

RELATIVE TRANSFER FUNCTION VECTOR ESTIMATION FOR ACOUSTIC SENSOR NETWORKS EXPLOITING COVARIANCE MATRIX STRUCTURE

Wiebke Middelberg, Henri Gode, Simon Doclo

Department of Medical Physics and Acoustics and Cluster of Excellence Hearing4all,
University of Oldenburg, Germany, {wiebke.middelberg, henri.gode, simon.doclo}@uni-oldenburg.de

ABSTRACT

In many multi-microphone algorithms for noise reduction, an estimate of the relative transfer function (RTF) vector of the target speaker is required. The state-of-the-art covariance whitening (CW) method estimates the RTF vector as the principal eigenvector of the whitened noisy covariance matrix, where whitening is performed using an estimate of the noise covariance matrix. In this paper, we consider an acoustic sensor network consisting of multiple microphone nodes. Assuming uncorrelated noise between the nodes but not within the nodes, we propose two RTF vector estimation methods that leverage the block-diagonal structure of the noise covariance matrix. The first method modifies the CW method by considering only the diagonal blocks of the estimated noise covariance matrix. In contrast, the second method only considers the off-diagonal blocks of the noisy covariance matrix, but cannot be solved using a simple eigenvalue decomposition. When applying the estimated RTF vector in a minimum variance distortionless response beamformer, simulation results for real-world recordings in a reverberant environment with multiple noise sources show that the modified CW method performs slightly better than the CW method in terms of SNR improvement, while the off-diagonal selection method outperforms a biased RTF vector estimate obtained as the principal eigenvector of the noisy covariance matrix.

Index Terms— Acoustic sensor networks, relative transfer function vector, beamforming, covariance whitening

1. INTRODUCTION

Acoustic sensor networks (ASNs) with multiple spatially distributed microphone nodes are of rising interest for speech communication applications due to their ability to capture spatially diverse information [1]. This allows ASNs to be deployed, e.g., for speech enhancement [2–7], and sound source localization [8], in applications such as smart speakers or hearing aids connected with external microphones. In these applications the desired speech signal is often corrupted by background noise. To achieve noise reduction, multi-microphone algorithms like the minimum variance distortionless response (MVDR) beamformer can be used [2, 6, 9–11], requiring an estimate of the noise covariance matrix and the relative transfer function (RTF) vector of the target speaker.

In this paper, we consider an ASN where the noise component between all nodes is assumed to be uncorrelated, which is for example the case in a diffuse noise field when the distance between the nodes is large or when different nodes capture different noise sources. This results in a block-diagonal structure of the noise covariance matrix. Exploiting this covariance matrix structure, we propose two methods to estimate the RTF vector of the target speaker in an ASN with at least three nodes. The first

method involves a modification of the state-of-the-art covariance whitening (CW) method [12–14]. Instead of using the entire estimated noise covariance matrix for whitening, the proposed CW-D method considers only the diagonal blocks, which allows for efficient inversion and square-root decomposition. The CW and CW-D methods both estimate the RTF vector as the best rank-1 approximation of the whitened noisy covariance matrix, which can be achieved via an eigenvalue decomposition (EVD).

The second method only requires the noisy covariance matrix and no estimate of the noise covariance matrix. Assuming uncorrelated noise between nodes, all information required for RTF vector estimation is contained in the off-diagonal blocks of the noisy covariance matrix. In the off-diagonal selection (ODS) method, an optimization problem is formulated to estimate the entire RTF vector using only the internode correlations of the noisy covariance matrix. Contrary to the first method, the solution of this optimization problem cannot be computed via an EVD, and we propose to use an iterative optimization procedure.

In the experimental evaluation with reverberant real-world recordings and multiple noise sources, the performance is evaluated in terms of RTF vector estimation accuracy and signal-to-noise ratio (SNR) improvement when applying the RTF vector estimates in an MVDR beamformer. The results show that the proposed CW-D method performs slightly better than the CW method. In addition, the proposed ODS method outperforms a biased estimator using the EVD of the entire noisy covariance matrix, especially at low input SNRs where the influence of noise on the diagonal blocks is most severe.

