



Building state-of-the-art distant speech recognition using the CHiME-4 challenge with a setup of speech enhancement baseline

Szu-Jui Chen, Aswin Shanmugam Subramanian, Hainan Xu, Shinji Watanabe

Center for Language and Speech Processing, Johns Hopkins University, Baltimore, MD 21218, USA

{schen146, asubra13, hxu31, shinjiw}@jhu.edu

Abstract

This paper describes a new baseline system for automatic speech recognition (ASR) in the CHiME-4 challenge to promote the development of noisy ASR in speech processing communities by providing 1) state-of-the-art system with a simplified single system comparable to the complicated top systems in the challenge, 2) publicly available and reproducible recipe through the main repository in the Kaldi speech recognition toolkit. The proposed system adopts generalized eigenvalue beamforming with bidirectional long short-term memory (LSTM) mask estimation. We also propose to use a time delay neural network (TDNN) based on the lattice-free version of the maximum mutual information (LF-MMI) trained with augmented all six microphones plus the enhanced data after beamforming. Finally, we use a LSTM language model for lattice and n-best re-scoring. The final system achieved 2.74% WER for the real test set in the 6-channel track, which corresponds to the 2nd place in the challenge. In addition, the proposed baseline recipe includes four different speech enhancement measures, short-time objective intelligibility measure (STOI), extended STOI (eSTOI), perceptual evaluation of speech quality (PESQ) and speech distortion ratio (SDR) for the simulation test set. Thus, the recipe also provides an experimental platform for speech enhancement studies with these performance measures.

Index Terms: Speech recognition, noise robustness, mask-based beamforming, lattice-free MMI, LSTM language modeling

1. Introduction

In recent years, multi-channel speech recognition has been applied on devices used in daily life, such as Amazon Echo and Google Home. The recognition accuracy is greatly improved by exploiting microphone arrays when compared to single channel microphone devices [1–3]. However, satisfactory performance is still not achieved in noisy everyday environments. Hence, the CHiME-4 challenge is designed to conquer this scenario by recognizing speech in challenging noisy environments [4]. Through the series of the challenge activities, several speech enhancement and recognition techniques are established as an effective method for this scenario including mask-based beamforming, multichannel data augmentation, and system combination with various front-end techniques [5–9].

Although many submitted systems in the CHiME-4 challenge have yielded a lot of outcomes in this multi-channel Automatic Speech Recognition (ASR) scenario [6–8], one of the drawbacks is that all top systems are highly complicated due to multiple systems and fusion techniques, and it is not easy for the other research groups to follow these outcomes. This paper aims to deal with the above drawback by building a new baseline to promote the development of noisy ASR in speech enhancement, separation, and recognition communities.

We propose a single ASR system to further push the border of this challenge. Most important of all, our system is reproducible since it is implemented in the Kaldi ASR toolkit and other opensource toolkits. All the scripts in our experiments can be downloaded from the official GitHub website¹. The original CHiME-4 baseline is described in [4], which uses a delay-and-sum beamformer (BeamformIt) [10], a deep neural network with state-level minimum Bayes Risk (DNN+sMBR) criterion [11], and recurrent neural network-based language model (RNNLM) [12]. On the contrary, our proposed system is shown in Figure 1. We adopt to use bidirectional long short-term memory (BLSTM) mask based beamformer (Section 3.2), which has been shown to be more effective [13, 14] than BeamformIt. For acoustic model, the DNN used in baseline is limited to represent long-term dependencies between acoustic characteristics. Hence, a sub-sampled time delay neural network (TDNN) [15] with the lattice-free version of the maximum mutual information (LF-MMI) is used for our acoustic model [16] (Section 3.3). This paper also shows the great improvement on the word error rate (WER) when we combine it with data augmentation in a multichannel scenario using all six microphones plus the enhanced data after beamforming. Then, we further use a LSTM language model (LSTMLM), which uses a new training criterion and importance sampling, and has been shown to be more efficient and better in performance [17], to re-score hypotheses.

