



Student-Teacher Learning for BLSTM Mask-based Speech Enhancement

Aswin Shanmugam Subramanian, Szu-Jui Chen, Shinji Watanabe

Center for Language and Speech Processing, Johns Hopkins University

{aswin, schen146, shinjiw}@jhu.edu

Abstract

Spectral mask estimation using bidirectional long short-term memory (BLSTM) neural networks has been widely used in various speech enhancement applications, and it has achieved great success when it is applied to multichannel enhancement techniques with a mask-based beamformer. However, when these masks are used for single channel speech enhancement they severely distort the speech signal and make them unsuitable for speech recognition. This paper proposes a student-teacher learning paradigm for single channel speech enhancement. The beamformed signal from multichannel enhancement is given as input to the teacher network to obtain soft masks. An additional cross-entropy loss term with the soft mask target is combined with the original loss, so that the student network with single-channel input is trained to mimic the soft mask obtained with multichannel input through beamforming. Experiments with the CHiME-4 challenge single channel track data shows improvement in ASR performance.

Index Terms: Speech enhancement, speech recognition, mask estimation, BLSTM, student-teacher learning

1. Introduction

The presence of background noise and reverberation degrades the performance of an automatic speech recognition (ASR) system. The performance of noise robust ASR was greatly improved by using multiple microphones instead of just using a single microphone [1, 2, 3]. Especially, mask-based beamforming techniques have shown outstanding results in this scenario [4, 5, 6], and many of the top systems in the CHiME-4 challenge [7] use these techniques for speech enhancement. For example, [5, 6] use a bidirectional long short-term memory (BLSTM) neural network to accurately predict speech (and noise) masks, given the noisy spectrogram. In [5], both speech and noise masks are in turn used to calculate cross power spectral density (PSD) matrices of the noise and the target speech, which are used by the generalized eigenvalue (GEV) beamformer to get the beamforming filter [8].

However, compared with the success of multichannel speech enhancement, single-channel speech enhancement is not well established, especially when the method is combined with other speech processing applications including ASR and speaker recognition [9]. Single-channel speech enhancement, especially based on deep neural networks (DNNs), has been studied by many research groups [10, 11, 12, 13] with promising performance improvement in terms of the speech enhancement metrics or hearing purposes. However, we often observe a degradation in performance when we use single-channel enhancement as a preprocessing step for ASR due to the spectral distortions induced by the enhancement. This paper focuses on the single-channel speech enhancement to overcome this issue.

Our idea is to fill out the gap between single-channel and multichannel speech enhancement by using a well known DNN

technique called student-teacher learning. We propose to train a *teacher* network, which takes the beamformed signal obtained from multichannel speech enhancement as input and predicts high-quality speech masks. Then, a *student* network, which takes the noisy signal as input, is trained to mimic the mask predicted by the teacher network with the corresponding beamformed signal as input. Conventionally student-teacher learning [14, 15] is used to reduce the complexity of the network. For example, a less complex student model with fewer parameters tries to mimic the soft targets of a more complex teacher model. In [16], student-teacher training is used for self-supervised learning of an ASR acoustic model, where the student model tries to mimic the effects of multichannel speech enhancement by using only single channel inputs. This was achieved by training the student network with noisy training data as input to mimic the soft posteriors of the teacher network trained on the enhanced speech. In addition, the benefits of using soft targets compared to hard targets for DNN acoustic models are also shown in [16], which is advocated by [17] as knowledge distillation.

This paper follows BLSTM-mask based beamforming proposed in [5], which has two cross-entropy loss terms - one for speech mask target and the other for noise mask target. The speech mask predicted by the model can also be used for single-channel enhancement, which this paper uses as the baseline. We have an additional loss term based on the cross entropy with the soft mask from the teacher network as a target. The effectiveness of the proposed method is investigated by using the single channel track of the CHiME-4 dataset [18]. The CHiME-4 dataset consists of both real and simulated data and it is common to use only the simulation data to prepare the clean and noise targets. As the proposed method uses the soft mask from the teacher network as the target, which can be obtained even for the real data, it can use both real and simulation data in the training stage. This is also a unique aspect of the proposed method. Additionally, four different speech enhancement metrics - perceptual evaluation of speech quality (PESQ) [19], short-time objective intelligibility measure (STOI) [20], extended STOI (eSTOI) [21] and speech distortion ratio (SDR) [22] are used as part of the experiments to discuss the performance of speech enhancement in addition to the ASR performance.

