Bridging the Gap Between Monaural Speech Enhancement and Recognition With Distortion-Independent Acoustic Modeling

Peidong Wang 🄾, *Student Member, IEEE*, Ke Tan 🄾, *Student Member, IEEE*, and De Liang Wang 🄾, *Fellow, IEEE*

Abstract—Monaural speech enhancement has made dramatic advances since the introduction of deep learning a few years ago. Although enhanced speech has been demonstrated to have better intelligibility and quality for human listeners, feeding it directly to automatic speech recognition (ASR) systems trained with noisy speech has not produced expected improvements in ASR performance. The lack of an enhancement benefit on recognition, or the gap between monaural speech enhancement and recognition, is often attributed to speech distortions introduced in the enhancement process. In this article, we analyze the distortion problem, compare different acoustic models, and investigate a distortion-independent training scheme for monaural speech recognition. Experimental results suggest that distortion-independent acoustic modeling is able to overcome the distortion problem. Such an acoustic model can also work with speech enhancement models different from the one used during training. Moreover, the models investigated in this paper outperform the previous best system on the CHiME-2 corpus.

Index Terms—Speech enhancement, speech recognition, speech distortion, distortion-independent acoustic modeling.

I. INTRODUCTION

FORMULATED as a supervised learning problem, speech enhancement has made major progress over the last few years with the use of data driven methods, particularly deep learning. Wang and Wang [31], [32] first introduced deep neural networks (DNNs) to perform time-frequency (T-F) masking for speech enhancement. Lu *et al.* and Xu *et al.* used a deep autoencoder (DAE) or DNN to map from the power spectrum of noisy speech to that of clean speech [14], [36], [37]. Many subsequent studies have been conducted to perform T-F masking or spectral mapping by employing a variety of deep learning models, acoustic features, and training targets [9], [13], [20], [34], [35], [38]. These studies have elevated the performance of

speech enhancement by a large margin [26]. DNN-based monaural speech enhancement has improved, for the first time, the intelligibility of noisy speech for human listeners with hearing impairment as well as those with normal hearing [10], [12], [26].

Along with the progress in speech enhancement, researchers have investigated using speech enhancement models as frontends for automatic speech recognition (ASR) systems. Narayanan *et al.* [18], [19] proposed to combine masking-based DNN speech enhancement with speech recognition. With a Gaussian mixture model (GMM) as backend, the enhancement frontend was shown to reduce word error rate (WER) significantly [18]. In a subsequent paper using DNN as backend, the benefit of speech enhancement is mixed, depending on training features [19]. For the acoustic model trained with cepstral features, speech enhancement still helps. With log-Mel features, however, the enhancement frontend causes performance degradation. Du *et al.* [4] applied mapping-based frontends to both GMM and DNN based recognition backends. Their observations are basically in line with those of Narayanan *et al.* The only difference is that their enhancement frontend can yield improvements on clean, noisy, and clean plus channel-mismatched conditions for the DNN acoustic model trained with noisy speech. In the fourth CHiME speech separation and recognition challenge (CHiME-4), Heymann *et al.* [11] noted that the harm of processing artifacts introduced during enhancement may outweigh the benefit brought by noise reduction. Based on these studies as well as our own attempts in applying monaural speech enhancement as a frontend for speech recognition on CHiME-4 corpus, the distortion to speech signals introduced in monaural speech enhancement is a major problem that can render enhancement useless or even harmful for robust ASR.

One way to alleviate the distortion problem is to reduce or eliminate speech distortions in enhancement frontends. Most papers in this direction use ASR information to guide speech enhancement model training. Erdogan *et al.* [6], [34] proposed to use the alignment information derived from decoded hidden Markov model (HMM) states. The mimic loss based method [1], [21] proposed by Bagchi *et al.* jointly trains enhancement frontends and recognition backends. It uses senone labels directly as the training target. Experimental results showed that such enhancement frontends can be effectively used with off-the-shelf ASR models in Kaldi [22] on the second CHiME speech separation and recognition corpus (CHiME-2). In Meng *et al.*'s study on applying adversarial training for speech enhancement [15],

senone level loss was also incorporated for better ASR performance. Gao *et al.* [8] proposed a progressive speech enhancement training scheme. It fine-tunes enhancement models in a multitask manner. Instead of using clean speech as the only target of the output layer, they add multiple layers in DNN treating speech with progressively decreased signal-to-noise ratios (SNRs) as labels. This way, the enhancement model is trained to reduce noise gradually, as well as the distortion in the output layer. Meng *et al.* [16] proposed cycle-consistent training for speech enhancement. It optimizes the noisy-to-clean mapping and a clean-to-noisy mapping simultaneously.

In addition to pursuing distortion reduction in speech enhancement models, designing more distortion tolerant acoustic model backends may be another direction. The joint training of speech enhancement frontends and ASR backends [4], [7], [33] would belong to this category. Acoustic model training using enhanced speech directly as acoustic model input has been explored in many previous studies [15], [16]. Distortion-independent acoustic modeling in our study differs from these methods in that we use a large variety of enhanced speech during training. Previous research in speech enhancement field shows that DNNs trained using a variety of noises have the ability to generalize to new noisy conditions [3]. A recent study performed by Narayanan *et al.* [17] investigated the generalization ability of acoustic models trained with various out-of-domain data (noises, bandwidths, codecs, and features). Their observation is that, through large-scale training, such acoustic models perform as well as acoustic models trained with in-domain data.

In this study, we investigate distortion tolerant acoustic model backends. In the context of the previous work that trains an acoustic model on enhanced speech, our study makes the following contributions. First, we analyze the distortion problem and view it as a noise mismatch between training and testing. After comparing six acoustic models, we find that the distortion-independent acoustic model, trained with a large variety of enhanced speech, is able to overcome the distortion problem. Second, we demonstrate that distortion-independent acoustic modeling can work with DNN based supervised speech enhancement frontends that are different from the one used in training. This is important for real-world applications as acoustic model retraining can thus be avoided when a different enhancement frontend is used. Distortion-independent training also allows distortion tolerant ASR backends and speech enhancement frontends to be trained separately. Finally, as a byproduct of this investigation, we advance the state-of-the-art performance on the CHiME-2 corpus.

Some preliminary results are presented in [27]. This paper expands [27] in several ways. First, we investigate the distortion-independent acoustic model on not only CHiME-2 but also Wall Street Journal (WSJ). Second, we add the noise-mismatched acoustic model and a joint-training based model as baseline systems. Finally, we evaluate the distortion-independent acoustic model on a minimum-mean square error (MMSE) based frontend.

