A New Metric to Evaluate Auditory Attention Detection Performance Based on a Markov Chain

Simon Geirnaert[†*], Tom Francart[*], and Alexander Bertrand[†]

[†]KU Leuven, Department of Electrical Engineering (ESAT),
STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, Belgium;
[*]KU Leuven, Department of Neurosciences, ExpORL, Belgium.

Abstract—**Auditory attention detection (AAD) is an essential building block for future generations of 'neuro-steered' hearing prostheses. In a multi-speaker scenario, it uses neural recordings to detect to which speaker the listener is attending and assists as such the noise reduction algorithm within the hearing device. Recently, a multitude of these AAD algorithms has been developed, based on electroencephalography (EEG) recordings. With the emergence of AAD algorithms, a standardized way of evaluating these AAD algorithms becomes paramount. However, this is not trivial due to an inherent trade-off between detection delay and accuracy. We propose a new performance metric to evaluate AAD algorithms that resolves this trade-off: the expected switch duration (ESD). The ESD is based on a Markov chain model of an adaptive gain control system in a hearing aid and quantifies the expected time needed to switch its operation from one speaker to another when the attention is switched. We validate the metric on simulated data and show on real EEG recordings that it is an interpretable metric that allows fair comparison between AAD algorithms, combining both the accuracy of the AAD algorithm and the time needed to make a decision.**

I. INTRODUCTION

The human brain is capable of focusing attention on a specific speaker in the presence of background noise, including competing speakers [1]. However, people with hearing impairments have major difficulties in understanding the attended speaker. Although current hearing aids are able to filter out the targeted speaker, and as such overcome the difficulties of the hearing impaired, these devices are not yet able to incorporate the attentional process. Suboptimal heuristics, such as speaker loudness, are currently used to select the targeted speaker. Recent advances have however shown that it is possible to decode attention directly from the brain, for example via the electroencephalogram (EEG) (e.g., [2]–[6]). These auditory attention detection (AAD) algorithms are essential as a building block in future generations of 'neuro-steered' hearing aids [7].

As there is an increase in efforts to design such AAD algorithms, it becomes important to have proper tools at our disposal to evaluate these AAD algorithms. Currently, AAD algorithms are evaluated by means of the detection accuracy: the percentage of decision windows of EEG and audio data

in which the attention is decoded properly. However, the accuracy depends on the length of the decision window, which consequently is an important parameter in the interpretation of the results. This method of evaluation inherently has some disadvantages. Firstly, the accuracy is often evaluated at different decision window lengths, resulting in multiple metrics (one for each window length), making it hard to draw a general conclusion. Moreover, a different choice of decision window lengths at which the accuracy is evaluated hampers comparison and could lead to different conclusions and thus to inconclusiveness. Furthermore, it is not clear which decision window lengths are more relevant in the context of neuro-steered hearing aids and are thus more important to base performance evaluation on.

There have been attempts to overcome some of these issues, for example by using the information transfer rate as metric, which quantifies the number of bits that can be transmitted per second [3]. However, there is not yet a performance metric that meets the following requirements:

- *Interpretable*: the performance metric should be interpretable in the hearing aid context. The information transfer rate lacks a clear interpretation in the context of AAD.
- *Single-number*: the performance metric should summarize the complete performance of the AAD algorithm in a single number, facilitating comparison between AAD algorithms.
- *Combining accuracy and decision time*: in the hearing-aid use case, the time needed to make an AAD decision (the decision window length), is equally important as the accuracy at that specific window length, as it also determines how fast a user can switch between two speakers. The performance metric should thus integrate both accuracy and decision time.
- *Independent of evaluated decision window lengths*: the performance metric should be as independent as possible of the used window lengths to evaluate the accuracy, in order to ease the comparison between scientific reports.

The lack of a performance metric that integrates the previous requirements motivates the design of a new metric: the *expected switch duration* (ESD). We build up the theory behind this metric in Section II, and validate the metric on simulations
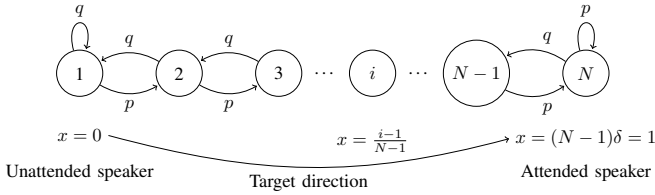
Fig. 1: This Markov chain models an adaptive gain control system in a neuro-steered hearing aid. Each state i is linked to a relative amplification x of the targeted speaker versus the background noise.

and show that it meets the listed requirements by evaluating it on real EEG recordings in Section III.