2. BLSTM mask-based speech enhancement

This section first describes a mask prediction method by using a binary cross entropy loss, which is a basic component of this paper. It also describes a mask-based beamformer to estimate beamforming filters based on the estimated masks.

2.1. Binary cross entropy loss

The noise-aware training of BLSTM mask proposed in [5] is explained in this section. The network estimates two masks:

the first is the speech mask which denotes the degree of how dominant the speech component is in each time-frequency bin, while the second is the noise mask which denotes that of the noise component [23]. The ideal binary speech mask target $\text{IBM}_X(t, b) \in \{0, 1\}$ at frame t in frequency bin b is defined based on the signal-to-noise (SNR) ratio with thresholding as:

$$\text{IBM}_X(t, b) = \begin{cases} 1, & \text{if } \frac{\|x(t, b)\|}{\|n(t, b)\|} > \text{threshold} \\ 0, & \text{otherwise} \end{cases}, \quad (1)$$

where, $\|x(t, b)\| \in \mathbb{R}_{\geq 0}$ is the power spectrum of the clean speech signal and $\|n(t, b)\| \in \mathbb{R}_{\geq 0}$ is the power spectrum of the noise signal at each time-frequency bin (t, b) . The ideal binary noise mask $\text{IBM}_N(t, b) \in \{0, 1\}$ is also calculated in a similar way.

Given a sequence of T -length noisy speech magnitude spectra $Y = (\|y(t, b)\|_{b=1}^B | t = 1, \dots, T)$, the BLSTM network predicts the speech mask $w_X(t, b) \in [0, 1]$ and noise mask $w_N(t, b) \in [0, 1]$ at each time-frequency bin (t, b) , as follows:

$$w_v(t, b) = \sigma(\text{Lin}_v(\text{BLSTM}(Y))), \text{ where } v \in \{X, N\}, \quad (2)$$

where $\sigma(\cdot)$, $\text{Lin}_v(\cdot)$, and $\text{BLSTM}(\cdot)$ are sigmoid activation, linear, and output of the masking network given in Table 1, respectively.

The binary cross entropy is used as the loss function to train the network. The overall loss function is defined by combining the speech and noise binary loss functions (loss_X and loss_N) as follows:

$$\text{loss} = \text{loss}_X + \text{loss}_N, \quad (3)$$

$$\triangleq \frac{1}{T * B} \sum_{t, b} \sum_{v \in \{X, N\}} \text{CE}(\text{IBM}_v(t, b), w_v(t, b)), \quad (4)$$

where $\text{CE}(a, a')$ is the binary cross entropy defined as follows:

$$\begin{aligned} \text{CE}(a, a') &\triangleq a \log a' + (1 - a) \log(1 - a') \\ &= \begin{cases} a \log a' & a = 0 \\ (1 - a) \log(1 - a') & a = 1 \end{cases}. \end{aligned} \quad (5)$$

Note that since $a \in \{0, 1\}$, either first or second term in the right hand side of Eq. (5) becomes zero. This type of target is particular called as a *hard label* in the student-teacher learning [14] or knowledge distillation [17] context.

When we apply the estimated speech mask $w_X(t, b)$ to the original signal $y(t, b)$, we can perform single-channel speech enhancement as follows:

$$x^{(\text{se})}(t, b) = w_X(t, b)y(t, b), \quad (6)$$

where $x^{(\text{se})}(t, b)$ is an enhanced signal (the superscript (se) denotes single-channel enhancement). Although the single channel enhancement only requires the prediction of the speech mask, the noise mask is used to obtain the noise PSD matrix in a multichannel scenario, which will be discussed in the next section.