2. SIGNAL MODEL AND NOTATION

We consider an ASN with N spatially distributed nodes, where node $n \in \{1, \dots, N\}$ contains M_n microphones, i.e., in total $M = \sum_{n=1}^N M_n$ microphones. The considered acoustic scene consists of a single target speaker and undesired ambient noise. The noisy m -th microphone signal of the n -th node can be written in the short-time Fourier transform (STFT) domain as

$$Y_{n,m}(k,l) = X_{n,m}(k,l) + V_{n,m}(k,l), \quad (1)$$

where k is the frequency bin index and l is the frame index, which - for the sake of brevity - are omitted in the remainder of this paper wherever possible. The speech and noise signal components are denoted by $X_{n,m}$ and $V_{n,m}$, respectively. The M_n -dimensional signal vector for the n -th node is defined as

$$\mathbf{y}_n = [Y_{n,1}, Y_{n,2}, \dots, Y_{n,M_n}]^T, \quad (2)$$

where $\{\cdot\}^T$ denotes the transpose operator. By stacking all node-wise signal vectors, the M -dimensional signal vector \mathbf{y} , containing all microphone signals in the ASN, is defined as

$$\mathbf{y} = [\mathbf{y}_1^T, \mathbf{y}_2^T, \dots, \mathbf{y}_N^T]^T. \quad (3)$$

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - Project ID 352015383 (SFB 1330 B2) and Project ID 390895286 (EXC 2177/1).

The speech vectors \mathbf{x}_n and \mathbf{x} and the noise vectors \mathbf{v}_n and \mathbf{v} are defined similarly to (2) and (3), respectively. For the speech component, we assume a multiplicative transfer function model [15], allowing the speech vector to be written as

$$\mathbf{x} = \mathbf{h}X_{\text{ref}}, \quad (4)$$

where $\mathbf{h} \in \mathbb{C}^M$ is the target RTF vector, which relates the speech component in the reference microphone X_{ref} to the speech component in all other microphones. Hence, the entry of \mathbf{h} corresponding to the reference microphone is equal to 1. It should be noted that the reference microphone is chosen for the entire ASN and not per node. The RTF vector for node n is defined as $\mathbf{h}_n \in \mathbb{C}^{M_n}$, such that $\mathbf{h} = [\mathbf{h}_1^T, \mathbf{h}_2^T, \dots, \mathbf{h}_N^T]^T$, where all \mathbf{h}_n are normalized to the same reference.

Assuming that the speech and noise signals are mutually uncorrelated, the noisy covariance matrix \mathbf{R}_y can be written in terms of the speech covariance matrix \mathbf{R}_x and the noise covariance matrix \mathbf{R}_v as

$$\mathbf{R}_y = \mathcal{E}\{\mathbf{y}\mathbf{y}^H\} = \mathbf{R}_x + \mathbf{R}_v, \quad (5)$$

where $\mathcal{E}\{\cdot\}$ denotes the expectation operator and $\{\cdot\}^H$ denotes the Hermitian transpose operator. The node-wise covariance matrices for the n -th node, $\mathbf{R}_{y,n}$, $\mathbf{R}_{x,n}$ and $\mathbf{R}_{v,n}$, are defined similarly to (5), using the node-wise vectors \mathbf{y}_n , \mathbf{x}_n and \mathbf{v}_n , respectively. Using (4), the speech covariance matrix \mathbf{R}_x can be written as a rank-1 matrix spanned by the RTF vector \mathbf{h} , i.e.,

$$\mathbf{R}_x = \phi_x \mathbf{h}\mathbf{h}^H, \quad (6)$$

where $\phi_x = \mathcal{E}\{|X_{\text{ref}}|^2\}$ denotes the speech power spectral density (PSD) in the reference microphone.

