BUILDING COMPETITIVE DIRECT ACOUSTICS-TO-WORD MODELS FOR ENGLISH CONVERSATIONAL SPEECH RECOGNITION

Kartik Audhkhasi, Brian Kingsbury, Bhuvana Ramabhadran, George Saon, Michael Picheny

IBM T. J. Watson Research Center, Yorktown Heights, New York

ABSTRACT

Direct acoustics-to-word (A2W) models in the end-to-end paradigm have received increasing attention compared to conventional subword based automatic speech recognition models using phones, characters, or context-dependent hidden Markov model states. This is because A2W models recognize words from speech without any decoder, pronunciation lexicon, or externally-trained language model, making training and decoding with such models simple. Prior work has shown that A2W models require orders of magnitude more training data in order to perform comparably to conventional models. Our work also showed this accuracy gap when using the English Switchboard-Fisher data set. This paper describes a recipe to train an A2W model that closes this gap and is at-par with stateof-the-art sub-word based models. We achieve a word error rate of 8.8%/13.9% on the Hub5-2000 Switchboard/CallHome test sets without any decoder or language model. We find that model initialization, training data order, and regularization have the most impact on the A2W model performance. Next, we present a joint wordcharacter A2W model that learns to first spell the word and then recognize it. This model provides a rich output to the user instead of simple word hypotheses, making it especially useful in the case of words unseen or rarely-seen during training.

Index Terms— End-to-end models, direct acoustics-to-word models, automatic speech recognition, deep learning

1. INTRODUCTION

Conventional sub-word based automatic speech recognition (ASR) typically involves three models - acoustic model (AM), pronunciation model and decision tree (PMT), and language model (LM) [1]. The AM computes the probability $p_{AM}(\mathbf{A}|\mathbf{s})$ of the acoustics **A** given the sub-word units **s**. The PMT models the probability $p_{PMT}(\mathbf{w}|\mathbf{s})$ of the word sequence **w** given the sub-word unit sequence **s**. The LM acts as a prior $p_{LM}(\mathbf{w})$ on the word sequence **w**. Hence, finding the most likely sequence of words given the acoustics **A** becomes a maximum aposteriori optimization problem over the following probability density function:

$$p_{\text{ASR}}(\mathbf{w}|\mathbf{A}) \propto p_{\text{LM}}(\mathbf{w})p(\mathbf{A}|\mathbf{w})$$
 (1)

$$= p_{\rm LM}(\mathbf{w}) \sum_{\mathbf{s}} p_{\rm AM}(\mathbf{A}|\mathbf{s}, \mathbf{w}) p_{\rm PMT}(\mathbf{s}|\mathbf{w}) \qquad (2)$$

$$\approx p_{\rm LM}(\mathbf{w}) \sum_{\mathbf{s}} p_{\rm AM}(\mathbf{A}|\mathbf{s}) p_{\rm PMT}(\mathbf{s}|\mathbf{w}) . \tag{3}$$

While both the AM and LM in modern ASR systems use deep neural networks (DNNs) and their variants [2], the PMT is usually based on

decision trees and finite state transducers. Training the AM requires alignments between the acoustics and sub-word units, and several iterations of model training and re-alignment. Recent end-to-end (E2E) models have obviated the need for aligning the sub-word units to the acoustics. Popular E2E models include recurrent neural networks (RNNs) trained with the connectionist temporal classification (CTC) loss function [3–10] and the attention-based encoder-decoder RNNs [11–15]. These approaches are not truly E2E because they still use sub-word units, and hence require a decoder and separately-trained LM to perform well.

In contrast, the recently-proposed direct acoustics-to-word (A2W) models [16, 17] train a single RNN to directly optimize $p_{\text{ASR}}(\mathbf{w}|\mathbf{A})$. This eliminates the need for sub-word units, pronunciation model, decision tree, decoder, and externally-trained LM, which significantly simplifies the training and decoding process. However, prior research on A2W models has shown that such models require several orders of magnitude more training data when compared with conventional sub-word based models. This is because the A2W models need sufficient acoustic training examples per word to train well. For example, [16] used more than 125,000 hours of speech to train an A2W model with a vocabulary of nearly 100,000 words that matched the performance of a state-of-the-art CD state-based CTC model. Our prior work [17] explored A2W models for the well-known English Switchboard task and presented a few initialization techniques to effectively train such models with only 2000 hours of data. However, we still observed a gap of around 3-4% absolute in WER between the Switchboard phone CTC and the A2W models on the Hub5-2000 evaluation set.

