



# Multiple-hypothesis RNN-T Loss for Unsupervised Fine-tuning and Self-training of Neural Transducer

Cong-Thanh Do, Mohan Li, and Rama Doddipatla

Cambridge Research Laboratory, Toshiba Europe Limited, Cambridge, U.K.

{cong-thanh.do, mohan.li, rama.doddipatla}@crl.toshiba.co.uk

## Abstract

This paper proposes a new approach to perform unsupervised fine-tuning and self-training using unlabeled speech data for recurrent neural network (RNN)-Transducer (RNN-T) end-to-end (E2E) automatic speech recognition (ASR) systems. Conventional systems perform fine-tuning/self-training using ASR hypothesis as the targets when using unlabeled audio data and are susceptible to the ASR performance of the base model. Here in order to alleviate the influence of ASR errors while using unlabeled data, we propose a multiple-hypothesis RNN-T loss that incorporates multiple ASR 1-best hypotheses into the loss function. For the fine-tuning task, ASR experiments on Librispeech show that the multiple-hypothesis approach achieves a relative reduction of 14.2% word error rate (WER) when compared to the single-hypothesis approach, on the `test_other` set. For the self-training task, ASR models are trained using supervised data from Wall Street Journal (WSJ), Aurora-4 along with CHiME-4 real noisy data as unlabeled data. The multiple-hypothesis approach yields a relative reduction of 3.3% WER on the CHiME-4's single-channel real noisy evaluation set when compared with the single-hypothesis approach.

**Index Terms:** RNN-T, end-to-end ASR, multiple-hypothesis, unsupervised fine-tuning, self-training

## 1. Introduction

RNN-T [1], or neural transducer [2], is a popular E2E ASR model which was found successful in both research and industry environments [3, 4]. This model provides state-of-the-art performance in a wide variety of streaming applications [4]. RNN-T extends connectionist temporal classification (CTC) to define a distribution over output sequences of all lengths, and jointly models both input-output and output-output dependencies [1]. The modeling in RNN-T overcomes the issue regarding the conditional independence between output predictions at different time steps given aligned inputs, introduced in CTC [5].

While RNN-T model is widely applied, it would be useful to be able to efficiently customize this model to new application domains [6]. However, speech data with manual transcriptions are not always available because the cost to obtain the labels is substantially high. How to leverage large-scale unlabeled speech data to improve ASR performance has been a longstanding research problem. Pseudo-labeling [7] is one of the approaches that can be used to customize RNN-T models to new domains using unlabeled data. The pseudo-labeling process starts with training a base model using available labeled data. The base model is then used to label the unlabeled data. The pseudo-labeled data are subsequently used to fine-tune or re-train the base model. The pseudo-labeling process can be repeated to improve the quality of the customized model. Pseudo-labeling has been a practically useful and extensively studied technique in ASR [8, 9].

In this paper, we present a new approach for unsupervised fine-tuning and self-training of RNN-T model using unlabeled

speech data and multiple ASR 1-best hypotheses. Multiple ASR 1-best hypotheses are used instead of a single ASR 1-best hypothesis which is conventionally used in pseudo-labeling. The ASR 1-best hypotheses are incorporated in the RNN-T loss function in order to alleviate the influence of errors in the ASR hypotheses to the computation of the RNN-T loss. This proposal is consistent with our previous work where multiple ASR 1-best hypotheses were incorporated into the CTC loss function to improve the adaptation of CTC-based E2E ASR [10]. In this respect, a multiple-hypothesis RNN-T loss is proposed to incorporate multiple ASR 1-best hypotheses during unsupervised fine-tuning and self-training. Experiments are performed using data from Aurora-4 [11], CHiME-4 [12], Wall Street Journal (WSJ) [13], and Librispeech [14] corpora. In both unsupervised fine-tuning and self-training scenarios, the approach using multiple ASR 1-best hypotheses is effective in reducing the WER and it outperforms the conventional approach using single ASR 1-best hypothesis.

