
Making Brain-Computer Interfaces robust, reliable and adaptive with Learning from Label Proportions

Pieter-Jan Kindermans
Machine Learning Group
TU Berlin
p.kindermans@tu-berlin.de

David Hübner
University of Freiburg
Brain State Decoding Lab
david.huebner@blbt.uni-freiburg.de

Thibault Verhoeven
Electronics and Information Systems
Ghent University
thibault.verhoeven@ugent.be

Konstantin Schmid
University of Freiburg
Brain State Decoding Lab
konstantin.schmid@uranus.uni-freiburg.de

Klaus-Robert Müller
TU-Berlin
Korea University
klaus-robot.mueller@tu-berlin.de

Michael Tangermann
University of Freiburg
Brain State Decoding Lab
tangermann@blbt.uni-freiburg.de

Abstract

To deliver optimal decoding performance of the ongoing brain state, a brain-computer interfaces (BCI) traditionally has to be calibrated at the start of every single session. Reasons are individual informative spatio-temporal patterns of brain signals measured by the electroencephalogram, a high noise level in this signal and non-stationarity over time. Adaptive decoding methods may help to tackle these challenges. Specifically, we propose to modify the traditional event-related potential (ERP) paradigm such that the experimental protocol and the decoding approach complement each other, transforming the decoding problem into a setting suitable for learning from label proportions. This results in a BCI where the decoder learns from unlabelled data only, is guaranteed to converge to the corresponding supervised solution but does not require actual calibration. As a proof of concept for the novel approach, we present preliminary data from three subjects, which are successfully using a visual ERP paradigm for spelling with the novel approach.

1 Introduction

One of the use scenarios of BCIs is to provide control directly through brain signals, e.g. to restore communication for patients in locked state [1]. As the informative content of the measured brain signals and even the background noise typically varies between subjects and sessions, and as the signal even within a session shows non-stationary behaviour, decoding the brain's activity from e.g. recordings of the electroencephalogram (EEG) is challenging. Thus in most BCI sessions, the user is required to go through at least a short calibration procedure, during which he/she is instructed to perform specific mental tasks such that labelled training data can be recorded. Unfortunately, this calibration must be repeated over time to achieve optimal performance. To avoid this we need robust and adaptive decoding models. We will address this issue for BCI based on visual event-related potentials (ERP) by modifying the paradigm such that a *learning from label proportions* (LLP) approach can be utilized.

Learning from label proportions can be used for problems where groups of data and the proportional

presence of each class is easily obtained but where getting precisely labelled data is difficult [7]. The key idea behind the mean map LLP algorithm lies in the fact that several loss functions, including the logistic and square loss, can be re-written in a form that only depends on the class means with respect to label information. While there have been several papers published on this topic, we are not aware of many practical applications for which LLP has been used. In most of the prior work on LLP, the evaluation is performed on artificially generated datasets.

The ERP speller is one of the most common types of BCI, and was already introduced in 1988 [2]. In general, an ERP-BCI works as follows. Each possible action of the BCI application is assigned to one of several stimuli presented to the user in a rapid randomly repeated sequence. Each time the stimulus corresponding to the actually *desired* action is presented to the user, a target ERP response will be elicited in the EEG. When a stimulus associated with a different action is presented, a non-target response will be generated. To make it more explicit, let us consider the original ERP spelling application. In this speller, a 6 by 6 grid of letters was shown on the screen. To enable the user to select one out of the 36 symbols, all 6 rows and all 6 columns were highlighted independently in a random sequence. The goal now is to figure out which row and which column generated a target response, as this information is sufficient to determine the desired target symbol. In practice, this highlighting of rows and columns is repeated multiple times to improve the reliability, but this comes at the cost of a reduced spelling speed. The repetitions introduce constraints on the labelling of the individual stimuli. This enabled Kindermans et al. to develop a probabilistic model which can be trained without labelled data using the expectation-maximization algorithm [4, 3]. While it works quite well in practice, there is no guarantee that their model will converge to a good solution.

In this abstract we present a different and novel BCI approach that combines the ideas from LLP and a novel stimulus presentation approach to facilitate unsupervised learning with guarantees. The key contributions are the following. First, we view it as a new conceptual idea. Many approaches have been suggested to enhance the signal to noise ratio in BCI paradigms, but in all of them the machine learning component is included at the end as a simple supervised learning problem. In contrast, we tune the BCI application to maximise the power of the machine learning algorithm. Second, thanks to the symbiosis between application and machine learning, we are able to present the first BCI application that can be used without prior subject-specific calibration that that still is guaranteed to converge to the same solution as the corresponding supervised classifier training would have delivered. Third, we demonstrate that our approach does not only work on paper but is successful in an online BCI experiment.

