

Improved training for online end-to-end speech recognition systems

Suyoun Kim^{*}, Michael L. Seltzer², Jinyu Li³, Rui Zhao³

¹Carnegie Mellon University ²Facebook ³Microsoft AI & Research

suyoun@cmu.edu, mikeseltzer@fb.com, {jinyli, ruzhao}@microsoft.com

Abstract

Achieving high accuracy with end-to-end speech recognizers requires careful parameter initialization prior to training. Otherwise, the networks may fail to find a good local optimum. This is particularly true for online networks, such as unidirectional LSTMs. Currently, the best strategy to train such systems is to bootstrap the training from a tied-triphone system. However, this is time consuming, and more importantly, is impossible for languages without a high-quality pronunciation lexicon. In this work, we propose an initialization strategy that uses teacher-student learning to transfer knowledge from a large, well-trained, offline end-to-end speech recognition model to an online end-to-end model, eliminating the need for a lexicon or any other linguistic resources. We also explore curriculum learning and label smoothing and show how they can be combined with the proposed teacher-student learning for further improvements. We evaluate our methods on a Microsoft Cortana personal assistant task and show that the proposed method results in a 19% relative improvement in word error rate compared to a randomly-initialized baseline system.

1. Introduction

Recently, several so-called end-to-end speech recognition systems have been proposed in which a neural network is trained to predict character sequences which can be converted directly to words, or even word sequences directly. Approaches to endto-end systems typically utilize a Connectionist Temporal Classification (CTC) framework [1-7], an attention-based encoderdecoder framework [8-10], or both [11-13]. These types of models bypass much of the complexity associated with a traditional speech recognition system, in which an acoustic model predicts context-dependent phonemes, which are then combined with a pronunciation lexicon and a language model to generate hypothesized word sequences. As such, they have the potential to dramatically simplify the system building process and even outperform traditional systems as the unified modeling framework avoids the disjointed training procedure of conventional systems.

Many prior studies on end-to-end models [2, 4, 7–9, 11, 14] showed promising results with bidirectional recurrent neural networks (RNNs) that exploit access to the entire utterance. However, such models are not appropriate for online applications when the speech signal needs to be processed in a streaming manner. When there are such requirements for an end-to-end system, the most appropriate architecture is a causal model, such as a unidirectional RNN, trained with a CTC objective function. Unfortunately, when such models are trained from

scratch, i.e. from randomly initialized parameters, performance is severely degraded.

As a result, many methods for initializing the parameters of online end-to-end speech recognition systems have been proposed in the literature. Most of these methods first train a tied-triphone acoustic model and then one or more intermediate models before finally having a model initialized well enough that character-based end-to-end training can be reliably performed. For phoneme-based CTC, bootstrapping the CTC models with the model trained with cross-entropy (CE) on fixed alignments has been proposed to address this initialization issue [15]. Recently, a method for training tied-triphone CTC models from scratch was proposed but it still requires a pronunciation lexicon [16]. Alternatively, training the models jointly using a weighted combination of the CTC and CE losses has been suggested [17]. In case of the grapheme-level or word-level CTC models, it was shown that starting from phoneme-based model was essential in obtaining better performance [18-20].

While these approaches significantly improve the performance of the resulting model, they are highly unsatisfying. First, it makes system development complex and time consuming, as multiple complete systems need to be trained, discarding one of the most appealing aspects of end-to-end models. Second, for lower resource languages, a high-quality pronunciation lexicon may not be available, making such initialization strategies impossible to use.

In this work, we propose a method for initializing and training a high accuracy online CTC-based speech recognition system entirely within the end-to-end framework. Our method uses teacher-student learning, curriculum learning, and label smoothing, to achieve significantly improved performance for an online end-to-end speech recognition system. We avoid any initialization that relies on a pronunciation lexicon, a senone or tied-triphone acoustic model, or any other artifacts from a traditional speech recognizer. Nevertheless, we obtain performance that approaches that of the best performing CTC systems that have been derived from tied-triphone models. We experimentally validate our approach on the Microsoft Cortana personal assistant task. A training set of 3,400 hours is used, demonstrating that the proposed approach is effective even when abundant training data is available.

The remainder of this paper is organized as follows. In Section 2, we briefly review CTC-based end-to-end speech recognition. We then describe our proposed initialization and training strategies in Section 3 including the use of teacher-student learning, curriculum learning, and label smoothing. In Section 4, we evaluate our approach on the Microsoft Cortana task. Finally, we summarize our findings in Section 5.