2.2. Mask-based beamformer

This section extends to deal with a multichannel signal with M as a number of channels. When the single channel mask estimation in the previous section is applied to the multichannel signal, the corresponding speech mask $w_{m, X}(t, b)$ and noise mask $w_{m, N}(t, b)$ for each channel m can be obtained. With

Table 1: *Masking network architecture*

Layer	Activation	Dimension
Input	-	513
BLSTM	Tanh	256
Feedforward 1	ReLU	513
Feedforward 2	clipped ReLU	513

these multichannel masks, the following time-frequency mask is obtained from a median operation:

$$\bar{w}_v(t, b) = \text{Median}(\{w_{m, v}(t, b)\}_{m=1}^M) \text{ where } v \in \{X, N\}. \quad (7)$$

With the speech and noise masks $\bar{w}_X(t, b)$ and $\bar{w}_N(t, b)$, the PSD matrices, which are $\Phi_X(b) \in \mathbb{C}^{M \times M}$ of speech at frequency bin b , and $\Phi_N(b) \in \mathbb{C}^{M \times M}$ of noise at frequency bin b , can be estimated as follows:

$$\Phi_v(b) = \sum_{t=1}^T \bar{w}_v(t, b) \mathbf{y}(t, b) \mathbf{y}(t, b)^H \text{ where } v \in \{X, N\} \quad (8)$$

where $\mathbf{y}(t, b) \in \mathbb{C}^M$ is the M -dimensional complex spectral value at time (frame) t and frequency bin b . The operator H is the Hermitian transpose.

From these PSD matrices, the M -dimensional complex beamforming filter $\mathbf{f}(b) \in \mathbb{C}^M$ can be estimated. This paper adopts the GEV beamformer, which is obtained by maximizing the expected SNR with respect to $\mathbf{f}(b)$ for each frequency bin as given by the equation below:

$$\mathbf{f}_{\text{GEV}}(b) = \underset{\mathbf{f}(b)}{\text{argmax}} \frac{\mathbf{f}^H(b) \Phi_X(b) \mathbf{f}(b)}{\mathbf{f}^H(b) \Phi_N(b) \mathbf{f}(b)}. \quad (9)$$

This optimization is equivalent to solving the following eigenvalue problem:

$$(\Phi_N(b))^{-1} \Phi_X(b) \mathbf{f}(b) = \lambda \mathbf{f}(b), \quad (10)$$

where $\mathbf{f}(b)$ is the eigenvector and λ is the corresponding eigenvalue.

Once we obtain the beamforming filter $\mathbf{f}_{\text{GEV}}(b)$, we can perform multichannel speech enhancement as follows:

$$x^{(\text{me})}(t, b) = \mathbf{f}_{\text{GEV}}^H(b) \mathbf{y}(t, b), \quad (11)$$

where $x^{(\text{me})}(t, b)$ is an enhanced signal (the superscript (me) denotes multichannel enhancement).

In general, $x^{(\text{me})}(t, b)$ can be well-denoised by making use of spacial information and beamforming operation, compared to the single-channel enhanced signal $x^{(\text{se})}(t, b)$ introduced in Section 2.1. This paper proposes to use these single- and multichannel speech enhancement properties, and designs a new objective function for single-channel enhancement, by using the mask obtained by multichannel enhancement as a better *soft label*.

3. Student-teacher model

This section explains in detail the proposed single-channel enhancement technique by using a student-teacher model.

3.1. Teacher model

Firstly, the beamformed signal $x^{(\text{me})}(t, b)$ is given as input to the teacher model. The architecture of the network is the same as