The rest of this paper is organized as follows. Section II gives an analysis of the distortion problem, an explanation of distortion-independent acoustic modeling, and a description of utterance-wise recurrent dropout for acoustic model training. Sections III and IV present the experiment setup and results, respectively. We make concluding remarks in Section V.

II. SYSTEM DESCRIPTION

A. An Analysis on the Distortion Problem

The distortion in this study refers to the alteration to clean speech signal introduced by speech enhancement that may cause performance degradation in an ASR system. More specifically, this paper tackles with the distortion problem of noise-independent speech enhancement. The input to a speech enhancement system is generated by mixing clean speech with an additive noise, as shown below:

$$y = s + n \tag{1}$$

where y denotes noisy speech, s clean speech, and n an additive noise.

The frequency domain representation of Eq. (1) can be written as (2) below:

$$Y = S + N \tag{2}$$

where Y, S, and N are the spectral representations of noisy speech, clean speech, and additive noise, respectively.

Speech enhancement typically operates on the magnitudes of frequency domain representations. Masking-based models generate a T-F mask, which is then element-wise multiplied with the magnitude of Y,

$$|\hat{S}| = |Y| \otimes M = |S + N| \otimes M \tag{3}$$

where $|\cdot|$ denotes magnitude, $\otimes$ element-wise multiplication, $\hat{S}$ enhanced speech, and M mask.

Depending on the T-F mask definition, M is typically a real-valued matrix with element values ranging from zero to one, e.g. the IRM [30]. For such masks, (3) can be written as below:

$$|\hat{S}| = |S + N| \otimes M = |S \otimes M + N \otimes M| \tag{4}$$

Thus, we have

$$|\hat{S}| = |S + S \otimes (M - A) + N \otimes M| \tag{5}$$

where A is an all-one matrix.

The distortion for ASR backends can be defined as:

$$D = S \otimes (M - A) + N \otimes M = N \otimes M - S \otimes \overline{M} \tag{6}$$

where $\overline{M}$ denotes the complement of M.

There are two special cases of D. First, if M is an all-one matrix, speech enhancement has no impact on noisy speech. The influence of S on D can also be ignored. Second, let us consider the case when M equals the IRM defined below:

$$IRM = \frac{|S|}{|S| + |N|} \tag{7}$$

In this case, D will be an all-zero matrix, and the distortion problem does not exist.

Other than the two cases above, the influence of S cannot be ignored and the second term in (6), i.e. $N \otimes M$, can be viewed
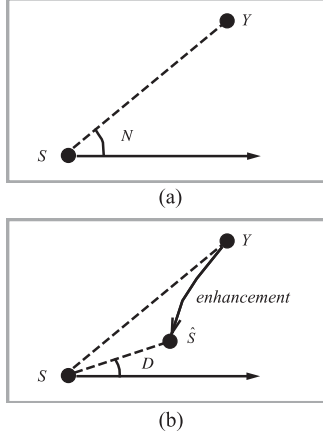
Fig. 1. Illustration of the signal distortion problem. (a) The polar coordinate system. (b) Clean, noisy, and enhanced speech.

as noise residue, which is different from N. Due to this residue, distortion is different from noise N.

Fig. 1 shows the deviation of D from N in an intuitive way. In this figure, spectral representations of different signals are plotted in a polar coordinate system. The center of the coordinate system denotes clean speech S. The distance between clean speech S and noisy speech Y indicates the intensity of noise, and the angle between SY and a predetermined axis indicates noise type N. Fig. 1(a) shows S and its mixture with N, and Fig. 1(b) illustrates the relative positions of Y and enhanced speech $\hat{S}$.

As is shown in Fig. 1(b), compared with Y, $\hat{S}$ is typically closer to S. This corresponds to the observation that the SNR of enhanced speech is typically higher than that of noisy speech. In fact, many enhancement models are explicitly designed to elevate the SNR.

Along with the shorter distance to S, enhanced speech $\hat{S}$ may deviate from line SY. Such a noise mismatch between Y and $\hat{S}$ may degrade the performance of ASR systems trained only on Y. This may be the main cause of the distortion problem. In fact, for two utterances mixed with the same kind of noise at different SNRs, experimental results suggest that the one with higher SNR typically yield higher recognition performance. Note that, because of the similarity of masking-based and mapping-based speech enhancement in terms of distortion, the analysis above is expected to be valid for mapping-based systems as well.

B. Distortion-Independent Acoustic Modeling

For ASR backends trained on noisy speech and evaluated on enhanced speech, the input data for training and evaluation can be expressed as (8) and (9), respectively,

$$|Y_{tr}| = |S_{tr} + N_{tr}| \tag{8}$$

$$|\hat{S}_{eval}| = |S_{eval} + D_{eval}| \tag{9}$$

where $D_{eval} = N_{eval} \otimes M_{eval} - S_{eval} \otimes \overline{M_{eval}}$. Subscripts tr and $eval$ denote training and evaluation, respectively. Y_{tr}, S_{tr}, and N_{tr} are the spectral representations of noisy speech, clean speech, and additive noise in training, respectively. $\hat{S}_{eval}$ is the

enhanced speech in evaluation. D_{eval} denotes the distortion in enhanced speech and M_{eval} the T-F mask in evaluation.

Based on our analysis in the previous subsection, the mismatch between N_{tr} and D_{eval} is the cause of the distortion problem given the consistency of S_{tr} and S_{eval}. The use of enhanced speech in both training and evaluation can be viewed as a way to reduce the mismatch [15], [16]. In speech recognition corpora such as Aurora and CHiME series, only a limited number of noises are provided for training. In addition, noise types are shared between training and evaluation on these corpora. ASR systems trained with such noisy speech may not perform well on enhanced speech, which contains mismatched interference D_{eval}. Moreover, same noise types between training and evaluation give an advantage to unenhanced evaluation data. This is likely a main reason why speech enhancement does not improve recognition performance on these tasks.

To alleviate the distortion problem, N_{tr} can be modified in two ways. If we view D_{eval} as a special type of additive noise, a straightforward way is to increase the scope of N_{tr}. Since this strategy typically uses a large variety of additive noises to train acoustic models, we denote it noise-independent acoustic modeling. An advantage of noise-independent training is that its efficacy is not influenced by speech enhancement frontends. This acoustic modeling strategy, however, does not account for the fact that additive noises may differ significantly from distortions. Another strategy to alleviate the distortion problem is to train the acoustic model directly with enhanced speech, i.e.,

$$|\hat{S}_{tr}| = |S_{tr} + D_{tr}| \tag{10}$$

where $\hat{S}_{tr}$ denotes enhanced training speech and D_{tr} refers to the distortion in it.

