

# SPATIAL DIFFUSENESS FEATURES FOR DNN-BASED SPEECH RECOGNITION IN NOISY AND REVERBERANT ENVIRONMENTS

*Andreas Schwarz, Christian Huemmer, Roland Maas, Walter Kellermann*

Multimedia Communications and Signal Processing  
Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)  
Cauerstr. 7, 91058 Erlangen, Germany  
{schwarz, huemmer, maas, wk}@lnt.de

## ABSTRACT

We propose a spatial diffuseness feature for deep neural network (DNN)-based automatic speech recognition to improve recognition accuracy in reverberant and noisy environments. The feature is computed in real-time from multiple microphone signals without requiring knowledge or estimation of the direction of arrival, and represents the relative amount of diffuse noise in each time and frequency bin. It is shown that using the diffuseness feature as an additional input to a DNN-based acoustic model leads to a reduced word error rate for the REVERB challenge corpus, both compared to logmelspec features extracted from noisy signals, and features enhanced by spectral subtraction.

**Index Terms**— Speech Recognition, Reverberation, Diffuse Noise, Deep Neural Networks

## 1. INTRODUCTION

In automatic speech recognizers (ASR) based on Gaussian Mixture Models and Hidden Markov Models (GMM-HMM), a wide variety of transformations and feature extraction steps is currently being employed with the aim of extracting and normalizing the information contained in the time-domain input signal as efficiently as possible. Recently, with the development of effective training methods for acoustic models based on multiple-layer neural networks, which are often summarized under the term “deep neural networks” (DNN) [1], it has become possible for the acoustic model to learn relationships between features and phonemes to a higher degree than it is possible with manually implemented feature transformation steps. For example, it has been found that simple filterbank features outperform mel-frequency cepstral coefficients (MFCCs) [2, 3], and it is conceivable that, given large amounts of training data and sufficiently complex network structures, time-domain signals may at some point even be directly used as inputs to a neural network.

Although the trend in ASR goes towards replacing explicit processing stages by implicit learning, for noise- and reverberation-robust ASR using microphone arrays, spatial information is still predominantly being exploited in a separate speech

enhancement preprocessor, e.g., in the form of beamforming [4], multichannel linear prediction [5], blocking matrix-based postfilters [6] or coherence-based postfilters [7]. The single-channel output of the preprocessor is then used to compute features for ASR. In some GMM-HMM-based systems, spatial information is exploited indirectly in uncertainty decoding-based approaches, e.g., in [8], where the feature uncertainty is derived from a noise estimate obtained in a multichannel signal enhancement stage. For DNN-based acoustic models, “noise-aware training” has been proposed [3], where a noise estimate is appended to the noisy feature vector. This has been evaluated for stationary noise estimates [3] and noise-estimates derived from time-frequency masking [9], but may in principle also be used for noise estimates obtained from spatial processing. In [10] and [11], feature vectors from multiple microphones are concatenated to form the input of a DNN-based acoustic model, however, no spatial phase information is exploited.

Inspired by the trend towards moving more explicit feature processing steps into the DNN, we propose to exploit spatial information about the diffuseness of the sound field directly by incorporating it into the acoustic model of a DNN-based speech recognizer. The diffuseness estimate is derived from the complex coherence between two omnidirectional microphones and has been used for signal enhancement based on the assumption that late reverberation and noise components can be modeled as diffuse noise [7]. Using the diffuseness as a feature is motivated by the fact that humans exploit similar spatial information for speech recognition in reverberant and noisy environments [12, 13], as it was found that the human auditory system treats spectro-temporal variations in the interaural coherence as “a perceptual surrogate for spectro-temporal variations in the energy of speech signals” [13]. The aim is to learn similar behavior in a DNN-based acoustic model.

We first describe the signal model for the estimation of the diffuseness from the instantaneous spatial coherence of a reverberated and noisy speech signal. Then, we show how this estimate is integrated into a feature extraction scheme for ASR, and describe the structure of the DNN-based speech recognizer. Finally, we evaluate the proposed feature on the two-channel task of the REVERB challenge [14], showing that the proposed approach outperforms both noisy multi-condition training and multichannel spectral subtraction-based signal enhancement.

