



Auxiliary Interference Speaker Loss for Target-Speaker Speech Recognition

Naoyuki Kanda¹, Shota Horiguchi¹, Ryoichi Takashima¹, Yusuke Fujita¹, Kenji Nagamatsu¹,
Shinji Watanabe²

¹Hitachi Ltd., Japan

²Johns Hopkins University, USA

{naoyuki.kanda.kn, shota.horiguchi.wk, ryoichi.takashima.dh, yusuke.fujita.su,
kenji.nagamatsu.dm}@hitachi.com, shinjiw@ieee.org

Abstract

In this paper, we propose a novel auxiliary loss function for target-speaker automatic speech recognition (ASR). Our method automatically extracts and transcribes target speaker's utterances from a monaural mixture of multiple speakers speech given a short sample of the target speaker. The proposed auxiliary loss function attempts to additionally maximize interference speaker ASR accuracy during training. This will regularize the network to achieve a better representation for speaker separation, thus achieving better accuracy on the target-speaker ASR. We evaluated our proposed method using two-speaker-mixed speech in various signal-to-interference-ratio conditions. We first built a strong target-speaker ASR baseline based on the state-of-the-art lattice-free maximum mutual information. This baseline achieved a word error rate (WER) of 18.06% on the test set while a normal ASR trained with clean data produced a completely corrupted result (WER of 84.71%). Then, our proposed loss further reduced the WER by 6.6% relative to this strong baseline, achieving a WER of 16.87%. In addition to the accuracy improvement, we also showed that the auxiliary output branch for the proposed loss can even be used for a secondary ASR for interference speakers' speech.

Index Terms: multi-talker speech recognition, deep learning

1. Introduction

Thanks to the recent advances in deep-learning [1–3], the accuracy of automatic speech recognition (ASR) for some datasets have become close to (e.g., Switchboard [4–6]) or beyond (e.g., LibriSpeech in [7] and [8]) the level of human transcribers. However, despite this progress, the accuracy of multi-talker speech recognition is still very limited [9–11] because of the difficulty of separating the target speaker's speech from other speech. One example is meeting speech recognition, which is known for having word error rates (WERs) around 30% (e.g. [9, 12]) even with state-of-the-art speech recognizers.

In this paper, we focus on ASR for monaural audio that contains overlapped speech uttered by multiple speakers. One direction to solve this problem is applying a speech separation method before ASR, such as deep clustering [13], deep attractor network [14], etc. However, a major drawback of this approach is that the training criteria for speech separation do not necessarily maximize ASR accuracy. If the goal is ASR, it will be better to use training criteria that directly maximize ASR accuracy.

Recently, multi-speaker ASR based on permutation invariant training (PIT) was proposed [15–19]. In the PIT scheme, the label permutation problem is solved by considering all possible permutations when calculating the loss [20]. PIT was first proposed for speech separation [20] and soon extended to ASR loss with promising results [15–19]. However, one hidden draw-

back of PIT arises when speaker tracing or speaker diarization is required. Namely, a PIT-ASR model produces transcriptions for each utterance of speakers in an unordered way, and it is no longer straightforward to solve speaker permutations across utterances. To make things worse, a PIT model trained with ASR-based loss normally does not produce separated speech waveforms, which makes speaker tracing more difficult. For applications in which speaker tracing or speaker diarization has an essential role (e.g. ASR for meeting recordings), it could become a serious drawback of the PIT model.

On the other hand, target-speaker ASR, which automatically extracts and transcribes only the target speaker's utterances given a short sample of the target speaker's speech, has been proposed [21, 22]. Zmolikova et al. proposed a target-speaker neural beamformer that extracts a target speaker's utterances given a short sample of the target speaker's speech [21]. This model was recently extended to deal with ASR-based loss to maximize ASR accuracy with promising results [22]. While the target-speaker models require additional input of a target speaker's speech sample, it can naturally solve the speaker permutation problem across utterances without using additional speaker identification after ASR. We believe this property of target-speaker ASR is attractive for many practical applications.

