



LUND UNIVERSITY

Real-time Bayesian control of reactive brain computer interfaces

Tufvesson, Pex; Gemborn Nilsson, Martin; Soltesz, Kristian; Bernhardsson, Bo

Published in:
IFAC Proceedings Volumes (IFAC-PapersOnline)

DOI:
[10.1016/j.ifacol.2023.10.1612](https://doi.org/10.1016/j.ifacol.2023.10.1612)

2023

Document Version:
Peer reviewed version (aka post-print)

[Link to publication](#)

Citation for published version (APA):
Tufvesson, P., Gemborn Nilsson, M., Soltesz, K., & Bernhardsson, B. (2023). Real-time Bayesian control of reactive brain computer interfaces. In H. Ishii, Y. Ebihara, J. Imura, & M. Yamakita (Eds.), *IFAC Proceedings Volumes (IFAC-PapersOnline)* (2 ed., Vol. 56, pp. 470-475). Elsevier B.V..
<https://doi.org/10.1016/j.ifacol.2023.10.1612>

Total number of authors:
4

General rights

Unless other specific re-use rights are stated the following general rights apply:
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Real-time Bayesian Control of Reactive Brain Computer Interfaces

Pex Tufvesson* Martin Gemborn Nilsson** Kristian Soltesz**
Bo Bernhardsson**

* *Department of Automatic Control, Lund, Sweden. Ericsson Research, Lund, Sweden. (e-mail: pex.tufvesson@control.lth.se).*

** *Department of Automatic Control, Lund, Sweden (e-mail: {martin.gemborn_nilsson,kristian.soltesz,bob}@control.lth.se).*

Abstract: This paper introduces an improved method for real-time brain computer interface control. We demonstrate how Bayesian optimization and feedback can be used to achieve faster statistical convergence by controlling the sequence of stimuli shown in a brain computer interface based on a visual oddball paradigm.

Keywords: Control in neuroscience, Biomedical system modeling, simulation and visualization, Bayesian methods, Parameter and state estimation, Input and excitation design

1. INTRODUCTION

Brain computer interfaces (BCIs) are devices that enable direct human-to-machine communication without using regular pathways such as peripheral nerves or muscles, Wolpaw et al. (2002). A distinction can be made between active and reactive BCIs. With an active BCI, the user is intentionally encoding mental states used as instructions for the system to act upon, for example by thinking *left, stop, forward*. This classification requires a long individual calibration time to obtain high accuracy.

In contrast, a reactive BCI analyzes the subject’s brain response to external stimuli. The typical example is to attempt classification of a stimulus based on the recorded response. In its simplest form, possible responses are partitioned into one target category, and one non-target category. This could be a user who is actively focusing on an infrequently shown target category (oddball paradigm) among a series of stimuli. Taking images from Fig. 1 as an example, displaying one person at a time, green persons could be the target category, while all non-green persons would be the non-target category.

One modality used to capture brain signals in wearable BCIs is the electroencephalogram (EEG), non-invasively collected through a number of electrodes placed along the scalp. When a subject is presented a sequence of stimuli and the subject is focusing on the target stimulus, for example by counting the number of occurrences, the recorded EEG will differ depending on whether a stimulus is target or non-target. To simplify the discussion, we will assume that some suitable algorithm transforms each stimulus response EEG time series to a real number, suitable for separating target from non-target responses. Representative examples of target and non-target distributions, estimated from real data using four different methods (see sections 2.6 and 3.1), are shown in Fig. 2.

* All authors are members of the ELLIIT Excellence Center at Lund University.



Fig. 1. Some of the visual color stimuli used in the Clear by Mind brain game from the BCI-HIL research framework published by Gemborn Nilsson et al. (2023).

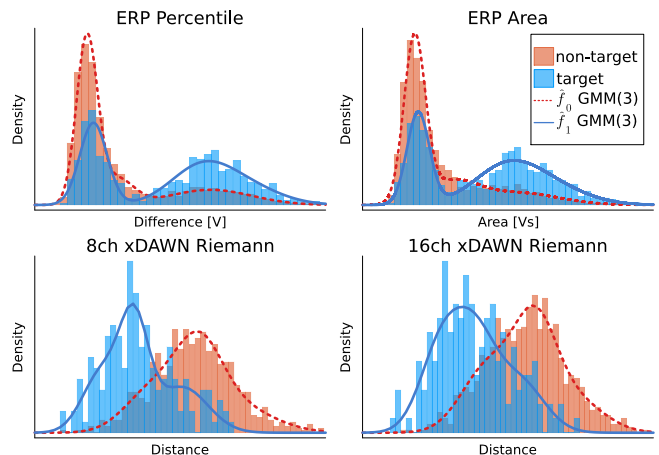


Fig. 2. Data distributions when mapping the multi-channel EEG epoch data to a real-valued output. Gaussian mixture models (GMMs) \hat{f}_0 and \hat{f}_1 are estimated from the dataset described in section 3.1.