The authors would like to thank the Deutsche Forschungsgemeinschaft (DFG) for supporting this work (contract number KE 890/4-2).

## 2. BLIND DIFFUSENESS ESTIMATION

We consider a reverberated and noisy speech signal recorded by two omnidirectional microphones. The signal  $x_i(t)$  recorded at the  $i$ -th microphone is composed of the desired signal component  $s_i(t)$  and the undesired noise component  $n_i(t)$  comprising additive noise and late reverberation, i.e.,  $x_i(t) = s_i(t) + n_i(t)$ ,  $i = 1, 2$ . The microphone, desired, and noise signals are represented in the short-time Fourier transform (STFT) domain by the corresponding uppercase letters, i.e.,  $X_i(k, f)$ ,  $S_i(k, f)$  and  $N_i(k, f)$ , respectively, with the discrete frame index  $k$  and continuous frequency  $f$ , and the auto- and cross-power spectra  $\Phi_{x_i x_j}(k, f)$ ,  $\Phi_{s_i s_j}(k, f)$ ,  $\Phi_{n_i n_j}(k, f)$ . Note that the continuous frequency  $f$  is used here for generality; in practice,  $f$  denotes discrete values along the frequency axis. It is assumed that the auto-power spectra of all signal components are identical at both microphones, i.e.,  $\Phi_{s_i s_i}(k, f) = \Phi_s(k, f)$ ,  $\Phi_{n_i n_i}(k, f) = \Phi_n(k, f)$ . The time- and frequency-dependent signal-to-noise ratio (SNR) of the microphone signals can then be defined as

$$SNR(k, f) = \frac{\Phi_s(k, f)}{\Phi_n(k, f)}. \quad (1)$$

The complex spatial coherence functions of the desired signal and noise components are given by

$$\Gamma_s(f) = \frac{\Phi_{s_1 s_2}(k, f)}{\Phi_s(k, f)}, \quad \Gamma_n(f) = \frac{\Phi_{n_1 n_2}(k, f)}{\Phi_n(k, f)}, \quad (2)$$

and are assumed to be time-invariant, i.e., dependent only on the spatial characteristics of the signal components. It is furthermore assumed that signal and noise components are orthogonal, such that  $\Phi_x(k, f) = \Phi_s(k, f) + \Phi_n(k, f)$ . The complex spatial coherence of the mixed sound field can then be written as a function of the SNR and the signal and noise coherence functions:

$$\Gamma_x(k, f) = \frac{SNR(k, f)\Gamma_s(f) + \Gamma_n(f)}{SNR(k, f) + 1}. \quad (3)$$

The direct sound is now modeled as a plane wave with an unknown direction of arrival (DOA) and therefore unknown time difference of arrival  $\Delta t$ , while the undesired noise and late reverberation component is modeled as a diffuse (spherically isotropic) sound field. The corresponding spatial coherence functions for the direct and diffuse sound components are then given by

$$\Gamma_s(f) = e^{j2\pi f \Delta t}, \quad (4)$$

$$\Gamma_n(f) = \Gamma_{\text{diffuse}}(f) = \text{sinc}(2\pi f \frac{d}{c}), \quad (5)$$

respectively. The direct signal coherence has a magnitude of one with an unknown phase determined by the DOA, while the diffuse noise coherence only depends on the known microphone spacing  $d$ .

The aim in the following is to estimate the SNR from the coherence of the mixed sound field  $\Gamma_x(k, f)$ . This coherence is first estimated as

$$\hat{\Gamma}_x(k, f) = \frac{\hat{\Phi}_{x_1 x_2}(k, f)}{\sqrt{\hat{\Phi}_{x_1 x_1}(k, f)\hat{\Phi}_{x_2 x_2}(k, f)}}, \quad (6)$$

where the spectral estimates  $\hat{\Phi}_{x_i x_j}(k, f)$  are obtained by recursive averaging:

$$\hat{\Phi}_{x_i x_j}(k, f) = \lambda \hat{\Phi}_{x_i x_j}(k-1, f) + (1-\lambda)X_i(k, f)X_j^*(k, f), \quad (7)$$

