



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Speaker-aware training of LSTM-RNNs for acoustic modelling

Citation for published version:

Tan, T, Qian, Y, Yu, D, Kundu, S, Lu, L, Xiao, X, SIM, KC & Zhang, Y 2016, Speaker-aware training of LSTM-RNNs for acoustic modelling. in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Institute of Electrical and Electronics Engineers (IEEE), pp. 5280-5284, 41st IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2016, Shanghai, China, 20/03/16. <https://doi.org/10.1109/ICASSP.2016.7472685>

Digital Object Identifier (DOI):

[10.1109/ICASSP.2016.7472685](https://doi.org/10.1109/ICASSP.2016.7472685)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



SPEAKER-AWARE TRAINING OF LSTM-RNNS FOR ACOUSTIC MODELLING

Tian Tan¹, Yanmin Qian^{1,2}, Dong Yu³, Souvik Kundu⁴, Liang Lu⁵, Khe Chai SIM⁴, Xiong Xiao⁶, Yu Zhang⁷

¹Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China

²Cambridge University Engineering Department, Cambridge, UK

³Microsoft Research, Redmond, USA

⁴School of Computing, National University of Singapore, Republic of Singapore

⁵Centre for Speech Technology Research, University of Edinburgh, UK

⁶Temasek Labs, Nanyang Technological University, Singapore.

⁷CSAIL, MIT, Cambridge, USA

ABSTRACT

Long Short-Term Memory (LSTM) is a particular type of recurrent neural network (RNN) that can model long term temporal dynamics. Recently it has been shown that LSTM-RNNs can achieve higher recognition accuracy than deep feed-forward neural networks (DNNs) in acoustic modelling. However, speaker adaption for LSTM-RNN based acoustic models is not well investigated. In this paper, we study the speaker-aware training for LSTM-RNNs that incorporate the speaker information during model training to normalise the speaker variability. We first present several different model architectures for this purpose, and then experimentally evaluate three different types of speaker representations, i.e. I-vectors, bottleneck speaker vectors and speaking rate. Furthermore, to factorize the variability in the acoustic signals caused by speakers and phonemes respectively, we investigate the speaker-aware and phone-aware joint training in the framework of multi-task learning. In AMI meeting speech transcription task, speaker-aware training of LSTM-RNNs improves word error rates by 6.5% relative over a very strong LSTM-RNN baseline, which has used fmlr features.

Index Terms— speaker-aware training, LSTM-RNNs, speaker adaptation, i-vector, speaking rate

1. INTRODUCTION

Deep learning has achieved tremendous success in acoustic modelling. With multiple hidden layers, the hybrid neural network and hidden Markov model (NN/HMM) [1] can obtain significant improvement in terms of recognition accuracy compared to the conventional Gaussian mixture models (GMMs) [2]. Previous studies mainly focus on the feed-forward neural networks using the acoustic features from a fixed-size context window, while recently, the recurrent neural networks (RNNs) have been demonstrated to be able to achieve higher recognition accuracy. The RNN has recurrent-connections on its hidden layers, which is expected to capture much longer temporal dynamics. However, training an RNN turns to be difficult due to the well-known gradient vanishing and exploding

problems [3]. In this context, Long Short-Term Memory (LSTM) [4] is proposed which overcomes these problems by introducing separate gate functions to control the flow of the information. For acoustic modelling, LSTM-RNNs are reported to outperform the DNNs on the large vocabulary tasks [5, 6].

However, a long standing problem in acoustic modelling is the mismatch between training and test data caused by speakers and environmental background. Although DNNs are more robust to the mismatch compared to GMMs, significant performance degradation has been observed [7]. There have been a number of studies to improve the robustness of DNN-HMM acoustic models, which can be divided into three categories [8], i.e. transformation based adaptation [9, 10], conservative training [11] and speaker-aware training [12, 13]. In particular, speaker dependent linear transformations for input or output layer were used for speaker adaptation in [9, 10], while in [11], KL distance was applied to adapt the DNN acoustic model with small amount of adaptation data without over-fitting. Finally, in [12, 13], a speaker auxiliary vector was introduced which allows DNNs to learn to normalise the speaker variability automatically. However, most of these studies have been limited to DNNs, and to our best knowledge, only the authors in [14] investigated speaker adaptation for LSTM-RNNs.