In this paper, we make the central assumption that the noise component is uncorrelated between different nodes. This assumption is realistic, e.g., for a diffuse noise field when the distance between the nodes is large enough [6, 7, 16] or when nodes capture different noise sources. For the noise correlation between different microphones within each node, no assumption is made, implying that within each node the noise component may be partially correlated. The node-wise noise covariance matrices $\mathbf{R}_{v,n}$ are assumed to be full-rank. Figure 1 schematically depicts the structure of the entire noise covariance matrix \mathbf{R}_v , where yellow indicates high correlation (within each node), and white indicates low correlation (between the nodes). To visualize the influence of such a block-diagonal noise covariance matrix on the noisy covariance matrix, Figure 1 also depicts the rank-1 speech covariance matrix \mathbf{R}_x and the resulting noisy covariance matrix \mathbf{R}_y . It can clearly be seen that \mathbf{R}_y only contains information about the target RTF vector in its off-diagonal blocks (orange and green), as they are unaffected by the noise. Note that for $N=2$ nodes, the off-diagonal block of \mathbf{R}_y only contains information about scaled versions of \mathbf{h}_1 and \mathbf{h}_2 , which cannot be unified into the vector $\mathbf{h} = [\mathbf{h}_1^T, \mathbf{h}_2^T]^T$, as the speech PSD ϕ_x (see (6)) evokes a scaling ambiguity for the RTF vector part that does not contain the reference microphone. This scaling ambiguity, however, can be lifted when $N \geq 3$, since direct information about the relative scaling of the different parts is contained in adjacent off-diagonal blocks.

To achieve noise reduction, we consider the MVDR beamformer, which requires an estimate of the noise covariance matrix $\hat{\mathbf{R}}_v$ and an estimate of the RTF vector $\hat{\mathbf{h}}$. The filter vector \mathbf{w} of the MVDR beamformer is given by [9, 10]

$$\mathbf{w} = \frac{\hat{\mathbf{R}}_v^{-1} \hat{\mathbf{h}}}{\hat{\mathbf{h}}^H \hat{\mathbf{R}}_v^{-1} \hat{\mathbf{h}}}, \quad (7)$$

yielding the output signal $Z = \mathbf{w}^H \mathbf{y}$ when applied to the noisy input signals. The filtered speech and noise components are defined as $Z_x = \mathbf{w}^H \mathbf{x}$ and $Z_v = \mathbf{w}^H \mathbf{v}$, respectively.

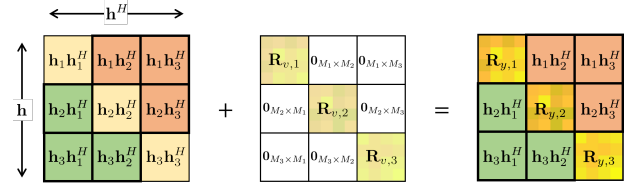


Figure 1: Visualization of the influence of noise on the noisy covariance matrix for $N=3$ nodes (with $\phi_x=1$). Left: Speech covariance matrix spanned by the target RTF vector. Middle: Noise covariance matrix, uncorrelated between nodes. Right: Noisy covariance matrix, containing only information about the target RTF vector on its off-diagonal blocks.

3. RTF VECTOR ESTIMATION METHODS

In this section, we present different RTF vector estimation methods, where we first discuss the general idea of the rank-1 approximation (Section 3.1), which is the basis for the biased estimator and the state-of-the-art CW method (Section 3.2). In Section 3.3, we propose a modification of the CW method by exploiting the assumed block-diagonal structure of the noise covariance matrix. In Section 3.4, a novel cost function is proposed, where only the off-diagonal blocks of the noisy covariance matrix are selected to compute the best fitting RTF vector.

3.1. Rank-1 Approximation and Biased Estimator

To motivate the RTF vector estimation methods in the following sections, we first consider the case where an estimate of the speech covariance matrix $\hat{\mathbf{R}}_x$ is available. In practice, the rank-1 model in (6) does not perfectly hold, e.g., due to an insufficient frame length. Hence, the RTF vector $\hat{\mathbf{h}}$ is often estimated as the best rank-1 approximation of $\hat{\mathbf{R}}_x$ [13], i.e., the vector solving the optimization problem

$$\min_{\hat{\mathbf{h}}'} \underbrace{\|\hat{\mathbf{R}}_x - \hat{\mathbf{h}}' \hat{\mathbf{h}}'^H\|_F^2}_{J(\hat{\mathbf{h}}')}, \quad (8)$$