with a constant forgetting factor  $\lambda$  between 0 and 1. In [15, 7], it was shown that (3) can be solved for the SNR without requiring knowledge of  $\Gamma_s$ , using only the assumption that the desired signal is fully coherent, i.e.,  $|\Gamma_s| = 1$ . This yields a “blind” estimator for the SNR (or coherent-to-diffuse ratio, CDR) from the mixture coherence  $\hat{\Gamma}_x(k, f)$  which does not require knowledge or estimation of the signal DOA. The estimator is given in (8) at the bottom of this page (the indices  $k$  and  $f$  are omitted for brevity). The CDR can be transformed into the diffuseness [16]

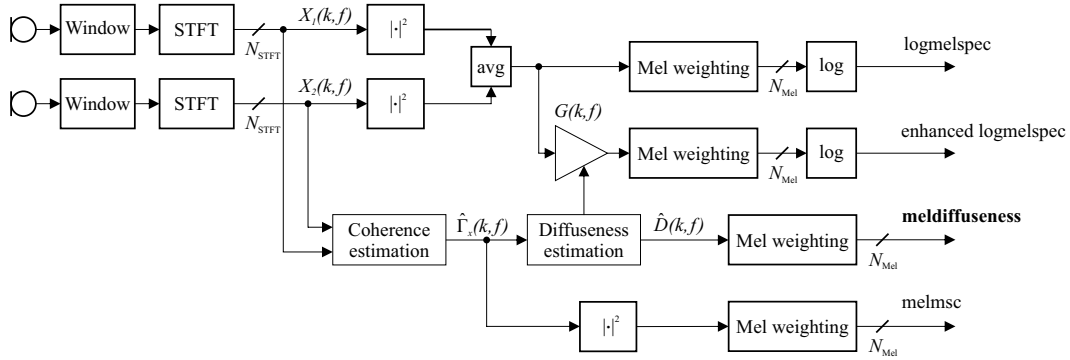
$$\hat{D}(k, f) = [\widehat{CDR}(k, f) + 1]^{-1}, \quad (9)$$

which can be thought of as the relative amount of diffuse signal power in the respective time and frequency bin. Since the diffuseness is bounded between 0 and 1, it is more convenient to use as basis for feature computation than the CDR itself.

## 3. FEATURE EXTRACTION FOR ASR

Fig. 1 shows the block diagram of the proposed feature extraction scheme. The microphone signals are first windowed and transformed into the STFT domain. The upper path then corresponds to a classical feature extraction of  $N_{\text{Mel}}$ -dimensional logmelspec (often termed “log-filterbank” or “Log FBank”) features, where the two microphone signals are combined by averaging the spectral powers computed from each microphone, and  $N_{\text{Mel}}$  triangular Mel-scaled weighting filters are applied. The second path shows the extraction of enhanced logmelspec features, where signal enhancement based on the diffuseness estimate is performed by multiplication in the STFT domain with a gain factor  $G(k, f)$ , which is computed as described in [7] according to the spectral magnitude subtraction rule. The third path illustrates the computation of the proposed “meldiffuseness” features: the diffuseness  $\hat{D}(k, f)$  is estimated as described in the previous section, and the same  $N_{\text{Mel}}$  triangular Mel weighting filters that are used in the logmelspec feature extraction are applied to create an output vector of the dimensionality  $N_{\text{Mel}}$ . Finally, for comparison, the Mel-weighted magnitude-squared coherence (“melmsc”) is computed as a feature. While the magnitude-squared coherence of a

$$\widehat{CDR}(k, f) = \frac{\Gamma_n \text{Re}\{\hat{\Gamma}_x\} - |\hat{\Gamma}_x|^2 - \sqrt{\Gamma_n^2 \text{Re}\{\hat{\Gamma}_x\}^2 - \Gamma_n^2 |\hat{\Gamma}_x|^2 + \Gamma_n^2 - 2\Gamma_n \text{Re}\{\hat{\Gamma}_x\} + |\hat{\Gamma}_x|^2}}{|\hat{\Gamma}_x|^2 - 1} \quad (8)$$



**Fig. 1.** Feature extraction of logmelspec, enhanced logmelspec, meldiffuseness and melmsc features from 2-channel signals.

mixed sound field is also related to the amount of diffuse noise, this relationship is strongly dependent on the signal DOA and the microphone spacing, therefore the melmsc feature is expected to perform worse than the proposed diffuseness estimate.