We also incorporate computation of four different speech enhancement measures in our recipe - perceptual evaluation of speech quality (PESQ) [18], short-time objective intelligibility measure (STOI) [19], extended STOI (eSTOI) [20] and speech distortion ratio (SDR) [21]. We include these measurements as part of the recipe for two reasons. First, the ASR performance shows only one aspect of the speech enhancement algorithm. Objective enhancement metrics can give an indication on how well the enhancement is with different aspects (e.g., intelligibility, signal distortions). Second, testing an enhancement algorithm with ASR takes a significant amount of computational time, whereas obtaining these scores is quite fast. Hence, it can give an initial indication of how good the enhancement is.

2. Related work

In [6], a fusion system in the DNN posterior domain is proposed to get the best result in the competition. [7–9] also use fusion systems in the decoding hypothesis domain with multiple systems mainly using different front-end techniques. Unlike these highly complicated systems, our proposed system is based on a *single system* without the above fusion systems, yet achieves comparable performance to these top systems in the challenge task. One of the unique technical aspects of our proposed system is to fully utilize the effectiveness of TDNN with

¹<https://github.com/kaldi-asr/kaldi/pull/2142>

LF-MMI by combining it with multichannel data augmentation techniques, which achieves significant improvement. Our new LSTMLM also contributes to boost the final performance.

3. Proposed system

Our system starts from BLSTM mask based beamformer and is followed by feature extraction. Phoneme to audio alignments are then generated by GMM acoustic model and are fed into TDNN acoustic model for training. Finally, the lattices after first pass decoding in TDNN are re-scored by a 5-gram LM and further re-scored by LSTMLM.

3.1. Data augmentation

Training with multichannel data has been shown to be effective for ASR systems [1, 8, 22]. This augmentation can increase the variety in the training data and help the generalization to test set. In our work, we not only use data from all 6 channels but also add the enhanced data generated by beamformer to training set.

Let $O = (\mathbf{o}(t) \in \mathbb{R}^D | t = 1, \dots, T)$ be a sequence of D -dimensional feature vectors with length T , which is a single channel speech recognition case. In our case, we deal with an M -channel input ($M = 6$), which is represented as $\mathbf{O} = (\mathbf{o}_m(t) \in \mathbb{R}^D | t = 1, \dots, T, m = 1, \dots, M)$. Then, the original training method only uses a particular channel input (e.g., m -th input) as training data to obtain acoustic model parameters Θ , as follows:

$$\hat{\Theta} = \arg \max_{\Theta} \mathcal{L}(\mathbf{O}_m), \quad (1)$$

where \mathcal{L} is an objective function (log likelihood for the GMM case and negative cross entropy for the DNN case), with reference labels as supervisions. Data augmentation approach tries to use training data of all channels, as follows:

$$\hat{\Theta} = \arg \max_{\Theta} \mathcal{L}(\mathbf{O} = \{\mathbf{O}_m\}_{m=1}^M) \quad (2)$$

Further, we extend to include an enhanced data $\mathbf{O}^{\text{enh}} = (\mathbf{o}^{\text{enh}}(t) \in \mathbb{R}^D | t = 1, \dots, T)$ with the above multichannel data, that is

$$\hat{\Theta} = \arg \max_{\Theta} \mathcal{L}(\{\mathbf{O}, \mathbf{O}^{\text{enh}}\}), \quad (3)$$

where the enhancement data \mathbf{O}^{enh} is obtained by a single-channel masking or beamformer method, which is described in Section 3.2.

3.2. BLSTM mask based beamformer

We use the BLSTM mask based Generalized Eigenvalue (GEV) beamformer described in [14]. The GEV beamforming procedure requires an estimate of the Cross-Power Spectral Density (PSD) matrix of the noise and the target speech. The BLSTM model estimates two masks: the first mask indicates the time frequency bin that is probably dominated by speech and the other indicates which is dominated by noise. With the combined speech and noise masks, we can estimate the PSD matrices of speech components $\Phi_{\text{speech}}(b) \in \mathbb{C}^{M \times M}$ at frequency bin b , and that of noise components $\Phi_{\text{noise}}(b) \in \mathbb{C}^{M \times M}$, as follows:

$$\Phi_v(b) = \sum_{t=1}^T w_v(t, b) \mathbf{y}(t, b) \mathbf{y}(t, b)^H \text{ where } v \in \{\text{speech, noise}\}, \quad (4)$$

