

Methods for Single Trial Analysis of Asynchronous EEG Patterns

THÈSE N° 5398 (2012)

PRÉSENTÉE LE 15 JUIN 2012

À LA FACULTÉ DES SCIENCES ET TECHNIQUES DE L'INGÉNIEUR
CHAIRE FONDATION DEFITECH EN INTERFACE NON-INVASIVE DE CERVEAU-MACHINE
PROGRAMME DOCTORAL EN INFORMATIQUE, COMMUNICATIONS ET INFORMATION

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Nicolas BOURDAUD

acceptée sur proposition du jury:

Prof. K. Aminian, président du jury
Prof. J. D. R. Millán Ruiz, directeur de thèse
Prof. F. Cincotti, rapporteur
Prof. M. Murray, rapporteur
Prof. D. Van De Ville, rapporteur



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2012

C'est le temps que tu as perdu pour ta rose
qui fait ta rose si importante.
— Antoine de Saint-Exupéry

Remerciements

Cette thèse a été un long voyage que je n'aurai pas pu mené à bien sans le soutien, les encouragements et l'aide de nombreuses personnes.

Tout d'abord je souhaiterais remercier chaleureusement mon directeur de thèse José qui m'a accueilli au sein de ses laboratoires à Martigny puis à Lausanne, qui a supervisé mon travail tout au long de ses années. Merci également pour avoir su créer un environnement agréable me permettant de me concentrer sur le travail.

Un grand merci à Ricardo qui a également supervisé mon travail de thèse. Même si son nom n'apparaît pas sur les documents officiels, il a eu un rôle primordial dans mon travail et su me poser les bonnes questions (et bons commentaires!).

Je voudrais également remercier mon jury de thèse : prof. Aminian, prof. Cincotti, prof. Murray et prof. Van de Ville. Merci pour vos questions et pour m'accepter officiellement dans la communauté scientifique!

Un grand merci à mes collègues mais en particulier à Robert, Ganga, Eileen, Semis, Hesam et Tom pour les commentaires sur ma thèse et certains pour me remonter le moral quand il le fallait! Je ne vous oublie pas les autres... travailler (directement ou non) avec vous a été un plaisir!

Maintenant, il faut avouer que je n'en serai pas là sans les encouragements de mes parents en particulier qui ont su me faire rouvrir les yeux sur ce qui comptait! Gros bisous à toute ma famille : vous avez énormément compté pour que je termine.

En plus de la famille un gros bisous à tous mes amis de Paris : Fanny, Raph, Flo, Guillaume, Vince, Touf, Mat, Eric, Waness, Raphi, Aurel, Greg, Jeanne, Kevin, Julia, Anais et Estelle et tant d'autres!! Et Valentine pour le soutien!!! Un gros remerciement à Fanny sans qui je n'aurais pas commencé la thèse. Mais aussi les Suisses : Sarah, Fanny, Tamara, Ferran, Olivier pour la guitare! Merci à tous!!!!

Lausanne, 1er juin 2012

Nicolas.

Abstract

Processing of electroencephalographic (EEG) signals has mostly focused on analysing correlates that are time-locked to an observable event. However, when the signal is acquired in less controlled environment, like in the context of a brain-computer interface operating in the real-world, this synchronous nature does not hold any longer. The analysis of such signal requires the design of methods that rely less on time-locked nature. These methods are also the requirements for study endogenous processes for which the ground truth of when the process take place in time is not available. In this thesis, we present methods to analyse brain signals, EEG in particular, that are not time-locked to observable events. This thesis documents three major contributions : (i) it proposes a Bayesian formalism to the problem of asynchronous EEG pattern classification, (ii) it shows the importance of generative models to achieve this task and (iii) it shows that such methods can be used to gather information and classify the EEG correlates of decision-making process while classical methods fail at it.

First, we propose methods to handle non-time-locked EEG patterns by making the hypothesis that, in each trial, only a part of the signal contains the relevant pattern of interest. This relevant part can appear at any time in the analysis window and differently for each trial. The rest of the trial corresponds to a non-informative part irrelevant to the targeted cognitive task. Starting from a discriminant asynchronous approach handling independently the time-samples in the trial, we extend this method to a generative Bayesian model where each part is formally modelled. This is a main difference compared to the classical approach which usually try to avoid to model the non-informative part. Then, making the assumption that the informative part can be modelled by a time sequence, we adapt the previous method to a Bayesian model of asynchronous template matching which allows the recognition of the time onset of the pattern of interest in each trial.

Second, we show the importance of the generative model which, thanks to the Bayesian approach allows us to alleviate the problem of choosing of the hyperparameters of the initial discriminative approach. Compared to the initial discriminant model, using a generative approach leads to use more parameters into the model but whose estimation is helped by the prior we provide. By doing so, we provide a more intuitive

Abstract

way for the experimenter to adapt the method to other problems. Using a generative model and a Bayesian estimation also enables us to improve the generalisation of the model of asynchronous template matching. This model has indeed been tested as benchmark on jittered evoked potential data and has shown to successfully improve the signal-to-noise ratio, recover the evoked response and classify better than classical methods.

Finally, we see the importance of asynchronous methods for classification of the EEG correlates of decision-making process. We test this in the context of the study of the exploration/exploitation contrast. Exploration is related to decision making in an uncertain environment. This situation arises a conflict between two opposing needs : gathering information about the environment and exploiting this knowledge in order to optimise the decision. Using an experimental setup that forces the subjects to switch between exploratory or exploitative actions, we show for the first time that it is possible to classify the EEG correlates of the exploratory behaviour. Moreover, we show that synchronous methods fail at classify this contrast thus requiring advanced asynchronous ones. The results also confirm that the brain areas relevant to this switch are mainly the left parietal and medial frontal cortex which is consistent with the neurophysiological findings based on functional magnetic resonance imagery. In addition we have been able to show the importance of alpha rhythm for this contrast.

In summary, this thesis provides a formal framework for classification of asynchronous EEG patterns using a generative Bayesian approach. It also provides a methodology to approach the study of EEG correlates of cognitive tasks when little is known about them and when the targeted pattern is reasonably assumed to be non time-locked.

Keywords : Classification, Bayesian approach, Generative model, Asynchronous pattern, Template matching, EEG, Decision-making, Exploration / Exploitation, Error-related Potential

Résumé

Le traitement du signal encéphalographique (EEG) s'est principalement concentré sur l'analyse de corrélés synchronisés dans le temps avec un événement observable. Cependant, quand le signal est acquis dans un environnement moins contrôlé, par exemple dans le contexte d'une interface cerveau-machine opérant dans un environnement de la vie de tous les jours, cette nature synchrone n'est plus de mise. L'analyse de tel signal requiert l'élaboration de méthodes qui s'appuie moins sur cette nature synchrone. Ces méthodes sont également le prérequis à l'étude de processus endogènes pour lesquels la connaissance de l'instant auquel le processus apparaît n'est pas disponible. Dans cette thèse, nous présentons des méthodes pour analyser les signaux cérébraux, EEG en particulier, qui ne sont pas synchronisés avec aucun événement observable. Cette thèse documente trois contributions majeures : (i) elle propose un formalisme bayésien au problème de classification de signaux EEG asynchrone, (ii) elle montre l'importance des modèles génératifs pour accomplir cette tâche et (iii) elle montre que de telles méthodes peuvent être utilisées pour assembler de l'information et classifier les corrélés EEG d'un processus de prise de décision, là où les méthodes classiques échouent.

Tout d'abord, nous proposons des méthodes pour traiter les signaux EEG non synchronisés en faisant l'hypothèse que, dans chaque essai, seule une partie contient le signal d'intérêt. Cette partie peut apparaître à tout instant dans la fenêtre d'analyse et différemment dans chaque essai. Le reste de l'essai correspond à la partie non-informative non pertinente pour la tâche ciblée. Partant d'une approche discriminative asynchrone traitant indépendamment les échantillons de l'essai, nous étendons cette méthode à un modèle génératif bayésien dans lequel chaque partie est modélisée. C'est une différence majeure avec les approches classiques qui habituellement cherchent à éviter de modéliser la partie non informative. Ensuite, en supposant que la partie informative peut être modélisée par une suite temporelle, nous adaptons la méthode précédente à un modèle bayésien asynchrone de filtrage par motif qui permet la reconnaissance du début du motif d'intérêt dans chaque essai.

Deuxièmement, nous montrons l'importance des modèles génératifs qui, grâce à l'approche bayésienne nous permet nous soulager du problème du choix des hyper-

Résumé

paramètres de l'approche discriminative initiale. Comparé au modèle initial, utiliser une approche générative implique d'utiliser plus de paramètres dans le modèle mais dont l'estimation est facilité par l'a priori que nous fournissons. Ce faisant, nous fournissons une manière plus intuitive pour l'expérimentateur d'adapter la méthode à d'autres problèmes. Avoir utiliser un modèle génératif et une estimation bayésienne nous a permis d'améliorer la généralisation du modèle asynchrone du filtrage par motif. Ce modèle a en effet été testé comme référence sur des données de potentiel évoqués et a montré sa capacité à améliorer le rapport signal-bruit, retrouver la réponse évoquée et à mieux classifier que les méthodes classiques.

Enfin, nous voyons l'importance des méthodes asynchrone pour la classification des corrélés EEG de processus de prise de décision. Nous testons ceci dans le contexte de l'étude du contraste exploration/exploitation. L'exploration est lié à la prise de décision en environnement incertain. Cette situation crée un conflit entre deux besoins : réunir de l'information à propos de l'environnement et exploiter cette connaissance pour optimiser la décision. En utilisant un protocole expérimental qui force les sujets à alterner entre actions exploratrice et exploitative, nous montrons pour la première fois qu'il est possible de classifier les corrélés EEG du comportement exploitative. Par ailleurs, nous montrons que les méthodes synchrones échoue dans la classification de ce contraste montrant ainsi la nécessité de méthodes asynchrones avancées. Ce résultat confirme également que les zones cérébrales impliquées dans cette alternance sont principalement le cortex pariétal gauche et médiofrontal ce qui est cohérent avec les résultats des études neurophysiologiques basé sur l'imagerie à résonance magnétique fonctionnelle. De plus nous avons pu montré l'importance du rythme alpha pour ce contraste.

Pour résumer, cette thèse fournit un cadre formel pour la classification de signaux EEG asynchrones en utilisant une approche générative bayésienne. Elle fournit aussi une méthodologie pour approcher l'étude des corrélés EEG de tâches cognitives peu connues et dont le signal d'intérêt est raisonnablement supposé désynchronisé.

Mots clé : Classification, approche bayésienne, modèle génératif, filtrage par motif, EEG, prise de décision, exploration/exploitation, potentiel d'erreur évoqué

Table des matières

Remerciements	v
Abstract (English/Français/Deutsch)	vii
List of figures	xvi
List of tables	xvii
Acronyms	xix
1 Introduction	1
1.1 Motivation	1
1.2 Thesis outline	4
2 State of the art and methods	7
2.1 Measuring the brain activity	7
2.2 The EEG signal used in neurosciences	10
2.2.1 Evoked potential	10
2.2.2 EEG correlates of cognitive processes	12
2.3 Methods used in EEG signal classification	13
2.3.1 Methods for time-locked patterns classification	14
2.3.2 Generative versus discriminative probabilistic models	15
2.3.3 Classification of asynchronous signal	16
2.4 Bayesian approach to estimation	17
2.4.1 Prior density and posterior density	17
2.4.2 Bayesian estimators	18
2.4.3 MCMC methods	19
3 Exploration/exploitation: protocol and synchronous approach	21
3.1 Introduction to exploitation/exploration paradigm	21
3.2 Experimental protocol	23
3.2.1 4-armed bandit	23
3.2.2 Distribution of the payoff	24
3.2.3 Signal acquisition and preprocessing	24

Table des matières

3.3	Behavioural model	25
3.3.1	Modelling subject estimation and selection	26
3.4	Synchronous approach	28
3.4.1	Classification	28
3.4.2	Results	29
3.5	Conclusion	29
4	Asynchronous approach of Exploration/Exploitation	35
4.1	Asynchronous discriminative approach	35
4.1.1	Classification context	35
4.1.2	Feature extraction: canonical variates analysis	36
4.1.3	Detection of discriminative samples	37
4.2	Results	39
4.2.1	Discriminant analysis	39
4.2.2	Classification	39
4.3	Discussion	42
4.4	Conclusion	44
5	Bayesian approach	45
5.1	Adaptation of the previous approach to a generative model	46
5.1.1	Signal model	46
5.1.2	Estimation of the parameters	47
5.2	Application to error related potentials	48
5.2.1	Model	49
5.2.2	Results	52
5.3	Discussion	55
6	Template based realignment	57
6.1	Introduction	57
6.2	Model	58
6.2.1	Template model	58
6.2.2	Non-informative samples	60
6.2.3	Training the models	61
6.3	Result on synthetic data	65
6.4	Test on real ErrP data	67
6.4.1	Methods	67
6.4.2	Results	68
6.5	Discussion	72
7	Template based classification	75
7.1	Estimation in a 2-classes problem	75
7.1.1	Model	75
7.1.2	Results	79

7.2	Classification	79
7.2.1	Model	79
7.2.2	Results	83
7.3	Discussion	84
8	Exploration/exploitation revisited	87
8.1	Introduction	87
8.2	Channels selection and prior	87
8.3	Classification results	88
8.4	Discussion	95
9	Conclusion and future work	99
9.1	Discussion	99
9.2	Perspectives	103
9.3	Future work	104
A	From Markov chains to MCMC methods	107
A.1	Markov Chains	107
A.2	MCMC	110
A.2.1	The Metropolis-Hastings algorithm	111
A.2.2	Variants and Gibbs sampler	114
B	Method for sampling a truncated normal distribution	119
C	notations	121
C.1	Mathematics	121
C.2	Probability	121
C.3	Distributions	122
	Bibliography	123
	Curriculum Vitae	133

Table des figures

1.1	Examples of class-specific patterns (synthetic data)	2
2.1	Schema of the generators of the EEG signal.	9
2.2	Examples of evoked potentials	11
2.3	Separating hyperplane of LDA.	13
2.4	Separating hyperplane of SVM.	14
3.1	Experimental protocol	23
3.2	Example of the payoff evolution	23
3.3	Electrodes used for analysis	25
3.4	Comparison a typical behaviour with the model predicted behaviour	25
3.5	Synchronous approach of classification on exploration/exploitation data	30
4.1	Example of trials of 2 classes (synthetic data)	36
4.2	Distribution of the two classes in the canonical space	37
4.3	Discriminant power	38
4.4	Performance of the combined classifier	43
5.1	Graphical model of the signal distribution	46
5.2	Grand average of an ErrP	49
5.3	Distribution beta as used for the prior of λ .	49
5.4	Distribution of the informative samples	53
5.5	overlapped plots of all trials	55
6.1	Graphical model of the signal distribution of the trial	59
6.2	Examples of trials	59
6.3	Distribution of the peaks of ErrP	60
6.4	Examples of skew normal distributions	61
6.5	100 trials of synthetic ErrP	66
6.6	Example of convergence in 3 synthetic trials	67
6.7	Prior used for the time-onset	67
6.8	Results of estimation on ErrP (subject 1-6)	69
6.9	Results of estimation on ErrP (subject 7-12)	70

Table des figures

7.1	Graphical model of the signal distribution of the trial	76
7.2	Results of estimation on ErrP, 2 classes (subject 1-6)	80
7.3	Results of estimation on ErrP, 2 classes (subject 7-12)	81
7.4	Examples of beta distribution with different parameters.	82
8.1	DP: subject 2	89
8.2	DP: subject 4	90
8.3	DP: subject 10	91
8.4	Prior on time onset	93
8.5	Estimated template and onset distribution: subject 8	94
8.6	Estimated template and onset distribution: subject 2	94
8.7	Estimated template and onset distribution: subject 1	94

