

# Robust coherence-based spectral enhancement for speech recognition in adverse real-world environments

Hendrik Barfuss, Christian Huemmer, Andreas Schwarz<sup>1</sup>,  
and Walter Kellermann

*Multimedia Communications and Signal Processing,  
Friedrich-Alexander University Erlangen-Nürnberg,  
Cauerstr. 7, 91058 Erlangen, Germany  
{barfuss,huemmer,schwarz,wk}@lnt.de*

---

## Abstract

Speech recognition in adverse real-world environments is highly affected by reverberation and nonstationary background noise. A well-known strategy to reduce such undesired signal components in multi-microphone scenarios is spatial filtering of the microphone signals. In this article, we demonstrate that an additional coherence-based postfilter, which is applied to the beamformer output signal to remove diffuse interference components from the latter, is an effective means to further improve the recognition accuracy of modern deep learning speech recognition systems. To this end, the recently updated 3rd CHiME Speech Separation and Recognition Challenge (CHiME-3) baseline speech recognition system is extended by a coherence-based postfilter and the postfilter's impact on the word error rates is investigated for the noisy environments provided by CHiME-3. To determine the time- and frequency-dependent postfilter gains, we use a Direction-of-Arrival (DOA)-dependent and a DOA-independent estimator of the coherent-to-diffuse power ratio as an approximation of the short-time signal-to-noise ratio. Our experiments show that incorporating coherence-based postfiltering into the CHiME-3 baseline speech recognition system leads to a significant reduction of the word error rate scores for the noisy and reverberant environments provided as part of CHiME-3.

---

<sup>1</sup>A. Schwarz was with the Friedrich-Alexander University of Erlangen-Nürnberg while the work has been conducted. He is now with Amazon, Aachen, Germany.

*Keywords:* Robust speech recognition, Postfiltering, Spectral enhancement, Coherence-to-diffuse power ratio, Wiener filter

---

## 1. Introduction

For a satisfying user experience of human-machine interfaces it is crucial to ensure a high accuracy in automatically recognizing the user’s speech. However, as soon as no close-talking microphone is used for capturing the desired  
5 speech signal, the recognition accuracy suffers from additional reverberation, background noise and active interfering speakers as picked up by the microphones [1, 2]. Signal processing techniques for robust speech recognition in such reverberant and noisy environments can be categorized into either front-end (e.g., speech enhancement [3, 4, 5]) or back-end (e.g., acoustic-model adapta-  
10 tion [6, 7, 8]) processing techniques.

The 3rd CHiME Speech Separation and Recognition Challenge (CHiME-3) [9] targets the performance of state-of-the-art Automatic Speech Recognition (ASR) systems in real-world scenarios. The primary goal is to improve the ASR performance of real recorded speech of a person talking to a tablet device  
15 in realistic everyday noisy environments by employing front-end and/or back-end signal processing techniques. For this, a baseline ASR system has been initially provided and recently updated (as follow-up of CHiME-3) to achieve a high recognition accuracy in everyday real-world scenarios. Front-end processing of the updated baseline now employs the BeamformIt toolkit [10] for pro-  
20 cessing the recorded microphone signals by a Weighted Delay-and-Sum (WDS) beamforming technique. The beamformer output is used as input to the ASR back-end system which contains a Deep Neural Network (DNN)-based acoustic model and a Recurrent Neural Network (RNN)-based language model.

In this contribution, we extend the updated CHiME-3 baseline system by a  
25 low-complexity coherence-based postfilter which is applied to the beamformer output signal to further remove reverberation and nonstationary background noise from the latter. The postfilter is realized as a Wiener filter, where, in




Figure 1: Overview of the overall signal processing pipeline system with beamformer and postfilter as acoustic front-end signal processing. The acoustic back-end system, including feature extraction/transformation, is equal to the baseline acoustic back-end system provided by CHiME-3 [9].

contrast to the classical Wiener filter, see, e.g., [11], we use an estimate of the Coherent-to-Diffuse Power Ratio (CDR) as an approximation of the short-time  
 30 Signal-to-Noise Ratio (SNR) to compute the time- and frequency-dependent Wiener filter gains. The CDR, which is the ratio of the power of direct and diffuse signal components, needs to be estimated from the microphone signals. In this article we compare and evaluate two different CDR estimators, namely a Direction-of-Arrival (DOA)-independent and a DOA-dependent estimator,  
 35 which have been proposed and shown to be very effective for dereverberation in [12, 13]. In contrast to this previous work, we now evaluate the dereverberation and noise reduction performance of both coherence-based Wiener filter realizations with respect to Word Error Rates (WERs) obtained with the state-of-the-art CHiME-3 baseline ASR system for simulated and real data. Furthermore,  
 40 we illustrate the impact of our proposed coherence-based Wiener filter in the short-time Fourier transform (STFT) domain.

