A COMPARISON OF LATTICE-FREE DISCRIMINATIVE TRAINING CRITERIA FOR PURELY SEQUENCE-TRAINED NEURAL NETWORK ACOUSTIC MODELS

Chao Weng, Dong Yu

Tencent AI Lab, Bellevue, USA

{cweng,dyu}@tencent.com

ABSTRACT

In this work, three lattice-free (LF) discriminative training criteria for purely sequence-trained neural network acoustic models are compared on LVCSR tasks, namely maximum mutual information (MMI), boosted maximum mutual information (bMMI) and state-level minimum Bayes risk (sMBR). We demonstrate that, analogous to LF-MMI, a neural network acoustic model can also be trained from scratch using LF-bMMI or LF-sMBR criteria respectively without the need of cross-entropy pre-training. Furthermore, experimental results on Switchboard-300hrs and Switchboard+Fisher-2100hrs datasets show that models trained with LF-bMMI consistently outperform those trained with plain LF-MMI and achieve a relative word error rate (WER) reduction of \sim 5% over competitive temporal convolution projected LSTM (TDNN-LSTMP) LF-MMI baselines.

Index Terms— acoustic modeling, sequence discriminative training, LF-MMI, boosted MMI, sMBR

1. INTRODUCTION

Sequence discriminative training has been proven to improve neural network (NN) acoustic models (AMs) significantly for LVCSR tasks [1, 2, 3, 4, 5]. Most widely used discriminative training criteria include maximum mutual information (MMI) [6, 7], minimum classification error (MCE) [8], boosted MMI (bMMI) [9], minimum phone/word error (MPE/MWE) [10] and state-level minimum Bayes risk (sMBR) [2, 4, 11]. Discriminative training of a NN is typically initialized from a frame-level cross-entropy pre-trained model which is also used to generate alignments and lattices, i.e, the reference and hypothesis sequences. Recently, lattice-free (LF) MMI [12] was proposed to train a NN from scratch using MMI criterion without initializing from a cross-entropy model and achieved state-of-the-art performances on various LVCSR tasks. To make LF-MMI training practically feasible, a phone-level, instead of word-level, language model (LM) is used to generate a unified denominator cyclic graph, eliminating the need of the lattice generation process. LF-MMI has gained its popularity lately as it not only saves the cross-entropy pre-training effort but also leads to a better sequence-trained model in general for LVCSR. LF-MMI and LF-bMMI were adopted in [13] and [14] respectively on English conversation transcription and multi-talker speech recognition task but in both work cross-entropy pre-training is still required. Later on, various model architectures [15, 16, 17] for LF-MMI have been explored. LF-sMBR was proposed and compared against lattice-based sMBR in [18] when initialized from a LF-MMI trained model. Very recently, various lattice-free discriminative training criteria are investigated for keyword spotting task [19]. However, there are few, if any, prior works having presented a comparison between different discriminative training criteria with which NN AMs are purely sequence-trained from scratch for LVCSR tasks.

In this work, three lattice-free discriminative training criteria for purely sequence-trained NN AMs are compared on LVCSR tasks, MMI, bMMI and sMBR. We demonstrate that, analogous to LF-MMI, a NN AM can also be trained from scratch using LF-bMMI or LF-sMBR criteria respectively without the need of cross-entropy pre-training. Furthermore, experimental results on Switchboard-300hrs and Switchboard+Fisher-2100hrs datasets show that models trained with LF-bMMI consistently outperform those trained with plain LF-MMI and achieve a relative word error rate (WER) improvement of $\sim 5\%$ over competitive TDNN-LSTMP LF-MMI baselines.

The remainder of the paper is organized as follows. In Section 2, we elaborate lattice-free discriminative training criteria and then describe implementation details of LF-bMMI and LF-sMBR. All the experimental setups and results are presented in Section 3. We conclude our work in Section 4.

