CVX-optimized Beamforming and Vector Taylor Series Compensation with German ASR employing Star-shaped Microphone Array

Juan A. Morales-Cordovilla[1], Hannes Pessentheiner[1], Martin Hagmüller[1], José A. González[2], and Gernot Kubin[1] *

[1]Signal Processing and Speech Communication Laboratory, Graz University of Technology, Austria,
[2]Dept. of Computer Science, University of Sheffield, UK,
{moralescordovilla,hannes.pessentheiner,hagmueller,gernot.kubin}@tugraz.at,j.gonzalez@sheffield.ac.uk

Abstract. This paper addresses the problem of distant speech recognition in reverberant noisy conditions employing a star-shaped microphone array and vector Taylor series (VTS) compensation. First, a beamformer yields an enhanced single-channel signal by applying convex (CVX) optimization over three spatial dimensions given the spatio-temporal position of the target speaker as prior knowledge. Then, VTS compensation is applied over the speech features extracted from the temporal signal obtained by the beamformer. Finally, the compensated features are used for speech recognition. Due to a lack of existing resources in German to evaluate the proposed enhancement framework, this paper also introduces a new speech database. In particular, we present a medium-vocabulary German database for microphone array made of embedded clean signals contaminated with real room impulsive responses and mixed in a 'natural' way with real noises. We show that the proposed enhancement framework performs better than other related systems on the presented database.

Keywords: distant speech recognition, cvx-optimized beamforming, vector Taylor series compensation, star-shaped microphone array, reverberant and noisy environment, natural mixing, German database.

1 Introduction

The distant interaction of a speaker with a dialogue system, which controls some mechanisms of a house, is a difficult challenge because of many reasons: the wake-up of the system (distinction between simple conversations and commands), the

* This work has been supported by the European project DIRHA FP7-ICT-2011-7-288121, ASD project and the Austrian Marshall Plan Foundation.

speech variations for the automatic speech recognition (ASR), and the degradation of the speech signal due to background noise, reverberation, or the speaker position. Different projects such as CHIL, DICIT, and the currently finalized CHiME [1] have been proposed to solve this challenge but the Distant-speech Interaction for Robust Home Applications (DIRHA) European project [2] (for people with disabilities) is different from the others in the use of the microphone array technology.

To address the above problems, we propose the enhancement framework depicted in Fig. 1 which is an improved version of the one presented in [12]. Also this framework is part of the distant speech recognition systems presented in [6, 7]. This consists of a spatio-temporal localizer (ST-Localizer) which tries to find when the user is speaking and where. Later, a novel convex (CVX)-optimization–based beamformer (BF) attenuates the interference signals different from the user's direction. Finally, a vector Taylor series compensation method further increases the robustness of the ASR on the still degraded signal provided by the beamformer. In this paper, we avoid the problem of the spatio-temporal localization and focuses on the beamformer and the compensation method justifying their proposed configuration with experimental results.

This paper also introduces a new and more realistic German speech database than presented in the previous work [12] to evaluate the proposed enhancement framework. In particular, we present a medium-vocabulary German database for microphone array configuration which contains embedded clean signals contaminated with real room impulsive responses and mixed in a 'natural' way [1] with real noises.

The paper is structured as follows: sections 2 and 3 describe the CVX beamforming and VTS compensation methods respectively. Section 4 explains the proposed BAS-embeded database and the ASR configuration. Section 5 presents and analyses the experimental results, and in section 6 we summarize the most important ideas presented along the paper together with some future works.

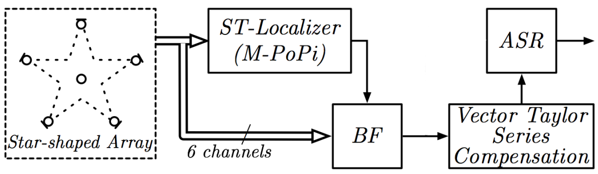

Fig. 1. Block diagram of the proposed system for distant speech recognition which consists of a 6-element star-shaped microphone array, a spatio-temporal localizer (ST-Localizer) of the speaker utterance, a beamformer (BF), a vector Taylor series compensation, and an automatic speech recognition (ASR) system.

