. The work reflects only the authors’ views and none of the funding parties is responsible for any use that may be made of the information it contains.

Experiments were partially performed with computing resources granted by RWTH Aachen University under project nova0003.

We thank Wei Zhou for help with generating lattices.

## References

- [1] G. Hinton, L. Deng, D. Yu, G. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, B. Kingsbury *et al.*, “Deep neural networks for acoustic modeling in speech recognition,” *IEEE Signal processing magazine*, vol. 29, 2012.
- [2] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *Proc. ICASSP*, Brisbane, Australia, Apr. 2015.
- [3] K. J. Han, A. Chandrashekar, J. Kim, and I. Lane, “The CAPIO 2017 conversational speech recognition system,” *arXiv preprint arXiv:1801.00059*, 2017.
- [4] N. Kanda, Y. Fujita, and K. Nagamatsu, “Lattice-free state-level minimum bayes risk training of acoustic models,” in *Proc. Interspeech*, Hyderabad, India, 2018.
- [5] A. Zeyer, K. Irie, R. Schlüter, and H. Ney, “Improved training of end-to-end attention models for speech recognition,” in *Interspeech*, Hyderabad, India, Sep. 2018.
- [6] A. Zeyer, A. Merboldt, R. Schlüter, and H. Ney, “A comprehensive analysis on attention models,” in *Interpretability and Robustness in Audio, Speech, and Language (IRASL) Workshop, Conference on Neural Information Processing Systems (NeurIPS)*, Montreal, Canada, Dec. 2018.
- [7] S. Sabour, W. Chan, and M. Norouzi, “Optimal completion distillation for sequence learning,” *arXiv preprint arXiv:1810.01398*, 2018.
- [8] N. Zeghidour, Q. Xu, V. Liptchinsky, N. Usunier, G. Synnaeve, and R. Collobert, “Fully convolutional speech recognition,” *arXiv preprint arXiv:1812.06864*, 2018.
- [9] K. Irie, R. Prabhavalkar, A. Kannan, A. Bruguier, D. Rybach, and P. Nguyen, “Model unit exploration for sequence-to-sequence speech recognition,” *preprint arXiv:1902.01955*, 2019.
- [10] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” *arXiv preprint arXiv:1904.08779*, 2019.
- [11] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [12] A. Graves, N. Jaitly, and A.-r. Mohamed, “Hybrid speech recognition with deep bidirectional LSTM,” in *Proc. ASRU*, 2013, pp. 273–278.
- [13] A. Zeyer, P. Doetsch, P. Voigtlaender, R. Schlüter, and H. Ney, “A comprehensive study of deep bidirectional lstm rnns for acoustic modeling in speech recognition,” in *Proc. ICASSP*, New Orleans, LA, USA, Mar. 2017.
- [14] R. Schlüter, I. Bezrukov, H. Wagner, and H. Ney, “Gamma-tone features and feature combination for large vocabulary speech recognition,” in *Proc. ICASSP*, Honolulu, HI, USA, Apr. 2007.
- [15] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [16] T. Dozat, “Incorporating nesterov momentum into Adam,” Stanford University, Tech. Rep., 2015. [Online]. Available: [http://cs229.stanford.edu/proj2015/054\\_report.pdf](http://cs229.stanford.edu/proj2015/054_report.pdf)
- [17] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: a simple way to prevent neural networks from overfitting,” *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [18] A. Neelakantan, L. Vilnis, Q. V. Le, I. Sutskever, L. Kaiser, K. Kurach, and J. Martens, “Adding gradient noise improves learning for very deep networks,” *arXiv preprint arXiv:1511.06807*, 2015.
- [19] M. Gibson and T. Hain, “Hypothesis spaces for minimum bayes risk training in large vocabulary speech recognition,” in *Proc. Interspeech*, Jan. 2006.
- [20] R. Kneser and H. Ney, “Improved backing-off for m-gram language modeling,” in *Proc. ICASSP*, Detroit, MI, USA, May 1995.