Liste des tableaux

2.1	The rhythms observed in EEG	8
3.1	A summary of number of trials	26
3.2	Estimation of the parameters of the behavioural model	27
3.3	Classification results (Subject 1 and 2)	31
3.4	Classification results (Subject 3, 4 and 5)	32
3.5	Classification results (Subject 6, 7 and 8)	33
4.1	Classification accuracies per subjects (EEG signal)	40
4.2	Classification accuracies per subjects (EOG signal)	41
5.1	Estimated lambda	52
5.2	Classification results on test set	54
6.1	Correlation of the grand averages	68
6.2	Signal to noise ratio of the evoked response	71
7.1	Classification results on test sets (MCC values)	83
8.1	Classification results on different electrodes	92

Acronyms

BCI	Brain Computer Interface
BOLD	Blood Oxygenation Level Dependent
BP	Bereitschaftspotential
CVA	Canonical Variates Analysis
DC	Direct Current
DP	Discriminant Power
EEG	Electroencephalography
ECoG	Electro-Corticography
EMG	Electromyogram
EOG	Electrooculogram
ERD	Event Related Desynchronization
ERP	Event Related Potential
ERN	Error Related Negativity
ErrP	Error Related Potential
fMRI	functional Magnetic Resonance Imaging
LDA	Linear Discriminant Analysis
LFP	Local Field Potential
MCC	Mathews Correlation Coefficient
MCMC	Monte Carlo Markov Chain
MEG	Magnetoencephalography
NIRS	Near-Infrared Spectroscopy
QDA	Quadratic Discriminant Analysis
rLDA	realigned Liner Disriminant Analysis
SNR	Signal-to-Noise Ratio
SSVEP	Steady State Visual Evoked Potential
SVM	Support Vector Machine
VEP	Visual Evoked Potential

1 Introduction

1.1 Motivation

Interacting with a machine using a Brain Computer Interface (BCI) relies on the recognition of cognitive processes based on their correlated brain activity which is most of time the Electroencephalography (EEG). One of the most popular ways to implement such systems is to exploit Event Related Potential (ERP) or evoked synchronisation of desynchronisation (i.e. automatic responses of the brain to external stimuli) [1]. Due to their time-locked nature evoked responses are usually well distinguished from background activity and easily picked out but require the subject being “synchronised to external machinery” (i.e. system-paced). Another commonly used approach is to identify the EEG correlates of spontaneous activity such as motor task imagination (e.g. imagination of right or left hand movement) [2]. In these asynchronous BCIs, the user provides mental commands at the pace he wants and the system assumes that the *current* measured EEG activity (maybe mixed with noise) relates to the task the user want to perform.

At this stage, we realise that the methods needed to detect and classify the EEG pattern for these type of BCIs do not require any complex method to handle the time information : either the pattern is synchronised to an observable event, either the time information is discarded by considering the time sample as independent and identically distributed, thus allowing to use methods similar to the first case. This has led to the development of many methods suitable for time-locked EEG. However, there exist many situations in real-world in which the assumption of time-locked pattern does not hold. First it may happens that the recording chain does not ensure perfect synchronisation and introduces delay variability between the stimulus and EEG response in the case of ERP-based BCI. Second the complexity of the task can also introduce variability in the reaction-time of the subject.

Another issue arises when we want to study induced EEG activity. This can be illustrated




FIGURE 1.1 – Examples (synthetic data) of trials of 2 classes : the class-specific part in each trial are highlighted in their corresponding colour. In y-axis : EEG potentials, x-axis : time (arbitrary units).

in the context of perception. When stimulus is presented to a subject, an evoked brain activity appear synchronised with the stimulus onset. But when the subject perceive the stimulus in its context (like face recognition), an induced brain activity appear with a jitter in latency different for each stimulus presentation [3]. The analysis of the induced activity requires asynchronous methods [3, 4]. Similarly, when studying endogenous processes like decision-making process, the information flow that lead from the decision to the observable action is quite complex and involves many brain structures. Thus the correlates of these is likely to be very loosely synchronised with the observable events (i.e. the decided action). In such cases, the classical methods used in EEG signal analysis are rather unadapted.

This thesis aimed at addressing the single classification of such asynchronous neural events. To address this question we have made the assumption that a class-specific pattern can appear in different in each trial (see illustration in Figure 1.1). The goal of this thesis is to develop a method that will learn the parameters of this pattern for each class, and given them, be able to classify each trial and estimate the time-onset of the class-specific pattern in each trial.

In order to develop such a single trial asynchronous classification method, we will approach the problem by modelling the EEG signal of each trial by two parts, one specific to the class which will call the informative part and one that is common to all the classes, the non informative part. As such, the methods developed in this thesis will differ from the classical methods in the sense that the non-informative part is specifically modelled. The introduction of this latter model allows the method to identify on single trial where the targeted pattern appear within the trial.

So at the opposite to them, we will tend to develop in this thesis a generative model of an asynchronous classifier. Increasing the number of components in the model may

rise problems in the estimation of its parameters. To mitigate the issue, we will use a Bayesian approach of the estimation of model parameter achieved with sampling techniques, i.e. Monte Carlo Markov Chain (MCMC) algorithm whose use in the context of EEG signal processing is quite novel. In spite of its complexity, we will see that this generative approach can be more robust than classical approach in the context of highly jittered EEG signal.

Given the complexity of the development of such methods, we use an incremental approach to build them. At first, we develop an asynchronous discriminant classifier based on the dichotomy of the informative and the non-informative set. The next step start to use a generative model to extend the previous classifier. In this new step, the method provides a formal probabilistic model of the informative part for each class. At this step we validate the parameter estimation method using the data of Error Related Potential (ErrP) which is an evoked potential. This allows us to test the estimation with an EEG signal that is known and whose the time of appearance is known. The next step starts modelling the informative part as time sequence that can appear at any latency within a trial. This step estimates only one informative pattern to test the estimation of the pattern model and its onset. Finally the previous method is extended to several classes and tested on ErrP data.

In order to illustrate the importance of such methods, we will approach the complex problem of classifying the EEG correlates of a decision-making process : the exploratory behaviour. Decision making in an uncertain environment raises a conflict between two opposing needs : gathering information about the environment and exploiting this knowledge in order to optimise the decision. These two needs are opposed because gathering information may not lead to the optimal decision, while exploiting the current knowledge precludes from testing other possible options that may lead to optimal long-term performance. Recent brain imaging studies using functional magnetic resonance imaging (fMRI) or positron emission tomography (PET) have focused on the identification of brain signatures of decision making. More specifically, recent intracranial recordings in primates and fMRI studies in humans have studied the neural substrates of the exploratory behaviour in different experimental paradigm, like maze navigation or gambling task [5, 6, 7]. Given its time resolution, EEG could give an information that fMRI does not provide, especially in terms of frequency components and, since both techniques measure different physical phenomena, the detection using one technique does not necessarily lead to the detection by the other one. To our knowledge, no studies have focused on the EEG correlates of the exploratory behaviour. Using the asynchronous classification methods developed in this thesis, we will show that it is possible.

1.2 Thesis outline

This thesis is composed of 9 chapters. The next chapter 2 focuses on the state-of-the-art in the field of EEG analysis and the Bayesian methodology. Chapter 3 and 4 shows the possibility of studying the EEG correlates of the exploratory behaviour and the need of asynchronous methods to achieve this. Chapter 5, 6 and 7 extend step by step the initial discriminant approach into a Bayesian generative classification method for asynchronous EEG patterns. In Chapter 8, we will test this method on the exploratory behaviour data. Finally, Chapter 9 provides an overview of the thesis along with future directions

Chapter 2 provides an overview of the state-of-the-art relevant to this thesis. We provide a brief introduction to the various measurements techniques used in the context of neuroscience and review the different EEG correlates used for classification. We also review the different classification methods used in the context of EEG analysis. Finally, we will describe the methodology of the Bayesian approach and some of the sampling techniques used to perform Bayesian inference.

Chapter 3 describes the protocol used for studying the exploration/exploitation contrast using EEG signal. We will describe the behavioural model used to study it as well as its fitting results. Finally we will show that the synchronous approach of classification do not allow to classify the EEG correlates of exploratory behaviour.

Chapter 4 describes a discriminant asynchronous classification method and applies it on the exploration/exploitation data. With this approach, we divide the data of each class into a set of informative sample and a set of non-informative sample based on their distribution in the canonical space by discarded the timing information. With this method, we show that it is possible to classify the EEG correlates of the exploratory behaviour.

Chapter 5 proposes an improved probabilistic formalism to the classification method used in the previous chapter. We model each informative set and the non informative set by a probability distribution and we model as well the probability of appear of an informative sample in a trial. By adapting it to a Bayesian approach, it allows us to get rid of the metaparameters used in the previous method. To method is tested on ERP data.

Chapter 6 describes the extension of the previous method to case of the informative part being modelled as a sequence in the time domain. While previous approach allows the informative part being scattered within a trial, it is assumed here that it has a temporal structure that can be used to recognise it. This method uses only one class and is tested for realignment of ERP response on both synthetic and real data. By using a direct model of the informative part, we are able to learn class specific pattern in jittered data.

Chapter 7 extends the Bayesian realignment method to a full Bayesian asynchronous

classification method. This method is tested on real ERP data and compared with classical methods on normal or highly jittered EEG signal.

Chapter 8 tests the method developed in the previous chapter on EEG data of the exploration/exploitation paradigm. By applying it channel by channel, we are able to classify EEG correlates of the exploratory behaviour.

Chapter 9 summarise the contributions of the thesis in the development of Bayesian classification methods for asynchronous EEG pattern. We also discuss of the future directions for improving these methods.

2 State of the art and methods

This chapter provides an overview of the different type of signal that can be used to measure the brain activity and how they differ. Focusing on the EEG signal, we will see what are the usual type of EEG activity mainly studied and the type of characteristics they exhibit. Then an overview of the classification methods used to analyse this EEG activity. The last section of this chapter will provide an overview of the Bayesian estimation which is the framework in which the classification presented in this thesis will be developed.

2.1 Measuring the brain activity

In vivo, the brain activity can be observed by different ways, which all correspond to measurable physical changes more or less directly due to the neuronal activity of either a single neuron or its interaction with other. This can manifest as change in its electrical, magnetic or metabolic activity. Different technologies are used to measure each of these activities and categorised mainly as, invasive, partially invasive and non-invasive measurements (for detailed reviews on these measurement techniques refer to [8, 9, 10, 11]).

With the advent of measuring technique and the computational power to process them in real time, the idea raised in the scientific community in the last 30 years to use the measured brain activity to directly control artificial devices while bypassing the natural end effectors of subjects, thus creating a BCI. As such the BCI technology is meant to provide a non-muscular communication and control channel for people with severe disabilities such as Amyotrophic Lateral Sclerosis, brainstem stroke, spinal cord injury and muscular dystrophy [12, 13, 14, 15, 16]. Different applications have been proposed like an environmental control [17, 18], a speller in which the BCI allows to type a text [19] or a continuous robotic control like a robotic arm control [20] or a wheelchair control [21] (see review on [22]). The development of BCI systems, the

TABLE 2.1 – The rhythms observed in human EEG.

Name	Range
delta	0.1–4.0 Hz
theta	4.0–8.0 Hz
alpha	8.0–12.0 Hz
beta	14.0–30.0 Hz
gamma	30.0–80.0 Hz

advance of neuroscience and the need of refined diagnostic methods when treating brain related diseases has led to the development of methods to measure and analyse brain activity. To do so, a compromise has to be found between the quality of the used brain signal, the invasiveness of the method, and the portability of the measuring device [23].

In the case of invasive measurements, single electrode [24] or electrode arrays (for recent developments, see [25]) can be implanted in the cortex (penetrating a few millimetres). It is also possible to measure electrical activity of deep brain structures using depth electrodes (e.g. [26, 27]). Such implanted electrodes record the activity of a single neuron, multiple neurons activity as well as so called Local Field Potential (LFP) that represent summed activity of a population of neurons within a volume of neural tissue [28]. Signals recorded from these electrodes are of high Signal-to-Noise Ratio (SNR) due to proximity to the sources, with a risk associated with surgical procedures and tissue damage (scar) [29, 27]. Besides this, a major drawback of invasive measurement is biocompatibility, that is implantation of electrodes have to be redone often because immune reactions tend to attack electrodes and impair dramatically their sensitivity [10]. However, the progress in electrode design could leverage advantages over risks. The invasive approach has been successfully used with monkeys to control a robotic arm [20] and a few attempts have been made with human subjects [30], but these are rare due to ethical and logistic reasons [9, 10, 11].

To reduce the risk of damaging the brain with implanted electrodes, an alternative approach is to place sheets of electrode arrays directly on the surface of the cortex but still below the skull, beneath the dura mater. As such this approach qualified as semi-invasive. The signal used is then the Electro-Corticography (ECoG), which represents only population activity (i.e. LFPs) but not single unit or multi unit activities [27]. Compared to non-invasive measurements of the electrical brain activity, the ECoG is superior in terms of SNR, can record both slow rhythms (delta, theta, alpha and beta bands; ranges are shown in Table 2.1) and fast rhythms (gamma), and has a higher spatial resolution [31, 32, 33]. The most frequent application of ECoG recordings is in the localisation of epileptic seizures for surgical planning [34]. At this occasion, the patients are sometime invited to participate to neuroscientific experiment. Due to its

2.1. Measuring the brain activity




FIGURE 2.1 – Schema of the generators of the EEG signal.

superior signal quality compared to non-invasive techniques and its low clinical risk, it is a promising intermediary modality [32]. However this approach still requires surgery to place the electrodes.

Completely non invasive (and risk-free for the subject), the EEG is used in both clinical and psychophysical context. In this recording technique, a number of electrodes (up to 256) are placed on the scalp. Early techniques of EEG measurement required scrapping the skin for good contact, but most laboratories currently use less aggressive procedures although still utilise gel-based electrodes. Driven by market potential, recently a few companies are producing dry electrodes that could be the ultimate recording methods for the needs of some of the BCI applications. The EEG recording technique is very attractive due to easy usage, cost and safety [11]. However, since the placement of electrodes is far from brain tissue and because of the smearing effect of the skull (see Figure 2.1), these recordings suffer low SNR and low spatial resolution [35]. Nevertheless, appropriate machine learning and signal processing techniques may significantly improve the features extracted from EEG signals [36, 37].

Another non-invasive measure is the Magnetoencephalography (MEG) that captures magnetic activity of the brain with a time resolution similar to EEG [38]. Although it can record structures deeper than EEG, the MEG measurement is extremely sensitive to any magnetic perturbation and requires a dedicated electromagnetic shielded room, thus making it unusable outside of the controlled environment of a laboratory. Finally, due to the causal relationship between neural and metabolic activity, other non-invasive recording techniques such as Blood Oxygenation Level Dependent (BOLD) functional Magnetic Resonance Imaging (fMRI) and Near-Infrared Spectroscopy (NIRS) are recently adopted for BCI applications [39, 40]. The upside of fMRI is that it provides

millimetre spatial resolution and better than NIRS but the set-up is bulky and expensive. Note that like for the MEG, using the fMRI signal cannot be considered outside of dedicated room. Both fMRI and NIRS suffer from poor temporal resolution compared to electrical or magnetic activity measurement techniques [41].

This thesis will concentrate only on the EEG signal processing.

2.2 The EEG signal used in neurosciences

The EEG reflects the synchronous activity of a large population of neurons (Figure 2.1). EEG is frequently used in experimentation because this technique is simple compared to other brain activity measurement techniques and the process is non-invasive. EEG is capable to detect changes in the brain electrical activity on the range of the millisecond. English physician Richard Caton already discovered the presence of electrical current in the brain in 1875, but it was not until 1924 that German neurologist Hans Berger made the first brain's electrical activity recording on graph paper [42]. Already at that time, Berger noticed that brain waves varied with the individual's state of consciousness.

