

IMPROVING NOISE ROBUSTNESS OF AUTOMATIC SPEECH RECOGNITION VIA PARALLEL DATA AND TEACHER-STUDENT LEARNING

Ladislav Mošner^{1*}, Minhua Wu², Anirudh Raju², Sree Hari Krishnan Parthasarathi², Kenichi Kumatani², Shiva Sundaram², Roland Maas², Björn Hoffmeister²

¹Brno University of Technology, Speech@FIT and IT4I Center of Excellence, Czechia
²Amazon.com, Inc. USA

imosner@fit.vutbr.cz, wuminhua@amazon.com

ABSTRACT

For real-world speech recognition applications, noise robustness is still a challenge. In this work, we adopt the teacher-student (T/S) learning technique using a parallel clean and noisy corpus for improving automatic speech recognition (ASR) performance under multimedia noise. On top of that, we apply a logits selection method which only preserves the k highest values to prevent wrong emphasis of knowledge from the teacher and to reduce bandwidth needed for transferring data. We incorporate up to 8000 hours of untranscribed data for training and present our results on sequence trained models apart from cross entropy trained ones. The best sequence trained student model yields relative word error rate (WER) reductions of approximately 10.1%, 28.7% and 19.6% on our clean, simulated noisy and real test sets respectively comparing to a sequence trained teacher.

Index Terms— automatic speech recognition, noise robustness, teacher-student training, domain adaptation

1. INTRODUCTION

With the exponential growth of big data and computing power, automatic speech recognition (ASR) technology has been successfully used in many applications. People can do voice search using mobile devices. They can also interact with smart home devices such as Amazon Echo or Google Home through distant speech recognition [1] [2] for entertainment, shopping or other personal assistance. For such real-world applications, noise robustness is important since the device needs to work well under various acoustic environments, and it still remains to be a challenging task [3, 4]. Although large-vocabulary speech recognition is of high accuracy by applying deep neural networks [5–7], it requires thousands of hours of transcribed data which is time-consuming and expensive to collect, and its performance under noisy environment may still suffer. Considerable efforts have been made to improve noise robustness by applying algorithms in the front-end feature domain [8–14] or in the back-end model [15–20]. Another natural way to deal with noise in the acoustic environment is to use multi-style training [21], which trains the acoustic model with noisy speech data. All these approaches require *supervision* where the speech data is manually transcribed.

In order to improve noise robustness of the distant speech recognition system in an *unsupervised* mode, it is desir-

able to scale multi-style training with an even larger training dataset. However, acquiring manually transcribed data for noisy speech can be slow and expensive. In this work, we explore the technique of teacher-student (T/S) learning using a parallel corpus of clean and noisy data. We focus on improving the ASR performance under multimedia noise which is commonly present at home.

T/S learning was at first explored in speech community [22] [23] to distill the knowledge from bigger models to a smaller one, and was successfully applied in the areas of ASR [24, 25] and keyword spotting [26] afterwards. Instead of knowledge distillation, we adopt the T/S learning for domain adaptation which was proposed in [27] to build an ASR system performing more robustly under multimedia noise. On top of this system, we apply logits selection keeping only the k highest values and experiment it with multiple settings of temperature T . This method was proposed in [28] to optimize storage and to parallelize the target generation for teacher-student training, and we find that it even helps improve performance of the adapted student model since it prevents over-emphasizing wrong senones by the teacher. Because the T/S learning technique applied in this work does not require transcribed data, we also explore how much the system performance can be further improved by gradually incorporating more training recordings. Finally, we study the effects of doing sequence training on top of the T/S learning for domain adaptation, which was not reported in [27].

2. METHOD

2.1. Teacher-student training for domain adaptation

In order to apply T/S training for domain adaptation (to adapt from clean teacher to noisy student domain in our use case), parallel clean and noisy corpus is needed and it could be generated by artificially adding noise on top of the clean data, which will be described in detail in section 3.2. As shown in Figure 1, a teacher network outputs discrete probability distribution $P_{teacher}(S|x_{clean})$ over senones S given a clean feature vector x_{clean} , and a student network estimates probabilities of senones $P_{student}(S|x_{noisy})$ given a noisy feature vector x_{noisy} . The objective is to minimize the Kullback-Leibler (KL) divergence between these two distributions in order to make behavior of the student network in the target domain approach that of the teacher network in the source domain.

