

Dereverberation of autoregressive envelopes for far-field speech recognition

Anurenjan Purushothaman^{a,b}, Anirudh Sreeram^b, Rohit Kumar^b,
Sriram Ganapathy^{b,*}

^a Department of ECE, College of Engineering, Thiruvananthapuram, 695016, India

^b Learning and Extraction of Acoustic Patterns (LEAP) lab, Electrical Engineering, Indian Institute of Science, Bangalore, 560012, India

ARTICLE INFO

Keywords:

Automatic speech recognition
Frequency domain linear prediction (FDLP)
Dereverberation
Neural speech enhancement
Joint learning

ABSTRACT

The task of speech recognition in far-field environments is adversely affected by the reverberant artifacts that elicit as the temporal smearing of the sub-band envelopes. In this paper, we develop a neural model for speech dereverberation using the long-term sub-band envelopes of speech. The sub-band envelopes are derived using frequency domain linear prediction (FDLP) which performs an autoregressive estimation of the Hilbert envelopes. The neural dereverberation model estimates the envelope gain which when applied to reverberant signals suppresses the late reflection components in the far-field signal. The dereverberated envelopes are used for feature extraction in speech recognition. Further, the sequence of steps involved in envelope dereverberation, feature extraction and acoustic modeling for ASR can be implemented as a single neural processing pipeline which allows the joint learning of the dereverberation network and the acoustic model. Several experiments are performed on the REVERB challenge dataset, CHiME-3 dataset and VOICES dataset. In these experiments, the joint learning of envelope dereverberation and acoustic model yields significant performance improvements over the baseline ASR system based on log-mel spectrogram as well as other past approaches for dereverberation (average relative improvements of 10–24% over the baseline system). A detailed analysis on the choice of hyper-parameters and the cost function involved in envelope dereverberation is also provided.

1. Introduction

Automatic speech recognition (ASR) is a challenging task in far-field conditions. This is particularly due to the fact that the speech signal will be reverberant and noisy. The word error rates (WER) in ASR have seen a dramatic improvement over the past decade due to the advancements in deep learning based techniques (Yu and Deng, 2016). Still the deterioration in performance in noisy and reverberant conditions persist (Hain et al., 2012). A relative increase in WER of 75% is reported by Peddinti et al. (2018) and Ganapathy and Peddinti (2018b) when the signal from headset microphone is replaced with far-field array microphone signals in the ASR systems. This deterioration is due to temporal smearing of time domain envelopes caused by reverberation (Yoshioka et al., 2012).

One common approach to suppress reverberation is to combine all channels by beamforming (Anguera et al., 2007) before feeding it to the ASR system. Recently, unsupervised neural mask estimator for generalized eigen-value beamforming is proposed (Kumar

* Corresponding author.

E-mail address: sriramg@iisc.ac.in (S. Ganapathy).

<https://doi.org/10.1016/j.csl.2021.101277>

Received 15 January 2021; Received in revised form 16 July 2021; Accepted 7 August 2021

Available online 24 August 2021

0885-2308/© 2021 Published by Elsevier Ltd.

et al., 2020). Traditional pre-processing also includes the weighted prediction error (WPE) (Nakatani et al., 2010) based dereverberation along with the beamforming in most state-of-art far-field ASR systems. Further, multi-condition training is usually used to alleviate the mismatch between training and testing (Seltzer et al., 2013). Here, either simulated reverberant data or real far-field data can be added to the training data. However, even with these techniques, the beamformed signal shows significant amount of temporal smoothing in sub-band envelopes. The temporal smearing is caused by the superposition of direct path signal and reflected signals and this leads to ASR performance degradation (Ganapathy, 2017).

In this paper, we analyze the effect of reverberation on sub-band Hilbert envelopes. We show that the effect of reverberation can be approximated as convolution of the long-term sub-band envelopes of clean speech with the envelope of room impulse response function. In order to compensate for the late reverberation component in the envelope, we explore a Wiener filtering approach where the Wiener filter gain is computed using a deep neural network (DNN). The gain estimation network is implemented using a convolutional-long short term memory (CLSTM) model. The gain is multiplied with the sub-band envelopes to suppress reverberation artifacts. The sub-band envelopes are converted to spectrographic features through integration and used for deep neural network based ASR. The sub-band envelopes are derived using the autoregressive modeling framework of frequency domain linear prediction (FDLP) (Thomas et al., 2008; Ganapathy and Harish, 2018).

The steps involved in envelope dereverberation, feature extraction and acoustic modeling for ASR can all be implemented as neural network layers. Therefore, we also propose an approach for joint learning of the speech dereverberation model with the ASR acoustic modeling network as a single neural model. Various ASR experiments are performed on the REVERB challenge dataset (Kinoshita et al., 2013) as well as the CHiME-3 dataset (Barker et al., 2015). In these experiments, we show that the proposed approach improves over the state-of-the-art ASR systems based on log-mel features as well as other past approaches proposed for speech dereverberation and denoising based on deep learning. In addition, we also extend the approach to large vocabulary speech recognition on VOICES dataset (Richey et al., 2018; Nandwana et al., 2019).

The rest of the paper is organized as follows. The related prior work is discussed in Section 2. This section also discusses the key contributions from the proposed work. Section 3 provides details regarding the reverberation artifacts and autoregressive envelope estimation using frequency domain linear prediction. In Section 4, we discuss the envelope dereverberation model, feature extraction as well as the joint approach to dereverberation with acoustic modeling for ASR. The ASR experiments and results are discussed in Section 5. Various model parameter choices and additional analyses are reported in Section 6. This is followed by a summary of the work in Section 7.

2. Related prior work

Xu et al. in Xu et al. (2014) attempted to find a mapping function from noisy and clean signals using supervised neural network, which is used for enhancement in the testing stage. In a similar manner, speech separation problem is also explored with ideal ratio mask based neural mapping (Wang and Chen, 2018). Zhao et al. proposed a LSTM model for late reflection prediction in the spectrogram domain for reverberant speech (Zhao et al., 2018). A spectral mapping approach using the log-magnitude inputs was attempted by Han et al. (2014). A mask based approach to dereverberation on the complex short-term Fourier transform domain was explored by Williamson and Wang (2017).

Speech enhancement for speech recognition based on neural networks has been explored in Wöllmer et al. (2013), Chen et al. (2015) and Weninger et al. (2015). In Maas et al. (2013), a recurrent neural network is used to map noise-corrupted input features to their corresponding clean versions. A context aware recurrent neural network based convolutional encoder-decoder architecture was used in Santos and Falk (2018) to map the power spectral features of noisy and clean speech. In a recent work by Pandey and Wang (2019), the speech enhancement is learned in the time domain itself, but using a matrix multiplication to convert the time domain signal into frequency domain and the frequency domain loss is used for training. This approach uses mean absolute error between the STFT frames of the clean and noisy speech for training.