Following the discussion above, in this paper, we focus on target-speaker ASR, and propose a novel auxiliary loss function to improve target-speaker ASR accuracy. Our proposed loss function attempts to additionally maximize interference speaker ASR accuracy (Fig. 1) and will regularize the network to achieve better representation, which, as a result, produces better target-speaker ASR accuracy in multi-speaker conditions. We evaluated our proposed method using two-speaker-mixed speech in various signal-to-interference (SIR)-ratio conditions to demonstrate its effectiveness. We also conducted various investigations on our proposed method and model architectures, including the possibility to use the auxiliary branch for the proposed loss for a secondary ASR. In this secondary ASR setting, our model can explicitly output not only transcriptions of the target speaker but also those of other speakers in a consistent order across utterances, which is another unique property of our model.

As an additional contribution of this paper, to the best of our knowledge, this is the first report applying a lattice-free maximum mutual information (LF-MMI)-based acoustic model (AM) [23] for target-speaker ASR¹. Thanks to the state-of-the-art performance given by LF-MMI, our results were fairly good. For example, we achieved a WER of 16.50% for the “wsj0-2mix” dataset [13] for which WERs in the range of 25 - 30% have been reported [18, 19, 24, 25].

¹LF-MMI has been applied for the PIT-ASR model [16].

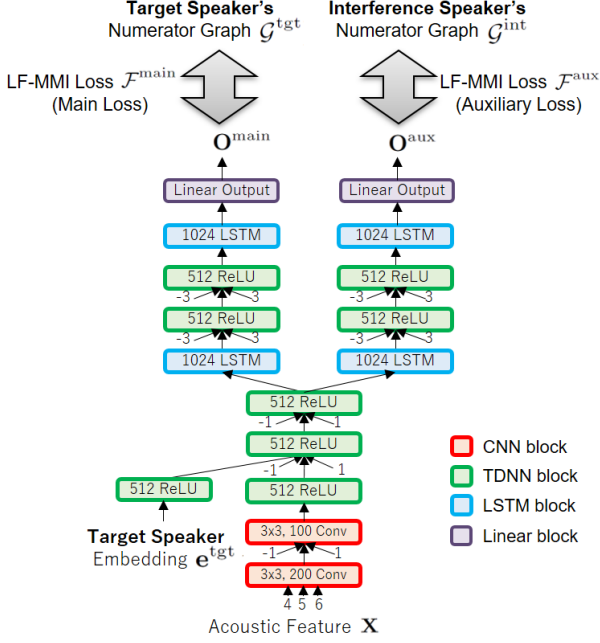


Figure 1: Overview of target-speaker AM architecture with auxiliary interference speaker loss. Auxiliary networks for interference speaker loss are only used in training and normally removed in the decoding phase. A number with an arrow indicates a time splicing index, which forms the basis of a time-delay neural network (TDNN) [26]. The input features were advanced by five frames, which has the same effect as reference label delay.

2. Auxiliary Interference Speaker Loss

In this section, we explain our proposed method and its use with an LF-MMI-based AM due to its state-of-the-art performance [23, 27, 28]. However, it should be noted that our work can be extended to any kind of ASR loss like cross entropy (CE) [1], state-level minimum Bayes risk (sMBR) [29, 30], and lattice-free sMBR [8]. In addition, the idea can even be extended to end-to-end models like connectionist temporal classification (CTC) [31] and attention encoder-decoder-based models [32, 33].

2.1. Overview of the proposed model

Fig. 1 describes our proposed model architecture for the LF-MMI AM. The network has two input branches. One branch accepts acoustic features \mathbf{X} as a normal AM while the other branch accepts an embedding \mathbf{e}^{tgt} that represents the characteristics of the target speaker. In this study, we used a log Mel-filterbank (FBANK) and i-vector [34, 35] for the acoustic features and target speaker embedding, respectively.

The key idea is in its output branch. The proposed model has multiple output branches which produce outputs \mathbf{O}^{main} and \mathbf{O}^{aux} for the main and proposed auxiliary loss functions, respectively. The main loss attempts to maximize the target speaker’s ASR accuracy, while the auxiliary loss attempts to maximize interference speaker ASR accuracy. In this study, we used LF-MMI for both criteria, details of which are explained at the latter section. Our assumption is as follows: By maximizing the ASR accuracy of the interference speaker as well as that of the target speaker, a better representation for speaker separation is learned in the shared layers of the network, resulting in improved accuracy on the target speaker’s ASR.

The network is trained with a mixture of multi-speaker speech given their transcriptions. We assume that, for each training data, (1) at least two speakers’ transcriptions are given, (2) at least the transcription for the target speaker is marked so that we can identify the target speaker’s transcription, and (3) a sample for the target speaker can be used to extract speaker embeddings. This assumption can be easily satisfied by artificially generating training data by mixing multiple speakers’ speech. After finishing the network training, the auxiliary output branches corresponding to the interference speakers are removed, and only the network branch for the target speaker is used for the decoding.