2 Convex-optimization–based Beamformer

In our experiments, we employ a novel CVX-optimization–based beamformer. The beamformer design, first reported in [12], exhibits an improved extension of the design mentioned in [5]. The remarkable improvements of our modified beamformer are null-steering, the compatibility with different array geometries, and an optimization to three spatial dimensions. The last one is a prerequisite to enable beamforming in three spatial dimensions and to reduce the influence of reflections from the ceiling and the floor discussed in [11]. The CVX constrains the white noise gain to be larger than a lower limit γ. It considers the three-dimensional undistorted capturing response with steering direction (φ_s, θ_s) and nulls placed in different directions as constraints. The beamformer design is based on least squares computations that approximate a desired three-dimensional directivity pattern

$$\hat{b}(\omega, \varphi, \theta) = \sum_{n=1}^{N} w_n(f) e^{i \frac{\omega}{c} r_n \cdot \eta(\varphi, \theta, \varphi_n, \theta_n)}$$

with

$$\eta(\varphi, \theta, \varphi_n, \theta_n) = \sin(\theta)\sin(\theta_n)\cos(\varphi - \varphi_n) + \cos(\theta)\cos(\theta_n),$$

or, in vector notation,

$$\hat{\mathbf{B}}(\omega) = \mathbf{G}(\omega) \cdot [\mathbf{w}(\omega) \otimes \mathbf{I}],$$

where f and ω represent the linear and angular frequency, φ and θ are steering-direction–dependent azimuthal and elevation angles, φ_n and θ_n are the angles of a microphone with index n, N is the number of microphones, c is the sound velocity, r_n is the distance between a microphone and the center of the coordinate system, and $\mathbf{w}(\omega) = (w_1(\omega), w_2(\omega), ..., w_N(\omega))^T$ is the beamformer coefficient vector. Moreover, $\mathbf{I}$ is the identity matrix, $\otimes$ denotes the Kronecker product, and $\mathbf{G}(\omega)$ is an $(N_\theta \times [N \cdot N_\varphi])$ capturing response matrix according to $G_{l,m,n}(\omega) = e^{i \frac{\omega}{c} r_n \cdot \eta(\varphi_m, \theta_l, \varphi_n, \theta_n)}$, where N_φ is the number of discretized azimuthal angles φ_m, and N_θ is the number of discretized elevation angles θ_l. The beamformer assumes the same desired response for all frequencies, i.e. $\hat{\mathbf{B}}(\omega) = \hat{\mathbf{B}}$, and

$$\arg\min_{\mathbf{w}(\omega)} \|\mathbf{G}(\omega) \cdot [\mathbf{w}(\omega) \otimes \mathbf{I}] - \hat{\mathbf{B}}\|_F$$

subjected to the white noise gain (WNG), the undistorted capturing response with steering direction (φ_s, θ_s), and the optional null-placement constraints

$$\frac{|\mathbf{w}^T(\omega)\mathbf{d}(\omega)|^2}{\mathbf{w}^H(\omega)\mathbf{w}(\omega)} \geq \gamma, \quad \mathbf{w}^H(\omega)\mathbf{d}(\omega) = 1, \quad \mathbf{w}^H(\omega)\mathbf{V}(\omega) = \mathbf{0},$$

where $\mathbf{d}(\omega) = (d_1(\omega), d_2(\omega), ..., d_N(\omega))^T$ represents the capturing response with steering direction (φ_s, θ_s), and $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, ..., \mathbf{v}_S]$ is a matrix which consists of vectors $\mathbf{v}(\omega) = (v_1(\omega), v_2(\omega), ..., v_{M-1}(\omega))^T$ that describe the capturing response of, e.g., competing speakers or other noise sources, S is the number of nulls, $(\cdot)^T$ is the transpose, $(\cdot)^H$ is the Hermitian-transpose, and $\| \cdot \|_F$ is the Frobenius norm. We set the lower limit γ and the desired response $\hat{\mathbf{B}}$ in a way that we were able to distribute the narrow null-lobe marked in Fig. 2 over frequencies below 1000 Hz. This yields a decreased main-lobe width at lower frequencies without increasing the width at higher ones. Although null-steering is one of the beamformer's big improvements, we did not consider it due to the assumption of unknown noise source positions in our experiments.