2. LATTICE-FREE SEQUENCE DISCRIMINATIVE TRAINING

2.1. LF-MMI

Given U pairs of the training speech utterance O_u and its reference word transcription W_u^r , the objective function of LF-MMI is given by,

$$\mathcal{F}_{\text{LF}-\text{MMI}} = \sum_{u=1}^{U} \log \frac{\sum_{\mathbf{S}_{u}^{r}} p(\mathbf{O}_{u} | \mathbf{S}_{u}^{r}) P(\mathbf{W}_{u}^{r})}{\sum_{\mathbf{S}_{u}} p(\mathbf{O}_{u} | \mathbf{S}_{u}) P(\mathbf{W}_{u})}, \quad (1)$$

where \mathbf{W}_u is a hypothesis sequence from the denominator graph. \mathbf{S}_u^r and \mathbf{S}_u are state sequences corresponding to \mathbf{W}_u^r and \mathbf{W}_u respectively. Note that there could be multiple \mathbf{S}_u corresponding to one \mathbf{W}_u due to multiple pronunciations and label tolerance, i.e., a tolerance which allows a state to occur slightly before or after where it appear in the original reference alignments or lattices [12]. Let \mathbf{o}_t and \mathbf{s}_t be the *t*th frame of \mathbf{O}_u and *t*th state of \mathbf{S}_u respectively, i.e., $\mathbf{O}_u =$ $\{\mathbf{o}_1, ..., \mathbf{o}_t, ..., \mathbf{o}_{T_u}\}$, $\mathbf{S}_u = \{\mathbf{s}_1, ..., \mathbf{s}_{T_u}\}$. The NN output at time *t* is interpreted as pseudo log-likelihood of \mathbf{o}_t given a particular state *j*, i.e., $\log p(\mathbf{o}_t | \mathbf{s}_t = j)$, and there is neither acoustic scaling or division by the prior in LF-MMI. The gradient required to perform LF-MMI training can be written as,

$$\frac{\partial \mathcal{F}_{\text{LF}-\text{MMI}}}{\partial \log p(\mathbf{o}_t | \mathbf{s}_t = j)} = \sum_{u=1}^{U} [\gamma_u^{\text{num}}(\mathbf{s}_t = j) - \gamma_u^{\text{den}}(\mathbf{s}_t = j)],$$
(2)

where $\gamma_u^{\text{num}}(\mathbf{s}_t = j)$ and $\gamma_u^{\text{den}}(\mathbf{s}_t = j)$ are posterior probabilities of *t*th state being *j* computed over the reference and hypothesis state sequences respectively given the input utterance \mathbf{O}_u .

2.2. LF-bMMI

In LF-bMMI, the objective function is modified to boost the likelihoods of those hypothesis sequences commit more errors,

$$\mathcal{F}_{\rm LF-bMMI} = \sum_{u=1}^{U} \log \frac{\sum_{\mathbf{S}_{u}^{r}} p(\mathbf{O}_{u} | \mathbf{S}_{u}^{r}) P(\mathbf{W}_{u}^{r})}{\sum_{\mathbf{S}_{u}} p(\mathbf{O}_{u} | \mathbf{S}_{u}) P(\mathbf{W}_{u}) e^{-bA(\mathbf{S}_{u}, \mathbf{S}_{u}^{r})}},$$
(3)

where b is the boosting factor and $A(\mathbf{S}_{u}^{r}, \mathbf{S}_{u})$ is the accuracy function which measures the number of matching labels between reference and hypothesis sequences. The gradient of LF-bMMI has the same form as LF-MMI in Eq. (2) whereas the difference from LF-MMI lies in the forward-backward implementation when calculating the $\gamma_{u}^{\text{den}}(\mathbf{s}_{t} = j)$ term. See more details in Subsection 2.4.