This paper further improves the state-of-the-art in A2W models for English conversational speech recognition. We present a training recipe that achieves WER of 8.8%/13.9% on the Switchboard/CallHome subsets of the Hub5-2000 evaluation set, compared to our previous best result of 13.0%/18.8% [17]. These new results are at par with several state-of-the-art models that use sub-word units, a decoder, and a LM. We quantify the gains made by each ingredient of our training recipe and conclude that model initialization, training data order, and regularization are the most important factors.

Next, we turn our attention to the issue of data sparsity while training A2W models. The conventional solution to this problem uses a sub-word unit-based model that needs a decoder and LM during testing. As an alternative, we propose the spell and recognize (SAR) CTC model that learns to first spell the word into its character sequence and then recognize it. Not only does this model retain all advantages of a direct A2W model, it also provides rich hypotheses to the user which are readable especially in the case of unseen or rarely-seen words. We illustrate the benefits of this model for out-of-vocabulary (OOV) words.

The next section discusses the baseline A2W model [17]. Section 3 discusses the proposed training recipe, an analysis of the impact of individual ingredients, and the results. Section 4 presents

The authors thank Ewout Van Den Berg and Andrew Rosenberg of IBM, and Florian Metze of Carnegie Mellon University for useful discussions.

our SAR model that jointly models words and characters. The paper concludes in Section 5 with a summary of findings.

2. BASELINE ACOUSTICS-TO-WORD MODEL

2.1. CTC Loss

Conventional losses used for training neural networks, such as crossentropy, require a one-to-one mapping (or alignment) between the rows of the $T \times D$ input feature vector matrix **A** and length-*L* output label sequence **y**. The connectionist temporal classification (CTC) loss relaxes this requirement by considering all possible alignments. It introduces a special *blank symbol* ϕ that expands the length-*L* target label sequence **y** to multiple length-*T* label sequences $\tilde{\mathbf{y}}$ containing ϕ , such that $\tilde{\mathbf{y}}$ maps to **y** after removal of all repeating symbols and ϕ . The CTC loss is then

$$p(\mathbf{y}|\mathbf{A}) = \sum_{\tilde{\mathbf{y}}\in\Omega(\mathbf{y})} p(\tilde{\mathbf{y}}|\mathbf{A}) = \sum_{\tilde{\mathbf{y}}\in\Omega(\mathbf{y})} \prod_{t=1}^{T} p(\tilde{y}_t|\mathbf{a}_t) , \qquad (4)$$

where the set $\Omega(\mathbf{y})$, is the set of CTC paths for \mathbf{y} , and \mathbf{a}_t/\tilde{y}_t denote the t^{th} elements of the sequences. A forward-backward algorithm efficiently computes the above loss function and its gradient, which is then back-propagated through the neural network [3]. The next section describes the baseline A2W model [17].

2.2. Baseline A2W Model

We used two standard training data sets for our experiments. The "300-hour" set contained 262 hours of segmented speech from the Switchboard-1 audio with transcripts provided by Mississippi State University. The "2000-hour" set contained an additional 1698 hours from the Fisher data collection and 15 hours from CallHome audio.

We extracted 40-dimensional logMel filterbank features over 25 ms frames every 10 ms from the input speech signal. We used stacking+decimation [6], where we stacked two successive frames and dropped every alternate frame during training. This resulted in a stream of 80-dimensional acoustic feature vectors at half the frame rate of the original stream. The baseline models also used 100-dimensional i-vectors [18] for each speaker, resulting in 180-dimensional acoustic feature vectors.