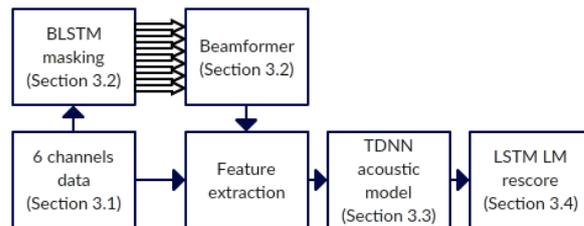


Figure 1: Diagram of speech recognition system.

where $\mathbf{y}(t, b) \in \mathbb{C}^M$ is an M -dimensional complex spectrum at time (frame) t in frequency bin b . \mathbf{y}^H denotes the conjugate transpose. $w_v(t, b) \in [0, 1]$ is the mask value.

The goal of GEV beamformer [23] is to estimate the beamforming filter $\mathbf{f}(b)$, which maximizes the expected SNR for each frequency bin b as given by the equation below:

$$\mathbf{f}_{\text{GEV}}(b) = \arg \max_{\mathbf{f}(b)} \frac{\mathbf{f}^H(b) \Phi_{\text{speech}}(b) \mathbf{f}(b)}{\mathbf{f}^H(b) \Phi_{\text{noise}}(b) \mathbf{f}(b)}. \quad (5)$$

Eq. (5) is equivalent to solve the following eigenvalue problem:

$$(\Phi_{\text{noise}}(b))^{-1} \Phi_{\text{speech}}(b) \mathbf{f}(b) = \lambda \mathbf{f}(b), \quad (6)$$

where $\mathbf{f}(b) \in \mathbb{C}^M$ at each frequency bin b is the M -dimensional complex eigenvector and λ is the eigenvalue.

3.3. Time delayed neural network with lattice-free MMI

For acoustic model, we use TDNN with LF-MMI training [16] instead of DNN+sMBR [11]. The architecture is similar to those described in [24]. The LF-MMI objective function is shown below, which is different from usual MMI training [25] in a sense that we use phoneme sequence L instead of a word sequence to narrow down a search space in the denominator:

$$\mathcal{L}_{\text{MMI}} = \sum_{n=1}^N \log \frac{p(\mathbf{O}^n | S^n)^C P(L^n)}{\sum_L p(\mathbf{O}^n | S^L)^C P(L)} \quad (7)$$

where $p(\mathbf{O}^n | S^L)$ is the likelihood function of a speech feature sequence \mathbf{O}^n given the state sequence S^L at n 'th utterance. $P(L)$ is the phoneme language model probability and C is the probability scale.

Note that when combined with the data augmentation technique (described in Section 3.1), TDNN is more effective than DNN.

3.4. LSTM language modeling

The LSTM based language model (LSTMLM) has been shown to be effective on language modeling [26]. It is better in finding a longer period of contextual information than conventional RNN. With this property, LSTMLM can predict the next word in a more accurate way than RNNLM. Hence, instead of using a vanilla RNNLM [12], we train an LSTMLM on WSJ data, which combines the use of subword features and one-hot encoding. An importance sampling method is used to speed up training. Most important of all, a new objective function \mathcal{L}_{LM} is used for LM training, which behaves like cross-entropy objective but trains the output to auto-normalize in order to speed up test time computation:

$$\mathcal{L}_{\text{LM}} = z_j + 1 - \sum_i \exp(z_i) \quad (8)$$

Table 1: *Speech Enhancement Scores*

Track	Enhancement Method	Dev (Simu)				Test (Simu)			
		PESQ	STOI	eSTOI	SDR	PESQ	STOI	eSTOI	SDR
1ch	No Enhancement	2.01	0.82	0.61	3.92	1.98	0.81	0.60	4.95
1ch	BLSTM Mask	2.52	0.88	0.73	9.26	2.46	0.87	0.71	10.76
2ch	BeamformIt	2.15	0.85	0.65	4.61	2.07	0.83	0.62	5.60
2ch	BLSTM Gev	2.13	0.87	0.69	2.86	2.12	0.87	0.69	3.10
6ch	BeamformIt	2.31	0.88	0.70	5.52	2.20	0.86	0.65	6.30
6ch	BLSTM Gev	2.45	0.88	0.75	3.57	2.46	0.87	0.73	2.92

where z is a pre-activation vector in the layer of neural network before the final softmax operation and j is an index for the correct word. More details can be found in [17].