We investigate a distortion-independent acoustic modeling method based on (10). The training set consists of a large variety of enhanced speech generated by a single well-trained speech enhancement frontend. The input to the speech enhancement model is noisy speech with various types of additive noise. An advantage of distortion-independent acoustic modeling is that D_{tr} in enhanced training speech is similar to D_{eval} during evaluation. The main concern is its generalization ability to other speech enhancement frontends. Since most supervised speech enhancement models can be viewed as nonlinear mapping from noisy speech to clean speech, distortion-independent acoustic model may be able to work with speech enhancement frontends different from the frontend used for training.

Fig. 2 illustrates distortion-independent acoustic modeling. The left diagram depicts the training stage and the right one testing. In the right diagram, speech enhancement blocks with dashed lines denote those not used during training. In this study, we evaluate three existing DNN based supervised speech enhancement models: gated residual network (GRN) [23], LSTM [2], and convolutional recurrent network (CRN) [24]. We also add the IRM as another enhancement frontend and a conventional MMSE based model [5]. The switch in the right diagram denotes the coupling between a distortion-independent acoustic model and various enhancement frontends.
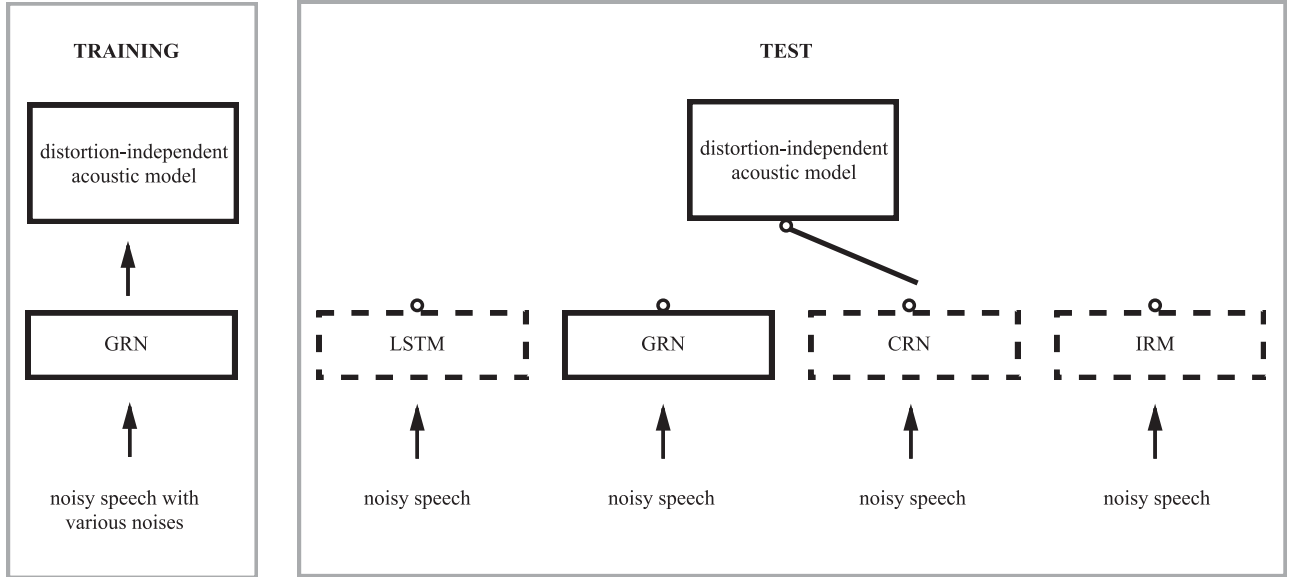
Fig. 2.   Illustration of distortion-independent acoustic modeling. See text for the meaning of acronyms.

C. Types of Acoustic Models

In addition to noise-independent and distortion-independent acoustic models, we investigate four other types of acoustic models: clean, noise-dependent, noise-mismatched, and joint training based. The clean acoustic model is trained using clean speech. In corpora containing both additive noise and reverberation, clean refers to reverberant speech without noise. The noise-dependent acoustic model is trained using only one type of noise and is tested on the same type of noise. This experimental setup represents typical robust speech recognition evaluations. The noise-mismatched acoustic model also uses a single type of noise during training, but it is tested on noises different from those for training. The joint training based acoustic model treats the speech enhancement frontend and ASR backend as a single model. It is initialized with the GRN based frontend and the well-trained noise-dependent acoustic model. The training and test noises for the joint training model are the same.

D. Utterance-Wise Recurrent Dropout for Acoustic Model Training

Utterance-wise recurrent dropout has been shown to be effective for acoustic model training on the CHiME-4 corpus [29]. A typical LSTM layer is described in Eqs. (11), (12), and (13) below:

$$\begin{pmatrix} \mathbf{i_t} \\ \mathbf{f}_t \\ \mathbf{o}_t \\ \mathbf{g}_t \end{pmatrix} = \begin{pmatrix} \sigma(\mathbf{W}_i\mathbf{x}_t + \mathbf{U}_i\mathbf{h}_{t-1} + \mathbf{b}_i) \\ \sigma(\mathbf{W}_f\mathbf{x}_t + \mathbf{U}_f\mathbf{h}_{t-1} + \mathbf{b}_f) \\ \sigma(\mathbf{W}_o\mathbf{x}_t + \mathbf{U}_o\mathbf{h}_{t-1} + \mathbf{b}_o) \\ f(\mathbf{W}_g\mathbf{x}_t + \mathbf{U}_g\mathbf{h}_{t-1} + \mathbf{b}_g) \end{pmatrix} \quad (11)$$

$$\mathbf{c}_t = \mathbf{f}_t \otimes \mathbf{c}_{t-1} + \mathbf{i}_t \otimes \mathbf{g}_t \quad (12)$$

$$\mathbf{h}_t = \mathbf{o}_t \otimes f(\mathbf{c}_t) \quad (13)$$

The dropout method can be expressed as follows:

$$\begin{pmatrix} \mathbf{i_t} \\ \mathbf{f}_t \\ \mathbf{o}_t \\ \mathbf{g}_t \end{pmatrix} = \begin{pmatrix} \sigma(\mathbf{W}_i d_{xit}(\mathbf{x}_t) + \mathbf{U}_i d_{hi}(\mathbf{h}_{t-1}) + \mathbf{b}_i) \\ \sigma(\mathbf{W}_f d_{xft}(\mathbf{x}_t) + \mathbf{U}_f d_{hf}(\mathbf{h}_{t-1}) + \mathbf{b}_f) \\ \sigma(\mathbf{W}_o d_{xot}(\mathbf{x}_t) + \mathbf{U}_o d_{ho}(\mathbf{h}_{t-1}) + \mathbf{b}_o) \\ f(\mathbf{W}_g d_{xgt}(\mathbf{x}_t) + \mathbf{U}_g d_{hg}(\mathbf{h}_{t-1}) + \mathbf{b}_g) \end{pmatrix} \quad (14)$$