The baseline A2W model consisted of a 5-layer bidirectional LSTM (BLSTM) RNN with a 180-dimensional input and 320dimensional hidden layers in the forward and backward directions. We picked words with at least 5 occurrences in the training data in the vocabulary. This resulted in a 10,000-word output layer for the 300-hour A2W system and a 25,000-output layer for the 2000-hour system. As noted in [17], initialization is crucial to training an A2W model. Thus, we initialized the A2W BLSTM with the BLSTM from a trained phone CTC model, and the final linear layer using word embeddings trained using GloVe [19].

Table 1 gives the WERs of the baseline A2W and phone CTC models reported in [17] on the Hub5-2000 Switchboard (SWB) and CallHome (CH) test sets. We performed the decoding of the A2W models via simple peak-picking over the output word posterior distribution, and removing repetitions and blank symbols. The phone CTC model used a full decoding graph and a LM. We observe that the 2000-hour A2W model lags behind the phone CTC model by 3.4%/2.8% absolute WER on SWB/CH, and the gap is much bigger for the 300-hour models. We next discuss our new training recipe.

Table 1. This table shows the Switchboard/CallHome WERs for our baseline 300/2000-hour A2W and phone CTC models from [17].

Data (hrs)	AM	LM	SWB	CH
300	A2W	-	20.8	30.4
300	Phone CTC	Small	14.5	25.1
2000	A2W	-	13.0	18.8
2000	Phone CTC	Big	9.6	16.0

3. UPDATED TRAINING RECIPE

Our prior experience with training A2W models led us to conclude that model initialization and regularization are important aspects of training such models. One key reason for this is the fact that A2W models attempt to solve the difficult problem of directly recognizing words from acoustics with a single neural network. Hence, our previously-proposed strategy of initializing the A2W BLSTM with the phone CTC BLSTM and the final linear layer with word embeddings gave WER gains. In this work, we started by exploring several other strategies in the same spirit. All our new experiments were conducted in PyTorch [20] with the following changes compared to [17]:

- We included delta and delta-delta coefficients because they slightly improved the WER. Hence, the total acoustic vector was of size 340 after stacking+decimation and appending the 100-dimensional i-vectors.
- We initialized all matrices to samples of a uniform distribution over (-ε, ε) where ε is the inverse of the √fan-in or the size of the input vector [21]. This takes the dimensionality of the input vectors at each layer in account, which improved convergence compared to our old strategy of using ε = 0.01.
- In place of new-bob annealing [17], we kept a fixed learning rate for the first 10 epochs and decayed it by $\sqrt{0.5}$ every epoch.

3.1. Training Data Order

Training data order is an important consideration for sequence-tosequence models such as E2E ASR systems because such models operate on the entire input and output sequences. All training sequences have to be padded to the length of the longest sequence in the batch in order to do GPU tensor operations. Random sequence order during batch creation is not memory-efficient because batches will contain a larger range of sequence lengths, which will lead to more wasteful padding on average. Hence, sequences have to be sorted before batch creation.

We compare the impact of sorting input acoustic sequences in order of ascending and descending length in Table 2. Our results show that ascending order gives significantly better WER than sorting in descending order. The intuition behind this result is that shorter sequences are easier to train on initially, which enables the network to reach a better point in the parameter space. This can be regarded as an instance of curriculum learning [22].

3.2. Momentum and Dropout

We also experimented with Nesterov momentum-based stochastic gradient descent (SGD), which has been shown to give better convergence compared to simple SGD on several tasks. We use the

Table 2. This table shows the impact of various training strategies on the 300-hour A2W model.

Model	SWB	CH
Descending Order	23.0	30.7
Ascending Order	18.3	28.1
+Momentum	18.0	27.4
+Dropout	17.4	26.9
+Output Projection	16.9	26.3
+Phone BLSTM Init.	14.9	23.8
+Bigger Model	14.6	23.6
Previous best [17]	20.8	30.4

following parameter updates:

$$\mathbf{v}_n = \rho \mathbf{v}_{n-1} + \lambda \nabla f(\Theta_{n-1} + \rho v_{n-1}) \tag{5}$$

$$\Theta_n = \Theta_{n-1} - \mathbf{v}_n , \qquad (6)$$

where \mathbf{v}_n is the *velocity* or a running weighted-sum of the gradient of the loss function f. The constant ρ is usually set to 0.9 and λ is the learning rate, set to 0.01 in our experiments. We also experimented with a dropout of 0.25 in order to prevent over-fitting. Table 2 shows that both momentum and dropout improve the WER.

