

# Exploiting Large-scale Teacher-Student Training for On-device Acoustic Models

Jing Liu<sup>1</sup>, Rupak Vignesh Swaminathan<sup>1</sup>, Sree Hari Krishnan Parthasarathi<sup>1</sup>,  
Chunchuan Lyu<sup>2</sup>, Athanasios Mouchtaris<sup>1</sup>, and Siegfried Kunzmann<sup>1</sup>

<sup>1</sup> Alexa Machine Learning, Amazon, USA

<sup>2</sup> School of Informatics, University of Edinburgh, Scotland, UK

{jlmk, swarupak, sparta, mouchta, kunzman}@amazon.com

**Abstract.** We present results from Alexa speech teams on semi-supervised learning (SSL) of acoustic models (AM) with experiments spanning over 3000 hours of GPU time, making our study one of the largest of its kind. We discuss SSL for AMs in a small footprint setting, showing that a smaller capacity model trained with 1 million hours of unsupervised data can outperform a baseline supervised system by 14.3% word error rate reduction (WERR). When increasing the supervised data to seven-fold, our gains diminish to 7.1% WERR; to improve SSL efficiency at larger supervised data regimes, we employ a step-wise distillation into a smaller model, obtaining a WERR of 14.4%. We then switch to SSL using larger student models in low data regimes; while learning efficiency with unsupervised data is higher, student models may outperform teacher models in such a setting. We develop a theoretical sketch to explain this behavior.

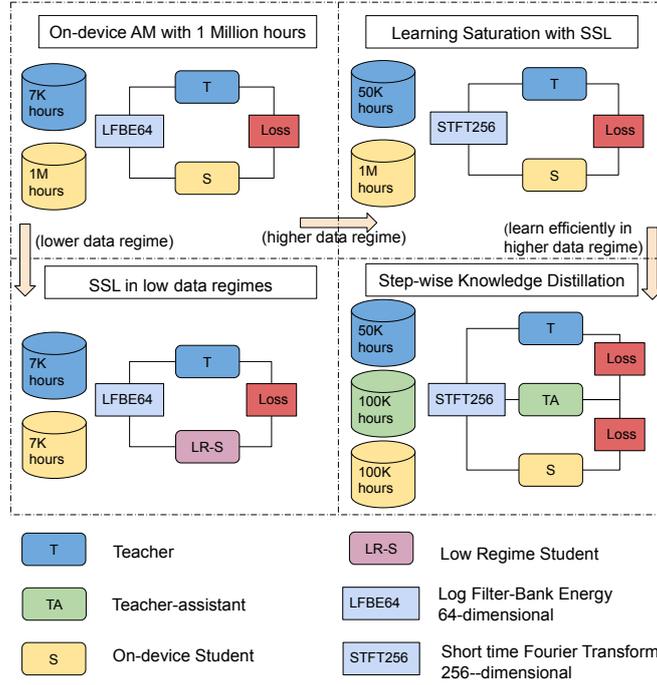
**Keywords:** speech recognition, acoustic models, edge computing, student-teacher learning, semi-supervised learning

## 1 Introduction

Semi-supervised learning (SSL) has a rich history in automatic speech recognition (ASR) [1, 11, 14, 17, 21]. Self-training is a commonly used technique employing confidence measures [32, 12]. Student-teacher distillation techniques [2, 9], foregoing full decoders and confidence models, have been shown to be effective for SSL [19]. SSL methods for end-to-end ASR have been studied in [16, 13, 36, 6, 25]. Furthermore, investigations with SSL in combination with data augmentation, pretraining and iterative self-training are done in [41] and [38].

Recently, student-teacher distillation techniques for hybrid HMM-LSTM models have been shown to scale to very large data sets (1 million hours) for models with high capacity [28, 27]. The efficacy of model compression using student-teacher distillation is well established [23, 35, 34]. In this context we study learning curves for AM for two tasks: (a) smaller footprint modeling, and (b) low training-data regimes.

