Probabilistic classification model for Brain Computer Interfaces

J. Mattout, G. Gibert, V. Attina, E. Maby & O. Bertrand
Brain Dynamics and Cognition, U821 INSERM, Lyon, France ; Lyon 1 – Université Claude Bernard, Lyon, France

www.irisa.fr/bunraku/OpenViBE

INTRODUCTION

To control an external device from brain signals directly, non-invasive Brain Computer Interfaces (BCI) rely on brain activity as measured with EEG, MEG or fMRI. To infer a command, any BCI requires to proceed with the following steps online (see figure 1): data preprocessing, extraction of the relevant signal features and classification into a decision.

A crucial aim in BCI is to optimize the information transfer rate. One has to find the best compromise between robustness and speed in the decision process. The main challenge is to accomodate the very few amount of data available online and the uncertainty due to noise and variability in individual responses. We adopt a probabilistic approach to account for both uncertainty and all our knowledge about the data generative process. This approach also yields an estimate of our confidence about each online decision.

We describe a probabilistic framework for offline learning and online inference. A first distinction in this type of models is between offline and online parameters. The former, refered to as hyperparameters, can be learned from a training set while the latter have to be infered in real-time. Furthermore, online parameters distinguish between locally generated hidden variables ($\mathbf{s}_t$) that condition the data ($\mathbf{y}_t$) and model inputs (θ) that pertain to the unknown user's choice. Both type of online parameters can be estimated using Bayesian Filtering [1].

We illustrate and evaluate this approach in one of the most well-known BCI paradigm: the P300 speller [2] (see figure 1 & 2).

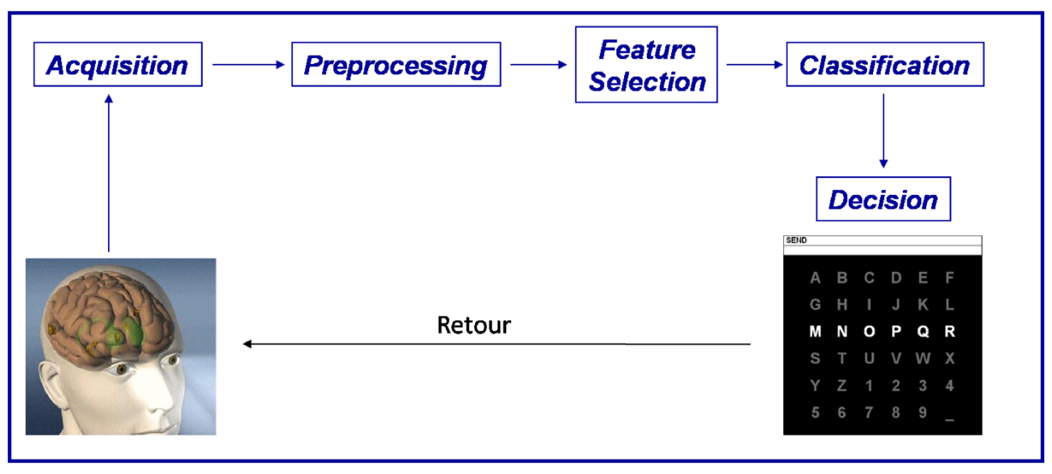

Figure 1: general BCI pipeline, with the P300-speller display [2]

METHODS

Bayesian filtering

We denote the current observation by $\mathbf{y}_t$, while $\mathbf{Y}_{t-1} = \{\mathbf{y}_1 \dots \mathbf{y}_{t-1}\}$ indicates the past observations. In the generative process of $\mathbf{y}_t$ we distinguish between time-variant $\mathbf{s}_t$ and time-invariant θ parameters. $\mathbf{s}_t$ refers to the unobserved causal states of observation $\mathbf{y}_t$, while θ indicates the user's intention common to current $\mathbf{y}_t$ and past observations $\mathbf{Y}_{t-1}$.