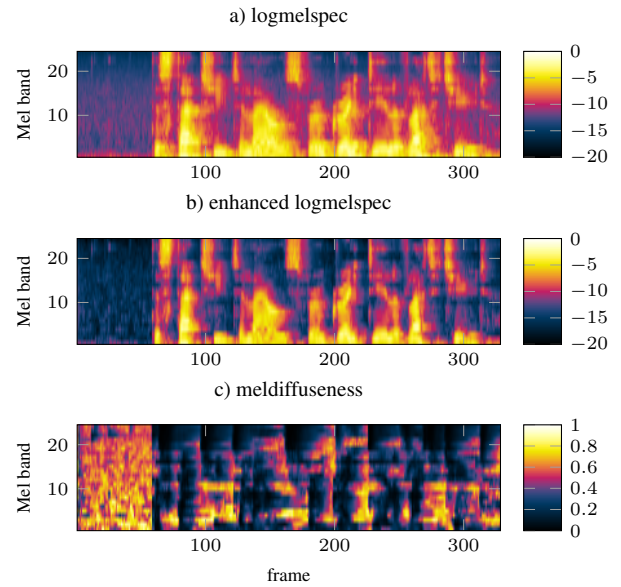
The interesting question is now how using concatenated logmelspec and meldiffuseness features as input to the neural network compares to using logmelspec features which have been enhanced in the STFT domain.

Since the trend in DNN-based acoustic modeling goes towards replacing explicit feature preprocessing and normalization steps by implicit learning, one might consider using the complex spatial coherence directly as feature. Note, however, that the proposed diffuseness feature has two significant advantages over the complex coherence. The complex coherence depends on two additional variables, namely the DOA and the microphone spacing, both of which would need to be sufficiently represented in the training data. Moreover, the diffuseness is a characteristic of the sound field which is independent of the microphone array geometry, and may therefore also be estimated from microphone arrays with other geometries, e.g., spherical arrays [17] or arrays consisting of directional microphones [18], without requiring adaptation of the acoustic model.

It is interesting to note that the additional temporal smoothing which is required for the estimation of the coherence (and therefore the diffuseness) has parallels in the human auditory system, where reaction to changes in interaural coherence was found to be more sluggish than reaction to changes in energy [19].

For the results presented in this paper, the time-domain signals (sampled at 16 kHz) are windowed using a 25 ms Hann window with a frame shift of 10 ms and transformed using a 512-point DFT, resulting in  $N_{STFT} = 257$  subbands in the STFT domain. The spatial coherence is estimated using the forgetting factor  $\lambda = 0.68$ .  $N_{Mel} = 24$  triangular Mel-scale weighting filters are used, covering a frequency range from 64 to 8000 Hz. MATLAB code for the feature computation is provided online<sup>1</sup>.

Fig. 2 illustrates the features computed from a noisy and reverberated speech signal taken from the multi-condition training set of the REVERB challenge corpus (LargeRoom2). The coherence-based spectral enhancement visibly reduces the noise floor and the smearing of the speech features over time. The meld-



**Fig. 2.** Features for the reverberated utterance “The statute allows for a great deal of latitude”.

diffuseness clearly highlights portions of the signal where noise or reverberation components are dominant.

#### 4. DNN-BASED SPEECH RECOGNITION

We employ the Kaldi toolkit [20] as ASR back-end system using the WSJ0 trigram 5k language model of the REVERB challenge and 3551 context-dependent triphone-states in the acoustic model.