Our motivation for low footprint AM comes from edge computing, where models are capacity restricted in terms of compute and memory [31, 40, 39]. We are interested in understanding if SSL, at very large data regimes for small models, can still yield gains in accuracy. In [24], mean squared errors between the teacher and the student



**Fig. 1.** Overview of our SSL approach progressing through 4 experiments. a) Top-left: Analysis of models using large unsupervised data. b) Top-right: Analysis of models in large supervised and unsupervised data regimes. c) Bottom-right: Improving learning efficiency in higher data regimes. d) Bottom-left: SSL studies in low data regimes.

hidden representations are explored as a regularization term in knowledge distillation. We demonstrate that a step-wise distillation approach, introduced in [23] can be effective, although this comes at the cost of more computation at training time. In low data regimes, SSL is an effective technique to reduce annotation costs [14, 17, 21, 5]. For our second task, using knowledge distillation for SSL, we find that to achieve a performance comparable to that of a fully supervised system, the proportion of required supervised data decreases as the amount of total data increases. However, we find that in low data regimes, students can be better than teachers.

Our contributions in this work are as follows: (1) we establish the robustness of small capacity semi-supervised models trained on 1 million hours of data; (2) we show an effective way to mitigate the learning saturation problem at higher data regimes for an on-device acoustic model; (3) we report results of a model distilled from a teacher trained on transcribed low-resource data, and present an empirical risk analysis. Figure 1 describes the approach we follow in this paper. We begin with a discussion on model configurations in Section 2. In Section 3, we present SSL for small footprint AMs. We show that such models can exploit a large amount of unsupervised data. How-

ever, when the amount of supervised data is increased from 7,000 hours to 52,000 hours the gains decrease. To mitigate this, we discuss a step-wise distillation into a smaller model. We then transition into experiments conducted in low resource settings in Section 4.

## 2 System Description

We now describe acoustic model configurations used in this work. It is a hybrid system, with an LSTM [10] estimating the senone posterior probabilities corresponding to clustered triphone HMM states. The HMMs are single state models using low-frame rate features [29], which are computed on speech signals every 10 ms, with a 25 ms analysis window. A running causal mean estimate is computed and subtracted from the features, and the resulting features are normalized by applying a global mean and variance normalization. The models are trained with the cross-entropy (CE) criterion, followed by sequence discriminative training using state-level minimum Bayes risk (sMBR) loss [15]. We follow an exponential learning rate decay for twelve epochs. More details on model configurations can be found in Table 1.

## 3 SSL for small footprint AM

In this section, we begin with an analysis of acoustic model complexity with supervised and unsupervised data for smaller and larger model footprints. We then investigate a method for efficient model distillation in larger supervised data regimes.

### 3.1 Learning Curves on Large Unsupervised Data

In Figure 2, on TST1 test data, we analyze the learning curves for on-device and cloud student models, taught by the same teacher network. Accuracy is reported as relative word error rate reduction (WERR) [26, 7]. Given model A’s WER ( $WER_A$ ) and a baseline B’s WER ( $WER_B$ ), the WERR of A over B is computed as  $WERR = (WER_B - WER_A) / WER_B$ .

The vertical axis is the relative WERR (%) against baseline LSTM AMs which are trained with CE criterion on the fully supervised 7,000 hour training data. The horizontal axis corresponds to the amount of unsupervised data. Each sub-epoch in the axis corresponds to about 60,000 hours of unsupervised data, totaling to 1 million hours. From Figure 2, we see that the relative WERR improves steadily for both cloud and on-device student AMs as we use increasing amount of unsupervised data. When using 1 million hours unsupervised data, the relative WERR for on-device and cloud student AMs are 14.3% and 14.6% respectively in CE comparing with its fully supervised models.