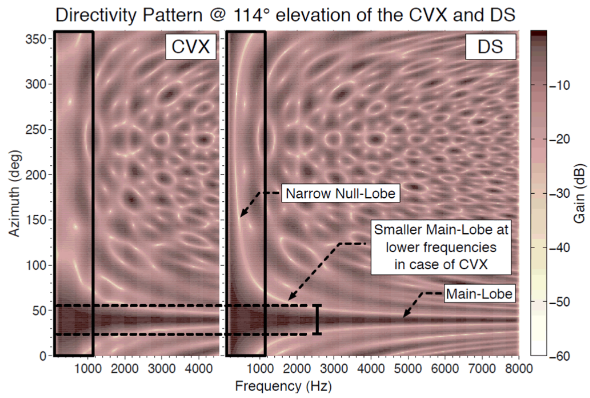

Fig. 2. The directivity patterns of the CVX without null-steering and the DS (delay-and-sum, [14]) based on a 6-element star-shaped array with steering direction $\phi_s = 40°$ and $\theta_s = 114°$.

3 Vector Taylor Series Compensation

After applying a beamformer, which yields a single-channel signal, a vector Taylor series (VTS) compensation [10] is used to further enhance the signal and the robustness of ASR. The reason of using VTS rather than other methods, such as marginalization missing data is that it let the final representation of the clean estimated signal be in the cepstral domain, which is a more appropriate representation for a medium or large vocabulary task. In this paper, we apply VTS in the log-Mel domain (i.e. the log-outputs of the Mel filters) and later we apply the cepstrum transformation (Sec. 4.2).

Let $\boldsymbol{y}_t$, $\boldsymbol{x}_t$ and $\boldsymbol{n}_t$ be the feature vectors at time t for the noisy speech, clean speech, and noise signals, respectively, expressed in this domain. Given the noisy observation $\boldsymbol{y}_t$, VTS estimates the clean feature vector as follows,

$$\hat{\boldsymbol{x}}_t = \boldsymbol{y}_t - \sum_{k=1}^{K} P(k|\boldsymbol{y}_t)\boldsymbol{g}\left(\boldsymbol{\mu}_X^{(k)}, \hat{\boldsymbol{n}}_t\right), \tag{1}$$

where $\hat{\boldsymbol{n}}_t$ is the noise estimate at time t and $\boldsymbol{g}(\boldsymbol{x}, \boldsymbol{n}) = \log(\mathbf{1} + \exp(\boldsymbol{n} - \boldsymbol{x}))$ is the so-called mismatch function. To derive the above estimator, a Gaussian mixture model (GMM) with K components is used as the prior speech model. Thus,

$$p(\boldsymbol{x}) = \sum_{k=1}^{K} \pi_X^{(k)} \mathcal{N}\left(\boldsymbol{x}; \boldsymbol{\mu}_X^{(k)}, \boldsymbol{\Sigma}_X^{(k)}\right), \tag{2}$$

with $\pi_X^{(k)}$, $\boldsymbol{\mu}_X^{(k)}$, and $\boldsymbol{\Sigma}_X^{(k)}$ being the parameters of the kth Gaussian component, i.e., its prior probability, mean vector and covariance matrix.

Finally, the noisy speech model $p(\boldsymbol{y}_t)$ is required for computing the posterior probabilities $P(k|\boldsymbol{y}_t)$ in (1). To obtain this model, the clean speech GMM is adapted as follows,

$$\boldsymbol{\mu}_{Y,t}^{(k)} = \boldsymbol{\mu}_X^{(k)} + \boldsymbol{g}\left(\boldsymbol{\mu}_X^{(k)}, \hat{\boldsymbol{n}}_t\right), \tag{3}$$

$$\boldsymbol{\Sigma}_{Y,t}^{(k)} = \boldsymbol{J}_t^{(k)} \boldsymbol{\Sigma}_X^{(k)} \boldsymbol{J}_t^{(k)} + (\boldsymbol{I} - \boldsymbol{J}_t^{(k)}) \boldsymbol{\Sigma}_{N,t} (\boldsymbol{I} - \boldsymbol{J}_t^{(k)}), \tag{4}$$

where $\boldsymbol{\Sigma}_{N,t}$ is the covariance matrix associated to the noise estimate $\hat{\boldsymbol{n}}_t$ and $\boldsymbol{J}_t^{(k)}$ is a diagonal matrix whose elements are given by,

$$\boldsymbol{J}_t^{(k)} = \mathrm{diag}\left(\frac{1}{1 + \exp\left(\hat{\boldsymbol{n}}_t - \boldsymbol{\mu}_X^{(k)}\right)}\right). \tag{5}$$