In this setting, a relevant problem is for the BCI system to rapidly guess, based on consecutive stimulus-response data, which stimulus category is the target. A few recent studies have proposed and evaluated heuristics for improving reactive BCI performance through feedback techniques. In Ma et al. (2021) an algorithm for adaptive

stimulus selection was developed and the speed of a BCI speller was reported to be increased by 70 percent by use of Thompson sampling, a classical method from the domain of so called multi-armed bandit problems.

In this paper, we continue work in this area and show how feedback and Bayesian optimization of stimulus selection can be applied to further improve classification efficiency.

To illustrate the ideas we choose to apply closed loop control on a reactive BCI based on visual stimuli, as described by Vidal (1973). The setup uses the P300 response, an event-related response potential found about 300 ms after stimuli that induce a "working memory update", described by Chapman and Bragdon (1964). It is one of the most studied evoked response potentials (ERPs) and can be used to separate target from non-target stimuli, see Kappenman and Luck (2011). A subject's evoked response to visual stimuli as seen in Fig. 1 is recorded as a fixed length time series containing multiple EEG channels, together forming a one-second *epoch*.

2. METHOD

2.1 Stochastic stimulus–response model

We denote the sequence of T consecutive stimuli

$$\mathbf{u} = [u_1 \dots u_T]^\top, \quad (1)$$

where each stimulus is coded as an integer, representing one out of C possible, mutually exclusive, categories, i.e.

$$u_t \in \{1, \dots, C\} \triangleq U, \quad t = 1, \dots, T. \quad (2)$$

To make the framework for the stimulus control problem general, we will assume that the evoked measured brain responses are mapped into real numbers denoted

$$\mathbf{y} = [y_1 \dots y_T]^\top. \quad (3)$$

Furthermore, events are for simplicity assumed to be independent in the sense that y_t is only affected by u_τ if $t = \tau$. One of the stimuli categories, $x \in U$, is denoted the target, while the remaining $C - 1$ categories are all non-targets. The aim is to determine this unknown latent variable x from input-output data (\mathbf{u}, \mathbf{y}) . The response y_t is assumed to follow a non-target distribution with Probability Density Function (PDF) f_0 if $u_t \neq x$, and a target distribution PDF f_1 if $u_t = x$. For ease of notation we introduce the compact notation

$$f(y_t|u_t, x) = \begin{cases} f_0(y_t) & \text{if } u_t \neq x, \\ f_1(y_t) & \text{if } u_t = x. \end{cases} \quad (4)$$

The likelihood of observing y_t as a response to u_t is also denoted $\mathcal{L}(y_t|u_t, x) = f(y_t|u_t, x)$. As further discussed in section 5, the PDF f is in this paper assumed to be static and known. In section 2.6, some alternatives for estimating these distributions from real data are explained. The, very relevant, problem of updating estimates of f_0 and f_1 from new data is not considered here.

Next, we express the likelihood of \mathbf{y} conditioned on the underlying stimuli \mathbf{u} and the target category, x , being k :

$$\mathcal{L}_k \triangleq \mathcal{L}(\mathbf{y}|\mathbf{u}, x = k) = \prod_{t=1}^T f(y_t|u_t, k). \quad (5)$$

The maximum likelihood target estimator is thus

$$\hat{x}_{ML} = \operatorname{argmax}_{k \in \{1, \dots, C\}} \mathcal{L}_k. \quad (6)$$

To compute the probability that a candidate k is the target, conditioned on the data, the known distributions, and an a priori assumption of equally likely targets, Bayes' formula gives

$$p_k|\mathbf{u}, \mathbf{y} \triangleq P(x = k|\mathbf{u}, \mathbf{y}) = \frac{\mathcal{L}_k}{\sum_{i=1}^C \mathcal{L}_i}. \quad (7)$$

Assume we have access to $\mathbf{p} = [p_1 \dots p_C]^\top$, choose the next stimulus u , and observe the resulting response y . The probability of k being the target p^k can then be updated as p_k^+ :

$$\begin{aligned} p_k^+(y|\mathbf{u}, \mathbf{p}) &= P(x = k|y, \mathbf{p}) \\ &= \frac{f(y|\mathbf{u}, k)\mathcal{L}_k}{\sum_{i=1}^C f(y|\mathbf{u}, i)\mathcal{L}_i} = \frac{f(y|\mathbf{u}, k)p_k}{\sum_{i=1}^C f(y|\mathbf{u}, i)p_i}. \end{aligned} \quad (8)$$

The vector \mathbf{p} of probabilities hence constitutes a sufficient statistic for updating the target probability distribution amongst the candidates after a new stimulus-response pair.