2.3. LF-sMBR

The objective function of LF-sMBR criterion is,

$$\mathcal{F}_{\rm LF-sMBR} = \sum_{u=1}^{U} \frac{\sum_{\mathbf{S}_{u}} p(\mathbf{O}_{u} | \mathbf{S}_{u}) P(\mathbf{W}_{u}) A(\mathbf{S}_{u}, \mathbf{S}_{u}^{r})}{\sum_{\mathbf{S}_{u}^{\prime}} p(\mathbf{O}_{u} | \mathbf{S}_{u}^{\prime}) P(\mathbf{W}_{u}^{\prime})}.$$
 (4)

The gradient required to perform LF-sMBR training can be written as,

$$\frac{\partial \mathcal{F}_{\rm LF-sMBR}}{\partial \log p(\mathbf{o}_t | \mathbf{s}_t = j)} = \sum_{u=1}^U \gamma_u^{\rm den}(\mathbf{s}_t = j) [\overline{A}_u(\mathbf{s}_t = j) - \overline{A}_u],$$
(5)

where $\overline{A}_u(\mathbf{s}_t = j)$ is the average state-level accuracy of those hypothesis state sequences that have state j at frame t and \overline{A}_u is the average accuracy of all state sequences in the denominator graph. $\gamma_u^{\text{den}}(\mathbf{s}_t = j)$ is the same posterior probability term as in Eq. (2).

Algorithm 1	LF-sMBR	Forward	Procedure	for Leak	y HMM
-------------	---------	---------	-----------	----------	-------

 $T \leftarrow$ number of frames; $S \leftarrow$ number of HMM states in the denominator graph; $\lambda \leftarrow$ leaky HMM coefficient for $s \leftarrow 1$ to S do 1: $P_0(s) \leftarrow$ initial probability of state s 2: 3: $\alpha(s,0) \leftarrow P_0(s)$ 4: $\alpha_{\rm mbr}(s,0) \leftarrow 0$ 5: end for for $s \leftarrow 1$ to S do 6: $\boldsymbol{\alpha}(s, 0) \leftarrow \boldsymbol{\alpha}(s, 0) + \lambda P_0(s) \sum_{1}^{S} \boldsymbol{\alpha}(s, 0)$ 7: $\alpha_{\mathrm{mbr}}(s,0) \leftarrow \alpha_{\mathrm{mbr}}(s,0) + \lambda P_0(s) \sum_{1}^{S} \alpha_{\mathrm{mbr}}(s,0)$ 8: $\alpha_{\mathrm{mbr}}(s,0) \leftarrow \alpha_{\mathrm{mbr}}(s,0) / \alpha(s,0)$ 9: 10: end for for $t \leftarrow 1$ to T do 11: for $s \leftarrow 1$ to S do 12: $\alpha(s,t) \leftarrow 0$ 13: 14: $\alpha_{\rm mbr}(s,t) \leftarrow 0$ for arc $a \in$ all arcs preceding state s do 15: $P_{\text{trans}} \leftarrow \text{transition probability of arc } a$ 16: $j \leftarrow$ state label (index of NN output) of arc a 17: $s_{\text{prev}} \leftarrow \text{starting state of arc } a$ 18: $\delta = \alpha(s_{\text{prev}}, t-1) * P_{\text{trans}} * p(\mathbf{o}_t | \mathbf{s}_t = j)$ 19: 20: $\alpha(s,t) \leftarrow \alpha(s,t) + \delta$ $\alpha_{\rm mbr}(s,t) \leftarrow [\alpha_{\rm mbr}(s_{\rm prev},t-1) + \gamma_u^{\rm num}(\mathbf{s}_t =$ 21: $j)] * \delta$ end for 22: end for 23: for $s \leftarrow 1$ to S do 24: $\begin{aligned} \alpha(s,t) &\leftarrow \alpha(s,t) + \lambda P_0(s) \sum_{1}^{S} \alpha(s,t) \\ \alpha_{\rm mbr}(s,t) &\leftarrow \alpha_{\rm mbr}(s,t) + \lambda P_0(s) \sum_{1}^{S} \alpha_{\rm mbr}(s,t) \end{aligned}$ 25: 26: $\alpha_{\rm mbr}(s,t) \leftarrow \alpha_{\rm mbr}(s,t)/\alpha(s,t)$ 27: end for 28: 29: end for

2.4. Implementation of LF-bMMI and LF-sMBR

Since LF-MMI implementation is available in the Kaldi toolkit [20], we mainly discuss the implementation details of LF-bMMI and LF-sMBR in this subsection. Due to multiple pronunciations and label tolerance, the supervision used in LF-MMI is a numerator graph which encodes multiple reference state sequences. In both our LF-bMMI and LF-sMBR implementations, the term $\gamma_u^{\text{num}}(\mathbf{s}_t = j)$ derived from the numerator graph is used to calculate per-frame state-level accuracy.