2.2. Main loss function

In the case of LF-MMI, the main loss function (to minimize) for the target speaker is defined as ²

$$\begin{aligned} \mathcal{F}^{\text{main}} &= \sum_u \text{LFMMI}(\mathbf{O}_u^{\text{main}}; \mathcal{G}_u^{\text{tgt}}), \\ &= \sum_u \sum_{\mathbf{S}} -P(\mathbf{S} | \mathbf{O}_u^{\text{main}}, \mathcal{G}_u^{\text{tgt}}) \log P(\mathbf{S} | \mathbf{O}_u^{\text{main}}, \mathcal{G}^D), \end{aligned}$$

where u is the index of training samples. The term $\mathcal{G}_u^{\text{tgt}}$ indicates a numerator (or reference) graph that represents a set of possible correct state sequences for the utterance of the target speaker of u -th training sample. The term \mathbf{S} denotes a hypothesis state sequence for u -th training sample. The term \mathcal{G}^D denotes a denominator graph, which represents a possible hypothesis space and normally consists of a 4-gram phone language model in LF-MMI training [23].

Note that one important technique known as cross entropy (CE)-regularization is normally used for LF-MMI training [23]. With this technique, an additional output layer that is optimized on the basis of the CE criterion is introduced. In training, parameters are optimized to minimize the weighted interpolation of LF-MMI and CE criteria. We used CE-regularization for our evaluation not only for the main loss but also for the auxiliary interference speaker loss. However, we omit the CE-regularization term here for the brevity of the equations.

2.3. Auxiliary interference speaker loss

The proposed auxiliary interference speaker loss is defined to maximize interference speaker ASR accuracy, which we expect to enhance the speaker separation ability of the network. In the case of using the LF-MMI criterion, the loss is defined as follows³

$$\mathcal{F}^{\text{aux}} = \sum_u \text{LFMMI}(\mathbf{O}_u^{\text{aux}}; \mathcal{G}_u^{\text{int}}),$$

where $\mathcal{G}_u^{\text{int}}$ denotes a numerator (or reference) graph that represents a set of possible correct state sequences for the utterance of the interference speaker of u -th training sample.

²This is a numerator-graph ($\mathcal{G}_u^{\text{tgt}}$)-based extension of a basic MMI-criterion $\mathcal{F}^{\text{MMI}} = \sum_u -\log P(\mathbf{S}_u | \mathbf{O}_u^{\text{main}})$ [29, 30], which uses a Viterbi-aligned reference state sequence \mathbf{S}_u instead of $\mathcal{G}_u^{\text{tgt}}$.

³The loss can be extended in the case of multiple interference speakers in accordance with the PIT principle, as follows.

$$\mathcal{F}^{\text{aux}} = \sum_u \min_{i \in \text{permu}(I)} \sum_n \text{LFMMI}(\mathbf{O}_{u,n}^{\text{aux}}; \mathcal{G}_{u,i[n]}^{\text{int}}),$$

where $\text{permu}(I)$ represents a set of permutations of interference speakers $\{1, \dots, I\}$, and n denotes the index of the permutation i .

Table 1: WERs (%) for the two-speaker mixed LibriSpeech corpus in various SIR conditions. Note that for clean single-speaker speech, a clean AM achieved WERs of 4.88% and 5.54% for dev-clean and test-clean, respectively.

SIR of the targeted speaker’s speech	dev-clean (two speakers mixed)						test-clean (two speakers mixed)					
	10	5	0	-5	-10	Avg.	10	5	0	-5	-10	Avg.
Clean AM	65.14	78.72	88.56	91.96	94.31	83.74	66.92	79.45	89.18	92.87	95.15	84.71
Target-Speaker AM w/o aux. loss	13.98	15.02	16.80	18.42	21.54	17.15	15.59	16.13	17.65	19.17	21.78	18.06
Target-Speaker AM w/ aux. loss	13.51	14.32	16.02	17.57	20.51	16.39	14.59	15.06	16.46	17.86	20.38	16.87

Finally, the loss function $\mathcal{F}^{\text{comb}}$ for training is defined as the combination of the main and auxiliary losses, as

$$\mathcal{F}^{\text{comb}} = \mathcal{F}^{\text{main}} + \alpha \cdot \mathcal{F}^{\text{aux}},$$

where α is the scaling factor for the auxiliary loss. In our evaluation, we set $\alpha = 1.0$, which gave us promising results.