where $\hat{\mathbf{h}}'$ is a scaled (non-normalized) version of $\hat{\mathbf{h}}$ and $\|\cdot\|_F$ denotes the Frobenius norm. Considering the gradient of the cost function $J(\hat{\mathbf{h}}')$ in (8), i.e.,

$$\nabla J(\hat{\mathbf{h}}') = -2(\hat{\mathbf{R}}_x - \hat{\mathbf{h}}' \hat{\mathbf{h}}'^H) \hat{\mathbf{h}}', \quad (9)$$

and setting it to zero, the solutions can be found by solving the eigenvalue problem

$$\hat{\mathbf{R}}_x \hat{\mathbf{h}}' = \hat{\mathbf{h}}' (\hat{\mathbf{h}}'^H \hat{\mathbf{h}}'). \quad (10)$$

Hence, the best rank-1 approximation of $\hat{\mathbf{R}}_x$ is a scaled version of the principal eigenvector \mathbf{v}_{max} , corresponding to the maximum eigenvalue σ_{max} . The RTF vector estimate can then be obtained as

$$\hat{\mathbf{h}} = \frac{\mathbf{v}_{\text{max}}}{\mathbf{e}_{\text{ref}}^T \mathbf{v}_{\text{max}}}, \quad (11)$$

where \mathbf{e}_{ref} is a selection vector containing all zeros except for the entry corresponding to the reference microphone, which equals 1.

If only the noisy covariance matrix $\hat{\mathbf{R}}_y$ is available, a biased estimate may be obtained as the best rank-1 approximation of $\hat{\mathbf{R}}_y$, i.e.,

$$\min_{\hat{\mathbf{h}}'} \|\hat{\mathbf{R}}_y - \hat{\mathbf{h}}' \hat{\mathbf{h}}'^H\|_F^2, \quad (12)$$

where the bias obviously is larger for lower SNR. The biased RTF vector estimate $\hat{\mathbf{h}}^{\text{B}}$ is obtained as in (11), where \mathbf{v}_{max} is the principal eigenvector of $\hat{\mathbf{R}}_y$.

3.2. Covariance Whitening (CW)

To compensate for the influence of the ambient noise on the RTF vector estimate, a frequently used approach is to perform whitening of the noisy covariance matrix using an estimate of the noise covariance matrix, i.e., $\hat{\mathbf{R}}_y^w = \hat{\mathbf{R}}_v^{-1/2} \hat{\mathbf{R}}_y \hat{\mathbf{R}}_v^{-H/2}$, where $\hat{\mathbf{R}}_v = \hat{\mathbf{R}}_v^{1/2} \hat{\mathbf{R}}_v^{H/2}$ corresponds to a square-root decomposition, e.g., the Cholesky decomposition [12–14].

Similarly to (8), the optimization problem in the whitened domain is given by

$$\min_{\hat{\mathbf{h}}'} \|\hat{\mathbf{R}}_y^w - \hat{\mathbf{R}}_v^{-1/2} \hat{\mathbf{h}}' \hat{\mathbf{h}}'^H \hat{\mathbf{R}}_v^{-H/2}\|_F^2, \quad (13)$$

i.e., the principal eigenvector \mathbf{v}_{\max} of the whitened noisy covariance matrix $\hat{\mathbf{R}}_y^w$ corresponds to a scaled version of the whitened RTF vector estimate. By de-whitening and normalizing \mathbf{v}_{\max} , the CW RTF vector estimate is obtained as

$$\hat{\mathbf{h}}^{\text{CW}} = \frac{\hat{\mathbf{R}}_v^{1/2} \mathbf{v}_{\max}}{\mathbf{e}_{\text{ref}}^T \hat{\mathbf{R}}_v^{1/2} \mathbf{v}_{\max}}. \quad (14)$$