The joint learning of the speech enhancement neural model and the acoustic model was attempted in Wang and Wang (2016). Here, a DNN based speech separation model is coupled with a DNN based acoustic model and the weights are adjusted jointly. Bo Wu et al. (2017) proposed to unify the speech enhancement neural model and the acoustic model trained separately, and then the joint model is further trained to improve the ASR performance. The power spectrum in the log domain was used as features in the enhancement stage. Bo Wu et al. (2017) also explored an end-to-end deep learning approach in, where the knowledge about reverberation time is incorporated in DNN based dereverberation front end. This reverberation time aware-DNN enhancement module and ASR acoustic module are further trained jointly to improve the ASR cost.

The key contributions from the current work can be summarized as follows,

- Deriving a signal model for reverberation effects on sub-band speech envelopes and posing the dereverberation problem as a gain estimation problem.
- Dereverberation of the autoregressive estimates of the sub-band envelope using a CLSTM model followed by feature extraction for ASR.
- Joint learning of the dereverberation model parameters and the acoustic model for ASR in a single neural pipeline.
- Illustrating the performance benefits of the proposed approach for multiple ASR tasks.

We use FDLP features (Purushothaman et al., 2020) for far-field speech. This paper extends the prior work done in (Purushothaman et al., 2020b) by proposing a joint neural dereverberation which forms an elegant neural learning framework. Further, several ASR experiments with the joint modeling approach are also conducted in this work.

3. Sub-band envelopes — Effect of reverberation and autoregressive estimation

We present the signal model for reverberation and the autoregressive model for estimating the sub-band envelopes (Ganapathy, 2017; Ganapathy et al., 2009).

3.1. Signal model

When speech is recorded in far-field reverberant environment, the data collected in the microphone is modeled as

$$r(t) = x(t) * h(t), \quad (1)$$

where $x(t)$, $h(t)$ and $r(t)$ denote the clean speech signal, the room impulse response and the reverberant speech respectively. The room response function $h(t) = h_e(t) + h_l(t)$, where $h_e(t)$ and $h_l(t)$ represent the early and late reflection components.

Let $x_q(n)$, $h_q(n)$ and $r_q(n)$ denote the decimated sub-band clean speech, room-response and the reverberant speech signal respectively. Here $q = 1, \dots, Q$ denotes the sub-band index and n denotes the decimated time-index (frame). Assuming an ideal band-pass filtering we can write (using Eq. (1)),

$$r_q(n) = x_q(n) * h_q(n) \quad (2)$$

In the proposed model, we explore the modeling of the sub-band temporal envelopes. In order to extract the envelopes, the analytic signal based demodulation is proposed. The analytic representation of a real-valued signal is the complex signal consisting of the original function (real part) and the Hilbert transform (imaginary part). The negative frequency components of the analytic signal are zero-valued. By representing the real-valued functions in analytic domain, the extraction of the modulation components (like envelopes and carrier signals) is facilitated.

Now, the analytic signal of the sub-band signal $r_q(n)$ is denoted as $r_{aq}(n)$, where $r_{aq}(n) = r_q(n) + j\mathcal{H}[r_q(n)]$. Here, $\mathcal{H}[\cdot]$ is the Hilbert operator. It can be shown that (Thomas et al., 2008; Ganapathy and Hermansky, 2012),

$$r_{aq}(n) = \frac{1}{2}[x_{aq}(n) * h_{aq}(n)], \quad (3)$$

If two signals have a modulating envelope on the same modulating sinusoidal carrier signal (single AM-FM signal), the convolution operation of the two signals will have an envelope which is the convolution of the two envelopes, i.e., the envelope of the convolution of the two signals is the convolution of the envelope of the signals. For sub-band speech signals, this envelope convolution model will form a good approximation if the sub-band signals are narrow-band.

Then, for band-pass filters with narrow band-width, we get the following approximation between the sub-band envelope (defined as the magnitude of the analytic signal) components of the reverberant signal and those of the clean speech signal.

$$m_{rq}(n) \simeq \frac{1}{2}m_{xq}(n) * m_{hq}(n), \quad (4)$$

where $m_{rq}(n)$, $m_{xq}(n)$, $m_{hq}(n)$ denote the sub-band envelopes of reverberant speech, clean speech and room response respectively. We can further split the envelope into early and late reflection coefficients.

$$m_{rq}(n) = m_{rqe}(n) + m_{rql}(n), \quad (5)$$

3.2. Autoregressive modeling of sub-band envelopes

Frequency domain linear prediction (FDLP) is the frequency domain dual of the conventional time domain Linear Prediction (TDLP). Just as the TDLP estimates the spectral envelope of a signal, FDLP estimates the temporal envelope of the signal (Ganapathy and Peddinti, 2018), i.e. square of its Hilbert envelope (Bedrosian, 1962). The Hilbert envelope is given by the inverse Fourier transform of the auto-correlation function of discrete cosine transform (DCT) (Ganapathy, 2012; Athineos and Ellis, 2003).

We use the auto-correlation of the DCT coefficients to model the temporal envelope of the signal. The autoregressive (AR) modeling property of linear prediction implies that the model preserves the peak location of the signal (which tend to be more robust in the presence of noise and reverberation) (Ganapathy et al., 2014). For the FDLP model, the sub-band AR model tries to preserve the peaks in temporal envelope (Ganapathy et al., 2009).

Let $x(t)$ denote an N -point discrete sequence. The type-I odd DCT (Martucci, 1994) $y[k]$ for $k = 0, 1, \dots, N-1$ is given by,

$$y[k] = \sum_{t=0}^{N-1} c_{t,k} x(t) \cos\left(\frac{2\pi tk}{M}\right) \quad (6)$$

where $c_{t,k} = 1$ for $t, k > 0$ and $c_{t,k} = \frac{1}{2}$ for $t, k = 0$ and $c_{t,k} = \frac{1}{\sqrt{2}}$ for the values of t, k where only one of the index is 0 and $M = 2N - 1$.