EEG recordings usually present rhythmical patterns which are often classified into 5 frequency bands (see Table 2.1). It is worth noting that this frequency bands correspond to what can be observed in EEG. The brain generates oscillatory activities higher than 80 Hz at the surface of the cortex but because of the skull and the dura mater, those frequencies are greatly attenuated when measured by the scalp electrodes (they can be measured by ECoG).

2.2.1 Evoked potential

ERP are signals generated by a population of neurons in response to a perceptual, cognitive or motor event, in opposition to spontaneous activity that reflects the brain activity related to volunteer self-paced tasks. Evoked potentials can be seen as a specific kind of ERP generated directly in response to external stimulus such as Visual Evoked Potential (VEP) and auditory evoked potentials. Reactions to stimuli or events lead to variations of the electrical activity of specific brain areas and the resulting EEG traces exhibit modifications called evoked potentials. In the case of external stimuli, for example discrete visual feedback, the precise time of the stimulus is known. It is therefore possible to extract averages of the stimulus-locked response of the brain. Time-locked averages allow elimination of random noise while keeping track of the ERP components. When the precise time of the stimulus is not available, it is much more complicated to extract ERP from the ongoing EEG. In this case, a specific action of the subject related to the nature of the stimulus can be used as trigger. In any case, the main challenge is no longer the detection of ERP in averages over many trials, but

2.2. The EEG signal used in neurosciences




FIGURE 2.2 – Example of evoked potentials. (c) extracted from [43].

directly at the level of the single trial. Here, some ERPs are presented to overview the signal characteristic of these potentials.

The so-called P300 component shows up as a very prominent positive deflection (usually with a parietal focus) of the ongoing EEG about 300 ms after the occurrence of an infrequent or particularly significant stimulus interspersed with frequent stimuli (Figure 2.2a). This stimulus can be of various nature : visual, auditory or somatosensory. Traditionally, P300 has been used in BCI research to develop virtual keyboards (i.e. P300-spellers) [44, 45, 13] with a typing speed of five letters per minute, but recently this ERP has also been the basis for brain-actuated control of a virtual reality system [46] and of a wheelchair [47].

VEP are EEG waveforms generated in response to visual stimuli that can be used to determine the direction of eye gaze. When stimuli are presented in a rapid succession, the evoked potentials overlap in time and the presentation rate is high enough to evoke a steady wave, this is referred to as Steady State Visual Evoked Potential (SSVEP). These signals are induced by a stimulus repeated at a rate higher than 6 Hz [48, 49, 50].

The ErrP described exhaustively since the early 1990s [51] is elicited in speeded reaction tasks 100 ms following an erroneous response or in reinforcement learning tasks 250 ms following presentation of a feedback that indicates incorrect performance [52]. The negative feedback, indicating incorrect performance, elicits an Error Related Negativity (ERN), whose main component is a negative deflection that occurs 250 ms after the feedback with central or frontocentral focus (Figure 2.2b).

Bereitschaftspotential (BP), also called readiness potential or pre-motor potential reflects activity in the motor cortex during voluntary muscle movement preparation. BP is a negative deflection which develops in a central area during the last second before limb movement.

Finally some EEG correlates of how a listener processed the structure of a spoken language have been identified. The closure positive shift is a large positive waveform

that is elicited at different timing depending on what constituent of the sentence the listener focused [43, 53] (Figure 2.2c).

2.2.2 EEG correlates of cognitive processes

Unlike low-level motor process or sensory processes which are exogenous activities, cognitive processes are by nature endogenous activities i.e. they are not elicited directly by the environment. To achieve a goal or interpret external stimuli, the brain has to integrate activities of different processes (both endogenous and exogenous) coming from different brain areas producing finally a coherent behaviour or perception. Since this integration involves numerous different cortical areas and is done on a large scale in the brain, it requires mechanisms through which brain areas can communicate to produce this coherent behaviour or perception. The current main theory is that transient phase synchrony of the oscillatory EEG activity reflects this large scale integration across different cortical areas [54].

Currently large scale integration has been mainly studied for the sensory inflow. Especially vision has become a paradigmatic example of this approach, and the successive stages of elaboration of the visual stimuli from retina to the various visual areas have been extensively studied [3, 55, 56]. Most of these studies have stressed the importance of the transient gamma phase synchrony for the conscious perception of the stimulus [56]. One of the main challenges in these studies is to detect induced responses to the stimuli. While evoked response is time-locked to the stimulus onset and can be seen in the averaged evoked response, the latency of induced activity varies across trials and is therefore cancelled out by averaging [3]. Specific methods are thus required to detect and characterise this induced activity, which is what the methods developed in this thesis are designed for.

Time-frequency decompositions and complex wavelet transforms in particular are well-suited to represent it [3]. These last years, many methods based on them have been proposed to measure or characterise the phase synchrony. While early methods aim to measure simply phase-locking of different frequency bands of different brain areas on short time-windows [57, 58], more elaborated techniques like Frequency Flows Analysis [59] appeared recently to track and characterise the non-stationary time-frequency dynamics of phase synchrony among groups of signals, interestingly allowing possible time varying frequencies of synchronisation. Apart from the large-scale integration which explain how different brain areas interact, spectral distribution of energy should reflect also the induced activity of cognitive processes. These two components should help to characterise and detect them. Some of them will be used in this thesis.

The exploratory behaviour is proposed as an example to build and test this method-

2.3. Methods used in EEG signal classification




FIGURE 2.3 – Separating hyperplane of LDA.

ology. Exploration is related to decision making in an uncertain environment. This situation arises a conflict between two opposing needs : gathering information about the environment and exploiting this knowledge in order to optimise the decision. These two needs are opposed because gathering information usually does not lead to the optimal decision, and exploiting prevents from testing other possible options so from accumulating or updating the knowledge. This exploration/exploitation dilemma is an important point of the reinforcement learning theory [60, 61]. While recent medical imagery studies have stressed the role of prefrontal in the estimation of uncertain states in a navigation task [5], other studies have found activation in the frontopolar cortex and intraparietal sulcus during exploratory decisions [7].

2.3 Methods used in EEG signal classification

There are several different techniques that have been employed to detect or classify patterns in the EEG signal. As mentioned in the previous section, we distinguished two kinds of problems : those in which the pattern to recognise/classify is synchronised (also said time-locked) with an *observable* event which then indicates precisely the analysis window. The other kind of problem is those in which the location in time domain is not known very well. We will refer the methods that deal with such problem as asynchronous methods. Because of their simplicity the methods that deal with time-locked EEG patterns have received a lot attention [37, 62]. Moreover, in a controlled environment, these methods are suitable for classifying ERP [45] or Event Related Desynchronization (ERD) [1].




FIGURE 2.4 – Separating hyperplane of SVM.

2.3.1 Methods for time-locked patterns classification

Many of these methods are variations of two classical algorithms that aim at solving the same linear classification problem with different criteria [63]. The first, the Linear Discriminant Analysis (LDA), is used very often in EEG signal classification [37]. Assuming that the features of the classes follows a Gaussian distribution and that they share the same covariance matrix, it tries to find hyperplanes in the feature space to separate the classes. In the case of a two-class problem, it corresponds to finding the hyperplane that separates the best the two classes (Figure 2.3). The space orthogonal to the hyperplane, called the canonical space, is the space in which, when projected on it, the classes are the most separated.

If the covariance matrix of each class cannot be approximated as equal, it is therefore useful to revert to an extension of the LDA that takes into account these difference called the Quadratic Discriminant Analysis (QDA). Although it is sometimes employed in the context of EEG processing [64], it is often avoided since the QDA requires the estimation of as many covariance matrices as the number of classes. Usually, in the context of EEG pattern classification, the number of dimension of the feature space is relatively high with respect to the number of trials used in the training sets. This often forces to employ LDA with regularisation techniques to favour for example spherical or diagonal covariance matrix [65].

Another popular method is the Support Vector Machine (SVM). Like LDA, it search the hyperplane (or set of hyperplanes for a multiple class problem) that separates most of the data. The major difference is the optimisation criterion of the SVM which tries

2.3. Methods used in EEG signal classification

to optimise the margin along the hyperplane between the class data Figure 2.4. We clearly see that the data follows a Gaussian distribution, the computed hyperplane will be close to the one computed by LDA. If the class data are not linearly separable (often the case when dealing with EEG data) a misclassification penalty function can be used but the SVM does not perform very well with non linear penalty function [66]. The use of SVM for classification of EEG data is quite common [67].

Both LDA and SVM can be used in the non linear case if the experimenter is able to find a transform of the feature space into a space with makes the class data linearly separable. However such transform is not necessarily easy to discover and can introduce hyperparameters making the training procedure more difficult if this transform is parameterisable.

2.3.2 Generative versus discriminative probabilistic models

It is worth noting that these methods are both discriminant : they model the conditional probability distribution of the observation given the unobserved variable (the class in our case), which can be used for predicting the class from the measured EEG signal. At the opposite, generative approaches model the joint probability distribution over observation and the unobserved variables. A generative model can be used, for example, to simulate (i.e. generate) values of any variable in the model, whereas a discriminative model allows only sampling of the target variables conditional on the observed quantities. Some EEG analysis studies uses generative approaches [68]. However, they stay a minority : most of them use discriminative approaches [37].

There are several compelling reasons for using discriminative rather generative classifiers, one of each, is that it is usually better to solve the problem directly than solving a more general problem as an intermediate step [69]. It is true that generative models require more parameters to estimate from the same data amount of data, which make lead overfitting problems if not handled carefully. However a Bayesian approach may temper these issue by incorporating prior knowledge into the estimation. Moreover, since a generative model learns the likelihood model of the data, it provides a natural way to update the estimated when new data becomes available which is not the case for discriminative model.

Even if it often requires a complex model with more parameters than using a discriminative approach, a generative model can be used to validate an incoming piece of data against it. This provides a measure of confidence in the decision made by the model or provides a way to notify when the model starts to be outdated. This is an interesting characteristic when doing adaptation. Although this thesis will not deal with the problem of adaptation, it is a characteristic that should be kept in mind when choosing a method.

2.3.3 Classification of asynchronous signal

When dealing with asynchronous signal, we should consider different cases. First case is the one where the pattern to classify has a shape over time which is supposed to be relatively similar over trials but for some reason, appearance onset is not constant in all trials. This is typically the example of ERPs whose response is time shifted because of reaction time of the subject or because of the recording instrumentation which is not very well synchronised by the stimulus generator. For these cases, when the time variability is known to be not too high, the usual approach is to realign on a *known* positive or negative peak that the expected response should contain [70, 71].

Another case, is when dealing with a pattern that could last for an undetermined amount of time and can happens anywhere in the analysis window. This is typically the case of a self-paced BCI. Methods that deal with this case often approach each time samples independently. After preprocessing to extract feature of interest (often frequency band power), these methods try to classify each time snapshots (sliding classifier). The methods used can then range from simple thresholding [72] to more complex neural network based discriminant approach [73], passing by LDA [74]. These methods do not differ much in principle from the methods used for time-locked EEG signal in the sense that they discard any time information by considering the time sample to be independent and identically distributed.

However they often face the problem of non class related nonstationarities in the signal. This is a typical problem with self-paced BCI with the use of discriminant methods which cannot model accurately the signal when the subject does not perform one of the possible task. This is often worked around by using a framework of "evidence accumulation" [75, 76] which will fuse the outputs of several successive classifiers (in time). This is can be done under the assumption that, if the subject performs consistently one of the possible task, the classifiers output will naturally be biased towards the target class, while if performing none of the task, the classifier output will not be biased towards any of the task.

Finally, more general approaches combining the problems of the two previous classification cases have also been proposed. Considering trials of the same tasks, the signal of interest is assumed to be non-time-locked (to a certain extent) and the length of the epoch relevant activity can be variable. This problem has been approached with the analysis of LFP activity recorded from rats [4]. They proposed a method to identify prominent oscillatory activities in the time-frequency maps (bumps in the maps) of each trial and then train a neural network to learn the shapes and the relative position of the bumps relevant for the target tasks. Although it is an interesting approach that deal with a very generic problem, first it is extremely expensive from the computation point of view. Second they suffer from the same issue that the discriminative approach, i.e. modelling the non-relevant part of the signal.

2.4 Bayesian approach to estimation

In this thesis, we propose to approach the problem of single trial analysis of asynchronous pattern. Moreover we propose to use generative model to achieve this. This often leads to use a Bayesian approach to estimate all the parameters of the model. To this end, this section details the framework of the Bayesian approach to the problem of estimation.

The context of an estimation problem is the following : considering a system whose behaviour (the response) depends on several parameters globally noted as θ , the purpose of estimation is to retrieve θ knowing the observation noted y .

2.4.1 Prior density and posterior density

From a probabilistic point of view, parameters are represented by a random variable Θ distributed by some probability laws to be determined. Two important probability densities are then considered. The first is the probability density of the parameters independently from any observation, called *prior* density. It represents the information that is known about the parameters without having done any measurements. The latter is the probability density with the knowledge of the observations, called the *posterior* density.

In a Bayesian approach, any estimation is deduced by examining the posterior density of the parameters of interest [77]. Or said otherwise, the estimation relies on their probability density conditionally to the observation *and* on the prior information at our disposal concerning the parametrisation and the shape of the model.

Let us suppose that y denotes the observations and θ the parameters to be estimated. The joint *posterior* density $p(\theta|y)$ of all parameters is given by the Bayes rules :

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{\int_{\Theta} p(y|\theta)p(\theta)d\theta} \quad (2.1)$$

with :

- $p(y|\theta)$ corresponds to the conditional density of the observations knowing the parameters. It represents the direct form of the model.
- $p(\theta)$ is the *a priori* density of the parameters.
- $\int_{\Theta} p(y|\theta)p(\theta)d\theta$ is a constant with respect to y .

One of the fundamental features of the Bayesian approach is the use of a probability density describing the state of knowledge or ignorance concerning the parameters without taking the observations into account. The choice of such prior density $p(\theta)$ is one of the most delicate and criticised issues of the Bayesian approach. It is indeed quite rare that the available prior information is precise enough to determine exactly

this density. Approximation should thus be done and this choice may have influence on the Bayesian statistical conclusions. However, this influence should not be overstated since, as the number of observations increases, the influence of the prior fades, to have no influence at all at the limits. Coming back to the choice of the prior, one might wonder how to deal with the choice of the density of those particular parameters whose prior information is not available. However, this issue is quite frequent and can be easily solved by choosing a non informative density function which in the case of finite samples (our case in this thesis) corresponds to the uniform law.

2.4.2 Bayesian estimators

Assuming the problem of the choice of the prior density solved, we still have to estimate the parameters from their joint posterior density. The knowledge of the density represents too much information, and as such, it is unexploitable. We indeed look for a probable value for the parameters with possibly a confidence interval and not a infinity of possible values. It is thus necessary to have a Bayesian estimator which will provide this value from the density.

However the use of such estimator gives an incomplete description of the posterior density. A unique statistic cannot characterise the whole density. The problem is then to choose the statistic characterising the posterior density at best.

There are many different estimator, but the mostly used are the maximum *a posteriori* (MAP) which searches the maximum of the posterior density or the expectation *a posteriori* (EAP) which computes its mean. The choice of the estimator depends on the problem to be solved and should be done carefully : one should keep in mind that depending on the problem some estimators might lead to degenerated solutions.

At this point, it should be noted that, although the computation of the posterior distribution is easy, the analytic calculus of the estimators is usually not easy and is very rarely possible. Numerical integration or optimisation methods can be considered but experience shows that when more than two dozens of parameters have to be estimated, those methods are not reliable any longer. Another approach is to simulate the distribution, i.e. by drawing a sufficient number of samples from the posterior distribution, a accurate approximation of the distribution will be build. This simulation will be realised by the mean of adapted methods like MCMC methods (standing for Monte Carlo Markov Chain). The general idea with these methods is to generate a number of samples large enough to have a good representation of the posterior distribution and then to deduce from them the Bayesian estimator.