3.3. Output Projection Layer

In contrast with phone or character-based CTC models, A2W models have a large output size equal to the size of the vocabulary. Prior research [23] has shown that decomposing the output linear layer of size $V \times D$ into two layers of sizes $V \times d$ and $d \times D$ with d < Dspeeds-up model training due to reduced number of parameters. We experimented with a projection layer of size 256 and found that it speeds-up training by a factor of 1.2x and also slightly improves the WER, which we attribute to a reduction in over-fitting.

3.4. Phone BLSTM Initialization and Bigger Model

We finally initialized our model with the phone CTC BLSTM which gave improvements in our previous work [17]. As expected, this initialization lowered the WER, despite the presence of all the above useful strategies. With dropout in place, we trained a bigger 6-layer model with 512-dimensional BLSTM and saw slight gains in WER.

3.5. Final Model and Results

We initialized the 2000-hour A2W model with the best 300-hour A2W model and used the same recipe for training. Table 3 shows the WER of the resulting system along with our previous best A2W WER and several other published results. We obtained a significant improvement of 4.2%/4.9% absolute WER compared to our previous result. We also see that our direct A2W is at par with most hybrid CD state-based and E2E models, while utilizing no decoder or LM. As noted in [18], the CallHome test set is more challenging than Switchboard because 36 out of 40 speakers in the latter appear in the training set. The results on the CallHome test set are especially good, where our A2W model matches the best result obtained using a hybrid BLSTM [18] that used exactly the same acoustic features¹.

Table 3. This table shows the WER of our current and previous-best A2W model trained on the 2000-hour Switchboard+Fisher set. We have also included several other published results for comparison. Results with * use data augmentation on the 2000-hour training set.

Model	Output	Decoder/	SWB	CH
	Units	LM		
BLSTM+LF MMI [24]	CD state	Y/Y	8.5	15.3
LACE+LF MMI [25]	CD state	Y/Y	8.3	14.8
Dilated Conv. [26]	CD state	Y/Y	7.7	14.5
BLSTM [18]	CD state	Y/Y	7.7	13.9
Iterated CTC [27]	Char	Y/Y	11.3	18.7
CTC* [28]	Char	Y/Y	9.0	17.7
Gram-CTC* [28]	Char N-gm	Y/Y	7.9	15.8
CTC+Gram-CTC* [28]	Char N-gm	Y/Y	7.3	14.7
RNN-T* [29]	Char	Y/N	8.5	16.4
Attention [*] [29]	Char	Y/N	8.6	17.8
CTC A2W [17]	Word	N/N	13.0	18.8
CTC A2W (current)	Word	N/N	8.8	13.9

We also note that our A2W model uses a vocabulary of 25,000 words which has an OOV rate of 0.5%/0.8% on the SWB/CH test sets. All the other models used much bigger vocabularies and hence did not suffer from OOV-induced errors.

3.6. Ablation Study

In order to understand the impact of individual components of our recipe, we conducted an ablation study on our best 300-hour A2W model. We removed each component of the recipe while keeping others fixed, trained a model, and decoded the test sets. Figure 1 shows the results of this experiment. We observed that changing the training data order from ascending to descending order of length resulted in the biggest drop in performance. The second biggest factor was dropout - excluding it leads to over-training because the heldout loss rises after epoch 10. Choosing a smaller 5-layer model instead of 6-layer led to the next largest drop in WER. Finally, as expected, excluding the projection layer had the least impact on WER.



Fig. 1. This figure shows the results of the ablation experiment on our best 300-hour A2W model. The numbers in parentheses in the legend give the SWB/CH WERs.

¹Adding additional FMLLR features gives a WER of 7.2%/12.7% [18].

Despite strong results, the A2W model does not give any meaningful output to the user in case of OOV words but simply emits an "UNK" tag. This is not a big problem for the Switchboard task because the OOV rates for the SWB/CH test sets are 0.5%/0.8% on the 25,000 word vocabulary. But other tasks might be affected by the limited vocabulary. As a solution, the next section discusses our joint word-character model that aims to provide the user with a richer output that is especially useful for unseen or rarely-seen words.