Since the teacher parameters are fixed when training the student model, minimizing the KL divergence is equivalent to minimizing the cross entropy between output distributions from the teacher and the student, where the teacher’s output

* Ladislav Mošner performed the work while he was a research intern at Amazon.

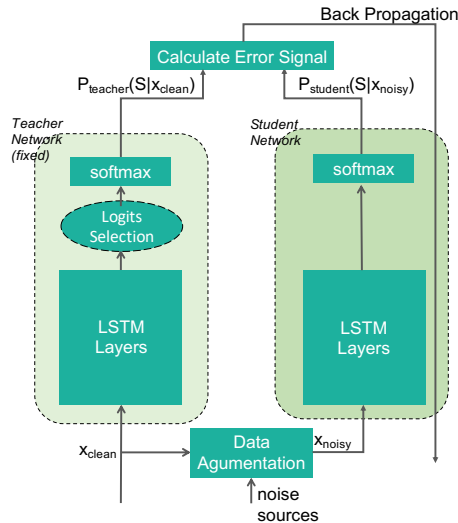


Fig. 1: The flow chart of teacher-student learning using parallel data for improving noise robustness of ASR.

distribution (“soft targets”) are considered as the ground truth. Therefore, student training relies completely on teacher outputs and no transcriptions are required.

2.2. Logits selection

The discrete probability distribution over senones at each time frame from the teacher will guide the learning of the student model. Since there are usually thousands of senones for a typical ASR system, a large number of values need to be loaded when processing just one frame. Moreover, the majority of the output probability mass is usually covered by just a few outputs, while the rest of output values are very small and noisy, which might confuse the student model instead. In this case, we adopt the logits selection method [28] with which only k highest logit values are preserved so that only the most reliable information from the teacher is preserved to train the student. At the same time, this method also reduces the bandwidth of transferring frames of soft targets for training. Let z_i be an activation before the softmax (logit), N be number of classified senones. Without any selection, the output probability of senone s_i is computed as follows:

$$P_{teacher}(S = s_i | x_{clean}) = q_i = \frac{\exp(z_i/T)}{\sum_{j=1}^N \exp(z_j/T)}, \quad (1)$$

where T is the temperature controlling softness of the distribution [23].

When considering only k highest logit values whose indices lie in the set \mathbb{K} , the k highest probabilities of senones are preserved with an emphasizing factor \mathbf{A} while the rest are suppressed to be zero.

$$q'_i = \begin{cases} \frac{\exp(z_i/T)}{\sum_{j \in \mathbb{K}} \exp(z_j/T)} = \mathbf{A}q_i & \text{if } i \in \mathbb{K}, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

$$\mathbf{A} = \frac{\sum_{j=1}^N \exp(z_j/T)}{\sum_{j \in \mathbb{K}} \exp(z_j/T)} \quad (3)$$

With this selection method, we are able to dramatically reduce storage space needed for soft targets and also I/O operations during training since k is significantly small compared with N . In fact, we notice that this method even helps to improve performance of the student model since it is boosting confidence of the teacher model and suppressing the confusing part as indicated in Equations (2) and (3).

3. EXPERIMENTAL SETUP

3.1. Model architecture

Both teacher and student have the same architecture for all the experiments. The neural network consists of three LSTM [29] layers each of which comprises 512 units. The last LSTM layer is followed by a fully connected layer outputting probability distribution over 3,010 senones. 64-dimensional feature vectors of Log Mel-Filter-Bank Energies (LFBE) are used as inputs to the network. There is no frame stacking, and the output HMM state label is delayed by 3 frames. We run decoding in a single-pass framework using a smoothed 4-gram language model (LM) trained on both internal and external data sources. The total text used to train the LM is over billion words.¹

3.2. Parallel training datasets

The original clean corpus consists of approximately 8,000 hours of beamformed speech recordings. However, only 800 hours are transcribed and thus are used to build the teacher model with supervised training. We perform data simulation [30] to obtain the parallel noisy training data. Since our primary aim is to deal with multimedia noise in room conditions, we collected a corpus consisting especially of music samples and acoustic video content. For every utterance, one to three additional noises are randomly selected. The notion of sound propagation in enclosures is obtained by means of room simulation. The image method is used to acquire artificial room impulse responses [31, 32]. Reverberation times are uniformly drawn from the interval of (500, 900) ms. The mixture of noises is combined with the clean signal at signal-to-noise ratio (SNR) ranging from 0 to 30 dB.