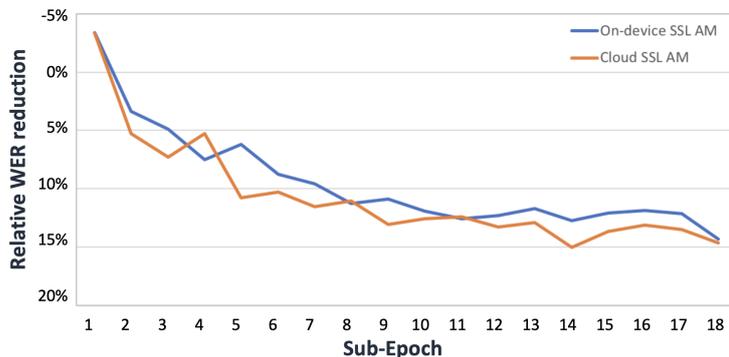
We perform sequence level discriminative training only on the 7,000 hour supervised dataset, demonstrating that the gains at the CE stage also carry over to the sMBR stage. Table 2 shows WERR for on-device and cloud student AMs using the same 4-gram LM. The WERRs for SSL CE + sMBR and supervised CE + sMBR are 25.7% and 18.2% respectively for on-device AMs.

**Table 1.** *Model Configurations*

<b>Common configurations for all experiments</b>
<ul style="list-style-type: none"> <li>• Features are stacked and subsampled to 33 Hz</li> <li>• 5 layers and 768 neurons/layer BLSTM teachers with 78 million parameters</li> <li>• Teacher model is cross-entropy and sequence trained</li> <li>• Teacher model trained with distributed trainer from [33]</li> <li>• Output posterior distribution over 3183 senones</li> <li>• Data selection for unsupervised data as in [28]</li> <li>• Training/test data are sampled from de-identified speech data</li> <li>• Student training strategy with compressed posteriors as in [28]</li> <li>• TST1: Test data consists of 100 hours of speech data</li> <li>• TST2: Test data consists of 30 hours of speech data</li> </ul>
<b>Section 3.1 Configurations for on-device AM</b>
<ul style="list-style-type: none"> <li>• 64-dimensional log filter-bank energies (LFBE)</li> <li>• 5 layers and 428 neurons/layer Uni-LSTM on-device student AMs with 8 million parameters</li> <li>• Cross entropy training of student models with distributed trainer from [4]</li> <li>• sMBR training of student model with distributed trainer from [33]</li> <li>• 7,000 hours of supervised training data with TST1 test data</li> <li>• 1 million hours of unsupervised speech data</li> </ul>
<b>Section 3.2 Configurations for studying learning saturation</b>
<ul style="list-style-type: none"> <li>• 256-dimensional STFT features [30, 18]</li> <li>• Teacher and student model architectures: same as above</li> <li>• Distributed training strategy: same as above</li> <li>• 52 K hours supervised training data with TST1 test data</li> <li>• 1 million hours of unsupervised speech data</li> </ul>
<b>Section 3.3 Configurations for step-wise distillation</b>
<ul style="list-style-type: none"> <li>• 256-dimensional STFT features</li> <li>• Teacher and student model architectures: same as above</li> <li>• Distributed training strategy: same as above</li> <li>• 50 K hours supervised training data with TST1 test data</li> <li>• 100 K hours unsupervised data</li> </ul>
<b>Section 4 Configurations for low resource experiments</b>
<ul style="list-style-type: none"> <li>• 64-dimensional LFBE features</li> <li>• 5 layers and 768 neurons/layer Uni-LSTM students with 24 million parameters</li> <li>• Teacher and student models trained with distributed trainer from [33]</li> <li>• Total training data is 7000 hours with TST2 test data</li> <li>• Supervised data ranges from 100 hours up to 7000 hours</li> <li>• The rest of the data is treated as unsupervised data</li> </ul>

### 3.2 Learning Saturation for On-device SSL System

In the following experiment, we increase the supervised data by about 7 times to 52,000 hours. We then gradually increase the amount of unsupervised data set to 1 million hours, computing WERR against a supervised model trained with 52,000 hours of supervised data. From Table 3, we get a relative WERR of 6.06% by adding up to 540 K hours of unsupervised data. If we further increase unsupervised data from 540 K hours to 1 million hours, the additional relative WERR is only 1%. We conclude that learning



**Fig. 2.** On TST1: relative WERR (%) per sub-epoch of the 1 million hour SSL models against baseline LSTM AMs that are trained with CE criterion on the fully supervised 7,000 hour training data.