The approaches investigated in [14] fall in the category of transformation based adaptation, where linear input features transformation and hidden activations transformation were applied to adapt LSTM-RNNs. However, in this work, we focus on the speaker-aware training for LSTM-RNNs. We show in our study that this is not trivial, because different from DNNs, LSTM-RNNs are dynamic systems that are sensitive to the static input. To deal with this problem, we investigated different model architectures in order to make the speaker-aware training effective for LSTM-RNNs. We also evaluated these model architectures with three different speaker representations, namely, I-vector [15], bottleneck speaker vector (BSV) [16] and the speaking rate. Finally, we also managed to incorporate the phone informations in our adaptation network, which helped to further reduce the word error rate. Our experiments were performed on the AMI meeting speech transcription dataset, and we obtained 6.5% relative improvement over a strong LSTM-RNN baseline.

The work reported here was started at JSALT 2015 in UW, Seattle, and was supported by JHU via grants from NSF (IIS), DARPA (LORELEI), Google, Microsoft, Amazon, Mitsubishi Electric, and MERL. Tian Tan is supported by the Program for Professor of Special Appointment (Eastern Scholar) at Shanghai Institutions of Higher Learning, the China NSFC project No. 61222208, JiangSu NSF project No. 201302060012.

2. SPEAKER-AWARE TRAINING ON LSTM-RNNs

2.1. LSTM-RNNs

Compared to the feedforward neural networks, RNNs have the advantage of learning complex temporal dynamics in sequential data. However, in practice, training RNNs to learn long-term temporal dependency can be difficult due to the gradient vanishing and exploding problem [3]. The LSTM architecture provides a solution that partially overcomes the weakness of simple RNNs and achieves the state-of-the-art performance in speech recognition [5, 6]. More on standard form of LSTM is presented in [4], while in [5], a projection layer is applied to project the memory cells' outputs to a lower-dimensional vector which is particularly effective for speech recognition where the amount of training data is not abundant. In this work, we used the same LSTM-RNN architecture.

2.2. Adaptation architectures for speaker-aware training

Speaker-aware training is an approach that incorporates the speaker information into the network training process in order to normalise the speaker variability. Typical forms of representations of the speaker information are I-vector [15], bottleneck speaker vector (BSV) [16] or joint trained speaker-code [12]. For feed-forward neural networks, it usually works well by simply concatenating the speaker representations with the acoustic features. Similar approach may be applied for LSTM-RNNs as shown in Figure 1(a). However, this approach may not work well because the speaker auxiliary vectors are constant for a given speaker, adding a static vector into a dynamical system may not helpful on the ASR such as in [17, 18].