The paper is organized as follows. Section 2 presents related works. The multiple-hypothesis RNN-T loss used for fine-tuning and self-training is introduced in section 3. Details of the unsupervised fine-tuning and self-training experiments together with the experimental results are presented in sections 4 and 5, respectively. Finally, section 6 concludes the paper.

## 2. Related works

Research on neural transducer and the customization of this model is an active research topic in ASR [6, 15]. In [16], the authors investigated which components of the RNN-T model is the most effective one to be adapted when the amount of adaptation data is limited, in supervised and unsupervised speaker adaptation. In [17], i-vector speaker embedding was applied for speaker adaptation in RNN-T E2E ASR. In [18], a small amount of textual data was used to adapt the RNN-T's prediction network to obtain fast adaptation of RNN-T model.

Multiple hypotheses were previously used in cross-system acoustic model adaptation where the transcriptions for adaptation were generated by several systems, which were built with various phoneme sets or acoustic front-ends [19, 20, 21]. Recently, hypotheses from the  $N$ -best list are increasingly used in ASR applications. For instance, in [22], the hypotheses in  $N$ -best list were used during the minimum WER training of RNN-T E2E ASR. In [23], an  $N$ -best list of pseudo-labels were used to generate lattice supervision in self-training using graph-based temporal classification objective function.

In the present paper, we propose to perform unsupervised fine-tuning and self-training of RNN-T E2E model using unlabeled data and multiple ASR 1-best hypotheses from different ASR systems. In fact, it was proved that better results are obtained when the adaptation process exploits a hypothesis generated by a system different than the one under adaptation [19], similar to in systems combination [24]. In our approach, the ASR 1-best hypotheses are incorporated into the RNN-T loss functions during fine-tuning and self-training.

### 3. Multiple-hypothesis RNN-T loss

Given a  $T$ -length acoustic feature vector sequence  $X = \{\mathbf{x}_t \in \mathbb{R}^d | t = 1, \dots, T\}$ , where  $\mathbf{x}_t$  is a  $d$ -dimensional feature vector at frame  $t$ , and a transcription  $C = \{c_l \in \mathcal{U} | l = 1, \dots, L\}$  which consists of  $L$  characters, where  $\mathcal{U}$  is a set of distinct characters, during the training of the RNN-T E2E ASR, the RNN-T loss function  $L_{\text{RNN-T}}$  is minimized.  $L_{\text{RNN-T}}$  is defined as follows:

$$L_{\text{RNN-T}} = -\log P_{\theta}(C|X), \quad (1)$$

where  $\theta$  are the network parameters and  $P_{\theta}(C|X) = \sum_{\hat{C}} P(\hat{C}|X)$ ,  $\hat{C} \in \mathcal{B}^{-1}(C)$ . In equation (1),  $C$  is the transcription of  $X$ , which can either be a manual transcription or an ASR hypothesis, and  $\mathcal{B}^{-1}(C)$  is the set of all possible alignments, including the blank symbol, between the transcription  $C$  and the input feature vector sequence  $X$ . In RNN-T, the posterior probability  $P(\hat{C}|X)$  is computed as:

$$P(\hat{C}|X) = \prod_{\hat{c}_i \in \hat{C}} P(\hat{c}_i | \mathbf{x}_1 \dots \mathbf{x}_{t_i}, \hat{c}_1 \dots \hat{c}_{i-1}). \quad (2)$$

From equations (1) and (2), we have:

$$L_{\text{RNN-T}} = -\log \left( \sum_{\hat{C} \in \mathcal{B}^{-1}(C)} \prod_{\hat{c}_i \in \hat{C}} P(\hat{c}_i | \mathbf{x}_1 \dots \mathbf{x}_{t_i}, \hat{c}_1 \dots \hat{c}_{i-1}) \right). \quad (3)$$

It can be observed in equations (2) and (3) that both input-output and output-output dependencies are modeled in the RNN-T loss, as mentioned in section 1. In this paper, we propose multiple-hypothesis RNN-T loss  $L_{\text{RNN-T}}^*$  to be used during fine-tuning and self-training as follows:

$$L_{\text{RNN-T}}^* = - \left( \sum_{i=1}^M \log P_{\theta}(\bar{C}_i | X) \right), \quad (4)$$

where  $\bar{C}_i, i = 1, \dots, M$  are the ASR 1-best hypotheses obtained by decoding the unlabeled speech data, which is available for fine-tuning or self-training, using  $M$  different base models. By combining multiple 1-best hypotheses in the computation of the RNN-T loss, the influence of ASR errors in the 1-best hypotheses to the computation of the RNN-T loss function could be alleviated [10].