4 Experimental Framework

4.1 Embedded-BAS Database

Due to a lack of existing resources in German to evaluate the proposed enhancement framework, this paper also introduces a new German database for a star-shaped microphone array. More precisely, this array consists of 6 microphones (1 at the center and 5 on the circle) placed on the ceiling of the living room of the ITEA apartment used by Fondazione Bruno Kessler (FBK) for the DIRHA project [2] (see Fig. 3).

Embedded noisy signals Each test multi-chanel signal of this database represents what the microphone array would record if a speaker, in the presence of noise, repeated the action of pronouncing an isolated utterance at a specific position in the room and later moved to another position to pronounce another utterance. We call to this connection of utterances with continuous background

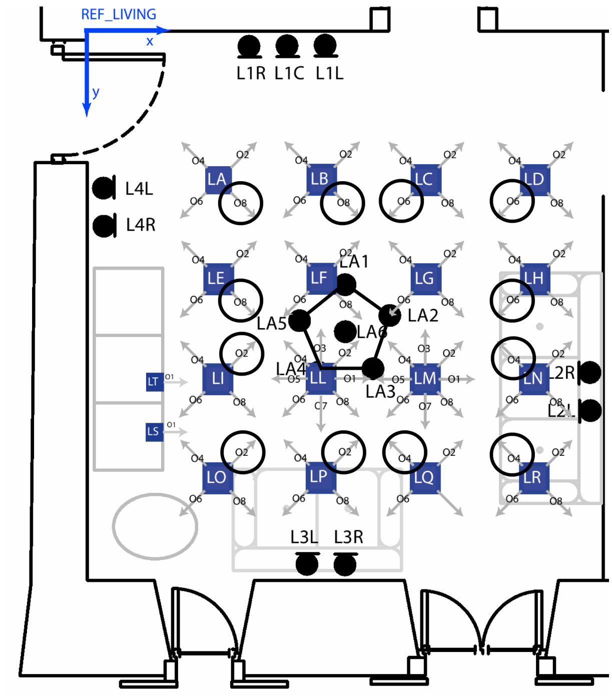

Fig. 3. Living room of the ITEA apartment of Fondazione Bruno Kessler (FBK) with the microphone array at the center and the 12 speaker position/directions employed in this work [provided by FBK].

Table 1. Word accuracies obtained by different configurations of the proposed systems tested over the presented Embedded-BAS database for different SNR values.

Systems	Clean	10 dB	0 dB	Average
Baseline (central microphone)	93.24	79.34	43.69	72.09
DS Beamforming	94.73	83.61	51.73	76.69
CVX Beamforming	95.34	83.65	51.98	76.99
Baseline + VTS (FLF noise)	91.60	84.00	53.61	76.40
DS Beamf. + VTS (FLF noise)	93.82	86.83	55.70	78.78
CVX Beamf. + VTS (FLF noise)	93.60	87.29	60.20	80.36
Baseline + VTS (Oracle noise)	92.93	91.75	79.10	87.93
DS Beamf. + VTS (Oracle noise)	94.67	92.42	79.04	88.71
CVX Beamf. + VTS (Oracle noise)	95.19	93.54	80.32	89.68

noise and with different reverberations, which depend on the speaker position, *embedded noisy signal* .

For the controllability of the experiments, the next 12 speaker position/directions, circled in Fig. 3, are only used: (LA/O8, LB/O8, etc.). To simulate the different SNR noisy conditions in the most possible 'natural' way, we follow the indications of SNR mixture of the CHiME corpus [1] by employing around 3-hours of real noise, recorded by the FBK group with this microphone array. The way to obtain an embedded noisy signal for a target SNR is summarized in the following steps:

1. We randomly select 7 isolated monaural clean (without reverberation) utterances of one speaker, convolve them with the corresponding impulse responses (obtained by the FBK group) of 7 random speaker position/directions and obtain a 6-channel embedded clean-reverberant signal by connecting them with a time gap in the middle. These gaps are randomly selected between 0.5 and 5 seconds.
2. We randomly select a segment from all available segments, of the 3-hours of noise, which yields the target SNR within an error of 1.5 dB. The following formula is used for the SNR:

$$SNR = 10log_{10}\frac{Ex_{central}}{En_{central}}(dB) \tag{6}$$

   where $Ex_{central}$ and $En_{central}$ are the whole energy of the central microphone of the embedded clean-reverberant signal and of the noise segment respectively. If no noise segment is found that yields the target SNR, all channels of the embedded clean-reverberant signal are multiplied by a gain (which depends on the closest found SNR to the target SNR) to find at least an appropriate noise segment.
3. The final *embedded noisy signal* is the sum of this embedded clean-reverberant signal with the selected noise segment. In addition, sometimes this sum can produce a saturated signal in some of the channels. In order to avoid

this problem we multiply all the channels of both, the embedded clean-reverberant signal and the noise, by a second factor which avoids this problem.

Database description The proposed *Embedded-BAS* database exhibits a sampling frequency of 16 kHz and employs the clean sentences of the Bavarian Archive for Speech Signals (BAS) PHONDAT-1 database [13] as its isolated monaural clean utterances (Sec. 4.1) due to their temporal similarity with house control commands. The database consists of the training and test sets.

The training set contains 4999 clean-reverberant isolated utterances corresponding to 50 different-gender speakers (around 100 sentences per speaker) with a reverberation that corresponds to position LA/O8 of Fig. 3. The inclusion of the reverberation in the training set is to reduce the mismatch with the test set. The test set consists of 100 embedded clean-reverberant signals (700 isolated utterances, Sec. 4.1) corresponding to 100 different speakers (half of them are in the training set) contaminated at 10 and 0 dB. Both, the training and test sets share the same medium-vocabulary lexicon and grammar and consist of 1504 words which belong to around 500 different phrases.

4.2 ASR system

Both, the front-end and the back-end, have been derived from the standard recognizer employed in Aurora-4 database [4].

The front-end takes the enhanced signal and obtains mel frequency cepstrum coefficients (MFCCs) using 16 kHz sampling frequency, frame shift and length of 10 and 32 ms, 1024 frequency bins, 26 Mel channels and 13 cepstral coefficients. Then we apply cepstral mean normalization to the MFCCs. Delta and delta-delta features are also appended, obtaining a final feature vector with 39 components.

The back-end employs a transcription of the training corpus based on 34 monophones to train triphone-HMMs. This transcription has been derived from a more detailed transcription (based on 44 SAMPA-monophones) by means of a careful clustering of the less common monophones. Each triphone is modeled by a HMM of 6 states and 8 Gaussian-mixtures/state. By means of a monophone classification (created with the help of a linguistic) a tree-based clustering of the states is also applied to reduce the complexity and a lack of training data. Tree-based clustering also allows to create triphones models for the test stage which have not been observed in the training stage. We train a bigram using the training word transcription. By means of an expansion based on the grammar, the triphone transcription of the test lexicon and the triphones, we obtain the final macro HMMs for the test stage. It is important to point out that only the central microphone of the clean-reverberant training set without any enhancement (beamforming and VTS) is used to train our HMMs-models.

5 Experimental Results

Tab. 1 shows the different Word Accuracies (WAcc, %) achieved by different configurations of the proposed systems tested over the presented Embedded-BAS database for different SNR values.

The *Baseline (central microphone)* results are obtained when no enhancement is performed over the speech signals, i.e., directly the performance of the signal captured by the central microphone of the microphone-array. *DS Beamforming* and *CVX Beamforming* are the results achieved by delay-and-sum [14] and convex-optimization beamformers (Sec. 2). As mentioned in Sec. 1 , we assume that the ST-Localizer of Fig. 1 provides the oracle spatial and temporal localization of the speaker, i. e., we cut the embedded noisy signal in pieces which correspond to the isolated utterances, then each of these pieces together with its spatial position are sent directly to the beamformer. Following, we can see the results of the three previous configurations but when the VTS compensation (Sec. 3) is applied with a First-Last-Frames (FLF) noise estimation. This estimation assumes that the first and last 20 frames of the cut signal correspond to noise and these frames are used to estimate the log-Mel noise (and its corresponding covariance matrix) by means of a linear interpolation to the remaining the frames as shown in [8].

