

EEG-informed speaker extraction from noisy recordings in neuro-steered hearing aids: linear versus deep learning methods

Neetha Das^{†*}, Jeroen Zegers[‡], Hugo Van hamme[‡], Tom Francart^{*}, Alexander Bertrand[†]

Abstract—Objective: A hearing aid’s noise reduction algorithm cannot infer to which speaker the user intends to listen to. Auditory attention decoding (AAD) algorithms allow to infer this information from neural signals, which leads to the concept of neuro-steered hearing aids. We aim to evaluate and demonstrate the feasibility of AAD-supported speech enhancement pipelines in challenging noisy conditions without access to clean speech signals. **Methods:** We evaluated a linear versus a deep neural network (DNN) based speaker separation pipeline, with same-gender speaker mixtures for 3 different speaker positions and 3 different noise conditions. **Results:** AAD results based on the linear approach were found to be at least on par and sometimes even better than pure DNN-based approaches in terms of AAD accuracy in all tested conditions. However, when extending the DNN with a linear data-driven beamformer, a performance improvement over the purely linear approach was obtained in the most challenging scenarios. The use of multiple microphones was also found to improve speaker separation and AAD performance over single-microphone systems. **Conclusion:** Our study shows that neuro-steered speech enhancement, combining the best of both worlds (linear and DNN), results in robust performance. **Significance:** Recent proof-of-concept studies in this context each focus on a different method in a different experimental setting, which makes it hard to compare them. Furthermore, their idealized experimental conditions only give a rather premature evidence on the viability of the AAD paradigm in a hearing aid context. This work provides a systematic proof-of-concept of neuro-steered speech enhancement in challenging conditions.

Index Terms—Neuro-steered hearing aids, source separation, auditory attention decoding, EEG processing, cocktail party, deep learning.

I. INTRODUCTION

In a noisy environment with multiple speakers talking simultaneously, i.e., the so-called cocktail-party problem, a person with normal hearing has the ability to focus attention on one speaker and ignore the other speakers and surrounding noise in a seemingly effortless manner. People with a hearing impairment, on the other hand, find such situations extremely challenging. Although modern hearing aids allow to suppress background noise, a major unsolved problem is to determine which of the speakers is the desired one, and which speakers should be treated as noise. Existing approaches use unreliable

heuristics based on speaker loudness or other acoustic features or simply assume that the target speaker is always in the frontal direction.

Recent developments in the field of neuroscience have shown that it is possible to decode the auditory attention of a listener in a multi-talker environment, from brain signals recorded with magneto- or electro-encephalography (M/EEG) [1]–[3]. This opens up new opportunities to design smarter hearing devices, augmented with electrodes to record neural signals, that decode in real time to which speaker a listener is attending thereby assisting the hearing device to determine and enhance the attended speaker.

The envelope tracking of speech streams in a listener’s cortical responses, and more so of the attended speech stream, is well established in the literature [4]–[6]. Furthermore, this differential tracking of attended and unattended streams is also present in hearing impaired listeners [7]–[9]. This fact can be exploited to design algorithms that perform auditory attention decoding (AAD). Various methods have been developed to achieve AAD, addressing different kinds of decoders [2], [10]–[13], stimulus features [3], [14], [15], data acquisition [16]–[18], electrode selection and miniaturization strategies [16], [18], [19] etc.

The aforementioned AAD studies in [2]–[19] assume access to the clean speech signals to perform AAD, which are not available in a practical setting. Indeed, a hearing aid only has access to noisy microphone recordings in which multiple speech sources and noise sources are mixed. The first study that aimed for AAD without access to the clean speech streams was published in [20], where a multiplicative non-negative ICA (M-NICA) algorithm [21] was used to extract the individual speech envelopes from a noisy multi-microphone input. The AAD method in [14] was then used to select the attended speaker, which was then extracted from the microphone array using a multi-channel Wiener filter (MWF) [20]. This pipeline was later extended in [22] with a twofold MWF (one for each speaker) leading to better decoding performance and less variation across different acoustic conditions. However, while in both of these studies, the performance of the speaker separation algorithm was evaluated for a range of background noise levels and speaker positions, the EEG signals used for AAD were collected from subjects who listened to a 2-speaker scenario with 180 degree separation between speakers and without background noise. Therefore, there was a mismatch between acoustic conditions in the AAD experiments and the audio processing modules. In [23], speech enhancement was achieved using binaural direction-of-arrival (DOA) estimators with linearly constrained minimum variance (LCMV) beamformers. The AAD module used EEG signals recorded in two relatively high signal-to-noise ratio (SNR) conditions and two

The work is funded by KU Leuven Special Research Fund C14/16/057, FWO project nr G0A4918N and the ERC (637424 and 802895) under the European Union’s Horizon 2020 research and innovation programme. Jeroen Zegers is supported by an SB PhD scholarship 1S66217N of FWO. The scientific responsibility is assumed by its authors.

[†] KU Leuven, Dept. Electrical Engineering (ESAT), Stadius Center for Dynamical Systems, Signal Processing and Data Analytics. Kasteelpark Arenberg 10, B-3001 Leuven, Belgium.

[‡] KU Leuven, Dept. Electrical Engineering (ESAT), Center for Processing of Speech and Images (PSI). Kasteelpark Arenberg 10, B-3001 Leuven, Belgium.

^{*} KU Leuven, Dept. Neurosciences, ExpORL. Herestraat 49 bus 721, B-3000 Leuven, Belgium.

reverberation conditions.

The audio processing pipelines in [20], [22], [23] are based on linear beamforming and linear source separation techniques, which do not require any training data and have the advantage of being computationally cheap. In recent years, DNN-based approaches have become a popular alternative to solve the speaker separation problem, particularly for the challenging single-microphone scenario [24]–[26] and the even more challenging scenario with additional background noise [27], [28]. When using multiple microphones, separation performance can be increased by making use of the spatial information of the sources [29], [30]. In [31], [32], a single-microphone DNN-based approach was used for speaker separation in an AAD context. However, both of these studies involved training and testing in noise-free conditions. Furthermore, AAD was performed on electro-corticography (ECoG) data, which is an invasive approach involving surgery. This is a limiting factor for hearing-aid users. Non-invasive techniques like EEG are more preferable in this regard.

It is evident that a detailed analysis of neuro-steered speech enhancement under challenging acoustic conditions is necessary to draw conclusions about its feasibility. Despite the promising results in the aforementioned studies, all of them have important caveats, which may lead to overoptimistic conclusions in a context of demonstrating viability of neuro-steered speech enhancement algorithms. To summarize: The audio pipeline of [20], [22] was tested in various noisy conditions with different angles between speakers, but the EEG recordings were performed in a noise-free condition with 180 degree separation between speakers, resulting in a mismatch between the acoustic conditions of the microphone and neural signals. [23] does not have this mismatch, but only considers mildly noisy conditions and always the same speaker positions. [32] uses ECoG instead of EEG, only reports result in noiseless conditions, only uses a single microphone (while most state-of-the-art hearing aids contain multiple microphones), and uses different-gender data¹.

In addition to these practical limitations, the different settings across the aforementioned studies make it impossible to draw conclusions about the impact on AAD performance of non-linear (DNN-based) methods like [31], [32], and linear methods like [20], [22], [23], as well as the effect of using spatial information from multiple microphones versus single-microphone methods.

The contribution of this paper is twofold. First, we avoided the aforementioned caveats and perform EEG-based AAD using microphones from a binaural hearing aid in challenging noisy conditions with same-gender competing talkers at various relative speaker positions. In each acoustic condition, the audio and neural data are matched to each other, i.e., the audio signals presented to the subjects during the EEG recording are the same as those used in the audio signal processing pipeline. Secondly, we investigated the potential improvement of using a DNN-based speech separation algorithm compared to a computationally cheap and training-free linear signal processing

¹Note that both AAD and speaker demixing are expected to be more difficult when the competing speakers are of the same gender.

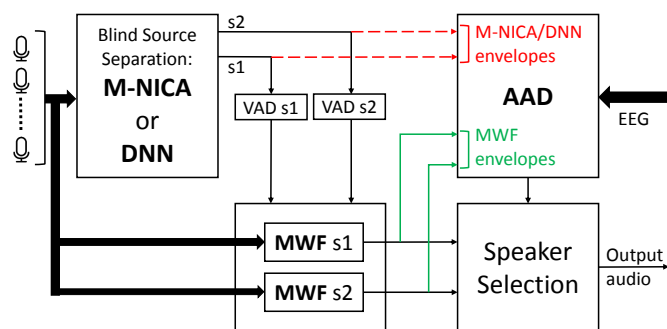


Fig. 1. Block diagram of the proposed neuro-steered speech enhancement algorithm for a 2 speaker scenario, consisting of modules for blind source separation (BSS), filtering for speech enhancement (MWF) and auditory attention decoding. The red (dashed) track and the green track correspond to two different implementations, as explained in the text.

algorithm. We show that the performance of the extended linear approach is at par or better than the purely DNN-based approach. We also investigated the potential advantage of using a combination of both approaches, i.e. a DNN-based speech separation in combination with a linear beamformer. We demonstrate the superior performance of such a system that combines ‘the best of both worlds’, both in terms of improvement in signal-to-interference-plus-noise ratio (SINR) and AAD accuracy, as well as possibilities of offsetting real-time delay by decoupling the blind source separation (BSS) part (and hence the delay caused by it) from the direct output to the listener using the beamformer as a go-between. We focused on two representative state-of-the-art algorithms, i.e., deep clustering [24] as the DNN-based pipeline and M-NICA-MWF [22] as the linear pipeline. In both cases, we used a multi-microphone input, which we will demonstrate substantially improves speech separation compared to single-microphone input.