An overview of the extended signal processing pipeline which is employed in this work is given in Figure 1. While the purpose of the beamformer is to spatially focus on the target source, i.e., to reduce the signal components from  
 45 interfering point sources, the postfilter shall remove diffuse interference components, e.g., reverberation, from the beamformer output signal. The output of the front-end signal enhancement (consisting of beamformer and postfilter) is further processed by the ASR back-end system of the CHiME-3 baseline, which provides an Hidden Markov Model (HMM)-DNN-based speech recognizer[9].

50 The remainder of this article is structured as follows: In Section 2, the

employed front-end signal enhancement is described in detail, followed by a brief review of the employed ASR system in Section 3. The performance of the front-end speech enhancement using two different CDR estimators is evaluated with respect to WERs of the baseline ASR system, and the results are presented in Section 4. Concluding remarks are provided in Section 5.

## 2. Front-end enhancement techniques

The front-end speech enhancement considered in this article consists of a WDS beamformer (based on the BeamformIt toolkit [10], provided by the CHiME-3 baseline [9]) and a single-channel coherence-based Wiener filter. In the following, we introduce the signal model which is used throughout this paper. Then, the baseline WDS beamformer is briefly reviewed, followed by a detailed presentation of the coherence-based Wiener filter based on DOA-independent and DOA-dependent CDR estimation.

### 2.1. Signal model

For a consistent presentation of the front-end speech enhancement considered in this work, we first introduce a signal model which will be used throughout this article.

The  $N$  microphone signals of the microphone array in the STFT domain at frame  $l$  and frequency  $f$  are given as:

$$\mathbf{x}(l, f) = \mathbf{h}(l, f)S(l, f) + \mathbf{n}(l, f), \quad (1)$$

where vector

$$\mathbf{x}(l, f) = [X_0(l, f), X_1(l, f), \dots, X_{N-1}(l, f)]^T \quad (2)$$

contains the microphone signals,  $S(l, f)$  denotes the clean source signal, and  $\mathbf{n}(l, f)$  includes sensor noise as well as interference and diffuse background noise components and is defined analogously to  $\mathbf{x}(l, f)$  in (2). Assuming free-field propagation of sound waves,  $\mathbf{h}(l, f)$  represents the steering vector modelling

the sound propagation between the target source DOA  $(\phi_d, \theta_d)$  (azimuth and elevation angle  $\phi$  and  $\theta$  are defined as in [14]) and all  $N$  microphones:

$$\mathbf{h}(l, f) = [e^{j2\pi f \tau_0}, e^{j2\pi f \tau_1}, \dots, e^{j2\pi f \tau_{N-1}}]^T, \quad (3)$$

where  $\tau_n$  is the Time Difference of Arrival (TDOA) for the target source direction of the  $n$ -th channel with respect to the reference microphone channel [14]:

$$\tau_n = \frac{\mathbf{a}^T \mathbf{p}_n}{c}. \quad (4)$$

Here, vector  $\mathbf{p}_n$  contains the position of the  $n$ -th microphone in Cartesian coordinates and vector  $\mathbf{a}$  is defined as [14]:

$$\mathbf{a} = -[\sin(\theta_d)\cos(\phi_d), \sin(\theta_d)\sin(\phi_d), \cos(\theta_d)]^T. \quad (5)$$

Moreover,  $c$  represents the speed of sound and operator  $(\cdot)^T$  denotes the transpose of a vector or matrix. The beamformer output  $Y_{\text{BF}}(l, f)$  is obtained by multiplying each microphone signal with a complex-valued beamformer coefficient  $W_n(l, f)$ , followed by a summation over all microphone channels:

$$Y_{\text{BF}}(l, f) = \mathbf{w}^T(l, f)\mathbf{x}(l, f), \quad (6)$$

where

$$\mathbf{w}(l, f) = [W_0(l, f), \dots, W_{N-1}(l, f)]^T \quad (7)$$

contains the beamformer coefficients  $W_n(l, f)$ .

Subsequently, the postfilter is applied to the beamformer output signal, yielding the overall output signal

$$Y(l, f) = G(l, f)Y_{\text{BF}}(l, f), \quad (8)$$

where  $G(l, f)$  describes the postfilter gain at frame  $l$  and frequency  $f$ . The enhanced signal  $Y(l, f)$  is used as input of the CHiME-3 baseline acoustic back-end system [9].

## 2.2. Weighted delay-and-sum beamformer

The employed beamformer is provided by the BeamformIt toolkit [10] and based on the WDS beamforming technique [15]. The  $n$ -th beamformer filter coefficient  $W_n(l, f)$  at frame  $l$  and frequency  $f$  is given as

$$W_n(l, f) = w_n(l)e^{-j2\pi f\tau_n}, \quad (9)$$

where  $w_n$  is the frequency-independent weight for the  $n$ -th channel.