In this work, the unlabeled speech data available for unsupervised fine-tuning and self-training are assumed to have no manual transcriptions. The proposed multiple-hypothesis RNN-T loss is used in both unsupervised fine-tuning and self-training to incorporate multiple ASR 1-best hypotheses.

## 4. Unsupervised fine-tuning

### 4.1. Method

In unsupervised fine-tuning experiments, we assume that labeled speech data is available for training a RNN-T base model. The base model is normally evaluated on test data which includes audio only. During unsupervised fine-tuning, the audio of the test data is used to fine-tune the base model. Since the labels of the test audio are not available, the base model is used to decode the test audio. The 1-best hypotheses obtained from the decoding are used as the labels for the test audio data during the fine-tuning of the base model.

During the unsupervised fine-tuning, back-propagation algorithm is used to fine-tune all parameters of the base model with respect to the loss function, using the available data. Stochastic gradient descent (SGD) algorithm [25] is used during

the minimization of the loss function. In conventional unsupervised fine-tuning, the standard RNN-T loss  $L_{\text{RNN-T}}$  (see equation (1)) and one single ASR 1-best hypothesis are used while in the multiple-hypothesis unsupervised fine-tuning, the proposed multiple-hypothesis RNN-T loss  $L_{\text{RNN-T}}^*$  (see equation (4)) and multiple ASR 1-best hypotheses are used. During the training of the base model, manual transcriptions of audio data are used in the standard RNN-T loss  $L_{\text{RNN-T}}$ .

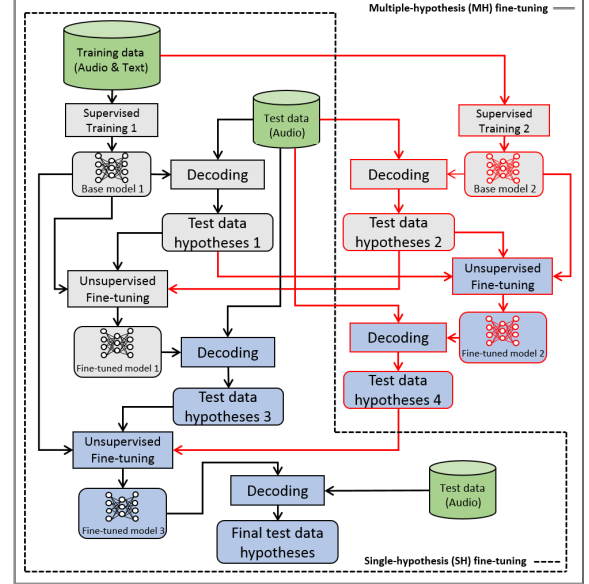


Figure 1: Unsupervised fine-tuning of RNN-T model using unlabeled audio data and multiple ASR 1-best hypotheses. The blocks surrounded by the dashed lines are used in conventional single-hypothesis fine-tuning. The red outline blocks and arrows are used only for multiple-hypothesis fine-tuning. The blocks filled with gray color are used in the first iteration of fine-tuning while those filled with pastel blue color are used in the second fine-tuning iteration. Arrows connecting the test audio data and unsupervised fine-tuning blocks are not shown to reduce the density of connections and improve clarity.

The unsupervised fine-tuning experiments are depicted in Fig. 1. The processing used by the conventional single-hypothesis unsupervised fine-tuning are surrounded by the dashed lines. Additional processing are required by the unsupervised fine-tuning using multiple ASR hypotheses and these processing are represented in red color. In this work, from the labeled training data we train different RNN-T base models, in addition to the one that needs to be fine-tuned, to decode the test audio data to obtain multiple ASR 1-best hypotheses.