The outline of this paper is as follows. The different neuro-steered speech enhancement pipelines are presented in section II. The various experiments used to validate the pipeline are presented in section III. In section IV, the results of the experiments are shown, and the various implications of these findings are discussed in section V.

II. NEURO-STEERED SPEECH ENHANCEMENT

The proposed neuro-steered speech enhancement pipeline is shown in the block diagram of Fig. 1 for a 2-speaker scenario². The pipeline consists of three steps:

- 1) In the M-NICA/DNN and the MWF modules, the speech sources are separated and denoised. This will be discussed in subsection II-B.
- 2) In the AAD module, the attended speech envelope from the listener’s EEG is reconstructed using a pre-trained decoder. This reconstructed speech envelope is correlated with the speech envelopes of the individual

²Although the analysis in this paper focuses on a 2-speaker scenario, the speech enhancement pipeline in Fig. 1 can straightforwardly be generalized to a larger number of speakers (denoted by N in the text).

speakers extracted in step 1, where the speaker with the highest correlation is selected as the attended speaker. To this end, either the envelopes from the M-NICA or DNN algorithm can be used directly (red dashed lines in Fig. 1), or a new set of envelopes can be extracted from the denoised MWF outputs (green lines in Fig. 1). Both cases will be investigated, where the envelopes in the red dashed lines will be referred to as M-NICA or DNN envelopes (depending on which algorithm is used in the BSS block) and the ones in the green lines as MWF envelopes. More details about the AAD are given in subsection II-A.

- 3) Once the speaker has been decoded, the corresponding MWF output signal is selected as the final output, which ideally consists of a denoised version of the speech signal corresponding to the attended speaker.

A. Auditory attention decoding (AAD)

In order to choose which enhanced speech estimate will be presented as output of the neuro-steered speech enhancement pipeline, it is necessary to know which speech signal is being attended to by the listener. Towards this goal, we use an AAD module with a pre-trained spatio-temporal decoder (subject-specific training) which takes the listener's EEG data as an input to reconstruct the speech envelope of the attended speaker. In the training data, the ground truth about the subject's attention is known, and the decoder is designed such that the difference between its output and the attended speech envelope is minimized in the minimum mean square error sense. The algorithm described in [14] was used for our decoder design. The reconstructed attended speech envelope is given by

$$\tilde{p}(k) = \sum_{\tau=0}^{L-1} \sum_{c=1}^C d_c(\tau) m_c(k + \tau), \quad (1)$$

where $m_c(k)$ is the value of the c -th EEG channel at sample time k , and $d_c(\tau)$ denotes the decoder weight at the c -th channel at time lag τ . We stack all the decoder weights in a vector $\mathbf{d} = [d_1(0), d_1(1), \dots, d_1(L-1), d_2(0), \dots, d_2(L-1), \dots, d_C(0), \dots, d_C(L-1)]^T$. Stacking EEG samples of all C channels for L time lags, we get $\mathbf{m}(k) = [m_1(k), m_1(k+1), \dots, m_1(k+L-1), m_2(k), \dots, m_2(k+L-1), \dots, m_C(k), \dots, m_C(k+L-1)]^T \in \mathbb{R}^{LC}$. Using this notation, (1) can be written as $\tilde{p}(k) = \mathbf{d}^T \mathbf{m}(k)$. The optimal decoder that minimizes the mean square error between the reconstructed attended envelope $\tilde{p}(k)$ and the attended speech envelope $p_a(k)$ is given by

$$\hat{\mathbf{d}} = \mathbf{R}_{mm}^{-1} \mathbf{R}_{mp_a}, \quad (2)$$

where $\mathbf{R}_{mm} = E\{\mathbf{m}(k)\mathbf{m}(k)^T\} \in \mathbb{R}^{LC \times LC}$ is the EEG covariance matrix, $\mathbf{R}_{mp_a} = E\{\mathbf{m}(k)p_a(k)\} \in \mathbb{R}^{LC}$ is the cross-correlation vector between the EEG data and the attended speech envelope, which is known during the training of the decoder and $E\{\cdot\}$ denotes the expected value operator. At test time, the correlation coefficients between $\tilde{p}(k)$ and the estimated speech envelope of each speaker are then computed.

The speech envelope that results in the higher correlation is selected as the attended speech stream.

B. Blind source separation and denoising

The microphone signal of a mixture of N speakers can be represented in the short-time Fourier transform (STFT) domain as

$$y(t, \omega) = \sum_{i=1}^N x_i(t, \omega) + n(t, \omega) = \sum_{i=1}^N h_i(\omega) s_i(t, \omega) + n(t, \omega), \quad (3)$$

where t indicates the time, ω the frequency index, $n(t, \omega)$ represents the noise recorded by the microphone and $x_i(t, \omega)$ represents the observed signal of the i -th speaker s_i , as recorded by the microphone. The transfer function that models the acoustic propagation path between the source s_i and the microphone is given by $h_i(\omega)$.

In a multi-microphone setting, such as a (binaural) hearing device, where M microphones are present, (3) generalizes to

$$\mathbf{y}(t, \omega) = \sum_{i=1}^N \mathbf{x}_i(t, \omega) + \mathbf{n}(t, \omega) = \sum_{i=1}^N \mathbf{h}_i(\omega) s_i(t, \omega) + \mathbf{n}(t, \omega), \quad (4)$$

where $\mathbf{y}(t, \omega) = [y_1(t, \omega), y_2(t, \omega), \dots, y_M(t, \omega)]^T$ represents the microphone components stacked over M microphone channels, $\mathbf{h}_i(\omega) = [h_{i,1}(\omega), h_{i,2}(\omega), \dots, h_{i,M}(\omega)]^T$ with $h_{i,j}(\omega)$ modeling the path between the source s_i and microphone j , and $\mathbf{n}(t, \omega) = [n_1(t, \omega), n_2(t, \omega), \dots, n_M(t, \omega)]^T$ represents the stacked components of noise in the M microphones. A multi-microphone mask (or a beamformer) $\mathbf{w}_i(t, \omega)$ can be estimated such that the speech signal can then be estimated as

$$\tilde{s}_i(t, \omega) = \mathbf{w}_i(t, \omega)^H \mathbf{y}(t, \omega), \quad (5)$$

where superscript H denotes the conjugate transpose operator. Note that, while we use the symbol \tilde{s} in (5) for the sake of clarity, we do not necessarily aim to perform a dereverberation, i.e., \tilde{s}_i may be closer to the recorded speech signal $x_i(t, \omega) = h_i(\omega) * s_i(t, \omega)$ than to the 'dry' speech signal s_i itself.

We present two methods to estimate the mask in (5): one that is fully linear, based on the M-NICA algorithm, and one that makes use of a DNN.

- 1) The linear multi-channel signal processing approach using M-NICA provides an estimate of the energy envelope of each individual speaker (see subsection II-B1). These envelopes can be directly fed to the AAD module (red dashed track in Fig. 1 as in [33], or alternatively, they can be used as a voice activity detection (VAD) mechanism to design an N -fold MWF which estimates per-speaker masks \mathbf{w}_i in (5), as explained in subsection II-B2. The envelope of \tilde{s}_i from (5) can then be fed to the AAD module as in [22] (green track in Fig. 1).
- 2) On the other hand, a DNN can be trained to separate speech mixtures into individual speech streams (see subsection II-B3). As opposed to M-NICA, the DNN directly estimates the masks \mathbf{w}_i to compute (5), after which the speech envelopes can be extracted and directly

fed to the AAD module as in [31], [32] (red dashed track in Fig. 1). Alternatively, the envelopes obtained from the DNN can also be used as a VAD to design MWF masks (green track in Fig. 1).

1) *Multiplicative non-negative independent component analysis (M-NICA)*: The M-NICA algorithm was introduced in [21], and has been used to blindly separate speech envelopes from a multi-microphone input [20], [33]. The M-NICA algorithm operates in the short-term energy domain, thereby transforming the problem into a (non-negative) energy envelope separation problem. To this end, we first compute the short-term energy signal from each microphone signal over a window of K samples and apply M-NICA on the resulting M energy signals (note that M-NICA operates at a sampling rate that is K times lower than the microphone signal sampling rate). For more details on the implementation of M-NICA for demixing speech envelopes, we refer to [21], [33].

M-NICA exploits the inter-channel differences in the energy distributions of the speech sources in order to extract the energy envelopes of each individual speaker. In the case of a binaural hearing aid, these differences are mostly due to head shadow effects. The non-negativity of the underlying sources facilitates the use of only second order statistics in its computations, in comparison to other source separation algorithms that use higher order statistics. This, together with the K -fold reduction in sampling rate, leads to a low computational complexity which is a desirable factor to incorporate such a source separation scheme in an actual hearing prosthesis.