2.4.3 MCMC methods

We aim at solving the following problem : given a probability density function noted ρ , how to draw samples from a random variable following this density ?

The principle of the MCMC methods is to use a Markov chain (a sequence of random variables) whose marginal law converges toward the target density (principle which gives the name to the method : Monte Carlo Markov Chain) [78]. The first section of Appendix A recalls some definitions and results about the Markov chains while its second section describes the mathematical foundations of the MCMC algorithms.

There are usually two algorithms that are considered when using MCMC methods although the second is a particular case of the first : the Metropolis-Hastings algorithm and the Gibbs sampler. The Metropolis-Hastings algorithm is the most general case in the sense that it can sample any type of density. It requires to choose an instrument density q which can be sampled, i.e. there are methods to generate a sample from the probability density function. Uniform and Gaussian density are examples of such densities that can be sampled. In addition, the support (the set of points on which the function is non null) of the instrument density should contain the support of the target density.

Under such a condition, we have an algorithm that defines a Markov chain $(\Theta[n])_{n=1:+\infty}$ which converges towards the target density ρ . Said otherwise, by noting as $\theta[n]$ a sample from $\Theta[n]$, the algorithm provides a way to generate $\theta[n]$ such as with n increasing, we tend to have $\theta[n] \sim \rho$:

Metropolis-Hastings

Given

- target density $\rho : \theta \mapsto \rho(\theta)$;
- instrument density $q : (\theta', \theta) \mapsto q(\theta'|\theta)$;
- Initialisation $\theta[0]$.

loop (over $n \geq 0$)

- draw $\theta^*[n + 1]$ following $q(\cdot|\theta[n])$;
- draw $u[n]$ following a uniform law between 0 and 1 ;
- calculate $\alpha[n] = \min \left(\frac{q(\theta[n]|\theta^*[n+1])\rho(\theta^*[n+1])}{q(\theta^*[n+1]|\theta[n])\rho(\theta[n])}, 1 \right)$;
- $\theta[n + 1] = \begin{cases} \theta^*[n + 1] & \text{if } \alpha[n] > u[n] \\ \theta[n] & \text{otherwise} \end{cases}$

However, the generality of this method may come at the price of efficiency. If the marginal of the instrument density $\int_{\theta} q(\theta'|\theta)d\theta$ is very different from the target density, the coefficient $\alpha[n]$ will be very small resulting in a very poor speed of convergence toward the target density. It is thus very desirable to choose an instrument density as

Chapitre 2. State of the art and methods

much similar to the target density as possible.

The Gibbs sampler is a possible solution to this problem when the Markov chain $(\Theta[n])$ can be decomposed in the following way :

$$\Theta[n] = \begin{bmatrix} \Theta_1[n] \\ \Theta_2[n] \\ \vdots \\ \Theta_d[n] \end{bmatrix}$$

In such case and if all conditional laws $\rho(\theta_{1:i-1}[n+1], \dots, \theta_{i+1:d}[n])$ can be sampled, there is a very efficient way to sample the target density ρ :

Gibbs (shifting update)

Given

- target density $\rho : \theta \mapsto \rho(\theta)$;
- initialisation $\theta_{2:d}[0]$.

loop (over $n \geq 0$)

- For i varying from 1 to d ; draw $\theta_i[n+1]$ following the law $q^{(i)}(\cdot | \theta_{1:i-1}[n+1], \theta_{i+1:d}[n]) \propto \rho(\theta_{1:i-1}[n+1], \dots, \theta_{i+1:d}[n])$

It might be argued that the situations in which the Gibbs sampler could be used are quite rare in practise. However, when considering a complex model, its complexity often comes from the fact that it takes into account many different and small interactions between the different variables of the model. So actually, what makes the model complex provides also a natural decomposition often suitable for the Gibbs sampler. In addition these small interactions are also modelled in a simple way, often by using usual distributions (Gaussian, uniform, beta, gamma...), thus resulting in conditional density functions that can be easily sampled.

3 Exploration/exploitation : protocol and synchronous approach

In this chapter we will start to study one high level cognitive process : the switch between exploratory and exploitative behaviour. We will study it using an experimental protocol inspired by other neuroimaging studies.

When dealing with the repetition of the same task, classical approaches of EEG signal processing tend to consider the EEG signal of a trial as a time locked measure of the same phenomenon. In other words, although there might be slight variations of the electrical potential measured at the same timing over all trials, the EEG response is assumed to have a similar waveform in all trials and to occur at the same latency in each trial. This assumption allows to average all trials to discover the mean pattern or to classify a trial by feeding its sequence of points as features to a classifier such as an LDA.

This chapter describes the experimental protocol used to study the EEG correlates of the switch between exploratory and exploitative behaviour and how suitable the classical approaches of EEG processing are for the data recorded in this paradigm of decision making.

3.1 Introduction to exploitation/exploration paradigm

Decision making in an uncertain environment raises a conflict between two opposing needs : gathering information about the environment and exploiting this knowledge in order to optimise the decision. These two needs are opposed because gathering information may not lead to the optimal decision, while exploiting the current knowledge precludes from testing other possible options that may lead to optimal long-term performance.

This conflict is an important point of the reinforcement learning theory and is classically tested by the N -armed bandit problem [60]. In this problem, a subject is faced

Chapitre 3. Exploration/exploitation: protocol and synchronous approach

repeatedly with a choice between different options. After each choice he receives a numerical reward chosen from a probability distribution that depends on the selected action. At each moment, the subject may either select the machine he expects to provide the highest payoff (i.e. to exploit) or another machine in order to improve his estimations (i.e. to explore).

Recent brain imaging studies using functional magnetic resonance imaging (fMRI) or positron emission tomography (PET) have focused on the identification of brain signatures of decision making. These studies link neural activity to external variables observed and manipulated [79, 80]. However, decision making is often based on internal decision variables not directly observable from the subject behaviour. Experiments that aim to study the correlates of these internal variables must be build upon a model of decision based on the observable variables [5, 81, 82]. Studying the differences of activation in the brain during exploratory decisions compared to exploitative decisions requires such a model.

Intracranial recordings in primates and fMRI studies in humans suggest that the anterior cingulate cortex (ACC) could control the balance between exploitative and exploratory behaviour [6]. Recently, Yoshida and Ishii [5] have reported, using fMRI techniques, lateral activation in the prefrontal cortex (PFC) and ACC activation when exploring a virtual maze. Using the same imagery techniques, Daw et al. [7] have shown that activations in the PFC and intraparietal sulcus are correlated with the differences between exploratory decisions and exploitative decisions.

To our knowledge, no electroencephalography (EEG) studies have focused on this issue. Given its time resolution, EEG could give an information that fMRI does not provide, especially in terms of frequency components. However, since both techniques measure different physical phenomena, the detection using one technique does not necessarily lead to the detection by the other one. Thus, *this study aims to determine whether a difference between exploration and exploitation can be detected in scalp EEG.*

Interestingly, studying the contrast between exploration and exploitation requires to estimate hidden variables. In a N -armed bandit paradigm, the exploratory behaviour is indeed not directly deduced from the actions of the subjects which indicates only which machine is selected. An action can be labelled as exploratory only if we can access to the subject's estimate of the mean reward of each machine and how confident he is with this estimate. Therefore a model of how the subject tracks the mean rewards of the machine and of how he bases his selection on this estimate will be necessary. The description of this behavioural model and its fitting to the subjects action is detailed in section 3.3.

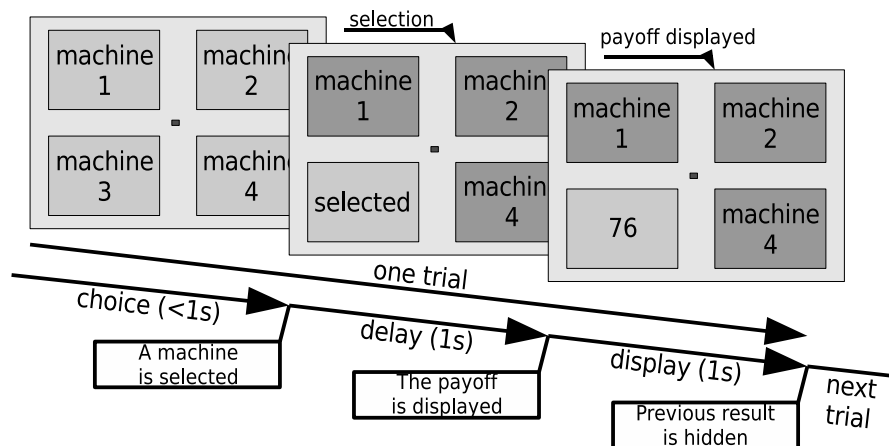


FIGURE 3.1 – Experimental protocol. (a) Visual stimuli : each trial is composed of three phases. (i) *Choice*, the 4 machines are presented. The subject has 1s to select a machine by pressing a key. (ii) *Delay*, once a machine has been chosen, the other machines are deactivated (greyed) and no result is displayed for 1s. (iii) *Display*, the payoff for the selected machine is displayed for 1s.




FIGURE 3.2 – Example of the payoff evolution for the 4 machines represented by different colors during one repetition of the experiment.

3.2 Experimental protocol

3.2.1 4-armed bandit

We adapt the experimental protocol described in Daw et al. [7]. Nine volunteer healthy human subjects (three females and six males; mean age 26) participated in the experiment. Each subject sits in front of a computer screen on which four squares are displayed representing four slot machines (see Figure 3.1). At each trial the subject chooses one machine by pressing a key using the index or middle finger on both hands (left hand for machines 1 and 3, and right hand for machines 2 and 4). One second after the key-press, the payoff of the selected machine is displayed for another second and then, a new trial starts. The subject is asked to select the machines in order to maximise the total gain (i.e. sum of individual payoffs) over a session of 400 trials.

Chapitre 3. Exploration/exploitation: protocol and synchronous approach

The payoff of each machine –a numerical value between 0 and 100– is drawn from a Gaussian distribution whose mean changes slowly from trial to trial. The use of different, non-stationary distributions for each machine requires the subject to regularly update his knowledge about the problem; i.e., he is obliged to explore. Before the experiment, nine examples of the payoff evolution for all machines are graphically shown to the subject (one of these examples is shown in Figure 3.2). The evolution of the payoffs is described in subsection 3.2.2.

Three sessions were recorded for all subjects, with the exception of one who completely only two sessions. During the experiment, subjects fixate on a red dot in the centre of the screen to reduce ocular artifacts. Moreover, they are specifically instructed not to move the arms during the experiment to reduce Electromyogram (EMG) artifacts.

3.2.2 Distribution of the payoff

The payoff $r_{i,k}$ associated with the i -th machine at trial k is drawn from a Gaussian distribution (mean $\mu_{i,k}$, SD : σ_0) and rounded to the nearest integer between [0 , 100]. At each time step, the mean $\mu_{i,k}$ is diffused in a decaying Gaussian random walk,

$$\mu_{i,k+1} = \lambda\mu_{i,k} + (1 - \lambda)\theta + e \quad (3.1)$$

$$\text{where } e \sim \mathcal{N}(0, \sigma_d^2)$$

$$r'_{i,k} \sim \mathcal{N}(\mu_{i,k}, \sigma_0^2) \quad (3.2)$$

$$r_{i,k} = \text{round}(r'_{i,k}) \quad (3.3)$$

where λ, θ controls the random walk of the mean $\mu_{i,k}$ and e corresponds to Gaussian noise with zero mean and standard deviation σ_d . The values of these parameters are reported in Table 3.2. The mean payoff $\mu_{i,0}$ at the beginning of the experiment was set to the result of computing 30 diffusion steps (equation (3.1)) with an initial value of 50.

3.2.3 Signal acquisition and preprocessing

EEG activity was recorded using a Biosemi Active II portable system. The signal was acquired with 64 electrodes according to the 10/20 international system at a sampling rate of 2048 Hz; then filtered by an eighth-order low-pass Chebyshev Type I filter with a cutoff frequency of 205 Hz and downsampled to 512 Hz. The data were filtered in both the forward and reverse directions to remove all phase distortion, effectively doubling the filter order. In addition, Electrooculogram (EOG) was recorded using two electrodes located below and at the outer canthus of the right eye.

Peripheral scalp electrodes (prone to artifact) were not taken into account for the study




FIGURE 3.3 – Electrodes used for analysis : highlighted in red for a total of 50 electrodes.




FIGURE 3.4 – Example of payoff received by a subject during the experiment (black line) and payoff obtained by the subject’s model (red line).

(Figure 3.3). For the remaining electrodes we extracted windows from 1.0s to 0.1s before the key press. The time window was defined to avoid EMG artifacts associated with the finger movements. We subtracted the continuous component using a fourth-order Chebychev high-pass filter with cutoff frequency of 2 Hz and the common average reference was removed. This referencing removes noise signals that are equally spread over the scalp. One subject was removed from the study because there were artifacts in his recording (strong powerline contamination).

3.3 Behavioural model

A behavioural model is required to label each trial corresponding to either an exploratory or exploitative decision. In order to compare our results with those reported

Chapitre 3. Exploration/exploitation: protocol and synchronous approach

TABLE 3.1 – A summary of number of trials

	s1	s2	s3	s4	s5	s6	s7	s8
num trial exploitation	938	852	730	511	686	929	887	828
num trial exploration	57	103	201	96	192	78	106	113
right/exploitation (%)	63.65	65.26	67.26	48.53	48.69	45.43	65.95	64.49
right/exploration (%)	52.63	57.28	51.24	55.21	54.17	52.56	54.72	49.56

using fMRI, we adopt the same behavioural model proposed in [7]. The model, described in subsection 3.3.1, assumes that the subject estimates the mean payoff of each machine using a Bayesian linear Gaussian rule (i.e. a Kalman filter) and, based on these estimations, selects a machine according to a softmax rule. We assume all the subjects share the same model for tracking the payoff means; thus, we compute these parameters using all the available data. In contrast, independent models of machine selection were built per subject to take into account inter-subject variability. Parameters of the model (for both mean tracking and machine selection) are estimated by maximising the model likelihood with respect to the subject's choices. A comparison of the choices taken by a subject and those given by his fitted model can be seen in Figure 3.4.

At any given trial, the behavioural model provides the mean payoff for all machines taking into account previous observations (i.e. the payoff obtained at previous trials). The comparison between the model's estimated payoff for all machines is used to label that trial as either exploration or exploitation. Those trials where the subject selects the machine with the highest estimated mean are labelled as corresponding to exploitative decisions. In order to increase the reliability of the labelling process, a threshold (with value 4) was set when computing the payoff difference between the machines, thus excluding trials that cannot be clearly tagged as either exploration or exploitation. Moreover, only exploratory trials (i.e. those trials where the selected machine does not have the highest estimated payoff) corresponding to a change of machine are kept for further analysis. An average of 22% of the trials are discarded at this stage. The total number of trials used in the analysis is shown in Table 3.1. In order to discard possible bias of the labelling, we also show the percentage of trials labelled as exploration that correspond to a selection with the right hand (machines 2 and 4). The same is done for the trials labelled as exploitation.