In Fig. 1, the “Unsupervised Fine-tuning” blocks can receive either single or multiple ASR hypotheses. When single hypothesis is received, the standard RNN-T loss is used. When multiple hypotheses are received, the multiple-hypothesis RNN-T loss is used to incorporate these hypotheses. In the fine-tuning using single hypothesis, the additional hypotheses, represented by red arrows in Fig. 1, are not used. Both single-hypothesis and multiple-hypothesis fine-tuning can do more than one iteration to improve further the performance. Indeed, the fine-tuned model can be used to decode the test audio data again to obtain better hypotheses. This process can be repeated until the performance improvement is saturated. In this work, depending on the experiment, we perform either one or two fine-tuning iterations, as shown in Fig. 1.

## 4.2. Data and model architecture

The unsupervised fine-tuning experiments are carried out with Librispeech corpus [14]. The Librispeech dataset contains approximately 1000 hours of speech data that is derived from read English audiobooks. The training data is split into 3 subsets, including two “clean” subsets with the duration of 100 and 360 hours, as well as an “other” subset with the duration of 500 hours. While the dev and test sets are also split into the “clean” and “other” categories with the duration of 5 hours and 20 minutes each, facilitating evaluation under different speech conditions. The total number of speakers in the training and dev/test sets are 2338 and 77, respectively. In our experiment, we used the full training data (960 hours) to train the ASR models. The dev set used for the training is a combination of the “clean” and “other” dev sets.

A word-level language model (LM) is trained to rescore the hypotheses produced by the RNN-T ASR system. The architecture of the LM is a 4-layer long short-term memory (LSTM) network [26], with each layer having 2048 memory blocks. We use the transcriptions of the Librispeech training set, as well as the texts of 14500 public domain books as the training data, and the vocabulary size is 89112. The LM is trained with SGD optimizer for 4 data epochs until it fully converges. This LM is used in all the decoding in the present paper.

The architecture of the RNN-T base model is as follows. The encoder consists of 6 bidirectional LSTM layers with 1024 memory blocks per layer per direction. The joint network is a feed-forward linear network which projects the 1024-dimensional stacked encoder vectors from the last encoder layer. The decoder is a 1-layer LSTM which contains 1024 memory blocks. Dropout [27] is applied in the decoder network. The RNN-T’s output units are English characters. The use of characters as output units helps to simplify and generalize the training of the RNN-T base models even though these models might not be as strong as those using context-dependent phonetic units, for instance word-pieces, as output units [28].

40-dimensional filter-bank are extracted as acoustic features. The filter-bank features are appended with 3-dimensional pitch features. Delta and acceleration features are computed from static features. Static and dynamic features are appended in one feature vector of 129 dimensions. The feature extraction is performed by using Kaldi toolkit [29]. The model training is done with 20 epochs and a beam size of 20 is used in the decoding. Training and decoding are performed by using ESPnet toolkit [30] and Pytorch [31].

## 4.3. Results

### 4.3.1. Multiple hypotheses with comparable WERs

We examine the effectiveness of the multiple-hypothesis unsupervised fine-tuning when multiple hypotheses with comparable WERs are used. These hypotheses can be obtained by creating a variant of the base model and use the variant model to decode the test audio data to obtain additional hypotheses. In this work, we simply change the dropout rate in the RNN-T’s decoder network to train the “Base model 2” which is a variant of the “Base model 1” (see Fig. 1). Decoding the test data with the base model 2 produces a variant of the ASR hypotheses compared to that produced by the base model 1 [32, 33]. Their WERs can be seen in the first two lines of Table 1.

Table 1 shows experimental results when two hypotheses are used in the multiple-hypothesis unsupervised fine-tuning. The results are obtained with the fine-tuned model 3 (see Fig. 1), after two fine-tuning iterations. The base model 2 applies a dropout rate of 0.5 while the base model 1 applies a dropout

rate of 0.1 in the decoder network. It can be seen in Table 1 that both the unsupervised fine-tunings using single hypothesis (SH) and multiple hypotheses (MH) yield WER reductions compared to the base models. The MH unsupervised fine-tuning outperforms the conventional SH unsupervised fine-tuning for both `test_clean` and `test_other` sets. The relative WER reductions obtained by the MH fine-tuning compared to the SH fine-tuning on `test_clean` and `test_other` sets are 12.7% and 14.2%, respectively.