3.3. Covariance Whitening Using Diagonal Blocks (CW-D)

To leverage the assumed block-diagonal structure of the noise covariance matrix (see Figure 1), we propose to only consider the diagonal blocks of $\hat{\mathbf{R}}_v$ in the CW method. The modified noise covariance matrix $\hat{\mathbf{R}}_v$ is constructed as a block-diagonal matrix containing the node-wise noise covariance matrices $\hat{\mathbf{R}}_{v,n}$ on its diagonal blocks (and all zeros in the off-diagonal blocks). For a block-diagonal matrix, the inverse and the square-root decomposition can be performed efficiently on the separate diagonal blocks, such that the matrix required for the whitening operation in (13) is given by

$$\hat{\mathbf{R}}_v^{-1/2} = \begin{bmatrix} \hat{\mathbf{R}}_{v,1}^{-1/2} & \mathbf{0}_{M_1 \times M_2} & \dots & \mathbf{0}_{M_1 \times M_N} \\ \mathbf{0}_{M_2 \times M_1} & \hat{\mathbf{R}}_{v,2}^{-1/2} & \dots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}_{M_N \times M_1} & \dots & \dots & \hat{\mathbf{R}}_{v,N}^{-1/2} \end{bmatrix}. \quad (15)$$

The whitened RTF vector estimate is obtained as the principal eigenvector of $\hat{\mathbf{R}}_v^{-1/2} \hat{\mathbf{R}}_y \hat{\mathbf{R}}_v^{-H/2}$, where de-whitening and normalization is performed similarly to (14) using $\hat{\mathbf{R}}_v^{1/2}$ to obtain the RTF vector estimate $\hat{\mathbf{h}}^{\text{CW-D}}$.

3.4. Off-Diagonal Selection (ODS)

In the optimization problem for the biased estimator in (12), it can directly be seen that biased information is used, as the diagonal blocks of $\hat{\mathbf{R}}_y$ contains both speech and noise information. To avoid this problem without compensating for the noise directly (as in the CW and CW-D methods), we propose the following optimization problem

$$\min_{\hat{\mathbf{h}}'} \underbrace{\|\mathbf{S} \odot (\hat{\mathbf{R}}_y - \hat{\mathbf{h}}' \hat{\mathbf{h}}'^H)\|_F^2}_{J(\hat{\mathbf{h}}')}, \quad (16)$$

where we only select the off-diagonal blocks of $\hat{\mathbf{R}}_y$ (see Figure 1) and its respective rank-1 approximation by means of the selection matrix \mathbf{S} and \odot denotes the Hadamard product, i.e., the element-wise multiplication of matrices. The selection matrix is defined as

$$\mathbf{S} = \begin{bmatrix} \mathbf{0}_{M_1 \times M_1} & \mathbf{1}_{M_1 \times M_2} & \dots & \mathbf{1}_{M_1 \times M_N} \\ \mathbf{1}_{M_2 \times M_1} & \mathbf{0}_{M_2 \times M_2} & \dots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{1}_{M_N \times M_1} & \dots & \dots & \mathbf{0}_{M_N \times M_N} \end{bmatrix}, \quad (17)$$

containing all ones except for the diagonal blocks, which contain zeros. Similarly to (9), the gradient of the cost function $J(\hat{\mathbf{h}}')$ in (16) is given by

$$\nabla J(\hat{\mathbf{h}}') = -2 \left(\mathbf{S} \odot (\hat{\mathbf{R}}_y - \hat{\mathbf{h}}' \hat{\mathbf{h}}'^H) \right) \hat{\mathbf{h}}'. \quad (18)$$

Setting the gradient equal to zero yields

$$(\mathbf{S} \odot \hat{\mathbf{R}}_y) \hat{\mathbf{h}}' = (\mathbf{S} \odot (\hat{\mathbf{h}}' \hat{\mathbf{h}}'^H)) \hat{\mathbf{h}}'. \quad (19)$$

In contrast to (10), this does not correspond to an eigenvalue problem, since the Hadamard product does not allow for further simplification. To the best of our knowledge, there is no closed-form solution or well-defined operation like the (generalized) EVD to solve (19). Nevertheless, iterative optimization procedures like gradient-descent or the quasi-Newton method can be used [17]. After solving the unconstrained optimization problem in (16), the RTF vector estimate $\hat{\mathbf{h}}^{\text{ODS}}$ is obtained by normalizing the solution to the reference entry.

In general, it should be noted that the optimization problems for the ODS method in (16) and the biased estimator in (12) only require the noisy covariance matrix $\hat{\mathbf{R}}_y$, whereas the optimization problem for the CW and CW-D methods in (13) also require an estimate of the noise covariance matrix $\hat{\mathbf{R}}_v$.