Table 2: WER of HMM-GMM ASR System

	Parameters				Train data (ASR)	BLSTM Mask	WER Dev (%)		WER Test (%)	
	λ_1	λ_2	λ_3	epoch			real	simu	real	simu
1	-	-	-	-	all 6ch noisy	-	21.40	23.22	35.63	31.98
2	-	-	-	-	all 6ch noisy	Baseline	28.99	28.05	40.98	35.50
3	-	-	-	7	all 6ch noisy	Teacher	24.91	26.00	40.26	35.73
4	1/3	1/3	1/3	6	all 6ch noisy	Student	25.95	24.66	35.50	29.98
5	0.25	0.25	0.5	12	all 6ch noisy	Student	26.56	26.19	36.33	31.36
6	0.50	0.25	0.25	5	all 6ch noisy	Student	25.82	25.17	35.86	30.17
7	0.45	0.05	0.50	3	all 6ch noisy	Student	25.54	24.77	34.78	29.78
8	0.35	0.15	0.50	3	all 6ch noisy	Student	23.34	23.11	33.11	28.30
9	0.05	0.45	0.50	3	all 6ch noisy	Student	24.07	24.01	34.88	30.00
10	0.15	0.35	0.50	3	all 6ch noisy	Student	25.01	24.66	35.66	30.10
11	0.35	0.15	0.50	3	all 6ch noisy	Student with real	23.42	23.55	32.64	28.88
12	-	-	-	-	all 6ch noisy + 5th ch enhanced data from baseline	Baseline	22.07	23.37	34.02	30.41
13	0.35	0.15	0.50	3	all 6ch noisy + 5th ch enhanced data from baseline	Student	19.78	20.76	30.66	26.60
14	0.35	0.15	0.50	3	all 6ch noisy + 5th ch enhanced data from baseline	Student with real	19.79	20.85	29.80	26.66

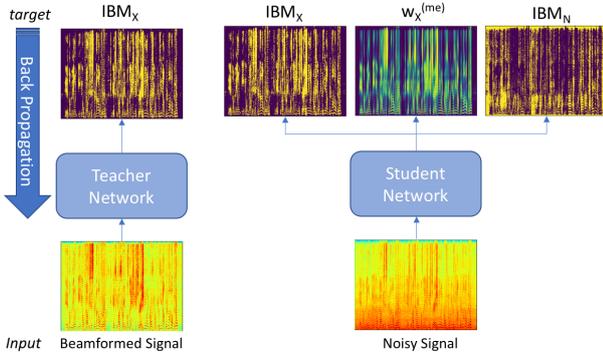


Figure 1: Student-Teacher Model

the one explained in Section 2.1 except that the output is only the speech mask $w_X^{(me)}(t, b)$ and there is no noise mask. Hence the loss function for the teacher network is

$$\text{loss} = \frac{1}{T * B} \sum_{t,b} \text{CE}(\text{IBM}_X(t, b), w_X^{(me)}(t, b)). \quad (12)$$

Since the beamformed signal is already well-denoised, the above teacher network provides a high-quality speech mask. Alternatively, the original clean signal $x(t, b)$ can be used as the input instead of the beamformed signal to train the teacher model. However, the problem of estimating a speech mask given a clean signal is too trivial, and the network is not well-trained in such a trivial condition.

3.2. Student model

With the speech mask $w_X^{(me)}(t, b)$ obtained by the above teacher network, we can additionally consider the following student-teacher loss function:

$$\text{loss}_{\text{st}} = \frac{1}{T * B} \sum_{t,b} \text{CE}(w_X^{(me)}(t, b), w_X^{(se)}(t, b)). \quad (13)$$

Compared with the binary cross entropy case in Eq. (5), the target $w_X^{(me)}(t, b) \in [0, 1]$ is a bounded continuous value, and provides richer supervisions to the student network.

The final loss function is expressed with the student-teacher loss and the original loss (Eq. (3)) as:

$$\text{loss} = \lambda_1 \text{loss}_{\text{st}} + \lambda_2 \text{loss}_X + \lambda_3 \text{loss}_N \quad (14)$$

where λ_1 , λ_2 , and λ_3 are linear interpolation weights. The student model is trained with this loss function, which expects to learn the multichannel enhancement ability within a single-channel enhancement framework. Figure 1 illustrates the proposed student-teacher training paradigm.

3.3. Student model with real data

All the above formulations are discussed given that we have parallel clean speech (x) and noise (n), which are not usually obtained in the real recording. However, the proposed network can obtain the soft target of the real data by simply applying multichannel speech enhancement for the real data. In this scenario, it has only one student-teacher loss term loss_{st} (Eq. (13)), which exactly mimics the teacher network. With this setup, the loss function can be obtained by switching the student-teacher loss (Eq. (13)) and the combined loss (Eq. (14)) for real and simulation data as follows:

$$\text{loss} = \begin{cases} \text{loss}_{\text{st}} & \text{for real} \\ \lambda_1 \text{loss}_{\text{st}} + \lambda_2 \text{loss}_X + \lambda_3 \text{loss}_N & \text{for simulation} \end{cases} \quad (15)$$

This is our final loss function, which fully makes use of the benefit of the student-teacher learning paradigm in speech enhancement.