II. EXPECTED SWITCH DURATION

The ESD performance metric is based on a Markov chain model of an adaptive gain control system in a hearing aid (Section II-A). Essentially, it quantifies the time that the hearing device requires to adapt its operation, when the attention is switched, between two predefined stable working regions, each corresponding to a different attended speaker (note that we further on assume a two-speaker setting for simplicity). By predefining the lower bound of the P_0-confidence interval to be above a comfort level for the user (representing the stable working region), the number of states of the Markov chain can be optimized (Section II-C). Based on this optimized Markov chain, the ESD can be rigorously defined (Section II-D).

A. Markov chain model

Figure 1 shows a Markov chain of N states as a model for a smooth gain control system, which is desired in a user-friendly neuro-steered hearing aid. Each state corresponds to a relative amplification $x \in [0, 1]$ of the targeted speaker versus the background noise (e.g., including a second speaker). $x = 1$ matches a certain target gain (maximal amplification) of the targeted speaker, while still allowing a switch of attention to the other speaker. $x = 0$ represents the symmetric case, where the unattended speaker is amplified at a similar level. In case the attention is switched to this speaker, the Markov chain is reversed, such that the previously unattended speaker corresponds to the target direction state. $x = 0.5$ implies an equal amplification of both speakers.

In a realistic setting, incoming EEG and audio data is buffered for τ seconds (the decision window length), such that every τ seconds, a new decision can be made. The application of an AAD algorithm in a real-time fashion then corresponds to a walk through the Markov chain where an incorrect AAD decision results in a step back, thereby increasing the amplification of the wrong speaker. The transition probability p is equal to the probability of a correct AAD decision, while $q = 1 - p$ is the probability of an incorrect AAD decision. We assume that $p > 0.5$, i.e., the applied AAD algorithm performs better than chance level.

The use of a Markov chain model increases (perceptual) comfort for the hearing aid user by switching from one speaker to another with smooth transitions, avoiding sudden changes in the dominantly perceived speaker. Furthermore, it enables the user to correct the behavior of the system when incorrect AAD decisions are made that make that the system starts switching to the unattended speaker.

B. Steady-state distribution

The steady-state distribution of the Markov chain in Figure 1 is defined by the per-state probability $\pi(i) = P(x = \frac{i-1}{N-1}), i \in \{1, \ldots, N\}$ to be in state i after an infinite number of steps, starting from any state, while the transition probability p is fixed. The global balance equations and a normalization condition can be used to compute this steady-state distribution [8]:

$$\begin{cases} \pi(i) = \sum_{l=1}^{N} \pi(l)p_{li}, & \text{(balance equations)} \\ \sum_{l=1}^{N} \pi(l) = 1 & \text{(normalization condition)} \end{cases},$$

where p_{li} corresponds to the transition probability from state l to state i. The steady-state distribution $\pi(i)$ can be found by recursively writing out the balance equations starting from $\pi(1)$ (derivation omitted due to space constraints). Defining $\frac{p}{q} = r$, we obtain $\pi(i) = r^{i-1}\pi(1)$, where $\pi(1)$ can be found from the normalization condition:

$$\sum_{l=1}^{N} \pi(l) = 1 \Rightarrow \pi(1) = \frac{r-1}{r^N - 1}.$$

Eventually, we find the following steady-state distribution:

$$\pi(i) = \frac{r-1}{r^N - 1}r^{i-1}, i \in \{1, \ldots, N\}. \tag{1}$$

C. P_0-confidence interval

The steady-state distribution (1) can be used to determine the P_0-confidence interval. The P_0-confidence interval corresponds to the smallest set of neighboring states in which the system resides, in a steady-state regime, for at least P_0 percent (e.g., 80%) of the time. In the context of a neuro-steered hearing aid, we interpret it as an optimal working region of relative gains $[\bar{x}, 1]$ in which the system operates for at least P_0 percent of the time, regardless of AAD errors that cause transitions opposite to the target direction.