The TDOAs are estimated using the Generalized Cross Correlation with  
75 Phase Transform (GCC-PHAT) localization technique, see, e.g., [16]. Before the signals are time-aligned a two-step post processing is applied to the estimated TDOAs: First, a noise threshold is estimated and employed to remove non-reliable TDOA estimates which may have been obtained from non-speech or noisy segments. Second, Viterbi decoding of the remaining TDOA values is  
80 performed to maximize the speaker continuity, i.e., to avoid steering the beam to noise sources which are only present at a very short time span.

The channel weights are chosen adaptively over time, starting with the classical delay-and-sum as initial value:  $W_n(0, f) = 1/N$ . Furthermore, automatic channel selection and elimination is performed to avoid using microphone  
85 signals of poor quality. Both, channel weight adaptation and channel selection and elimination are based on an average cross-correlation value of the microphone channels.

For a more detailed explanation of the baseline beamformer, as provided by the BeamformIt toolkit, we refer the reader to [10, 17].

## 90 2.3. Coherence-based postfilter

As illustrated in Figure 1, we apply a postfilter to remove diffuse noise components from the beamformer output signal. The postfilter gain  $G(l, f)$  at frame  $l$  and frequency  $f$  is given as [11, 18]:

$$G(l, f) = \max \left\{ 1 - \mu \frac{1}{1 + \text{SNR}(l, f)}, G_{\min} \right\}, \quad (10)$$

with overestimation factor  $\mu$ , and gain floor  $G_{\min}$ . The postfilter in (10) is a Wiener filter which uses the short-time SNR to compute the filter gain  $G(l, f)$ . In this work, we approximate the short-time SNR in (10) by an estimate of the so-called CDR, which is the ratio between direct and diffuse signal components. From (10), it can be seen that a low CDR value, which corresponds to strong diffuse signal components being present at the input of the system, leads to low filter gains and vice versa.

The CDR between two omnidirectional microphones is defined as [12]:

$$\text{CDR}(l, f) = \frac{\Gamma_n(l, f) - \Gamma_x(l, f)}{\Gamma_x(l, f) - \Gamma_s(l, f)}, \quad (11)$$

where  $\Gamma_x(l, f)$ ,  $\Gamma_s(l, f)$ ,  $\Gamma_n(l, f)$  denote the coherence of the observations, the direct-path signals, and the noise between two observation points (microphones), respectively. In the following, the two microphones are indexed by the variables  $p = 1, \dots, N$  and  $q = 1, \dots, N$ , respectively. To this end, the spatial coherence functions for the direct-path signals and diffuse noise components are given as

$$\Gamma_s(l, f) = e^{j2\pi f(\tau_p - \tau_q)}, \quad (12)$$

$$\Gamma_n(l, f) = \frac{\sin(2\pi f \frac{d_{pq}}{c})}{2\pi f \frac{d_{pq}}{c}}, \quad (13)$$

respectively, with TDOAs  $\tau_p, \tau_q$  calculated in (4) and microphone spacing  $d_{pq}$ . Moreover, a short-time estimate  $\hat{\Gamma}_x(l, f)$  of the coherence function of both microphone signals  $\Gamma_x(l, f)$  in (11) can be obtained using

$$\hat{\Gamma}_x(l, f) = \frac{\hat{\Phi}_{x_p x_q}(l, f)}{\sqrt{\hat{\Phi}_{x_p x_p}(l, f) \hat{\Phi}_{x_q x_q}(l, f)}} \quad (14)$$

by estimating the auto- and cross-power spectra from the microphone signals  $X_p(l, f)$  and  $X_q(l, f)$  based on recursive averaging

$$\hat{\Phi}_{x_p x_q}(l, f) = \lambda \hat{\Phi}_{x_p x_q}(l, f) + (1 - \lambda) X_p(l, f) X_q^*(l, f) \quad (15)$$

with forgetting factor  $\lambda$ . Operator  $(\cdot)^*$  creates the conjugate complex of  $(\cdot)$ . However, inserting the coherence estimate  $\hat{\Gamma}_x(l, f)$  into (11) is not feasible due to the mismatch between coherence models and actual acoustic conditions as

$\text{CDR}(l, f)$  might become a complex-valued quantity [13]. Thus, a variety of different CDR estimators have been proposed in the literature, see, e.g., [19, 20, 21]. In this work, we use two CDR estimators, which have been proposed and shown to be especially effective in [12], and which are given by

$$\widehat{\text{CDR}}_{\text{DOAindep}} = \frac{\Gamma_n \text{Re}\{\hat{\Gamma}_x\} - |\hat{\Gamma}_x|^2 - \sqrt{\Gamma_n^2 \text{Re}\{\hat{\Gamma}_x\}^2 - \Gamma_n^2 |\hat{\Gamma}_x|^2 + \Gamma_n^2 - 2 \Gamma_n \text{Re}\{\hat{\Gamma}_x\} + |\hat{\Gamma}_x|^2}}{|\hat{\Gamma}_x|^2 - 1} \quad (16)$$