For LF-bMMI, the only extra computation as opposed to plain LF-MMI is to boost per-frame likelihood according to $\gamma_u^{\text{num}}(\mathbf{s}_t=j)$ when forward-backward is performed on the denominator graph, specifically,

$$p'(\mathbf{o}_t | \mathbf{s}_t = j) = p(\mathbf{o}_t | \mathbf{s}_t = j) \times e^{-b \times \gamma_u^{\text{num}}(\mathbf{s}_t = j)}.$$
 (6)

For LF-sMBR, the implementation closely follows latticebased sMBR but there are two notable variations. The first one lies in the leaky hidden Markov model (HMM) regularization [12] used in LF-MMI where an ϵ transition from each state a to each other state b is added to the denominator HMM graph with probability equal to a coefficient times the initial probability of state b. Since the leaky HMM is enabled in both our LF-MMI and LF-bMMI experiments, we also incorporate it in our LF-sMBR implementation. We list the implementation details of LF-sMBR forward procedure on a denominator leaky HMM graph in Algorithm 1 where Line 7, 8, 25 and 26 correspond to the forward computations for leaked ϵ transitions. Note that per-frame state-level accuracy for all ϵ transitions is zero. The second variation is we only need one single forward-backward pass during LF-sMBR training for efficiency as opposed to the lattice based sMBR implementation in the Kaldi [20] which requires two forward-backward passes. The backward procedure can be easily derived given the forward procedure listed in Algorithm 1.

3. EXPERIMENTS

We conduct our experiments on both Switchboard-300hrs and Switchboard+Fisher-2100hrs datasets to compare LF-MMI, LF-bMMI and LF-sMBR for LVCSR. Speed perturbation is used to augment the training data three times in both Switchboard-300hrs and Switchboard+Fisher-2100hrs experiments. We use the same set of numerator and denominator lattices/graphs for all three training criteria. The numerator graphs are generated using a GMM-HMM system trained on LDA+STC+FMLLR MFCC features described in [2] and the denominator graph is generated using a phone language model estimated from the phone alignments on all training transcriptions using the GMM-HMM system. TDNN-LSTMP architecture [15, 17] is adopted in all our experiments as it was found to outperform bidirectional LSTMs and achieve state-of-the-art performances on the Switchboard-300hrs dataset when trained with LF-MMI. We listed the detailed model configuration used in all of our experiments in Table 1. The input features to the TDNN-LSTMP are 40-dim MFCCs and 100-dim online i-vectors. All TDNN-LSTMPs are trained from scratch without any type of pre-training but we apply cross-entropy regularization during training and the smooth factor is set to 0.025 unless otherwise indicated. The leaky HMM coefficient is set to 0.1 through all our experiments. The initial and final learning rates are set to 0.001 and 0.0001 respectively for all three criteria. We adopt the default parallel optimizer [21] in the Kaldi toolkit [20] for all training experiment and use 3 GPUs initially and ramp up to 8 in the end to train TDNN-LSTMPs. Note that there is no early

Table 1. The model configuration used in all our experiments

Layer	Context	Layer-type
1	[-2,-1,0,1,2]	TDNN
2	[-1,0,1]	TDNN
3	[-1,0,1]	TDNN
4	[0]	LSTMP
5	[-3, 0, 3]	TDNN
6	[-3, 0, 3]	TDNN
7	[0]	LSTMP
8	[-3, 0, 3]	TDNN
9	[-3, 0,3]	TDNN
10	[0]	LSTMP
11	[0]	Projection

Table 2. WERs of TDNN-LSTMP LF-MMI baseline and LF-bMMI on Eval2000 using Switchboard-300hrs data

Models	h	WERs(%)		
Models	0	SWB	Total	
LF-MMI	0.0	9.1	14.2	
LF-bMMI	0.05	8.7	13.7	
LF-bMMI	0.15	8.7	13.7	
LF-bMMI	0.20	8.7	13.6	
LF-bMMI	0.25	8.8	13.6	
LF-bMMI	0.30	8.8	13.7	

stop for this optimizer, one has to specify a fixed number of training epochs before each experiment. Both tri-gram and 4-gram LMs are estimated on Switchboard+Fisher-2100hrs transcriptions and all WER results are reported using the 4-gram LM to re-score the lattices that are generated with the trigram LM.