3.3. Test datasets

In order to evaluate the performance of our acoustic models, we used three test datasets.

- “clean”: It is of similar domain to that of the clean training data ($\sim 41k$ utterances).
- “noisy”: It is derived from the “clean” dataset by the same simulation method described in section 3.2, but different multimedia noise sources are selected and reverberation times are drawn from the different interval of (520, 920) ms.
- “realistic”: It is collected in a real room ($\sim 2k$ utterances). Clean speech and multimedia noises are played by loudspeakers and recorded by multiple microphones varying in their positions.

¹Our experimental ASR system does not reflect the performance of the production Alexa system.

4. RESULTS

4.1. Multi-condition versus teacher student training

A teacher acoustic model is trained using the transcribed 800-hour clean corpus and treated as the baseline. The conventional multi-condition [21] trained acoustic model is also built using the transcribed portion of the noisy training dataset. As displayed in Table 1, the multi-condition trained acoustic model outperforms the teacher by a significant margin on the noisy test set. At the same time, its performance on the clean test set is comparable to the performance of the baseline.

When training the student acoustic model, we at first initialize it using the baseline teacher model. During the training process, the cross entropy between $P_{student}(S|x_{noisy})$ and $P_{teacher}(S|x_{clean})$ are being minimized.

When the temperature is set to $T = 1$ (standard softmax), we observe improvements in performance over the multi-condition trained model on all the datasets. Matching the student output probability distribution with that of the teacher enhances generalization ability of the student and we see an average relative reduction of about 1.9% in WER even on the clean test dataset. The increased temperature does not seem to help achieve further improvements but leads to worse performance. A temperature of $T = 5$ even results in higher WER in comparison with the multi-condition trained model, which indicates that a flatter output distribution from the teacher may result in confusions for the student to learn effectively.

Table 1: Comparison of multi-condition and teacher-student training with different temperatures (no logits selection). Results are reported in relative word error rate reduction (WERR) [%]. Minus sign indicates improvement.

Acoustic model	Clean test set	Noisy test set	Realistic test set
Baseline/teacher (clean 800h)	0.00	0.00	0.00
Multi-condition (noisy 800h)	0.69	-15.20	-4.26
Student, $T = 1$ (parallel 800h)	-1.93	-17.20	-6.82
Student, $T = 2$ (parallel 800h)	-0.46	-16.43	-6.35
Student, $T = 5$ (parallel 800h)	1.16	-13.50	-2.20

4.2. Number of candidate logits and temperature

To prevent needless emphasis of wrong senones from the teacher and reduce the bandwidth needed for transferring soft targets, we explore the logits selection approach explained in section 2.2. In our experiments, we slightly changed the computation of output probabilities previously defined by Equations (2) and (3). This modification is performed for the convenience of training but does not break the general idea. Instead of assigning zero probability to the non- k -best senones, we assign a sufficiently high negative constant C to the corresponding logits. The output probability now becomes

$$\tilde{q}_i = \begin{cases} \frac{\exp(z_i/T)}{(N-k)\exp(C/T) + \sum_{j \in \mathbb{K}} \exp(z_j/T)} = \tilde{\mathbf{A}}q_i & \text{if } i \in \mathbb{K}, \\ \frac{\exp(C/T)}{(N-k)\exp(C/T) + \sum_{j \in \mathbb{K}} \exp(z_j/T)} \approx 0 & \text{otherwise.} \end{cases} \quad (4)$$

$$\tilde{\mathbf{A}} = \frac{\sum_{j=1}^N \exp(z_j/T)}{(N-k)\exp(C/T) + \sum_{j \in \mathbb{K}} \exp(z_j/T)} \quad (5)$$

We explore the resulting recognition accuracy for multiple settings of temperature T and candidate logits selection k , which are summarized in Figure 2.