To estimate θ at time **t**, Bayesian filtering treats the state variables as missing data and forms the complete data likelihood $\mathbf{p}(\mathbf{y}_t, \mathbf{s}_t | \boldsymbol{\theta}, \mathbf{Y}_{t-1})$. As in Expectation Maximization, integrating out the hidden states yields the marginalized observation model $\mathbf{p}(\mathbf{y}_t | \boldsymbol{\theta}, \mathbf{Y}_{t-1}) = \int \mathbf{p}(\mathbf{y}_t, \mathbf{s}_t | \boldsymbol{\theta}, \mathbf{Y}_{t-1}) \mathbf{d}\mathbf{s}_t$. However in practice, the form of the observation model prevents for exact Bayesian inference. Therefore, we resort to the Variational approximation, assuming that the joint posterior factorizes as $\mathbf{q}(\boldsymbol{\theta}, \mathbf{s}_t | \mathbf{Y}_t) \approx \mathbf{q}(\boldsymbol{\theta} | \mathbf{Y}_t) \mathbf{q}(\mathbf{s}_t | \mathbf{Y}_t)$. This yields an iterative algorithm where the conditional densities on $\mathbf{s}_t$ and θ are alternatively estimated to compute the approximmate marginal likelihood $\mathbf{q}(\mathbf{y}_t | \boldsymbol{\theta}, \mathbf{Y}_{t-1})$ and to infer θ, respectively. Note that the parametric form of the posterior on θ becomes time-independent if it can be chosen as conjugate to the approximate marginal likelihood, since $\mathbf{q}(\boldsymbol{\theta} | \mathbf{Y}_t) \approx \mathbf{q}(\mathbf{y}_t | \boldsymbol{\theta}, \mathbf{Y}_{t-1}) \mathbf{q}(\boldsymbol{\theta} | \mathbf{Y}_{t-1})$.

P300 speller

This paradigm involves a permanently displayed matrix of symbols. In random order, columns and rows of the matrix are flashed (see figure 2). When the symbol to which the subject is paying attention is flashed, an automated P300 EEG wave is generated. Repetitive flashes of the same row/column are needed to identify the P300 response, hence the target symbol, reliably.

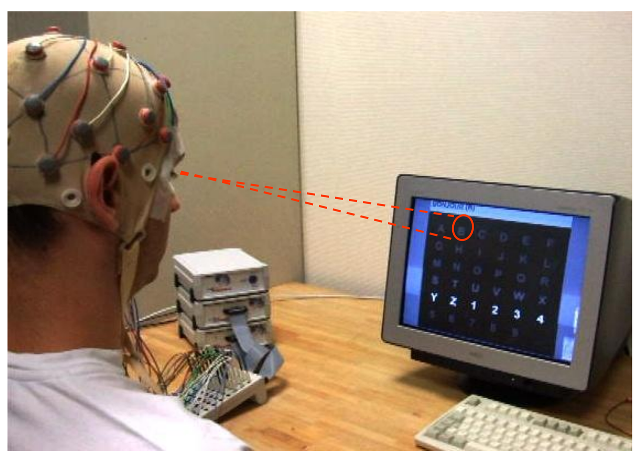

Figure 2: classical P300 speller interface [2]

Obviously, the fewer repetitions we need, the faster the communication.

Generative model

Here is the graphical model for the responses to the flashing of either the last 6 rows or 6 columns of the symbol matrix. The single responses are generated by a mixture of two Gaussian distributions, one for each class of signals ('P3' and 'No').

- Data likelihood: $\mathbf{p}(\mathbf{y}_t | \mathbf{D}_t, \mathbf{X}, \boldsymbol{\Sigma}) = \prod_{m=1}^{2} \prod_{k=1}^{6} [\mathbf{Normal}(\mathbf{vec}(\mathbf{y}_{t,k}); \mathbf{X}_m, \boldsymbol{\Sigma}_m)]^{\mathbf{D}_{t,k,m}}$
- Priors on switches: $\mathbf{p}(\mathbf{D}_t | \mathbf{P}_t) = \prod_{m=1}^{2} \prod_{k=1}^{6} \mathbf{P}_{t,k,m}^{\mathbf{D}_{t,k,m}}$
- Priors on mixing proportions: $\mathbf{p}(\mathbf{P}_{t,k} | \mathbf{R}, \mathbf{a}_k) = \prod_{q=1}^{6} [\mathbf{Beta}(\mathbf{a}_{k,q,m=1}, \mathbf{a}_{k,q,m=2})]^{\mathbf{R}_q}$
- Priors on decision switches: $\mathbf{p}(\mathbf{R} | \mathbf{L}) = \prod_{q=1}^{6} \mathbf{L}_q^{\mathbf{R}_q}$
- Priors on decision proportions: $p(L) = Dirichlet(b)$