2) *MWFs for source separation and/or denoising*: In this paper, we will use an N -fold MWF (one for each speaker) to estimate each speech signal in real time. The MWF is a data-driven and adaptive linear beamforming algorithm which allows to efficiently denoise and extract a single speaker from a mixture by computing the mask in (5) which results in the best speech estimate in the linear minimum mean squared error (LMMSE) sense. It is a semi-supervised method as it requires a voice activity detector (VAD) for the target speaker, i.e., it needs to know at which times the target speaker is not active in order to estimate the second-order statistics of the background noise and interfering speakers. The VAD information can be straightforwardly extracted by thresholding the speech envelopes obtained from the M-NICA algorithm (or alternatively using the deep learning algorithm as explained in subsection II-B3). The i -th MWF computes the mask $\mathbf{w}_i(t, \omega)$ in (5) that yields the LMMSE estimate $\tilde{s}_i(t, \omega)$ that is closest to $x_{i, m_{ref}}(t, \omega)$, i.e., the i -th speaker's contribution in an arbitrarily-chosen reference microphone m_{ref} . The MWF $\mathbf{w}_i(t, \omega)$ is defined as

$$\hat{\mathbf{w}}_i(t, \omega) = \arg \min_{\mathbf{w}_i} E\{|x_{i, m_{ref}}(t, \omega) - \mathbf{w}_i(t, \omega)^H \mathbf{y}(t, \omega)|^2\}. \quad (6)$$

The solution for this LMMSE problem is given by [34]:

$$\hat{\mathbf{w}}_i(t, \omega) = \mathbf{R}_{yy}^{-1} \mathbf{R}_{x_i x_i} \mathbf{e}_{m_{ref}}, \quad (7)$$

where $\mathbf{R}_{yy} = E\{\mathbf{y}(t, \omega) \mathbf{y}(t, \omega)^H\}$, $\mathbf{R}_{x_i x_i} = E\{\mathbf{x}_i(t, \omega) \mathbf{x}_i(t, \omega)^H\}$, and $\mathbf{e}_{m_{ref}}$ denotes the m_{ref} -th column of an $M \times M$ identity matrix, which selects the column of $\mathbf{R}_{x_i x_i}$ corresponding to the reference microphone.

The ‘speech plus interference’ autocorrelation matrix \mathbf{R}_{yy} can be estimated during the time periods when the i -th speaker is active³. Assuming independence between all sources, $\mathbf{R}_{x_i x_i}$ can be estimated as $\mathbf{R}_{x_i x_i} = \mathbf{R}_{yy} - \mathbf{R}_{vv}$. Here, $\mathbf{R}_{vv} = E\{\mathbf{n}(t, \omega) \mathbf{n}(t, \omega)^H\} + \sum_{j \neq i} P_j(\omega) \mathbf{h}_j(\omega) \mathbf{h}_j(\omega)^H$ is the interference autocorrelation matrix, where $P_j(\omega) = E\{|s_j(t, \omega)|^2\}$. \mathbf{R}_{vv} can be estimated by averaging over all STFT frames in which speaker i is not active. Estimating $\mathbf{R}_{x_i x_i}$ as $\mathbf{R}_{yy} - \mathbf{R}_{vv}$ can result in poor filters, particularly, if \mathbf{R}_{vv} contains non-stationary sources. Therefore, in our computations, $\mathbf{R}_{x_i x_i}$ was estimated based on a generalized eigenvalue decomposition (GEVD) of \mathbf{R}_{yy} and \mathbf{R}_{vv} (details in [35]) to ensure robustness. Note that, while speech is known to be highly non-stationary, the expected values here are calculated over a long-term window, thereby capturing the average long-term power spectrum of the speech. Therefore, the MWF will not react to short-term changes in the speech spectra, it mainly focuses on the (fixed or slowly varying) spatial coherence across the microphones. To also exploit the fast spectral changes in the speech, a single-channel Wiener filter which uses short-term statistics can be added as a postfilter [36], yet this is beyond the scope of this paper.

To separately estimate \mathbf{R}_{yy} and \mathbf{R}_{vv} , the MWFs require the voice activity information of the speaker it must enhance. The energy envelopes from the output of the source separation algorithm (M-NICA or a DNN approach, see subsections II-B1 and II-B3, respectively) are used to identify the active and silent periods of each speaker, for e.g., through a thresholding operation⁴, and hence form the VAD tracks that each MWF needs.

Note that the N MWFs work in parallel to enhance the speech streams of N different speakers. It is important to note that these MWFs can share a large part of the computations, as the estimation of the matrix inverse \mathbf{R}_{yy}^{-1} is a common requirement across them. The real-time aspects of such a system will be discussed in section V-D.

3) *Deep learning for source separation*: Deep learning solutions for source separation make a distinction between single-microphone and multi-microphone approaches as well as between methods that can and cannot cope with additional background noise. The deep learning based separation technique that will be used in this paper is called Deep Clustering (DC) [24], which can be extended to cope with noise [27], [28] and to use multi-microphone information [29].

In DC, a D -dimensional embedding vector $\mathbf{v}(t, \omega)$ is constructed for every time-frequency bin as $\mathbf{v}(t, \omega) = f_{t\omega}(|Y|)$, where Y is used to refer to the complete input signal across all time-frequency bins, $\mathbf{v}(t, \omega)$ is a D -dimensional embedding vector normalized to unit length, and the function f is chosen to be a DNN. A $(T\Omega \times D)$ -dimensional matrix \mathbf{V} is then constructed from these embedding vectors, where T and Ω denote the total number of time and frequency bins,

³For the sake of conciseness, we omit the frequency variable ω for autocorrelation matrices in the in this section.

⁴There is also the possibility of using other parameters, speech presence probability [37] for instance, to make softer VADs, but this is not within the scope of this paper.

respectively. Similarly, a $(T\Omega \times N)$ -dimensional target matrix \mathbf{Z} is defined. If target speaker i is the dominant speaker for bin (t, ω) , then $z_i(t, \omega) = 1$, otherwise $z_i(t, \omega) = 0$. Speaker i is dominant in a bin (t, ω) if $i = \arg \max_j (|s_j(t, \omega)|)$. A permutation independent loss function is then defined as

$$\begin{aligned} \mathcal{L} &= \|\mathbf{V}\mathbf{V}^T - \mathbf{Z}\mathbf{Z}^T\|_F^2 \\ &= \sum_{t_1, \omega_1, t_2, \omega_2}^{T, \Omega, T, \Omega} (\langle \mathbf{v}(t_1, \omega_1), \mathbf{v}(t_2, \omega_2) \rangle - \langle \mathbf{z}(t_1, \omega_1), \mathbf{z}(t_2, \omega_2) \rangle)^2, \end{aligned} \quad (8)$$

where $\|\cdot\|_F^2$ is the squared Frobenius norm. Minimizing this loss function pushes embeddings together if their time-frequency bin are dominated by the same speaker and otherwise the embeddings should be perpendicular.

After estimating \mathbf{V} , all embedding vectors are clustered into N clusters using k-means clustering. Speaker masks are then constructed as follows

$$\hat{w}_i(t, \omega) = \begin{cases} 1, & \text{if } \mathbf{v}(t, \omega) \in c_i \\ 0, & \text{otherwise} \end{cases}, \quad (9)$$

with c_i a cluster. These masks can be used to estimate the source signals by applying them to the original speech mixture:

$$\tilde{s}_i(t, \omega) = w_i(t, \omega)y(t, \omega), \quad (10)$$

where (10) can be seen as the single channel simplification of (5). Since binary masks are estimated, it is assumed that only one speaker is active per time-frequency bin. Since speech spectra are sparse, this assumption holds to some extent [38].

Deep clustering, as presented above, is not well suited for noisy environments. The binary masks in (9) will cause the noise energy to be distributed over the speaker signal estimates. The noise energy can be filtered out by estimating an additional speech-to-noise ratio mask $\alpha(t, \omega)$ [27]. Both the embeddings $\mathbf{v}(t, \omega)$ and the ratio mask $\alpha(t, \omega)$ can be estimated by the DNN from the magnitude spectrum of the microphone signal: $[\mathbf{v}(t, \omega), \alpha(t, \omega)] = f_{t\omega}(|Y|)$. During training of the network, we set the target noise suppression mask to be

$$\alpha^{target}(t, \omega) = \sqrt{\frac{\sum_j^N |s_j(t, \omega)|^2}{|n(t, \omega)|^2 + \sum_j^N |s_j(t, \omega)|^2}}, \quad (11)$$

as it was shown to perform well in many noise scenarios [39]. The loss function of (8) can then be extended to:

$$\mathcal{L} = \|\mathbf{V}\mathbf{V}^T - \mathbf{Z}\mathbf{Z}^T\|_F^2 + \gamma \sum_{t, \omega}^{T, \Omega} (\alpha(t, \omega) - \alpha^{target}(t, \omega))^2, \quad (12)$$

where the hyper-parameter γ can be tuned to weigh the importance of speech separation and noise suppression. The estimated mask of (9) is then changed to

$$\hat{w}_i(t, \omega) = \begin{cases} \alpha(t, \omega), & \text{if } \mathbf{v}(t, \omega) \in c_i \\ 0, & \text{otherwise} \end{cases}. \quad (13)$$

DC has been developed for single microphone setups, but can easily be extended to include spatial information \mathbf{Y}_{sp}

from a multi-microphone setup [29]. We simply append this spatial information to the input of the model to estimate the embeddings (and the noise mask): $[\mathbf{v}(t, \omega), \alpha(t, \omega)] = f_{t\omega}(|\mathbf{Y}|, \mathbf{Y}_{sp})$. The interchannel phase difference (IPD) between two microphones $m1$ and $m2$ is chosen to represent the spatial information [29]:

$$y_{sp}(t, \omega) = [\cos(\theta_{m1, m2}(t, \omega)), \sin(\theta_{m1, m2}(t, \omega))] \quad (14)$$

with $\theta_{m1, m2}(t, \omega) = \angle y_{m1}(t, \omega) - \angle y_{m2}(t, \omega)$ and \mathbf{Y}_{sp} is the matrix containing the spatial information $y_{sp}(t, \omega)$ over all time-frequency bins (t, ω) . In case there are more than two microphones ($M > 2$), the spatial features can be determined for all microphone pairs and stacked as a single feature representation.