We are now faced with an experiment design choice: How should we select the next stimulus u ?

2.2 Naive candidates for stimuli selection

Naive candidates for this experiment design step are:

- *Round robin*: Given the previous input u^- , the next input is chosen as $u = \operatorname{mod}(u^-, C) + 1$.
- *Uniform random*: The next input u is drawn from a discrete uniform distribution over U .
- *Thompson sampling*: Assign u randomly, stratified by belief, so that $P(u = i) = p_i$.
- *Favorite*: Choose the u we currently believe is the most likely target, the one with the highest p .

In section 4 we will compare these naive alternatives to the one we propose in this paper:

- *Target expectation maximization (TEM)*: Choose u to maximize the next true target probability p_x^+ .

A survey of strategies for multi-armed bandit problems can be found in Heskebeck et al. (2022).

2.3 Analysing a simple but non-trivial case

Given $C \geq 3$ (the other cases being trivial) categories and the possibility to show only $T = 2$ consecutive stimuli u_1 and u_2 , what is the optimal decision strategy for maximizing the expectation of the true target probability p_x based on stimuli u_1, u_2 , responses y_1, y_2 , and knowing the distributions f_0, f_1 ? We will assume that we are allowed to use information conveyed by the first stimulus-response pair u_1, y_1 to decide the second stimulus u_2 .

Using (5) and (7), while assuming a uniform initial probability vector \mathbf{p} , the true target probability is

$$\begin{aligned}
p_x|(\mathbf{u}_1, \mathbf{u}_2, y_1, y_2) &= \frac{\mathcal{L}_x(\mathbf{y}|\mathbf{u}, x)}{\sum_{i=1}^C \mathcal{L}_i(\mathbf{y}|\mathbf{u}, i)} \\
&= \frac{f(y_1|\mathbf{u}_1, x)f(y_2|\mathbf{u}_2, x)}{\sum_{i=1}^C f(y_1|\mathbf{u}_1, i)f(y_2|\mathbf{u}_2, i)}. \quad (9)
\end{aligned}$$

At start we consider it equally likely that the target is either of the C categories, and therefore can choose \mathbf{u}_1 randomly. After having shown \mathbf{u}_1 and observed y_1 we define $w(y_1)$ as

$$w(y_1) \triangleq P(u_1 = x|y_1) = \frac{f_1(y_1)}{(C-1)f_0(y_1) + f_1(y_1)}, \quad (10a)$$

$$P(u_1 \neq x|y_1) = 1 - w(y_1). \quad (10b)$$

The two strategies we can choose between are $u_2 = u_1$ and $u_2 \neq u_1$. For the latter strategy, we do not have information to make an educated choice between the $C-1$ stimuli candidates. This symmetry motivates us to treat $u_2 \neq u_1$ as one case, rather than $C-1$ statistically identical cases.

Since \mathbf{u}_1 carries no information, the only decision support we have for choosing u_2 is the observed y_1 , and our complete knowledge of f_0 and f_1 .

Let us begin with investigating the strategy $u_2 = u_1$. With probability $w(y_1)$ we have $x = u_1$, resulting in

$$\mathcal{L}_x|(u_1 = u_2 = x) = f_1(y_1)f_1(y_2). \quad (11a)$$

With probability $1-w(y_1)$ we instead have $u_1 \neq x$, in which case the $u_1 = u_2$ strategy gives

$$\mathcal{L}_x|(u_1 \neq x, u_2 = u_1) = f_0(y_1)f_0(y_2). \quad (11b)$$

With $B = C-1$ to shorten equations, we have that

$$\sum_{i=1}^C \mathcal{L}_i|(u_2 = u_1) = Bf_0(y_1)f_0(y_2) + f_1(y_1)f_1(y_2). \quad (12)$$