and

$$\widehat{\text{CDR}}_{\text{DOAdep}} = \left| \frac{\Gamma_s^*(\Gamma_n - \hat{\Gamma}_x)}{\text{Re}\{\Gamma_s^* \hat{\Gamma}_x\} - 1} \right|, \quad (17)$$

respectively, where  $\text{Re}\{\cdot\}$  and  $|\cdot|$  represent the real part and magnitude of  $(\cdot)$ , respectively. Note that frame and frequency index have been omitted in (16) and (17) for brevity. As can be seen from (16), this estimator does not require the TDOA of the target source, since  $\Gamma_s(l, f)$  is not required for calculating  $\widehat{\text{CDR}}_{\text{DOAindep}}(l, f)$ . On the other hand, this is not true for the estimator in (17). Hence,  $\widehat{\text{CDR}}_{\text{DOAdep}}$  is DOA-dependent. A more detailed investigation of the employed CDR estimation and a comparison to other estimators with respect to bias, robustness, and dereverberation performance for a two-channel microphone array can be found in [12, 13].

When applying the coherence-based postfilter to the output of a beamformer, two aspects need to be considered: First, since the microphone array of the CHiME-3 challenge consists of five forward-facing microphones, the CDR estimator (initially designed for a pair of microphones) has to be adapted to exploit all available microphone signals. To do so, we apply the CDR estimator in (16) or (17) to every pair of non-failing microphones to obtain the CDR estimate of each microphone pair. From each of these estimates, we calculate the respective diffuseness values as [13, 22]

$$D(l, f) = (1 + \widehat{\text{CDR}}(l, f))^{-1}. \quad (17)$$

Subsequently, we take the arithmetic average of all microphone pair-specific



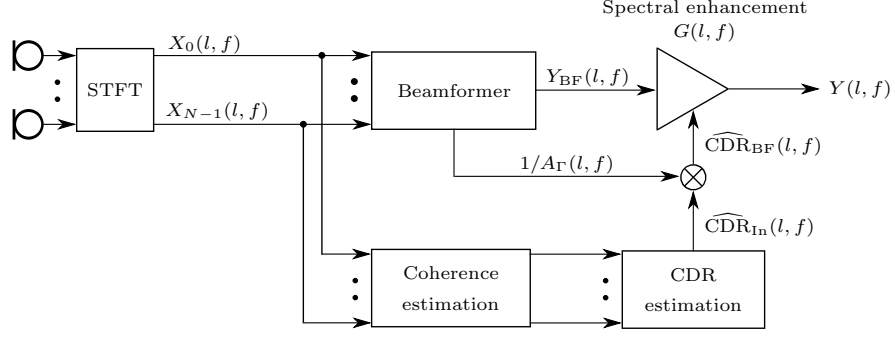


Figure 2: Illustration of the front-end signal processing consisting of beamformer and coherence-based postfilter, where the latter is applied to the beamformer output signal.

diffuseness values to obtain an average diffuseness estimate  $\overline{D}(l, f)$ . We then calculate the final CDR estimate as:

$$\widehat{\text{CDR}}_{\text{In}}(l, f) = \frac{1 - \overline{D}(l, f)}{\overline{D}(l, f)}. \quad (18)$$

Note that  $\widehat{\text{CDR}}_{\text{In}}(l, f)$  is a CDR estimate at the input of the signal enhancement system, i.e., at the microphones. However, what we need is the CDR at the output of the beamformer. This can be obtained by applying a correction factor  $A_{\Gamma}(l, f)$  to  $\widehat{\text{CDR}}_{\text{In}}(l, f)$ . Thus, the CDR estimate at the output of the beamformer  $\widehat{\text{CDR}}_{\text{BF}}(l, f)$  is defined as

$$\widehat{\text{CDR}}_{\text{BF}}(l, f) = \frac{\widehat{\text{CDR}}_{\text{In}}(l, f)}{A_{\Gamma}(l, f)}, \quad (19)$$

where  $A_{\Gamma}(l, f)$  is the beamformer gain for diffuse noise given by [23]

$$A_{\Gamma}(l, f) = \mathbf{w}^H(l, f) \mathbf{J}_{\text{diff}}(f) \mathbf{w}(l, f), \quad (20)$$

where  $(\cdot)^H$  denotes the Hermitian of a vector or matrix and  $\mathbf{J}_{\text{diff}}(f)$  is the  $N \times N$  spatial coherence matrix of a diffuse noise field with the  $(p, q)$ -th element given by (13).

110 Figure 2 shows the block-diagram of the employed front-end enhancement system, consisting of beamformer and coherence-based postfilter.