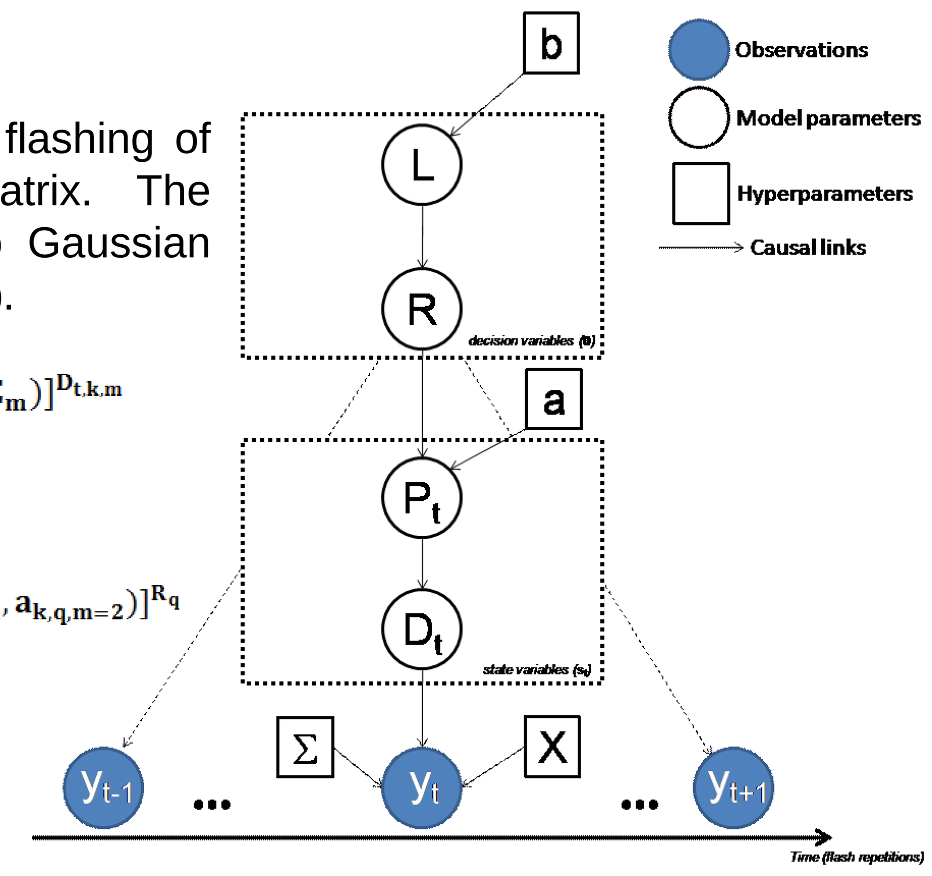

APPLICATION

Offline estimation of model hyperparameters (Learning)

To reduce the dimensionality of the data and to seperate the two classes in an optimal fashion, we used spatial filtering [3]. Spatial filters are learned from the data and given by the main eigenvalues of the generalized eigendecomposition of the Rayleigh quotient (or LDA criterion) for the two classes. Learning operates similarly and independently for the row and column models.

For evaluation, we used data available on the web from the BCI Competition III in 2004 [4]. Training set from subject B yielded the three orthogonal spatial filters shown in figures 3 and 4. From these filters, means ($\mathbf{X}$) and variances ($\boldsymbol{\Sigma}$) of each class features were computed. For simplicity here, parameters **a** of the Beta prior on $\mathbf{P}_t$ were not learned but set as non informative. Finally, parameters **b** were set in order to inform the model that only one row/column among the six belongs to the 'P3' class.