An even symmetric version of the input signal $x(t)$ is the signal $q(t)$ of length $M = 2N - 1$ where,

$$\begin{aligned} q(t) &= x(t), \quad t = 0..N-1 \\ q(t) &= x(M-t), \quad t = N..M \end{aligned}$$

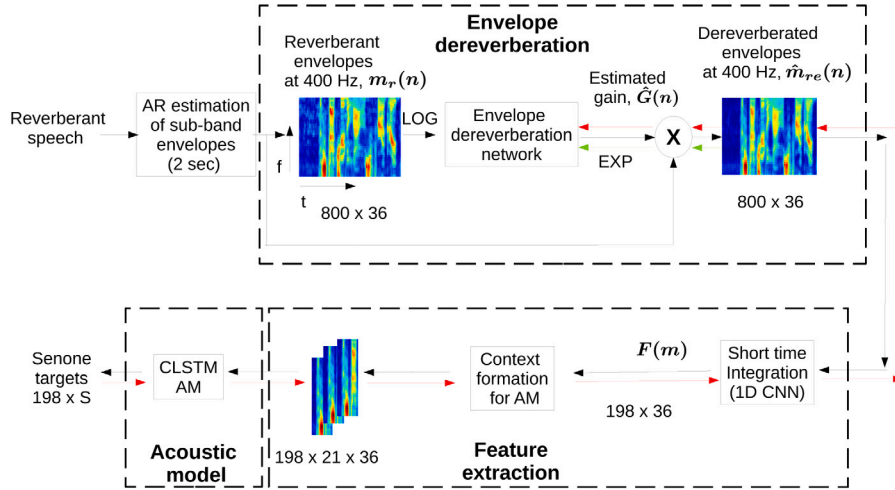


Fig. 1. Block schematic of envelope dereverberation model, the feature extraction module and acoustic model (Here $m_r(n)$, $\hat{G}(n)$, $\hat{m}_{re}(n)$ and $F(m)$ are given by Eqs. (5), (8)–(11)), the subscript q is dropped to indicate the fact that, signals from all the bands are considered). The entire model can be constructed as an end-to-end neural framework. The black arrows denote the forward pass, the red arrows represent backward propagation with ASR loss (E_{CE}), and green arrows denote the backward propagation with mean square error loss (E_{MSE}). Here, S is the total number of senone targets. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The analytic signal of a discrete time sequence can be defined using the one-sided discrete Fourier transform (DFT) (Ganapathy, 2012). The analytic signal $q_a(t)$ of the even-symmetric signal $q(t)$ can be shown to be (Ganapathy et al., 2009) the zero-padded DCT (up to scale) $\hat{y}[k]$, where $\hat{y}[k] = y[k]$, $k = 0..N - 1$ and $\hat{y}[k] = 0$, $k = N, \dots, M$.

Further, it can be shown that (Ganapathy, 2012), the auto-correlation of the zero-padded DCT signal $\hat{y}[k]$ and the squared magnitude of the analytic signal (Hilbert envelope) of the even-symmetric signal $|q_a(t)|^2$ are Fourier transform pairs (Ganapathy and Peddinti, 2018). Hence, the application of linear prediction on the zero-padded DCT signal yields the AR model of the Hilbert envelope of signal.

Let the linear prediction coefficients obtained from the zero-padded DCT signal be denoted as $\{a_k\}_{k=0}^p$, where p is the order of the LP. The FDLP model for the envelope is given by,

$$E(n) = \frac{\sigma}{|\sum_{k=0}^p a_k e^{-2\pi i k n}|^2} \quad (7)$$

where σ denotes the LP gain. The envelope estimated in above equation represents the autoregressive model of the temporal envelopes. Note that, when the model is applied on sub-band DCT coefficients, the envelope estimated will be the sub-band temporal envelope.

In this work, the sub-band envelopes of speech in mel-spaced bands are estimated using FDLP. Specifically, the discrete cosine transform (DCT) of sub-band signal $r_q(t)$ is computed and a linear prediction (LP) is applied on the DCT components. The LP envelope estimated using the prediction on the DCT components provides an all-pole model of the sub-band envelopes $m_{rq}(n)$.

4. Envelope dereverberation and joint modeling

The proposed framework (Fig. 1), consists of three modules, (i) envelope dereverberation, (ii) feature extraction and (iii) ASR acoustic model.

4.1. Neural dereverberation network

As seen in Eq. (5), the FDLP envelope of reverberant speech can be expressed as sum of the direct component (early reflection) and those with the late reflection. In the envelope dereverberation model, our aim is to input the envelope of the reverberant sub-band temporal envelope $m_{rq}(n)$ to predict the late reflection components $m_{rl}(n)$. Once this prediction is achieved, the late reflection component can be subtracted from the sub-band envelope to suppress the artifacts of reverberation. A similar analogy to this envelope subtraction approach is the spectral subtraction model where the noise and clean power spectral density (PSD) gets added in noisy speech PSD. If Gaussian assumptions are made for PSD components (Martin, 2005), the Wiener filtering approach to noisy speech enhancement provides the minimum mean squared error, where the noisy PSD is multiplied by the gain of the filter. In a similar manner, we pose the dereverberation problem as an envelope gain estimation problem.

The envelope gain (G_q) is defined as,

$$G_q(n) = \frac{\hat{m}_{r_{qe}}(n)}{\hat{m}_{r_{qe}}(n) + \hat{m}_{r_{ql}}(n)} \quad (8)$$

The gain $G_q(n)$ is estimated using the input sub-band envelope $m_{r_q}(n)$. With the gain estimate, the dereverberated envelope can be computed as,

$$\hat{m}_{r_{qe}}(n) = G_q(n)m_{r_q}(n) \quad (9)$$

The product model of enhancement is inspired by Wiener filtering principles. This sub-band envelope gain estimation is achieved using a deep neural network model in the proposed work. Following the model training, the dereverberation is achieved by multiplying the estimated sub-band envelope gain with the sub-band envelope of reverberant speech.

The block schematic of the envelope dereverberation model is shown in Fig. 1. The input to the dereverberation model is the FDLP sub-band envelope of the reverberant speech. The model is trained to learn the sub-band envelope gain which is the ratio of the clean envelopes (direct component) with the reverberant envelopes. During the model training, the model inputs are either far-field microphone recordings or the simulated reverberant recordings. The model targets are the envelope gain (Eq. (8)) computed using either the close talking/near-room microphone corresponding to the far-field microphone data, or the clean close-talking microphone data for the simulated reverberant training data. Thus, model is trained with paired data to estimate the gain.

As the envelopes and the gain parameters are positive in nature, the model implementation in the neural architecture uses a logarithmic transform at the input and the estimated gain is transformed by an exponential operation. Specifically, the input to the dereverberation model is the set of sub-band envelopes $\{\log(m_{r_q}(n))\}_{q=1}^Q$, where Q is the number of sub-bands. The model is trained to predict the log-gain $\{\log(G_q)\}_{q=1}^Q$. The sub-band dereverberated envelope is,

$$\hat{m}_{r_{qe}} = \exp[(\log(\hat{G}_q(n))) + \log(m_{r_q}(n))] \quad (10)$$

where $\hat{G}_q(n)$ is the estimate of the gain from the model.

The entire model developed in Section 3 is applicable only on long analysis windows (which are typically greater than the T60 of the room response function). Hence, the proposed approach operates on long temporal envelopes of the order of 2 s duration. From the reverberant speech and the corresponding clean speech, the FDLP sub-band envelopes corresponding to 2 non-overlapping segments are extracted. If the input sampling rate is 16 kHz, a 2s segment will correspond to 32k samples ($t = \{1..32000\}$). The FDLP envelopes are extracted at a sampling rate of 400 Hz. Thus, 2s segment of audio corresponds to 800 envelope samples ($n = \{1..800\}$) for each sub-band.