Table 1: *Unsupervised fine-tuning experiments on Librispeech. SH denotes single hypothesis and MH denotes multiple hypotheses. Two ASR hypotheses are used in the fine-tuning using multiple ASR hypotheses.*

Fine-tuning method	WER (in %)	
	Test_clean	Test_other
Base model 1	8.5	22.9
Base model 2	8.1	23.1
SH fine-tuning (1 hypothesis)	7.9	22.6
MH fine-tuning (2 hypotheses)	6.9	19.4

### 4.3.2. Multiple hypotheses with different WERs

We also examine the effectiveness of the multiple-hypothesis fine-tuning when the additional model used to decode the test audio data has a substantial lower WER than the model that needs to be fine-tuned. This situation can happen when, for instance, a model from one site is used to help fine-tuning a model at another site and the performance of the two or more models are quite different.

In this experiment, a second hypothesis with substantially lower WER compared to the one produced by the model being fine-tuned is used. By using this setting, we verify whether the fine-tuned model can achieve lower WER compared to the base models participated in the fine-tuning. In this experiment, we perform one iteration of fine-tuning instead of two.

Table 2: *Unsupervised fine-tuning experiments on Librispeech when 2 and 4 hypotheses are used in the fine-tuning using multiple ASR 1-best hypotheses.*

Fine-tuning method	WER (in %)	
	Test_clean	Test_other
Base model 1	8.5	22.9
Base model 3	6.8	20.0
MH fine-tuning (2 hypotheses)	6.6	18.5
Base model 2	8.1	23.1
Base model 4	7.0	20.7
MH fine-tuning (4 hypotheses)	5.7	17.2

Experimental results are shown in Table 2. In the fine-tuning using 2 hypotheses, we replace the hypotheses produced by the base model 2 used in Table 1 by the hypotheses produced by another model, the base model 3. The base model 1 is the model to be fine-tuned. The hypotheses produced by the base model 3 have substantial lower WER compared to those produced by the base model 1. The base model 3 is trained similarly to the base model 1 where the dropout rates in the decoder networks of both models equal 0.1, but with 3-fold speed perturbation [34] applied during training. The base model 1 is trained on 960 hours of Librispeech training data without speed perturbation. It can be seen in Table 2 that the MH fine-tuning using two hypotheses creates a fine-tuned model which yields lower WERs on `test_clean` and `test_other` sets compared to the base model 3.

We include two additional hypotheses, one produced by the base model 2, and another produced by a base model 4 which is trained similarly to the base model 3 but with the decoder network’s dropout rate equals 0.5. When including all 4 hypotheses in the fine-tuning using multiple hypotheses, we obtain further WER reductions on both `test_clean` and `test_other` sets. The fine-tuned model, obtained by the fine-tuning using 4 hypotheses, yields 16.2% and 14.0% relative WER reductions on `test_clean` and `test_other` sets, respectively, compared to the best base model (base model 3).

## 5. Self-training

### 5.1. Method

Similar to the unsupervised fine-tuning experiment, in self-training, training data with manual transcriptions are available to train a base model. We study the scenario where additional audio data for training is available but the manual transcriptions, or labels, of these data are not available. In conventional self-training using pseudo-labeling, a base model (base model 1, see Fig. 2) is trained using the supervised audio training data. This model is then used to decode the additional training data. In this work, the base model 2 is trained similarly to the base model 1. The only difference between these models is the dropout rate of the decoder network which equals 0.1 in the base model 1 and 0.5 in the base model 2. The additional audio data and ASR 1-best hypotheses are then combined with the supervised training data to do semi-supervised training of the final model.