A single mask is produced, in the same way as was done in (13) which is then copied and applied to each microphone. The mask is applied to the average of all microphone spectrograms, or equivalently⁵

$$\hat{\mathbf{w}}_i(t, \omega) = \left[\frac{\hat{w}_i(t, \omega)}{M}, \frac{\hat{w}_i(t, \omega)}{M}, \dots, \frac{\hat{w}_i(t, \omega)}{M} \right]^T \quad (15)$$

when using (5).

So far we have focused on estimating masks to filter out the interfering speaker (and noise). The estimated $\tilde{s}_i(t, \omega)$ can be further enhanced in a secondary stage as in [40]. To this end, we train a new DNN denoted by the function g' , which estimates an enhanced mask for speaker i based on $|Y|$ and the previous estimate \tilde{s}_i : $\hat{w}_i^{enh}(t, \omega) = g'_{t\omega}(|Y|, \tilde{s}_i)$. Notice that the network used for g' is shared over the speakers. In the multi-microphone case, we again include the IPD to obtain $\hat{\mathbf{w}}_i^{enh}(t, \omega)$ similarly to (14)-(15).

In the next step, the envelopes of the speech sources as estimated from the mask (15) can be used for auditory attention decoding (red dashed lines in Fig. 1). There is also the option of using these envelopes to first compute the VADs for the N -fold MWFs as discussed in subsection II-B2, and then using the envelopes from these MWF outputs for attention decoding instead (green lines in Fig. 1). A final option, which will not be further discussed in this paper, is to use the masks estimates $\hat{w}_i(t, \omega)$ as a weight when determining \mathbf{R}_{yy} and \mathbf{R}_{vv} in section II-B2, similarly as was done in [41]. Initial experiments with this method gave no significant improvements in AAD accuracies.

III. EXPERIMENTS

A. AAD experiment setup

1) *AAD setup*: 18 normal-hearing subjects (of which 6 males) between 20 and 25 years old participated in the experiment, where each subject had to focus on the speech stream from a particular direction, in the presence of a competing speaker from another direction and background noise. Their 64-channel EEG was recorded at a sampling rate of 8192 Hz. This dataset consists of 138 minutes of EEG recordings per

⁵This corresponds to applying the mask to the output of a forward-steering beamformer. Other choices are possible (e.g. select only one microphone, or summing magnitude spectra and adding the phase of one of the microphones), but these were empirically found to not have a significant impact on the results.

subject. Further details can be found in [42]. The stimuli and the corresponding acoustic conditions used in this experiment are described below.

RBM Dataset: The RadioBoeken Mix (RBM) dataset was built by mixing parts of 8 Dutch stories narrated by 8 different female narrators [43], taken 2 at a time as competing speech streams in the presence of different levels of background multi-talker ('babble') noise. The background babble noise consisted of 9 different 4-talker babble streams (2-male and 2-female) perceived to be coming from 9 equidistant positions (from -180° to 140° in steps of 40°) around the listener. The babble was constructed from 36 audiobooks (18 male and 18 female narrators) from LibriVox (a public domain collection of audiobooks). Anechoic head-related transfer functions for 6 behind-the-ear microphones ($M = 6$) in a binaural hearing aid set-up [44] (with approximately 7.5 mm distance between neighbouring microphones) were used for making directional audio. All audio was sampled at 44.1 kHz. The long term spectrum of the babble sources was matched with the average spectrum of all the target speech streams. The experiments included a condition without background babble noise (referred to as the noise-free case) and two noisy conditions with SNRs⁶ of -1.1 dB and -4.1 dB. Note that both SNRs are negative, resulting in challenging conditions where the signal power of the background babble noise is higher than the signal power of the attended speaker. We simulated different scenarios regarding the relative position of both speakers, namely (-90° (utmost left), 90° (utmost right)), (30° , 90°), (-30° , -90°), and (-5° , 5°) with angle between speakers 180° , 60° , 60° and 10° respectively.

2) **DNN training dataset:** A separate audio dataset was needed to train the model for the speaker independent deep learning based source separation as explained in II-B3. A novel dataset, called *Corpus Gesproken Nederlands Mix* (CGNM), was created. It used the *Corpus Gesproken Nederlands component-o-Flemish* (CGN-o-VL) [45], which contains 150 Dutch speakers of which 102 (51 female and 51 male) were randomly chosen to build up the training set, containing 27 hours of unique single-speaker speech in total (16 minutes per speaker on average). 20,000 2-speakers mixtures were artificially created using randomly selected single-speaker audio from CGN-o-VL. These mixtures totalled 60 hours of overlapping speech, to be used as training set. Even though in the RBM dataset (section III-A1) only female-female mixtures are used, we include male-female and male-male mixtures in our training set as well for better generalizability. Furthermore, both speakers can be located anywhere between -100° and 100° , again to make the network more general. Two types of datasets were made: one containing the 20,000 mixtures without background babble noise (CGNM-noisefree), and the same mixtures with background babble noise (CGNM-noise). The SNR for the background noise for each mixture was uniformly sampled between -4dB and 4 dB. Babble speakers were taken from the LibriSpeech corpus [46] and were posi-

tioned in the same way as in III-A1. Both CGNM-noisefree and CGNM-noise were used to train the DNN models, unless specified otherwise. To assess separation quality, 12 held-out female speakers (totalling 3.2 hours of unique speech) were chosen to create a test set. The test set consisted of 728 2-speaker mixtures with babble noise (totalling 2.2 hours of mixed speech).

B. Data preprocessing and design choices

1) **EEG preprocessing:** All EEG data was filtered using an equiripple bandpass filter with passband between 0.5 Hz and 10 Hz, and with passband attenuation of at most 0.5 dB and stopband attenuation of 20 dB (lower) and 15 dB (upper). Previous studies [1], [5], [47] have shown that cortical envelope tracking is best within this frequency range. EEG data was then downsampled to 40 Hz.

2) **M-NICA audio preprocessing:** The microphone signals were low pass filtered with an 800 Hz cut-off and downsampled to 8 kHz before computing the energy envelopes, since it was found to be beneficial for the source separation process under low SNR conditions [22]. Energy was computed every 200 samples (25 ms), bringing down the sampling rate to 40 Hz.

3) **Deep learning design choices:** For the deep learning model utterances were also downsampled to 8 kHz. A STFT with a 32 ms window length (256 samples) and a hop size of 8 ms was used. A bidirectional long short-term memory recurrent neural network (BLSTM-RNN) with 2 layers and 600 hidden units each, was used to estimate the embeddings which were then normalized. The embedding dimension was chosen at $D = 20$ and since the number of frequency bins was $\Omega = 129$, the total number of output nodes was $D\Omega = 20 * 129 = 2580$. This is identical to the setup in [24]. For the enhancement network g' another BLSTM-RNN was used with 2 layers using 300 hidden units each. The output layer had $\Omega = 129$ nodes and a sigmoid function was applied on top to estimate the final masks. We set γ in (12) to 1.

Curriculum learning was applied by first training the networks on segments of 100 time frames, before training on the full mixture [24], [48]. The weights and biases were optimized with the Adam learning algorithm [49] and early stopping was used on the validation set. The log-magnitude of the STFT coefficients were used as input features and were mean and variance normalized. Zero mean Gaussian noise with standard deviation 0.2 was applied to the training features for better generalization [50]. For the single-channel scenario a single microphone was randomly selected. For the multi-channel scenario the input consisted of the average of the left-ear and the average of the right-ear microphone signals, as well as the spatial features (see (14)) from 3 microphone pairs (with one pair using a left-ear and a right-ear microphone and pair 2 and 3 using two left-ear and two right-ear microphones, respectively). All networks were trained using TensorFlow [51] and the code for the deep learning based source separation can be found here: <https://github.com/JeroenZegers/Nabu-MSSS>.

⁶In this paper (input) SNR is defined as the ratio of the power of the target speaker to the power of background babble. Note that the actual SNR is even lower due to the presence of an interfering speaker, which is not included in this SNR metric.