3.1. Experiments on Switchboard-300hrs

After building a tree using the special two-state HMM topology described in [12], the output dimension of TDNN-LSTMPs on Switchboard-300hrs is 5994 which leads the number of TDNN-LSTMP's total parameters to 39.5 million. This is a pretty large model on 300hrs training data. Therefore, we use dropout and follow the same schedule in [22] to overcome over-fitting. We first compare TDNN-LSTMP LF-MMI baseline with LF-bMMI using different boosting factors b from 0.05 to 0.30 in Table 2. The number of training epochs for both LF-MMI and LF-bMMI is set to 4, which appears to be optimal according to the Kaldi toolkit [20]. It shows that TDNN-LSTMPs trained with LF-bMMI outperform the LF-MMI baseline consistently and the best LF-bMMI model achieves an absolute WER improvement of 0.4% and 0.6% on the Switchboard subset and full set of Eval2000 respectively. Then LF-sMBR is compared against LF-MMI in Table 3. During LF-sMBR experiments, it is found that if we use

 Table 3. WERs of TDNN-LSTMP LF-MMI baseline and LF-sMBR with different configurations on Eval2000 using Switchboard-300hrs data

Models -	Label	x-ent	Sil-acc.	WER	Ls(%)
#Epochs	tolerance	smooth	Scale	SWB	Total
LF-MMI - 4	± 5	0.025	NA	9.1	14.2
LF-sMBR - 4	± 5	0.025	0.001	20.3	23.9
LF-sMBR - 5	± 1	0.025	0.001	17.1	20.9
LF-sMBR - 8	± 1	0.025	0.001	16.2	19.7
LF-sMBR - 8	± 0	0.025	0.01	12.0	16.9
LF-sMBR - 12	± 0	0.025	0.01	11.6	16.2
LF-sMBR - 12	± 0	0.025	0.013	11.2	15.8
LF-sMBR - 12	± 0	0.025	0.015	11.5	16.3
LF-sMBR - 12	± 0	0.050	0.013	11.2	16.0

the original $\gamma_u^{\text{num}}(\mathbf{s}_t = j)$ term for the silence state accuracy, the model diverges after a few tens of training iterations. So we simply apply a scale, eg., $0 < \mu < 1$, to $\gamma_{\mu}^{\text{num}}(\mathbf{s}_t = j)$ if state j belongs to silence. Note that we also tried zeroing the gradients for silence states, i.e., $\gamma_u^{\text{den}}(\mathbf{s}_t = j) = 0$ if $j \in$ silence, as in [18], but the number of insertion errors is disproportionately high due to the fact that the silence model is not being trained well. Another critical configuration for LF-sMBR training is the label tolerance, which allows a state drifting away from where it appear in the original alignment. As shown in Table 3, the setup using the numerator graph derived from the original alignments, i.e., no label tolerance, achieves better performances. We experiment different silence state accuracy scales from 0.001 to 0.02 and find small scales like 0.001 will lead to too many insertion errors due to insufficient silence modeling and a scale of 0.02 will make the training diverged. It is also found that LF-sMBR converges slower than LF-MMI or LF-bMMI, and there are improvements when we increase the number of training epochs from 8 to 12. The best LF-sMBR setup achieves WERs of 11.2% and 15.8% on the Switchboard subset and full set of Eval2000 which unfortunately still shows an inferior performance to the LF-MMI baseline. Therefore, we will only compare LF-MMI and LF-bMMI on the full Switchboard+Fisher-2100hrs dataset.