### 3. ASR back-end system

As indicated in Figure 1, we employ the ASR back-end system provided by the CHiME-3 baseline. This includes an HMM-Gaussian Mixture Model (GMM) system, consisting of 2500 tied triphone HMM states which are modelled by 15000 Gaussians, as well as an HMM-DNN-based ASR system providing state-of-the-art ASR performance for real-world scenarios. The HMM-DNN system employs a seven-layer DNN with 2048 neurons per hidden layer and is based on “Karel’s implementation“ of the Kaldi toolkit [24]: the DNN training process includes pre-training using restricted Boltzmann machines, cross entropy training, and sequence discriminative training using the state-level minimum Bayes risk (sMBR) criterion. In the updated version, the baseline system uses an RNN-based language model. For a more detailed presentation of the baseline ASR systems, we refer to [9].

### 4. Experiments

In the following, we evaluate the signal enhancement performance of our proposed front-end. First, we give an overview over the evaluation setup and our choice of front-end parameters. After this, we illustrate the impact of the proposed front-end enhancement on the STFT spectra of a noisy speech utterance. Finally, we evaluate the speech recognition accuracy achieved by the CHiME-3 baseline ASR system, incorporating our proposed front-end with DOA-independent and DOA-dependent coherence-based postfiltering.

#### 4.1. Setup and parameters

To obtain the STFT-representation, we use a Discrete Fourier Transform (DFT)-based uniform filterbank with window length 1024, Fast Fourier Transform (FFT) size 512, and downsampling factor 128 [25]. The signals were processed at a sampling rate of 16 kHz. DOA estimation and WDS beamforming were provided by the BeamformIt toolkit. For realizing the coherence-based

postfilter, we chose a gain floor  $G_{\min} = 0.1$  and optimize the overestimation factor  $\mu$ . The short-time coherence estimates  $\hat{\Gamma}_x(l, f)$  were obtained by recursive averaging of the auto- and cross-power spectra with forgetting factor  $\lambda = 0.68$ , as in [12, 13].

The ASR task included sets of real and simulated noisy utterances in four different environments: café (CAF), street junction (STR), public transport (BUS), and pedestrian area (PED). For each environment, a training set, a development set, and an evaluation set consisting of real and simulated data were provided [9].

#### 4.2. Illustration of front-end impact in the STFT domain

In Figure 3, an exemplary illustration of the impact of our proposed front-end, including WDS beamformer and coherence-based postfilter with overestimation factor  $\mu = 1.3$ , on the STFT spectra of a noisy utterance is shown, with frame  $l$  and frequency  $f$  on the horizontal and vertical axis, respectively. The coherence-based postfilter was realized using DOA-dependent CDR estimation (17). As a reference signal, the spectrum of the close-talking microphone (channel 0) is shown in Figure 3(a). It contains the desired speech signal plus little background noise. The desired signal is a male speaker saying “*Our guess is no*” in a café environment. The spectrum of microphone channel 1 is illustrated in Figure 3(b). As can be seen, low- as well as high-frequency noise is acquired by the microphone, whereas most of the noise is present in the frequency range of speech. Applying the baseline beamformer already leads to a significant reduction of the interfering components, as illustrated in Figure 3(c). A thorough comparison of Figure 3(c) with Figure 3(d) shows that applying the coherence-based postfilter to the beamformer output further reduces the interference across the entire frequency range. The estimated diffuseness  $\overline{D}(l, f)$  at the microphones is illustrated in Figure 3(e). Comparing Figure 3(e) with the spectrogram of the reference signal in 3(a) shows that  $\overline{D}(l, f)$  exhibits low values whenever the target source is active. Furthermore, it can be seen that  $\overline{D}(l, f)$  is large whenever the target source is not active, indicating a high level of diffuse noise

components. A final comparison of Figures 3(a) and 3(d) reveals the similarity  
 170 between the front-end output signal  $Y(l, f)$  and the close-talking microphone  
 signal  $S(l, f)$ , which indicates the effectiveness of the proposed front-end signal  
 enhancement technique.

In Figure 4, we illustrate the estimated diffuseness  $\overline{D}(l, f)$  obtained with the  
 DOA-dependent estimator (17) in Figure 4(b) and with the DOA-independent  
 175 estimator (16) in Figure 4(c), to highlight the difference between the two esti-  
 mators. As a reference, the spectrogram of the reference signal is shown in Fig-  
 ure 4(a), again. It can be seen that both estimators yield an estimated diffuse-  
 ness with a similar structure, i.e., the diffuseness is of low value when the target  
 signal is active and vice versa. However, it can be seen that the DOA-dependent  
 180 estimator in general attributes higher diffuseness values to time-frequency re-  
 gions where the target signal is not active. A closer look also reveals that in  
 the diffuseness estimate which was obtained with the DOA-dependent estimator  
 (Figure 4(b)), target source regions in the higher frequency range ( $f > 4\text{kHz}$ ) ex-  
 hibit a larger diffuseness than in the lower frequency range ( $f \leq 4\text{kHz}$ ). Finally,  
 185 it can be observed that the diffuseness obtained with the DOA-independent es-  
 timator (Figure 4(c)) takes larger values in the very low frequencies than the  
 diffuseness obtained with the DOA-dependent estimator (Figure 4(b)).