Hence the expectation of the correct target probability, as a function of y_1 , conditioned on $u_1 = u_2$ is

$$\begin{aligned}
\mathbb{E} p_x(y_1)|(u_1 = u_2) &= \\
&w(y_1) \int \frac{f_1(y_1)f_1(y_2)}{Bf_0(y_1)f_0(y_2) + f_1(y_1)f_1(y_2)} f_1(y_2) dy_2 \\
&+ (1-w(y_1)) \int \frac{f_0(y_1)f_0(y_2)}{Bf_0(y_1)f_0(y_2) + f_1(y_1)f_1(y_2)} f_0(y_2) dy_2 \\
&= \frac{1}{Bf_0(y_1) + f_1(y_1)} \int \frac{f_1^2(y_1)f_1^2(y_2) + Bf_0^2(y_1)f_0^2(y_2)}{Bf_0(y_1)f_0(y_2) + f_1(y_1)f_1(y_2)} dy_2. \quad (13)
\end{aligned}$$

For the case where $u_2 \neq u_1$ we instead have three possible cases. First, with probability $w(y_1)$ we have $u_1 = x$, in which case the $u_1 \neq u_2$ strategy gives

$$\mathcal{L}_x|(u_1 = x, u_2 \neq u_1) = f_1(y_1)f_0(y_2). \quad (14)$$

Second, with probability

$$w_2(y_1) \triangleq \frac{f_0(y_1)}{(C-1)f_0(y_1) + f_1(y_1)} = \frac{1-w(y_1)}{C-1}, \quad (15)$$

we have $u_1 \neq x$ and $u_2 = x$, giving

$$\mathcal{L}_x|(u_1 \neq x, u_2 = x) = f_0(y_1)f_1(y_2). \quad (16)$$

Finally, with probability $(C-2)w_2(y_1) = (C-2)(1-w(y_1))/(C-1)$ we have $u_1 \neq x$ and $u_2 \neq x$, giving

$$\mathcal{L}_x|(u_1 \neq x, u_2 \neq u_1) = f_0(y_1)f_0(y_2). \quad (17)$$

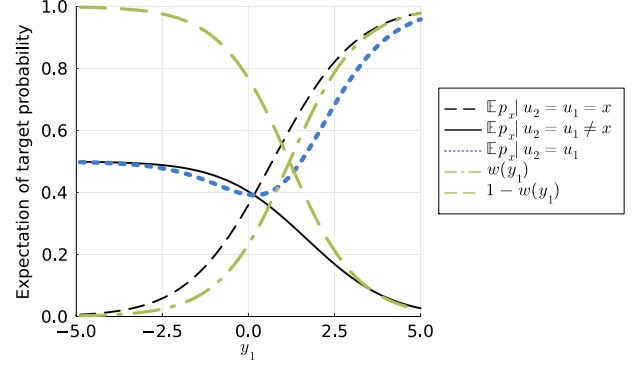


Fig. 3. Expectation of target probability $\mathbb{E} p_x$ when showing the same stimuli $u_2 = u_1$, as a function of y_1 .

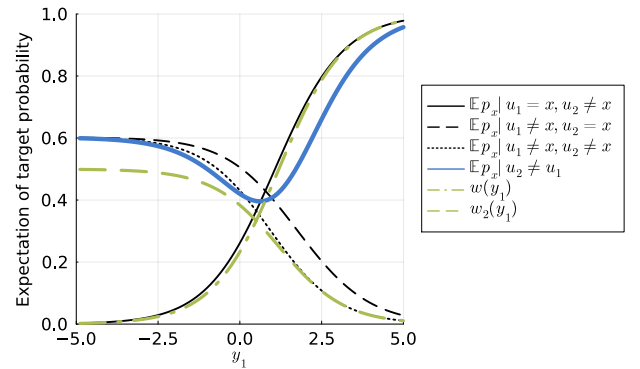


Fig. 4. Expectation of target probability $\mathbb{E} p_x$ when showing different stimuli $u_2 \neq u_1$, as a function of y_1 .

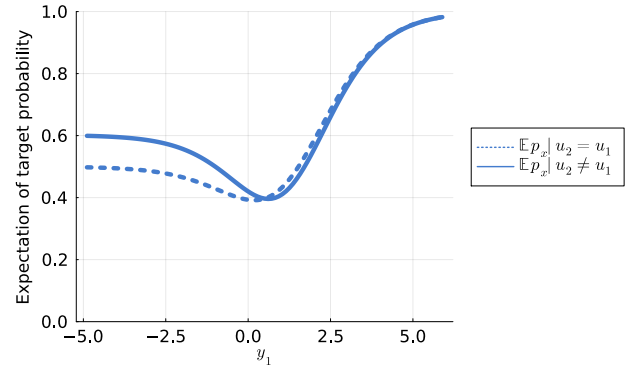


Fig. 5. Summary of Fig 3 and 4 showing the TEM algorithm decision boundary: The optimal choice when $y_1 < 0.5$ is to choose the next stimuli as $u_2 \neq u_1$. When $y_1 \geq 0.5$, choose $u_2 = u_1$, since this will maximize the resulting certainty and lead to a faster decision on guessing the target class.