where $\mathbf{i}_t$, $\mathbf{f}_t$, and $\mathbf{o}_t$ are the input, forget, and output gates at step t; $\mathbf{g}_t$ is the vector of cell updates and $\mathbf{c}_t$ denotes updated cell vector; $\mathbf{c}_t$ is used to update hidden state $\mathbf{h}_t$; σ is a sigmoid function and f is typically chosen to be $tanh$. $\mathbf{W}$ and $\mathbf{U}$ are the weight matrices for the input vector $\mathbf{x}_t$ and hidden vector $\mathbf{h}_{t-1}$, respectively. $\mathbf{b}$ denotes the bias term. The dropout function is denoted as $d()$. Subscripts x and h refer to the two corresponding feature vectors and i, f, o, g correspond to the four LSTM components. Dropout functions with subscript t are conventional frame-wise dropout, and those without t are recurrent, i.e., they use the same dropout mask at different time steps.

Utterance-wise recurrent dropout is designed to be both recurrent and have little temporal information loss. Four independently sampled utterance-wise masks are applied to $\mathbf{h}_{t-1}$. For the dropout on $\mathbf{x}_t$, we opt for a conventional frame-wise method since utterance-wise dropout may completely lose the information in some feature dimensions.

III. EXPERIMENTAL SETUP

A. Datasets

We use two corpora in our experiments. One of them is designed specially for this study and the other one follows the official CHiME-2 recipe.

1) WSJ: We compose a corpus by mixing clean speech in WSJ with additive noise. Although such simulated corpora are

not commonly used in speech recognition, they are common in speech enhancement [3], [23].

Training sets for the six acoustic models are designed in the following way. For the clean acoustic model, the clean utterances in the original WSJ corpus are used directly. The noise-dependent acoustic model has two instances, each corresponding to a different noise. The noise-mismatched acoustic model also has two instances, but it differs from the noise-dependent acoustic model in that its training and testing noises are mismatched. The training sets for the noise-dependent and noise-mismatched acoustic models are the same. It contains 7138 utterances generated by mixing clean utterances with a training noise (ADTbabble or ADTcafeteria1) at SNRs randomly chosen from {9 dB, 6 dB, 3 dB, 0 dB, −3 dB, −6 dB}. ADTbabble and ADTcafeteria1 (available at http://www.auditec.com) are commonly used in speech enhancement tasks [3], [23]. For the noise-independent acoustic model, the training set is generated by adding noise segments from a 10000 noise database (available at https://www.soundideas.com) to clean utterances at SNRs randomly chosen from the above six levels. The size of the noise-independent training set is 157036, 22 times that of the clean training set. The distortion-independent acoustic model is trained using GRN enhanced speech. GRN takes as input the noisy speech used for noise-independent acoustic model training. The distortion-independent training set thus also contains 157036 utterances. For the joint training based acoustic model, the training set is the same as that for the noise-dependent and noise-mismatched acoustic model.

A validation set is shared among the six acoustic models. It contains 1206 clean utterances from 10 speakers different from those used in training sets. Note that clean utterances are used directly in the validation set, avoiding biases to any specific noise.

The six acoustic models also share the same test set. It consists of 330 noisy utterances for each of the two test noises (ADT-babble and ADTcafeteria1) and at each of the six SNRs (i.e. {9 dB, 6 dB, 3 dB, 0 dB, −3 dB, −6 dB}). The total number of utterances is 3960. These utterances are from 12 speakers different from those in the training and validation sets.

Note that although ASR backends and enhancement frontends both use the 10 k noise database, their actual noise segments are different. First, ASR backends only use the first halves of noises, and enhancement frontends the second halves. Second, noise segments are randomly selected for recognition and enhancement.

2) CHiME-2: CHiME-2 is a commonly used corpus for robust speech recognition. Different from WSJ, utterances in CHiME-2 contain room reverberation. We treat reverberant speech in CHiME-2 as clean speech.

Training sets for the six acoustic models are designed based on the official recipe of the CHiME-2 challenge. The reverberant acoustic model is trained using reverberant utterances. Since each recording in CHiME-2 has two channels, we apply an average operation to get the corresponding monaural utterance. The noise-dependent acoustic model exactly follows the

CHiME-2 recipe. The noise-mismatched acoustic model tests the noise-dependent acoustic model on ADT noises (ADTbabble and ADTcafeteria1) rather than the CHiME-2 noises. Due to the limited number of noises provided in the CHiME-2 corpus, the noise-independent acoustic model is trained with noises from the 10 k noise database. We mix reverberant utterances with noise segments. The SNR levels are the same as those for WSJ. The noise-independent training set contains 157036 utterances in total. For the distortion-independent acoustic model, the training set consists of 157036 utterances enhanced by GRN. The training set for the joint training based acoustic model follows the official CHiME-2 recipe.

For the noise-dependent and joint training based acoustic model, we apply a validation set consisting of noisy utterances. For the other four acoustic models, we use reverberant utterances.

In addition to the official CHiME-2 test set, we generate two other test sets containing ADT noises. The average results on ADT noises are reported in this paper.

Due to reverberation, speech enhancement models for CHiME-2 are trained to map from reverberant-noisy speech to reverberant speech. The training data for speech enhancement models are generated similarly to those for the noise-independent acoustic model.

B. Implementation Details

We use a wide residual bidirectional LSTM network (WRBN) as the DNN architecture of acoustic models [11], [28], [29]. For speech enhancement frontends, we adopt a conventional MMSE model, the IRM, and three DNN based supervised models. For the three supervised models, GRN is the main frontend in our experiments and is used to generates both training data and test data. The other two models, LSTM and CRN, generate additional test data for distortion-independent acoustic modeling. These three frontends use different training targets. GRN applies the phase sensitive mask (PSM), LSTM uses the IRM, and CRN is mapping based.