**Table 2.** On TST1: relative WERR (%) for sequence training of SSL students. The baseline LSTM AMs are trained with CE criterion in a fully supervised setting with 7,000 hours of data

WERR (%)		
System	On-device	Cloud
Baseline supervised CE	0	0
Baseline supervised CE + sMBR	18.2	9.7
SSL CE	14.3	14.6
SSL CE + sMBR	25.7	24.0

starts to saturate at 540 K hours. Adding additional unsupervised data helps little with accuracy improvement.

### 3.3 Improving the Efficiency of Knowledge Distillation

In this section, we study a distillation method to improve the learning efficiency at higher data regimes. Specifically, we reduce the gap between the teacher and student models through an intermediate teacher assistant [23].

The supervised part of the training data consists of about 50 K hours of supervised US English speech, and the unsupervised part consists of about 100 K hours of unsupervised data. The unsupervised training data is used for training the teacher assistant. We perform step-wise teacher assisted knowledge distillation. The first-step knowledge distillation happens through a 78-million-parameter bidirectional LSTM teacher model and a 28-million-parameter teacher assistant model. The second-step knowledge distillation occurs between the 28-million-parameter teacher assistant model and an 8-million-parameter student network using only the soft-targets from the teacher assistant.

Table 4 shows relative WERR against a baseline cross-entropy student model (row 4) trained with 50,000 hours of supervised data. Rows 1, 2 and 3 show WERR of the

**Table 3.** On TST1: relative WERR (%) for SSL student AM with 8M parameters. Both the baseline and SSL AM are trained with 52k hours of fully supervised data.

WERR(%)	Unsupervised (hours)
6.05%	540 K
6.62%	720 K
6.86%	900 K
7.05%	1 M

**Table 4.** On TST1: relative WERR (%) for teacher assisted knowledge distillation method. The baseline model is the cross-entropy trained model on supervised data (row 4). WERR values are computed against baseline model’s WER.

Index	AM	Params (M)	Stage	WERR (%)
1	Teacher	78	sMBR	33.0
2	Cloud	28	sMBR	21.67
3	TA	28	sMBR	25.97
4	On-device	8	CE	0
			sMBR	17.7
5	KD	8	CE	9.2
			sMBR	15.7
6	TAKD	8	CE	14.4
			sMBR	19.0

teacher, supervised cloud model, and SSL-trained teacher assistant models evaluated against baseline. Sequence training with sMBR (in row 4) shows 17.7% relative WERR in a fully supervised system. The distillation results are presented in rows 5 and 6. The improvement over direct Knowledge Distillation (KD) is 9.2% relative. For Teacher Assisted Knowledge Distillation (TAKD), there is a substantial improvement of 14.4% compared to baseline supervised system, thus showing that step-wise distillation with a teacher assistant indeed helps with the efficiency for smaller models in large data regimes.

## 4 SSL in Low Resource Settings

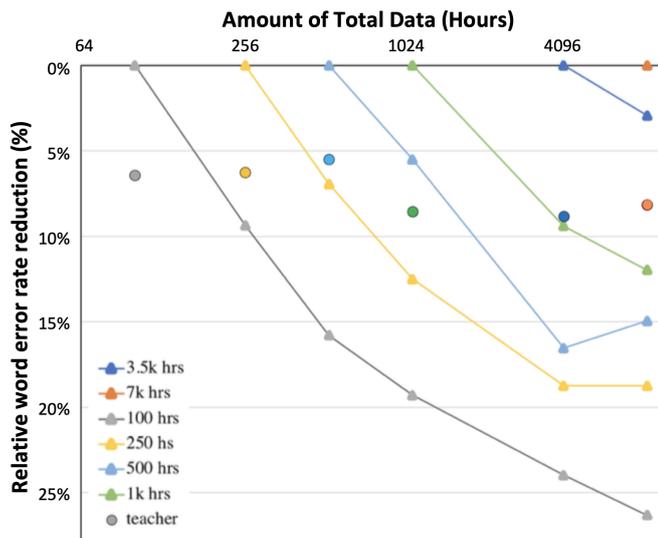
In this section, we present our results on low resource settings with SSL; specifically, now the student model is larger, but the overall amount of data is restricted to 7,000 hours.