Note that the original target-speaker ASR corresponds to the model without the proposed auxiliary loss; Namely, it is the case of $\alpha = 0.0$.

3. Evaluation

3.1. Experiment with LibriSpeech

3.1.1. Experimental settings

For our primary evaluation, we used the LibriSpeech corpus [36], which consists of about 1,000 hours of read-aloud English speech. In this study, we used 100 hours of the clean portion of the dataset for training the AMs. For evaluation, we used “dev-clean” and “test-clean” in accordance with the Kaldi recipes [37] as the basic materials for development and evaluation sets, respectively.

For all training, development, and evaluation data, we artificially generated two-speaker mixed speech in accordance with the protocol below.

1. Prepare a list of speech samples (main list), which is the main target of ASR.
2. Shuffle the main list to create a second list under the constraint that the same speaker does not appear in the same line in the main and second lists.
3. Mix the audio in the main list and the second list one-by-one, with a specific SIR. For training data, we randomly sampled an SIR value from uniform distribution between -10 dB and 10 dB for each mixture. For the development and evaluation data, we generated data with an SIR of {10, 5, 0, -5, -10} dB.
4. Only in the case of the training data, the volume of each mixed speech was randomly changed to enhance robustness for the volume difference.

Note that, in accordance with the protocol above, the speech of the target speaker could be much shorter or much longer than that of the interference speaker. We intentionally selected this protocol because we believe it is also important to evaluate the ability to correctly select the portion of the target speaker’s speech when a single speaker, who could be either target speaker or interference speaker, is speaking.

We trained an acoustic model consisting of a convolutional neural network (CNN), time-delay neural network (TDNN) [38], and long short-term memory (LSTM) [39] as shown in Fig. 1. Input acoustic features for the network were 40-dimensional log-Mel-filterbank (FBANK) without normalization. In addition, a 100-dimensional i-vector was extracted and used for the target speaker embedding to indicate the target

speaker. For extracting i-vector, we randomly selected an utterance of the same speaker. We conducted 8 epochs of training on the basis of LF-MMI, where the initial learning rate was set to 0.001 and exponentially decayed to 0.0001 by the end of the training. We applied l_2 -regularization and CE-regularization [23] with scales of 0.00005 and 0.1, respectively. The leaky hidden Markov model coefficient was set to 0.1. In addition, a backstitch technique [40] with a backstitch scale of 1.0 and backstitch interval of 4 was used.

For comparison purposes, we trained the AM without the proposed auxiliary loss, which corresponds to the original target-speaker model. We also trained a “clean AM” using clean, non-speaker-mixed speech. For this model, we also used a model architecture without the auxiliary output branch, and i-vector was extracted every 100 msec to realize online speaker/environment adaptation⁴.

In decoding, we used an officially provided large 4-gram LM. For each test utterance, a sample of the target speaker was randomly selected from the other utterances in the test set⁵. The average duration of the target speaker’s sample was 7.2 sec and 7.4 sec for dev-clean and test-clean, respectively. All decoding parameters were tuned by the development set, and the best parameters were used for the evaluation set. All of our experiments were conducted on the basis of the Kaldi toolkit [37].

3.1.2. Effect of auxiliary interference speaker loss

The first row of Table 1 shows the result for the AM trained with clean data (“clean AM”). Note that WERs for clean speech with the clean AM was 4.88% and 5.54% for dev-clean and test-clean, respectively. As shown in the table, WERs were severely degraded by mixing the two speakers’ speech, and the clean AM produced WER averages of 83.74% and 84.71% for dev-clean and test-clean, respectively.

The second row of Table 1 shows the result for the original target-speaker AM without auxiliary loss. This model dramatically improved the accuracy, achieving WER averages of 17.15% and 18.06% for dev-clean and test-clean, respectively.

The third row of Table 1 shows the result for the proposed auxiliary loss function. It achieved the best WER among all SIR conditions. The proposed auxiliary loss achieved WERs of 16.39% and 16.87% for dev-clean and test-clean, respectively; a relative WER reduction of 6.6%.