The most significant conclusions which can be drawn from the table are the follows:

1. Using beamformers, specially the CVX, always improves the recognition results (compare the 72.09 of the *Baseline* with the 76.99 % of the *CVX Beamf.*).
2. Considering VTS after applying beamformers additionally improves the results (compare the 76.99 of the *CVX Beamf.* with the 80.36 % of the *CVX Beamf. + VTS (FLF noise)*).

The results with oracle noise are only displayed to show the upper performance of this framework. We can see that we should further improve the noise estimation at 0 dB.

Other compensation mechanisms (such as missing data (MD) imputation based on binary mask) and types of noise estimations (such as pitch-based noise estimations) have been employed in [9]. Due to the techniques' sensitivities to the MD mask and to the pitch estimation errors, the performance of these techniques have been lower.

6 Conclusion and Future Work

This paper presented a system for distant speech recognition in reverberant and noisy conditions, intended to control a room with commands. The proposed system is an improved version of the system presented in [12]. The improvement consists of a recently presented beamformer based on convex optimization, the application of a single-channel enhancement algorithm based on VTS compensation and the presentation of a more realistic database for evaluations. The database consists of embedded noisy signals which represent, with 'natural' noise mixing, what the microphone array would record if the speaker was emitting German commands at different positions of the room. This database is a very suitable challenge for the spatio-temporal localization algorithm of the utterance which is our next future objective. To do it we plan to make use of the pitch information provided by the M-PoPi algorithm [3] .

References

1. H. Christensen, J. Barker, N. Ma, and P. Green. The chime corpus: A resource and a challenge for computational hearing in multisource environments. In *Interspeech*, 2010.
2. European project FP7. Distant-speech interaction for robust home applications (dirha). In *http://dirha.fbk.eu*, March 2012-2015.
3. T. Habib and H. Romsdorfer. Concurrent speaker localization using multi-band position-pitch (m-popi) algorithm with spectro-temporal pre-processing. In *Interspeech*, 2010.
4. H. G. Hirsch. Experimental framework for the performance evaluation of speech recognition front-ends of large vocabulary task. Technical report, STQ AURORA DSR, Working Group, 2002.
5. E. Mabande, A. Schad, and W. Kellermann. Design of robust superdirective beamformers as a convex optimization problem. In *ICASSP*, 2009.
6. J. A. Morales-Cordovilla, M. Hagmüller, H.Pessentheiner, and G. Kubin. Distant speech recognition in reverberant noisy conditions employing a microphone array. In *EUSIPCO*, 2014.
7. J. A. Morales-Cordovilla, H. Pessentheiner, M. M. Hagmüller, and G. Kubin. Room localization for distant speech recognition. In *Interspeech*, 2014.
8. Juan A. Morales-Cordovilla, Ning Ma, Victoria Sánchez, José L. Carmona, Antonio M. Peinado, and Jon Barker. A pitch based noise estimation technique for robust speech recognition with missing data. In *ICASSP*, pages 4808–4811, May, 22-27 2011.
9. Juan A. Morales-Cordovilla, H. Pessentheiner, M. Hagmller, P. Mowlaee, F. Pernkopf, and G. Kubin. A german distant speech recognizer based on 3d beamforming and harmonic missing data mask. In *AIA-DAGA*, 2013.
10. P. Moreno. *Speech Recognition in Noisy Environments.* PhD thesis, Carnegie Mellon University, 1996.
11. H. Pessentheiner, G. Kubin, and H. Romsdorfer. Improving beamforming for distant speech recognition in reverberant environments using a genetic algorithm for planar array synthesis. In *10th ITG Symposium on Speech Communication*, 2012.
12. H. Pessentheiner, S. Petrik, and H. Romsdorfer. Beamforming using uniform circular arrays for distant speech recognition in reverberant environments and double-talk scenarios. In *Interspeech*, 2012.
13. F. Schiel and A. Baumann. Phondat 1, corpus version 3.4. Technical report, Bavarian Archive for Speech Signals (BAS). http://www.bas.uni-muenchen.de/Bas/BasFormatseng.html, 2006.
14. I. Tashev. *Sound Capture and Processing: Practical Approaches.* John Wiley and Sons, 2009.