4. Experiments

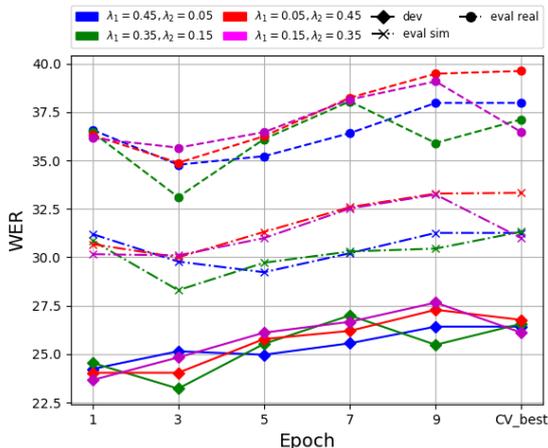
To verify the effectiveness of the proposed single-channel experiments, the 1 channel track in the CHiME-4 challenge [7] was used.

4.1. Speech recognition

The Kaldi speech recognition toolkit [24] was used for ASR experiments. The GMM baseline provided by the CHiME-4 challenge [7] was used for comparing different mask models. The mask network was implemented using the Chainer neural network toolkit [25]. The cross-validation loss was calculated

Table 3: *Speech Enhancement Scores.*

Parameters				BLSTM Mask	Dev (Simu)				Test (Simu)			
λ_1	λ_2	λ_3	epoch		PESQ	STOI	eSTOI	SDR	PESQ	STOI	eSTOI	SDR
-	-	-	-	-	2.01	0.82	0.61	3.92	1.98	0.81	0.60	4.95
-	-	-	-	Baseline	2.52	0.88	0.73	9.26	2.46	0.87	0.71	10.76
-	-	-	7	Teacher	2.36	0.86	0.68	8.14	2.33	0.85	0.67	9.51
1/3	1/3	1/3	6	Student	2.54	0.88	0.73	9.13	2.49	0.87	0.71	10.64
0.50	0.25	0.25	5	Student	2.52	0.88	0.73	9.00	2.47	0.87	0.71	10.44
0.35	0.15	0.50	3	Student	N/A	0.88	0.72	8.93	N/A	0.87	0.70	10.35

Figure 2: *WER vs epoch for different parameter combinations*

using the development data after every epoch during the training of the teacher network. The training was stopped when the loss did not decrease anymore after 5 epochs of patience and the model with the least cross-validation loss was chosen. The 5th channel data was used to estimate the target masks for the beamformed data. The student network was trained in a batch mode where a minibatch comprised the frames of all 6 channels of one utterance.

To compare the performance of different masking models, we first prepared an ASR baseline that would be neutral to the enhancement method by using only the noisy data of all 6 channels as the training data. We performed enhancement on the development and evaluation sets of the official 1-channel track CHiME-4 dataset, with different models and decoded their enhanced speech with the above ASR systems. The results of these ASR experiments are shown in Table 2 (rows 1–3). The performance of the teacher model was better than that of the original BLSTM mask as expected, although both degraded the performance from the non-enhanced noisy speech.

The second experiments focused on our proposed student models (rows 4–11 in Table 2), and we found that most of the different configurations of student models performed better than that of the teacher model. Note that in this experiment, we fixed the number of epochs for several models as 3. This is because we additionally investigated the dependency between the validation loss and WER, as shown in Figure 2, and we found that it did not necessarily give us the best results for the best validation score with different hyper-parameter λ , as introduced in Eq. (14). We empirically found that 3 epochs in most of the cases seemed to give good performance. This result indicates that choosing the epoch based on the WER of the development data seems to be a better criterion than the enhancement-driven

cross-validation loss (related discussions will be shown in the next section). With this setup (3 epochs), we found that the parameter choice of $\lambda_1 = 0.35$, $\lambda_2 = 0.15$ and $\lambda_3 = 0.5$ gave the best performance amongst the different values we tried. This is an interesting finding, since this result indicates that soft targets obtained by the teacher model would effect the performance more than that of supervised targets. We also trained the student model with additional real data, as discussed in Section 3.3 (row 11 in Table 2), and this shows the improvement for the real test set owing to the inclusion of the real training data, which is a desired property in real environments.