The author performed the work during an internship at Microsoft Research.

2. End-to-end speech recognition with character-based CTC

In this work, we use Connectionist Temporal Classification (CTC) [1] with character outputs for our end-to-end model, as it is well-suited to online applications.

CTC allows the network to learn a variable-length sequence labeling problem where the input-output alignment is unknown. Given input features $\boldsymbol{x} = (x_1, \ldots, x_T)$, and the corresponding grapheme label sequence $\boldsymbol{y} = (y_1, \ldots, y_U)$, CTC trains the neural network according to a maximum-probability training criterion computed over all possible alignments. The probability of the possible label sequence is modeled as being conditionally independent by the product of each label probability. The loss function of the CTC model is computed as:

$$\mathcal{L}_{\text{CTC}} \triangleq -\ln P(\boldsymbol{y}|\boldsymbol{x}) \approx -\ln \sum_{\boldsymbol{\pi} \in \Phi} \prod_{t=1}^{T} P(k = \pi_t | \boldsymbol{x}) \quad (1)$$

where $\boldsymbol{\pi}$ is a label sequence in all possible expanded CTC path alignments Φ which have the length T, and $P(k = \pi_t | \boldsymbol{x})$ is a label distribution at time step t.

CTC uses a softmax output of the network to define a label distribution $P(k|\mathbf{x})$, where k represents each label among the K graphemes and a blank symbol ϵ representing no emission of output label. Deep recurrent neural networks (RNNs) are generally used to model the distribution over labels, $P(k|\mathbf{x})$. Each RNN hidden layer computes the sequence $\mathbf{h}^{l} = (h_{1}^{l}, \cdots, h_{T}^{l})$ and multiple RNN layers are stacked on top of each other. The hidden vector in the *l*-th layer, \mathbf{h}^{l} , is computed iteratively from l = 1 to L. The $P(k|\mathbf{x})$ is defined by the softmax of the final hidden layer h_{t}^{L} as follows:

$$P_t(k|\mathbf{x}) = \frac{exp(h_t^L(k))}{\sum_{i=1}^{K+1} exp(h_t^L(i))}$$
(2)

For offline applications, it is common to use bidirectional RNNs capable of modeling future context [21].

3. Improved character-based training for online CTC systems

3.1. Teacher-student learning from offline to online models

Teacher-student learning is an approach that was originally proposed to transfer knowledge from a large deep expert model (teacher) to a smaller shallower model (student) [22-25]. The student network is trained to minimize the difference between its own output distributions and those of the teacher network. This approach has been shown to be successful in several tasks including acoustic modeling [26], speech enhancement [27], and domain adaptation [28]. Inspired by these studies, we aim to improve the accuracy of a unidirectional online model by transferring knowledge from a bidirectional offline model that does not have any latency or complexity constraints. The recent work in [29] proposed a similar teacher-student method and they use the criteria as an additional objective. Unlike their frame-level teacher-student method degrades the performance, our method improves the performance by using the teacherstudent criteria in a separate training step as follows.

3.1.1. BLSTM-CTC

The first step is to build an offline end-to-end model as our *teacher* model. Since there is no latency restriction, we use a



Figure 1: Our proposed teacher-student learning for online CTC models.

deep bidirectional RNN with LSTM units (BLSTM) to predict the correct label sequence y given the entire utterance x. Each BLSTM layer computes h^l , the hidden representation in layer lby combining the two hidden sequences: the forward hidden sequence \overrightarrow{h}^l computed from processing the input in the forward direction, from $t = 1 \rightarrow T$, and the backward hidden sequence \overleftarrow{h}^l computed from processing the inputs in the backward direction from $t = T \rightarrow 1$ as follows,

$$h_t^l = W_{fw} \,\overline{h}_t^l + W_{bw} \,\overline{h}_t^l + b \tag{3}$$

$$\overrightarrow{h}_{t}^{l+1} = \text{LSTM}(h_{t}^{l}, \overrightarrow{h}_{t-1}^{l+1}) \tag{4}$$

$$\overleftarrow{h}_{t}^{l+1} = \text{LSTM}(h_{t}^{l}, \overleftarrow{h}_{t+1}^{l+1}) \tag{5}$$

where W_{fw}, W_{bw}, b are trainable parameters, and $h_t^0 = x_t$.