4) *MWF design choices*: For the estimation of MWF coefficients an STFT with a 64 ms window length (512 samples) and a hop size of 32 ms (256 samples) was applied to the microphone signals, after they were downsampled to 8 kHz. The thresholds for estimating the VAD tracks for the different approaches were determined empirically. For energy envelopes extracted from the M-NICA and DNN-based source separation, the 25th percentile of the per sample amplitude of the envelope was used as the threshold, above which the speaker was considered to be active.

5) *Audio envelope extraction*: We used the envelope extraction method from [14]. Each of the speech streams used as stimulus were filtered with a gammatone filterbank [52], splitting the signal into 15 filter bands. In each gammatone band, the absolute value of each sample was taken after which a power law compression with an exponent of 0.6 was applied. The power law-compressed samples were then bandpass filtered in the same manner as was done for the EEG signals resulting in subband envelopes. The subband envelopes were added together and downsampled to 40 Hz to form a single ‘powerlaw subband’ envelope. Power law compression aims to replicate the non-linear transformation of the stimulus in the human auditory system, with relatively higher attenuation of higher amplitude signals [53]. This method of auditory inspired envelope extraction has been found to result in significantly better attention decoding accuracies than other envelope extraction methods [14].

Power law subband envelopes were only used in cases where broadband speech streams were available, i.e., the DNN and MWF envelopes built from the speech streams obtained with the DNN-based mask (15) or the MWF-based mask (7), respectively. Since M-NICA directly extracts energy envelopes from the broadband speech streams, we could not apply the gammatone filterbank here. When applying AAD on the M-NICA-envelopes, we used the square-root of the extracted energy envelopes to transform them to an amplitude envelope, since energy envelopes are sub-optimal to perform AAD [14].

6) *AAD design choices*: In order to perform AAD, EEG data was split into trials of 30 s, resulting in 276 trials per subject for each of the 3 different angles between speakers (separations of 180° for (-90°,90°), 60° for (30°,90°) and (-30°,-90°), and 10° for (-5°,5°) respectively). Decoders were trained (using leave-one-trial-out (LOO) cross validation) to reconstruct the envelope of the attended speech stream. To improve decoding accuracy and eliminate the need for regularization, we computed a single decoder over the entire training set data as in [14], instead of averaging over per-trial decoders in the training set as in [2].

IV. RESULTS

A. Single-channel and multi-channel source separation

The deep learning approaches to noise-free source separation, both for single-channel and multi-channel setup, were trained on the CGNM-noisefree train set and evaluated on the CGNM-noisefree test set and the RBM noise-free set. Results are expressed in Signal-to-Distortion Ratio (SDR) [54] and are shown in table I.

TABLE I: SDR results for a single-channel and multi-channel DNN without noise.

	noise-free CGNM test	noise-free RBM
single-channel	10.1 dB	3.8 dB
multi-channel	12.4 dB	12.3 dB

TABLE II: SDR results for a single-channel DNN for the noise-free RBM set, depending on the speaker positions.

	10°	60°	180°	avg.
single-channel	2.5 dB	1.7 dB	7.3 dB	3.8 dB

Notice that for the single-channel case, performance on the noise-free RBM dataset is much worse than on the CGNM-noisefree test set, even though both consist of held-out female speakers. A possible reason for this could be differences in the recording set-up, and the fact that the CGNM-noisefree test set consists of volunteers, while the RBM dataset is made with professional story-tellers. We conclude that the single-channel setup does not generalize well across data sets. Multi-channel separation quality is considerably better than the single-channel counter part. Furthermore, SDR results on the CGNM-noisefree test set and the RBM dataset are also more alike, compared to the single-channel results. The network uses spatial information to characterize the speakers and does not need to solely rely on the difference in speaker characteristics to separate the speakers, as is necessary in the single-channel case [55]. In table II, we split up the single-channel result for the RBM dataset per angle between speakers and notice that for the 180° case, the model performs much better than the other cases. For the 180° the amplitudes of both speakers will differ more due to head shadow effects, which may explain why it is easier for the network to separate both. Further research is needed to support these claims/assumptions.

The absolute envelopes, extracted from the outputs of both single-channel and multi-channel DNNs performing speaker separation of the stimuli in the RBM dataset for the noise-free condition, were used for auditory attention decoding of the EEG data collected under matching acoustic conditions and




Fig. 2. AAD accuracies (on 30 s trials) of 18 subjects listening in noise-free conditions when envelopes from the output of multi-channel DNN and single-channel DNN were used. Comparisons between the different approaches were done using Wilcoxon’s signed-rank test with Holm correction. : ‘***’ for $p < 0.001$, ‘*’ for $p < 0.05$.

stimuli (ref. section III-A1). Fig. 2 shows AAD accuracies over 18 subjects under the noise-free condition, for the 3 different angles between the speakers. For all angles between speakers, the multi-channel DNN resulted in significantly higher AAD accuracies compared to the single-channel DNN, as was expected from the results in table I. We therefore continue with the multi-channel approach for all further analyses.

B. Linear versus non-linear source separation for AAD

The AAD accuracies for each subject were computed for each of the 3 different angles between speakers and the 3 SNRs. For each 30 s trial, we computed the correlation of the reconstructed attended envelope with 5 pairs of envelopes extracted using different speech enhancement methods:

- ‘Oracle’ - powerlaw subband envelopes extracted from the original speech signals (not mixed)
- ‘M-NICA’ - amplitude envelopes extracted from the output of the M-NICA source separation algorithm (red dashed track in Fig. 1)
- ‘DNN’ - powerlaw subband envelopes extracted from the output of the DNN-based source separation algorithm (also red dashed track in Fig. 1)
- ‘M-NICA+MWF’ - powerlaw subband envelopes extracted from the output of MWFs when using the output of the M-NICA based source separation algorithm to build VAD tracks (green track in Fig. 1).
- ‘DNN+MWF’ - powerlaw subband envelopes extracted from the output of MWFs when using the output of the DNN based source separation algorithm to build VAD tracks (also green track in Fig. 1).

Figure 3 shows the resulting AAD accuracies for the 18 subjects for the noise-free condition, for -1.1 dB and for -4.1 dB. For each angle between speakers, comparison between the different approaches were done using Wilcoxon’s signed-rank test with Holm correction. The green boxplots correspond to the green lines (MWF envelopes), and the red box plots correspond to the dashed red lines (M-NICA/DNN envelopes) respectively in Fig. 1.

As can be seen in Fig. 3, in the noise-free condition, for 180° and 10° separations, the M-NICA envelopes result in accuracies that match those when using oracle speech envelopes for AAD while the DNN envelopes perform significantly poorer. On the other hand, for 60° separation between speakers, the DNN envelopes outperform M-NICA envelopes, yet upon including the MWF filtering step, there are no more significant differences between the 2 source separation approaches. However, the results are still significantly lower than with oracle speech envelopes. It must be noted that in the 60° case, the competing speaker positions are on the same side of the listener’s head ((-30°, -90°) or (30°, 90°)). It has already been observed in [22] that this is a challenging acoustic condition for M-NICA. We also observe that, for all three angles between speakers, the performance of the fully linear pipeline of M-NICA+MWF is on par with the pure DNN approach.

Fig. 3 also shows AAD accuracies from trials in a noisy condition with -1.1 dB SNR. For 180° separation between

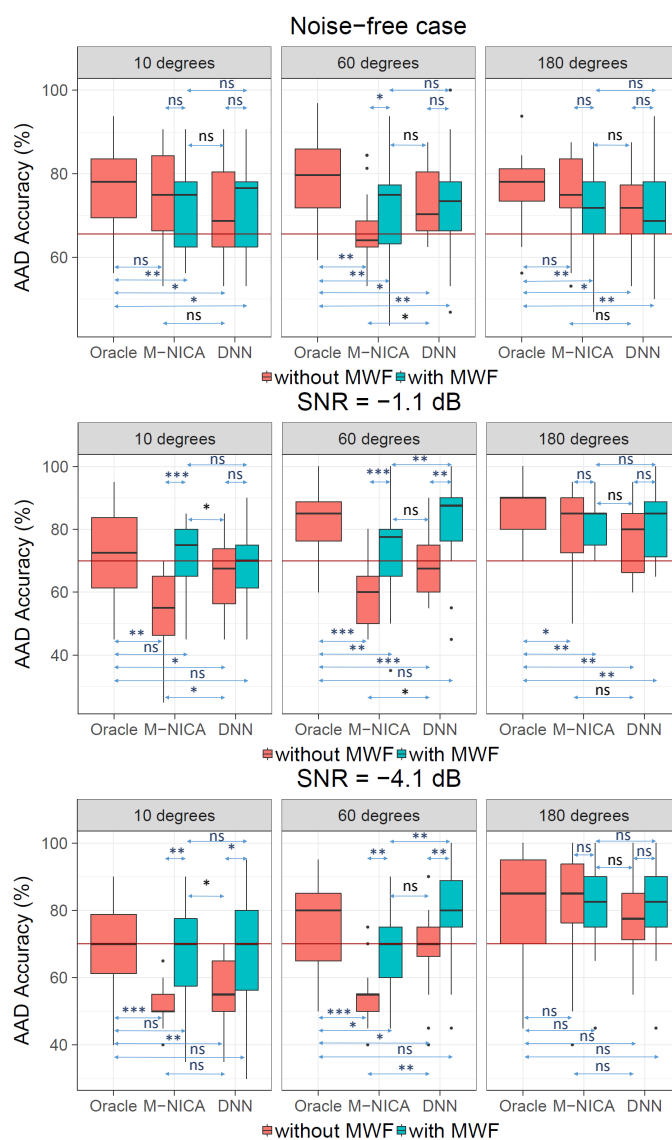


Fig. 3. AAD accuracies (on 30 s trials) of 18 subjects for different speech separation approaches for 3 cases: the noise-free case, -1.1 dB SNR, and -4.1 dB SNR. The results of using the M-NICA and DNN approaches are compared with the results when using oracle speech envelopes for AAD. The separated streams were either directly used for AAD (red boxes here and dashed red lines in Fig. 1), or used to generate VAD tracks for MWFs (green boxes here and green lines in Fig. 1), the outputs of which were then used for AAD. The red line indicates the significance level. Comparison between the different approaches were done using Wilcoxon’s signed-rank test with Holm correction. : ‘***’ for $p < 0.001$, ‘**’ for $p < 0.01$, ‘*’ for $p < 0.05$, ‘ns’ for no significant difference.