The P_0-confidence interval is defined by a lower bound state $\bar{k}$, which is equal to the largest i for which:

$$\sum_{j=i}^{N} \pi(j) \geq P_0$$

$$\overset{(1)}{\Leftrightarrow} \frac{r-1}{r^N - 1}\sum_{j=i}^{N} r^{j-1} \geq P_0$$

$$\Leftrightarrow \frac{r^N - r^{i-1}}{r^N - 1} \geq P_0$$

$$\overset{r \geq 1}{\Leftrightarrow} r^N - r^N P_0 + P_0 \geq r^{i-1}$$

$$\Leftrightarrow \frac{\log(r^N - r^N P_0 + P_0)}{\log(r)} + 1 \geq i.$$

The last steps are valid because $r > 1$ due to the assumption that $p > 0.5$ and because the log-function is a monotonically increasing function. The lower bound state $\bar{k}$ of the P_0-confidence interval is thus equal to:

$$\bar{k} = \left\lfloor \frac{\log(r^N - r^N P_0 + P_0)}{\log(r)} + 1 \right\rfloor, \quad (2)$$

with $\lfloor \cdot \rfloor$ the flooring operation. As a result, the P_0-confidence interval is defined as $[\bar{x}, 1] = [\frac{\bar{k}-1}{N-1}, 1]$

Eventually, the number of states N in the Markov chain has to be minimized to minimize the duration of the switch. However, in order to model a realistic neuro-steered hearing aid system, extra constraints on N, relating to the P_0-confidence interval and the smoothness, are needed:

- $\bar{x} \in [c, 1]$: the lower bound of the P_0-confidence interval should be at least equal to a predefined desired minimum relative gain c, which ensures a comfortable level for the listener to sufficiently understand the target speaker. We thus choose the region in which the hearing aid should operate for at least P_0 percent of the time. This results in a constraint when minimizing N:

$$\bar{x} = \frac{\bar{k}-1}{N-1} \geq c. \quad (3)$$

It can be proven that there exists always such an N (proof omitted).

- $N \geq N_{\min}$: in order to fully realize the potential of the Markov chain as an adaptive gain control system to implement a user-friendly hearing aid (see Section II-A), a minimal number of states is needed. As a minimal 'smoothness' constraint, we want to prevent that the system *always* (for every p and P_0 such that $p < P_0$) crosses $x = 0.5$ (i.e., the unattended speaker is dominantly amplified) when leaving the P_0-confidence interval due to an AAD error. This corresponds to putting $N_{\min} = 5$.

The optimal number of states N can thus be found by iterating over $N = N_{\min}, N_{\min} + 1, N_{\min} + 2, \ldots$, in this specific order, until an N is found that satisfies (3).

D. Expected switch duration

In order to quantify the duration of the switch in the Markov chain model, we use the *mean hitting time* as a basic metric. The mean hitting time $h_j(i)$ quantifies the expected number of steps needed to hit a target state j for the first time, starting from a certain state i. It is defined as follows:

$$h_j(i) = \mathbb{E}\{s|i \to j\} = \sum_{s=0}^{+\infty} sP(s|i \to j), \quad (4)$$

with $i, j \in \{1, \ldots, N\}$, $\mathbb{E}\{\cdot\}$ denoting the expectation operator and where $P(s|i \to j)$ is the probability that target state j is reached for the first time after s steps, when starting in state i. Using the recursive definition in [8] results in the following expression for the mean hitting time, when $i < j$ (derivation omitted due to space constraints):

$$h_j(i) = \frac{j-i}{2p-1} + \frac{p(r^{-j} - r^{-i})}{(2p-1)^2}, \text{ for } i \leq j. \quad (5)$$

We define the duration of the switch as the expected number of steps needed to enter the optimal working region when starting outside of that region. In other words, it is the expected number of steps to go from *any* state $i < k$ to state $k = \lceil c(N-1)+1 \rceil$, corresponding to the first state above or equal to the predefined comfort level c:

$$\mathbb{E}\{s|i \to k, \forall\, i < k\} = \sum_{s=0}^{+\infty} sP(s|i \to k, \forall\, i < k),$$

with $P(s|i \to k, \forall\, i < k)$ the probability that target state k is reached for the first time after s steps, when starting from *any* state $i < k$. By marginalizing in the initial state i, we find:

$$\mathbb{E}\{s|i \to k, \forall\, i < k\} = \sum_{s=0}^{+\infty} s \sum_{i=1}^{N} P(s|i \to k)P(i|i < k).$$