With the notation

$$\begin{aligned}
d(y_1, y_2) &= \sum_{i=1}^C \mathcal{L}_i|(u_2 \neq u_1) \\
&= (C-2)f_0(y_1)f_0(y_2) + f_0(y_1)f_1(y_2) + f_1(y_1)f_0(y_2),
\end{aligned}$$

we get the corresponding expected value of p_x , marginalized over y_2 , as function of y_1

$$\begin{aligned}
\mathbb{E} p_x(y_1) | (u_1 \neq u_2) = & \\
& w(y_1) \int \frac{f_1(y_1)f_0(y_2)}{d(y_1, y_2)} f_0(y_2) dy_2 \\
& + w_2(y_1) \int \frac{f_0(y_1)f_1(y_2)}{d(y_1, y_2)} f_1(y_2) dy_2 \\
& + (C - 2)w_2(y_1) \int \frac{f_0(y_1)f_0(y_2)}{d(y_1, y_2)} f_0(y_2) dy_2.
\end{aligned} \tag{18}$$

Fig. 3, Fig. 4 and Fig. 5 show results for the case where non-target responses are distributed as $\mathcal{N}(0, 1)$ and target responses as $\mathcal{N}(1, 1)$ so

$$f_0(y) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}y^2}, \quad f_1(y) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(y-1)^2}. \tag{19}$$

2.4 Target expectation maximization (TEM)

When we are to decide on which of the possible stimulus candidates u to present, we do of course not yet have access to the resulting response y . All we know is that y will follow the target distribution with probability p_u and the non-target distribution with probability $1 - p_u$. Presenting the stimulus u , the likelihood of the resulting response y is

$$\mathcal{L}(y|u) = p_u f(y|u, x = u) + (1 - p_u) f(y|u, x \neq u). \tag{20}$$

Since $\mathcal{L}(y|u)$ integrates to unity over y , we can simply say that y will follow a distribution with PDF

$$g(y|u) = \mathcal{L}(y|u). \tag{21}$$

If we treat p_k^* of (8) as an ordinary function of the stochastic variable y , we can thus compute the expectation

$$p_k^*(u) \triangleq \mathbb{E} p_k^*(y|u) = \int p_k^*(y|u) g(y|u) dy, \tag{22}$$

and treat it as an ordinary function of u .

One candidate for this choice is to maximize the expectation of the updated true target probability p_x . Since we (obviously) do not know which candidate is the target, we will rely on our prior belief \mathbf{p} and choose

$$u^* = \operatorname{argmax}_{u \in \{1, \dots, C\}} \sum_{k=1}^C p_k p_k^*(u). \tag{23}$$

2.5 One Step Optimization

An alternative criterion for choice of stimulus u is to maximize the success probability for maximum likelihood to produce the correct choice after the next experiment result (u, y) has been obtained. This is the optimal strategy if it is decided that only one more stimulus is to be shown before a classification decision must be taken.

Notice that, after showing a stimuli $u = i$, the pairwise quotient of all other p_j where $j \neq i$ will remain constant. Thus, with categories labeled such that $p_1 \geq p_2 \geq \dots \geq p_C$, after showing $u = 1$ the only two candidates for \hat{x}_{ML} will become 1 and 2. Similarly, when showing $u = i \neq 1$, the only two candidates for \hat{x}_{ML} will become i and 1. When displaying $u = i$ we get a decision boundary (as a function of y) between the candidates i, j , given by

$$\hat{x}_{ML} = \begin{cases} i & \text{if } p_i f_1(y) \geq p_j f_0(y) \\ j & \text{otherwise.} \end{cases} \tag{24}$$

To simplify further analysis, we assume that (19) holds, and consider each of the two cases one at a time.

Case 1: When category $u = 1$ is used, (6) gives

$$\hat{x}_{ML} = \begin{cases} 1 & \text{if } y \geq r\left(\frac{p_1}{p_2}\right) \\ 2 & \text{otherwise,} \end{cases}$$

where the decision threshold r is given by $r(\alpha) = \frac{1}{2} - \log(\alpha)$. The probability of a successful ML decision, $P(\hat{x}_{ML} = x)$, is in this case (where Φ denotes the CDF of $\mathcal{N}(0, 1)$)

$$P_1 := \left(1 - \Phi\left(r\left(\frac{p_1}{p_2}\right) - 1\right)\right) p_1 + \Phi\left(r\left(\frac{p_1}{p_2}\right)\right) p_2.$$