The input 2-D data of sub-band envelopes (800 samples from 36 mel sub-bands) are fed to a set of convolutional layers where the first two layers have 32 filters each with kernels of size of 41×5 . The next two CNN layers have 64 filters with 21×3 kernel size. All the CNN layer outputs with ReLU activations are zero padded to preserve the input size and no pooling operation is performed. The output of the CNN layers are reshaped to perform time domain recurrence using 3 layers of LSTM cells. The first two LSTM layers have 1024 cells while the last layer has 36 cells corresponding to the size of the target signal (envelope gain). The training criteria is based on the mean square error between the target and predicted output. The model is trained with stochastic gradient descent using Adam optimizer (Kingma and Ba, 2014).

4.2. Feature extraction and acoustic modeling

For feature extraction, the sub-band envelopes are integrated in short Hamming shaped windows of size 25 ms with a shift of 10 ms (Ganapathy, 2012). A 25 ms slice corresponds at 400 Hz sampling (FDLP envelopes are sampled at 400 Hz) to 10 samples and the hop size of 10 ms corresponds to 4 samples.

The windowed FDLP envelopes are multiplied with a Hamming shaped window (size of 10) and accumulated. This window is shifted by 4 samples. A log compression is applied to limit the dynamic range of values. Following this integration, a 2 Schunk of 800×36 sub-band FDLP envelopes becomes 198×36 .

In particular, let $\hat{m}_{r_{qe}}(n)$ denote the dereverberated sub-band envelope obtained using Eq. (10). Further, let $w(n)$ denote a Hamming window of size 10 (corresponding to 25 ms at 400 Hz sampling). Then, the features for ASR are extracted as,

$$F_q(m) = \log(\hat{m}_{r_{qe}}(m) * w(m)) \quad (11)$$

where $*$ is the convolution operation, and F_q denotes the scalar feature of q th sub-band. Here, m denotes the feature frame index at 10 ms sampling (100 Hz). The features for all the Q sub-bands are spliced to form the final feature vector for ASR model training.

The set of operations described above for short-term integration can be implemented as a 1-D CNN layer with a fixed Hamming shaped kernel size of 10 and a stride 4. A log non-linearity is applied on the convolution output.

The integrated envelopes are used as time–frequency representations for ASR training. A context of 21 frames, with 10 frames on the left and 10 frames on the right is used in the acoustic model training.

4.3. Acoustic model

The architecture of the acoustic model is based on convolutional long short term memory (CLSTM) networks (Fig. 1). The acoustic model corresponds to 2-D CLSTM network described in Purushothaman et al. (2020), consisting of 4 layers of CNN, a layer of LSTM with 1024 units performing recurrence over frequency and 3 fully connected layers with batch normalization.

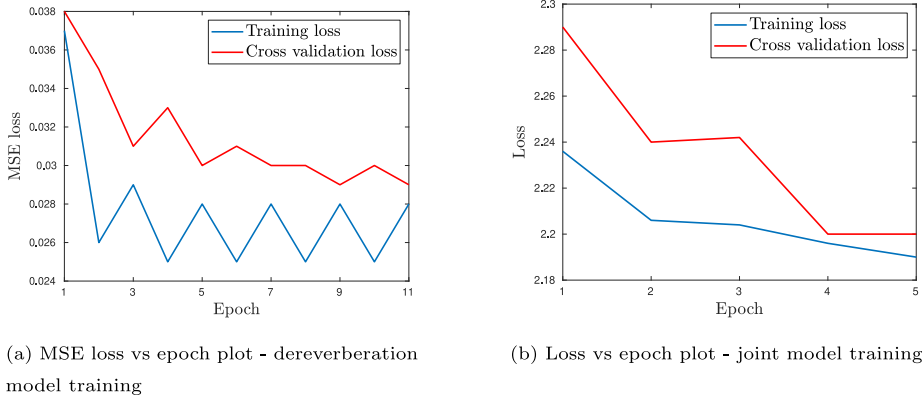


Fig. 2. Variation of training loss and cross validation loss with training epochs. The envelope dereverberation model and the ASR model are pre-trained before the joint learning step.

4.4. Joint learning

As shown in Fig. 1, the three modules of (i) envelope dereverberation, (ii) feature extraction and context formation and (iii) the ASR acoustic modeling can be combined into a single neural end-to-end framework.¹ The intermediate envelope integration step is implemented as a 1-layer of 1-D convolutions with Hamming shaped kernel and log non-linearity. The context creation for acoustic features in the given segment is also performed as a fixed 1-D convolution layer. In this manner, the entire processing pipeline can be performed using an elegant joint learning approach.

For generating mini-batches in the model training, a 2s speech segment is read along with the corresponding frame level targets (198 frames of senone targets for the 2s segment). The entire joint neural network is trained using a combination of ASR cross entropy training criterion and mean square loss between the clean and reverberant envelopes. The gradients from the ASR loss at the input of the acoustic model (computed for each senone target) is accumulated over all the frames in the given 2s segment. This accumulated gradient is of size 198×36 which corresponds to the size of the integrated envelopes. This ASR loss function when further back-propagated through fixed 1-D CNN layer provides a gradient matrix of size 800×36 . The gradient w.r.t. mean square error (MSE) between the target envelopes and the dereverberation model outputs is combined with the ASR based gradient for training the joint model. The two gradients are indicated by two different backward arrows in Fig. 1.

4.4.1. Joint loss function

The separate dereverberation model is trained to minimize the mean square error loss, E_{MSE} , which is the squared error between the reverberant envelope and the clean counter part. For joint training, we have two loss functions, one is the mean square error loss, E_{MSE} for a mini-batch and the cross-entropy loss, E_{CE} between the senone targets and the corresponding posteriors for the same mini-batch. We use a combination of these two losses. Thus the final joint loss, E_{Total} is given by,

$$E_{Total} = E_{CE} + \mu \times E_{MSE}, \quad (12)$$

where μ is a regularization parameter, which decides the share of E_{MSE} in the joint loss, E_{Total} . In all our ASR experiments, we have used regularization parameter $\mu = 0.4$. The absolute value of the two loss functions (different dynamic range in Fig. 2) does not have an impact as the model is trained with the gradient of the losses. Note that the MSE loss changes by 0.014 over the course of joint training while the combined loss changes by 0.04, so they generally have comparable dynamic ranges, even if these occur at different offsets from 0. The regularization constant μ controls the trade-off between the two loss functions.

The variation of the MSE loss in the envelope dereverberation network is shown in Fig. 2. The joint loss function on the training and validation data is also shown in this Figure. While the MSE loss trained with a higher learning rate exhibits oscillatory behavior, the joint loss function is relatively smooth. The final joint model is used in our ASR experiments.