3.3.1 Modelling subject estimation and selection

We model the subject's strategy for tracking the payoff of each machine by a Kalman filter and assume that parameters do not change over trials. *After* a machine j has been selected at trial k , the estimated payoff distribution $(\hat{\mu}_{j,k}^{post}, \hat{\sigma}_{j,k}^{post})$ can be updated

TABLE 3.2 – Estimation of the parameters of the behavioural model

	λ	θ	σ_d	σ_0
Real values	0.9836	50	2.8	4
Est. values	0.92	51.37	8.12	N/A

subj.	1	2	3	4	5	6	7	8
β	0.37	0.28	0.19	0.21	0.19	0.29	0.29	0.23

given the received payoff r_k and the estimation *before* the observation $(\hat{\mu}_{j,k}^{pre}, \hat{\sigma}_{j,k}^{pre 2})$,

$$\hat{\mu}_{j,k}^{post} = \hat{\mu}_{j,k}^{pre} + \kappa_k (r_k - \hat{\mu}_{j,k}^{pre}) \quad (3.4)$$

$$\hat{\sigma}_{j,k}^{post 2} = (1 - \kappa_k) \hat{\sigma}_{j,k}^{pre 2} \quad (3.5)$$

$$\text{where } \kappa_k = \frac{\hat{\sigma}_{j,k}^{pre 2}}{\hat{\sigma}_{j,k}^{pre 2} + \hat{\sigma}_0^2} \quad (3.6)$$

Since the subject cannot observe the payoff of the remaining machines, we assume the mean estimation for these machines does not change as a result of the choice. That is,

$$\forall i \neq j \quad \begin{cases} \hat{\mu}_{i,k}^{post} = \hat{\mu}_{i,k}^{pre} \\ \hat{\sigma}_{i,k}^{post} = \hat{\sigma}_{i,k}^{pre} \end{cases} \quad (3.7)$$

Then, the estimations are updated in time according to the diffusion rule seen in (3.1),

$$\hat{\mu}_{i,k+1}^{pre} = \hat{\lambda} \hat{\mu}_{i,k}^{post} + (1 - \hat{\lambda}) \hat{\theta} \quad (3.8)$$

$$\hat{\sigma}_{i,k+1}^{pre 2} = \hat{\lambda}^2 \hat{\sigma}_{i,k}^{post 2} + \hat{\sigma}_d^2 \quad (3.9)$$

We model the choice of the subjects by a softmax rule, i.e. at each trial k the probability of choosing the machine i is :

$$P_{i,k} = \frac{\exp(\beta \hat{\mu}_{i,k}^{pre})}{\sum_j \exp(\beta \hat{\mu}_{j,k}^{pre})} \quad (3.10)$$

The parameters of the behavioural model are estimated by maximising the log-likelihood under constraints ($\hat{\lambda} \in [0, 1]$, $\hat{\theta} \in [0, 100]$, $\hat{\sigma}_d \in [0, 100]$, $\beta \geq 0$). The parameters of the mean payoff tracking ($\hat{\sigma}_0$, $\hat{\lambda}$, $\hat{\theta}$ and $\hat{\sigma}_d$) are shared by all subjects, while the parameter of the selection (β) is specific to each subject. To speed up convergence, parameters $\hat{\sigma}_0$, $\hat{\mu}_{i,0}^{pre}$ and $\hat{\sigma}_{i,0}^{pre}$ are fixed to the real values σ_0 , $\mu_{i,0}$ and $\sigma_{i,0}$. Fixing these last two parameters does not significantly affect the estimation of the others because their

influence vanishes quickly within a few trials. Table 3.2 shows the estimated values of the model, which are consistent with the real values of the machines. Thus the model fits sufficiently well.

3.4 Synchronous approach

3.4.1 Classification

The continuous wavelet transform was computed on all EEG channels using a Morlet mother wavelet on logarithmically scaled frequencies : 7.5 Hz, 10.0 Hz, 13.2 Hz, 17.6 Hz, 23.3 Hz, 31.0 Hz, 41.1 Hz, 54.6 Hz, 72.6 Hz, 96.4 Hz and 128.0 Hz. This scale regularly covers the full spectrum from 7.5 to 128.0 Hz avoiding redundancy among the different frequency channels. The logarithm of the energy of the transformed EEG data (the coefficients of the wavelet transform) were computed.

For every trial the 900 ms length EEG data were divided into 8 intervals of 110 ms each and the mean log transform energy for each interval. For each subject, a LDA classifier was built based on the mean of the log-energy for each couple interval-frequency.

To take into account the different sizes of the data sets for both classes (c.f. Table 3.1) we use the normalised Mathews Correlation Coefficient (MCC) [83] that takes into account the rate of correctly classified samples for each class :

$$MCC = \frac{t_1 t_2 - f_1 f_2}{\sqrt{(t_1 + f_1 t)(t_1 + f_2)(t_2 + f_1)(t_2 + f_2)}} \quad (3.11)$$

where t_1 and f_1 denote correctly and incorrectly classified *exploratory* trials respectively and the same for t_2 and f_2 for the *exploitative* trials. A MCC coefficient of 1 corresponds to perfect classification for both exploration and exploitation samples, while a value of zero corresponds to random performance. If all samples are misclassified, the resulting coefficient is equal to -1.

A 10-fold cross-validation procedure was used to assess the performance of the classification (using the MCC values) : trials are sorted chronologically and divided into 10 consecutive and approximately equal subsets of trials. Each fold in the cross validation procedure is formed by one of these subsets which is used as the testing set and by the remaining four subsets used as the training set. This procedure differs from the usual cross-validation procedure in the sense that the trials in test set do not come from time period that have been used to train the classifiers : the time period use in training and testing are disjoint. Thus this cross-validation further tests the generalization strength of the method. As such, it is a less optimistic measure than the usual cross-validation but also a more realistic one for single trial analysis.

3.4.2 Results

The complete classification results are available in Table 3.3, 3.4 and 3.5. A synthesised figure (see Figure 3.5) shows that the maximum of classification is obtained with a couple of frequency-interval that is different for each subject. In addition, this maximum is barely better than random classification. The best results we can obtain with this synchronous approach of classification are for two subjects (subject 2 and subject 8) for whom some combinations of frequency-interval seem to present classification results significantly better than random. This might indicate some initial success of the approach. However their locations are quite different for both subjects even if they are all below 31.0 Hz. Although there might be some variability among subjects, it is expected a better consistency between subjects to validate the results.

The low classification results added to the absence of consistency of the time-frequency location of the pair providing best classification results simply suggest that with this synchronous classification approach, it is difficult to catch the differences between exploration and exploitation. At best the results of two subjects (s2 and s8) suggest that the discriminant time-frequency locations might be found over time in the lower frequencies.

3.5 Conclusion

We have seen that studying the exploration/exploitation contrast requires the use of a behavioural model of the subject's decision in order to label trials as exploration or exploitation. The fitting of the parameters of the model has been shown to be consistent with the actual statistical parameters of the machines. More refined models have been proposed to explicitly include the subject's need for exploring [61]. However for the sake of simplicity we will keep using this model for later studies on exploratory behaviour.

Since the classification performances using a classical approach (time-locked) are very poor, this suggests that either there is no information in the EEG signal that can help to discriminate the exploratory from the exploitative behaviour, either the information is contained in a way that is inaccessible to a time-locked approach, maybe spread in time over trials. Under this hypothesis, an asynchronous approach where relevant information is not assumed to be time-synchronised over all the trials might catch those differences between exploration/exploitation EEG activity. The next chapter proposed such asynchronous approach for classifying the exploration/exploitation contrast.




FIGURE 3.5 – Classification performance on the test set for each subject per frequency and time interval. The values reported are the mean MCC value over all folds with its standard deviation subtracted. The values that are less than 0 are reported with the same colour as 0 since they all refers to random classification.

3.5. Conclusion

TABLE 3.3 – Classification results (Subject 1 and 2) : reported as MCC values, the standard deviation over the folds in the cross-validation procedure are reported in parenthesis

(a) subject1								
bins	1	2	3	4	5	6	7	8
7.5Hz :	-0.00(0.03)	-0.03(0.07)	-0.05(0.06)	-0.00(0.08)	-0.04(0.04)	-0.07(0.04)	-0.02(0.08)	-0.03(0.02)
10.0Hz :	-0.01(0.08)	0.06(0.08)	0.00(0.07)	-0.01(0.08)	0.03(0.06)	-0.01(0.06)	0.02(0.05)	-0.08(0.02)
13.2Hz :	-0.07(0.03)	-0.03(0.06)	-0.02(0.06)	0.03(0.08)	0.04(0.04)	0.04(0.09)	-0.01(0.08)	0.01(0.07)
17.6Hz :	0.03(0.02)	-0.05(0.08)	-0.05(0.07)	0.00(0.06)	-0.02(0.07)	-0.07(0.03)	-0.01(0.05)	0.05(0.06)
23.3Hz :	0.04(0.06)	-0.01(0.06)	0.02(0.07)	0.03(0.03)	-0.02(0.07)	-0.06(0.03)	-0.07(0.06)	0.00(0.09)
31.0Hz :	0.07(0.10)	0.07(0.07)	-0.04(0.05)	0.06(0.04)	-0.02(0.06)	-0.02(0.03)	0.04(0.04)	0.01(0.05)
41.1Hz :	0.03(0.04)	-0.02(0.02)	-0.04(0.09)	0.01(0.10)	0.00(0.06)	-0.02(0.07)	-0.02(0.10)	0.06(0.04)
54.6Hz :	0.10(0.12)	0.02(0.06)	-0.09(0.08)	0.06(0.11)	0.01(0.09)	-0.00(0.04)	-0.02(0.04)	-0.01(0.06)
72.6Hz :	-0.02(0.05)	0.03(0.06)	0.05(0.11)	0.05(0.08)	-0.06(0.05)	0.01(0.08)	0.06(0.04)	-0.06(0.09)
96.4Hz :	0.02(0.07)	0.11(0.05)	-0.06(0.05)	0.02(0.05)	-0.07(0.05)	-0.04(0.05)	-0.01(0.09)	0.02(0.04)
128.0Hz :	0.09(0.06)	0.06(0.06)	0.01(0.10)	0.07(0.04)	0.01(0.05)	-0.05(0.07)	0.07(0.02)	-0.07(0.04)

(b) subject2								
bins	1	2	3	4	5	6	7	8
7.5Hz :	-0.03(0.07)	0.01(0.08)	0.01(0.09)	-0.06(0.09)	-0.01(0.05)	-0.00(0.06)	-0.04(0.06)	0.03(0.05)
10.0Hz :	0.01(0.06)	0.00(0.10)	0.10(0.09)	0.09(0.05)	0.15(0.05)	0.17(0.02)	0.18(0.07)	0.16(0.09)
13.2Hz :	0.02(0.04)	0.05(0.05)	0.08(0.08)	0.07(0.06)	-0.02(0.08)	0.11(0.03)	0.10(0.09)	0.06(0.05)
17.6Hz :	0.04(0.08)	0.05(0.09)	-0.02(0.09)	0.04(0.03)	0.12(0.07)	0.17(0.08)	0.18(0.09)	0.14(0.07)
23.3Hz :	0.02(0.09)	-0.05(0.06)	0.02(0.05)	0.13(0.05)	0.16(0.09)	0.13(0.06)	0.11(0.08)	0.21(0.12)
31.0Hz :	0.08(0.07)	-0.04(0.09)	0.05(0.05)	-0.04(0.06)	0.03(0.09)	0.06(0.10)	0.01(0.05)	0.05(0.04)
41.1Hz :	-0.03(0.08)	-0.04(0.04)	-0.02(0.08)	-0.02(0.09)	0.01(0.04)	0.06(0.08)	0.08(0.11)	0.06(0.10)
54.6Hz :	-0.05(0.08)	-0.03(0.14)	0.04(0.07)	-0.03(0.05)	0.01(0.05)	0.03(0.05)	0.00(0.02)	-0.03(0.10)
72.6Hz :	0.04(0.05)	-0.02(0.11)	-0.07(0.06)	0.03(0.13)	0.02(0.06)	0.02(0.10)	0.03(0.08)	-0.01(0.08)
96.4Hz :	0.04(0.05)	0.06(0.04)	0.02(0.13)	0.03(0.09)	0.04(0.09)	0.07(0.04)	-0.02(0.05)	0.04(0.02)
128.0Hz :	0.06(0.06)	0.01(0.03)	0.01(0.06)	0.02(0.07)	-0.00(0.03)	-0.04(0.07)	-0.04(0.06)	0.09(0.07)

Chapitre 3. Exploration/exploitation: protocol and synchronous approach

TABLE 3.4 – Classification results (Subject 3, 4 and 5) : reported as MCC values, the standard deviation over the folds in the cross-validation procedure are reported in parenthesis

(a) subject3

bins	1	2	3	4	5	6	7	8
7.5Hz :	-0.01(0.07)	0.04(0.02)	0.05(0.10)	0.08(0.05)	0.04(0.04)	0.02(0.07)	0.02(0.09)	0.08(0.04)
10.0Hz :	-0.04(0.09)	0.03(0.11)	0.08(0.07)	0.06(0.10)	0.08(0.08)	0.04(0.07)	0.06(0.02)	0.03(0.07)
13.2Hz :	0.01(0.08)	0.02(0.05)	0.00(0.11)	0.06(0.09)	0.02(0.01)	0.04(0.08)	0.02(0.12)	-0.03(0.07)
17.6Hz :	0.00(0.07)	0.07(0.06)	0.01(0.11)	0.02(0.04)	0.04(0.05)	-0.07(0.06)	0.00(0.07)	0.00(0.04)
23.3Hz :	0.06(0.08)	0.07(0.07)	0.07(0.02)	0.06(0.04)	-0.00(0.06)	0.03(0.05)	0.07(0.08)	-0.02(0.09)
31.0Hz :	0.02(0.09)	0.05(0.06)	0.05(0.03)	-0.01(0.05)	0.02(0.08)	0.00(0.06)	-0.03(0.10)	-0.04(0.04)
41.1Hz :	0.03(0.06)	-0.04(0.07)	-0.02(0.07)	-0.03(0.04)	-0.04(0.07)	-0.03(0.07)	-0.01(0.06)	0.00(0.06)
54.6Hz :	-0.01(0.07)	-0.02(0.08)	-0.00(0.08)	0.01(0.06)	-0.05(0.09)	0.02(0.07)	-0.02(0.06)	-0.04(0.07)
72.6Hz :	0.00(0.05)	0.04(0.07)	-0.06(0.06)	0.04(0.06)	-0.03(0.03)	-0.06(0.09)	0.00(0.07)	0.11(0.10)
96.4Hz :	-0.02(0.04)	0.02(0.06)	0.02(0.05)	0.01(0.09)	0.00(0.05)	0.02(0.05)	-0.04(0.06)	0.00(0.06)
128.0Hz :	-0.02(0.02)	-0.01(0.06)	0.04(0.06)	0.02(0.06)	0.01(0.08)	0.03(0.07)	-0.02(0.02)	-0.03(0.05)

(b) subject4

bins	1	2	3	4	5	6	7	8
7.5Hz :	0.02(0.06)	0.04(0.10)	0.07(0.09)	0.01(0.09)	0.01(0.07)	0.10(0.05)	0.04(0.02)	0.12(0.09)
10.0Hz :	0.02(0.12)	0.06(0.11)	0.06(0.06)	0.02(0.04)	-0.01(0.07)	0.03(0.08)	0.08(0.03)	0.09(0.08)
13.2Hz :	0.06(0.10)	0.07(0.08)	-0.03(0.06)	0.04(0.05)	0.01(0.05)	0.02(0.09)	0.11(0.09)	0.10(0.12)
17.6Hz :	0.00(0.08)	-0.03(0.05)	0.10(0.07)	0.07(0.04)	0.06(0.04)	-0.03(0.16)	-0.02(0.13)	0.11(0.06)
23.3Hz :	0.01(0.07)	0.04(0.10)	0.05(0.09)	-0.03(0.05)	-0.05(0.12)	-0.09(0.10)	0.03(0.06)	0.03(0.04)
31.0Hz :	-0.01(0.12)	-0.02(0.09)	0.00(0.13)	-0.05(0.09)	0.00(0.10)	0.01(0.12)	-0.01(0.10)	-0.09(0.08)
41.1Hz :	0.02(0.09)	0.08(0.06)	0.09(0.14)	0.03(0.06)	0.08(0.04)	0.05(0.07)	0.00(0.09)	0.08(0.08)
54.6Hz :	0.08(0.07)	-0.04(0.10)	0.02(0.06)	-0.01(0.08)	0.06(0.07)	0.01(0.07)	0.10(0.12)	0.03(0.11)
72.6Hz :	0.06(0.11)	-0.01(0.07)	0.06(0.08)	0.05(0.09)	-0.01(0.04)	0.14(0.07)	-0.04(0.07)	0.05(0.07)
96.4Hz :	0.01(0.07)	0.04(0.04)	0.02(0.11)	0.09(0.08)	-0.02(0.10)	0.10(0.11)	0.01(0.06)	-0.04(0.09)
128.0Hz :	-0.02(0.09)	0.06(0.07)	0.09(0.16)	0.00(0.04)	0.06(0.09)	0.05(0.12)	0.07(0.12)	0.11(0.14)