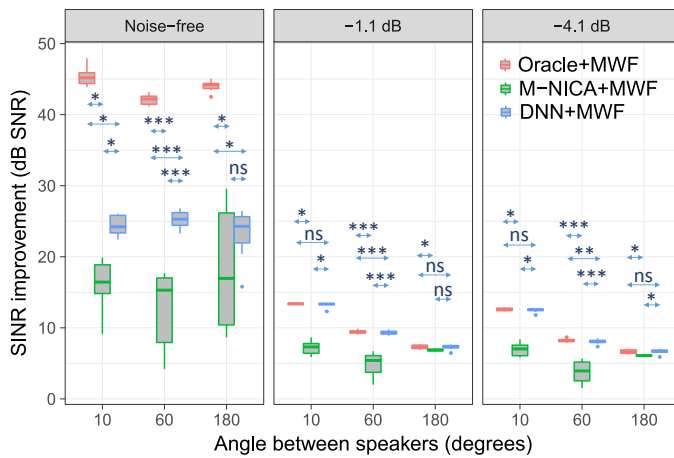


Fig. 4. Boxplots showing improvement in signal to interference-plus-noise ratios of 8 story parts when the MWFs use VAD tracks extracted from the different speech separation approaches. The results of using M-NICA and DNN outputs for the VAD are compared with the results when using oracle speech for the VAD.

speakers, it can be seen that all source separation approaches, with or without the MWF filtering step, result in similar AAD performances, but do not match up to that of oracle speech envelopes. Nevertheless, the performance decrease is minimal, in particular when an MWF is used. For 10° and 60° separation between speakers, the DNN envelopes result in better performance than the M-NICA envelopes. MWF filtering substantially improves results for both approaches for the 60° case, where the DNN+MWF combination even matches performance with the oracle speech envelopes. For the 10° case, only the M-NICA approach requires MWF filtering to match oracle speech performance. For all the angles, we can see that the performance of the purely linear M-NICA+MWF pipeline is either on par or better than the purely non-linear DNN approach without the linear beamformer extension with MWF.

For the -4.1 dB SNR (Fig. 3), which is a highly challenging condition, we observe that for 180° separation, the performances of all the approaches do not differ significantly from each other. However, for the 10° condition, the M-NICA and DNN envelopes result in accuracies below chance level, and only with the additional MWF filtering, do the results match up to those of oracle speech envelopes. In the 60° case, the MWF filtering does improve results for both the source separation approaches, however only the DNN+MWF approach's results match up to those of oracle speech envelopes. In addition, we again observe that the purely linear approach of M-NICA+MWF performs either at par or better than the purely non-linear DNN approach without the beamformer extension.

C. Speech enhancement performance

To estimate the effectiveness of the MWF for speech enhancement, we looked at the improvement in the signal to interference-plus-noise ratios (SINRs) at the output of the MWFs when using VAD tracks generated from 3 sources: the oracle speech envelopes, the M-NICA envelopes, and the DNN

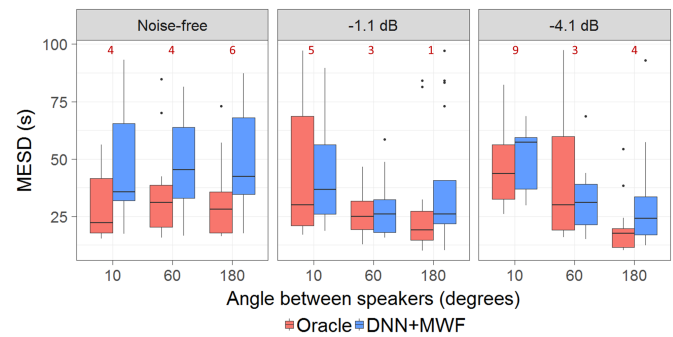


Fig. 5. Boxplots showing MESDs for two cases: when attention decoding was performed using ‘Oracle’ speech envelopes versus when using ‘DNN+MWF’ speech envelopes. The numbers in red indicate the number of outliers (MESD > 100s) which were removed from the plot (such high values are due to near-chance level performances, which are too low to achieve an acceptable gain control, leading to extremely high MESDs). They were included in the statistical analysis.

envelopes. The SINR is defined as the ratio of the power of the attended speaker to the total of the power of the unattended speaker and the background noise. The reference input SINR is taken as the highest SINR among the 6 microphones (note that the achieved SINR improvement will be larger for the other microphones). For each SNR and angle between speakers, we computed the improvement in SINR for 8 story parts to which the subjects had to listen without interruption. Fig. 4 shows the improvement in SINRs of the 8 story parts, for the 3 SNRs and 3 different angles between speakers. For the noise-free case, we observe that the oracle speech envelopes result in more speech enhancement than M-NICA or DNN envelopes in all separation angles (although perceptually there is not much difference: the difference between 40dB and 20dB SINR is hard to notice). However, in the presence of background noise, the DNN+MWF approach significantly outperforms the M-NICA+MWF approach, and matches up to the performance with the oracle speech envelopes in the 10° and 180° cases. For the acoustically more challenging case of 60° (since the speakers are on the same side of the head), we observe that the DNN+MWF approach results in significantly lower SINRs than those of oracle+MWF, but they are significantly higher than the M-NICA+MWF approach.

D. Gain control performance

AAD accuracies are dependent on evaluated trial lengths, which was set to 30 s in our experiments. In the context of gain control in a neuro-steered hearing device, longer trial lengths can result in higher accuracies, but at the cost of slow decisions affecting the speed with which the device can follow a switch in attention of the user. The minimal expected switching duration (MESD) [56] is a performance metric for AAD algorithms that takes this trade-off into account. This metric finds the optimal operating point to minimize the expected time required for such a hearing device to switch between speakers in a robust fashion after the user switches attention. We used a toolbox [57] to compute the MESDs of the best

performing algorithm (DNN+MWF) and the oracle case for the 3 different angles between speakers and the 3 SNRs. The results are shown in Fig. 5. Using Wilcoxon’s signed-rank test with Holm correction, we found that the MESDs of the oracle case were not significantly lower than those of the DNN+MWF pipeline, except for all speaker separations in the noise-free condition ($p = 0.05$), and for 180° at -1.1 dB SNR ($p < 0.05$). For all other cases, there were no significant differences between the two approaches.

V. DISCUSSION

Previous studies that combine speech enhancement/demixing with AAD investigated only one type of algorithm pipeline (either linear [20], [22], [23] or DNN-based [31], [32]), mostly in relatively mild acoustic conditions. Furthermore each study has its own specific set-up (single- vs. multi-microphone, ECoG vs. EEG, matched or mismatched acoustic conditions during EEG recordings, etc.) which does not allow to directly compare the results achieved in these studies. The study in this paper was designed to investigate the feasibility of a neuro-steered speech enhancement pipeline in challenging acoustic conditions with negative SNRs and various speaker positions, while also comparing different source separation approaches to produce the speech envelopes used by the attention decoding block. We investigated the AAD performance of a training-free and purely linear audio signal processing algorithm pipeline (M-NICA+MWF) on the one hand, and a non-linear DNN on the other hand, as well as a combination of both (DNN+MWF). In the remainder of this section, we will discuss the main findings of our study.

A. Multi-microphone outperforms single-microphone

Existing AAD studies using DNNs for the acoustic source separation were based on single-microphone recordings [31], [32]. We observed that a single-channel approach generalizes less well across datasets and that a multi-channel approach is more robust in terms of separation quality. In addition, envelopes extracted from the output of the multi-channel neural network resulted in AAD accuracies that were more robust to varying speaker positions compared to those from the single-channel neural network. We conclude that a multi-microphone set-up is crucial to obtain a sufficiently high AAD accuracy in practical settings. Thus we chose the multi-channel neural network approach for our comprehensive analysis of the neuro-steered source separation pipeline.