A visualization of the dereverberation, achieved for the sub-band envelope of one single sub-band (10 th mel-band), is shown in Fig. 3. The sub-band envelopes of reverberant signal deviate from their clean signal counterparts (as explained in Section 3). Using the dereverberation model proposed in this paper, we find that the FDLP envelopes are more closely matched with the clean signal envelopes. In Section 5, we compare the performance of the CLSTM acoustic model architecture with other model architectures for dereverberation.

¹ The implementation of the work can be found in https://github.com/iiscleap/FDLP_Envelope_Dereverberation.




Fig. 3. Comparison of temporal envelopes, FDLP envelopes for clean, reverberant speech and reverberant speech after proposed joint learning based dereverberation, recordings from the REVERB Challenge dataset.

Table 1

Word Error Rate (%) in REVERB dataset for different features and proposed dereverberation method. Here prop denotes proposed work in this paper.

Model Features	Dev			Eval		
	Real	Simu	Avg	Real	Simu	Avg
BF-FBANK	19.1	6.1	12.6	14.7	6.5	10.6
BF-FDLP (Purushothaman et al., 2020)	17.8	6.8	12.3	14.0	7.0	10.5
BF-FBANK + CLSTM derevb. (prop)	17.3	5.5	11.4	13.1	6.9	10.0
BF-FBANK + spectral mapping derevb.(Han et al., 2015)	15.8	5.2	10.5	12.8	6.7	9.7
BF-FBANK + context aware derevb.(Santos and Falk, 2018)	19.6	6.9	13.2	17.5	9.0	13.2
BF-FBANK + end to end derevb.(Wu et al., 2017)	–	–	–	24.8	7.9	16.4
BF-FDLP + CLSTM derevb. (prop)	16.3	5.6	10.9	13.4	7.1	10.2
BF-FDLP + CLSTM derevb. + joint (prop)	15.2	5.6	10.4	12.1	7.1	9.6

5. Experiments and results

The experiments are performed on REVERB challenge (Kinoshita et al., 2013) and CHiME-3 (Barker et al., 2015) datasets. For the baseline model, we use WPE enhancement (Nakatani et al., 2010) along with unsupervised GEV beamforming (Kumar et al., 2020). This signal is processed with filter-bank energy features (denoted as BF-FBANK). The FBANK features are 36 band log-mel spectrogram with frequency range of 200–6500 Hz. This is the same frequency decomposition used in the FDLP and FDLP-dereverberation experiments. The acoustic model is the 2-D CLSTM network described in Purushothaman et al. (2020).

5.1. ASR framework

We use the Kaldi toolkit (Povey et al., 2011) for deriving the senone alignments used in the PyTorch deep learning framework for acoustic modeling. A hidden Markov model–Gaussian mixture model (HMM–GMM) system is trained with MFCC (Mel Frequency Cepstral Coefficients) features (Logan et al., 2000) to generate the alignments for training the CLSTM acoustic model. A tri-gram language model (Brown et al., 1992) is used in the ASR decoding and the best language model weight obtained from development set is used for the evaluation set.

5.2. REVERB challenge ASR

The REVERB challenge dataset (Kinoshita et al., 2016) for ASR consists of 8 channel recordings with real and simulated reverberation conditions. The simulated data is comprised of reverberant utterances generated (from the WSJCAM0 corpus (Robinson et al., 1995)) by artificially convolving clean WSJCAM0 recordings with the measured room impulse responses (RIRs) and adding noise at an SNR of 20 dB. The simulated data has six different reverberation conditions. The real data, which is comprised of utterances from the MC-WSJ-AV corpus (Lincoln et al., 2005), consists of utterances spoken by human speakers in a noisy reverberant room. The training set consists of 7861 utterances from the clean WSJCAM0 training data by convolved with 24 measured RIRs.

5.2.1. Discussion

Table 1 shows the WER results for experiments on REVERB challenge dataset. The WPE along with unsupervised GEV beamformed signal is used for all the ASR experiments (denoted as BF). The BF-FDLP baseline by itself is better than the BF-FBANK baseline (average relative improvements of 2% on the development set and about 1% on the evaluation set). For a fair comparison of the proposed approach, we have applied a similar dereverberation method on BF-FBANK baseline. Here, we have trained the neural model with log-mel features corresponding to 2 duration with all the 36 mel-bands jointly. This approach is denoted as BF-FBANK + CLSTM derevb. (prop). Average relative improvements of 10% on the development set and about 6% on the evaluation set is achieved compared to the BF-FBANK baseline.

BF-FBANK + spectral mapping derevb. Han et al. (2015) corresponds to the work by Kun Han et al. Here a 3-layer deep neural network of 2570 units is used as the dereverberation neural model. The network is fed with 257-dimensional log-magnitude STFT

Table 2

Word Error Rate (%) in CHiME-3 dataset for different features and proposed dereverberation method. Here prop denotes proposed work in this paper.

Model Features	Dev			Eval		
	Real	Simu	Avg	Real	Simu	Avg
BF-FBANK	7.8	8.0	7.9	14.0	9.7	11.8
BF-FDLP	7.0	8.1	7.5	12.0	10.0	11.0
BF-FBANK + CLSTM derevb. (prop)	7.2	8.3	7.7	12.9	9.8	11.4
BF-FBANK + spectral mapping derevb.(Han et al., 2015)	8.0	10.0	9.0	14.3	12.3	13.3
BF-FBANK + context aware derevb.(Santos and Falk, 2018)	7.7	9.9	8.8	13.4	13.3	13.3
BF-FDLP + CLSTM derevb. (prop)	7.2	7.9	7.5	13	9.6	11.3
+ spec. reg. (prop)	6.9	8.0	7.4	11.8	9.8	10.8
+ spec. reg. + joint (prop)	7.0	7.7	7.3	11.7	9.3	10.5

Table 3

WER in VOICES dataset for different features and proposed dereverberation method. Here prop denotes proposed work in this paper.

Model Architecture	Dev	Eval
BF-FBANK	55.5	66.6
BF-FDLP	51.5	62.6
BF-FDLP + CLSTM derevb. (prop)	52.8	62.4
+ joint. (prop)	49.9	59.8

features from a frame of 25 ms. A context window of 10-frames (5-left and 5-right) is selected and the network tries to predict the central frame. The work by Santos et al. is implemented as BF-FBANK + context aware derevb. Santos and Falk (2018). A CNN-GRU based encoder–decoder model sees the entire utterance at the 257-STFT magnitude level features and trained to predict the clean utterance. The results for end-to-end dereverberation network (joint learning) proposed in Wu et al. (2017) is also compared with the proposed work in Table 1.