2 Methods

2.1 Learning from label proportions

In supervised classification we typically optimise a loss function that depends on the data \mathbf{x}_i and the labels $y_i \in \{-1, 1\}$ where $i = 1 \dots N$ indicates the sample. For a specific subset of losses called *symmetric proper scoring losses*, which include the logistic loss and the square loss, we can re-write this in a form that depends on the class means and the input data [6] only, but not on the individual label of every single data point. As an example, the squared error can be re-written as follows:

$$\sum_{i=1}^N (\mathbf{w}^T \mathbf{x}_i - y_i)^2 = \sum_{i=1}^N ((\mathbf{w}^T \mathbf{x}_i)^2 + 1) - 2\mathbf{w}^T \left(\sum_{i_+} \mathbf{x}_{i_+} - \sum_{i_-} \mathbf{x}_{i_-} \right).$$

It is clear that the first term $\sum_{i=1}^N ((\mathbf{w}^T \mathbf{x}_i)^2 + 1)$ does not depend on label information. The second term still contains the class means and thus some form of label information. It can be rewritten as follows:

$$2\mathbf{w}^T \left(\sum_{i_+} \mathbf{x}_{i_+} - \sum_{i_-} \mathbf{x}_{i_-} \right) = 2\mathbf{w}^T (N_+ \hat{\boldsymbol{\mu}}_+ - N_- \hat{\boldsymbol{\mu}}_-),$$

where $\hat{\boldsymbol{\mu}}_+$ is the estimated average feature vector of the positive class and N_+ is the number of data points in the training set for this class. Hence, for the optimisation of the loss function knowledge of the class means suffices and explicit label information does not need to be known.

The idea behind the mean map algorithm lies in the fact that the empirical class means $\hat{\boldsymbol{\mu}}_+, \hat{\boldsymbol{\mu}}_-$, which are used in the objective function, converge to the true class means $\boldsymbol{\mu}_+, \boldsymbol{\mu}_-$. Hence, any

unbiased estimator of the class means will enable us to approximate the loss function such that in the limit it will converge to the supervised loss function. Now consider the two-class case where we have G groups of data where the proportions of each class in a group are given by Π , which in our case is known by experimental design. Then the expected value of the feature vectors in the groups μ_1, \dots, μ_G can be expressed in terms of the class means μ_+, μ_- as

$$\begin{bmatrix} \mu_1 \\ \vdots \\ \mu_G \end{bmatrix} = \Pi \begin{bmatrix} \mu_+ \\ \mu_- \end{bmatrix}, \quad \Pi = \begin{bmatrix} \pi_+^1 & \pi_-^1 \\ \vdots & \vdots \\ \pi_+^G & \pi_-^G \end{bmatrix}.$$

To obtain an empirical estimate of the group means μ_1, \dots, μ_G we do not need label information. These quantities can be computed directly. Hence, by solving the resulting system of linear equations we can obtain an approximation $\tilde{\mu}_+, \tilde{\mu}_-$ of the true class means μ_+, μ_- .

2.2 A modified ERP paradigm

In the original ERP speller with a 6 by 6 grid, each row and each column was highlighted once. Since the user was focussing his attention on a single symbol, only the corresponding row and column can be expected to elicit a target ERP response. We have a $\frac{1}{6}$ target proportion. However, for the learning with label proportions algorithm to work, this is not sufficient since the resulting proportion matrix (actually a vector) would not be invertible. To make the LLP idea applicable to ERP-based BCIs, we need to elicit ERP responses in multiple groups, which differ from each other with respect to their target-to-non-target ratio. This is non-trivial since modulating the target to non-target interval/ratio changes the strength of the ERP components. For this reason, we propose to combine two different stimulus sequences that are mixed on a stimulus by stimulus level. In our experiment, one stimulus sequence contained 2 targets for every 18 stimuli, while the other sequence contained 3 targets for every 8 stimuli, resulting in a 2×2 matrix Π with $\pi_+^1 = 3/8, \pi_-^1 = 5/8$ and $\pi_+^2 = 2/18, \pi_-^2 = 16/18$. In order to make these sequences applicable, we increased the spelling interface to 42 entries (Fig. 1(B)). In addition to letters, space and punctuation symbols, we included 10 hash symbols (#) which served as uninformative blanks. The blanks are included to match the brightness levels of the two sequences, but were never attended by users. Without them, the ERP responses from the different sequences would not necessarily be homogeneous, thus leading to a violation of the assumed model. The actual stimulus sequences were optimised using the algorithm by Verhoeven et al. which aims to maximise the signal to noise ratio of the EEG signal by structuring the sequences intelligently [9]. The actual visual stimulus applied to 12 symbols per event was optimised in a previous study by Tangermann et al. [8] and was expected to lead to improved ERP responses compared to the traditional brightness highlighting.