This model is trained using the CTC objective function described in Section 2 using character labels. Because every prediction the model makes is based on observing the entire utterance (in either the forward or backward direction), the BLSTM can be reliable trained from randomly-initialized model parameters.

3.1.2. LSTM-KL

Once the BLSTM-CTC model is trained, the next step is to transfer the predictive "knowledge" of this offline model to a model that can operate in an online manner, without access to the future input frames. To do so, we adopt a teacher-student approach in order to train the LSTM model to minimize the Kullback-Leibler (KL) divergence between the output distributions of the offline BLSTM-CTC model and the online LSTM-CTC model.

Let θ_{BLSTM} be the optimized parameters of BLSTM-CTC (*teacher*), and $P_t(k|\boldsymbol{x}; \theta_{\text{BLSTM}})$ be the output distribution at time step t generated from the BLSTM-CTC *teacher* model. Let $Q_t(k|\boldsymbol{x}; \theta_{\text{LSTM}})$ be the output distribution at time step t generated from LSTM-KL *student* model. The goal is to find the parameters θ_{LSTM} that minimizes the KL divergence D_{KL} between these distributions,

$$D_{KL}(P_t \mid\mid Q_t) = \sum_{k=1}^{K+1} P_t(k \mid \boldsymbol{x}; \theta_{\text{BLSTM}}) \ln \frac{P_t(k \mid \boldsymbol{x}; \theta_{\text{BLSTM}})}{Q_t(k \mid \boldsymbol{x}; \theta_{\text{LSTM}})}$$
(6)

$$=H(P_t, Q_t) - H(P_t) \tag{7}$$

where $k \in \{1, \ldots, K, \epsilon\}$ are the CTC labels including blank, $H(P_t, Q_t)$ is the cross entropy (CE) term and $H(P_t)$ is the entropy term. We can ignore the entropy term, $H(P_t)$, since its gradient is zero with respect to θ_{LSTM} . Thus, minimizing the KL divergence is equivalent to the CE between the two distributions P_t and Q_t :

$$\mathcal{L}_{\mathrm{KL}} \triangleq \sum_{t=1}^{T} H(P_t, Q_t) \tag{8}$$

$$= -\sum_{t=1}^{T} \sum_{i} P_t(k|\boldsymbol{x}; \theta_{\text{BLSTM}}) \ln Q_t(k|\boldsymbol{x}; \theta_{\text{LSTM}})$$
(9)

Figure 1 illustrates teacher-student learning between the BLSTM-CTC model and the LSTM-KL model.

3.1.3. LSTM-CTC

The last step in the training process is to optimize the online LSTM-KL model using the CTC objective function, as described in Section 2

$$\mathcal{L}_{\text{CTC}}(\theta_{\text{LSTM}}) \triangleq -\ln P(\boldsymbol{y}|\boldsymbol{x};\theta_{\text{LSTM}})$$
(10)

In the remainder of this section, we describe two enhancements that can further improve CTC training.

3.2. Curriculum learning

Curriculum learning has been proposed as a means of improving training stability, particularly in the early stages of learning. It was proposed to address the challenges of training deep neural networks under non-convex training criteria [30]. The main idea is to present the network a "curriculum" of tasks, starting with simple tasks and gradually increasing the task difficulty until the primary task is presented to the model. This approach is motivated by the observation that humans learn better with the examples organized in a meaningful order from simple to complex. Many prior studies [14,31,32] have shown improved performance as well as learning speed in various tasks, including language modeling, task memorization, and speech recognition. Specifically, in [14], they organized the training set in increasing order of the length and showed improved performance.

In this work, we explore two different curriculum strategies. First, we train the network on subset of the training set that consists of shorter length utterances, similar to prior study [14]. Because CTC training considers all possible sequences using the forward-backward algorithm, shorter utterances which have fewer possible paths are simpler to learn. After the network is able to learn the subset of short utterances, the full training set is introduced and training continues on the complete training set.

We also propose a new curriculum, where we reduce the number of categories, and simplify the classification to only four symbols: vowel, consonant, space, and blank. As the simpler classification task is learned, the full character-based label set is restored and training proceeds.

3.3. Label smoothing

Label smoothing is a general way to improve generalization by adding label noise, which has the effect of penalizing lowentropy output distributions, i.e. overly confident predictions. The authors in [33] showed the effectiveness of such a strategy on several well-known benchmark tasks, including image classification, language modeling, machine translation, and speech recognition. However, they have only explored the effectiveness of label smoothing in an encoder-decoder framework [8,9].