Case 2: If $u > 1$ is used, then instead we have

$$\hat{x}_{ML} = \begin{cases} u & \text{if } y \geq r\left(\frac{p_u}{p_1}\right) \\ 1 & \text{otherwise.} \end{cases}$$

The probability of a successful ML decision is in this case

$$P_u := \left(1 - \Phi\left(r\left(\frac{p_u}{p_1}\right) - 1\right)\right) p_u + \Phi\left(r\left(\frac{p_u}{p_1}\right)\right) p_1.$$

Using that $1 - \Phi(y) = \Phi(-y)$ and $r(\alpha) = 1 - r(\alpha^{-1})$ it is easy to see that $P_1 = P_2$. Further analysis also shows that $P_1 \geq P_u$ for all u . Any of the two top candidates corresponding to p_1 and p_2 are hence optimal for one step optimization.

2.6 From EEG epochs to Gaussian mixture models

The stimulus control algorithm presented assumes the existence of some algorithm that reduces the dimension of the evoked EEG response into one single number $y \in \mathbb{R}$, to be used for target classification. This can be done in several ways. We have investigated some alternatives based on *Riemannian distance* between covariance matrices, averaging of *percentile amplitudes*, and *absolute area under the curve*, all with the goal of obtaining good separation between distributions of y 's for the target and non-target classes. Due to the inherently noisy nature of EEG signals, the separation will not be perfect, and the probability density function of non-targets f_0 and targets f_1 may overlap.

When estimating f_0 and f_1 using *Riemannian geometry*, we first reduce the number of channels by applying a spatial filter, *xDAWN*, introduced by Rivet et al. (2009). Then, augmented *ERP-covariance matrices* are constructed, as introduced by Congedo et al. (2013). Finally, using the Riemannian metric for symmetric positive definite matrices presented by Moakher (2005), for target-class data, the geometric mean covariance matrix can be computed iteratively as in Fletcher et al. (2004). Using the Riemannian metric, distances from ERP-covariance matrices for each epoch to the mean target ERP-covariance matrix is computed. The processing steps described above was done using either all 16 available EEG channels (FP1, FP2, F5, AFz, F6, T7, Cz, T8, P7, P3, Pz, P4, P8, O1, Oz and O2), or an eight channel subset (AFz, Cz, Pz, Oz, P7, P3, P4 and P8) resembling the set used in Hoffmann et al. (2008).

Notice that xDAWN, construction of ERP-covariance matrices, and the computation of the mean-covariance matrices for the target class, all need labeled training data to fit parameters. In our case we used the first six sessions out of eight as training dataset. Accordingly, the data used

to estimate the target and non-target PDFs of y are the epochs from the two remaining sessions.

Using *percentile amplitudes*, the mapping from one EEG-epoch to a scalar value is done in two steps. First, for each epoch and channel, the difference between the 95th and 5th percentiles for the distribution of samples are computed. Second, for each epoch, these differences are averaged across all EEG channels, resulting in a scalar, y , representing each epoch. This was done for all eight available sessions for one subject with 480 epochs per session, 3840 in total.

Using *absolute area under the curve*, for each epoch and channel, the sum of the absolute value of the time-series is computed. Then, for each epoch, these sums are averaged across all EEG channels, resulting in a scalar, y representing each epoch. This was done for all eight available sessions for one subject with 480 epochs per session, 3840 in total.

For each method listed above, the expectation maximization (EM) algorithm was used to fit a Gaussian mixture model (GMM) with three components for the target and non-target distributions respectively. Histograms and fitted GMMs for each method can be seen in Fig. 2.

3. SIMULATIONS

3.1 Dataset and analysis

The dataset used for empirical estimation of the target and non-target distributions (see Figure 2) was the "Brain Invaders 2013" dataset developed at the GIPSA-lab by Congedo et al. (2011). The data was accessed with the *moabb* Python package by Jayaram and Barachant (2018). For the estimation of probability distributions we used data from eight sessions recorded on a different days with the first subject. Each session consists of 480 trials (80 target, 400 non-target), resulting in a total of 3840 epochs (640 target, 3200 non-target). Each epoch contains one second of EEG data starting from the stimuli onset.

The probability distributions, f_0 and f_1 , used for simulations was estimated using the method based on xDAWN, ERP-covariance matrices and the Riemannian metric, here applied to eight EEG channels. The resulting histograms are shown in the lower left plot of Fig. 2 leading to f_0 and f_1 as seen in Fig. 6. Using these, we now simulate different stimuli selection algorithms to understand their statistical convergence properties.