In a first step, we set up a GMM-HMM baseline system based on Weninger et al. [4] by extracting 13 mean and variance normalized MFCCs (including the zeroth cepstral coefficient), followed by  $\pm 4$  frame splicing, linear discriminant analysis (LDA), maximum likelihood linear transform (MLLT), and feature-space maximum likelihood linear regression (fMLLR) (see [4, 21] for a detailed description). After conventional maximum likelihood training, discriminative training is performed with the boosted maximum mutual information (bMMI) criterion [4]. The GMM-

<sup>1</sup><http://www.lms.lnt.de/files/publications/icassp2015-diffuseness.zip>

**Table 1.** ASR Word Error Rate for the REVERB challenge evaluation and development test sets.

Recognizer	Feature	Evaluation Set										Development Set	
		SimData								RealData		SimData	RealData
		Room 1		Room 2		Room 3		Avg	Room 1		Avg	Avg	Avg
		near	far	near	far	near	far		near	far			
GMM-HMM	MFCC-LDA-MLLT-fMLLR	6.61	7.50	9.42	16.59	11.05	20.69	11.98	31.17	30.15	30.66	12.14	31.61
DNN-HMM	logmelspec+Δ+ΔΔ	5.74	6.67	7.65	13.92	8.65	14.62	9.54	28.45	29.14	28.80	9.68	24.93
	enhanced logmelspec+Δ+ΔΔ	6.61	7.12	7.65	12.18	8.32	14.57	9.41	28.46	29.07	28.77	9.13	25.29
	logmelspec+Δ+meldiffuseness	5.91	6.06	6.94	10.96	8.24	12.88	8.50	27.82	26.27	27.05	7.92	24.19
	logmelspec+Δ+melmisc	6.17	6.32	7.04	12.25	8.17	13.87	8.97	27.34	27.99	27.67	8.68	24.73

HMM system is trained on the clean WSJCAM0 Cambridge Read News REVERB corpus [22]. The alignment of the training data to the HMM states is then extracted from the clean training data and used for the later multi-condition training of the DNN-HMM system. This technique is known to yield better results than a multi-condition state-frame alignment [9, 23].

The hybrid DNN-HMM Kaldi system is based on “Dan’s implementation” [20] using a maxout network with 2-norm nonlinearities/activation functions and 4 hidden layers, each one with an input dimension of 2000 and an output dimension of 400.

In accordance with [2, 3], and as described in the previous section, we extract  $N_{\text{Mel}} = 24$  static logmelspec coefficients, with or without applying coherence-based spectral subtraction enhancement in the STFT domain. Depending on the particular setup in Table 1, also Delta ( $\Delta$ ), acceleration ( $\Delta\Delta$ ), melmsc, and/or the proposed meldiffuseness features are derived. Mean and variance normalization and  $\pm 5$  frame splicing is applied to the entire resulting feature vector. The training is performed on the REVERB multi-condition training set [14], consisting of 7861 noisy and reverberated utterances from the WSJCAM0 corpus, using greedy layer-wise supervised training, preconditioned stochastic gradient descent, “mixing up” [24] as well as final model combination [24].

## 5. EVALUATION RESULTS

We evaluate the proposed system using the two-channel task of the REVERB challenge [14]. The REVERB evaluation test set consists of  $\sim 5000$  reverberated and noisy utterances, partially created by convolution of clean WSJCAM0 utterances with impulse responses and mixing with recorded noise sequences (“SimData”), and partially consisting of multichannel recordings of speakers in a reverberant and noisy room from the MC-WSJ-AV corpus (“RealData”). For SimData, the reverberation times of the three rooms are approx. 0.25 s, 0.5 s and 0.7 s and the source-microphone spacing is 0.5 m (near) or 2 m (far). For RealData, the reverberation time is approx 0.7 s and the source-microphone distance is 1 m (near) or 2.5 m (far). In both cases, an 8-channel circular microphone array with a diameter of 20 cm was used, of which two microphones with a spacing of  $d = 8$  cm are selected for the two-channel recognition task which is evaluated here.

First, we evaluate the word error rate (WER) obtained from the GMM-based recognizer with MFCC features, which is used to obtain the alignment. For the DNN-based recognizer, we compare logmelspec features extracted from the noisy signals, and

enhanced logmelspec features. In both cases, the feature vector is extended by first- ( $\Delta$ ) and second-order ( $\Delta\Delta$ ) derivatives. Then, we evaluate the combination of noisy logmelspec features with spatial meldiffuseness or melmsc features; in this case, only first-order derivatives ( $\Delta$ ) are computed for the logmelspec features, in order to keep the overall dimension of the feature vectors the same ( $3N_{\text{Mel}}$ ).