#### 4.3. Evaluation of recognition accuracy

In the following, we evaluate the recognition accuracy obtained with the  
 190 CHiME-3 baseline ASR system incorporating our proposed coherence-based  
 Wiener filter, realized using ...

- DOA-independent CDR estimation in (16), termed “WF w/o DOA”,
- DOA-dependent CDR estimation in (17), referred to as “WF with DOA”.

In Figure 5, the resulting WERs for the CHiME-3 task (averaged over all real  
 195 and simulated data sets) are illustrated for different values of the overestimation  
 factor  $\mu$ . It is obvious that the recognition accuracy of the CHiME-3 baseline  
 ASR system without postfilter (“No WF” in Figure 5) is consistently improved




Figure 3: Illustration of the impact of front-end signal processing on the recorded noisy microphone signal. Spectrograms of (a) recorded close-talking desired signal  $20 \log_{10}(|S(l, f)|)$ , of (b) microphone signal  $20 \log_{10}(|X_1(l, f)|)$ , of (c) baseline beamformer output signal  $20 \log_{10}(|Y_{BF}(l, f)|)$ , and of (d) postfilter output signal  $20 \log_{10}(|Y(l, f)|)$ . Figure (e) shows the average diffuseness  $\overline{D}(l, f)$ , estimated from the microphone signals using the DOA-dependent estimator (17).

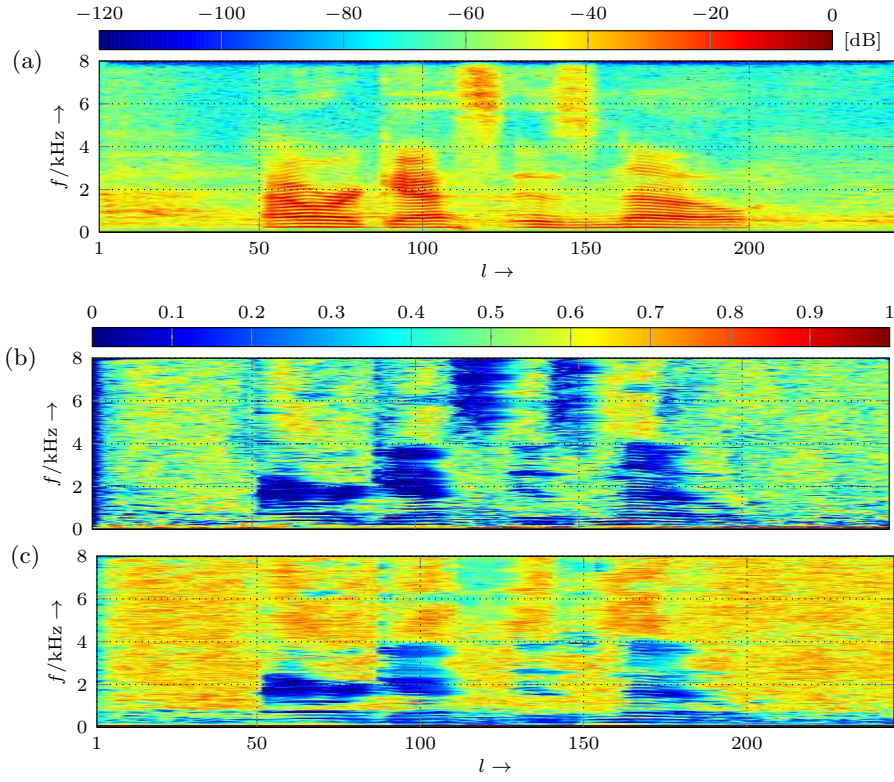


Figure 4: Comparison of the behaviour of the two different CDR estimators. Figure (a) shows the spectrograms of the recorded close-talking signal  $20 \log_{10}(|S(l, f)|)$ . Figures (b) and (c) show the average diffuseness, estimated using the DOA-independent and DOA-dependent estimator in (16) and (17), respectively.




Figure 5: Average WER scores achieved by the CHiME-3 baseline ASR system without (“No WF”) and with coherence-based postfiltering of the beamformer output signal using DOA-independent (“WF w/o DOA”) and DOA-dependent (“WF with DOA”) CDR estimation.

by incorporating coherence-based postfiltering. Moreover, it can be seen that the Wiener filter based on DOA-dependent (“WF with DOA”) CDR estimation significantly outperforms the Wiener filter based on DOA-independent (“WF w/o DOA”) CDR estimation. The best WERs for both coherence-based postfilter are achieved with overestimation factor  $\mu = 1.3$ .