We observed that the label distribution from the CTC models is primarily dominated by the blank symbol. Motivated by this observation, we penalize low-entropy output distributions, similar to prior study [33]. To do so, we add a regularization term to the CTC objective function which consists of the KL divergence between the network's predicted distribution P and a uniform distribution U over labels.

$$\mathcal{L}(\theta_{\text{online}}) \triangleq (1 - \alpha)\mathcal{L}_{\text{CTC}} - \alpha \sum_{t=1}^{T} D_{KL}(P_t||U) \qquad (11)$$

where α is tunable parameter for balancing the weight regularization term and CTC loss.

4. Experiments

4.1. Datasets

We investigated the performance of our proposed training strategies on Microsoft's U.S. English Cortana personal assistant task. The training set has approximately 3,400 hours of utterances, the validation set has 10 hours of utterances, and the test set has 10 hours of utterances. As acoustic input features, we used 80-dimensional log mel filterbank coefficients extracted from 25 ms frames of audio every 10 ms. We employed a frame-skipping approach [34], where three consecutive frames are concatenated to obtain a 240-dimensional feature vector. None of our experiments used a pronunciation lexicon except when we compared to a system initialized with a tied-triphone model. Following [5], we used a label symbol inventory consisting of the individual characters and their doubleletter units. An initial capitalized letter rather than a space symbol was used to indicate word boundaries. This results in 81 distinct labels derived from 26 letters.

4.2. Training and Decoding

Our online LSTM-CTC model was a 5-layer LSTM [35, 36] with 1024 cells in each layer. Each LSTM layer has a linear projection to 512 dimensions. The offline BLSTM-CTC was a 5-layer network with 1024 cells in each direction, forward and backward. The hidden layer outputs in both directions were projected down to a 512-dimensional vector. When the models were randomly initialized, all the weights of our models were initialized with a uniform distribution in the range [-0.05, 0.05]. Parameters were trained using stochastic gradient descent with momentum. We use an initial learning rate of 0.0001 per sample. After each epoch, the criterion of the development set is evaluated, and if performance has degraded, the learning rate was decreased by a factor of 0.7.

For the decoding, the most likely sequence of characters was generated by the model in a greedy manner. The final output sequence was then obtained by removing any blank symbols or repetitions of characters from the output and replacing any capital letter with a space and its lowercase counterpart. Note that we did not use any lexicon or language models.

4.3. Results

We first compared the performance of an online LSTM model trained from random initialization to an LSTM trained by bootstrapping from a tied-triphone model. The tied-triphone initialization procedure was as follows. First, the LSTM was trained using cross-entropy computed on frame-level tied-triphone labels. Then, the model was retrained using CTC on tied-triphone outputs. Finally, the output layer was removed and replaced with character labels and then trained using a character-based CTC criterion.

As shown in Table 1, a randomly initialized system obtains a WER of 41.0%, while a tied-triphone-initialized system performs significantly better, with a WER of 30.8%, a 25% relative improvement. This difference in performance demonstrates that proper initialization is critical in training online end-to-end systems. Our goal is to close this gap in performance without relying on any of the linguistic resources required to build a tied-triphone acoustic model.

We next evaluated the three proposed training strategies individually, and all three resulted in substantial improvements in WER over the randomly-initialized baseline system. Curriculum learning, label smoothing, and teacher-student model initialization provided 7.8%, 13.7%, and 12.2% relative improvement in WER, respectively. The label smoothing technique showed largest improvement when a value of $\alpha = 0.05$ was used in Equation (11). We also found that curriculum that focused on a simplified label set also improved performance but did not work as effectively as one that focused on short utterances, so the results are not shown.

When we combined the two top performing strategies, teacher-student initialization and label smoothing, additional improvement is obtained, showing a 33.9 % WER. In this configuration, we first trained an LSTM-KL model, as described in Section 3.1.2, and then used it as an initial model for training with the CTC loss and the uniform distribution regularization. Lastly, we evaluated the combination of all three training strategies. This method performed best and obtained 33.2 % in WER, a 19% relative improvement over a randomly-initialized LSTM model. For reference, the table also shows the performance that can be obtained using an offline bidirectional model. Interestingly, even the performance of the BLSTM is improved considerably from label smoothing.