We couple enhancement frontends and ASR backends with enhanced waveforms. The preprocessing steps for the enhancement frontends include windowing and Fourier transform. We apply the Hamming window with window width 20 ms and shift 10 ms. The windowed waveform signals are then converted to 320-dimensional short-time Fourier transform (STFT) features. Speech enhancement models take as input the STFT magnitudes and generate masks or enhanced magnitudes. We combine enhanced magnitudes with the phase of noisy speech to resynthesize enhanced waveform signals. As for the feature preprocessing for ASR backends, we make modifications to the recipe in Kaldi and our previous experiments [28], [29]. In order to avoid manually added interferences to enhanced speech, we skip most of preprocessing steps, including pre-emphasizing, dithering, and direct currency offset removal. Similar to speech enhancement frontends, we extract spectral features from enhanced waveform signals by applying the Hamming window and performing 320-dimensional STFT. We then apply Mel filters

TABLE I
WERs of the Six Acoustic Models on WSJ. *bab* and *caf* Denote ADTbabble and ADTcafeteria1, Respectively. *w/o* Refers to Noisy Evaluation Data Without Speech Enhancement (i.e. Unenhanced Speech), and *w/* Evaluation Data With Enhancement

SNR	clean				noise-dependent				noise-mismatched				noise-independent				distortion-independent				joint training			
	bab		caf		bab		caf		bab		caf		bab		caf		bab		caf		bab		caf	
	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/
9dB	11.92	**3.08**	12.83	**3.53**	**3.62**	4.35	**4.28**	5.04	6.31	**5.01**	4.95	**4.05**	4.89	**4.00**	4.97	**4.04**	4.18	**3.10**	3.81	**3.29**	**4.50**	9.02	**6.46**	9.19
6dB	22.19	**4.11**	22.06	**6.15**	**4.28**	5.04	**5.55**	6.39	9.83	**5.94**	7.98	**5.77**	7.14	**4.86**	7.17	**5.55**	5.10	**4.00**	5.59	**4.80**	**5.55**	9.38	**8.13**	9.51
3dB	38.26	**6.67**	38.88	**9.15**	**5.12**	6.31	**8.11**	8.93	17.07	**7.85**	14.16	**8.59**	10.59	**6.65**	11.06	**8.09**	7.17	**5.23**	8.85	**7.08**	**7.32**	10.63	**10.65**	11.60
0dB	60.32	**12.46**	58.25	**17.34**	**7.55**	9.64	**12.07**	14.68	28.41	**11.68**	26.13	**14.05**	18.23	**10.74**	17.56	**14.18**	12.87	**9.19**	15.21	**12.85**	**11.04**	13.38	16.87	**16.23**
-3dB	82.44	**23.24**	79.15	**32.51**	**12.55**	18.03	**21.93**	27.72	46.16	**21.39**	48.48	**26.43**	31.89	**19.71**	31.14	**26.64**	24.30	**17.13**	27.82	**24.58**	**18.92**	21.86	28.49	**26.15**
-6dB	93.16	**44.76**	91.44	**56.25**	**22.66**	34.34	**40.13**	48.40	71.64	**38.05**	74.67	**49.52**	54.06	**36.19**	53.99	**47.94**	45.41	**33.55**	50.68	**45.17**	**33.83**	37.49	48.29	**43.55**
avg	51.4	**15.7**	50.4	**20.8**	**9.3**	13.0	**15.3**	18.5	29.9	**15.0**	29.4	**18.1**	21.1	**13.7**	21.0	**17.7**	16.5	**12.0**	19.4	**16.3**	**13.5**	17.0	19.8	**19.4**

to STFT magnitudes to generate Mel frequency features. The dimension of Mel features is 80. In order to avoid underflow, we add a small value e^{-40} to Mel features and apply logarithm to the summation. The delta and delta-delta of log-Mel features are then generated, tripling the size of feature dimensionality. We calculate the mean value along time for each utterance and subtract it from the features. ASR backends take as input the normalized features and generate log posterior probabilities for senones. There are 1965 senones in our experiments. Subtracting log priors from log posteriors, we feed log likelihoods to the decoder in CHiME-2 to generate transcriptions.

In training the noise-independent and distortion-independent acoustic models, we monitor validation results after every 7138 utterances. This technique is commonly used in speech and language processing experiments.

1) WSJ: During training, most hyper-parameters for the six acoustic models are the same. The optimizer is Adam and dropout rate is 0.2. Initial learning rate is set to 10^{-3} for all acoustic models except the joint training based one, whose learning rate is 10^{-4}.

2) CHiME-2: For experiments on CHiME-2, the noise-mismatched acoustic model uses the well-trained noise-dependent acoustic model. Therefore, there are only five acoustic models on this corpus. The optimizer and dropout rate are the same as those on WSJ. All but the joint training based acoustic models share the same initial learning rate of 10^{-4}. The initial learning rate of the joint training based model is 10^{-5}.

IV. EVALUATION RESULTS AND ANALYSIS

This section presents and analyzes our evaluation results on the six acoustic models. The results are provided separately for the WSJ and CHiME-2 corpora.

A. Results on WSJ

Table I shows the WERs of the six acoustic models on WSJ. We use ADTbabble and ADTcafeteria1 as the noises for evaluation. The clean acoustic model clearly benefits from speech enhancement, as enhanced speech has a higher SNR than the corresponding noisy speech.

For the noise-dependent acoustic model, consistent with previous observations [4], [11], [19], the results on unenhanced speech are better. Based on our analysis on the distortion problem, the performance degradation on enhanced speech is caused by the mismatch between N_{tr} and D_{eval}.

The noise-mismatched acoustic models are able to benefit from speech enhancement in our experiments on WSJ. Such an ability, however, is influenced by the type of noise used for testing. We will discuss this more after presenting the results on CHiME-2. In Table I, we observe that the results of noise-dependent acoustic models are much better than those of noise-mismatched acoustic models on unenhanced speech. This indicates that acoustic models trained on one noise cannot generalize to untrained noises. This performance degradation caused by noise mismatch supports our analysis on the distortion problem.

The noise-independent acoustic model also benefits from speech enhancement on WSJ. This indicates that 10 k additive noises can capture a lot of the distortions on WSJ. The efficacy of noise-independent acoustic modeling, however, may be influenced by factors such as reverberation, as will be shown in the results on the CHiME-2 corpus.

For the distortion-independent acoustic model, the results on enhanced speech are better than those on unenhanced speech. This shows that distortion-independent acoustic models are able to alleviate the distortion problem caused by GRN. Moreover, the results of the distortion-independent acoustic model are better than those of the noise-independent model. Note that both noise-independent and distortion-independent acoustic models are tested on noises different from those used during training. The strong performance of our distortion-independent acoustic model shows that large-scale training with various distortions generalizes well to untrained distortions.