Table 2: *Experimental configurations*

BLSTM mask estimation	
input layer dimension	513
L1 - BLSTM layer dimension	256
L2 - FF layer 1 (ReLU) dimension	513
L3 - FF layer 2 (clipped ReLU) dimension	513
L4 - FF layer (Sigmoid) dimension	1026
$p_{dropout}$ for L1, L2 and L3	0.5
TDNN acoustic model	
input layer dimension	40
hidden layer dimension	750
output layer dimension	2800
l2-regularize	0.00005
num-epochs	6
initial-effective-lrate	0.003
final-effective-lrate	0.0003
shrink-value	1.0
num-chunk-per-minibatch	128,64
LSTM language model	
layers dimension	2048
recurrent-projection-dim	512
N-best list size	100
RNN re-score weight	1.0

4. Experiments

4.1. Speech Enhancement Experiments

First experiments describe the speech enhancement performance of BLSTM-based speech enhancement. For the single channel track, we used the BLSTM masking technique [27] trained on the 6 channel data and took only the speech mask after the forward propagation.

We took a Hadamard product of the single channel spectrogram with the speech mask and used it as the enhanced signal to compare it with the original signal without any enhancement. For the 2ch and 6ch tracks, we used the BLSTM based GEV beamformer described in Section 3.2 and compared it with BeamformIt. Four different scores as described in Section 1 - PESQ, STOI, eSTOI and SDR are computed. The BLSTM architecture used in the experiments is listed in Table 2.

The enhancement scores are shown in Table 1. The 5th channel clean signal from the 6ch data convolved with room

impulse response was used as the reference signal for computing all the four metrics. For the 1ch track, the BLSTM mask gives significantly better scores in all four metrics compared to the noisy data without any enhancement. However, this is contrary to the ASR results, which will be discussed in the next section. BeamformIt has better SDR scores compared to BLSTM GEV in both the multi-channel tracks. Also, for both the multi-channel track data, eSTOI is slightly better for BLSTM GEV. In the 6ch track experiments, BLSTM GEV has a significantly better PESQ score. Overall, BLSTM-based speech enhancement shows improvement in most of conditions except for the case of the multichannel SDR metric.

4.2. Speech Recognition Experiments

Our system is trained on the speech recognition toolkit Kaldi [28]. For TDNN acoustic model training, backstitch optimization method [29] is used. The decoding is based on 3-gram language models with explicit pronunciation and silence probability modeling as described in [30]. The model is re-scored by a 5-gram language model first. Then the Kaldi-RNNLM [17] is used for training the LSTMLM, and n-best re-scoring is used to improve performance. We got our best result in 6 channel experiments by averaging forward and backward LSTMLM. The RNN re-score weight is set to be 1.0, which means the results of 5-gram LM is completely discarded. All the results in this section are reported in terms of word error rate (WER). We also provide the parameters used in our system in Table 2.

Table 3: *WER of adding enhanced data when using TDNN with BeamformIt and RNNLM in the 6 channel track experiment*

Data Augmentation	Dev (%)		Test (%)	
	real	simu	real	simu
all 6ch data	3.97	4.33	7.04	7.39
all 6ch and enhanced data	3.74	4.31	6.84	7.49

Table 3 shows the effectiveness of the data augmentation for the system using TDNN with BeamformIt and RNNLM, which are described in Section 2, in the 6 channel track experiment. We confirmed the improvement by adding enhanced data in almost all cases except for the simulation test data. This is also found in 2 channels experiment when using TDNN (i.e. row 3 and row 4 in table 5).