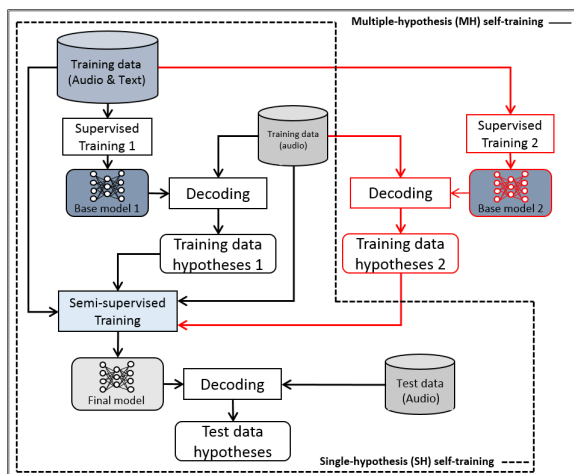


Figure 2: Self-training of RNN-T model using additional unlabeled audio data and multiple ASR 1-best hypotheses. The blocks surrounded by the dashed lines are used in conventional single-hypothesis self-training. The red outline blocks and arrows are used only for multiple-hypothesis self-training.

In self-training using multiple ASR hypotheses, a variant of the base model is trained using the supervised training data to create the base model 2 (see Fig. 2). This model is used to decode the additional training data and the obtained hypotheses are used in the multiple-hypothesis semi-supervised training. The SH and MH self-training are depicted in Fig. 2.

### 5.2. Data and system

We use the WSJ’s clean training data and Aurora-4’s multi-condition training data to train the base models, with 96 hours of speech in total. Speed perturbation is applied to increase the training data to 288 hours. Around 9 hours of dev data is taken from Aurora-4’s dev data and is used during training. The system is evaluated on two test sets, one of 4.8 hours

extracted from Aurora-4’s 14 test sets, and one of 2.2 hours which is the CHiME-4’s single-channel real noisy evaluation set, `et05_real_isolated_1ch_track`. 17.5-hour additional training data is taken from the `tr05_real_noisy` training set of CHiME-4 which consists of 9600 utterances recorded in real noisy environments [12]. When being added to the original training data, speed perturbation is applied to the additional training data to increase the amount of the added training data to 52.5 hours. The RNN-T model in this experiment has the same architecture and specifications as that used in section 4.2. The RNN-LM mentioned in section 4.2 is used in the decoding. The WERs of the base model 1 and base model 2 on the additional training data are 52.9% and 54.7%, respectively.

### 5.3. Results

Table 3 shows the results of the self-training experiments. The results show that both SH and MH self-training yield WER reductions on the test data from Aurora-4 and CHiME-4. The MH self-training yields larger WER reductions compared to the SH self-training. On Aurora-4 test data, the MH self-training yields 6.8% and 6.0% relative WER reductions compared to the base model 1, which is the better base model, and the SH self-training, respectively. On CHiME-4 evaluation data, the MH self-training yields 7.8% and 3.3% relative WER reductions compared to the base model 1 and the SH self-training, respectively. The “Supervised training” in which the manual transcriptions of the additional training data are used yields the lowest WERs on both Aurora-4 and CHiME-4 test data, as can be seen in Table 3.

Table 3: Self-training experiments using training data from WSJ, Aurora-4, and CHiME-4. Evaluation data are from Aurora-4 and CHiME-4.

Self-training method	WER (in %)	
	Aurora-4	CHiME-4
Base model 1	23.6	60.4
Base model 2	24.1	61.1
SH self-training	23.4	57.6
MH self-training	22.0	55.7
Supervised training	21.0	44.2

## 6. Conclusion

We have proposed a new approach to perform unsupervised fine-tuning and self-training of RNN-T E2E ASR model. The proposed approach used unlabeled speech data and multiple ASR 1-best hypotheses. A multiple-hypothesis RNN-T loss was proposed to incorporate multiple ASR 1-best hypotheses into the loss function which was used during fine-tuning and self-training. In the fine-tuning experiments performed with Librispeech, the approach using multiple ASR 1-best hypotheses yielded 14.2% relative WER reductions compared to the conventional single-hypothesis approach, on the `test_other` set. In the self-training experiments performed with WSJ, Aurora-4, and CHiME-4 data, the multiple-hypothesis approach yielded 3.3% relative WER reduction compared to the conventional single-hypothesis approach, on the CHiME-4’s single-channel real noisy evaluation set.