Another way to interpret these results is by considering how much the gap in performance between the randomly-initialized system and the tied-triphone-initialized system was narrowed by the proposed approach. There is a difference of 10.2% absolute between the randomly-initialized system and the tiedtriphone-initialized system and the proposed approach reduces this gap to 2.4%. Thus, the proposed approach has closed the gap by over 75% in a way that does not rely on a pronunciation lexicon.

Figure 2 shows the log-mel feature vectors for an utterance and the corresponding character probabilities using models trained with different training strategies. The black line shows the probability of the blank symbol and the other colors correspond to other letters. We manually chose an utterance that was recognized incorrectly by the randomly-initialized LSTM but correctly recognized by both the LSTM trained with the proposed approach and a BLSTM trained with label smoothing.

The figure shows that an LSTM initialized with teacherstudent training has consistently higher probability on the most likely output character compared to the randomly-initialized where several of the most-likely characters only have a posterior probability of 0.6 or less. It is also interesting to note the differences in spike locations between the online models and the offline BLSTM models. The BLSTM models can often have spikes that precede the acoustic observations of the corresponding letter as a result of the model processing information in the forward and backward directions simultaneously.

Table 1: WER on test set of models with various training strategies, curriculum learning (CL), label smoothing regularization (LS), and teacher-student approach (TS). None of our experiments used any language model or lexicon information.

Training Strategy	WER(%) w/o LM
LSTM + random initialization	41.0
LSTM + tied-triphone pre-training	30.8
LSTM + CL	37.8
LSTM + LS	35.4
LSTM + TS	36.0
LSTM + TS + LS	33.9
LSTM + TS + CL + LS	33.2
BLSTM	27.8
BLSTM + LS	25.5



Figure 2: Comparison of the CTC label probabilities over the input frames of the utterance "Is it supposed to snow tonight". for various models: a) an LSTM with random initialization, b) our proposed model with teacher-student learning, LSTM + TS + CL + LS, and c) an offline model BLSTM + LS.

5. CONCLUSION

We proposed a method to improve the accuracy for an online end-to-end speech recognizer via teacher-student learning. We transfer knowledge from a large, well-trained offline BLSTM model to an online LSTM model, by minimizing the KL divergence between the label distributions of the online model and the offline model. Because our method does not require bootstrapping from a tied-triphone system, it simplifies the training procedure and is beneficial for the languages where a pronunciation lexicon may be unavailable. We also found that our method can be easily combined with curriculum learning strategies and label smoothing. Our model was shown to outperform models that are trained from random initialization, and approach the performance of models bootstrapped from tiedtriphone systems.

6. References

- A. Graves, S. Fernández, F. Gomez, and J. 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.
- [2] A. Graves and N. Jaitly, "Towards end-to-end speech recognition with recurrent neural networks," in *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, 2014, pp. 1764–1772.
- [3] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates *et al.*, "Deep speech: Scaling up end-to-end speech recognition," *arXiv* preprint arXiv:1412.5567, 2014.
- [4] Y. Miao, M. Gowayyed, and F. Metze, "EESEN: End-to-end speech recognition using deep RNN models and WFST-based decoding," in 2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU). IEEE, 2015, pp. 167–174.
- [5] G. Zweig, C. Yu, J. Droppo, and A. Stolcke, "Advances in allneural speech recognition," in *Proc. ICASSP 2017.*
- [6] K. Audhkhasi, B. Kingsbury, B. Ramabhadran, G. Saon, and M. Picheny, "Building competitive direct acoustics-to-word models for english conversational speech recognition," in *submitted to ICASSP*, 2018.
- [7] J. Li, G. Ye, A. Das, R. Zhao, and Y. Gong, "Advancing Acousticto-Word CTC Model," in *Proc. ICASSP*, 2018.
- [8] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," in Advances in Neural Information Processing Systems, 2015, pp. 577– 585.
- [9] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, "Listen, attend and spell," arXiv preprint arXiv:1508.01211, 2015.
- [10] 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.*, "Stateof-the-art speech recognition with sequence-to-sequence models," in *Proc. ICASSP*, 2018.
- [11] S. Kim, T. Hori, and S. Watanabe, "Joint ctc-attention based endto-end speech recognition using multi-task learning," in Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on. IEEE, 2017, pp. 4835–4839.
- [12] T. Hori, S. Watanabe, and J. Hershey, "Joint CTC/attention decoding for end-to-end speech recognition," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics* (Volume 1: Long Papers), vol. 1, 2017, pp. 518–529.
- [13] A. Das, J. Li, R. Zhao, and Y. Gong, "Advancing connectionist temporal classification with attention modeling," in *Proc. ICASSP*, 2018.
- [14] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen *et al.*, "Deep speech 2: End-to-end speech recognition in english and mandarin," in *International Conference on Machine Learning*, 2016, pp. 173–182.
- [15] H. Sak, A. Senior, K. Rao, O. Irsoy, A. Graves, F. Beaufays, and J. Schalkwyk, "Learning acoustic frame labeling for speech recognition with recurrent neural networks," in *Acoustics, Speech* and Signal Processing (ICASSP), 2015 IEEE International Conference on. IEEE, 2015, pp. 4280–4284.
- [16] A. Senior, H. Sak, and K. Rao, "Flatstart-CTC: a new acoustic model training procedure for speech recognition," in *ICASSP* 2016, 2016.
- [17] E. Battenberg, R. Child, A. Coates, C. Fougner, Y. Gaur, J. Huang, H. Jun, A. Kannan, M. Kliegl, A. Kumar *et al.*, "Reducing bias in production speech models," *arXiv preprint arXiv*:1705.04400, 2017.
- [18] H. Soltau, H. Liao, and H. Sak, "Neural speech recognizer: Acoustic-to-word lstm model for large vocabulary speech recognition," arXiv preprint arXiv:1610.09975, 2016.