4.2.1. Temperature $T = 1$

While using a temperature of 1, no significant differences in WER can be seen even for a very aggressive logits selection (5 senones out of 3010). Inspecting the average output distribution after application of softmax, we find out that the highest probability is close to one and the rest close to zero. Therefore, a sum of probabilities of senones that do not belong to the 5-best is still small. Redistribution of this mass among 5-best senones then does not affect the distribution much.

4.2.2. Temperature $T = 2$

Interestingly, the combination of temperature 2 with logits selection bring accuracy improvements. We observe that taking 5-best values into consideration is not sufficient. However, the difference between results obtained using 20-best and 40-best logits is minimal.

4.2.3. Temperature $T = 5$

When the output distribution gets flatter (temperature 5), the difference among highest probabilities diminishes. Then the effect of multiplication by constant is similar for all k senones and it becomes more difficult to estimate the correct one. The fewer candidates are taken into account, the more severe degradation occurs.

Based on the analysis, the most promising hyperparameters for our student models are: temperature $T = 2.0$ and $k = 20$ for logits selection. These values are fixed for the following experiments.

4.3. Size of the training dataset

Building both the baseline and the multi-condition system require transcribed data. Compared to those methods, the major advantage of the teacher-student training approach is that it does not require transcripts once the teacher model is trained. Therefore, we could explore how much the WER can be further reduced by incorporating even larger number of utterances in the training set to build the student model. The student training dataset is gradually increased up to ten times more audio compared with the original amount. The relative word error rate reduction (WERR) is displayed in Figure 3 as a function of training data size. Similar trends are observed for all the test datasets. As expected, the accuracy improves with the increasing amount of data. However, the minimum WER is achieved when using approximately 4800 hours ($6 \times$). It could be possible that we only have about 4800 hours of unique noise resources and the model overfits after adding repeated noise examples. This hypothesis, however, requires further investigation. Alternatively, the fact that the teacher itself is erroneous could also affect dependence of accuracy on training data size.




Fig. 2: Relative word error rate reduction (WERR) [%] (with respect to the cross-entropy trained teacher) from different student models when tuning the distillation temperature and logits selection. *max kbest* means preserving all logits.




Fig. 3: Relative word error rate reduction (WERR) [%] (with respect to the cross-entropy trained teacher) from different student models incorporating more untranscribed training data. “1x” means 1×800 hours of training data. Fix $T = 2, k = 20$ for building all the models.

Table 2: Effects of applying sequence training on both teacher and student models. Results are expressed in terms of relative word error rate reduction (WERR) [%] (with respect to the cross-entropy trained teacher). Minus sign indicates improvement. Relative WERR inside the bracket is computed with respect to the sequence trained teacher. $T = 2, k = 20$ for building all the student models.

Acoustic model	Training objective	Training data	Clean test set	Noisy test set	Realistic test set
Teacher	xent	800h clean	0.00	0.00	0.00
Student	xent	6×800 h parallel	-9.50	-26.08	-19.57
Teacher	xent, sMBR	800h clean, 800h clean	-5.79 (0.00)	-0.98 (0.00)	-3.68 (0.00)
Student	xent	6×800 h parallel	-11.89 (-6.48)	-26.80 (-26.08)	-21.63 (-18.63)
Student	xent, sMBR	6×800 h parallel, 800h noisy	-15.29 (-10.08)	-29.36 (-28.67)	-22.54 (-19.58)

4.4. Sequence training

Sequence training has been shown to be effective in improving ASR performance in general [33]. Wong and Gales also investigated the usefulness of the combination of sequence and teacher-student training [34]. In our experiment, we at first fine-tune the original cross-entropy trained teacher using state-level minimum Bayes risk (sMBR) criterion [35]. As displayed in Table 2, the new teacher outperforms the original one on all our test datasets. We then train a new student network on top of the new sequence-trained teacher using the parallel corpus. We used size $6 \times$ of the parallel datasets to train the student as it was shown to be the best option in the previous experiment. The new cross-entropy trained student network is able to make use of the improved teacher, since its performance is better than that of the student taught by the weaker teacher. Finally, the new student is further optimized by means of sMBR training using only the transcribed portion of noisy data ($1 \times$).