$P(i|i < k)$ can be found using Bayes' law:

$$\mathbb{E}\{s|i \to k, \forall\, i < k\} = \sum_{s=0}^{+\infty} s \sum_{i=1}^{N} P(s|i \to k) \frac{P(i < k|i)P(i)}{P(i < k)}.$$

The last expression can be simplified by using:

- $P(i) = \pi(N - i + 1)$, where the order in the steady-state distribution (1) is reversed, as the target direction reversed direction as well. Before the attention switch, the new interfering speaker was equal to the targeted speaker, such that as steady-state distribution the one before the currently indexed situation in Figure 1 should be taken. The steady-state distribution in the previous situation is equal to the reversed current steady-state distribution.
- $P(i < k|i) = 1$ when $i < k$, 0 otherwise.
- $P(i < k) = \sum_{l=1}^{k-1} \pi(N - l + 1)$, which we define as C.

Using these expressions and (4), the expected number of steps eventually becomes equal to:

$$\mathbb{E}\{s|i \to k, \forall\, i < k\} = \frac{1}{C} \sum_{i=1}^{k-1} \pi(N - i + 1)h_k(i). \quad (6)$$

The time needed to take a step, τ, which is equal to the decision window length in the context of an AAD algorithm, can be used to convert the number of steps into a time metric, the *expected Markov transit time* (EMTT) T [s]:

$$T(p(\tau), \tau, N) = \frac{\tau}{C} \sum_{i=1}^{k-1} \pi(N - i + 1)h_k(i), \quad (7)$$

with $\pi(N - i + 1)$ given by (1), $C = \sum_{l=1}^{k-1} \pi(N - l + 1)$ and $h_k(i)$ given by (5).

One of the main motivations to define a new metric is the need for a metric that quantifies the performance of an AAD algorithm regardless of the specific accuracies p at evaluated window lengths τ. To use a single EMTT as metric for an AAD algorithm, we first construct a $p(\tau)$-performance curve by piecewise linearly interpolating through the points $(\tau_i, p_i), i \in \{1, \ldots, I\}$ for which the AAD accuracy is actually

evaluated based on real data. The EMTT T (7) can then be minimized over the performance curve. To this end, we have to minimize the number of states for each (possibly interpolated) (τ, p)-pair using the guidelines in Section II-C, as the EMTT T (7) is monotonically non-decreasing with N. We call the resulting EMTT the *expected switch duration* (ESD) of an AAD algorithm, which is formally defined below:

Definition (Expected switch duration)
The expected switch duration (ESD) is the expected time required to reach a predefined stable working region defined via the comfort level c, after an attention switch, in an optimized Markov chain as a model for an adaptive gain control system in a neuro-steered hearing prosthesis. Formally, it is the expected time to reach the comfort level c in the fastest Markov chain with at least $N_{\min}$ states for which $\bar{x} \geq c$, i.e., the lower bound $\bar{x}$ of the P_0-confidence interval is above c:

$$\text{ESD} = \min_{N, \tau} \quad T(p(\tau), \tau, N)$$
$$\text{s.t.} \quad \bar{x} \in [c, 1] \quad (8)$$
$$N \geq N_{\min}$$

where $T(p(\tau), \tau, N)$ is defined in (7) and $\bar{x} = \frac{\bar{k}-1}{N-1}$, with $\bar{k}$ defined in (2).

First minimizing the number of states N, obeying the inequality constraints of (8) (see Section II-C), results in an optimal number of states $\hat{N}_\tau$. Plugging $\hat{N}_\tau$ afterwards into (8) results in an unconstrained optimization problem in τ:

$$\text{ESD} = \min_\tau \quad T(p(\tau), \tau, \hat{N}_\tau).$$

The ESD can finally be found as the minimal EMTT over all sampled window lengths τ on the $p(\tau)$-performance curve.

III. EXPERIMENTS

The lower bound $\bar{k}$ (2) and EMTT T (7) are validated via a simulation study (Section III-A). Afterwards, real EEG and audio data is used to show the computation of the ESD metric (Section III-B).