2.3 Experimental setup




Figure 1: (A) Decoding accuracy (online experiment). (B) Spelling matrix with visual highlighting.

In our proof of concept study, three subjects attempted to spell the sentence "FRANZY JAGT IM KOMPLETT VERWAHRLUSTEN TAXI QUER DURCH FREIBURG." with the BCI three times. Prior to the start of each of those three sentences, the decoding model was reset. We recorded EEG signals from 31 gel-based passive Ag/AgCl electrodes at 1 kHz sampling frequency and subsequently subsampled it to 100 Hz. The electrodes were placed according to the extended 10–20 system and impedances were kept below 20 kOhm. All channels were referenced against the nose. For classification we removed the channels Fp1 and Fp2 to impede the potential use of eye movements to control the BCI, which resulted in 29 channels. We computed the average EEG at 6 intervals for classification, resulting in

$6 \times 29 = 174$ features. Intervals used were [50 120; 120 200; 201 280; 281 380; 381 530; 531 700] milliseconds post stimulus. The stimulus duration was 100 ms while the stimulus onset asynchrony was 250 ms, such that the inter-stimulus-interval was 150 ms. To spell a symbol we presented 6 stimulus sequences, 4 sequences with the 3/8 ratio and 2 with the 2/18 ratio.

The classifier used for decoding target from non-target ERP responses was a least squares classifier where the targets were rescaled according to the global ratio between targets to non-targets such that $X^T \mathbf{y} = \tilde{\boldsymbol{\mu}}_+ - \tilde{\boldsymbol{\mu}}_-$. This classifier was regularised using the analytical ledoit-wolf shrinkage for the covariance matrix [5]. The class means for the optimisation of the loss function were re-estimated and the classifier was re-trained in each trial, after making the prediction.

3 Results and Discussion

Due to space constraints we will focus on the accuracy values obtained by the LLP approach as in Fig. 1(A). A blue square indicates a correctly spelled symbol, a yellow square indicates an incorrectly spelled symbol. The time increases from left to right. During the initial trials, the classifier typically misclassifies the desired symbol, but after 7-10 trials the classifier is able to decode the brain-signals reliably. Averaged over all runs and subjects the character accuracy was 88.36%.

4 Conclusion

In this abstract we have demonstrated how the ERP-based BCI protocol can be modified to suit the learning from label proportions setting. By integrating LLP into the BCI we are able to build the first unsupervised BCI, where the decoder is guaranteed to converge to a supervised solution without actually knowing the user's intention.

Acknowledgments: This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement NO 657679. This work was partially funded by the BrainLinks-BrainTools Cluster of Excellence funded by the German Research Foundation (DFG, grant number EXC 1086).

References

- [1] Guido Dornhege, Jose del R. Millan, Thilo Hinterberger, Dennis J. McFarland, and Klaus-Robert Müller. *Towards Brain-Computer Interfacing*. MIT Press, 2007.
- [2] Lawrence Ashley Farwell and Emanuel Donchin. Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalography and Clinical Neurophysiology*, 70(6):510 – 523, 1988.
- [3] Pieter-Jan Kindermans, Hannes Verschore, David Verstraeten, and Benjamin Schrauwen. A P300 BCI for the masses: Prior information enables instant unsupervised spelling. In *Advances in Neural Information Processing Systems (NIPS)*, pages 719–727. 2012.
- [4] Pieter-Jan Kindermans, David Verstraeten, and Benjamin Schrauwen. A Bayesian model for exploiting application constraints to enable unsupervised training of a P300-based BCI. *PLoS ONE*, 7(4):e33758, 04 2012.
- [5] Olivier Ledoit and Michael Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88:365–411, 2004.
- [6] Giorgio Patrini, Richard Nock, Tiberio Caetano, and Paul Rivera. (almost) no label no cry. In *Advances in Neural Information Processing Systems*, pages 190–198, 2014.
- [7] Novi Quadrianto, Alex J Smola, Tiberio S Caetano, and Quoc V Le. Estimating labels from label proportions. *Journal of Machine Learning Research*, 10(Oct):2349–2374, 2009.
- [8] Michael Tangermann, Martijn Schreuder, Sven Dähne, Johannes Höhne, Sebastian Regler, Andrew Ramsay, Melissa Quek, John Williamson, and Roderick Murray-Smith. Optimized stimulation events for a visual ERP BCI. *International Journal of Bioelectromagnetism*, 13(3):119–120, 2011.
- [9] Thibault Verhoeven, Pieter Buteneers, JR Wiersema, Joni Dambre, and PJ Kindermans. Towards a symbiotic brain–computer interface: exploring the application–decoder interaction. *Journal of Neural Engineering*, 12(6):066027, 2015.