4. EVALUATION

In this section, we evaluate the performance of the presented RTF vector estimation methods using real-world recordings. The considered performance measures are the Hermitian angle between the ground truth RTF vector and the estimated RTF vector, and the intelligibility-weighted SNR improvement of the MVDR beamformer using the respective RTF vector estimates.

4.1. Setup and Implementation

Figure 2 depicts the considered acoustic scene for the evaluation, consisting of a target speaker in an acoustically treated laboratory with dimensions $7 \times 6 \times 2.7 \text{ m}^3$ and a reverberation time $T_{60} \approx 500 \text{ ms}$. As speech material, four different talkers (two male and two female) from the EBU SQAM CD [18] and the VCTK corpus [19] were used. All utterances had a duration of 20 s. The ambient noise was generated by four loudspeakers in the corners of the room. Different versions of multi-talker babble noise were played back at approximately the same level by the four loudspeakers.

The acoustic sensor network for the evaluation consisted of four nodes with uniform linear arrays, placed at about 0.5 m distance from the noise loudspeakers. Nodes 1-3 contained four microphones each, while node 4 contained three microphones, giving a total of $M=15$ microphones. For all arrays, two different microphone spacings of 1 cm and 3 cm were considered. The first microphone of node 1 was chosen as the reference microphone. The speech and noise components were recorded separately at a sampling rate of 16 kHz and mixed subsequently at an SNR of $\text{SNR}_{\text{in}} = \{-5, 0, 5\} \text{ dB}$ in the reference microphone.

For the implementation of the algorithms, an STFT framework with a frame length of 512 samples (corresponding to 32 ms), a frame




Figure 2: Schematic setup of acoustic scene with four different speaker target positions, four microphone nodes and four distinct noise sources.

overlap of 50%, and a square-root-Hann window for analysis and synthesis was used. The covariance matrices were estimated in batch, where for each frequency bin $\hat{\mathbf{R}}_y$ was estimated during speech activity and $\hat{\mathbf{R}}_v$ was estimated during speech pauses. For each frequency bin, speech-plus-noise and noise-only frames were determined by means of a speech presence probability (SPP) estimator [20], which was computed on one microphone per node and subsequently averaged.

The MVDR beamformer was computed according to (7), where the estimated noise covariance matrix was used in conjunction with one of the four presented RTF vector estimates:

- CW: EVD of $\hat{\mathbf{R}}_y$ whitened with $\hat{\mathbf{R}}_v$.
- CW-D: CW using block-diagonal noise covariance matrix $\hat{\mathbf{R}}_v$.
- Biased estimator: EVD of $\hat{\mathbf{R}}_y$.
- ODS: Iterative optimization method using only off-diagonal blocks of $\hat{\mathbf{R}}_y$. Optimized using MATLAB's `fminunc` function [21] supplied with the gradient and initialized on a random vector.

As a measure of RTF vector estimation accuracy, we use the Hermitian angle [22] between the ground truth RTF vector \mathbf{h} and the estimated RTF vector $\hat{\mathbf{h}}$, i.e., $\theta = \arccos\left(\frac{|\mathbf{h}^H \hat{\mathbf{h}}|}{\|\mathbf{h}\|_2 \|\hat{\mathbf{h}}\|_2}\right)$, averaged over all frequency bins. For each target position, the ground truth RTF vector was computed via the EVD of the oracle speech covariance matrix, obtained using the measured room impulse response convolved with white Gaussian noise. The second performance measure is the intelligibility-weighted SNR improvement [23] $\Delta\text{SNR} = \text{SNR}_{\text{out}} - \text{SNR}_{\text{in,max}}$, where the output SNR SNR_{out} is computed using the filtered speech and noise signal components and $\text{SNR}_{\text{in,max}}$ is the highest input SNR among all microphones.

4.2. Results and Discussion

For different input SNRs, Figure 3 depicts the Hermitian angle and the SNR improvement for the considered RTF vector estimation methods, where the bars represent the mean over 32 conditions (four target speaker positions, four speakers, two microphone spacings) and the error bars depict the standard deviation. First, it can be observed that in general the Hermitian angle in the upper panel of Figure 3 decreases with increasing input SNR, implying more accurate estimation at higher SNRs. Although at low SNRs the CW and CW-D methods achieve a lower Hermitian angle than the biased estimator and ODS method, these differences become negligible at higher input SNRs.