Finally applying the proposed neural model based dereverberation on BF-FDLP baseline (denoted as BF-FDLP + CLSTM derevb. (prop)) yields average relative improvements of 13% on the development set and about 4% on the evaluation set, compared to the BF-FBANK baseline. After joint training this further improves to 17% and 9% respectively. The improvement in real condition is much more than that of simulated data. Average relative improvements of 20% on the real development set and about 18% on the real evaluation set, compared to the BF-FBANK baseline is achieved by the proposed method. This suggests that, even though the jointly learned neural model is trained only with simulated reverberation, it generalizes well on unseen real data.

5.3. CHiME-3 ASR

The CHiME-3 dataset (Barker et al., 2015) for the ASR has multiple microphone tablet device recording in four different environments, namely, public transport (BUS), cafe (CAF), street junction (STR) and pedestrian area (PED). For each of the above environments, real and simulated data are present. The real data consists of 6 channel recordings from WSJ0 corpus sampled at 16 kHz spoken in the four varied environments. The simulated data was constructed by mixing clean utterances with the environment noise. The training dataset consists of 1600 (real) noisy recordings and 7138 (simulated) noisy recordings from 83 speakers.

5.3.1. Discussion

The WER results for experiments on CHiME-3 dataset are shown in Table 2. The FDLP baseline, denoted as BF-FDLP is better than the FBANK baseline (BF-FBANK). We observe average relative improvements of 8% on the development set and about 12% on the evaluation set when comparing BF-FDLP and BF-FBANK baseline systems. It can also be seen from Table 2 that the proposed dereverberation method improves the FBANK-baseline system. The results based on the implementation of works done by Han et al. (2015) and Santos and Falk (2018) degrade the word error rates further compared to the BF-FBANK baseline.

In the CHiME-3 dataset, we observed that the significant cause of degradation in the signal quality came from the additive noise sources. On further investigation, we found that the dereverberation model also resulted in smoothing of the spectral variations in the FDLP spectrogram. In order to circumvent this issue, we regularized the MSE loss with a term that encouraged the spectral channels to be uncorrelated. The regularization parameter was kept at 0.05. Using this regularized MSE loss, we further improved the BF-FDLP-Dereverberation system results over the dereverberation approach with MSE loss alone. These experiments suggest that even when the audio data does not have significant late reflection components (like CHiME-3 dataset), the proposed approach improves significantly over the baseline method (average relative improvements of 10.3% over the baseline BF-FBANK system in the real dev condition and 23.5% on real eval condition).

Table 4
Word Error Rate (%) in REVERB dataset using different model architectures for dereverberation(without joint training).

Architecture	Dev			Eval		
	Real	Simu	Avg	Real	Simu	Avg
BF-FDLP (Purushothaman et al., 2020)	17.8	6.8	12.3	14.0	7.0	10.5
Neural Dereverberation						
3 layer DNN	17.3	5.4	11.3	14.2	6.9	10.5
5 layer CNN	16	5.9	11	13.5	7.2	10.3
4 layer CNN + 3 layer DNN	17.9	5.6	11.7	14.4	6.7	10.5
7 layer LSTM (1024 units each)	17	5.3	11.1	14.2	7.4	10.8
7 layer Resnet	17.4	7.9	12.6	14.8	10.3	12.5
CNN + DNN + LSTM [2,2,3]	19.1	6.8	13.0	15.5	7.9	11.7
CLSTM (4-CNN + 3-LSTM)	16.3	5.6	10.9	13.4	7.1	10.2

Table 5
WER for various regularization to alleviate spectral smearing CHiME3 dataset. The regularization term is the cross correlation of the spectral bands. In absence of reverberation, alleviating spectral smearing improves the WER.

Regularizer weight, λ	Dev			Eval		
	Real	Simu	Avg	Real	Simu	Avg
$\lambda = 0.3$	7	8.1	7.5	12.3	9.9	11.1
$\lambda = 0.1$	7	8.2	7.6	12.5	10	11.2
$\lambda = 0.05$	6.9	8	7.4	11.8	9.8	10.8
$\lambda = 0.02$	7.2	8.4	7.8	12.5	10	11.2
$\lambda = 0$	7.2	7.9	7.5	13	9.6	11.3
BF-FBANK	7.8	8.0	7.9	14.0	9.7	11.8

5.3.2. VOICES corpus ASR

Since the REVERB challenge dataset and CHiME-3 dataset are relatively smaller datasets, we wanted to establish the efficacy of the proposed dereverberation method in a larger dataset. Thus we experimented with VOICES challenge dataset. VOICES corpus (Richey et al., 2018) is released as part of ‘‘The voices from a distance challenge 2019’’ (Nandwana et al., 2019) of Interspeech 2019. For the ASR fixed conditions track, the training set consists of 80-hours subset of LibriSpeech corpus (Panayotov et al., 2015). The training set has close talking microphone recordings from 427 different speakers from quiet environment. The development and evaluation sets consists of 19 h and 20 h of distant microphone recordings of varying room, environment and noise conditions. The significant difference between the training set and development/evaluation set makes the challenge even more difficult. We have used the same acoustic model configurations and hence these results reflect the true acoustic mismatch condition in ASR.

5.3.3. Discussion

The WER results for VOICES corpus is given in Table 3. As seen, the baseline FDLP, denoted by BF-FDLP, provides a better WER compared to the baseline FBANK, denoted as BF-FBANK. This is further improved with joint learning based dereverberation. The final WER shows improvement in both development set and evaluation set. A relative WER improvement of 10% in both development set and evaluation set over the baseline FBANK system is observed in these experiments.

6. Analysis

In this section the effect of different neural network architectures and various parameters like regularization parameter, λ , FDLP model order, p on WER are reported in Tables 4–6 and Fig. 4.

6.1. Architecture of dereverberation model

Table 4 shows the WER for different neural network architectures. We initially explore a DNN of three feed forward layers. A slight improvement in development set is seen over the FDLP baseline, BF-FDLP. The relative improvement in WER becomes appreciable as we move to 5 layer CNN. The architecture with LSTM alone is promising. We also explore a Resnet (He et al., 2016) style architecture which was successful in image recognition. A combination of CNN, DNN and LSTM did not perform well compared to baseline. Finally the CNN + LSTM combination provides the best WER.

6.2. Spectral correlation loss

As reported in Table 2 on CHiME-3 dataset, an extra loss function which encourages the spectral bands to be uncorrelated improves the ASR performance on noisy data when the data is corrupted by additive noise with minimal reverberation artifacts. Tables 5 and 6 shows the effect of the regularization weight, λ on WER in CHiME-3 and REVERB datasets respectively for the

Table 6

WER for various regularization to alleviate spectral smearing REVERB dataset. The regularization term is the cross correlation of the spectral bands. In the presence of significant reverberation, the extra regularization did not improve the performance.