5. CONCLUSION

In this paper, we explore the teacher-student learning approach using parallel clean and noisy corpus to improve speech recognition performance under multimedia noise. We gradually optimize this system by applying logits selection and incorporating larger amount of untranscribed training data. With a temperature of $T = 2$ and logits selection of $k = 20$ highest values, we obtain the best student model using a parallel clean and noisy corpus which is about 6 times more of the original clean training data. By applying standard sequence training on both the teacher and student model, the final student brings relative WER reductions of about 10.1%, 28.7% and 19.6% on the clean, simulated noisy and real test sets, respectively.

In the future, a larger corpus of noise sources will be collected to prevent the multimedia samples from being repeated. We will attempt to scale up architecture and training datasets. Soft targets selection based on teacher’s certainty could be explored as well.

6. REFERENCES

- [1] Kenichi Kumatani, John McDonough, and Bhiksha Raj, "Microphone array processing for distant speech recognition: From close-talking microphones to far-field sensors," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 127–140, 2012.
- [2] Minhua Wu, Kenichi Kumatani, Shiva Sundaram, Nikko Ström, and Björn Hoffmeister, "Frequency domain multi-channel acoustic modeling for distant speech recognition," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2019, IEEE.
- [3] Jinyu Li, Li Deng, Yifan Gong, and Reinhold Haeb-Umbach, "An overview of noise-robust automatic speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 4, pp. 745–777, 2014.
- [4] Brian King, I-Fan Chen, Yonatan Vaizman, Yuzong Liu, Roland Maas, Sree Hari Krishnan Parthasarathi, and Björn Hoffmeister, "Robust speech recognition via anchor word representations," *INTERSPEECH-2017*, pp. 2471–2475, 2017.
- [5] Tara N Sainath, Brian Kingsbury, Bhuvana Ramabhadran, Petr Fousek, Petr Novak, and Abdel-rahman Mohamed, "Making deep belief networks effective for large vocabulary continuous speech recognition," in *Automatic Speech Recognition and Understanding (ASRU), 2011 IEEE Workshop on*. IEEE, 2011, pp. 30–35.
- [6] George E Dahl, Dong Yu, Li Deng, and Alex Acero, "Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition," *IEEE Transactions on audio, speech, and language processing*, vol. 20, no. 1, pp. 30–42, 2012.
- [7] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman 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," *IEEE Signal processing magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [8] Steven Boll, "Suppression of acoustic noise in speech using spectral subtraction," *IEEE Transactions on acoustics, speech, and signal processing*, vol. 27, no. 2, pp. 113–120, 1979.
- [9] Bishnu S Atal, "Effectiveness of linear prediction characteristics of the speech wave for automatic speaker identification and verification," *the Journal of the Acoustical Society of America*, vol. 55, no. 6, pp. 1304–1312, 1974.
- [10] Alejandro Acero, "Acoustical and environmental robustness in automatic speech recognition," in *Proc. of ICASSP*, 1990.
- [11] Hynek Hermansky and Nelson Morgan, "Rasta processing of speech," *IEEE transactions on speech and audio processing*, vol. 2, no. 4, pp. 578–589, 1994.
- [12] Hynek Hermansky, "Perceptual linear predictive (PLP) analysis of speech," *the Journal of the Acoustical Society of America*, vol. 87, no. 4, pp. 1738–1752, 1990.
- [13] Duncan Macho, Laurent Mauuary, Bernhard Noé, Yan Ming Cheng, Douglas Ealey, Denis Jovet, Holly Kelleher, David Pearce, and Fabien Saadoun, "Evaluation of a noise-robust dsr front-end on aurora databases," in *INTERSPEECH*, 2002.
- [14] Xiaodong Cui, Markus Iseli, Qifeng Zhu, and Abeer Alwan, "Evaluation of noise robust features on the Aurora databases," in *Seventh International Conference on Spoken Language Processing*, 2002.