Note that according to the identified spatial features below, not only the P300 response but also the early visual response to the flash seems to help classifying the signals.

Figure 3: Column model - Spatial filter maps (bottom) and corresponding average temporal features (top) for class 'P3' (red curves) and class 'No' (black curves).

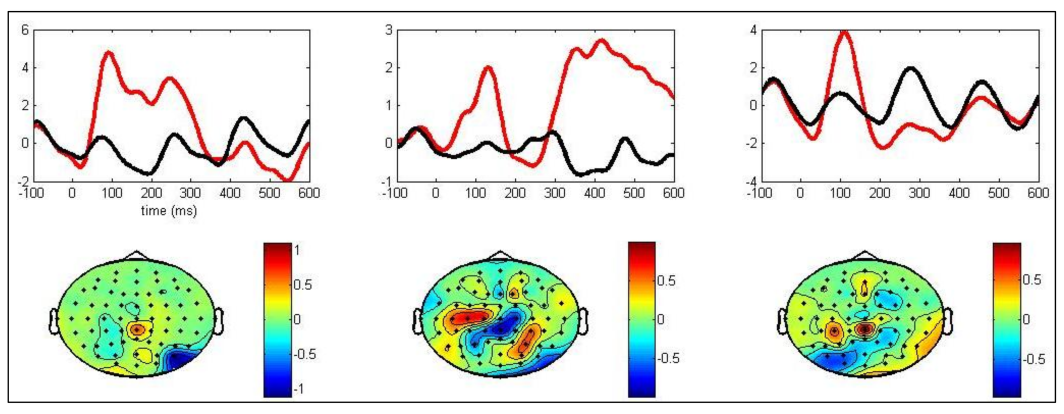

Figure 4: Row model - Spatial filter maps (bottom) and corresponding average temporal features (top) for class 'P3' (red curves) and class 'No' (black curves).

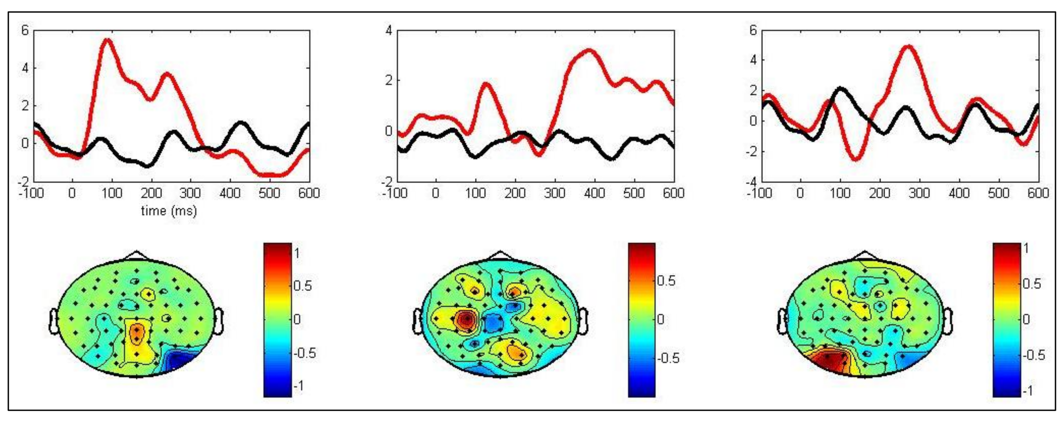

Online estimation of model parameters (Inference)

Simulations

Based on the above features, we used our model to simulate data from 50 randomly selected symbols. For each symbol, 15 trial observations (flash repetitions) were generated. Online estimation was mimicked and decision was made as soon as the estimated probability for one symbol to be the target was greater than 0.15.

All symbols were estimated correctly and Figure 5 shows the histogram of the humber of needed trials before a could be made. Importantly, when more trials were needed, the model proved able to 'change his mind' while accumulating evidence.