### 4.1 Accuracy Gains With Unsupervised Data

Figure 3 shows learning curves for student models at fixed amount of supervised data. The actual amounts of supervised data are: 100, 250, 500, 1000, 3500 and 7000 hours. Triangular markers on each curve correspond to student models trained with increasing amounts of unsupervised data by fixing the amount of supervised data.

Relative WERR is computed against a baseline model trained only on the supervised data corresponding to that curve of the same color. The circular solid dots in this plot show the relative WERR for the teacher models. These dots line up vertically against the triangles corresponding to the baselines models.

For each curve, the slope of the curve becomes less steep as the amount of unsupervised data increases, meaning diminishing returns for additional data. Note that the student models can outperform the corresponding teacher model, once the student model observes unsupervised data.



**Fig. 3.** On TST2: relative WERR (%) for different amounts of unsupervised data: each curve corresponds to a fixed amount of supervised data. Markers on each curve correspond to student models trained with increasing amounts of unsupervised data. Relative WERR is computed against the model trained only on the supervised data corresponding to that curve. The circular dots correspond to the BLSTM teacher models for the curves of the same color.

## 4.2 An Analysis of Empirical Risk with Student-Teacher Learning

To simplify analysis, we restrict our setting to binary classification, and adopt notations from [3]. Let  $Z = \{(x_i, y_i)\}_{i=1, \dots, N}$  be a supervised set drawn from a distribution  $D$  on  $\mathcal{X} \times \{-1, 1\}$ , where  $\mathcal{X} = \{x_i\}_{i=1..N}$ . Let  $\mathbb{1}[A] : A \subset \mathcal{X} \rightarrow \{0, 1\}$  be an indicator function defined as  $\mathbb{1}[A](x) = \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{if } x \notin A \end{cases}$ . Teacher and student network outputs are denoted  $h^t(x)$  and  $h^s(x)$  respectively, and are drawn from  $\{-1, 1\}$ . The empirical

teacher-student training risk is

$$R_Z^{h^t}(h^s) := \frac{1}{N} \sum_{i=1, \dots, N} \mathbb{1}[h^t(x_i) \neq h^s(x_i)],$$

and the actual empirical risk of the student is

$$R_Z(h^s) := R_Z^y(h^s) = \frac{1}{N} \sum_{i=1, \dots, N} \mathbb{1}[y_i \neq h^s(x_i)].$$

**Lemma 1 (Decomposition).** *Given  $h^t, h^s, y \in \{-1, 1\}$ , we have*

$$R_Z(h^s) = R_Z(h^t)(1 - R_{acc(h^t)}^{h^t}(h^s) - R_{err(h^t)}^{h^t}(h^s)) + R_{acc(h^t)}^{h^t}(h^s)$$

where  $acc(h^t) := \{x, y \in Z | h^t(x) = y\}$  and  $err(h^t) := \{x, y \in Z | h^t(x) \neq y\}$ .

*Proof.* Decomposing the student risk by partitioning  $Z$  on whether the teacher makes a mistake

$$\begin{aligned} R_Z(h^s) &= R_Z(h^t)(1 - R_{err(h^t)}^{h^t}(h^s)) + (1 - R_Z(h^t))R_{acc(h^t)}^{h^t}(h^s) \\ &= R_Z(h^t)(1 - R_{acc(h^t)}^{h^t}(h^s) - R_{err(h^t)}^{h^t}(h^s)) + R_{acc(h^t)}^{h^t}(h^s) \end{aligned}$$

**Theorem 1 (Truth Over Teacher).**

$$R_Z(h^s) \leq R_Z(h^t) \iff \frac{R_{err(h^t)}^{h^t}(h^s)}{R_{acc(h^t)}^{h^t}(h^s)} \geq \frac{1}{R_Z(h^t)} - 1$$