- [21] M. Sundermeyer, R. Schlüter, and H. Ney, “LSTM neural networks for language modeling,” in *Proc. Interspeech*, Portland, OR, USA, Sep. 2012.
- [22] E. Beck, W. Zhou, R. Schlüter, and H. Ney, “Lstm language models for lvsr in first-pass decoding and lattice-rescoring,” *arXiv preprint arXiv:1907.01030*, 2019.
- [23] A. Zeyer, T. Alkhouli, and H. Ney, “RETURNN as a generic flexible neural toolkit with application to translation and speech recognition,” in *Annual Meeting of the Assoc. for Computational Linguistics*, Melbourne, Australia, Jul. 2018.
- [24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Proc. NeurIPS*, Long Beach, CA, USA, Dec. 2017.
- [25] K. Irie, A. Zeyer, R. Schlüter, and H. Ney, “Language modeling with deep Transformers,” in *Proc. Interspeech*, Graz, Austria, Sep. 2019.
- [26] M. Sundermeyer, Z. Tüske, R. Schlüter, and H. Ney, “Lattice decoding and rescoring with long-span neural network language models,” in *Proc. Interspeech*, Singapore, Sep. 2014.
- [27] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [28] M.-T. Luong, H. Pham, and C. D. Manning, “Effective approaches to attention-based neural machine translation,” *arXiv preprint arXiv:1508.04025*, 2015.
- [29] M. X. Chen, O. Firat, A. Bapna, M. Johnson, W. Macherey, G. Foster, L. Jones, M. Schuster, N. Shazeer, N. Parmar *et al.*, “The best of both worlds: Combining recent advances in neural machine translation,” in *Proc. ACL*, vol. 1, 2018.
- [30] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *ICASSP*, 2016.
- [31] P. Doetsch, A. Zeyer, and H. Ney, “Bidirectional decoder networks for attention-based end-to-end offline handwriting recognition,” in *International Conference on Frontiers in Handwriting Recognition*, Shenzhen, China, Oct. 2016.
- [32] C.-C. Chiu, T. N. Sainath, Y. Wu, R. Prabhavalkar, P. Nguyen, Z. Chen, A. Kannan, R. J. Weiss, K. Rao, K. Gonina *et al.*, “State-of-the-art speech recognition with sequence-to-sequence models,” *arXiv preprint arXiv:1712.01769*, 2017.
- [33] E. Battenberg, J. Chen, R. Child, A. Coates, Y. Gaur, Y. Li, H. Liu, S. Satheesh, A. Sriram, and Z. Zhu, “Exploring neural transducers for end-to-end speech recognition,” in *Proc. ASRU*, Okinawa, Japan, Dec. 2017.
- [34] R. Sennrich, B. Haddow, and A. Birch, “Neural machine translation of rare words with subword units,” in *ACL*, Berlin, Germany, August 2016.
- [35] Ç. Gülçehre, O. Firat, K. Xu, K. Cho, L. Barrault, H.-C. Lin, F. Bougares, H. Schwenk, and Y. Bengio, “On using monolingual corpora in neural machine translation,” *Computer Speech & Language*, vol. 45, pp. 137–148, Sep. 2017.
- [36] S. Toshniwal, A. Kannan, C.-C. Chiu, Y. Wu, T. N. Sainath, and K. Livescu, “A comparison of techniques for language model integration in encoder-decoder speech recognition,” in *Proc. SLT*, Athens, Greece, Dec. 2018.
- [37] A. Hannun, A. Lee, Q. Xu, and R. Collobert, “Sequence-to-sequence speech recognition with time-depth separable convolutions,” *arXiv preprint arXiv:1904.02619*, 2019.
- [38] S. Wiesler, A. Richard, P. Golik, R. Schlüter, and H. Ney, “RASR/NN: The RWTH neural network toolkit for speech recognition,” in *Proc. ICASSP*, Florence, Italy, May 2014.
- [39] P. Doetsch, A. Zeyer, P. Voigtlaender, I. Kulikov, R. Schlüter, and H. Ney, “RETURNN: the RWTH extensible training framework for universal recurrent neural networks,” in *Proc. ICASSP*, New Orleans, LA, USA, Mar. 2017, pp. 5345–5349.
- [40] <https://github.com/rwth-i6/returnn-experiments/tree/master/2019-librispeech-system>, [Online; accessed 1-July-2019].