Table 1 shows the WER results for the REVERB challenge evaluation test set, and the average WER for the development test set. As expected, the DNN-based acoustic model achieves a lower WER than the GMM-based model. The diffuseness-based signal enhancement has a negligible effect on WER. This seems to contradict [15], where the same signal enhancement method led to a significantly lower WER, however, there, acoustic models were trained on clean speech. Apparently the effect of the multichannel spectral subtraction for signal enhancement is compensated by noisy multi-condition training. Using the combined noisy logmelspec and diffuseness features as input to the neural network however yields a significantly reduced WER. This confirms that the spatial information extracted from the coherence can be exploited more successfully by the DNN than by speech enhancement using spectral subtraction, even though, in this case, the frequency resolution of the meldiffuseness features is reduced compared to the diffuseness estimate used for spectral subtraction. The melmsc feature also leads to a reduced WER compared to noisy logmelspec features, although the improvement is smaller than with meldiffuseness features.

## 6. CONCLUSION

It has been shown that spatial information extracted from multiple microphones does not necessarily have to be exploited in a signal enhancement front-end, but may be used more effectively as an additional feature input for a DNN-based speech recognizer. The proposed approach has a number of properties which make it highly suitable for practical applications like cloud-based speech recognition for smartphones. First, the diffuseness feature is normalized with respect to the microphone array geometry, and can therefore be used for speech recognition with features extracted from a variety of multichannel recording devices without requiring adaptation of the acoustic model. Second, the feature can be computed in real-time (as opposed to batch processing) and “blindly” in the sense that knowledge or estimation of the direction of arrival is not required. Finally, the evaluation shows that consistent improvements in recognition accuracy can be achieved.