(c) subject5

bins	1	2	3	4	5	6	7	8
7.5Hz :	0.04(0.06)	-0.04(0.06)	-0.05(0.04)	0.02(0.07)	-0.04(0.07)	0.07(0.06)	0.09(0.05)	0.02(0.05)
10.0Hz :	0.02(0.06)	0.02(0.07)	-0.01(0.03)	0.02(0.05)	0.01(0.07)	0.11(0.09)	-0.00(0.12)	0.03(0.07)
13.2Hz :	0.00(0.10)	-0.04(0.07)	0.08(0.11)	0.06(0.08)	0.08(0.07)	0.04(0.05)	0.06(0.05)	0.03(0.06)
17.6Hz :	0.05(0.07)	0.03(0.11)	0.00(0.09)	0.06(0.07)	-0.03(0.07)	0.10(0.04)	0.07(0.03)	-0.06(0.04)
23.3Hz :	-0.01(0.04)	0.05(0.04)	-0.03(0.08)	-0.05(0.03)	-0.01(0.05)	0.02(0.05)	-0.04(0.10)	0.05(0.08)
31.0Hz :	0.02(0.11)	-0.03(0.04)	0.03(0.06)	0.07(0.12)	-0.05(0.02)	-0.04(0.11)	-0.05(0.05)	-0.02(0.10)
41.1Hz :	0.01(0.09)	-0.04(0.10)	-0.00(0.06)	0.03(0.07)	-0.04(0.07)	-0.02(0.03)	-0.05(0.09)	-0.00(0.09)
54.6Hz :	-0.13(0.07)	-0.02(0.06)	0.00(0.08)	0.00(0.04)	0.02(0.12)	0.02(0.05)	-0.03(0.03)	0.00(0.05)
72.6Hz :	0.01(0.07)	-0.03(0.08)	0.04(0.08)	0.08(0.05)	0.01(0.02)	0.00(0.07)	-0.02(0.06)	0.00(0.09)
96.4Hz :	0.04(0.06)	-0.03(0.07)	-0.03(0.07)	-0.04(0.08)	-0.04(0.11)	-0.05(0.03)	-0.04(0.08)	0.01(0.04)
128.0Hz :	-0.02(0.10)	0.03(0.06)	0.00(0.08)	-0.10(0.08)	0.05(0.05)	-0.04(0.06)	-0.03(0.03)	-0.00(0.12)

3.5. Conclusion

TABLE 3.5 – Classification results (Subject 6, 7 and 8) : reported as MCC values, the standard deviation over the folds in the cross-validation procedure are reported in parenthesis

(a) subject6								
bins	1	2	3	4	5	6	7	8
7.5Hz :	0.01(0.04)	0.06(0.05)	-0.05(0.07)	-0.02(0.07)	-0.05(0.04)	0.03(0.09)	-0.02(0.07)	-0.06(0.05)
10.0Hz :	0.03(0.03)	-0.02(0.05)	-0.03(0.04)	-0.02(0.08)	-0.15(0.04)	0.00(0.09)	-0.02(0.08)	-0.02(0.07)
13.2Hz :	0.01(0.06)	-0.07(0.03)	0.04(0.08)	0.05(0.05)	0.06(0.04)	-0.04(0.08)	0.03(0.05)	-0.04(0.05)
17.6Hz :	-0.06(0.02)	0.01(0.08)	0.01(0.06)	0.01(0.08)	0.02(0.03)	0.06(0.07)	0.06(0.03)	-0.03(0.05)
23.3Hz :	-0.10(0.05)	-0.00(0.02)	0.01(0.05)	-0.02(0.03)	0.03(0.06)	-0.02(0.09)	0.07(0.08)	-0.04(0.07)
31.0Hz :	0.07(0.05)	0.03(0.12)	0.02(0.06)	0.02(0.07)	0.01(0.07)	0.02(0.09)	0.03(0.04)	0.04(0.10)
41.1Hz :	-0.03(0.06)	0.02(0.06)	-0.01(0.04)	0.06(0.04)	-0.01(0.06)	-0.03(0.07)	0.10(0.02)	-0.04(0.05)
54.6Hz :	-0.04(0.06)	-0.04(0.04)	-0.02(0.06)	-0.04(0.09)	-0.02(0.04)	-0.01(0.06)	-0.04(0.04)	-0.01(0.07)
72.6Hz :	0.05(0.06)	-0.02(0.04)	-0.05(0.04)	0.02(0.03)	-0.02(0.06)	0.02(0.04)	0.02(0.05)	0.00(0.10)
96.4Hz :	0.01(0.07)	-0.03(0.03)	-0.04(0.04)	-0.02(0.05)	-0.04(0.07)	0.03(0.04)	-0.03(0.04)	0.00(0.05)
128.0Hz :	-0.07(0.06)	0.05(0.07)	-0.04(0.09)	0.02(0.03)	0.03(0.07)	0.01(0.08)	-0.03(0.09)	-0.04(0.14)

(b) subject7								
bins	1	2	3	4	5	6	7	8
7.5Hz :	0.01(0.05)	0.04(0.08)	0.05(0.04)	0.07(0.08)	-0.02(0.05)	0.00(0.06)	0.06(0.07)	0.04(0.02)
10.0Hz :	0.06(0.06)	0.05(0.07)	-0.02(0.07)	-0.02(0.13)	0.07(0.05)	0.01(0.11)	-0.04(0.04)	0.04(0.06)
13.2Hz :	0.03(0.06)	0.02(0.05)	-0.00(0.06)	0.06(0.07)	0.06(0.04)	0.07(0.07)	0.10(0.08)	0.12(0.09)
17.6Hz :	0.00(0.05)	-0.03(0.05)	-0.01(0.09)	0.05(0.04)	0.04(0.08)	0.07(0.06)	0.02(0.08)	0.02(0.06)
23.3Hz :	0.00(0.06)	0.03(0.06)	0.02(0.03)	0.03(0.06)	-0.06(0.06)	0.05(0.07)	0.01(0.05)	0.04(0.09)
31.0Hz :	-0.01(0.14)	-0.02(0.04)	0.01(0.08)	-0.02(0.07)	0.02(0.09)	0.03(0.04)	0.04(0.03)	-0.00(0.08)
41.1Hz :	-0.04(0.05)	0.07(0.07)	0.05(0.09)	-0.02(0.07)	-0.00(0.06)	0.01(0.03)	0.04(0.06)	-0.03(0.05)
54.6Hz :	-0.04(0.06)	-0.02(0.06)	0.03(0.10)	0.02(0.06)	0.09(0.06)	0.05(0.03)	0.06(0.10)	0.01(0.07)
72.6Hz :	-0.00(0.04)	-0.03(0.10)	0.00(0.07)	0.04(0.05)	-0.06(0.08)	0.02(0.08)	-0.02(0.04)	0.01(0.09)
96.4Hz :	-0.01(0.08)	0.01(0.02)	0.00(0.05)	-0.07(0.03)	0.02(0.03)	0.02(0.02)	-0.00(0.05)	-0.05(0.04)
128.0Hz :	-0.01(0.05)	-0.01(0.04)	0.01(0.11)	-0.03(0.04)	-0.05(0.07)	-0.04(0.04)	0.06(0.04)	-0.04(0.03)

(c) subject8								
bins	1	2	3	4	5	6	7	8
7.5Hz :	0.11(0.10)	0.08(0.05)	0.17(0.07)	0.11(0.09)	0.05(0.08)	0.15(0.09)	0.05(0.14)	0.12(0.06)
10.0Hz :	0.05(0.06)	0.03(0.07)	0.05(0.08)	0.10(0.07)	0.11(0.15)	0.12(0.06)	0.10(0.08)	0.10(0.04)
13.2Hz :	0.12(0.10)	0.11(0.04)	0.10(0.05)	0.08(0.04)	0.11(0.06)	0.04(0.07)	0.08(0.06)	0.07(0.06)
17.6Hz :	0.10(0.03)	-0.01(0.06)	0.09(0.04)	0.08(0.09)	0.15(0.02)	0.17(0.07)	0.07(0.03)	0.11(0.05)
23.3Hz :	0.02(0.03)	0.03(0.07)	0.04(0.08)	0.10(0.10)	0.07(0.07)	0.17(0.12)	0.08(0.07)	0.09(0.09)
31.0Hz :	0.15(0.06)	0.01(0.04)	0.07(0.07)	0.11(0.08)	0.02(0.04)	0.06(0.06)	0.06(0.03)	0.04(0.05)
41.1Hz :	-0.07(0.03)	0.01(0.09)	0.02(0.04)	-0.04(0.02)	0.06(0.08)	0.02(0.07)	0.03(0.08)	-0.02(0.06)
54.6Hz :	0.04(0.06)	-0.00(0.05)	0.01(0.04)	-0.03(0.10)	0.00(0.04)	0.01(0.06)	-0.07(0.05)	0.01(0.07)
72.6Hz :	-0.00(0.07)	-0.05(0.04)	-0.01(0.04)	-0.01(0.03)	0.03(0.06)	-0.01(0.09)	0.02(0.04)	-0.00(0.06)
96.4Hz :	-0.00(0.05)	0.01(0.05)	0.02(0.08)	0.01(0.02)	0.06(0.07)	-0.02(0.09)	0.03(0.08)	0.01(0.05)
128.0Hz :	0.00(0.06)	0.08(0.07)	0.07(0.08)	0.10(0.07)	-0.03(0.07)	0.08(0.07)	-0.01(0.07)	0.05(0.04)

4 Asynchronous approach of Exploration/Exploitation

Studying the EEG correlates of exploratory behaviour raises the problem of knowing when a decision is made. Because of the complexity of the task, a decision is unlikely to be made at the same instant over all the trials. This makes the search of difference of activity particularly difficult since we cannot synchronise the EEG signals well across trials as it has been assumed in the previous chapter. The discriminant analysis employed in this study has been developed specifically in order to address this issue.

In order to deal with the asynchronicity of the data, we propose a method which studies the patterns of the EEG signal over the scalp by considering the distributions of each class (exploration and exploitation) as a “bag of time-samples”, meaning that any information about the time location of the sample within the trial is discarded. The rationale of this approach is to detect regions in the feature space that correspond to the most discriminant samples irrespective of the event occurrence within each trial. This chapter, describes the feature extraction mechanism and then the feature selection method and how it performs on exploration/exploitation data set.

4.1 Asynchronous discriminative approach

4.1.1 Classification context

We want to address the problem of asynchronous trials classification. By asynchronous, we consider that we cannot find any strong time-cue in the trials indicating the beginning the temporal location of the signal of interest.

This implies that the relevant part of the trial is not located at the same instant from one trial to another (see Figure 4.1). In the context of classification of exploratory decision, these assumptions can be justified by the fact that while making their decision, subjects can hesitate, considering alternatives for an undetermined amount of time, thus causing variability of the time course of the decision from one trial to another.




FIGURE 4.1 – example (synthetic data) of trials of 2 classes : trials of class blue are represented in (a) and in (b) of class red. The discriminant parts in each trial are highlighted in their corresponding colour. In y-axis : EEG potentials, x-axis : time (arbitrary units).

Because of the absence of strong time-cue, usual techniques based on averaging the EEG response across trials may not be suitable as we have seen in Chapter 3.

Finally, it is assumed that the distribution of the features of each class are stationary. This means that the features characterising each class will not change with time and that the performance of the classification procedure should not change significantly by changing the partition of the data used to estimate the parameters of the model.

4.1.2 Feature extraction : canonical variates analysis

In order to detect the aforementioned regions that correspond to the most discriminant sample, we need first to find a transformation of the feature space that separate the best between the two classes. Under the hypothesis of linearity and that the features follows a Gaussian distribution, the Canonical Variates Analysis (CVA) (also referred to as multidimensional discriminant analysis, MDA) provides such transformation. The original feature space is obtained by computing the power signal for all electrodes per frequency band as described in section 3.2.3. The new features are then extracted using CVA [85, 86]. The CVA computes the subspace on which the linear projection of the data maximizes the discriminability of the classes. Such analysis provides the projection of each time sample onto a uni-dimensional space (since there is two classes), as well as a measure of the Discriminant Power (DP) of each electrode. The DP ranks the electrodes according to the correlation between the original signal and the features in the projected space. For our purpose, the discriminant power provides information about which electrodes on the scalp EEG convey more information to distinguish between exploratory and exploitative trials.

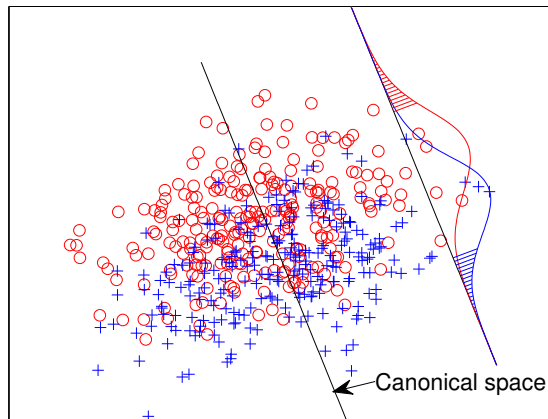


FIGURE 4.2 – Schematic representation of the processing method. Red (o) and blue (+) symbols correspond to samples of two different classes. The distributions of the projected samples on the canonical space of both class are also presented. The hatched red and blue areas on the distribution function show the frame sets for the red (o) and blue (+) classes.

4.1.3 Detection of discriminative samples

We propose a method that relies on the detection of the most discriminative samples for each class based on the sample distribution on the canonical projection. Under this approach, we attempt to recognize informative phenomena by identifying samples that lie on non-overlapping regions of the canonical space. We then perform classification solely based on those samples. Following Freeman's theory, we will use the term *frame* to denote these informative samples [87, 86].

Using this approach, the canonical transformation is computed for each subject on a subset of the data (i.e. train set). A sample will be considered as a frame if it lies on the non-overlapping tail of the samples feature distribution in the canonical space (c.f. Figure 4.2). In this study we use the opposing 5-percentiles of the class distributions as threshold to define the frame sets of the corresponding classes.

In order to evaluate whether a new test sample corresponds to a frame, its canonical projection is compared to these thresholds, i.e. only these samples are taken into account for classification. The classification of a trial uses a voting scheme based on the number of identified frames of each class. In case of equality or if no frame has been detected, the trial is marked as unknown.

Separate classifiers are built for each frequency band i.e. specific projections are computed per frequency band (for a total of 11 bands). In addition, a combined classifier is built using the data from several selected frequency bands. A frequency band is selected for this combined classifier if the classifier based on the single band has a MCC higher than 0.16 on the training set (equivalent to a classification accuracy of 58%




FIGURE 4.3 – Discriminant power (DP) of the electrodes activity using : (a) data of all subjects, (b)–(d) data of three different subjects.

with equally sized sets). The choice of this threshold is based on a trade-off between selecting more than one band and a reasonable confidence that the classifier based on this single band performs better than random. To classify a trial, this combined classifier attributes a label to a trial according to the total number of frames detected in *all selected* bands of each class. In other words, to classify a trial using the combined classifier, the total number for frames detected in all selected frequency bands is computed for each class and the class is determined through majority voting.

4.2 Results

4.2.1 Discriminant analysis

We computed the canonical space projection and the electrodes DP as described in subsection 4.1.2 using data for all subjects altogether (representing 3.8×10^6 time samples) as described in subsection 4.1.2. Figure 4.3a shows the electrode discriminant power for the different frequency bands. The figure shows that the most informative scalp areas (high DP) correspond to left frontal and bilateral parietal : 11 of the 15 most discriminant electrodes are located in these areas.