3.1.3. Comparison of model architectures

To better understand the proposed method, we investigated the effect of the model architecture for the auxiliary loss. Although we added the auxiliary output branch at the middle point of the network in Fig. 1, we could add the auxiliary output branch at a different point as shown in Fig. 2.

The results for each model architecture are shown in Fig. 3. We first found that the early splitting (Fig. 2 (a)) was ineffective

⁴This is one of the standard use cases for a clean AM in the Kaldi toolkit.

⁵We used the same random seed over all experiments so that all experiments were conducted on the basis of the same conditions.




Figure 2: Model architectures with early, middle, and late splitting. Note that the middle splitting model was used as the default model in other experiments.




Figure 3: Comparison among model architectures in Fig. 2.

in all cases. We believe this to be natural because if we split the network early, only a limited number of network parameters can be used for speech separation.

The comparison of the middle splitting (Fig. 2 (b)) and the late splitting (Fig. 2 (c)) produced some complicated results. As shown in Fig. 3, the middle splitting model showed robust improvements over broad SIR conditions. On the other hand, the late splitting model showed very good WERs when the SIR was higher than 5 dB, but the improvement became marginal when the SIR became low. Our interpretation is as follows. In the case of the late splitting model, the responsibility for the output prediction for the target speaker was concentrated on the last layer. Although this did not result in any severe problem when the evaluation condition was simple (i.e. a high SIR), problems did occur when the condition became more complicated. On the basis of these interpretation, we concluded that the middle splitting model is the most robust architecture.

3.1.4. Interference speaker’s ASR by the auxiliary network

So far, we explained that the auxiliary output branch is removed in decoding. However, the auxiliary output branch could be used for interference speaker ASR. One possible scenario is that there is one target speaker whose utterance should be recognized with a mark of “presenter”, and other audiences’ utterances are recognized via the auxiliary branch with a mark of “questioners.”

Therefore, we evaluated the ability of an auxiliary network for the secondary ASR. In this evaluation, we provided the target speaker’s embeddings for the network, and evaluated the WERs of the ASR results using the output of the auxiliary output branch. The results are shown in Table 2. From this table, we found that the auxiliary output branch worked very well for the secondary ASR. This result clearly indicated that the shared layers of the network were learned to separate speaker separation as we expected. In addition, we believe this secondary ASR itself is attractive as exemplified above. Different from PIT-ASR models, our model can explicitly output not only transcriptions

Table 2: WERs (%) for two-speaker-mixed test-clean. Main output branch was used for the target speaker’s ASR and auxiliary output branch was used for the interference speaker’s ASR.

SIR of target spk.	SIR of interference spk.	WER (%)	
		target spk.	interference spk.
10	-10	14.59	26.22
5	-5	15.06	19.90
0	0	16.46	16.23
-5	5	17.86	15.05
-10	10	20.38	14.50
Avg.		16.87	18.38

Table 3: WERs (%) for WSJ corpus with clean AM and target-speaker (TS) AMs.

Model	Dev93	Eval92	wsj0-2mix
Clean AM	77.51	78.03	79.81
TS-AM w/o aux. loss	12.24	11.32	16.78
TS-AM w/ aux. loss	11.31	11.38	16.50

of the target speaker but also those of other speakers in a consistent order across utterances.

3.2. Experiment with Wall Street Journal corpus

For our final evaluation, we conducted experiments on the Wall Street Journal (WSJ) corpus. In accordance with [19], we used the WSJ SI284 for the training data and Dev93 for the development data. For evaluation, we used Eval92 and “wsj0-2mix” [13]. When testing Dev93 and Eval92, we mixed two speakers’ speech with SIRs of 0 dB. Other experimental settings were the same with those in the previous section.

The results are shown in Table 3. We observed a significant improvement for Dev93 and moderate improvement for wsj0-2mix, while a marginal degradation of WER was observed for EVAL92. Note that the results by our AM were fairly good thanks to the state-of-the-art accuracy given by the LF-MMI. Our result for “wsj0-2mix” (16.50% WER) is significantly better than the results reported by the conventional PIT-based ASR method, which achieved about a WER in the range of 25 - 30% [18, 19, 24, 25]. It is also significantly better than the result reported in the target-speaker ASR paper [22], which reported a WER of 34.0% for the WSJ corpus in 0-dB-mixture settings.