Regularizer weight, λ	Dev			Eval		
	Real	Simu	Avg	Real	Simu	Avg
$\lambda = 0.0$	16.3	5.6	10.9	13.4	7.1	10.2
$\lambda = 0.05$	16.6	5.5	11.0	14	6.9	10.4
$\lambda = 0.1$	16.9	5.6	11.2	13.7	6.9	10.3
$\lambda = 0.2$	17.2	5.5	11.3	14.2	6.8	10.5
BF-FBANK	19.1	6.1	12.6	14.7	6.5	10.6




Fig. 4. WER (%) for various model-order, p in FDLP model for REVERB dataset.

spectral correlation loss used in the model learning. The introduction of the spectral correlation loss improves the WER in CHiME-3 dataset. The best results are obtained for a choice of $\lambda = 0.05$.

The introduction of spectral correlation loss does not benefit the REVERB challenge dataset. We hypothesize that this may be due to the more dominant effect of temporal smearing seen in the REVERB challenge dataset. For the experiments on the VOICES corpus, the spectral correlation loss is not used.

6.3. Choice of FDLP model order

Fig. 4 shows the effect of model order, p used in the FDLP envelope estimation on the WER for the REVERB challenge dataset. The model order p is the number of “past” samples used in the auto-regressive modeling of the sub-band DCT signal for a 2s window. While the WER results on the simulated conditions improve with higher model order of the FDLP, the performance on the real conditions is observed to be the best for about 100 poles per 2s of audio in each sub-band. All the other experiments reported in the paper use the 100 poles per 2s window of the audio signal.

6.4. Discussion on performance gains

All the results reported in **Tables 1–3** use a strong baseline system with GEV based beamforming and weighted prediction error (WPE) based enhancement. Hence, we note that all systems use the same pre-processing pipeline and the gains observed over the baseline system are in addition to these enhancement steps. In addition, we also ensure that the baseline FBANK based system, neural enhancement methods explored in the past and the proposed approach have the same sub-band decomposition, feature normalization, acoustic model and language model settings. In this way, the results reported highlight the effectiveness of the proposed work in suppressing reverberation distortions.

The methods proposed previously based on neural enhancement and dereverberation improve the performance of the baseline system on the REVERB challenge dataset. However, as seen in **Table 2**, in the presence of additive noise conditions on the CHiME-3 dataset, most of these prior works degrade the performance compared to the BF-FBANK baseline system. In this regard, the method proposed in this paper provides significant performance improvements on all three datasets. Further, the results consistently highlight the performance gains of using the joint neural learning framework.

7. Summary

In this paper, we propose a new neural model for dereverberation of temporal envelopes and joint learning of the acoustic model to improve the ASR cost. The joint learning framework combines the envelope dereverberation framework, feature pre-processing and acoustic modeling into a single neural pipeline. This framework is hence elegant and the model can be learned using a joint loss function. Using the proposed neural dereverberation approach and joint learning, we perform speech recognition experiments on the REVERB challenge dataset as well as on the CHiME-3 dataset. These experiments indicate that the proposed neural dereverberation approach generalizes well on unseen conditions. The analysis of results also highlight the incremental benefits achieved for different choice of hyper-parameters and model architecture settings. The application of the proposed approach for large vocabulary speech recognition experiments on VOICES dataset further emphasizes the performance benefits.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

This work was funded by grants from Samsung Research India, Bangalore.

References

- Anguera, X., Wooters, C., Hernando, J., 2007. Acoustic beamforming for speaker diarization of meetings. *IEEE Trans. Audio Speech Lang. Process.* 15 (7), 2011–2022.
- Athineos, M., Ellis, D.P., 2003. Frequency-domain linear prediction for temporal features. *IEEE Trans. Audio Speech Lang.*
- Barker, J., Marxer, R., Vincent, E., Watanabe, S., 2015. The third 'CHiME' speech separation and recognition challenge: Dataset, task and baselines. In: *IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pp. 504–511.
- Bedrosian, E., 1962. The analytic signal representation of modulated waveforms. *Proc. IRE* 50 (10), 2071–2076.
- Brown, P.F., Desouza, P.V., Mercer, R.L., Pietra, V.J.D., Lai, J.C., 1992. Class-based n-gram models of natural language. *Comput. Linguist.* 18 (4), 467–479.
- Chen, Z., Watanabe, S., Erdogan, H., Hershey, J.R., 2015. Speech enhancement and recognition using multi-task learning of long short-term memory recurrent neural networks. In: *Sixteenth Annual Conference of the International Speech Communication Association*.
- Ganapathy, S., 2012. *Signal Analysis Using Autoregressive Models of Amplitude Modulation*. (Ph.D. thesis). Johns Hopkins University.
- Ganapathy, S., 2017. Multivariate autoregressive spectrogram modeling for noisy speech recognition. *IEEE Signal Process. Lett.* 24 (9), 1373–1377.
- Ganapathy, S., Harish, M., 2018. Far-field speech recognition using multivariate autoregressive models. In: *Interspeech*. pp. 3023–3027.
- Ganapathy, S., Hermansky, H., 2012. Temporal resolution analysis in frequency domain linear prediction. *J. Acoust. Soc. Am.* 132 (5), EL436–EL442.
- Ganapathy, S., Mallidi, S.H., Hermansky, H., 2014. Robust feature extraction using modulation filtering of autoregressive models. *IEEE/ACM Trans. Audio Speech Lang. Process.* 22 (8), 1285–1295.
- Ganapathy, S., Motlicek, P., Hermansky, H., 2009. Autoregressive models of amplitude modulations in audio compression. *IEEE Trans. Audio Speech Lang. Process.* 18 (6), 1624–1631.
- Ganapathy, S., Peddinti, V., 2018. 3-D CNN models for far-field multi-channel speech recognition. In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5499–5503.
- Ganapathy, S., Peddinti, V., 2018b. 3-D CNN models for far-field multi-channel speech recognition. In: *2018 IEEE International Conference on Acoustics, Speech and Signal Processing. ICASSP, IEEE*, pp. 5499–5503.
- Hain, T., Burget, L., Dines, J., Garner, P.N., Grézl, F., El Hannani, A., Huijbregts, M., Karafiat, M., Lincoln, M., Wan, V., 2012. Transcribing meetings with the AMIDA systems. *IEEE Trans. Audio Speech Lang. Process.* 20 (2), 486–498.
- Han, K., Wang, Y., Wang, D., 2014. Learning spectral mapping for speech dereverberation. In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4628–4632.
- Han, K., Wang, Y., Wang, D., Woods, W.S., Merks, I., Zhang, T., 2015. Learning spectral mapping for speech dereverberation and denoising. *IEEE/ACM Trans. Audio Speech Lang. Process.* 23 (6), 982–992.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778.
- Kingma, D.P., Ba, J., 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kinoshita, K., Delcroix, M., Gannot, S., Habets, E.A., Haeb-Umbach, R., Kellermann, W., Leutnant, V., Maas, R., Nakatani, T., Raj, B., et al., 2016. A summary of the REVERB challenge: state-of-the-art and remaining challenges in reverberant speech processing research. *EURASIP J. Adv. Signal Process.* 2016 (1), 7.
- Kinoshita, K., et al., 2013. The REVERB challenge: A common evaluation framework for dereverberation and recognition of reverberant speech. In: *IEEE WASPPA*. pp. 1–4.
- Kumar, R., Sreeram, A., Purushothaman, A., Ganapathy, S., 2020. Unsupervised neural mask estimator for generalized eigen-value beamforming based ASR. In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7494–7498.
- Lincoln, M., McCowan, I., Vepa, J., Maganti, H.K., 2005. The multi-channel wall street journal audio visual corpus (MC-WSJ-AV): Specification and initial experiments. In: *IEEE Workshop on Automatic Speech Recognition and Understanding, 2005*, pp. 357–362.
- Logan, B., et al., 2000. Mel frequency cepstral coefficients for music modeling. In: *ISMIR*. 270, pp. 1–11.
- Maas, A.L., O'Neil, T.M., Hannun, A.Y., Ng, A.Y., 2013. Recurrent neural network feature enhancement: The 2nd CHiME challenge. In: *Proceedings the 2nd CHiME Workshop on Machine Listening in Multisource Environments Held in Conjunction with ICASSP*, pp. 79–80.
- Martin, R., 2005. Speech enhancement based on minimum mean-square error estimation and supergaussian priors. *IEEE Trans. Speech Process.* 13 (5), 845–856.
- Martucci, S.A., 1994. Symmetric convolution and the discrete sine and cosine transforms. *IEEE Trans. Signal Process.* 42 (5), 1038–1051.
- Nakatani, T., Yoshioka, T., Kinoshita, K., Miyoshi, M., Juang, B.-H., 2010. Speech dereverberation based on variance-normalized delayed linear prediction. *IEEE Trans. Audio Speech Lang. Process.* 18 (7), 1717–1731.
- Nandwana, M.K., Van Hout, J., McLaren, M., Richey, C., Lawson, A., Barrios, M.A., 2019. The voices from a distance challenge 2019 evaluation plan. *arXiv preprint arXiv:1902.10828*.