The same analysis was done independently per subject (using 4.1×10^5 time samples per subject on average, standard deviation : 5.4×10^4). See Figures 4.3b, 4.3c and 4.3d for examples. Comparing among subjects, we report a high inter-subject variability in the precise location of the source of discrimination. But consistently with the global analysis, left frontal and bilateral parietal areas are often discriminant : left frontal electrodes were found to be in the 15 most discriminant electrodes for 7 subjects, right parietal area for 7 subjects and left parietal area for 6 subjects. In addition, right frontal electrodes were found among the 15 most discriminant for all the subjects. Finally, bilateral frontal and parietal activities seem to be, in spite of the inter-subject variability, the most discriminant activities of the exploratory behavior.

From the analysis on all EEG data of all subjects, we can observe that frontal and parietal electrodes are not discriminant in the same frequency bands. Discriminant frontal activity is mainly found between 7.5 and 41 Hz while parietal activity is found in the full spectrum of analysis.

4.2.2 Classification

A 10-fold cross-validation procedure was used to assess the performance of the classification : trials are sorted chronologically and divided into 5 consecutive and approximately equal subsets of trials. Each fold in the cross validation procedure is formed by one of these subsets which is used as the testing set and by the remaining four subsets used as the training set. As explained in subsection 3.4.1, this procedure differs from the usual cross-validation procedure in the sense that the trials in the test set do not come from time epochs that have been used to train the classifiers. The classification performances reported in the study are the average and standard deviation of the MCC value of all folds.

The classification performances on train and test sets are reported in Table 4.1. For all subjects, the best performances are mainly obtained below 23.3 Hz (low beta and alpha band). Considering the performance on the test set, we have observed that the classifiers based on the low frequency bands also perform the best. For 6 subjects,

Chapitre 4. Asynchronous approach of Exploration/Exploitation

TABLE 4.1 – Classification accuracies per subjects (EEG signal). The performances are reported as “mean MCC±SD(UNK)” of the 10-fold cross validation where UNK is the percentage of trials labelled as unknown by the classifier. CC refers to the combined classifiers. The value in bold corresponds to the maximum performance accross frequency per subject. See text for details.

(a) training set

freq.	s1	s2	s3	s4	s5	s6	s7	s8
7.5	0.31±0.07(34)	0.26±0.02(28)	0.24±0.02(21)	0.29±0.04(31)	0.31±0.02(32)	0.18±0.02(27)	0.21±0.01(19)	0.31±0.02(47)
10.0	0.22±0.01(26)	0.41±0.03(59)	0.26±0.02(25)	0.24±0.03(24)	0.22±0.02(20)	0.15±0.02(21)	0.27±0.02(23)	0.32±0.02(45)
13.2	0.20±0.02(21)	0.26±0.03(40)	0.21±0.02(10)	0.26±0.02(18)	0.29±0.02(17)	0.19±0.02(21)	0.21±0.03(13)	0.25±0.02(29)
17.6	0.17±0.03(6)	0.25±0.02(18)	0.23±0.01(4)	0.19±0.03(8)	0.22±0.02(8)	0.17±0.01(12)	0.20±0.02(6)	0.36±0.02(28)
23.3	0.14±0.03(5)	0.36±0.01(21)	0.21±0.03(3)	0.24±0.02(4)	0.17±0.03(3)	0.16±0.02(9)	0.11±0.02(3)	0.31±0.02(19)
31.0	0.16±0.02(4)	0.17±0.02(4)	0.17±0.03(2)	0.19±0.02(3)	0.23±0.03(2)	0.14±0.02(6)	0.21±0.02(4)	0.26±0.02(5)
41.1	0.13±0.02(1)	0.12±0.01(2)	0.13±0.02(7)	0.18±0.02(4)	0.03±0.07(3)	0.14±0.02(3)	0.18±0.03(4)	0.08±0.02(1)
54.6	0.13±0.03(4)	0.16±0.02(4)	0.17±0.01(9)	0.18±0.03(18)	0.09±0.07(8)	0.12±0.03(4)	0.18±0.02(11)	0.11±0.03(5)
72.6	0.10±0.02(1)	0.09±0.01(2)	0.11±0.02(3)	0.24±0.02(2)	0.12±0.03(2)	0.11±0.01(3)	0.19±0.02(2)	0.12±0.02(1)
96.4	0.15±0.02(2)	0.15±0.01(2)	0.12±0.03(10)	0.22±0.02(2)	0.07±0.07(5)	0.11±0.02(4)	0.20±0.02(2)	0.08±0.03(2)
128	0.07±0.02(1)	0.09±0.02(1)	0.01±0.10(4)	0.14±0.02(1)	0.02±0.08(1)	0.10±0.01(1)	0.19±0.02(1)	0.11±0.01(2)
CC	0.30±0.04(1)	0.36±0.03(0)	0.36±0.02(0)	0.39±0.02(0)	0.43±0.02(0)	0.24±0.01(1)	0.31±0.02(0)	0.38±0.01(0)

(b) test set

freq.	s1	s2	s3	s4	s5	s6	s7	s8
7.5	0.08±0.14(36)	0.15±0.10(26)	0.13±0.08(19)	0.08±0.12(30)	0.18±0.10(30)	-0.06±0.11(28)	0.05±0.14(22)	0.20±0.15(45)
10.0	0.07±0.14(28)	0.30±0.22(59)	0.16±0.08(24)	0.12±0.09(23)	0.10±0.10(18)	-0.00±0.10(19)	0.15±0.13(21)	0.27±0.09(45)
13.2	0.01±0.14(19)	0.18±0.13(38)	0.12±0.11(10)	0.15±0.16(17)	0.19±0.12(18)	0.03±0.12(21)	0.04±0.09(14)	0.16±0.11(26)
17.6	0.01±0.13(6)	0.19±0.08(18)	0.13±0.11(5)	0.07±0.10(8)	0.14±0.12(9)	0.07±0.08(12)	0.00±0.17(8)	0.31±0.11(28)
23.3	-0.02±0.09(4)	0.30±0.10(21)	0.11±0.15(4)	0.05±0.13(3)	0.02±0.15(3)	0.01±0.10(9)	-0.09±0.12(4)	0.26±0.15(19)
31.0	0.06±0.12(3)	0.07±0.13(5)	-0.01±0.13(2)	0.02±0.13(3)	0.08±0.10(3)	0.05±0.09(5)	0.06±0.15(7)	0.16±0.14(5)
41.1	0.06±0.11(1)	0.04±0.08(1)	0.02±0.06(5)	-0.04±0.11(4)	-0.06±0.17(1)	0.06±0.10(4)	-0.02±0.15(4)	-0.02±0.13(1)
54.6	0.08±0.13(3)	0.08±0.11(4)	-0.01±0.06(8)	-0.01±0.10(20)	-0.08±0.19(6)	0.06±0.08(4)	-0.00±0.17(5)	0.02±0.09(3)
72.6	0.02±0.10(1)	0.05±0.13(2)	0.04±0.08(3)	0.12±0.14(3)	0.04±0.08(1)	0.03±0.14(2)	0.02±0.12(2)	-0.01±0.13(1)
96.4	0.09±0.15(2)	0.10±0.12(1)	-0.04±0.05(12)	0.05±0.17(2)	0.01±0.08(2)	0.08±0.10(4)	0.05±0.16(2)	-0.00±0.07(1)
128	0.03±0.13(1)	0.06±0.10(2)	0.01±0.03(3)	0.04±0.12(1)	-0.02±0.04(0)	0.05±0.08(1)	0.01±0.16(2)	0.00±0.09(2)
CC	0.05±0.14(1)	0.29±0.09(0)	0.16±0.13(0)	0.11±0.19(0)	0.21±0.14(0)	0.03±0.08(1)	0.04±0.15(0)	0.29±0.13(0)

TABLE 4.2 – Classification accuracies per subjects (EOG signal). The performances are reported as “mean MCC \pm SD(UNK)” of the 10-fold cross validation where UNK is the percentage of trials labelled as unknown by the classifier. CC refers to the combined classifiers.

(a) train set								
freq.	s1	s2	s3	s4	s5	s6	s7	s8
7.5	0.34 \pm 0.10(44)	-0.01 \pm 0.02(39)	0.04 \pm 0.01(36)	0.00 \pm 0.01(38)	0.04 \pm 0.01(32)	0.03 \pm 0.01(37)	0.10 \pm 0.03(27)	0.04 \pm 0.01(40)
10.0	0.19 \pm 0.09(35)	0.06 \pm 0.06(31)	0.03 \pm 0.04(27)	-0.02 \pm 0.02(31)	0.01 \pm 0.02(25)	0.00 \pm 0.02(25)	0.08 \pm 0.02(28)	0.09 \pm 0.02(27)
13.2	0.08 \pm 0.05(19)	-0.05 \pm 0.07(15)	0.00 \pm 0.01(16)	0.03 \pm 0.02(15)	0.09 \pm 0.02(14)	0.02 \pm 0.07(15)	0.10 \pm 0.03(11)	0.01 \pm 0.05(17)
17.6	0.03 \pm 0.04(8)	0.01 \pm 0.01(9)	0.02 \pm 0.01(10)	0.05 \pm 0.02(12)	0.04 \pm 0.02(11)	0.05 \pm 0.02(8)	0.00 \pm 0.02(9)	0.00 \pm 0.03(12)
23.3	-0.01 \pm 0.01(7)	0.05 \pm 0.01(4)	0.02 \pm 0.03(6)	0.03 \pm 0.04(7)	0.05 \pm 0.02(10)	0.05 \pm 0.01(6)	0.03 \pm 0.03(8)	0.03 \pm 0.01(5)
31.0	-0.01 \pm 0.01(4)	-0.01 \pm 0.01(4)	0.00 \pm 0.03(3)	0.03 \pm 0.04(3)	0.09 \pm 0.03(3)	0.06 \pm 0.01(4)	0.05 \pm 0.04(5)	0.02 \pm 0.03(3)
41.1	-0.02 \pm 0.01(1)	0.00 \pm 0.01(2)	0.00 \pm 0.03(1)	0.02 \pm 0.02(2)	0.08 \pm 0.02(3)	0.06 \pm 0.01(3)	0.05 \pm 0.02(3)	0.02 \pm 0.04(1)
54.6	0.00 \pm 0.01(1)	0.00 \pm 0.02(1)	-0.04 \pm 0.03(2)	0.01 \pm 0.02(1)	0.06 \pm 0.02(1)	0.06 \pm 0.01(2)	0.03 \pm 0.06(2)	0.06 \pm 0.01(2)
72.6	-0.03 \pm 0.03(2)	0.03 \pm 0.01(1)	0.04 \pm 0.06(1)	0.03 \pm 0.06(1)	0.04 \pm 0.02(1)	0.05 \pm 0.01(1)	0.05 \pm 0.02(1)	0.02 \pm 0.07(1)
96.4	-0.02 \pm 0.02(1)	0.01 \pm 0.02(1)	-0.03 \pm 0.11(1)	0.02 \pm 0.03(1)	0.07 \pm 0.02(1)	0.07 \pm 0.01(1)	0.02 \pm 0.02(1)	0.01 \pm 0.02(1)
128.0	-0.01 \pm 0.03(1)	0.05 \pm 0.02(1)	-0.01 \pm 0.07(1)	0.01 \pm 0.02(1)	0.07 \pm 0.02(1)	0.05 \pm 0.01(1)	-0.00 \pm 0.05(1)	0.03 \pm 0.02(1)
CC	0.21 \pm 0.14(20)	-0.06 \pm 0.04(15)	-0.10 \pm 0.02(4)	-0.04 \pm 0.02(14)	0.01 \pm 0.02(23)	-0.02 \pm 0.03(21)	-0.02 \pm 0.04(4)	-0.03 \pm 0.03(7)

(b) test set								
freq.	s1	s2	s3	s4	s5	s6	s7	s8
7.5	0.18 \pm 0.34(44)	-0.04 \pm 0.12(39)	0.04 \pm 0.08(36)	-0.02 \pm 0.08(38)	0.04 \pm 0.10(32)	0.02 \pm 0.13(36)	0.04 \pm 0.16(27)	0.03 \pm 0.12(39)
10.0	0.09 \pm 0.24(35)	0.01 \pm 0.18(31)	0.05 \pm 0.13(27)	-0.02 \pm 0.14(31)	0.02 \pm 0.10(23)	-0.08 \pm 0.11(25)	0.05 \pm 0.14(28)	0.08 \pm 0.16(27)
13.2	0.01 \pm 0.15(19)	-0.04 \pm 0.18(15)	-0.02 \pm 0.12(14)	0.01 \pm 0.16(16)	0.09 \pm 0.09(15)	-0.02 \pm 0.13(15)	0.08 \pm 0.13(11)	-0.03 \pm 0.10(17)
17.6	-0.05 \pm 0.09(8)	-0.02 \pm 0.11(9)	0.01 \pm 0.10(10)	0.01 \pm 0.14(11)	0.03 \pm 0.10(11)	0.03 \pm 0.07(8)	-0.00 \pm 0.09(9)	-0.03 \pm 0.11(11)
23.3	-0.00 \pm 0.07(7)	-0.01 \pm 0.07(4)	0.04 \pm 0.07(6)	0.02 \pm 0.19(8)	0.01 \pm 0.12(8)	0.02 \pm 0.12(6)	0.07 \pm 0.11(8)	0.01 \pm 0.10(5)
31.0	0.03 \pm 0.10(4)	-0.00 \pm 0.12(4)	0.01 \pm 0.10(3)	0.01 \pm 0.13(3)	0.08 \pm 0.09(3)	0.05 \pm 0.10(4)	0.08 \pm 0.08(5)	0.00 \pm 0.10(3)
41.1	-0.02 \pm 0.10(1)	-0.02 \pm 0.09(3)	0.04 \pm 0.11(1)	-0.04 \pm 0.15(2)	0.06 \pm 0.09(4)	0.06 \pm 0.09(3)	0.05 \pm 0.09(3)	-0.02 \pm 0.13(1)
54.6	0.01 \pm 0.10(1)	-0.03 \pm 0.13(1)	-0.01 \pm 0.14(1)	-0.01 \pm 0.09(1)	0.02 \pm 0.08(1)	0.06 \pm 0.09(2)	0.03 \pm 0.11(2)	0.01 \pm 0.08(2)
72.6	-0.01 \pm 0.13(2)	0.01 \pm 0.06(2)	-0.00 \pm 0.10(1)	-0.02 \pm 0.15(1)	-0.03 \pm 0.07(1)	0.05 \pm 0.10(1)	0.06 \pm 0.11(1)	0.01 \pm 0.13(1)
96.4	-0.00 \pm 0.13(1)	-0.04 \pm 0.09(1)	-0.02 \pm 0.08(1)	-0.01 \pm 0.17(2)	0.04 \pm 0.08(0)	0.07 \pm 0.09(0)	0.07 \pm 0.09(1)	0.03 \pm 0.11(1)
128.0	-0.03 \pm 0.09(0)	0.02 \pm 0.11(1)	-0.02 \pm 0.08(1)	0.05 \pm 0.14(3)	0.01 \pm 0.10(0)	0.05 \pm 0.10(1)	-0.03 \pm 0.13(1)	-0.02 \pm 0.08(1)
CC	0.07 \pm 0.27(19)	-0.04 \pm 0.15(14)	-0.02 \pm 0.07(4)	-0.04 \pm 0.16(16)	0.04 \pm 0.12(21)	-0.05 \pm 0.16(19)	-0.05 \pm 0.12(4)	-0.04 \pm 0.09(6)

these classifiers produce results significantly better than random (the mean MCC value minus the standard deviation over folds is higher than 0). Despite having a high performance, classifiers centered around 7.5 and 10 Hz have a high rejection rate (more than 20% of the trials are classified as unknown). On average, 6.4 frequency bands have been selected to build the combined classifier for each subject.