4. Conclusions

In this paper, we proposed a novel auxiliary loss function for target-speaker ASR, in which it attempts to maximize interference speaker ASR accuracy. We evaluated our proposed method using two-speaker-mixed speech in various SIR conditions. We first built a strong target-speaker ASR baseline based on the state-of-the-art LF-MMI, achieving a WER of 18.06% on the test set while a normal ASR trained with clean data produced a completely corrupted result (WER of 84.71%). Then, our proposed loss further reduced the WER by 6.6% relative to this strong target-speaker ASR baseline, achieving a WER of 16.87%. By investigating model architectures for the proposed loss, we determined that adding an auxiliary output branch from the middle point of the network worked the most robustly. We also showed that the auxiliary output branch can be used for a secondary ASR for interference speakers’ speech.

5. References

- [1] F. Seide, G. Li, and D. Yu, “Conversational speech transcription using context-dependent deep neural networks,” in *Proc. INTER-SPEECH*, 2011, pp. 437–440.

- [2] G. E. Dahl, D. Yu, L. Deng, and A. Acero, "Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition," *IEEE Trans. on ASLP*, vol. 20, no. 1, pp. 30–42, 2012.
- [3] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, and B. Kingsbury, "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.
- [4] W. Xiong, J. Droppo, X. Huang, F. Seide, M. Seltzer, A. Stolcke, D. Yu, and G. Zweig, "Achieving human parity in conversational speech recognition," *arXiv preprint arXiv:1610.05256*, 2016.
- [5] G. Saon, G. Kurata, T. Sercu, K. Audhkhasi, S. Thomas, D. Dimitriadis, X. Cui, B. Ramabhadran, M. Picheny, L.-L. Lim *et al.*, "English conversational telephone speech recognition by humans and machines," *Proc. INTERSPEECH*, pp. 132–136, 2017.
- [6] W. Xiong, J. Droppo, X. Huang, F. Seide, M. L. Seltzer, A. Stolcke, D. Yu, and G. Zweig, "Toward human parity in conversational speech recognition," *IEEE/ACM Trans. on ASLP*, vol. 25, no. 12, pp. 2410–2423, 2017.
- [7] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen *et al.*, "Deep speech 2: End-to-end speech recognition in English and Mandarin," in *Proc. ICML*, 2016, pp. 173–182.
- [8] N. Kanda, Y. Fujita, and K. Nagamatsu, "Lattice-free state-level minimum Bayes risk training of acoustic models," in *Proc. INTERSPEECH*, 2018, pp. 2923–2927.
- [9] T. Yoshioka, H. Erdogan, Z. Chen, X. Xiao, and F. Alleva, "Recognizing overlapped speech in meetings: A multichannel separation approach using neural networks," in *Proc. INTERSPEECH*, 2018, pp. 3038–3042.
- [10] J. Barker, S. Watanabe, E. Vincent, and J. Trmal, "The fifth CHiME speech separation and recognition challenge: Dataset, task and baselines," in *Proc. INTERSPEECH*, 2018, pp. 1561–1565.
- [11] N. Kanda, R. Ikeshita, S. Horiguchi, Y. Fujita, K. Nagamatsu, X. Wang, V. Manohar, N. E. Y. Soplin, M. Maciejewski, S.-J. Chen *et al.*, "The Hitachi/JHU CHiME-5 system: Advances in speech recognition for everyday home environments using multiple microphone arrays," in *Proc. CHiME-5*, 2018, pp. 6–10.
- [12] N. Kanda, Y. Fujita, S. Horiguchi, R. Ikeshita, K. Nagamatsu, and S. Watanabe, "Acoustic modeling for distant multi-talker speech recognition with single- and multi-channel branches," in *Proc. ICASSP*, 2019.
- [13] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *Proc. ICASSP*, 2016, pp. 31–35.
- [14] Z. Chen, Y. Luo, and N. Mesgarani, "Deep attractor network for single-microphone speaker separation," in *Proc. ICASSP*, 2017, pp. 246–250.
- [15] D. Yu, X. Chang, and Y. Qian, "Recognizing multi-talker speech with permutation invariant training," in *Proc. INTERSPEECH*, 2017, pp. 2456–2460.
- [16] Z. Chen, J. Droppo, J. Li, W. Xiong, Z. Chen, J. Droppo, J. Li, and W. Xiong, "Progressive joint modeling in unsupervised single-channel overlapped speech recognition," *IEEE/ACM Trans. on ASLP*, vol. 26, no. 1, pp. 184–196, 2018.