B. Linear pipeline outperforms DNN

We found that the AAD performance with the linear M-NICA+MWF pipeline from [20], [22] is at least on par and often outperforms that of a non-linear DNN approach in all investigated conditions. In addition, the linear pipeline is computationally cheaper than the DNN approach (in terms of operations per second and memory size), while also being training-free thereby not depending on representative training data. Furthermore, implementing the MWF as an adaptive, causal, low-latency filter is relatively straightforward, and

latencies can be limited to less than 5 ms [58] (see also section V-D below). Speech separation based on deep clustering on the other hand, usually uses a non-causal network where the computation of the embedding vector at present time is based on all past and future time samples, although recently also causal DNNs for speech enhancement have been proposed [32], [59].

C. Combining the best of both worlds: DNN and MWF

While linear and non-linear approaches both individually result in good speech separation, a combination of both - as in the DNN+MWF approach - resulted in the best AAD performance, particularly in challenging acoustic conditions. In this case the (non-linear) DNN is used as a VAD mechanism to inform the (linear) MWF which performs the actual speaker separation and denoising, thereby combining the best of both worlds. This results in the best AAD performance as well as the highest SINR improvement.

D. Real-time aspects

When using a weighted overlap-add (WOLA) procedure [60], the algorithmic delay of the MWF is equal to the window length used in the STFT computations, i.e., the number of samples over which the discrete Fourier transform is computed. For example, for a hearing aid using a sampling rate of 20480 Hz and a 96-point STFT windowing [58], the algorithmic delay is less than 5 ms. Note that a delay in the VAD information (e.g. due to possible non-causality of the BSS block in Fig. 1), does not have an impact on the overall input-output delay of the stimulus presented to the user. This is because the adaptation of the MWF filter coefficients can be decoupled from the actual filtering operation itself, where the former can lag behind on the latter. This implies that any delay or non-causality in the M-NICA or DNN pipeline will not create a bottleneck towards real-time speech processing. It will only add some inertia in the updating of the MWF coefficients to adapt to changes in the acoustic scenario.

E. EEG-based AAD is feasible in challenging low-SNR conditions

The SNRs in this study were chosen such that the speech intelligibility varied across conditions. The estimated speech recognition threshold (SRT), at which a normal-hearing subject can understand 50% of the attended story on average, was found to be -7.1 dB (see [42] for details of the estimation procedure). The SNRs -4.1 dB and -1.1 dB were chosen by adding 3 dB steps from the 50% speech intelligibility point. The -7.1 dB condition was excluded from this study due to the difficulty subjects faced in focusing on the attended speaker (as reflected in the AAD accuracies under this condition reported in [42]). Subjective speech intelligibility reported during the recording sessions for the different angles between speakers show that -4.1 dB and -1.1 dB SNRs also are not easy listening conditions. Compared to [23] where the noise conditions are relatively mild (4 dB and 9 dB SNR), [32] where there is no background noise, and [22] where the EEG

recording conditions did not match the acoustic conditions, our current study provides a holistic view on the performance of a neuro-steered speech enhancement pipeline under realistic to challenging acoustic conditions. We find that not only EEG-based AAD with no access to clean-speech sources is feasible under challenging acoustic conditions, but also that the AAD accuracies and improvement in SINRs for a DNN+MWF pipeline get remarkably close to oracle performance.

F. Future outlook

The proposed neuro-steered speech enhancement pipeline and its analysis over a range of acoustic conditions is a step towards effectively incorporating attention decoding and source separation algorithms in neuro-steered hearing prostheses. With the field of deep learning expanding and exerting its presence in various scientific domains, we have tried to address the possibility of deep learning based source separation in this context and systematically compare it with a classical signal processing approach. While the realisation of neuro-steered hearing prostheses is still far from being a reality, there are ongoing advancements that can accelerate this process. In addition to hardware requirements such as miniaturisation of neural recording equipment, better computational and storage capabilities in hearing devices etc., also newer approaches for improved attention decoding itself based on DNNs [11]–[13], canonical correlation analysis [61] or state-space modelling [3], [10] can contribute to faster and more robust attention decoding and hence better neuro-steered speech enhancement pipelines.

VI. CONCLUSION

In a multi-speaker ‘cocktail party’ scenario, a hearing aid’s noise reduction algorithms does not have the knowledge of which speaker the user intends to listen to. Incorporating AAD algorithms to detect the attention of the user, leads to the concept of so-called neuro-steered hearing aids. In this study we present a neuro-steered speech enhancement pipeline, with no access to clean speech signals in challenging acoustic scenarios.

We analyzed the performance of a linear as well as DNN-based speaker separation approach for the pipeline. We found that a purely linear approach often outperforms the DNN-based approach. However, the best performance, under challenging scenarios, was obtained when combining the best of both worlds, i.e., extending the DNN with a linear data-driven beamformer. In addition, for the DNN-based speaker separation, we have demonstrated the benefit of using a multi-microphone system compared to a single-microphone system. An additional advantage of using a linear data-driven beamformer is that it decouples the algorithmic delay in the blind source separation of the audio signals from the total system delay, thus facilitating robust real-time speech processing. With this study, we present a proof-of-concept of feasibility of a neuro-steered speech enhancement pipeline in challenging acoustic conditions.

REFERENCES

- [1] N. Ding and J. Z. Simon, “Neural coding of continuous speech in auditory cortex during monaural and dichotic listening,” *Journal of neurophysiology*, vol. 107, no. 1, pp. 78–89, 2012.
- [2] J. A. O’Sullivan, A. J. Power, N. Mesgarani, S. Rajaram, J. J. Foxe, B. G. Shinn-Cunningham, M. Slaney, S. A. Shamma, and E. C. Lalor, “Attentional selection in a cocktail party environment can be decoded from single-trial EEG,” *Cerebral Cortex*, p. bht355, 2014.
- [3] S. Akram, A. Presacco, J. Z. Simon, S. A. Shamma, and B. Babadi, “Robust decoding of selective auditory attention from MEG in a competing-speaker environment via state-space modeling,” *NeuroImage*, vol. 124, pp. 906–917, 2016.
- [4] N. Ding and J. Z. Simon, “Emergence of neural encoding of auditory objects while listening to competing speakers,” *Proc. National Academy of Sciences*, vol. 109, no. 29, pp. 11 854–11 859, 2012.
- [5] E. M. Z. Golombic, N. Ding, S. Bickel, P. Lakatos, C. A. Schevon, G. M. McKhann, R. R. Goodman, R. Emerson, A. D. Mehta, J. Z. Simon *et al.*, “Mechanisms underlying selective neuronal tracking of attended speech at a “cocktail party”,” *Neuron*, vol. 77, no. 5, pp. 980–991, 2013.
- [6] N. Mesgarani and E. F. Chang, “Selective cortical representation of attended speaker in multi-talker speech perception,” *Nature*, vol. 485, no. 7397, pp. 233–236, 2012.
- [7] E. B. Petersen, M. Wöstmann, J. Obleser, and T. Lunner, “Neural tracking of attended versus ignored speech is differentially affected by hearing loss,” *Journal of neurophysiology*, vol. 117, no. 1, pp. 18–27, 2016.
- [8] L. Decruy, J. Vanthornhout, and T. Francart, “Hearing impairment is associated with enhanced neural tracking of the speech envelope,” *bioRxiv*, p. 815530, 2019.
- [9] A. Presacco, J. Z. Simon, and S. Anderson, “Speech-in-noise representation in the aging midbrain and cortex: Effects of hearing loss,” *PLoS one*, vol. 14, no. 3, p. e0213899, 2019.
- [10] S. Miran, S. Akram, A. Sheikhattar, J. Z. Simon, T. Zhang, and B. Babadi, “Real-time tracking of selective auditory attention from M/EEG: A bayesian filtering approach,” *Frontiers in neuroscience*, vol. 12, 2018.
- [11] T. de Taillez, B. Kollmeier, and B. T. Meyer, “Machine learning for decoding listeners’ attention from electroencephalography evoked by continuous speech,” *European Journal of Neuroscience*, 2017.
- [12] L. Deckers, N. Das, A. H. Ansari, A. Bertrand, and T. Francart, “EEG-based detection of the attended speaker and the locus of auditory attention with convolutional neural networks,” *bioRxiv*, p. 475673, 2018.
- [13] G. Ciccarelli, M. Nolan, J. Perricone, P. T. Calamia, S. Haro, J. O’Sullivan, N. Mesgarani, T. F. Quatieri, and C. J. Smalt, “Comparison of two-talker attention decoding from EEG with nonlinear neural networks and linear methods,” *Scientific reports*, vol. 9, no. 1, pp. 1–10, 2019.
- [14] W. Biesmans, N. Das, T. Francart, and A. Bertrand, “Auditory-inspired speech envelope extraction methods for improved EEG-based auditory attention detection in a cocktail party scenario,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 5, pp. 402–412, 2017.
- [15] G. M. Di Liberto, J. A. O’Sullivan, and E. C. Lalor, “Low-frequency cortical entrainment to speech reflects phoneme-level processing,” *Current Biology*, vol. 25, no. 19, pp. 2457–2465, 2015.
- [16] B. Mirkovic, M. G. Bleichner, M. De Vos, and S. Debener, “Target speaker detection with concealed EEG around the ear,” *Frontiers in neuroscience*, vol. 10, p. 349, 2016.
- [17] R. Zink, A. Baptist, A. Bertrand, S. Van Huffel, and M. De Vos, “Online detection of auditory attention in a neurofeedback application,” in *Proc. 8th International Workshop on Biosignal Interpretation*, Osaka, Japan, 2016, pp. 1–4.
- [18] L. Fiedler, M. Woestmann, C. Graversen, A. Brandmeyer, T. Lunner, and J. Obleser, “Single-channel in-ear-EEG detects the focus of auditory attention to concurrent tone streams and mixed speech,” *Journal of Neural Engineering*, vol. 14, no. 3, p. 036020, 2017.
- [19] A. M. Narayanan and A. Bertrand, “Analysis of miniaturization effects and channel selection strategies for EEG sensor networks with application to auditory attention detection,” *IEEE Transactions on Biomedical Engineering*, 2019.
- [20] S. Van Eyndhoven, T. Francart, and A. Bertrand, “EEG-informed attended speaker extraction from recorded speech mixtures with application in neuro-steered hearing prostheses,” *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 5, pp. 1045–1056, 2017.