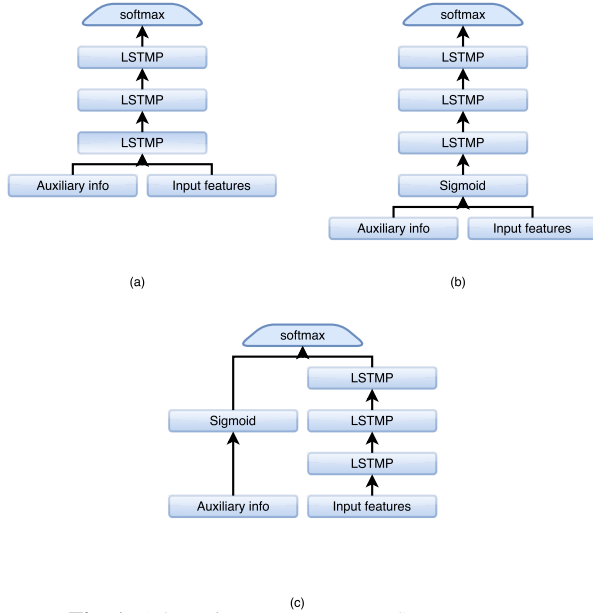


Fig. 1. Adaptation structures on LSTM-RNNs

To solve this problem, two structures are explored. Shown in figure 1(b), the auxiliary vector and acoustic feature are passing through a non-linear transformation first so that the input for LSTM-RNN becomes different among frames. Another structure is shown in figure 1(c), the auxiliary vector will go through a shallow NN first and then be concatenated to the output of the LSTM-RNN.

3. SPEAKER AUXILIARY VECTOR

In this paper, we investigate three kinds of auxiliary speaker representations, i.e. i-vector, bottleneck speaker vector and speaking rate. The methods to obtain these speaker representations are detailed as follows.

3.1. I-vector

I-vector is a popular technique for speaker verification and recognition [15]. It can capture the most important information of a speaker in a low-dimensional representation. Furthermore it has been shown that i-vectors can be used in speaker-aware training on DNNs for speech recognition [13, 19].

3.2. Bottleneck speaker vector

Bottleneck speaker vector, first proposed in [7], has obtained competitive performance to the i-vector based adaptation on DNNs. Figure 2(a) is the standard network for extracting BSVs. It is a 3-layer DNN trained to classify speakers. A bottleneck layer is set at the third layer to extract the bottleneck vectors. To obtain the BSV, we averaged all the bottleneck vectors from each speaker and then normalise its L2-norm to be 1.

Motivated by the work in [20] which uses an ASR-DNN system to integrate the speech content and extract a more powerful i-vector, two more structures are proposed to incorporate BSV extraction with phone knowledges.

- Figure 2(b) shows the phone-aware training (PAT) used in this paper. The posterior of a mono-phone based DNN is fed as input for the speaker recognition network. Different from the traditional speaker-aware training, in this work, these two networks are trained jointly. The criterion used for model optimization is

$$E(\theta) = \lambda E_{\text{mono}}(\theta) + E_{\text{spk}}(\theta) \quad (1)$$

where θ denotes model parameters, $E_{\text{mono}}(\theta)$, $E_{\text{spk}}(\theta)$ are the cross-entropy for mono-phone and speaker based DNNs respectively. λ is a *mixing factor*.

- Another structure is shown in figure 2(c), where multi-task joint-training is used for training mono-phone and speaker DNNs jointly. Multi-task joint-training framework, that more than one criteria is used in model training, can yield better performance than normal cross entropy training. Such as in [21], triphones and trigraphemes are used to train acoustic modeling. Phone and speaker combined training [22] is another example. The criterion in this framework is the same as Eq. (1).
- Note that the criteria used in previous two methods are the same, however, the model architectures are different. In PAT, phone information acts as an auxiliary information for the network classify speaker better. However, in multi-task training, phone information is acting as a regularizer, which will make the network more robust.

3.3. Speaking rate

It has been known that the speaking rate impacts the accuracy of automatic speech recognition (ASR). A relatively low or high speaking rate usually increases the word error rate [23, 24]. The speaking rate may influence the speech signal in two ways: on the temporal