- [15] Mark JF Gales and Steve J Young, "Robust continuous speech recognition using parallel model combination," *IEEE Transactions on Speech and Audio Processing*, vol. 4, no. 5, pp. 352–359, 1996.
- [16] Pedro J Moreno, "Speech recognition in noisy environments," 1996.
- [17] Christopher J Leggetter and Philip C Woodland, "Maximum likelihood linear regression for speaker adaptation of continuous density hidden markov models," *Computer speech & language*, vol. 9, no. 2, pp. 171–185, 1995.
- [18] Jian Wu and Qiang Huo, "Supervised adaptation of MCE-trained CDHMMs using minimum classification error linear regression," in *Acoustics, Speech, and Signal Processing (ICASSP), 2002 IEEE International Conference on*. IEEE, 2002, vol. 1, pp. 1–605.
- [19] Xiaodong He and Wu Chou, "Minimum classification error linear regression for acoustic model adaptation of continuous density HMMs," in *Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03). 2003 IEEE International Conference on*. IEEE, 2003, vol. 1, pp. 1–1.
- [20] Kai Yu, Mark Gales, and Philip C Woodland, "Unsupervised adaptation with discriminative mapping transforms," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 17, no. 4, pp. 714–723, 2009.
- [21] R Lippmann, Edward Martin, and D Paul, "Multi-style training for robust isolated-word speech recognition," in *Acoustics, Speech, and Signal Processing, IEEE International Conference on ICASSP'87*. IEEE, 1987, vol. 12, pp. 705–708.
- [22] Jinyu Li, Rui Zhao, Jui-Ting Huang, and Yifan Gong, "Learning small-size dnn with output-distribution-based criteria," in *Fifteenth annual conference of the international speech communication association*, 2014.
- [23] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [24] Yevgen Chebotar and Austin Waters, "Distilling knowledge from ensembles of neural networks for speech recognition.," in *Interspeech*, 2016, pp. 3439–3443.
- [25] Liang Lu, Michelle Guo, and Steve Renals, "Knowledge distillation for small-footprint highway networks," in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 4820–4824.
- [26] George Tucker, Minhua Wu, Ming Sun, Sankaran Panchapagesan, Gengshen Fu, and Shiv Vitaladevuni, "Model compression applied to small-footprint keyword spotting.," in *INTERSPEECH*, 2016, pp. 1878–1882.
- [27] Jinyu Li, Michael L. Seltzer, Xi Wang, Rui Zhao, and Yifan Gong, "Large-scale domain adaptation via teacher-student learning," in *INTERSPEECH*, 2017.
- [28] Sree Hari Krishnan Parthasarathi and Nikko Ström, "Lessons from building acoustic models with a million hours of speech," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2019, IEEE.
- [29] Haşim Sak, Andrew Senior, and Françoise Beaufays, "Long short-term memory recurrent neural network architectures for large scale acoustic modeling," in *Fifteenth annual conference of the international speech communication association*, 2014.
- [30] Anirudh Raju, Sankaran Panchapagesan, Xing Liu, Arindam Mandal, and Nikko Ström, "Data augmentation for robust keyword spotting under playback interference," Aug. 2018, arXiv:1808.00563 e-prints.
- [31] Jont B. Allen and David A. Berkley, "Image method for efficiently simulating smallroom acoustics," *The Journal of the Acoustical Society of America*, vol. 65, no. 4, pp. 943–950, 1979.
- [32] Robin Scheibler, Eric Bezzam, and Ivan Dokmanić, "Pyroomacoustics: A python package for audio room simulations and array processing algorithms," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2018, IEEE.
- [33] Karel Veselý, Arnab Ghoshal, Lukáš Burget, and Daniel Povey, "Sequence-discriminative training of deep neural networks," in *Proceedings of Interspeech 2013*. 2013, number 8, pp. 2345–2349, International Speech Communication Association.
- [34] Jeremy H.M. Wong and Mark J.F. Gales, "Sequence student-teacher training of deep neural networks," in *INTERSPEECH 2016*, pp. 2761–2765.
- [35] Matthew Gibson and Thomas Hain, "Hypothesis spaces for minimum bayes risk training in large vocabulary speech recognition," in *Ninth International Conference on Spoken Language Processing*, 2006.