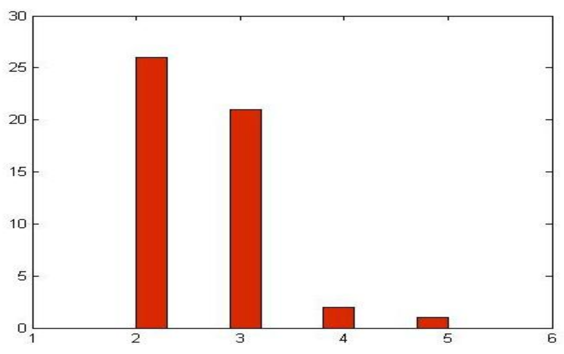

Figure 5: Histogram of the required number of data to take a reliable decision. For some of the symbols, more observations were needed to decode the user's intent.

Real Data

These training and test sets contain 85 and 100 symbols each, with 15 trial observations per symbol. Figures 6 and 7 show the obtained online performance for the training set (validation) and test set (generalization) respectively.

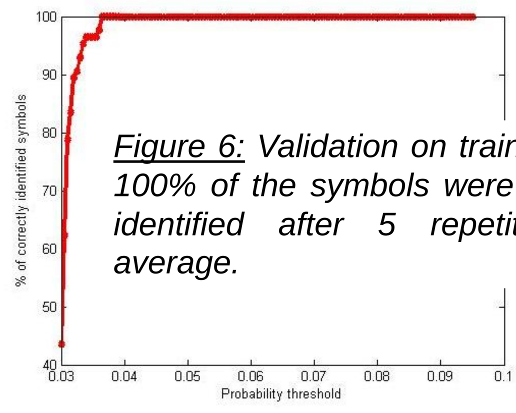

Figure 6: Validation on training set – 100% of the symbols were correctly identified after 5 repetitions on average.

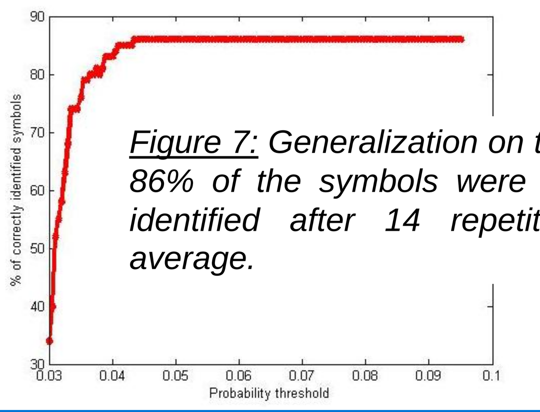

Figure 7: Generalization on test set – 86% of the symbols were correctly identified after 14 repetitions on average.

CONCLUSION

We described a probabilistic framework for offline learning and online inference using Variational Bayes filtering. This general dynamic approach enables to quantify and optimise online decision.

Within this scheme, a first simple model has been derived for the well-known P300-speller BCI. Performance proved excellent in both simulated and real data However, generalization performance could be further improved.

Future directions include:
- The fine learning of parameters **a**;
- The online update of the model hyperparameters;
- The bit rate quantification, in comparison with other approaches;
- The extension and evaluation of this approach in other BCI paradigms.

REFERENCES

[1] Smidl V., Quinn A., *The Variational Bayes Method in Signal Processing*, Springer, 2006.

[2] Farwell L.A., Donchin E., *Taking off the top of your head: toward a mental prothesis utilizing event-related brain potentials,* Electroenceph. Clin. Neurophys., 70:510-23, 1988.

[3] Hoffmann U., Vesin J., Ebrahimi T., *Spatial filters for the classification of event-related potentials*, In European Symposium on Artificial Neural Networks, 2006.

[4] BCI competition III dataset IIb, IEEE Trans. Biomed. Eng., Jun;51(6), 2004. (http://ida.first.fraunhofer.de/projects/bci/competition_iii/)

ACKNOWLEDGMENTS

This work is supported by the OpenVibe project, funded by the French Research National Agency (www.irisa.fr/bunraku/OpenVIBE).