*Proof.*

$$\begin{aligned} R_Z(h^s) &\leq R_Z(h^t) \\ &\iff \\ R_Z(h^t)(1 - R_{acc(h^t)}^{h^t}(h^s) - R_{err(h^t)}^{h^t}(h^s)) + R_{acc(h^t)}^{h^t}(h^s) &\leq R_Z(h^t) \\ &\iff \\ R_{acc(h^t)}^{h^t}(h^s) &\leq R_Z(h^t)(R_{acc(h^t)}^{h^t}(h^s) + R_{err(h^t)}^{h^t}(h^s)) \\ &\iff \\ \frac{1}{R_Z(h^t)} &\leq 1 + \frac{R_{err(h^t)}^{h^t}(h^s)}{R_{acc(h^t)}^{h^t}(h^s)} \\ &\iff \\ \frac{R_{err(h^t)}^{h^t}(h^s)}{R_{acc(h^t)}^{h^t}(h^s)} &\geq \frac{1}{R_Z(h^t)} - 1 \end{aligned}$$

The better the teacher is, the easier a student model fits the true label compared to a false label. In particular, if the prediction error the teacher makes  $R_Z(h^t)$  is better

than random guess, namely  $R_Z(h^t) < 0.5$ , then it is necessary that  $\frac{R_{err(h^t)}(h^s)}{R_{acc(h^t)}(h^s)} > 1$  so that  $R_Z(h^s) \leq R_Z(h^t)$ . From Theorem 1, for the student to be better than its teacher, the student’s risk evaluated on the error set of the teacher (the erroneous teacher labels) has to be greater than the risk evaluated on the accuracy set of the teacher (the correct teacher labels) by a factor of  $\frac{1}{R_Z(h^t)} - 1$ .

## 5 Discussion

In low data regimes, our experiments surprisingly showed that a student model’s performance is not upper bounded by the teacher model used in learning. Our theoretical analysis indicates that this can happen, for example, when the student makes fewer errors on data with high teacher errors. We speculate this is because a capacity-restricted student model is able to generalize better than the teacher using the unsupervised data.

In high data regimes, we observed that the TAKD method facilitated efficient learning. We speculate that this could be due to step-wise distillation providing better calibrated posteriors for the eventual student to learn from. Indeed, it has been observed that in modern over-parameterized neural networks, posterior probabilities can become less calibrated [8]; using large amounts of supervised data could be making teacher models’ estimates of posteriors less calibrated.

## 6 Conclusions And Future Work

Using BLSTM and LSTM as teacher and student models respectively, we studied SSL for AMs for two tasks: (a) for small footprint on-device models; and (b) a larger footprint, but lower training data regime. Despite smaller model capacity, on-device models were able to exploit 1 million hours of unsupervised data and achieve a 14.3% relative WER improvement after cross-entropy training and the gains could carry over to the sMBR stage. However, when we increased the supervised data from 7,000 hours to 52,000 hours, the learning efficiency decreased yielding only 7.1% relative improvements in WER. We utilized a step-wise distillation and recovered the WER gains. For SSL in low data regimes, using knowledge distillation, we found that to achieve a performance comparable to that of a fully supervised system, the proportion of required supervised data decreased as the amount of total data increased. However, we found that in low data regimes, students can be better than teachers. In future work, we would like to extend this large scale study to distill using sequence-based discriminative criterion [37, 22, 20] instead of the frame-level cross-entropy criterion. Finally, it would be informative to understand if the analyses carried out in this paper using hybrid models would carry over to end-to-end ASR systems.

## 7 Acknowledgements

We would like to thank Minhua Wu, Jangwon Kim, Srinivas Parthasarathy, Kishore Nandury and Brian King for their helpful discussions.