3.2 Runtime and download

The computational complexity of the TEM algorithm is low, making real-time decisions on stimuli selection possible with embedded processors. We conducted 2048 simulations for each algorithm, using different random seeds and thus different samples from the PDFs. The code is written in Julia, takes less than a minute to run, and can be downloaded from bci.lu.se/bayesian

4. RESULTS

Depending on how the next stimuli is chosen, the probability for correctly classifying the target category differs. Al-

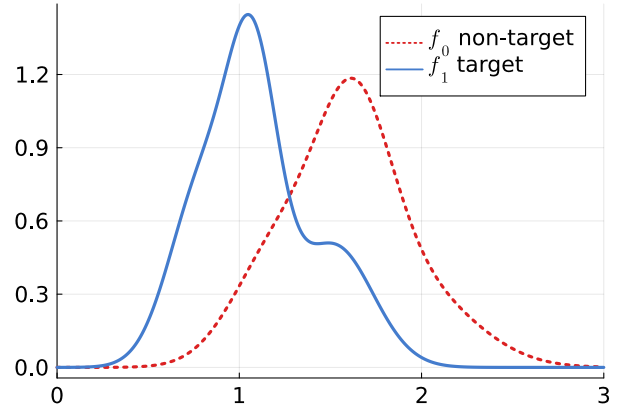


Fig. 6. Probability density functions $f_0(y)$ and $f_1(y)$ for non-target respective target responses to a visual stimuli. A Riemannian distance method and an eight channel xDAWN spatial filter has been used to produce the real-valued coding y of the EEG response.

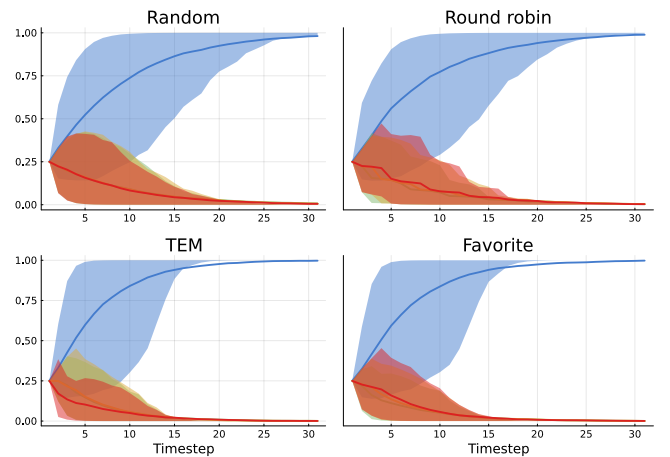


Fig. 7. Mean probability of correct target classification as a function of number of timesteps as a solid line for different stimuli selection algorithms. The target is represented by the upper blue area, which shows the 10th to 90th percentile of its distribution. The non-targets are represented by the overlapping red, yellow and green areas.

gorithms that don't know the responses gotten so far generally perform the worst, as when making random or round robin choices. The algorithms that perform the best adapts to the responses gotten so far, like repeatedly choosing to show the favorite category that we currently assess having the highest probability, or the target expectation maximization algorithm described in this paper. Fig. 7 shows how the Bayesian probability evolves for the target category using 4 different stimuli selection algorithms.

The upper limit for any algorithm is the Oracle, which already knows the correct category and cheats by choosing that one all the time. We can compare and rate the performance of algorithms by their statistical convergence, as can be seen in Fig. ??.

Another way of presenting the performance of algorithms is to register how many stimuli are needed to reach a 95% confidence in the target category. By running

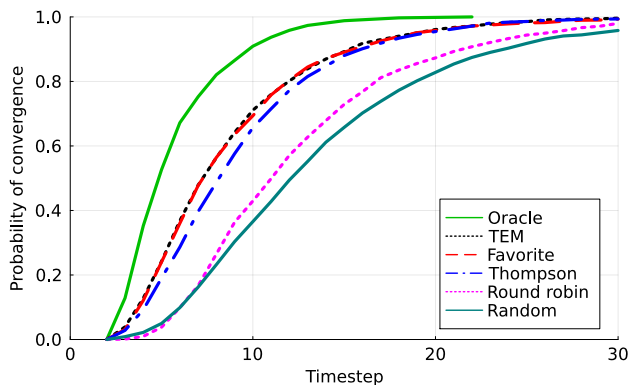


Fig. 8. Probability that a session has reached a correct guess with a certainty of 0.95 as a function of number of timesteps, using six different stimuli selection algorithms. Oracle performs the best while breaking the rules of the game. Random performs the worst. TEM performs slightly better than the Thompson sampling method used in Ma et al. (2021).

many simulations with the same algorithm with responses sampled from the PDFs f_0 and f_1 , we get the probability that a certain algorithm has reached a 95% confidence after a certain number of shown stimuli, as seen in Fig. 8.