- [19] K. Rao and H. Sak, "Multi-accent speech recognition with hierarchical grapheme based models," in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on.* IEEE, 2017, pp. 4815–4819.
- [20] K. Audhkhasi, B. Ramabhadran, G. Saon, M. Picheny, and D. Nahamoo, "Direct acoustics-to-word models for english conversational speech recognition," arXiv preprint arXiv:1703.07754, 2017.
- [21] A. Graves, A.-r. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks."
- [22] J. Li, R. Zhao, J.-T. Huang, and Y. Gong, "Learning small-size dnn with output-distribution-based criteria," in *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [23] J. Ba and R. Caruana, "Do deep nets really need to be deep?" in Advances in neural information processing systems, 2014, pp. 2654–2662.
- [24] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [25] Y. Kim and A. M. Rush, "Sequence-level knowledge distillation," arXiv preprint arXiv:1606.07947, 2016.
- [26] W. Chan, N. R. Ke, and I. Lane, "Transferring knowledge from a rnn to a dnn," arXiv preprint arXiv:1504.01483, 2015.
- [27] S. Watanabe, T. Hori, J. Le Roux, and J. R. Hershey, "Studentteacher network learning with enhanced features," in Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on. IEEE, 2017, pp. 5275–5279.
- [28] J. Li, M. L. Seltzer, X. Wang, R. Zhao, and Y. Gong, "Largescale domain adaptation via teacher-student learning," *Proc. Interspeech*, 2017.
- [29] R. Takashima, S. Li, and H. Kawai, "An investigation of a knowledge distillation method for CTC acoustic models," in *Proc. ICASSP*, 2018.
- [30] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th annual international conference on machine learning*. ACM, 2009, pp. 41–48.
- [31] T. Mikolov, A. Deoras, D. Povey, L. Burget, and J. Černocký, "Strategies for training large scale neural network language models," in *Automatic Speech Recognition and Understanding* (ASRU), 2011 IEEE Workshop on. IEEE, 2011, pp. 196–201.
- [32] W. Zaremba and I. Sutskever, "Learning to execute," arXiv preprint arXiv:1410.4615, 2014.
- [33] G. Pereyra, G. Tucker, J. Chorowski, L. Kaiser, and G. Hinton, "Regularizing neural networks by penalizing confident output distributions," arXiv preprint arXiv:1701.06548, 2017.
- [34] H. Sak, A. Senior, K. Rao, and F. Beaufays, "Fast and accurate recurrent neural network acoustic models for speech recognition," *arXiv preprint arXiv:1507.06947*, 2015.
- [35] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [36] A. Graves, N. Jaitly, and A.-r. Mohamed, "Hybrid speech recognition with deep bidirectional lstm," in *Automatic Speech Recognition and Understanding (ASRU)*, 2013 IEEE Workshop on. IEEE, 2013, pp. 273–278.