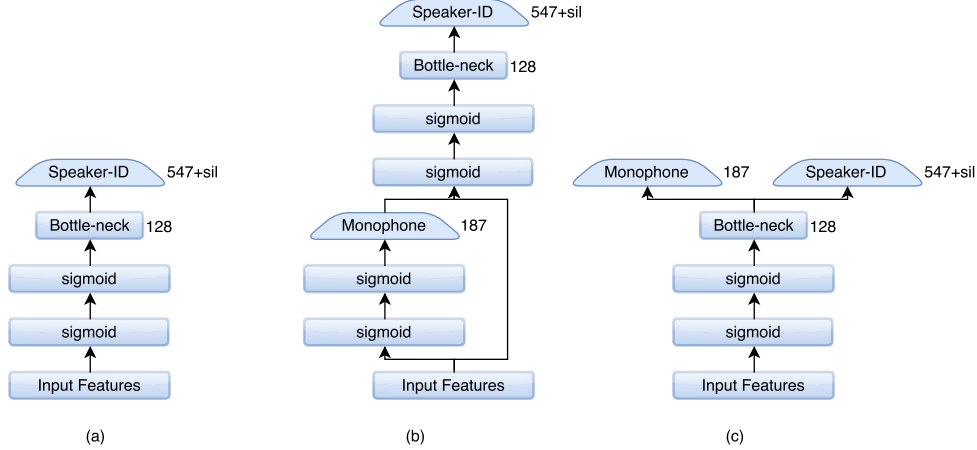


Fig. 2. Structures for extracting bottleneck speaker vectors

properties or on the spectrum distortion. In previous works, frame rate normalization [25] was applied on GMM-HMM and speaking rate assisted training was tried on DNN-HMM [26]. In this work, speaking rate was implemented for speaker-aware training on LSTM-RNNs at the first time. Considering that LSTM-RNNs are combination of a dynamic system and a neural network, it may compensate for the speaking rate in both the temporal properties and the spectrum distortion together. The speaking rate was extracted as follows:

- An alignment is generated with a baseline DNN system.
- The speaking rate is then calculated by dividing the number of phones by the duration of all phones.

$$f(i) = N / \sum_{j=1}^N t_{i,j}$$

where $f(i)$ is the speaking rate for utterance i , $t_{i,j}$ is the duration of phone j , and there are N phones in utterance i .

4. EXPERIMENT

Our experiments are conducted on AMI corpus, which contains around 100 hours of meetings recorded in specifically equipped instrumented meeting rooms at three sites in Europe (Edinburgh, IDIAP, TNO) [27]. Acoustic signal is captured and synchronized by multiple microphones including individual head microphones (IHM, close-talk), lapel microphones, and one or more microphone arrays. In this work, the IHM (close-talk) data are used. Our experiments adopted the suggested AMI corpus partition that contains about 80 hours and 8 hours in training and evaluation sets respectively [28]. For fast turnarounds, we also select 10000 utterances from the training set and create a lighter set-up, the training procedures and test sets are identical in the sub- and full-set experiments. For decoding, the 50K-word AMI dictionary and a trigram language model interpolated from the one created using the AMI training transcripts and the Fisher English corpus were used.

4.1. Baseline set up

The GMM-HMM system was built using the standard Kaldi AMI recipe [29]. 39-dimensional MFCC (plus deltas and double deltas) features with CMVN was used to train the initial ML model. Then 7

frames of MFCCs were spliced and projected to 40 dimensions with linear discriminant analysis (LDA). A maximum likelihood linear transform (MLLT) was estimated on the LDA features to generate the LDA+MLLT model. After that, speaker adaptive training was performed with one FMLLR transform per speaker. The fmlr features were then used for training DNN, LSTM-RNN, i-vector and BSV.

CNTK [30] was used to train the DNN and LSTM-RNN. The DNN has 6 hidden layers each of which contains 2048 neurons. The input feature for DNNs contains 11 frames (5 frames on each side of the center frame). Cross-entropy (CE) was used. The learning rate started from 0.1 and changed to 1 for the second epoch, then it was decayed by a factor of 0.5 when the cross entropy on a cross-validation set between two consecutive epochs increases.

The LSTM-RNN has 3 projected LSTM layers which are followed by the soft-max layer. Each LSTM layer has 1024 memory cells and 512 output units in projection. Input to the LSTM-RNN is a single acoustic frame with 5 frames shift, the truncated version of BPTT was used for training LSTM-RNN. [5] 40 utterances were processed in parallel and the BPTT truncation size was set to 20. To ensure training stability, the gradient was clipped to the range of [-1, 1] in parameters update.