## 7. REFERENCES

- [1] G. Hinton, L. Deng, D. Yu, G. Dahl, A.-R. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. Sainath, and B. Kingsbury, "Deep neural networks for acoustic modeling in speech recognition," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 82–97, Nov. 2012.
- [2] L. Deng, J. Li, J.-T. Huang, K. Yao, D. Yu, F. Seide, M. Seltzer, G. Zweig, X. He, J. Williams, and others, "Recent advances in deep learning for speech research at Microsoft," in *Proc. ICASSP*, Vancouver, Canada, 2013, pp. 8604–8608.
- [3] M. Seltzer, D. Yu, and Y. Wang, "An investigation of deep neural networks for noise robust speech recognition," in *Proc. ICASSP*, Vancouver, Canada, 2013, pp. 7398–7402.
- [4] F. Weninger, S. Watanabe, J. Le Roux, J. R. Hershey, Y. Tachioka, J. Geiger, B. Schuller, and G. Rigoll, "The MERL/MELCO/TUM system for the REVERB challenge using deep recurrent neural network feature enhancement," in *Proc. REVERB Workshop*, Florence, Italy, 2014.
- [5] M. Delcroix, T. Yoshioka, A. Ogawa, Y. Kubo, M. Fujimoto, N. Ito, K. Kinoshita, M. Espi, T. Hori, and T. Nakatani, "Linear prediction-based dereverberation with advanced speech enhancement and recognition technologies for the REVERB challenge," in *Proc. REVERB Workshop*, Florence, Italy, 2014.
- [6] R. Maas, A. Schwarz, Y. Zheng, K. Reindl, S. Meier, A. Sehr, and W. Kellermann, "A two-channel acoustic front-end for robust automatic speech recognition in noisy and reverberant environments," in *Proc. International Workshop on Machine Listening in Multisource Environments (ChiME)*, Florence, Italy, 2011, pp. 41–46.
- [7] A. Schwarz and W. Kellermann, "Coherent-to-diffuse power ratio estimation for dereverberation," *IEEE/ACM Trans. on Audio, Speech and Language Processing*, 2015, under review, preprint available: <http://arxiv.org/abs/1502.03784>.
- [8] R. F. Astudillo, D. Kolossa, A. Abad, S. Zeiler, R. Saeidi, P. Mowlae, J. P. da Silva Neto, and R. Martin, "Integration of beamforming and uncertainty-of-observation techniques for robust ASR in multi-source environments," *Computer Speech & Language*, vol. 27, no. 3, pp. 837–850, May 2013.
- [9] A. Narayanan and D. Wang, "Joint noise adaptive training for robust automatic speech recognition," in *Proc. ICASSP*, Florence, Italy, 2014, pp. 2504–2508.
- [10] P. Swietojanski, A. Ghoshal, and S. Renals, "Hybrid acoustic models for distant and multichannel large vocabulary speech recognition," in *Proc. ASRU*, Olomouc, Czech Republic, 2013, pp. 285–290.
- [11] Y. Liu, P. Zhang, and T. Hain, "Using neural network front-ends on far field multiple microphones based speech recognition," in *Proc. ICASSP*, Florence, Italy, 2014, pp. 5542–5546.
- [12] L. Danilenko, "Binaurales Hören im nichtstationären diffusen Schallfeld," *Kybernetik*, vol. 6, no. 2, pp. 50–57, June 1969.
- [13] J. F. Culling, B. A. Edmonds, and K. I. Hodder, "Speech perception from monaural and binaural information," *The Journal of the Acoustical Society of America*, vol. 119, no. 1, pp. 559–565, 2006.
- [14] K. Kinoshita, M. Delcroix, T. Yoshioka, T. Nakatani, A. Sehr, W. Kellermann, and R. Maas, "The REVERB challenge: A common evaluation framework for dereverberation and recognition of reverberant speech," in *Proc. WASPAA*, New Paltz, NY, USA, 2013, pp. 1–4.
- [15] A. Schwarz and W. Kellermann, "Unbiased coherent-to-diffuse ratio estimation for dereverberation," in *Proc. IWAENC*, Antibes - Juan Les Pins, France, 2014, pp. 6–10.
- [16] G. Del Galdo, M. Taseska, O. Thiergart, J. Ahonen, and V. Pulkki, "The diffuse sound field in energetic analysis," *The Journal of the Acoustical Society of America*, vol. 131, no. 3, pp. 2141–2151, Mar. 2012.
- [17] D. P. Jarrett, O. Thiergart, E. A. P. Habets, and P. A. Naylor, "Coherence-based diffuseness estimation in the spherical harmonic domain," in *Proc. IEEEI*, Eilat, Israel, 2012, pp. 1–5.
- [18] O. Thiergart, T. Ascherl, and E. A. P. Habets, "Power-based signal-to-diffuse ratio estimation using noisy directional microphones," in *Proc. ICASSP*, Florence, Italy, 2014, pp. 7440–7444.
- [19] J. F. Culling and H. S. Colburn, "Binaural sluggishness in the perception of tone sequences and speech in noise," *The Journal of the Acoustical Society of America*, vol. 107, no. 1, pp. 517–527, Jan. 2000.
- [20] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembe, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, and others, "The Kaldi speech recognition toolkit," in *Proc. ASRU*, Waikoloa, HI, USA, 2011.
- [21] S. P. Rath, D. Povey, K. Vesely, and J. Cernocky, "Improved feature processing for deep neural networks," in *Proc. Interspeech*, Lyon, France, 2013, pp. 109–113.
- [22] T. Robinson, J. Fransen, D. Pye, J. Foote, S. Renals, P. Woodland, and S. Young, "WSJCAM0 Cambridge read news for REVERB LDC2013E109," Web Download. Philadelphia: Linguistic Data Consortium, 2013.
- [23] M. Delcroix, Y. Kubo, T. Nakatani, and A. Nakamura, "Is speech enhancement pre-processing still relevant when using deep neural networks for acoustic modeling?," in *Proc. Interspeech*, Lyon, France, 2013, pp. 2992–2996.
- [24] X. Zhang, J. Trmal, D. Povey, and S. Khudanpur, "Improving deep neural network acoustic models using generalized maxout networks," in *Proc. ICASSP*, 2014, pp. 215–219.