- [21] A. Bertrand and M. Moonen, "Blind separation of non-negative source signals using multiplicative updates and subspace projection," *Signal Processing*, vol. 90, no. 10, pp. 2877–2890, 2010.
- [22] N. Das, S. Van Eyndhoven, T. Francart, and A. Bertrand, "EEG-based attention-driven speech enhancement for noisy speech mixtures using N-fold multi-channel Wiener filters," in *25th European Signal Processing Conference (EUSIPCO)*, Kos, Greece, 2017, pp. 1660–1664.
- [23] A. Aroudi and S. Doclo, "Cognitive-driven binaural LCMV beamformer using EEG-based auditory attention decoding," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Brighton, UK: IEEE, 2019, pp. 406–410.
- [24] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Shanghai, China: IEEE, 2016, pp. 31–35.
- [25] M. Kolbæk, D. Yu, Z.-H. Tan, J. Jensen, M. Kolbaek, D. Yu, Z.-H. Tan, and J. Jensen, "Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks," *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, vol. 25, no. 10, pp. 1901–1913, 2017.
- [26] Z. Chen, Y. Luo, and N. Mesgarani, "Deep attractor network for single-microphone speaker separation," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. New Orleans, USA: IEEE, 2017, pp. 246–250.
- [27] P. Appeltans, J. Zegers, and H. Van hamme, "Practical applicability of deep neural networks for overlapping speaker separation," in *Interspeech*. Graz, Austria: ISCA, 2019, pp. 1353–1357.
- [28] G. Wichern, J. Antognini, M. Flynn, L. R. Zhu, E. McQuinn, D. Crow, E. Manilow, and J. L. Roux, "Wham!: Extending speech separation to noisy environments," in *Interspeech*. Graz, Austria: ISCA, 2019, pp. 1368–1372.
- [29] Z.-Q. Wang, J. Le Roux, and J. R. Hershey, "Multi-channel deep clustering: Discriminative spectral and spatial embeddings for speaker-independent speech separation," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Calgary, Canada: IEEE, 2018, pp. 1–5.
- [30] Z. Chen, J. Li, X. Xiao, T. Yoshioka, H. Wang, Z. Wang, and Y. Gong, "Cracking the cocktail party problem by multi-beam deep attractor network," in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. Okinawa, Japan: IEEE, 2017, pp. 437–444.
- [31] J. O'Sullivan, Z. Chen, J. Herrero, G. M. McKhann, S. A. Sheth, A. D. Mehta, and N. Mesgarani, "Neural decoding of attentional selection in multi-speaker environments without access to clean sources," *Journal of neural engineering*, vol. 14, no. 5, p. 056001, 2017.
- [32] C. Han, J. O'Sullivan, Y. Luo, J. Herrero, A. D. Mehta, and N. Mesgarani, "Speaker-independent auditory attention decoding without access to clean speech sources," *Science advances*, vol. 5, no. 5, p. eaav6134, 2019.
- [33] A. Bertrand and M. Moonen, "Energy-based multi-speaker voice activity detection with an ad hoc microphone array," in *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*. Dallas, Texas, USA: IEEE, 2010, pp. 85–88.
- [34] S. Doclo and M. Moonen, "GSVD-based optimal filtering for single and multimicrophone speech enhancement," *IEEE Transactions on Signal Processing*, vol. 50, no. 9, pp. 2230 – 2244, Sep. 2002.
- [35] 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. Audio, Speech, and Language Processing*, vol. 22, no. 4, pp. 785–799, 2014.
- [36] B. Cornelis, M. Moonen, and J. Wouters, "Reduced-bandwidth multi-channel Wiener filter based binaural noise reduction and localization cue preservation in binaural hearing aids," *Signal Processing*, vol. 99, pp. 1–16, 2014.
- [37] T. Gerkmann and R. C. Hendriks, "Noise power estimation based on the probability of speech presence," in *2011 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. New Paltz, NY, USA: IEEE, 2011, pp. 145–148.
- [38] S. T. Roweis, "One microphone source separation," in *Advances in neural information processing systems*, 2001, pp. 793–799.
- [39] Y. Wang, A. Narayanan, and D. Wang, "On training targets for supervised speech separation," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 22, no. 12, pp. 1849–1858, 2014.
- [40] Y. Isik, J. Le Roux, Z. Chen, S. Watanabe, and J. R. Hershey, "Single-channel multi-speaker separation using deep clustering," in *Interspeech*. San Francisco, California: ISCA, 2016, pp. 545–549.
- [41] T. Yoshioka, H. Erdogan, Z. Chen, and F. Alleva, "Multi-microphone neural speech separation for far-field multi-talker speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5739–5743.
- [42] N. Das, A. Bertrand, and T. Francart, "EEG-based auditory attention detection: boundary conditions for background noise and speaker positions," *Journal of neural engineering*, vol. 15, no. 6, p. 066017, 2018.
- [43] "Radioboeken," <http://www.radioboeken.eu/radioboeken.php?lang=NL>, 2017, [Online; accessed: May-2017].
- [44] H. Kayser, S. D. Ewert, J. Anemüller, T. Rohdenburg, V. Hohmann, and B. Kollmeier, "Database of multichannel in-ear and behind-the-ear head-related and binaural room impulse responses," *EURASIP Journal on Advances in Signal Processing*, vol. 2009, p. 6, 2009.
- [45] L. Van Eerten, "Over het corpus gesproken nederlands," *Nederlandse Taalkunde*, vol. 12, no. 3, pp. 194–215, 2007.
- [46] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Queensland, Australia: IEEE, 2015, pp. 5206–5210.
- [47] N. Das, W. Biesmans, A. Bertrand, and T. Francart, "The effect of head-related filtering and ear-specific decoding bias on auditory attention detection," *Journal of neural engineering*, vol. 13, no. 5, p. 056014, 2016.
- [48] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proc. 26th annual international conference on machine learning*. Montreal, Quebec, Canada: ACM, 2009, pp. 41–48.
- [49] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [50] R. Reed and R. J. Marks II, *Neural smithing: supervised learning in feedforward artificial neural networks*. MIT Press, 1999.
- [51] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin *et al.*, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," *arXiv preprint arXiv:1603.04467*, 2016.
- [52] R. D. Patterson, M. H. Allerhand, and C. Giguere, "Time-domain modeling of peripheral auditory processing: A modular architecture and a software platform," *The Journal of the Acoustical Society of America*, vol. 98, no. 4, pp. 1890–1894, 1995.
- [53] S. S. Stevens, "The measurement of loudness," *The Journal of the Acoustical Society of America*, vol. 27, no. 5, pp. 815–829, 1955.
- [54] E. Vincent, R. Gribonval, and C. Févotte, "Performance measurement in blind audio source separation," *IEEE transactions on audio, speech, and language processing*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [55] J. Zegers and H. Van hamme, "Memory time span in lstms for multi-speaker source separation," in *Interspeech*. Hyderabad, India: ISCA, 2018, pp. 1477–1481.
- [56] S. Geirnaert, T. Francart, and A. Bertrand, "An interpretable performance metric for auditory attention decoding algorithms in a context of neuro-steered gain control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 28, pp. 307–317, 2019. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/8895794/>
- [57] —, "MESD toolbox," August 2019. [Online]. Available: <https://github.com/exporl/mesd-toolbox>
- [58] T. Van den Bogaert, S. Doclo, J. Wouters, and M. Moonen, "Speech enhancement with multichannel wiener filter techniques in multimicrophone binaural hearing aids," *The Journal of the Acoustical Society of America*, vol. 125, no. 1, pp. 360–371, 2009.
- [59] C. Han, Y. Luo, and N. Mesgarani, "Online deep attractor network for real-time single-channel speech separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Brighton, UK: IEEE, 2019, pp. 361–365.
- [60] R. Crochiere, "A weighted overlap-add method of short-time Fourier analysis/synthesis," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 28, no. 1, pp. 99–102, 1980.
- [61] A. de Cheveigné, D. D. Wong, G. M. D. Liberto, J. Hjortkjaer, M. Slaney, and E. Lalor, "Decoding the auditory brain with canonical component analysis," *NeuroImage*, vol. 172, pp. 206 – 216, 2018. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1053811918300338>