These results show that incorporating multiple ASR 1-best from different systems in the RNN-T loss is beneficial in alleviating the influence of ASR errors and helps improving the unsupervised fine-tuning and self-training performance. Future works will explore new ways of integrating multiple ASR hypotheses into the loss and solutions to reduce the computational complexity of the current approach, including the use of  $N$ -best hypotheses from one ASR system.

## 7. References

- [1] A. Graves, “Sequence transduction with recurrent neural networks,” in *arXiv preprint arXiv: 1211.3711*, 2012.
- [2] 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. 2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, Okinawa, Japan, December 2017, pp. 206–213.
- [3] J. Li, R. Zhao, H. Hu, and Y. Gong, “Improving RNN transducer modeling for end-to-end speech recognition,” in *Proc. 2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, Sentosa, Singapore, December 2019, pp. 114–121.
- [4] J. Li, Y. Wu, Y. Gaur, C. Wang, R. Zhao, and S. Liu, “On the comparison of popular end-to-end models for large scale speech recognition,” in *Proc. INTERSPEECH*, Shanghai, China, October 2020, pp. 1–5.
- [5] A. Graves, S. Fernandez, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proc. of the 23rd International Conference on Machine Learning*, Pittsburgh, USA, June 2006, pp. 369–376.
- [6] J. Li, R. Zhao, Z. Meng, Y. Liu, W. Wei, S. Parthasarathy, V. Mazalov, Z. Wang, L. He, S. Zhao, and Y. Gong, “Developing RNN-T models surpassing high-performance hybrid models with customization capability,” in *Proc. INTERSPEECH*, Shanghai, China, October 2020, pp. 3590–3594.
- [7] D.-H. Lee, “Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks,” in *Proc. ICML 2013 Workshop: Challenge in Representation Learning (WREPL)*, Atlanta, USA, June 2013.
- [8] J. Kahn, A. Lee, and A. Hannun, “Self-training for end-to-end speech recognition,” in *Proc. 2020 IEEE ICASSP*, Barcelona, Spain, May 2020, pp. 7084–7088.
- [9] G. Synnaeve, Q. Xu, J. Kahn, T. Likhomanenko, E. Grave, V. Pratap, A. Sriram, V. Liptchinsky, and R. Collobert, “End-to-end ASR: from supervised to semi-supervised learning with modern architectures,” in *Proc. ICML 2020 Workshop: Self-supervision in Audio and Speech (SAS)*, Vienna, Austria, July 2020.
- [10] C.-T. Do, R. Doddipatla, and T. Hain, “Multiple-hypothesis CTC-based semi-supervised adaptation of end-to-end speech recognition,” in *Proc. 2021 IEEE ICASSP*, Toronto, Canada, June 2021, pp. 6978–6982.
- [11] N. Parihar and J. Picone, *Aurora working group: DSR front end LVCSR evaluation: AU/384/02*. Institute for Signal and Information Processing Technical Report, 2002.
- [12] E. Vincent, S. Watanabe, A. A. Nugraha, J. Barker, and R. Marxer, “An analysis of environment, microphone and data simulation mismatches in robust speech recognition,” *Computer Speech and Language*, vol. 46, pp. 535–557, September 2017.
- [13] D. B. Paul and J. M. Barker, “The design for the Wall Street Journal-based CSR corpus,” in *HLT ’91 Proceedings of the workshop on Speech and Natural Language*, New York, USA, February 1992, pp. 357–362.
- [14] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *Proc. 2015 IEEE ICASSP*, Brisbane, Australia, April 2015, pp. 5206–5210.
- [15] P. Bell, J. Fainberg, O. Klejch, J. Li, S. Renals, and P. Swietojanski, “Adaptation algorithms for neural network-based speech recognition: an overview,” *IEEE Open Journal of Signal Processing*, vol. 2, pp. 33–66, 2021.