Tables 4 and 5 show the WER of 6 channel and 2 channel experiments. We change our experimental condition incrementally to compare the effectiveness of each method described in Section 2. In most of the situations, every method improved the WER steadily. We observed that the performance was degraded

Table 4: WER of 6 channel track experiments

Data Augmentation	Method			Dev (%)		Test (%)	
	Acoustic Model	Beamforming	Language Model	real	simu	real	simu
only 5th channel	DNN+sMBR	BeamformIt	RNNLM	5.79	6.73	11.50	10.92
all 6ch data	DNN+sMBR	BeamformIt	RNNLM	5.05	5.82	9.50	9.24
all 6ch and enhanced data	DNN+sMBR	BeamformIt	RNNLM	5.62	6.46	10.27	9.41
all 6ch and enhanced data	TDNN with LF-MMI	BeamformIt	RNNLM	3.74	4.31	6.84	7.49
all 6ch and enhanced data	TDNN with LF-MMI	BLSTM Gev	RNNLM	2.83	2.94	4.01	3.80
all 6ch and enhanced data	TDNN with LF-MMI	BLSTM Gev	LSTMLM	1.90	2.10	2.74	2.66

Table 5: WER of 2 channel track experiments

Data Augmentation	Method			Dev (%)		Test (%)	
	Acoustic Model	Beamforming	Language Model	real	simu	real	simu
only 5th channel	DNN+sMBR	BeamformIt	RNNLM	8.23	9.50	16.58	15.33
all 6ch data	DNN+sMBR	BeamformIt	RNNLM	6.87	8.06	13.33	12.57
all 6ch data	TDNN with LF-MMI	BeamformIt	RNNLM	5.57	6.08	10.53	9.90
all 6ch and enhanced data	TDNN with LF-MMI	BeamformIt	RNNLM	5.03	6.02	10.20	10.35
all 6ch and enhanced data	TDNN with LF-MMI	BLSTM Gev	RNNLM	3.79	5.03	6.93	6.07
all 6ch and enhanced data	TDNN with LF-MMI	BLSTM Gev	LSTMLM	2.85	3.94	5.40	5.03

Table 6: WER of 1 channel track experiments

Data Augmentation	Acoustic Model	Beamforming	Language Model	Dev (%)		Test (%)	
				real	simu	real	simu
only 5th channel	DNN+sMBR	-	RNNLM	11.57	12.98	23.70	20.84
all 6ch data	DNN+sMBR	-	RNNLM	8.97	11.02	18.10	17.31
all 6ch data	TDNN with LF-MMI	-	RNNLM	6.64	7.78	12.92	13.54
all 6ch data	TDNN with LF-MMI	-	LSTMLM	5.58	6.81	11.42	12.15
all 6ch data	TDNN with LF-MMI	BLSTM masking	RNNLM	13.15	15.62	22.47	21.61
all 6ch and enhanced data	TDNN with LF-MMI	BLSTM masking	LSTMLM	6.78	9.10	13.64	14.95

if we applied enhanced data on the system using DNN+sMBR (i.e. row 2 and row 3 in table 4), while TDNN with LF-MMI could make use of the enhanced data, as discussed above. In addition, comparing with the speech enhancement results in Table 1, it shows that better speech enhancement scores do not necessarily give lower WER. Especially, there always seems to be a negative correlation between the ASR performance and the SDR scores.

Table 6 illustrates the results of the 1 channel track experiment. We found that BLSTM masking was not effective if we only used one microphone although it scores better in terms of all four speech enhancement metrics in Table 1. From row 3 and row 5 of 6, the WER with BLSTM masking was degraded more than twice when compared to the system without BLSTM masking. However, we also discovered that after adding the enhanced data into the system with BLSTM masking, the WER became closer to the best setup without masking, which can be seen in row 4 and row 6 of 6. Thus, adding the enhanced data seems to be a good strategy to mitigate the degradation of speech enhancement.

Finally, Table 7 presents the comparison with the official baseline and top systems in the CHiME-4 challenge. We can see that all of these systems use a fusion technique to get their best WER. On the other hand, our proposed single system achieved 76% relative improvement from the official baseline, and achieved the 2nd best performance.

Table 7: Final WER comparison for the real test set.

System	# systems	WER (%)
CHiME-4 baseline [4]	1	11.51
Proposed system	1	2.74
USTC-iFlytek [6]	5	2.24
RWTH/UPB/FORTH [7]	5	2.91
MERL [8]	6	2.98

5. Conclusion

This paper describes our single ASR system for CHiME-4 speech separation and recognition challenge. The system consists of BLSTM masked GEV beamformer (Section 3.2), TDNN with LF-MMI as acoustic model (Section 3.3) and re-scoring using LSTMLM (Section 3.4), which trained on all 6 channels data plus enhanced data generated by beamformer (Section 3.1). The system finally achieved 2.74% WER, which outperforms the 2nd place result in the challenge. The system is publicly available through the Kaldi speech recognition toolkit.