Data set	Model	WER
Sub Set	DNN	34.9
	LSTM-RNN	32.4
Full Set	DNN	26.5
	LSTM-RNN	26.0

Table 1. WER (%) of DNN and LSTM-RNN on AMI IHM condition

Table 4.1 shows the results of baseline DNNs and LSTM-RNNs trained on the sub- and full-set. On both sets, the LSTM-RNN performs better than the DNN, demonstrating its advantage in acoustic modeling. However, the improvement becomes smaller when more data was used, which might because the fmlr is not very suitable for LSTM-RNNs. The transformation may break the inherent temporal dependency between neighbor frames and this also indicates that other speaker adaptation methods exploration is necessary for LSTM-RNNs.

4.2. Experiment on different structure for LSTM adaptation

Since the efficacy of i-vector for adaptation has been shown on DNNs, we firstly choose i-vector as the speaker vector for speaker-aware training on LSTM-RNNs. In this work, a 128 dimensional i-vector was extracted for each speaker by using GMM (the universal background model) with 2048 components, and then the i-vector was length normalized to one. Three different structures presented in section 2 were investigated to do speaker adaptation on LSTM-RNN.

Model	Struct.	WER
LSTM-RNN	—	32.4
LSTM+DNN-RNN	—	35.1
+i-vector	(a)	31.1
	(b)	33.2
	(c)	34.5

Table 2. WER (%) of different structures for LSTM-RNN adaptation on AMI IHM Subset

In Table 2, (a),(b),(c) are structures shown in Figure 1. These structures did not perform according to our expectation. The straightforward structure gained the best performance, outperforming all other structures. This is partly because the additional sigmoid function introduced into LSTM-RNNs makes the network hard to train well. We verified our conjecture using an LSTM-RNN followed by a 1-hidden-layer DNN. This structure performed worse than normal LSTM-RNNs. However, this conflicts with results in [31]. The key reason might be that in [31], the DNN part and the LSTM part are trained with different learning rates. We will investigate advanced learning rate strategy in the future work.

4.3. Experiment on different BSV for speaker adaptation

In this section, we compare the WER of the DNN and LSTM-RNN adaptation using different bottleneck speaker vectors obtained by standard training (ST), phone aware training (PAT) and multi-task training (MT). The mixing factor λ for BSV-MT and BSV-PAT were set to 1.0 in our experiments.

Model	Struct.	WER
DNN	—	34.9
+i-vector	—	34.4
+BSV	ST	34.6
	PAT	34.7
	MT	34.3
LSTM-RNN	—	32.4
+i-vector	—	31.1
+BSV	ST	31.9
	PAT	31.8
	MT	31.7

Table 3. WER (%) of different BSV networks trained on AMI IHM Subset

From Table 3, we can observe that using BSV feature can lead to improvements on both DNN and LSTM-RNN systems. BSV with multi-task training is slightly better than standard BSV on both DNNs and LSTM-RNNs. On DNNs, it was even slightly better than using i-vector. However, phone-aware training is not helpful for BSV extraction on both DNN and LSTM models.

4.4. Experiment on different speaker auxiliary vectors

In this subsection we compare different speaker vectors. As shown in Table 4, all auxiliary vectors achieved 1.5%-4% relative improvement over the baseline LSTM-RNN model. Moreover these auxiliary vectors are complementary. The best performance is obtained using the i-vector+BSV-MT+spk-rate setup. The WER is improved from 32.4% to 30.6% (relative 6%).

Feature	WER
FMLLR	32.4
+i-vector	31.1
+BSV-MT	31.7
+spk-rate	31.9
+i-vector+BSV-MT	30.9
+i-vector+spk-rate	30.9
+i-vector+BSV-MT+spk-rate	30.6

Table 4. WER (%) of different auxiliary information on AMI IHM Subset

4.5. Experiment on full set

Finally speaker-aware training on LSTM-RNNs is evaluated on the AMI IHM full corpus. For adaptation on LSTM-RNNs, we adopt the best configurations discovered in the previous section, i.e., BSV-MT and combining all auxiliary vectors.