A. Simulation study

To validate the lower bound $\bar{k}$ (2) of the P_0-confidence interval, 10^6 Monte-Carlo runs are performed. In each run, 1000 decisions (0 or 1) are drawn from the Bernoulli distribution with a predefined probability of success p (the accuracy). We randomly select the initial state over a uniform distribution and then perform a walk through a Markov chain of 5 states. The final state after 1000 steps is considered as a sample from the steady-state distribution. Based on this sampled steady-state distribution over 10^6 runs, the state corresponding to the lower bound of the 80%-confidence interval is identified and compared with the theoretical lower bound $\bar{k}$ (2). The result is shown in Figure 2 between brackets, for different p, while $N = 5$ remains fixed. It confirms the validity of (2).

The EMTT T (7) is validated in similar way. In each of the 10^6 runs, 1000 steps are taken with transition probability
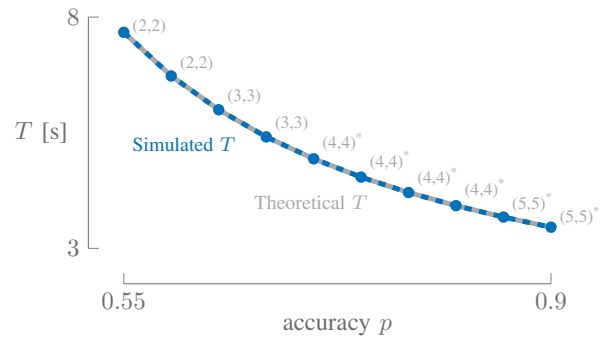

Fig. 2: The simulated EMTTs match the theoretical EMTTs. For every evaluated p, the theoretical and simulated lower bound $\bar{k}$ are shown as well $(\bar{k}_{\text{theo}}, \bar{k}_{\text{sim}})$. An asterisk shows if $\bar{x} \geq c = 0.65$.

p. After 1000 steps, the target direction in the Markov chain (Figure 1) is reversed and the number of steps s is registered before state $k = \lceil c(N-1)+1 \rceil$ is hit for the first time. $k = 4$, given that we choose $c = 0.65$ and $N = 5$. The conditioning on $i < k$, i.e., only switching from outside the predefined working region, is taken into account by removing all simulations in which the process arrived within the predefined working region $[c, 1]$ after 1000 steps. Each simulation results in one sample $s|i \rightarrow k, i < k$ of (6). The resulting hitting times are averaged over all simulations to obtain a sample of the EMTT T (7), where $\tau = 1\,\text{s}$. This experiment is performed for different p, while $N = 5, k = 4$ and $\tau = 1\,\text{s}$. Figure 2 shows that the simulations correspond to the theoretical formula of T (7). The relative error is for every p smaller than 10^{-3}.

B. Experiment on real EEG and audio data

To show the computation of the ESD metric, we apply it on the results of the minimal mean-squared error (MMSE)-based AAD algorithm [2], [9] on real EEG and audio data. A trained linear spatio-temporal decoder can be applied to new EEG data, in order to predict the attended speech envelope. The resulting speech envelope is then correlated with the recorded speech envelopes presented to the subject within a decision window length of length τ. The recorded speech envelope that correlates best with the predicted envelope, is identified as the attended speaker. Note that using a larger decision window length τ results in more accurate estimates of the correlation coefficients, thereby improving the AAD performance.

The data originates from an AAD experiment in which 16 subjects listened to Dutch short stories of 24 minutes in a competing two-speaker situation [9]. Each subject was instructed to target attention on one of the stories presented to one ear. Details about the experiment, data and preprocessing can be found in [9]. It is noted that for the training and testing of the AAD algorithm, we use the envelopes of the original speech signals. In practice, these envelopes have to be extracted from the hearing aid's microphones [7].

The recorded EEG and audio data of 72 minutes are split into trials of $60\,\text{s}$ and per-subject decoders are trained and
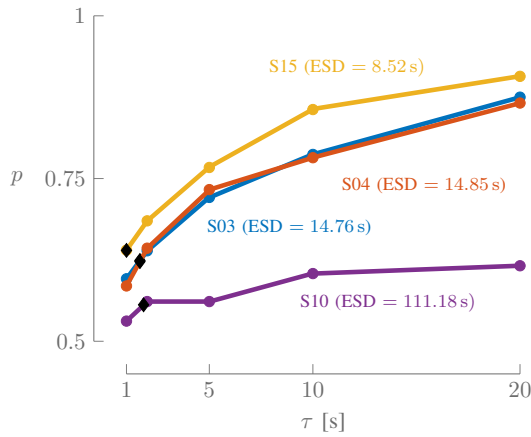
Fig. 3: The ESD summarizes the shown $p(\tau)$-performance curves for different subjects in a single performance metric. The optimal working points are indicated with a diamond (♦).

tested in a leave-one-trail-out fashion. A decoder is trained on 71 minutes of data and tested on the left-out trial. This trial is further split into sub-trials of smaller lengths to evaluate the trained decoder on smaller decision window lengths τ. The AAD decisions are registered and compared with ground-truth attention labels of the test trials, resulting in an average accuracy over the complete recorded dataset per subject, for each decision window length.