Second, it can be observed that in terms of SNR improvement the dif-




Figure 3: Hermitian angle between ground truth RTF vector and different RTF vector estimates (upper panel) and SNR improvement (lower panel) for different input SNRs.

ferences between the methods are more noticeable than in terms of Hermitian angle. At all input SNRs, the CW and CW-D methods consistently achieve around 12 dB of SNR improvement and outperform the biased estimator and the ODS method. This indicates a clear benefit of compensating for the noise using the estimated noise covariance matrix $\hat{\mathbf{R}}_v$ instead of using biased or selected information from the covariance matrix $\hat{\mathbf{R}}_y$. The performance of the CW and CW-D methods is similar, although the CW-D method attains a slightly higher SNR improvement, particularly at higher input SNRs. These results indicate a good validity of the block-diagonal model for the noise covariance matrix for the considered scenario. Hence, inverting only the diagonal blocks, cf. (15), seems to be sufficient or even beneficial, as it may reduce estimation errors of $\hat{\mathbf{R}}_v$.

Comparing the SNR improvement of the biased estimator with the ODS method, it can be observed that at an input SNR of -5 dB, the ODS method significantly outperforms the biased estimator (by about 1.5 dB). At higher input SNRs, the advantage of using only the off-diagonal blocks vanishes, and it seems more beneficial to use all information as the influence of noise diminishes. This indicates a higher robustness of the EVD of the full covariance matrix towards deviations from the rank-1 speech model compared to selecting only the off-diagonal blocks. At low SNRs, however, it seems more advantageous to exclude biased information and use only the off-diagonal blocks, which are affected less by noise, leading to a better performance of the ODS method compared to the biased estimator.

5. CONCLUSION

In this paper, we presented and compared different RTF vector estimation methods leveraging the assumed block-diagonal structure of the noise covariance matrix in an acoustic sensor network with multiple nodes. In an evaluation with real-world recordings, the modified CW method, which only considers the diagonal blocks of the noise covariance matrix, showed equal or even slightly better results than the original CW method at a lower complexity. Furthermore, we proposed a novel optimization problem for RTF vector estimation by selecting only the off-diagonal blocks of the noisy covariance matrix which are assumed not to be affected by noise. The evaluation results showed that the ODS method clearly outperforms a biased estimator in terms of SNR improvement, especially at low SNRs, showing that the selection of only unbiased information is beneficial if the influence of noise is large.