Feature	WER
fmlr	26.0
+i-vector	24.3
+BSV-MT	25.0
+spk-rate	25.7
+i-vector+BSV-MT+spk-rate	24.3

Table 5. WER (%) of different auxiliary information on AMI IHM Full set

From Table 5, we can see that for each auxiliary vector, the adapted LSTM-RNN performed consistently better than the baseline LSTM-RNN. However, when using all speaker representative vectors together, no further improvement can be obtained. The reason may be that the amount of speakers on full set is large enough to capture all variability, and on full set the LSTM-RNN model can encode richer speaker variability by itself and thus the efficacy of speaker adaptation decreases. Even that the speaker-aware trained LSTM-RNN still outperforms the baseline LSTM-RNN using fmlr feature with 6.5% relative improvement.

5. CONCLUSION

In this paper, the speaker-aware training was studied on LSTM-RNNs. A simple but effective structure was found for LSTM-RNNs adaptation and two structures are proposed to incorporate phone information into BSV. In addition different auxiliary vectors were explored to test the efficiency for LSTM-RNN adaptation. Our experiments with the AMI corpus show that adaptation with i-vector significantly improves the performance of LSTM-RNN models and the improvement also can be obtained by using BSV and speaking rate. For the future work, we will focus on joint training of the speaker information extractor and the speech recognition model.