3.2. Experiments on Switchboard+Fisher-2100hrs

We then use the full Switchboard+Fisher2100hrs dataset as training data to compare LF-MMI and LF-bMMI criteria. The output size of TDNN-LSTMP is 7266, a bit larger than the one used in the 300hrs experiment. As we have much larger amount of training data, the dropout is disabled for all 2100hrs experiments. For both LF-MMI and LF-bMMI experiments on 2100hrs, we set the number of training epochs to 4 and report the WER results in Table 4. The LF-MMI base-line obtain WERs of 8.1% and 15.5% on Switchboard and

Table 4. WERs of TDNN-LSTMP LF-MMI baseline and LF-bMMI on Eval2000 using Switchboard+Fisher-2100hrs data

Models	h	WERs(%)			
widdels	0	SWB	CH	Total	
LF-MMI	0.0	8.1	15.5	12.0	
LF-bMMI	0.05	8.1	15.2	11.7	
LF-bMMI	0.10	7.7	14.7	11.3	
LF-bMMI	0.15	7.9	14.9	11.5	

Callhome subsets of Eval2000 which are very competitive to those results reported in [15] and [22]. Three LF-bMMI with different boosting factors are explored and the best setup achieves WERs of 7.7% and 14.7% on Switchboard and Callhome subsets respectively, a \sim 5% relative WER improvement over the TDNN-LSTMP LF-MMI baseline. Finally we compare our best LF-bMMI trained TDNN-LSTMP with previous published systems. Note that we only consider the reported results using similar n-gram LMs. As can be seen in Table 5, our best LF-bMMI trained TDNN-LSTMP is highly competitive to those state-of-the-art systems.

Table 5. Comparing our best model to previous publishedsystems built on the Switchboard+Fisher-2100hrs.

Systems	WERs(%)	
Systems	SWB	CH
BLSTM+LFMMI [12]	8.5	15.3
ResNet+xent+LFMMI [13]	8.6	15.2
LACE+xent+LFMMI [13]	8.5	15.2
ResNet+xent+MMI [23]	8.3	14.9
TDNN-NormOPGRU+LFMMI [17]	8.3	14.7
TDNN-LSTMP+LFbMMI (current)	7.7	14.7

4. CONCLUSION

In this work, three lattice-free (LF) discriminative training criteria for purely sequence-trained NN AMs are compared on LVCSR tasks, MMI, bMMI and sMBR. We demonstrate that, analogous to LF-MMI, a NN AM can also be trained from scratch using LF-bMMI or LF-sMBR criteria respectively without the need of cross-entropy pre-training. Furthermore, experimental results on Switchboard-300hrs and Switchboard+Fisher-2100hrs datasets show that models trained with LF-bMMI consistently outperform those trained with plain LF-MMI and achieve a WER reduction of $\sim 5\%$ over competitive TDNN-LSTMP LF-MMI baselines. Future work includes investigating why LF-sMBR falls behind both LF-MMI and LF-bMMI significantly and combining LF-sMBR with LF-MMI for multi-task sequence-trained NN AMs.