- Panayotov, V., Chen, G., Povey, D., Khudanpur, S., 2015. Librispeech: an asr corpus based on public domain audio books. In: 2015 IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE, pp. 5206–5210.
- Pandey, A., Wang, D., 2019. A new framework for CNN-based speech enhancement in the time domain. *IEEE/ACM Trans. Audio Speech Lang. Process.* 27 (7), 1179–1188.
- Peddinti, V., Wang, Y., Povey, D., Khudanpur, S., 2018. Low latency acoustic modeling using temporal convolution and LSTMs. *IEEE Signal Process. Lett.* 25 (3), 373–377.
- Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., Hannemann, M., Motlicek, P., Qian, Y., Schwarz, P., et al., 2011. The kaldı speech recognition toolkit. In: *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding (CONF)*, IEEE Signal Processing Society.
- Purushothaman, A., Sreeram, A., Ganapathy, S., 2020. 3-D acoustic modeling for far-field multi-channel speech recognition. In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6964–6968.
- Purushothaman, A., Sreeram, A., Kumar, R., Ganapathy, S., 2020b. Deep learning based dereverberation of temporal envelopes for robust speech recognition. *arXiv preprint arXiv:2008.03339*.
- Richey, C., Barrios, M.A., Armstrong, Z., Bartels, C., Franco, H., Graciarena, M., Lawson, A., Nandwana, M.K., Stauffer, A., van Hout, J., et al., 2018. Voices obscured in complex environmental settings (voices) corpus. *arXiv preprint arXiv:1804.05053*.
- Robinson, T., Franssen, J., Pye, D., Foote, J., Renals, S., 1995. WSJCAMO: A British english speech corpus for large vocabulary continuous speech recognition. In: *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Vol. 1, pp. 81–84.
- Santos, J.F., Falk, T.H., 2018. Speech dereverberation with context-aware recurrent neural networks. *IEEE/ACM Trans. Audio Speech Lang. Process.* 26 (7), 1236–1246.
- Seltzer, M.L., Yu, D., Wang, Y., 2013. An investigation of deep neural networks for noise robust speech recognition. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 7398–7402.
- Thomas, S., Ganapathy, S., Hermansky, H., 2008. Recognition of reverberant speech using frequency domain linear prediction. *IEEE Signal Process. Lett.* 15, 681–684.
- Wang, D., Chen, J., 2018. Supervised speech separation based on deep learning: An overview. *IEEE/ACM Trans. Audio Speech Lang. Process.* 26 (10), 1702–1726.
- Wang, Z.-Q., Wang, D., 2016. A joint training framework for robust automatic speech recognition. *IEEE/ACM Trans. Audio Speech Lang. Process.* 24 (4), 796–806.
- Weninger, F., Erdogan, H., Watanabe, S., Vincent, E., Le Roux, J., Hershey, J.R., Schuller, B., 2015. Speech enhancement with LSTM recurrent neural networks and its application to noise-robust ASR. In: *International Conference on Latent Variable Analysis and Signal Separation*. Springer, pp. 91–99.
- Williamson, D.S., Wang, D., 2017. Time-frequency masking in the complex domain for speech dereverberation and denoising. *IEEE/ACM Trans. Audio Speech Lang. Process.* 25 (7), 1492–1501.
- Wöllmer, M., Zhang, Z., Weninger, F., Schuller, B., Rigoll, G., 2013. Feature enhancement by bidirectional LSTM networks for conversational speech recognition in highly non-stationary noise. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 6822–6826.
- Wu, B., Li, K., Ge, F., Huang, Z., Yang, M., Siniscalchi, S.M., Lee, C., 2017. An end-to-end deep learning approach to simultaneous speech dereverberation and acoustic modeling for robust speech recognition. *IEEE J. Sel. Top. Sign. Process.* 11 (8), 1289–1300.
- Wu, B., Li, K., Huang, Z., Siniscalchi, S.M., Yang, M., Lee, C.-H., 2017. A unified deep modeling approach to simultaneous speech dereverberation and recognition for the REVERB challenge. In: *2017 Hands-Free Speech Communications and Microphone Arrays. HSCMA*, IEEE, pp. 36–40.
- Xu, Y., Du, J., Dai, L.-R., Lee, C.-H., 2014. A regression approach to speech enhancement based on deep neural networks. *IEEE/ACM Trans. Audio Speech Lang. Process.* 23 (1), 7–19.
- Yoshioka, T., et al., 2012. Making machines understand us in reverberant rooms: Robustness against reverberation for automatic speech recognition. *IEEE Signal Process. Mag.* 29 (6), 114–126.
- Yu, D., Deng, L., 2016. *Automatic Speech Recognition*. Springer.
- Zhao, Y., Wang, D., Xu, B., Zhang, T., 2018. Late reverberation suppression using recurrent neural networks with long short-term memory. In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5434–5438.