6. REFERENCES

- [1] Herve A Bourlard and Nelson Morgan, *Connectionist speech recognition: a hybrid approach*, vol. 247, Springer Science & Business Media, 2012.
- [2] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdelrahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al., “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *Signal Processing Magazine, IEEE*, vol. 29, no. 6, pp. 82–97, 2012.
- [3] Yoshua Bengio, Patrice Simard, and Paolo Frasconi, “Learning long-term dependencies with gradient descent is difficult,” *Neural Networks, IEEE Transactions on*, vol. 5, no. 2, pp. 157–166, 1994.
- [4] Sepp Hochreiter and Jürgen Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [5] Hasim Sak, Andrew Senior, and Françoise Beaufays, “Long short-term memory recurrent neural network architectures for large scale acoustic modeling,” in *INTERSPEECH*, 2014.
- [6] Xiangang Li and Xihong Wu, “Constructing long short-term memory based deep recurrent neural networks for large vocabulary speech recognition,” in *ICASSP. IEEE*, 2015, pp. 4520–4524.
- [7] Yan Huang, Dong Yu, Chaojun Liu, and Yifan Gong, “A comparative analytic study on the gaussian mixture and context dependent deep neural network hidden markov models,” in *INTERSPEECH*, 2014.
- [8] Yu Dong and Deng Li, *Automatic speech recognition, A deep learning approach*, Springer.
- [9] Frank Seide, Gang Li, Xie Chen, and Dong Yu, “Feature engineering in context-dependent deep neural networks for conversational speech transcription,” in *ASRU. IEEE*, 2011, pp. 24–29.
- [10] Kaisheng Yao, Dong Yu, Frank Seide, Hang Su, Li Deng, and Yifan Gong, “Adaptation of context-dependent deep neural networks for automatic speech recognition,” in *SLT*, 2012, pp. 366–369.
- [11] Dong Yu, Kaisheng Yao, Hang Su, Gang Li, and Frank Seide, “KL-divergence regularized deep neural network adaptation for improved large vocabulary speech recognition,” in *ICASSP. IEEE*, 2013, pp. 7893–7897.
- [12] Ossama Abdel-Hamid and Hui Jiang, “Fast speaker adaptation of hybrid nn/hmm model for speech recognition based on discriminative learning of speaker code,” in *ICASSP. IEEE*, 2013, pp. 7942–7946.
- [13] George Saon, Hagen Soltau, David Nahamoo, and Michael Picheny, “Speaker adaptation of neural network acoustic models using i-vectors,” in *ASRU. IEEE*, 2013, pp. 55–59.
- [14] Yajie Miao and Florian Metze, “On speaker adaptation of long short-term memory recurrent neural networks,” in *INTERSPEECH*, 2015.
- [15] Najim Dehak, Patrick Kenny, Réda Dehak, Pierre Dumouchel, and Pierre Ouellet, “Front-end factor analysis for speaker verification,” *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 19, no. 4, pp. 788–798, 2011.
- [16] Hengguan Huang and Khe Chai Sim, “An investigation of augmenting speaker representations to improve speaker normalisation for dnn-based speech recognition,” in *ICASSP. IEEE*, 2015, pp. 4610–4613.
- [17] X Chen, T Tan, Xunying Liu, Pierre Lanchantin, M Wan, Mark JF Gales, and Philip C Woodland, “Recurrent neural network language model adaptation for multi-genre broadcast speech recognition,” in *INTERSPEECH*, 2015.
- [18] Tomas Mikolov and Geoffrey Zweig, “Context dependent recurrent neural network language model,” in *SLT*, 2012, pp. 234–239.
- [19] Yajie Miao, Hao Zhang, and Florian Metze, “Towards speaker adaptive training of deep neural network acoustic models,” in *INTERSPEECH*, 2014.
- [20] Yun Lei, Nicolas Scheffer, Luciana Ferrer, and Moray McLaren, “A novel scheme for speaker recognition using a phonetically-aware deep neural network,” in *ICASSP. IEEE*, 2014, pp. 1695–1699.
- [21] Dongpeng Chen, Brian Mak, Cheung-Chi Leung, and Sunil Sivasadas, “Joint acoustic modeling of triphones and tri-grammes by multi-task learning deep neural networks for low-resource speech recognition,” in *ICASSP. IEEE*, 2014, pp. 5592–5596.
- [22] Nanxin Chen, Yanmin Qian, and Kai Yu, “Multi-task learning for text-dependent speaker verification,” in *INTERSPEECH*, 2015.
- [23] Matthew Siegler, Richard M Stern, et al., “On the effects of speech rate in large vocabulary speech recognition systems,” in *ICASSP. IEEE*, 1995, vol. 1, pp. 612–615.
- [24] Nelson Morgan, Eric Fosler-Lussier, and Nikki Mirghafori, “Speech recognition using on-line estimation of speaking rate,” in *Eurospeech*, 1997, vol. 97, pp. 2079–2082.
- [25] Stephen M Chu and Daniel Povey, “Speaking rate adaptation using continuous frame rate normalization,” in *ICASSP. IEEE*, 2010, pp. 4306–4309.
- [26] Zeng Xiangyu, Yin Shi, and Wang Dong, “Learning speech rate in speech recognition,” in *INTERSPEECH*, 2015, pp. 528–532.
- [27] Thomas Hain, Luká Burget, John Dines, Philip N Garner, František Grézl, Asmaa El Hannani, Marijn Huijbregts, Martin Karafiat, Mike Lincoln, and Vincent Wan, “Transcribing meetings with the amida systems,” *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 20, no. 2, pp. 486–498, 2012.
- [28] Pawel Swietojanski, Arnab Ghoshal, and Steve Renals, “Hybrid acoustic models for distant and multichannel large vocabulary speech recognition,” in *ASRU. IEEE*, 2013, pp. 285–290.
- [29] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukáš Burget, Ondřej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlíček, Yanmin Qian, Petr Schwarz, et al., “The kaldı speech recognition toolkit,” 2011.
- [30] Dong Yu, Adam Eversole, Mike Seltzer, Kaisheng Yao, Zhiheng Huang, Brian Guenter, Oleksii Kuchaiev, Yu Zhang, Frank Seide, Huaming Wang, et al., “An introduction to computational networks and the computational network toolkit,” Tech. Rep., Microsoft Research, 2014.
- [31] Tara N Sainath, Oriol Vinyals, Andrew Senior, and Hasim Sak, “Convolutional, long short-term memory, fully connected deep neural networks,” in *ICASSP. IEEE*, 2015.