## References

1. Amodei, D., Ananthanarayanan, S., et al.: Deep speech 2: End-to-end speech recognition in English and Mandarin. In: Proc. of ICML (2016)
2. Ba, J., Caruana, R.: Do deep nets really need to be deep? In: Advances in neural information processing systems. pp. 2654–2662 (2014)
3. Bégin, L., Germain, P., Laviolette, F., Roy, J.F.: Pac-bayesian theory for transductive learning. In: Proc. of AISTATS (2014)
4. Chen, K., Huo, Q.: Scalable training of deep learning machines by incremental block training with intra-block parallel optimization and blockwise model-update filtering. In: Proc. of ICASSP (2016)
5. Chen, L., Leutnant, V.: Acoustic Model Bootstrapping Using Semi-Supervised Learning. In: Proc. Interspeech 2019. pp. 3198–3202 (2019). <https://doi.org/10.21437/Interspeech.2019-2818>, <http://dx.doi.org/10.21437/Interspeech.2019-2818>
6. Chen, Y., Wang, W., Wang, C.: Semi-supervised asr by end-to-end self-training (2020)
7. Garimella, S., Mandal, A., Strom, N., Hoffmeister, B., Matsoukas, S., Parthasarathi, S.H.K.: Robust i-vector based adaptation of DNN acoustic model for speech recognition. In: Proc. of Interspeech (2015)
8. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. arXiv preprint arXiv:1706.04599 (2017)
9. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
10. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Computation* **9**(8), 1735–1780 (1997). <https://doi.org/10.1162/neco.1997.9.8.1735>, <https://doi.org/10.1162/neco.1997.9.8.1735>
11. Huang, Y., Wang, Y., Gong, Y.: Semi-supervised training in deep learning acoustic model. In: Proc. of Interspeech (2016)
12. Huang, Y., Yu, D., Gong, Y., Liu, C.: Semi-supervised GMM and DNN Acoustic Model Training with Multi-system Combination and Confidence Re-calibration. In: Proc. of Interspeech (2013)
13. Kahn, J., Lee, A., Hannun, A.: Self-training for end-to-end speech recognition. ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (May 2020). <https://doi.org/10.1109/icassp40776.2020.9054295>, <http://dx.doi.org/10.1109/ICASSP40776.2020.9054295>
14. Kemp, T., Waibel, A.: Unsupervised training of a speech recognizer: recent experiments. In: Proc. of Eurospeech (1999)
15. Kingsbury, B.: Lattice-based optimization of sequence classification criteria for neural-network acoustic modeling. In: 2009 IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 3761–3764 (2009). <https://doi.org/10.1109/ICASSP.2009.4960445>
16. Kurata, G., Audhkhasi, K.: Guiding CTC posterior spike timings for improved posterior fusion and knowledge distillation. *CoRR* **abs/1904.08311** (2019), <http://arxiv.org/abs/1904.08311>
17. Lamel, L., Gauvain, J.L., Adda, G.: Lightly supervised and unsupervised acoustic model training. *Computer Speech & Language* (2002)
18. Li, J., Mohamed, A., Zweig, G., Gong, Y.: LSTM time and frequency recurrence for automatic speech recognition. In: 2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU). pp. 187–191 (2015). <https://doi.org/10.1109/ASRU.2015.7404793>
19. Li, J., Zhao, R., Huang, J.T., Gong, Y.: Learning small-size DNN with output-distribution-based criteria. In: Proc. of Interspeech (2014)