5. DISCUSSION

In this paper we have investigated how different mechanisms for choice of presented stimuli affect accuracy within a BCI prediction scenario. In particular, we have used Bayesian analysis for the evolution of the probability vector \mathbf{p} for some different methods and compared performance by simulations. We determined the optimal strategy for the case of a sequence of two consecutive stimuli, $T = 2$. We also proposed a method which can be a good candidate for longer time horizons, where multistep optimization remains computationally intractable.

Probability distribution functions f_0 and f_1 used in our simulations are based on measured responses from EEG captured during an experimental setup with a human subject. The processing steps in the chosen Riemannian method are fitted on data from sessions different than the ones used for estimating the distributions, but still gives good results, indicating robustness towards inter-session variability. Distributions fitted to data from other subjects show similar results, however inter-subject analysis are left for future work.

It is also likely that faster convergence could be achieved by more advanced mappings from EEG epochs to the target and non-target distributions, for example by allowing for vector-valued representations.

An issue worthy of further consideration is better models for the stimuli to response memory, since evoked response potentials, like the P300 signal, are not memory-less. In fact response strength typically increases, the more rare the target stimulus is. This introduces an interesting trade-off, reducing the efficiency of some methods.

The proposed TEM and the Favorite algorithms perform equally well using known static PDFs for one subject. However, when there is little or no previous EEG data

available for a new subject, we can neither assume that the PDFs are known nor static. Transfer learning could initially use known PDFs from other subjects and/or sessions, but would need to update the PDFs while capturing the responses of this new subject.

REFERENCES

- Chapman, R. and Bragdon, H. (1964). Evoked responses to numerical and non-numerical visual stimuli while problem solving. *Nature*, 203(6), 1155–1157.
- Congedo, M., Barachant, A., and Andreev, A. (2013). A new generation of brain-computer interface based on riemannian geometry.
- Congedo, M., Goyat, M., Tarrin, N., Ionescu, G., Varnet, L., Rivet, B., Phlypo, R., Jrad, N., Acquadro, M., and Jutten, C. (2011). "Brain Invaders": a prototype of an open-source P300- based video game working with the OpenViBE platform. In *BCI 2011 - 5th International Brain-Computer Interface Conference*, 280–283.
- Fletcher, P., Lu, C., Pizer, S., and Joshi, S. (2004). Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions on Medical Imaging*, 23(8), 995–1005.
- Gemborn Nilsson, M., Tufvesson, P., and Heskebeck, F. (2023). An open-source human-in-the-loop BCI research framework: Method and design. Images by Kirsty Pargeter/Freepik.
- Heskebeck, F., Bergeling, C., and Bernhardsson, B. (2022). Multi-armed bandits in brain-computer interfaces. *Frontiers in Human Neuroscience*, 16.
- Hoffmann, U., Vesin, J., Ebrahimi, T., and Diserens, K. (2008). An efficient P300-based brain-computer interface for disabled subjects. *Journal of Neuroscience Methods*, 167(1), 115–125.
- Jayaram, V. and Barachant, A. (2018). MOABB: trustworthy algorithm benchmarking for BCIs. *Journal of Neural Engineering*, 15(6), article no. 066011.
- Kappenman, E. and Luck, S. (2011). *The Oxford Handbook of Event-Related Potential Components*. Oxford University Press.
- Ma, T., Huggins, J., and Kang, J. (2021). Adaptive sequence-based stimulus selection in an ERP-based brain-computer interface by Thompson sampling in a multi-armed bandit problem. In *2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 3648–3655.
- Moakher, M. (2005). A differential geometric approach to the geometric mean of symmetric positive-definite matrices. *SIAM J. Matrix Analysis Applications*, 26, 735–747.
- Rivet, B., Souloumiac, A., Attina, V., and Gibert, G. (2009). xDAWN algorithm to enhance evoked potentials: application to brain-computer interface. *IEEE Trans. on bio-medical engineering*, 56(8), 2035–2043.
- Vidal, J. (1973). Toward direct brain-computer communication. *Annual review of biophysics and bioengineering*, 2, 157–80.
- Wolpaw, J., Birbaumer, N., McFarland, D., Pfurtscheller, G., and Vaughan, T. (2002). Brain-computer interfaces for communication and control. *Clinical Neurophysiology*, 113(6), 767–791.