Generally speaking, the joint training based acoustic model, similar to the noise-dependent model, does not benefit from speech enhancement. For ADTcafeteria1, enhancement yields an improvement at low SNRs. Note that in our experiments, the

TABLE II
WERs OF THE DISTORTION-INDEPENDENT ACOUSTIC MODEL ON WSJ WITH OTHER FRONTENDS. SEE TABLE I CAPTION FOR NOTATIONS

SNR	unenhanced		LSTM		CRN		IRM		MMSE	
	bab	caf	bab	caf	bab	caf	bab	caf	bab	caf
9dB	4.18	3.81	3.36	3.21	3.27	4.09	2.73	2.73	4.45	5.36
6dB	5.10	5.59	4.24	4.88	4.13	4.65	2.75	2.88	5.96	7.17
3dB	7.17	8.85	5.70	7.64	5.03	7.53	2.84	2.65	9.42	11.40
0dB	12.87	15.21	9.02	13.17	8.39	11.53	2.88	2.86	17.13	19.30
-3dB	24.30	27.82	17.95	25.14	14.37	22.38	3.05	2.76	31.14	34.88
-6dB	45.41	50.68	34.99	47.80	28.19	40.82	2.91	2.93	53.63	61.20
avg	16.5	19.4	12.5	17.0	10.6	15.2	2.9	2.8	20.3	23.2

TABLE III
WERs OF THE SIX ACOUSTIC MODELS ON CHiME-2. *CHIME*-2 DENOTES THE OFFICIAL CHiME-2 EVALUATION SET. *ADT* REFERS TO THE AVERAGE WER OF ADTbabble AND ADTcafeteria1. SEE TABLE I CAPTION FOR OTHER NOTATIONS

SNR	reverberant				noise-dependent		noise-mismatched		noise-independent				distortion-independent				joint training	
	chime-2		ADT		chime-2		ADT		chime-2		ADT		chime-2		ADT		chime-2	
	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/
9dB	31.27	**10.50**	31.03	**11.40**	**5.49**	5.81	**7.82**	8.33	6.63	**6.37**	**6.59**	7.98	7.42	**5.51**	10.20	**6.60**	**6.03**	7.36
6dB	38.69	**13.67**	47.53	**19.00**	**6.26**	7.98	**10.28**	11.76	**7.72**	7.92	**8.66**	11.04	8.61	**6.54**	13.27	**8.64**	**6.67**	7.55
3dB	46.85	**17.26**	67.50	**31.68**	**6.78**	8.33	**18.03**	20.46	8.82	**8.78**	**14.00**	19.35	10.01	**7.10**	20.99	**14.73**	**7.75**	9.81
0dB	57.33	**23.73**	85.96	**50.89**	**8.95**	11.26	**30.07**	34.38	**10.69**	11.62	**23.72**	30.99	12.93	**9.70**	32.30	**22.76**	**9.81**	12.25
-3dB	62.94	**29.91**	93.49	**71.89**	**9.98**	14.48	**50.34**	55.30	**13.06**	13.30	**39.04**	51.60	14.85	**11.04**	51.02	**37.74**	**12.44**	16.64
-6dB	72.31	**39.87**	95.58	**88.79**	**14.83**	19.05	**75.65**	78.83	**17.45**	19.80	**60.73**	76.08	21.80	**15.45**	75.54	**58.24**	**16.78**	22.32
avg	51.6	**22.5**	70.2	**45.6**	**8.7**	11.2	**32.0**	34.8	**10.7**	11.3	**25.5**	32.8	12.6	**9.2**	33.9	**24.8**	**9.9**	12.7

joint training based model is initialized with the well-trained noise-dependent model and it does not perform better than the noise-dependent acoustic model. The reason may be that, in order to keep the same experimental setup as the other models, the number of trainable parameters in the joint model is overly large for a relatively small corpus like WSJ.

Along each column of Table I, there is a clear performance degradation as SNR reduces, consistent with our analysis on the cause of the distortion problem.

In Table II, we present the results of the distortion-independent acoustic model when coupled with speech enhancement frontends different from the one used during training. From the table, we observe that both LSTM and CRN yield better results than unenhanced speech. This shows that the distortion-independent acoustic model is able to generalize to different enhancement frontends. This also suggests that there may be a common pattern in the distortions introduced by DNN based supervised speech enhancement models.

Comparing the results of LSTM, CRN, and IRM, we find that for the distortion-independent acoustic model, the improvement of speech enhancement quality results in the improvement of recognition performance. In real-world applications, this suggests that a distortion-independent model need not to be re-trained when a more advanced speech enhancement frontend is applied. In addition, the distortion-independent acoustic model

on WSJ may be used to provide an indicator on the modeling ability of different speech enhancement frontends. Note that at different SNRs, the IRM results slightly vary, which may be due to the waveform resynthesis during speech enhancement. When the distortion-independent acoustic model is evaluated on clean speech, the average WER is 2.7%. For the clean acoustic model evaluated on clean speech, the WER is 2.0%.

The WERs on the MMSE enhanced test set are higher than those on unenhanced speech. It appears that MMSE has trouble handling non-stationary noises like ADTbabble and ADTcafeteria1 [37]. In fact, according to short-time objective intelligibility (STOI), which is a standard metric for speech enhancement, the average score of MMSE enhanced speech over the two test noises is 70.90%, lower than that of unenhanced speech, which is 76.02%.

B. Results on CHiME-2

Table III presents the WERs of the six acoustic models on CHiME-2. The noises used for evaluation include chime-2 noises and ADT. WERs on ADT are the averages of those on ADTbabble and ADTcafeteria1. The reverberant acoustic model on CHiME-2 corresponds to the clean acoustic model on WSJ. It is clear that the reverberant acoustic model benefits from speech enhancement.

TABLE IV
WERs of the Distortion-Independent Acoustic Model on CHiME-2 With Other Frontends. See Table III Caption for Notations

SNR	unenhanced		LSTM		CRN		IRM		MMSE	
	chime-2	ADT	chime-2	ADT	chime-2	ADT	chime-2	ADT	chime-2	ADT
9dB	7.42	10.20	5.79	7.61	6.65	7.50	3.40	3.66	8.46	10.65
6dB	8.61	13.27	7.47	10.21	7.68	10.09	3.44	3.64	10.09	15.13
3dB	10.01	20.99	8.63	17.57	9.04	15.57	3.34	3.62	12.40	22.35
0dB	12.93	32.30	11.36	28.47	11.25	25.73	3.38	3.73	15.36	34.52
-3dB	14.85	51.02	14.16	44.83	13.51	41.12	3.74	3.95	18.14	53.80
-6dB	21.80	75.54	19.41	67.08	18.06	62.33	3.31	4.10	27.59	79.29
avg	12.6	33.9	11.1	29.3	11.0	27.1	3.4	3.8	15.3	36.0

The noise-dependent acoustic model follows the official training recipe of the CHiME-2 challenge. Similar to prior observations [4], [18], the noise-dependent acoustic model does not benefit from speech enhancement, which is in line with the results in Table I.