When the 5th channel of the training data was also enhanced in the same way as the development and evaluation data using the baseline BLSTM mask and included as part of the ASR training (row 12), the performance was slightly better than using the original noisy data for the evaluation data but it was slightly worse for the development data. When the development and evaluation data was enhanced using our best student models (rows 13 and 14), the performance improved significantly compared to using the original noisy data in the all conditions. Also, we observed the similar improvement for the real test set again when we trained the student model with real data.

4.2. Speech enhancement scores

The four different scores described in Section 1 - PESQ, STOI, eSTOI and SDR were computed for several enhancement methods, as shown in Table 3. The 5th channel clean signal from the 6ch track convolved with room impulse response was used as the reference signal. All the different masking models gave significantly better scores in all four metrics compared to using the noisy data without any enhancement although we did not observe any considerable difference in the scores amongst the models. In addition, by comparing Tables 2 and 3, we did not observe clear correlations between the speech enhancement scores and WERs, which also suggests that we need some careful investigation on the objective function of BLSTM-based enhancement for the ASR purpose.

5. Conclusion

We proposed a new training paradigm for mask estimation using BLSTMs based on student-teacher training for single channel speech enhancement. We showed that the proposed student-teacher technique improved the ASR performance from the original noisy speech and the enhanced speech obtained by conventional BLSTM masking. Our future work is to evaluate our speech enhancement techniques with a strong ASR backend including a time delay neural network (TDNN) [26, 27] with the lattice-free version of the maximum mutual information (LF-MMI) [28].

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*, 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.
- [4] T. Higuchi, N. Ito, T. Yoshioka, and T. Nakatani, "Robust MVDR beamforming using time-frequency masks for online/offline ASR in noise," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 5210–5214.
- [5] J. Heymann, L. Drude, and R. Haeb-Umbach, "Neural network based spectral mask estimation for acoustic beamforming," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 196–200.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] X. Lu, Y. Tsao, S. Matsuda, and C. Hori, "Speech enhancement based on deep denoising autoencoder," in *Interspeech*, 2013, pp. 436–440.
- [11] Y. Xu, J. Du, L.-R. Dai, and C.-H. Lee, "An experimental study on speech enhancement based on deep neural networks," *IEEE Signal processing letters*, vol. 21, no. 1, pp. 65–68, 2014.
- [12] A. Narayanan and D. Wang, "Ideal ratio mask estimation using deep neural networks for robust speech recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2013, pp. 7092–7096.
- [13] H. Erdogan, J. R. Hershey, S. Watanabe, and J. Le Roux, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 708–712.
- [14] J. Ba and R. Caruana, "Do deep nets really need to be deep?" in *Advances in Neural Information Processing Systems 27 (NIPS)*, 2014, pp. 2654–2662.
- [15] J. Li, R. Zhao, J.-T. Huang, and Y. Gong, "Learning small-size DNN with output-distribution-based criteria," in *Interspeech*, 2014.
- [16] S. Watanabe, T. Hori, J. Le Roux, and J. Hershey, "Student-teacher network learning with enhanced features," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 5275–5279.
- [17] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [18] 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.
- [19] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2001, pp. 749–752.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] 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 Workshop on Automatic Speech Recognition and Understanding (ASRU)*, 2011.
- [25] S. Tokui, K. Oono, S. Hido, and J. Clayton, "Chainer: a next-generation open source framework for deep learning," in *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Twenty-ninth Annual Conference on Neural Information Processing Systems (NIPS)*, 2015.
- [26] 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.
- [27] V. Peddinti, D. Povey, and S. Khudanpur, "A time delay neural network architecture for efficient modeling of long temporal contexts," in *Interspeech*, 2015.
- [28] 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.