6. REFERENCES

- [1] A. Bertrand, "Applications and trends in wireless acoustic sensor networks: A signal processing perspective," in *Proc. IEEE Symposium on Communications and Vehicular Technology in the Benelux (SCVT)*, Ghent, Belgium, Nov. 2011, pp. 1–6.
- [2] S. Markovich-Golan, A. Bertrand, M. Moonen, and S. Gannot, "Optimal distributed minimum-variance beamforming approaches for speech enhancement in wireless acoustic sensor networks," *Signal Processing*, vol. 107, pp. 4–20, 2015.
- [3] V. M. Tavakoli, J. R. Jensen, M. G. Christensen, and J. Benesty, "A framework for speech enhancement with ad hoc microphone arrays," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 6, pp. 1038–1051, 2016.
- [4] A. I. Koutrouvelis, T. W. Sherson, R. Heusdens, and R. C. Hendriks, "A low-cost robust distributed linearly constrained beamformer for wireless acoustic sensor networks with arbitrary topology," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 8, pp. 1434–1448, 2018.
- [5] J. Zhang, R. Heusdens, and R. C. Hendriks, "Relative acoustic transfer function estimation in wireless acoustic sensor networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 10, pp. 1507–1519, 2019.
- [6] N. Gößling, W. Middelberg, and S. Doclo, "RTF-steered binaural MVDR beamforming incorporating multiple external microphones," in *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, New Paltz, USA, Oct. 2019, pp. 368–372.
- [7] R. M. Corey and A. C. Singer, "Adaptive binaural filtering for a multiple-talker listening system using remote and on-ear microphones," in *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, New Paltz, USA, Oct. 2021, pp. 1–5.
- [8] M. Cobos, F. Antonacci, A. Alexandridis, A. Mouchtaris, and B. Lee, "A survey of sound source localization methods in wireless acoustic sensor networks," *Wireless Communications and Mobile Computing*, vol. 2017, 2017.
- [9] S. Doclo, W. Kellermann, S. Makino, and S. E. Nordholm, "Multichannel signal enhancement algorithms for assisted listening devices: Exploiting spatial diversity using multiple microphones," *IEEE Signal Processing Magazine*, vol. 32, no. 2, pp. 18–30, Mar. 2015.
- [10] S. Gannot, E. Vincent, S. Markovich-Golan, and A. Ozerov, "A consolidated perspective on multi-microphone speech enhancement and source separation," *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, vol. 25, pp. 692–730, Apr. 2017.
- [11] B. D. Van Veen and K. M. Buckley, "Beamforming: A versatile approach to spatial filtering," *IEEE ASSP Magazine*, vol. 5, no. 2, pp. 4–24, Apr. 1988.
- [12] S. Markovich, S. Gannot, and I. Cohen, "Multichannel eigenspace beamforming in a reverberant noisy environment with multiple interfering speech signals," *IEEE Trans. on Audio, Speech, and Language Processing*, vol. 17, no. 6, pp. 1071–1086, Aug. 2009.
- [13] R. Serizel, M. Moonen, B. van Dijk, and J. Wouters, "Low-rank approximation based multichannel Wiener filter algorithms for noise reduction with application in cochlear implants," *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, vol. 22, no. 4, pp. 785–799, Apr. 2014.
- [14] S. Markovich-Golan and S. Gannot, "Performance analysis of the covariance subtraction method for relative transfer function estimation and comparison to the covariance whitening method," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Brisbane, Australia, Apr. 2015, pp. 544–548.
- [15] Y. Avargel and I. Cohen, "On multiplicative transfer function approximation in the short-time Fourier transform domain," *IEEE Signal Processing Letters*, vol. 14, no. 5, pp. 337–340, 2007.
- [16] W. Middelberg and S. Doclo, "Bias analysis of spatial coherence-based RTF vector estimation for acoustic sensor networks in a diffuse sound field," in *Proc. International Workshop on Acoustic Signal Enhancement (IWAENC)*, Bamberg, Germany, Sep. 2022, pp. 1–5.
- [17] R. Fletcher, *Practical methods of optimization*. New York: John Wiley & Sons, 1987.
- [18] European Broadcasting Union, "Sound quality assessment material — recordings for subjective tests: User's handbook for the EBU SQUAM CD," Tech. Rep. EBU – TECH 3253, European Broadcasting Union, 2008.
- [19] C. Veaux, J. Yamagishi, and K. MacDonald, "CSTR VCTK corpus: English multi-speaker corpus for CSTR voice cloning toolkit," University of Edinburgh. The Centre for Speech Technology Research (CSTR), 2017. [Online]. Available: <https://doi.org/10.7488/ds/2645>
- [20] T. Gerkmann and R. C. Hendriks, "Unbiased MMSE-Based Noise Power Estimation With Low Complexity and Low Tracking Delay," *IEEE Trans. on Audio, Speech, and Language Processing*, vol. 20, no. 4, pp. 1383–1393, May 2012.
- [21] The MathWorks, Inc., *Optimization Toolbox User's Guide*, The MathWorks, Inc., 2023.
- [22] R. Varzandeh, M. Taseska, and E. A. P. Habets, "An iterative multichannel subspace-based covariance subtraction method for relative transfer function estimation," in *Proc. Joint Workshop on Hands-free Speech Communication and Microphone Arrays*, San Francisco, USA, Mar. 2017, pp. 11–15.
- [23] J. E. Greenberg, P. M. Peterson, and P. M. Zurek, "Intelligibility-weighted measures of speech-to-interference ratio and speech system performance," *Journal of the Acoustical Society of America*, vol. 94, pp. 3009–3010, 1993.