- [17] S. Settle, J. Le Roux, T. Hori, S. Watanabe, and J. R. Hershey, "End-to-end multi-speaker speech recognition," in *Proc. ICASSP*, 2018, pp. 4819–4823.
- [18] H. Seki, T. Hori, S. Watanabe, J. Le Roux, and J. R. Hershey, "A purely end-to-end system for multi-speaker speech recognition," in *Proc. ACL*, 2018, pp. 2620–2630.
- [19] X. Chang, Y. Qian, K. Yu, and S. Watanabe, "End-to-end monaural multi-speaker asr system without pretraining," in *Proc. ICASSP*, 2019.
- [20] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *Proc. ICASSP*, 2017, pp. 241–245.
- [21] K. Zmolikova, M. Delcroix, K. Kinoshita, T. Higuchi, A. Ogawa, and T. Nakatani, "Speaker-aware neural network based beamformer for speaker extraction in speech mixtures," in *Proc. INTERSPEECH*, 2017.
- [22] M. Delcroix, K. Zmolikova, K. Kinoshita, A. Ogawa, and T. Nakatani, "Single channel target speaker extraction and recognition with speaker beam," in *ICASSP*, 2018, pp. 5554–5558.
- [23] D. Povey, V. Peddinti, D. Galvez, P. Ghahramani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely sequence-trained neural networks for ASR based on lattice-free MMI," in *Proc. INTERSPEECH*, 2016, pp. 2751–2755.
- [24] Y. Isik, J. Le Roux, Z. Chen, S. Watanabe, and J. R. Hershey, "Single-channel multi-speaker separation using deep clustering," in *Proc. INTERSPEECH*, 2016, pp. 545–549.
- [25] Y. Qian, X. Chang, and D. Yu, "Single-channel multi-talker speech recognition with permutation invariant training," *Speech Communication*, vol. 104, pp. 1–11, 2018.
- [26] V. Peddinti, D. Povey, and S. Khudanpur, "A time delay neural network architecture for efficient modeling of long temporal contexts," in *Proc. INTERSPEECH*, 2015, pp. 3214–3218.
- [27] N. Kanda, Y. Fujita, and K. Nagamatsu, "Investigation of lattice-free maximum mutual information-based acoustic models with sequence-level Kullback-Leibler divergence," in *Proc. ASRU*, 2017, pp. 69–76.
- [28] N. Kanda, Y. Fujita, and K. Nagamatsu, "Sequence distillation for purely sequence trained acoustic models," in *Proc. ICASSP*, 2018, pp. 5964–5968.
- [29] K. Vesely, A. Ghoshal, L. Burget, and D. Povey, "Sequence-discriminative training of deep neural networks," in *Proc. INTERSPEECH*, 2013, pp. 2345–2349.
- [30] H. Su, G. Li, D. Yu, and F. Seide, "Error back propagation for sequence training of context-dependent deep networks for conversational speech transcription," in *Proc. ICASSP*, 2013, pp. 6664–6668.
- [31] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proc. ICML*, 2006, pp. 369–376.
- [32] J. Chorowski, D. Bahdanau, K. Cho, and Y. Bengio, "End-to-end continuous speech recognition using attention-based recurrent NN: First results," *Proc. NIPS workshop on Deep Learning and Representation Learning*, 2014.
- [33] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *Proc. ICASSP*, 2016, pp. 4960–4964.
- [34] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE Trans. on ASLP*, vol. 19, no. 4, pp. 788–798, 2011.
- [35] G. Saon, H. Soltau, D. Nahamoo, and M. Picheny, "Speaker adaptation of neural network acoustic models using i-vectors," in *Proc. ASRU*, 2013, pp. 55–59.
- [36] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an ASR corpus based on public domain audio books," in *Proc. ICASSP*, 2015.
- [37] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlíček, Y. Qian, P. Schwarz *et al.*, "The Kaldi speech recognition toolkit," in *Proc. ASRU*, 2011.
- [38] A. Waibel, T. Hanazawa, G. Hinton, K. Shikano, and K. J. Lang, "Phoneme recognition using time-delay neural networks," *IEEE Trans. ASSP*, vol. 37, no. 3, pp. 328–339, 1989.
- [39] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [40] Y. Wang, V. Peddinti, H. Xu, X. Zhang, D. Povey, and S. Khudanpur, "Backstitch: Counteracting finite-sample bias via negative steps," in *Proc. INTERSPEECH*, 2017, pp. 1631–1635.