4. SPELL AND RECOGNIZE MODEL

The advantage of the A2W model is that it directly emits word hypotheses by forward-passing acoustic features through a RNN without needing a decoder or externally-trained LM. However, its vocabulary is fixed and OOV words cannot be recognized by the system. Furthermore, words infrequently seen in the training data are not recognized well by the network due to insufficient training examples. Prior approaches to dealing with the above limitations completely rely on sub-word units. This includes work on character models [9, 27] and N-grams [28], RNN-Transducer [29], RNN-Aligner [30].

In contrast, our approach is to have the best of both worlds by combining the ease of decoding a A2W model with the flexibility of recognizing unseen/rarely-seen words with a character-based model. One natural candidate is a multi-task learning (MTL) model containing a shared lower network, and two output networks corresponding to the two tasks - recognizing words and characters. However, such an MTL network is not suitable for our purpose because the recognized word and character sequences for an input speech utterance are not guaranteed to be synchronized in time. This is because the CTC loss does not impose time-alignment on the output sequence.

The proposed spell and recognize (SAR) model circumvents this alignment problem by presenting training examples that contain both words and characters. This allows us to continue to leverage an A2W framework without resorting to more complex graph-based decoding methodologies employed in, for example, word-fragment based systems [31–33]. Consider the output word sequence "THE CAT IS BLACK" The SAR model uses the following target sequence:

b-t h e-e THE b-c a e-t CAT b-i e-s IS b-b l a c e-k BLACK

where lowercase alphabets are the character targets, and b-/edenote special prefixes for word beginning and end. Hence, the model is trained to first spell the word and then recognize it. In contrast with a MTL model, the SAR model has a single softmax over words+characters in the output layer.

4.1. Choice of Character Set

We experimented with two character sets for the SAR model. The first one is the simple character set consisting of a total of 41 symbols - alphabets a-z, digits 0-9, whitespace _, and other punctuations. The second character set is the one used in [5], and includes separate character variants depending on position in a word - beginning, middle, and end. It also includes special symbols for repeated characters, e.g. a separate symbol for *ll*. The intuition behind this character set is that its symbols capture more context as compared to the simple set, and also disambiguate legitimate character repetitions from double peaks emitted by the CTC model. We observed that the performance with the latter character set is slightly better than using simple characters. Hence, we present results only for this case.

4.2. Experiments and Results

We restricted ourselves to the 300-hour set for experiments on the SAR model because it uses a 10,000 word vocabulary, leading to

a higher OOV rate than the 2000-hour set and also contains several rare words. We trained a 6-layer BLSTM with joint word and character targets after preparing the output training sequences as described in the previous section. We initialized the SAR BLSTM using the A2W BLSTM. The training recipe was the same as for the A2W model presented previously. The SAR model permits three decodes:

- Word: Use only word predictions, similar to the A2W model.
- **Characters**: Use only character predictions, and combine them into words using the word-begin characters.
- **Switched**: Use the character predictions up to the previous word when the model predicts an "UNK" symbol, and use the word prediction everywhere else.

 Table 4. This table shows performance of the 300-hour SAR model for three types of decodes.

Decode	SWB	CH
Word	14.5	23.9
Characters	18.9	30.9
Switched	14.4	24.0

We observe that the SAR model gives comparable performance to the baseline A2W model, but additionally gives meaningful output for OOVs, as illustrated by the following test set examples:

<i>REF</i> : SUCH AS LIKE (%HESITATION) THE MURDERING OF A COP OR <i>HYP</i> : b-s u c e-h SUCH _ b-a e-s AS _ b-l i k e-e LIKE _ b-t h e-e THE _ b-u e-u UH _ b-m u r d e r i n e-g UNK _ b-o e-f OF _ b-a A _ b-s o e-p COP _ b-o e-r OR
REF: THAT IS RIGHT WE ARE WE ARE FURTHERING HIGHER
<i>HYP</i> : b-t h a e-s THAT'S _ b-r i g h e-t RIGHT _ b-r h I RIGHT _ b-w e ' r e-e
WE'RE _ b-w e ' r e-e WE'RE _ b-f u r t h e r i n e-g UNK _ b-h i g h e e-r
HIGHER
<i>REF</i> : BUT SOMETIMES LIKE I JUST HAD TO DO THIS SUMMARY OF
THIS ONE YOU KNOW THESE SCHOLARLY JOURNALS AND STUFF
HYP: b-b u e-t BUT _ b-s o m e t i m e e-s SOMETIMES _ b-l i k e-e LIKE _ b-i I
b-j u s e-t JUST _ b-h a e-d HAD _ b-t e-o TO _ b-d e-o DO _ b-t h e-e THIS _ b-s
u mm e r e-y SUMMARY _ b-t h i e-s THIS _ b-o n e-e ONE _ b-y o e-u YOU _
b-k n o e-w KNOW _ b-t h e s e-e THESE _ b-c o l a r l e-y UNK _ b-j o u r n a l
e-s UNK _ b-a n e-d AND _ b-s t u e-2f STUFF

The words in bold are OOVs. We observe that the SAR model emits the UNK tag in these cases, but the characters preceding it contain the spelling of the word. In some cases, this spelling is incorrect, e.g. "SCHOLARLY \rightarrow COLARLY", but still is more meaningful to the user than the UNK tag. Future research will try to fix these errors using data-driven methods.

5. CONCLUSION

Conventional wisdom and prior research suggests that direct acousticto-word (A2W) models require orders of magnitude more data than sub-word unit-based models to perform competitively. This paper presents a recipe to train an A2W model on the 2000-hour Switchboard+Fisher data set that performs at-par with several stateof-the-art hybrid and end-to-end models using sub-word units. We conclude that data order, model initialization, and regularization are crucial to obtaining a competitive A2W model with a WER of 8.8%/13.9% on the Switchboard/CallHome subsets of the Hub5-2000 test set. Next, we present a spell and recognize (SAR) model that learns to first spell a word and then recognize it. The proposed SAR model gives a rich and readable output to the user while maintaining the training/decoding simplicity and performance of a A2W model. We show some examples illustrating the SAR model's benefit for utterances containing OOV words.

6. REFERENCES

- [1] F. Jelinek, *Statistical methods for speech recognition*, MIT press, 1997.
- [2] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, and B. Kingsbury, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [3] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proc. ICML*, 2006, pp. 369–376.
- [4] A. Graves and N. Jaitly, "Towards end-to-end speech recognition with recurrent neural networks," in *Proc. ICML*, 2014, vol. 14, pp. 1764–1772.
- [5] G. Zweig, C. Yu, J. Droppo, and A. Stolcke, "Advances in allneural speech recognition," in *Proc. ICASSP*, 2017, pp. 4805– 4809.
- [6] H. Sak, A. Senior, K. Rao, and F. Beaufays, "Fast and accurate recurrent neural network acoustic models for speech recognition," in *Proc. Interspeech*, 2015.
- [7] Y. Miao, M. Gowayyed, and F. Metze, "EESEN: End-to-end speech recognition using deep RNN models and WFST-based decoding," in *Proc. ASRU*, 2015, pp. 167–174.
- [8] Y. Miao, M. Gowayyed, X. Na, T. Ko, F. Metze, and A. Waibel, "An empirical exploration of CTC acoustic models," in *Proc. ICASSP*, 2016, pp. 2623–2627.
- [9] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates, and A. Y. Ng, "Deep speech: Scaling up end-to-end speech recognition," arXiv preprint arXiv:1412.5567, 2014.
- [10] D. Amodei et al., "Deep speech 2: End-to-end speech recognition in English and Mandarin," in *Proc. ICML*, 2016, pp. 173–182.
- [11] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, "End-to-end attention-based large vocabulary speech recognition," in *Proc. ICASSP*, 2016, pp. 4945–4949.
- [12] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *Proc. ICLR*, 2015.
- [13] A. L. Maas, Z. Xie, D. Jurafsky, and A. Y. Ng, "Lexicon-free conversational speech recognition with neural networks," in *Proc. HLT-NAACL*, 2015, pp. 345–354.
- [14] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *Proc. ICASSP*, 2016, pp. 4960–4964.
- [15] W. Chan, Y. Zhang, Q. Le, and N. Jaitly, "Latent sequence decompositions," arXiv preprint arXiv:1610.03035, 2016.
- [16] H. Soltau, H. Liao, and H. Sak, "Neural Speech Recognizer: Acoustic-to-Word LSTM Model for Large Vocabulary Speech Recognition," in *Proc. Interspeech*, 2017.
- [17] K. Audhkhasi, B. Ramabhadran, G. Saon, M. Picheny, and D. Nahamoo, "Direct acoustics-to-word models for English conversational speech recognition," in *Proc. Interspeech*, 2017, pp. 959–963.