We test the noise-mismatched acoustic model on ADT noises. Different from the results on WSJ, the noise-mismatched acoustic model does not perform better on enhanced speech. Note that the experiments on CHiME-2 use CHiME-2 noises for training, whereas the experiments on WSJ use ADT noises. The inconsistent results on the two corpora indicate that the ability of the noise-mismatched acoustic model to overcome the distortion problem may depend on the noise used for testing.

The noise-independent acoustic model on CHiME-2 does not gain performance improvement on enhanced speech. This is again different from the corresponding results on WSJ. On the CHiME-2 corpus, room impulse responses (RIRs) are different between training and testing [25]. Although we use a large variety of additive noises to train the noise-independent acoustic model, the RIR mismatch still exists. During testing, distortions introduced by speech enhancement thus deviate from the 10 k additive noises used for training. Note that at SNR level 9 dB and 3 dB, enhanced speech performs better than unenhanced speech on the CHiME-2 corpus.

The distortion-independent acoustic model is able to benefit from speech enhancement. On both CHiME-2 and ADT noises, distortion-independent acoustic model outperforms noise-independent acoustic model. The ability of distortion-independent acoustic modeling to benefit from speech enhancement shows that large-scale training on a variety of distortions generalizes to untrained distortions.

Similar to the situation on WSJ, the joint training based acoustic model on CHiME-2 does not benefit from speech enhancement. It also does not perform better than the noise-dependent acoustic model, likely due to overfitting.

Table IV shows the results of the distortion-independent acoustic model when used with different speech enhancement frontends. Similar to the experiments on WSJ, distortion-independent acoustic model is tested on LSTM and CRN enhanced speech. The results of both LSTM and CRN enhanced speech are better than those of unenhanced speech. This indicates the ability of the distortion-independent acoustic model to

TABLE V
WER Comparisons Between The Proposed Models and Prior Work on CHiME-2

models	9dB	6dB	3dB	0dB	-3dB	-6dB	avg
Wang and Wang [33]	6.61	6.86	8.67	10.39	13.02	18.23	10.6
Plantinga et al. [21]	-	-	-	-	-	-	9.3
distortion-independent	5.51	6.54	7.10	9.70	11.04	15.45	9.2
noise-dependent	**5.49**	**6.26**	**6.78**	**8.95**	**9.98**	**14.83**	**8.7**

work with various DNN based supervised speech enhancement frontends. This also suggests that distortions introduced by different supervised enhancement models have certain similarities.

On IRM enhanced speech, the distortion-independent acoustic model performs very well. This suggests that as speech enhancement research progresses, speech recognition performances of the distortion-independent acoustic model should also improve. The average WER of the distortion-independent acoustic model on reverberant speech is 3.4%. For the reverberant acoustic model evaluated on reverberant speech, the WER is 2.8%.

Like in Table II, the WERs on MMSE enhanced speech are higher than those on unenhanced speech. For the average STOI score over the three test noises, MMSE enhanced speech gets 67.44%, whereas unenhanced speech achieves 70.95%.

Table V shows a comparison of ASR systems in this study with those in prior work. It is worth noting that our distortion-independent acoustic model achieves a 9.2% WER, which is better than the previous best systems on the CHiME-2 corpus [21], [33]. For the noise-dependent acoustic model, we achieve an average WER of 8.7%, outperforming the previous best system by 6.5% relatively. Note that, in order to avoid the influence of model adaptation on our analysis of the distortion problem, we do not apply speaker adaptation to our models. The good results of our proposed models suggest that the observations in this study are likely valid for real world systems.

V. CONCLUDING REMARKS

The distortion problem occurs when we apply speech enhancement as a frontend for ASR tasks. This study treats the

distortion problem as a noise mismatch between training and testing. We categorize acoustic models into six types and examine each of them for their ability to overcome the distortion problem. Distortion-independent acoustic modeling emerges as the best among the six acoustic models. Its ability to generalize to untrained noises suggests the utility of large-scale training for acoustic modeling. We also show that the distortion-independent acoustic model is able to work with various DNN based supervised speech enhancement frontends. In addition, the WERs of our proposed distortion-independent and noise-dependent acoustic models both outperform the previous best system on the CHiME-2 corpus.

Future work on the distortion problem includes using ASR features as the training target of speech enhancement models, applying time-domain speech enhancement frontends, improving the generalization ability of the distortion-independent acoustic model by using both supervised and conventional speech enhancement models during training, and investigating distortion-independent training for end-to-end ASR systems.

REFERENCES

[1] D. Bagchi, P. Plantinga, A. Stiff, and E. Fosler-Lussier, "Spectral feature mapping with mimic loss for robust speech recognition," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2018, pp. 5609–5613.

[2] J. Chen and D. Wang, "Long short-term memory for speaker generalization in supervised speech separation," *J. Acoust. Soc. Am.*, vol. 141, pp. 4705–4714, 2017.

[3] J. Chen, Y. Wang, S. Yoho, D. Wang, and E. Healy, "Large-scale training to increase speech intelligibility for hearing-impaired listeners in novel noises," *J. Acoust. Soc. Am.*, vol. 139, pp. 2604–2612, 2016.

[4] J. Du, Q. Wang, T. Gao, Y. Xu, L. Dai, and C. Lee, "Robust speech recognition with speech enhanced deep neural networks," in *Proc. 15th Annu. Conf. Int. Speech Commun. Assoc.*, 2014, pp. 616–620.

[5] Y. Ephraim and D. Malah, "Speech enhancement using a minimum mean-square error short-time spectral amplitude estimator," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-32, no. 6, pp. 1109–1121, Dec. 1984.

[6] H. Erdogan, J. Hershey, S. Watanabe, and J. Le Roux, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2015, pp. 708–712.

[7] T. Gao, J. Du, L. Dai, and C. Lee, "Joint training of front-end and back-end deep neural networks for robust speech recognition," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2015, pp. 4375–4379.

[8] T. Gao, J. Du, L. Dai, and C. Lee, "SNR-based progressive learning of deep neural network for speech enhancement," in *Proc. INTERSPEECH*, 2016, pp. 3713–3717.

[9] T. Gao, J. Du, Y. Xu, C. Liu, L. Dai, and C. Lee, "Improving deep neural network based speech enhancement in low SNR environments," in *Proc. Int. Conf. Latent Variable Anal. Signal Sep.*, 2015, pp. 75–82.