Table 1 presents the WERs (in %) for the CHiME-3 development and evaluation data set, obtained with baseline front-end enhancement (“No PF”) and with the extended front-end containing the DOA-independent (“WF w/o DOA”) and DOA-dependent (“WF with DOA”) postfilter, respectively. For both realisations of the postfilter, the overestimation factor was set to  $\mu = 1.3$ . The results show that the proposed coherence-based postfiltering of the beamformer output signal consistently improves the recognition of the CHiME-3 baseline ASR system for both, simulated and real data from the development set as well as from the evaluation set. It can also be seen that the DOA-dependent postfilter on average always outperforms the DOA-independent postfilter, as illustrated by the boldface numbers.

Table 2 compares the scenario-specific WERs (in %) for the CHiME-3 evaluation set (real and simulated data), obtained with the baseline front-end as well as the extended front-end using DOA-independent and DOA-dependent

Table 1: Average WERs (in %) for the CHiME-3 development and evaluation set achieved by the CHiME-3 baseline ASR system without (“No WF”) and with coherence-based post-filtering of the beamformer output signal using DOA-independent (“WF w/o DOA”) and DOA-dependent (“WF with DOA”) CDR estimation ( $\mu = 1.3$  for both cases).

	Development set		Evaluation set	
	Real data	Sim. data	Real data	Sim. data
No WF	6.2	7.5	12.6	12.7
WF w/o DOA	6.0	7.3	12.0	12.2
WF with DOA	<b>5.8</b>	<b>6.8</b>	<b>11.9</b>	<b>11.2</b>

Table 2: Room-specific WERs (in %) for the CHiME-3 evaluation set achieved by the CHiME-3 baseline ASR system without (“No WF”) and with coherence-based postfiltering of the beamformer output signal using DOA-independent (“WF w/o DOA”) and DOA-dependent (“WF with DOA”) CDR estimation ( $\mu = 1.3$  for both cases).

	Real data				Sim. data			
	BUS	CAF	PED	STR	BUS	CAF	PED	STR
No WF	18.5	11.4	10.5	10.3	8.6	13.4	13.5	14.7
WF w/o DOA	18.0	9.8	10.0	<b>10.2</b>	8.4	13.3	11.9	15.3
WF with DOA	<b>17.3</b>	<b>9.6</b>	<b>9.9</b>	10.8	<b>7.8</b>	<b>12.2</b>	<b>11.2</b>	<b>13.7</b>

postfilter. As for the results in Table 1, for both postfilter realisations the over-estimation factor was set to  $\mu = 1.3$ . It can be seen that except for real data recorded in the STR scenario, coherence-based postfiltering significantly reduces the WERs of the baseline front-end. In addition, including information on the DOA of the target source into the CDR estimation proves to be beneficial for the majority of all scenarios as well.



## 5. Conclusion

We proposed to extend the front-end speech enhancement of the CHiME-3 baseline ASR system by a coherence-based postfilter. The postfilter is realized as a Wiener filter, where an estimate of the ratio between direct and diffuse signal components at the output of the baseline beamformer is used as an approximation of the short-time SNR to compute the time- and frequency-dependent postfilter gains. To estimate the ratio between direct and diffuse signal components, we used a DOA-independent and a DOA-dependent estimator, which can be efficiently realized by estimating the auto- and cross-power spectra at the microphone signals. As a consequence, the postfilter has a very low computational complexity. Both the baseline and the extended front-end speech enhancement have been evaluated on real and simulated data with respect to WERs using the baseline HMM-DNN-based ASR system. The results confirmed that the proposed coherence-based postfilter significantly improves the recognition accuracy of the enhanced speech compared to the CHiME-3 baseline system for both, simulated and real data. The improved recognition accuracy in addition to the low computational complexity makes the proposed postfilter very suitable for real-time robust distant speech recognition. As future work, using spatial diffuseness features as an additional input to a DNN-based acoustic model, as proposed in [26], should be investigated.

## 6. Acknowledgement

We would like to thank Stefan Meier and Christian Hofmann for their continuous support and fruitful discussions.

The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 609465 and from the Deutsche Forschungsgemeinschaft (DFG) under contract number KE 890/4-2.

## 250 References

- [1] M. Delcroix, Y. Kubo, T. Nakatani, A. Nakamura, Is speech enhancement pre-processing still relevant when using deep neural networks for acoustic modeling?, in: Interspeech, 2013, pp. 2992–2996.
- [2] T. Yoshioka, M. Gales, Environmentally robust ASR front-end for deep  
255 neural network acoustic models, *Computer Speech and Language (CSL)* 31 (1) (2015) 65–86.
- [3] I. Cohen, Noise spectrum estimation in adverse environments: improved minima controlled recursive averaging, *IEEE Trans. Speech Audio Process. (SAP)* 11 (5) (2003) 466–475.
- [4] A. Krueger, R. Haeb-Umbach, Model-based feature enhancement for re-  
260 verberant speech recognition, *IEEE Trans. Audio, Speech Lang. Process. (ASLP)* 18 (7) (2010) 1692–1707.
- [5] M. Gales, Y.-Q. Wang, Model-based approaches to handling additive noise in reverberant environments, in: *Proc. IEEE Joint Workshop Hands-free  
265 Speech Comm. Microphone Arrays (HSCMA)*, IEEE, 2011, pp. 121–126.
- [6] X. Li, J. Bilmes, Regularized adaptation of discriminative classifiers, in: *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Process. (ICASSP)*, IEEE, 2006, pp. 237–240.
- [7] H. Liao, Speaker adaptation of context dependent deep neural networks,  
270 in: *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Process. (ICASSP)*, IEEE, 2013, pp. 7947–7951.
- [8] D. Yu, K. Yao, H. Su, G. Li, F. Seide, KL-divergence regularized deep neural network adaptation for improved large vocabulary speech recognition, in: *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Process. (ICASSP)*,  
275 IEEE, 2013, pp. 7893–7897.