Figure 3 shows the constructed performance curves based on the accuracies p evaluated at window lengths $\tau = \{1, 2, 5, 10, 20\}$s of the decoders of four subjects. The confidence level P_0 is chosen equal to 0.8, while comfort level c is chosen equal to 0.65, based on a subjective listening test (details omitted).

The ESDs reported in Figure 3 capture the global behavior of the corresponding performance curves. The accuracies are very similar for subject 3 (S03) and S04, which is reflected in the ESD. Note that the optimal working point corresponds to decision window lengths around $\tau = 1 - 2$s. These small decision window lengths are also more relevant in the hearing aid use case. S15 has higher accuracies for all evaluated window lengths. This is reflected in the large decrease of $\approx 40\%$ in EMTT from 14.76s (of the second-best decoder of S03) to 8.52s for S15. For all three subjects, the optimal Markov chain contains 7 states and has approximately the same transition probability of $p \approx 63\%$. This working point is however reached for a smaller decision window length for S15, resulting in the lower ESD.

S10 is an outlier: the MMSE decoder globally performs much worse than the other subject-dependent decoders. This is reflected in an increase with a factor ≈ 7.5 with respect to the second-worst decoder of S04. As the accuracies are overall very low, the ESD metric focuses more on minimizing the window length and thus the decision time to bring down the EMTT, rather than picking a working point at a higher accuracy, resulting in fewer states. Selecting the working point at a smaller accuracy results in a Markov chain with 19 states,

which is much more than the other decoders.

It is clear that the ESD metric succeeds in capturing the global performance of different decoders based on a relevant and interpretable criterion. It is observed to focus more on smaller window lengths, which has more practical relevance as well and allows comparison of decoders where other methods would fail. Although the ESD values appear high for practical hearing aid use cases, it should not be confused with the time it takes to *perceive* the shift towards the attended speaker. This generally happens earlier than indicated by the ESD, which corresponds to the time needed to realize a *stable* switch.

IV. Conclusion

We proposed a new metric to evaluate the performance of AAD algorithms: the expected switch duration. It represents the expected time needed to switch the operation of the hearing aid from one speaker to another when it is not yet operating in the desired stable working region. First, the number of states of the Markov chain is for each window length τ optimized to obey certain design requirements based on the P_0-confidence interval. The ESD is then defined as the optimal transit time from any state outside the working region to this predefined working region, over all window lengths τ on an interpolated $p(\tau)$-performance curve.

The simulations validated the formulas behind the ESD metric and experiments have shown that it is an *interpretable*, *single-number* metric that *combines both accuracy and decision time*. Based on the ESD, it is easy to compare different AAD algorithms or decoders based on a relevant criterion, *independent of the evaluated window lengths*.

References

[1] E. C. Cherry, "Some Experiments on the Recognition of Speech, with One and with Two Ears," *The Journal of the Acoustical Society of America*, vol. 25, no. 5, pp. 975–979, 1953.

[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*, vol. 25, no. 7, pp. 1697–1706, 2014.

[3] 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.

[4] A. de Cheveigné, D. D. Wong, G. M. Di Liberto, J. Hjortkjær, M. Slaney, and E. Lalor, "Decoding the auditory brain with canonical component analysis," *NeuroImage*, vol. 172, pp. 206–216, 2018.

[5] 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, p. 262, 2018.

[6] 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.

[7] 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 Trans. on Biomedical Engineering*, vol. 64, no. 5, pp. 1045–1056, 2017.

[8] P. Brémaud, *Markov chains: Gibbs fields, Monte Carlo Simulation, and Queues*, ser. Texts in Applied Mathematics. New York: Springer Science & Business Media, 2013, vol. 31.

[9] 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 Trans. on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 5, pp. 402–412, 2017.