[10] E. Healy, S. Yoho, Y. Wang, and D. Wang, "An algorithm to improve speech recognition in noise for hearing-impaired listeners," *J. Acoust. Soc. Am.*, vol. 134, pp. 3029–3038, 2013.

[11] J. Heymann, L. Drude, and H. Reinhold, "Wide residual BLSTM network with discriminative speaker adaptation for robust speech recognition," in *Proc. 4th Int. Workshop Speech Process. Everyday Environ.*, 2016, pp. 12–17.

[12] M. Kolbk, Z. Tan, and J. Jensen, "Speech intelligibility potential of general and specialized deep neural network based speech enhancement systems," *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 25, no. 1, pp. 153–167, Jan. 2017.

[13] K. Li, Z. Huang, Y. Xu, and C. Lee, "DNN-based speech bandwidth expansion and its application to adding high-frequency missing features for automatic speech recognition of narrowband speech," in *Proc. 16th Annu. Conf. Int. Speech Commun. Assoc.*, 2015, pp. 2578–2582.

[14] X. Lu, Y. Tsao, S. Matsuda, and C. Hori, "Speech enhancement based on deep denoising autoencoder," in *Proc. INTERSPEECH*, 2013, pp. 436–440.

[15] Z. Meng, J. Li, Y. Gong, and B. H. Juang, "Adversarial feature-mapping for speech enhancement," in *Proc. INTERSPEECH*, 2018, pp. 3259–3263.

[16] Z. Meng, J. Li, Y. Gong, and B. H. Juang, "Cycle-consistent speech enhancement," in *Proc. INTERSPEECH*, 2018, pp. 1165–1169.

[17] A. Narayanan *et al.*, "Toward domain-invariant speech recognition via large scale training," in *Proc. IEEE Spoken Language Technol. Workshop (SLT)*, 2018, pp. 441–447.

[18] A. Narayanan and D. Wang, "Ideal ratio mask estimation using deep neural networks for robust speech recognition," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2013, pp. 7092–7096.

[19] A. Narayanan and D. Wang, "Investigation of speech separation as a frontend for noise robust speech recognition," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 22, pp. 826–835, Apr. 2014.

[20] S. Nie, H. Zhang, X. Zhang, and W. Liu, "Deep stacking networks with time series for speech separation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2014, pp. 6667–6671.

[21] P. Plantinga, D. Bagchi, and E. Fosler-Lussier, "An exploration of mimic architectures for residual network based spectral mapping," in *Proc. IEEE Spoken Language Technol. Workshop (SLT)*, 2018, pp. 550–557.

[22] D. Povey *et al.*, "The Kaldi speech recognition toolkit," in *Proc. IEEE Workshop Autom. Speech, Recognit. Understanding*, no. EPFL-CONF-192584, 2011.

[23] K. Tan, J. Chen, and D. Wang, "Gated residual networks with dilated convolutions for monaural speech enhancement," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 1, pp. 189–198, Jan. 2019.

[24] K. Tan and D. Wang, "A convolutional recurrent neural network for real-time speech enhancement," in *Proc. INTERSPEECH*, 2018, pp. 3229–3233.

[25] E. Vincent, J. Barker, S. Watanabe, J. Le Roux, F. Nesta, and M. Matassoni, "The second CHIME speech separation and recognition challenge: Datasets, tasks and baselines," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2013, pp. 126–130.

[26] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 26, no. 10, pp. 1702–1726, Oct. 2018.

[27] P. Wang, K. Tan, and D. Wang, "Bridging the gap between monaural speech enhancement and recognition with distortion-independent acoustic modeling," in *Proc. INTERSPEECH*, 2019, pp. 471–475.

[28] P. Wang and D. Wang, "Filter-and-convolve: A CNN based multichannel complex concatenation acoustic model," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 5564–5568.

[29] P. Wang and D. Wang, "Utterance-wise recurrent dropout and iterative speaker adaptation for robust monaural speech recognition," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 4814–4818.

[30] Y. Wang, A. Narayanan, and D. Wang, "On training targets for supervised speech separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 22, no. 12, pp. 1849–1858, Dec. 2014.

[31] Y. Wang and D. Wang, "Boosting classification based speech separation using temporal dynamics," in *Proc. 13th Annu. Conf. Int. Speech Commun. Assoc.*, 2012, pp. 1528–1531.

[32] Y. Wang and D. Wang, "Cocktail party processing via structured prediction," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 224–232.

[33] Z. Wang and D. Wang, "A joint training framework for robust automatic speech recognition," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 24, no. 4, pp. 796–806, Apr. 2016.

[34] F. Weninger *et al.*, "Speech enhancement with LSTM recurrent neural networks and its application to noise-robust ASR," in *Proc. Int. Conf. Latent Variable Anal. Signal Separation*, 2015, pp. 91–99.

[35] Y. Xu, J. Du, L. Dai, and C. Lee, "Dynamic noise aware training for speech enhancement based on deep neural networks," in *Proc. 15th Annu. Conf. Int. Speech Commun. Assoc.*, 2014, pp. 2670–2674.

[36] Y. Xu, J. Du, L. Dai, and C. Lee, "An experimental study on speech enhancement based on deep neural networks," *IEEE Signal Process. Lett.*, vol. 21, no. 1, pp. 65–68, Jan. 2014.

[37] Y. Xu, J. Du, L. Dai, and C. Lee, "A regression approach to speech enhancement based on deep neural networks," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 23, no. 1, pp. 7–19, Jan. 2015.

[38] Y. Xu, J. Du, Z. Huang, L. Dai, and C. Lee, "Multi-objective learning and mask-based post-processing for deep neural network based speech enhancement," in *Proc. INTERSPEECH*, 2015, pp. 1508–1512.

Peidong Wang received the B.E. degree in electronic information engineering from the University of Science and Technology of China, Hefei, China, in 2015. He is currently working toward the Ph.D. degree with the Department of Computer Science and Engineering, The Ohio State University, Columbus, OH, USA. His research interests include robust automatic speech recognition, speech and language processing, microphone array processing, multimodal speech recognition, and machine learning.

Ke Tan received the B.E. degree in electronic information engineering from the University of Science and Technology of China, Hefei, China, in 2015. He is currently working toward the Ph.D. degree with the Department of Computer Science and Engineering, The Ohio State University, Columbus, OH, USA. His research interests include speech enhancement and separation, microphone array processing, robust speech recognition, and deep learning.

DeLiang Wang, photograph and biography not available at the time of publication.