- [16] Y. Huang, J. Li, L. He, W. Wenning, W. Gale, and Y. Gong, “Rapid RNN-T adaptation using personalized speech synthesis and neural language generator,” in *Proc. INTERSPEECH*, Shanghai, China, October 2020, pp. 1256–1260.
- [17] G. Saon, Z. Tuske, D. Bolanos, and B. Kingsbury, “Advancing RNN transducer technology for speech recognition,” in *Proc. 2021 IEEE ICASSP*, Toronto, Canada, June 2021, pp. 5654–5658.
- [18] J. Pylkkonen, A. Ukkonen, J. Kilpikoski, S. Tamminen, and H. Heikinheimo, “Fast text-only domain adaptation of RNN-transducer prediction network,” in *Proc. INTERSPEECH*, Brno, Czechia, August-September 2021, pp. 1882–1886.
- [19] D. Giuliani and F. Brugnara, “Experiments on cross-system acoustic model adaptation,” in *Proc. IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, Kyoto, Japan, Dec. 2007, pp. 117–122.
- [20] S. Stueker, C. Fuegen, S. Burger, and M. Woelfel, “Cross-system adaptation and combination for continuous speech recognition: the influence of phoneme set and acoustic front-end,” in *Proc. INTERSPEECH*, Pittsburgh, USA, September 2006, pp. 521–524.
- [21] M. Gibson and T. Hain, “Hypothesis spaces for minimum Bayes risk training in large vocabulary speech recognition,” in *Proc. INTERSPEECH*, Pittsburgh, USA, September 2006, pp. 2406–2409.
- [22] J. Guo, G. Tiwari, J. Droppo, M. Van Segbroeck, C.-W. Huang, A. Stolcke, and R. Maas, “Efficient minimum word error rate training of RNN-transducer for end-to-end speech recognition,” in *Proc. INTERSPEECH*, Shanghai, China, October 2020, pp. 2807–2811.
- [23] N. Moritz, T. Hori, and J. Le Roux, “Semi-supervised speech recognition via graph-based temporal classification,” in *Proc. 2021 IEEE ICASSP*, Toronto, Canada, June 2021, pp. 6548–6552.
- [24] H. Meinedo and J. P. Neto, “Combination of acoustic models in continuous speech recognition hybrid systems,” in *Proc. 6th International Conference on Spoken Language Processing (ICSLP 2000)*, Beijing, China, October 2000, pp. 931–934.
- [25] S. Ruder, “An overview of gradient descent optimisation algorithms,” in *arXiv preprint arXiv: 1609.04747*, 2016.
- [26] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, pp. 1735–1780, November 1997.
- [27] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: a simple way to prevent neural networks from overfitting,” *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [28] K. Irie, R. Prabhavalkar, A. Kannan, A. Bruguier, D. Rybach, and P. Nguyen, “On the choice of modeling unit for sequence-to-sequence speech recognition,” in *Proc. INTERSPEECH*, Graz, Austria, September 2019, pp. 3800–3804.
- [29] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Vesely, “The Kaldi speech recognition toolkit,” in *Proc. IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, Hawaii, USA, December 2011.
- [30] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, “ESPnet: end-to-end speech processing toolkit,” in *Proc. INTERSPEECH*, Hyderabad, India, September 2018, pp. 2207–2211.
- [31] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [32] S. Dey, P. Motlicek, T. Bui, and F. Deroncourt, “Exploiting semi-supervised training through a dropout regularization in end-to-end speech recognition,” in *Proc. INTERSPEECH*, Graz, Austria, September 2019, pp. 734–738.
- [33] S. Khurana, N. Moritz, T. Hori, and J. Le Roux, “Unsupervised domain adaptation for speech recognition via uncertainty driven self-training,” in *Proc. 2021 IEEE ICASSP*, Toronto, Canada, June 2021, pp. 6553–6557.
- [34] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, “Audio augmentation for speech recognition,” in *Proc. INTERSPEECH*, Dresden, Germany, September 2015, pp. 3586–3589.