20. Li, J., Zhao, R., Sun, E., Wong, J.H.M., Das, A., Meng, Z., Gong, Y.: High-accuracy and low-latency speech recognition with two-head contextual layer trajectory lstm model. In: ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 7699–7703 (2020). <https://doi.org/10.1109/ICASSP40776.2020.9054387>
21. Ma, J., Matsoukas, S., Kimball, O., Schwartz, R.: Unsupervised training on large amounts of broadcast news data. In: Proc. of ICASSP (2006)
22. Manohar, V., Ghahremani, P., Povey, D., Khudanpur, S.: A teacher-student learning approach for unsupervised domain adaptation of sequence-trained asr models. In: 2018 IEEE Spoken Language Technology Workshop (SLT). pp. 250–257 (2018). <https://doi.org/10.1109/SLT.2018.8639635>
23. Mirzadeh, S.I., Farajtabar, M., Li, A., Ghasemzadeh, H.: Improved knowledge distillation via teacher assistant: Bridging the gap between student and teacher. arXiv preprint arXiv:1902.03393 (2019)
24. Moriya, T., Kanagawa, H., Matsui, K., Fukutomi, T., Shinohara, Y., Yamaguchi, Y., Okamoto, M., Aono, Y.: Efficient building strategy with knowledge distillation for small-footprint acoustic models. In: 2018 IEEE Spoken Language Technology Workshop (SLT). pp. 21–28 (2018). <https://doi.org/10.1109/SLT.2018.8639545>
25. Munim, R.M., Inoue, N., Shinoda, K.: Sequence-level knowledge distillation for model compression of attention-based sequence-to-sequence speech recognition. In: ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 6151–6155. IEEE (2019)
26. Parthasarathi, S.H.K., Hoffmeister, B., Matsoukas, S., Mandal, A., Strom, N., Garimella, S.: fmlr based feature-space speaker adaptation of DNN acoustic models. In: Proc. of Interspeech (2015)
27. Parthasarathi, S.H.K., Sivakrishnan, N., Ladkat, P., Strom, N.: Realizing petabyte scale acoustic modeling. IEEE Journal on Emerging and Selected Topics in Circuits and Systems (2019)
28. Parthasarathi, S.H.K., Strom, N.: Lessons from building acoustic models from a million hours of speech. In: Proc. of ICASSP (2019)
29. Pundak, G., Sainath, T.: Lower frame rate neural network acoustic models. In: Proc. of Interspeech (2016)
30. Sak, H., Senior, A., Rao, K., Beaufays, F.: Fast and accurate recurrent neural network acoustic models for speech recognition. In: INTERSPEECH (2015)
31. Shi, W., Cao, J., Zhang, Q., Li, Y., Xu, L.: Edge computing: Vision and challenges. IEEE internet of things journal (2016)
32. Siu, M.H., Gish, H., Richardson, F.: Improved estimation, evaluation and applications of confidence measures for speech recognition. In: Proc. of European Conference on Speech Communication and Technology (1997)
33. Strom, N.: Scalable distributed DNN training using commodity GPU cloud computing. In: Proc. of Interspeech (2015)
34. Watanabe, S., Hori, T., Le Roux, J., Hershey, J.R.: Student-teacher network learning with enhanced features. In: 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 5275–5279 (2017). <https://doi.org/10.1109/ICASSP.2017.7953163>
35. Waters, A., Chebotar, Y.: Distilling knowledge from ensembles of neural networks for speech recognition. In: Interspeech (2016)
36. Weninger, F., Mana, F., Gemello, R., Andres-Ferrer, J., Zhan, P.: Semi-supervised learning with data augmentation for end-to-end asr (2020)
37. Wong, J.H.M., Gales, M.: Sequence student-teacher training of deep neural networks. In: INTERSPEECH (2016)

38. Xie, Q., Dai, Z., Hovy, E., Luong, M.T., Le, Q.V.: Unsupervised data augmentation for consistency training (2020)
39. Yanzhang, H., Sainath, T., Prabhavalkar, R., McGraw, I., Alvarez, R., Zhao, D., Rybach, D., Kannan, A., Wu, Y., Pang, R., Liang, Q., Bhatia, D., Shangguan, Y., Li, B., Pundak, G., Sim, K.C., Bagby, T., Chang, S., Rao, K., Gruenstein, A.: Streaming end-to-end speech recognition for mobile devices. arXiv preprint arXiv:1811.06621 (2018)
40. Yu, W., Liang, F., He, X., Hatcher, W.G., Lu, C., Lin, J., Yang, X.: A survey on the edge computing for the internet of things. IEEE access (2017)
41. Zhang, Y., Qin, J., Park, D.S., Han, W., Chiu, C.C., Pang, R., Le, Q.V., Wu, Y.: Pushing the limits of semi-supervised learning for automatic speech recognition (2020)