Our future work will explore different architectures for TDNN and LSTM networks for further improvement. Furthermore, this system can be applied to other multichannel tasks such as AMI [31], and the CHiME-5 challenge [32].

6. References

- [1] J. Barker, R. Marxer, E. Vincent, and S. Watanabe, "The third CHiMEspeech separation and recognition challenge: Dataset, task and baselines," in *IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, 2015, pp. 504–511.
- [2] K. Kinoshita, M. Delcroix, S. Gannot, E. A. Habets, R. Haeb-Umbach, W. Kellermann, V. Leutnant, R. Maas, T. Nakatani, B. Raj *et al.*, "A summary of the REVERB challenge: state-of-the-art and remaining challenges in reverberant speech processing research," *EURASIP Journal on Advances in Signal Processing*, vol. 2016, no. 1, p. 7, 2016.
- [3] B. Li, T. Sainath, A. Narayanan, J. Caroselli, M. Bacchiani, A. Misra, I. Shafran, H. Sak, G. Pundak, K. Chin *et al.*, "Acoustic modeling for google home," *INTERSPEECH-2017*, pp. 399–403, 2017.
- [4] E. Vincent, S. Watanabe, A. A. Nugraha, J. Barker, and R. Marxer, "An analysis of environment, microphone and data simulation mismatches in robust speech recognition," *Computer Speech & Language*, vol. 46, pp. 535–557, 2017.
- [5] T. Yoshioka, N. Ito, M. Delcroix, A. Ogawa, K. Kinoshita, M. Fujimoto, C. Yu, W. J. Fabian, M. Espi, T. Higuchi *et al.*, "The NTT CHiME-3 system: Advances in speech enhancement and recognition for mobile multi-microphone devices," in *Automatic Speech Recognition and Understanding (ASRU)*. IEEE, 2015, pp. 436–443.
- [6] J. Du, Y.-H. Tu, L. Sun, F. Ma, H.-K. Wang, J. Pan, C. Liu, J.-D. Chen, and C.-H. Lee, "The USTC-iFlytek system for CHiME-4 challenge," *Proc. CHiME*, pp. 36–38, 2016.
- [7] T. Menne, J. Heymann, A. Alexandridis, K. Irie, A. Zeyer, M. Kitzka, P. Golik, I. Kulikov, L. Drude, R. Schlüter *et al.*, "The RWTH/UPB/FORTH system combination for the 4th CHiME challenge evaluation," in *CHiME-4 workshop*, 2016.
- [8] H. Erdogan, T. Hayashi, J. R. Hershey, T. Hori, C. Hori, W.-N. Hsu, S. Kim, J. Le Roux, Z. Meng, and S. Watanabe, "Multi-channel speech recognition: LSTMs all the way through," in *CHiME-4 workshop*, 2016.
- [9] Y. Fujita, T. Homma, and M. Togami, "Unsupervised network adaptation and phonetically-oriented system combination for the chime-4 challenge," *Proc. CHiME*, pp. 49–51, 2016.
- [10] X. Anguera, C. Wooters, and J. Hernando, "Acoustic beamforming for speaker diarization of meetings," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 7, pp. 2011–2022, 2007.
- [11] K. Vesely, A. Ghoshal, L. Burget, and D. Povey, "Sequence-discriminative training of deep neural networks," in *Interspeech*, 2013, pp. 2345–2349.
- [12] T. Mikolov, M. Karafiát, L. Burget, J. Černocký, and S. Khudanpur, "Recurrent neural network based language model," in *Interspeech*, 2010.
- [13] H. Erdogan, J. R. Hershey, S. Watanabe, M. I. Mandel, and J. Le Roux, "Improved MVDR beamforming using single-channel mask prediction networks," in *INTERSPEECH*, 2016, pp. 1981–1985.
- [14] J. Heymann, L. Drude, and R. Haeb-Umbach, "Neural network based spectral mask estimation for acoustic beamforming," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2016, pp. 196–200.
- [15] A. Waibel, T. Hanazawa, G. Hinton, K. Shikano, and K. J. Lang, "Phoneme recognition using time-delay neural networks," in *Readings in speech recognition*. Elsevier, 1990, pp. 393–404.