- [18] G. Saon, G. Kurata, T. Sercu, K. Audhkhasi, S. Thomas, D. Dimitriadis, X. Cui, B. Ramabhadran, M. Picheny, L. Lim, B. Roomi, and P. Hall, "English conversational telephone speech recognition by humans and machines," in *Proc. Interspeech*, 2017.
- [19] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global Vectors for Word Representation," in *Proc. EMNLP*, 2014, vol. 14, pp. 1532–1543.
- [20] "PyTorch," https://github.com/pytorch/ pytorch.
- [21] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. AISTATS*, 2010, pp. 249–256.
- [22] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proc. ICML*, 2009.
- [23] T. N. Sainath, B. Kingsbury, V. Sindhwani, E. Arisoy, and B. Ramabhadran, "Low-rank matrix factorization for deep neural network training with high-dimensional output targets," in *Proc. ICASSP*, 2013, pp. 6655–6659.
- [24] D. Povey, V. Peddinti, D. Galvez, P. Ghahremani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely Sequence-Trained Neural Networks for ASR Based on Lattice-Free MMI," in *Proc. Interspeech*, 2016, pp. 2751–2755.
- [25] W. Xiong, J. Droppo, X. Huang, F. Seide, M. Seltzer, A. Stolcke, D. Yu, and G. Zweig, "Achieving human parity in conversational speech recognition," *arXiv preprint arXiv:1610.05256*, 2016.
- [26] T. Sercu and V. Goel, "Dense prediction on sequences with time-dilated convolutions for speech recognition," arXiv preprint arXiv:1611.09288, 2016.
- [27] W. Xiong, J. Droppo, X. Huang, F. Seide, M. Seltzer, A. Stolcke, D. Yu, and G. Zweig, "The Microsoft 2016 conversational speech recognition system," in *Proc. ICASSP*, 2017, pp. 5255– 5259.
- [28] H. Liu, Z. Zhu, X. Li, and S. Satheesh, "Gram-CTC: Automatic Unit Selection and Target Decomposition for Sequence Labelling," in *Proc. ICML*, 2017.
- [29] E. Battenberg, J. Chen, R. Child, A. Coates, Y. Gaur, Y. Li, H. Liu, S. Satheesh, D. Seetapun, A. Sriram, and Z. Zhu, "Exploring neural transducers for end-to-end speech recognition," *arXiv preprint arXiv:1707.07413*, 2017.
- [30] H. Sak, M. Shannon, K. Rao, and F. Beaufays, "Recurrent Neural Aligner: An Encoder-Decoder Neural Network Model for Sequence to Sequence Mapping," *Proc. Interspeech 2017*, pp. 1298–1302, 2017.
- [31] O. Siohan and M. Bacchiani, "Fast vocabulary-independent audio search using path-based graph indexing," in *Proc. European Conference on Speech Communication and Technology*, 2005.
- [32] B. Ramabhadran, A. Sethy, J. Mamou, B. Kingsbury, and U. Chaudhari, "Fast decoding for open vocabulary spoken term detection," in *Proc. NAACL-HLT*, 2009, pp. 277–280.
- [33] A. Rastrow, A. Sethy, B. Ramabhadran, and F. Jelinek, "Towards using hybrid word and fragment units for vocabulary independent LVCSR systems," in *Interspeech*, 2009, pp. 1931– 1934.