- [9] J. Barker, R. Marxer, E. Vincent, S. Watanabe, The third 'CHiME' speech separation and recognition challenge: Dataset, task and baselines., in: IEEE Workshop Automat. Speech Recog., Understanding (ASRU), IEEE, 2015, pp. 504–511.
- 280 [10] X. Anguera, C. Wooters, J. Hernando, Acoustic beamforming for speaker diarization of meetings, IEEE Trans. Audio, Speech Lang. Process. (ASLP) 15 (7) (2007) 2011–2022.
- [11] E. J. Diethorn, Acoustic signal processing for telecommunication, Kluwer Academic Publishers, Hingham, MA, USA, 2000, Ch. Subband Noise Reduction Methods for Speech Enhancement, pp. 155–179.
- 285 [12] A. Schwarz, W. Kellermann, Unbiased coherent-to-diffuse ratio estimation for dereverberation, in: Proc. IEEE Int. Workshop Acoustic Echo, Noise Control (IWAENC), IEEE, 2014, pp. 6–10.
- [13] A. Schwarz, W. Kellermann, Coherent-to-diffuse power ratio estimation for dereverberation, IEEE Trans. Audio, Speech Lang. Process. (ASLP) (2015) 1006–1018.
- 290 [14] H. Van Trees, Detection, Estimation, and Modulation Theory, Optimum Array Processing, Detection, Estimation, and Modulation Theory, Wiley, 2004.
- [15] B. D. V. Veen, K. M. Buckley, A versatile approach to spatial filtering, IEEE ASSP Magazine 5 (2) (1988) 4–24.
- 295 [16] M. S. Brandstein, H. F. Silverman, A robust method for speech signal time-delay estimation in reverberant rooms, in: Proc. IEEE Int. Conf. Acoustics, Speech, Signal Process. (ICASSP), 1997, pp. 375–378.
- [17] X. Anguera, Robust speaker diarization for meetings, Ph.D. thesis, UPC Barcelona (2006).
- 300

- [18] E. Haensler, G. Schmidt, *Acoustic Echo and Noise Control: A Practical Approach*, Wiley-Interscience, 2004.
- [19] M. Jeub, C. M. Nelke, C. Beaugeant, P. Vary, Blind estimation of the  
 305 coherent-to-diffuse energy ratio from noisy speech signals, in: *Proc. 19th European Signal Processing Conference (EUSIPCO)*, 2011, pp. 1347–1351.
- [20] O. Thiergart, G. Del Galdo, E. Habets, Signal-to-reverberant ratio estimation based on the complex spatial coherence between omnidirectional microphones, in: *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Process. (ICASSP)*, IEEE, 2012.  
 310
- [21] O. Thiergart, G. Del Galdo, E. Habets, On the spatial coherence in mixed sound fields and its application to signal-to-diffuse ratio estimation, *J. Acoust. Soc. Am. (JASA)* 132 (2012) 2337.
- [22] G. Del Galdo, M. Taseska, O. Thiergart, J. Ahonen, V. Pulkki, The diffuse  
 315 sound field in energetic analysis, *J. Acoust. Soc. Am. (JASA)* 131 (3) (2012) 2141–2151.
- [23] K. U. Simmer, J. Bitzer, C. Marro, Post-filtering techniques, in: *Microphone Arrays, Digital Signal Processing*, Springer Berlin Heidelberg, 2001, pp. 39–60.
- [24] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel,  
 320 M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, others, The Kaldi speech recognition toolkit.
- [25] M. Harteneck, S. Weiss, R. W. Stewart, Design of near perfect reconstruction oversampled filter banks for subband adaptive filters, *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*  
 325 46 (8) (1999) 1081–1085.
- [26] A. Schwarz, C. Huemmer, R. Maas, W. Kellermann, Spatial diffuseness features for DNN-based speech recognition in noisy and reverberant en-

vironments, in: Proc. IEEE Int. Conf. Acoustics, Speech, Signal Process.  
(ICASSP), IEEE, 2015, pp. 4380–4384.

330