- [16] D. Povey, V. Peddinti, D. Galvez, P. Ghahremani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely Sequence-Trained Neural Networks for ASR Based on Lattice-Free MMI," in *Interspeech*, 2016, pp. 2751–2755.
- [17] H. Xu, K. Li, Y. Wang, J. Wang, S. Kang, X. Chen, D. Povey, and S. Khudanpur, "Neural network language modeling with letter-based features and importance sampling," 2018.
- [18] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual Evaluation of Speech Quality (PESQ)-a New Method for Speech Quality Assessment of Telephone Networks and Codecs," in *Proceedings of the Acoustics, Speech, and Signal Processing, 2000. On IEEE International Conference - Volume 02*, ser. ICASSP '01. IEEE Computer Society, 2001, pp. 749–752.
- [19] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, "An algorithm for intelligibility prediction of time-frequency weighted noisy speech," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 7, pp. 2125–2136, Sept 2011.
- [20] J. Jensen and C. H. Taal, "An algorithm for predicting the intelligibility of speech masked by modulated noise maskers," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 11, pp. 2009–2022, Nov 2016.
- [21] E. Vincent, R. Gribonval, and C. Fevotte, "Performance measurement in blind audio source separation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 4, pp. 1462–1469, July 2006.
- [22] T. Hori, Z. Chen, H. Erdogan, J. R. Hershey, J. Le Roux, V. Mitra, and S. Watanabe, "Multi-microphone speech recognition integrating beamforming, robust feature extraction, and advanced DNN/RNN backend," *Computer Speech & Language*, vol. 46, pp. 401–418, 2017.
- [23] E. Warsitz and R. Haeb-Umbach, "Blind acoustic beamforming based on generalized eigenvalue decomposition," *IEEE Transactions on audio, speech, and language processing*, vol. 15, no. 5, pp. 1529–1539, 2007.
- [24] V. Peddinti, D. Povey, and S. Khudanpur, "A time delay neural network architecture for efficient modeling of long temporal contexts," in *Interspeech*, 2015.
- [25] D. Povey, "Discriminative training for large vocabulary speech recognition," Ph.D. dissertation, University of Cambridge, 2005.
- [26] M. Sundermeyer, R. Schlüter, and H. Ney, "LSTM neural networks for language modeling," in *Interspeech*, 2012.
- [27] "Speech enhancement with LSTM recurrent neural networks and its application to noise-robust ASR, author=Weninger, Felix and Erdogan, Hakan and Watanabe, Shinji and Vincent, Emmanuel and Le Roux, Jonathan and Hershey, John R and Schuller, Björn, booktitle=International Conference on Latent Variable Analysis and Signal Separation, pages=91–99, year=2015, organization=Springer."
- [28] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The Kaldi speech recognition toolkit," in *IEEE 2011 workshop on automatic speech recognition and understanding*. IEEE Signal Processing Society, 2011.
- [29] Y. Wang, V. Peddinti, H. Xu, X. Zhang, D. Povey, and S. Khudanpur, "Backstitch: Counteracting finite-sample bias via negative steps," in *Interspeech*, 2017.
- [30] G. Chen, H. Xu, M. Wu, D. Povey, and S. Khudanpur, "Pronunciation and silence probability modeling for ASR," in *Interspeech*, 2015.
- [31] I. McCowan, J. Carletta, W. Kraaij, S. Ashby, S. Bourban, M. Flynn, M. Guillemot, T. Hain, J. Kadlec, V. Karaiskos *et al.*, "The AMI meeting corpus," in *Proceedings of the 5th International Conference on Methods and Techniques in Behavioral Research*, vol. 88, 2005, p. 100.
- [32] J. Barker, S. Watanabe, E. Vincent, and J. Trmal, "The fifth CHiME speech separation and recognition challenge: Dataset, task and baselines," in *Interspeech*, 2018, (submitting).