5. REFERENCES

- Brian Kingsbury, "Lattice-based optimization of sequence classification criteria for neural-network acoustic modeling," in *Proceedings of ICASSP*, 2009, ICASSP '09, pp. 3761–3764.
- [2] Karel Veselý, Arnab Ghoshal, Lukás Burget, and Daniel Povey, "Sequence-discriminative training of deep neural networks," in *INTERSPEECH 2013*, 2013, pp. 2345– 2349.
- [3] Hang Su, Gang Li, Dong Yu, and Frank Seide, "Error back propagation for sequence training of contextdependent deep networks for conversational speech transcription," in *ICASSP 2013*. January 2013, ICASSP.
- [4] G. Heigold, E. McDermott, V. Vanhoucke, A. Senior, and M. Bacchiani, "Asynchronous stochastic optimization for sequence training of deep neural networks," in *ICASSP 2014*, May 2014, pp. 5587–5591.
- [5] Matt Shannon, "Optimizing expected word error rate via sampling for speech recognition," *CoRR*, vol. abs/1706.02776, 2017.
- [6] L. R. Bahl, P. F. Brown, P. V. de Souza, and R. L. Mercer, "Maximum mutual information estimation of hidden Markov model parameters for speech recognition," in *ICASSP*, Tokyo, 1986, pp. 49–52.
- [7] P.C. Woodland and D. Povey, "Large scale mmie training for conversational telephone speech recognition," 2000.
- [8] B. Juang and S. Katagiri, "Discriminative learning for minimum error classification (pattern recognition)," *IEEE Transactions on Signal Processing*, vol. 40, no. 12, pp. 3043–3054, Dec 1992.
- [9] Daniel Povey, Dimitri Kanevsky, Brian Kingsbury, Bhuvana Ramabhadran, George Saon, and Karthik Visweswariah, "Boosted mmi for model and featurespace discriminative training.," in *ICASSP*. 2008, pp. 4057–4060, IEEE.
- [10] Dan Povey, Discriminative Training for Large Vocabulary Speech Recognition, Ph.D. thesis, University of Cambridge, 2003.
- [11] Daniel Povey and Brian Kingsbury, "Evaluation of proposed modifications to mpe for large scale discriminative training," *ICASSP '07*, vol. 4, pp. IV–321–IV–324, 2007.
- [12] Daniel Povey, Vijayaditya Peddinti, Daniel Galvez, Pegah Ghahremani, Vimal Manohar, Xingyu Na, Yiming Wang, and Sanjeev Khudanpur, "Purely sequencetrained neural networks for ASR based on lattice-free MMI," in *Interspeech 2016*, 2016, pp. 2751–2755.

- [13] Wayne Xiong, Jasha Droppo, Xuedong Huang, Frank Seide, Mike Seltzer, Andreas Stolcke, Dong Yu, and Geoffrey Zweig, "Achieving human parity in conversational speech recognition," *CoRR*, vol. abs/1610.05256, 2016.
- [14] Zhehuai Chen, Jasha Droppo, Jinyu Li, and Wayne Xiong, "Progressive joint modeling in unsupervised single-channel overlapped speech recognition," *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 26, no. 1, pp. 184–196, Jan. 2018.
- [15] Vijayaditya Peddinti, Yiming Wang, Daniel Povey, and Sanjeev Khudanpur, "Low latency acoustic modeling using temporal convolution and lstms," *IEEE Signal Process. Lett.*, vol. 25, no. 3, pp. 373–377, 2018.
- [16] Daniel Povey, Hossein Hadian, Pegah Ghahremani, Ke Li, and Sanjeev Khudanpur, "A time-restricted selfattention layer for ASR," in *ICASSP*. 2018, pp. 5874– 5878, IEEE.
- [17] Gaofeng Cheng, Daniel Povey, Lu Huang, Ji Xu, Sanjeev Khudanpur, and Yonghong Yan, "Output-gate projected gated recurrent unit for speech recognition," in *Interspeech*. 2018, pp. 1793–1797, ISCA.
- [18] Naoyuki Kanda, Yusuke Fujita, and Kenji Nagamatsu, "Lattice-free state-level minimum bayes risk training of acoustic models," in *Interspeech 2018*, 2018, pp. 2923– 2927.
- [19] Zhehuai Chen, Yanmin Qian, and Kai Yu, "Sequence discriminative training for deep learning based acoustic keyword spotting," *CoRR*, vol. 102, 100-111, 2018, 2018.
- [20] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Nagendra Goel, Mirko Hannemann, Yanmin Qian, Petr Schwarz, and Georg Stemmer, "The kaldi speech recognition toolkit," in *In IEEE 2011 workshop*, 2011.
- [21] Daniel Povey, Xiaohui Zhang, and Sanjeev Khudanpur, "Parallel training of deep neural networks with natural gradient and parameter averaging," *CoRR*, vol. abs/1410.7455, 2014.
- [22] Gaofeng Cheng, Vijayaditya Peddinti, Daniel Povey, Vimal Manohar, Sanjeev Khudanpur, and Yonghong Yan, "An exploration of dropout with lstms," in *INTER-SPEECH*, 2017.
- [23] George Saon, Tom Sercu, Steven J. Rennie, and Hong-Kwang Jeff Kuo, "The IBM 2016 english conversational telephone speech recognition system," *CoRR*, vol. abs/1604.08242, 2016.