For all subjects, performance of the combined classifier is comparable to the one of the best classifiers based on a single frequency band. However, its rejection rate is very low on the two sets (train and test) and for four subjects (s2, s3, s5 and s8), its performances are significantly better than random while it rejects only 0.45% of the trials on average. The combination of selected frequency bands is mostly below 31 Hz. Since the rejection rate of the combined classifier is significantly lower than that of any other classifier, this shows that it does not rely on the data from only one band but instead, its performance is effectively based on the combination of several bands.

The same analysis was done using EOG data instead of EEG data to rule out signal contamination due to eye or facial movements. No classifier based on any single frequency band nor any combined classifier performed better than random for every subject even on the training set (see Table 4.2). This shows that no EOG artifacts had contaminated the EEG analysis.

To compare the cross-validation procedure used in this study with the traditional cross validation (that is, by mixing in the training and test set trials from different time instants), we report in Figure 4.4 the performance –estimated by the two methods– of the combined classifier for each subject. The traditional cross-validation presents slightly better classification accuracies with a smaller variance in the set of exploitation trials.

4.3 Discussion

In this study, we were able to find scalp EEG activity discriminant between exploratory and exploitative behaviour. This activity was mainly located in bilateral frontal and parietal areas, which is consistent with the intracranial activity reported by previous studies using fMRI. Daw et al. [7], using the same protocol, have reported activity in prefrontal cortex (PFC) and in parietal areas as discriminative of exploratory behaviour. Similarly, Yoshida and Ishii [5] found lateral PFC and parietal activity, although they suggest the latter activity to be related to the maintenance of the spatial information of the task.

Our results show that the discriminant frontal activity is mainly found in the alpha and beta band, whereas the parietal activity is not restricted to a particular sub-band. Moreover, the performance of classifiers based on single frequencies confirms the importance of the alpha and beta band for recognition of exploratory behaviour from EEG. However, the increased performance of the combined classifiers suggests that discriminant EEG correlates are unlikely to be restricted to a single frequency band.

We achieved test classification performance above random levels in four of the subjects (s2, s3, s5, s8). In particular, classifiers using single frequency bands equal or below 23.3 Hz have the highest performance, although they have a high rejection rate. The use of combined classifiers increases the classification performance while dramatically reducing the number of trials labelled as unknown. It should also be noticed that we measure the classification performance using a cross-validation method that preserves the temporal order of the data –as opposed to traditional data partitioning for cross-validation– thus giving a better approximation of the method’s ability to generalise non-stationary data (c.f. Figure 4.4). In addition, classifier generalisation capabilities may also be affected by the difference in the number of trials corresponding to each class (c.f. Table 3.1 in chapter 3).

It has been argued that EEG signals above 20 Hz are highly affected by EMG artifacts [88, 89]. However, distant EMG activity (i.e. generated by limb movements) is unlikely to bias the results since the time window used in the study is chosen before the actual movement. Moreover, the number of trials corresponding to movements done by right and left hand are roughly equal in both classes (see Table 3.1). EMG artifacts from

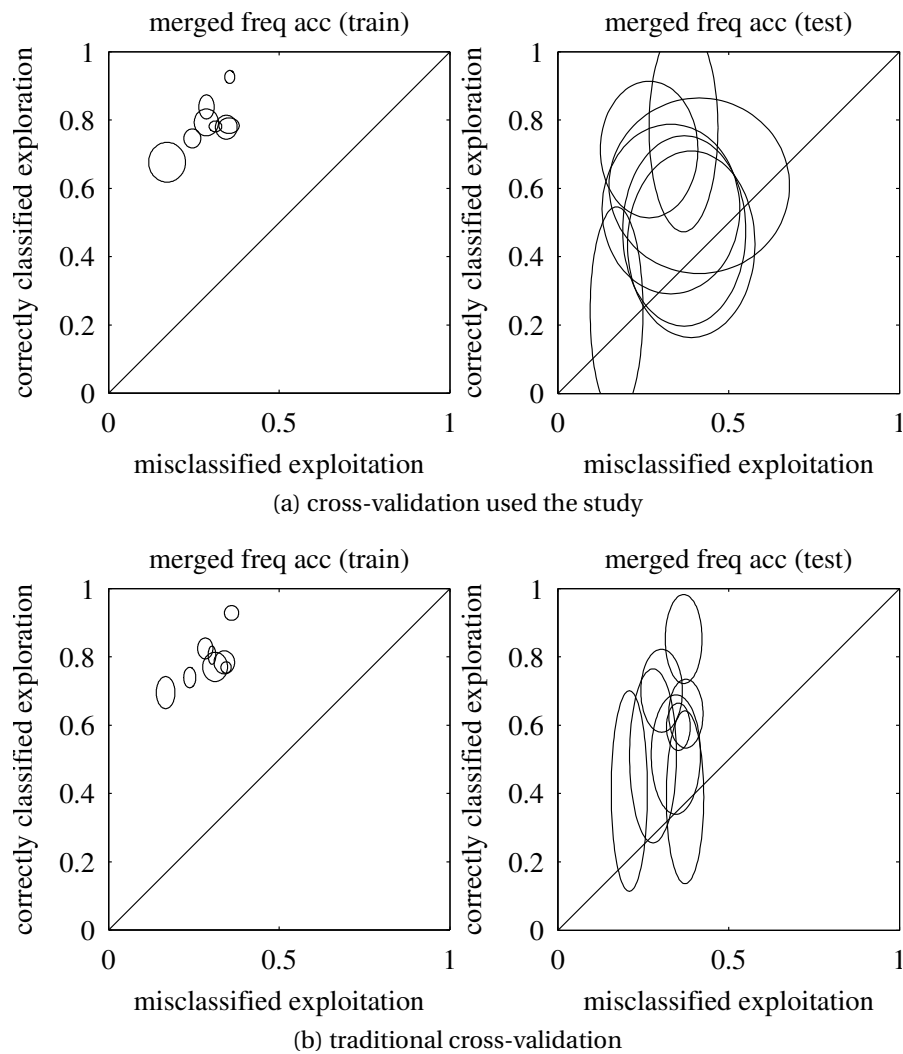


FIGURE 4.4 – Performance of the combined classifier using : (a) the cross validation procedure used in the study (preserving data time order), (b) the traditional cross validation. Each plot shows the exploration classification accuracy versus the exploitation error rate. A perfect classifier would get a point in the upper-left corner of the plot and random classification results would lie on the diagonal line. In addition, the standard deviations of the estimation of the accuracy and the error rate in cross-validation procedure are reported, thus each classification performance is represented by an ellipse.

eye, face or neck movements could have affected the EEG signal, but the classification results using EOG data and visual inspection of the signal energy distribution across different frequencies contradict this hypothesis.

In this study, we assume that discriminant features are not synchronised to any observable event (e.g. key press to select a machine). While a synchronous classification method failed to classify the the same data, as we have seen in the previous chapter, the use of a detection approach allowed us to overcome the issue of the non-time-locked signal, as suggested by the classification results. However, the applied technique is not able to capture temporal relationships between several discriminative patterns of activity (potentially in different frequency bands). These might be even more relevant since our results show that taking into account combined information from different bands helps to better discriminate the two types of decision.

One issue with the current approach is the definition of the informative sets. As of now, we define the tail of a Gaussian distribution as informative set. This leads to two issues. First the threshold must be chosen. In this study, the choice of the 5-percentile was suitable for the data set, but it is a meta-parameter that is problem specific and in absence of any knowledge about it, in the general case, it should be trained.

The other issue with the percentile definition of informative set is that it is sensitive to the variation of the threshold. Under the assumption of Gaussian distribution, most of the samples recognised as informative are located close to the threshold (See Figure 4.2). If the threshold is slightly over- or under-estimated, this will results in many uninformative samples incorrectly recognised as informative or the reverse : small variation of the threshold leads to bigger variation in the informative sets. This issue could be alleviated by using a different model for the informative sets. For example, we could model it as a Gaussian distribution itself, not the tail of Gaussian distribution.

4.4 Conclusion

To sum up, this study has shown that scalp EEG conveys discriminative information between exploratory and exploitative decisions. The spatial pattern of these signals (i.e. most discriminative electrodes) was found to be consistent with previous fMRI studies. Moreover, using a feature detection approach we achieve classification performance above random levels in four out of eight subjects. However, the level of performance of classification shows the weaknesses of the method. While keeping the concept of informative samples, we might improve the model using a suitable model, for example by directly modelling the informative set in a probabilistic way. This will be addressed in the next chapter.

5 Bayesian approach

In the previous chapter, we have designed a method that was able to classify data with trials that were not time locked by defining which samples in the trials should help to classify it as one class or another. However the method showed its limitations in the definition of those subset of samples that help classification.

In this chapter, we will adapt the previous method by modelling, in a probabilistic way, those informative and non-informative sets thus using a generative approach. Because of the number of hidden parameters we introduce with it, we will use a Bayesian approach to learn all of the model parameters. However since the method has been less tested, it is necessary to validate the results of convergence in a controlled case. If we kepted using the problem of classifying the exploration/exploitation contrast, we would face the issue of validating the results : as reminded in its introduction in section 3.1, besides the work we have done, up to our knowledge, no EEG studies have focused on the exploratory behaviour paradigm.

Thus, in order to test this new approach, we will use a widely known brain signal on which references using EEG is more abundant : the ErrP. It is not only an evoked potential whose waveform has been studied extensively and it is known which electrodes are supposed to elicit the strongest response [90], but the ErrP is a synchronous EEG signal (or at least close to be). It might be a little bit paradoxal to test a method that handle classification of asynchronous signal with a synchronous signal. However it is actually a good idea since using such signal, we know where to *expect* the informative part. In addition, a synchronous signal can always to be turned into an asynchronous one by introducing an artificial jitter different for each trial (but in a controllable way). So using the ErrP provides a context in which we know the ground truth and thus facilitates the validation of the method.

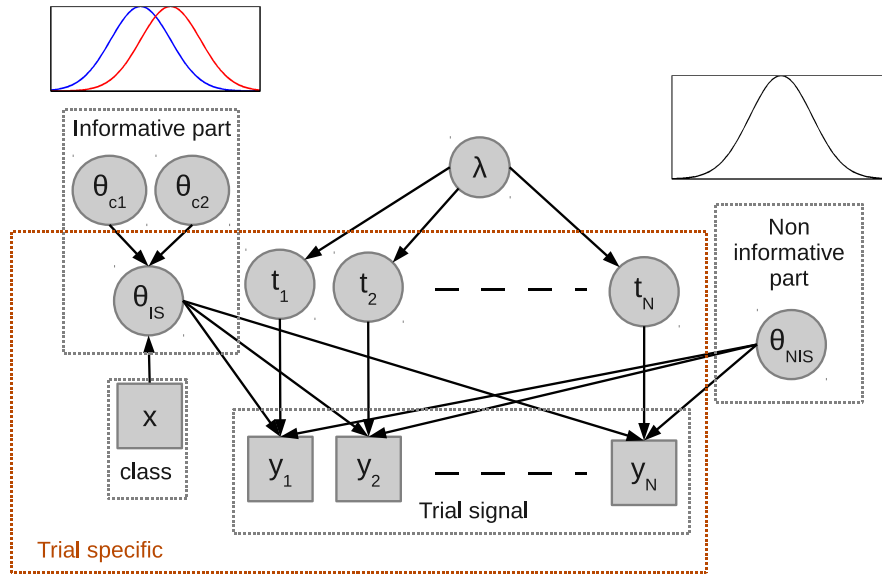


FIGURE 5.1 – Graphical model of the signal distribution. θ denotes collectively the parameters of the distribution of an element.

5.1 Adaptation of the previous approach to a generative model

5.1.1 Signal model

The trials are assumed to contain single occurrence of the task that the subject is instructed to perform : the condition (class) does not change during a trial. In the case of the protocol for studying the contrast between exploration and exploitation, the window of analysis was chosen to contain one decision, i.e. the EEG signal extracted prior to the machine selection at the trial t could not overlapped with the previous machine selection at time $t - 1$. With this constraint, we assume that the T time samples of a trial are independent observations y_t of the same state x . Under those conditions we can express the posterior probability :

$$p(C = x|y_{1:T}) = p(C = c) \frac{\prod_{t=1}^T p(y_t|C = x)}{\sum_{c=0}^1 p(C = c) \prod_{t=1}^T p(y_t|C = c)} \quad (5.1)$$

As explained previously, we assume that not all the time-samples in a trial are informative. More formally, we define two sets of samples in a trial following two different distributions : the non-informative and the informative sets (see Figure 5.1). The distribution of the non-informative set is the same for both classes whereas the *distribution of the informative set is specific to each class*. It is specifically their difference that will enable us to classify a trial. Moreover, we define the probability λ of a sample to belong to the informative set of any class. This probability is shared by both classes. We can then model the distribution of the samples as a mixture of the distribution of the

5.1. Adaptation of the previous approach to a generative model

informative (p_i) and non-informative (p_n) sets :

$$p(y_t|X = x) = \lambda p_i(y_t|X = x) + (1 - \lambda)p_n(y_t) \quad (5.2)$$

By combining with the equation 5.1, we can express the posterior probability based on the informative and non-informative distribution

$$p(X = x|y_{1:T}) = \frac{1}{1 + \frac{p(X \neq x)}{p(X = x)} \prod_{t=1}^T \frac{\lambda p_i(y_t|X \neq x) + (1 - \lambda)p_n(y_t)}{\lambda p_i(y_t|X = x) + (1 - \lambda)p_n(y_t)}} \quad (5.3)$$

At this point, if we know λ , p_i and p_n , given any recorded data $y_{1:T}$ of a trial, we can compute the probability of belonging to a particular class. Now, in order to go further and to do useful prediction with $p(X = x|y_{1:T})$, it is necessary to choose a model of p_i and p_n parametrised by θ . Then by inferring the estimation of θ and λ by a Bayesian approach from the data collected in the training phase, we will be able to use this model to classify each trial as one of the classes.

5.1.2 Estimation of the parameters

The model $p(y_t|x)$ introduces several free parameters, which makes the estimation procedure difficult. Following a Bayesian approach (see section 2.4 : Bayesian approach to estimation), we will explicitly define the prior of the parameters of our model. This prior represents all the knowledge we have about the problem, and will reduce the space of possible solutions. A sensible choice of the priors can also favour a discriminant solution : for example, we can safely say that λ should not be too small (i.e. there is at least part of the trial that is class specific) neither too big (i.e. only a subpart of the trial is informative). Depending on the problem, we may incorporate knowledge about the parameters of the non-informative and informative sets as well : for example the mean of the distribution of the two informative sets have to be separated. If we note collectively y_{tr} the data in the trials used in the training phase and θ the parameters to estimate, from the Bayes theorem :

$$p(\theta|y_{tr}) \propto p(y_{tr}|\theta)p(\theta) \quad (5.4)$$

We estimate the different parameters by using the posterior mean, Bayesian estimator which minimises the quadratic cost :

$$\hat{\theta} = \int_{\theta} \theta p(\theta|y_{tr}) d\theta \quad (5.5)$$

In order to compute this integral, we use a Monte Carlo Markov Chain (MCMC) algorithm which allows to sample the distribution $p(\theta|y_{tr})$ [77]. The integral is then

approximated by averaging the samples θ_i of the distribution :

$$\theta_i \sim p(\theta|y_{tr}) \quad (5.6)$$

$$\hat{\theta} = \sum_{i=0}^{+\infty} \theta_i \quad (5.7)$$

One might argue that the estimation is then biased by the choice of the prior. However, given a finite number of observations, making an inference on a parameter, from a mathematical point of view, requires defining a prior [77]. If absolutely no prior information is known about the parameter, it is possible to use a non informative prior (the term is not related to what we call the informative/non informative sample in this chapter) which would correspond in our case to a uniform distribution. However, since we have it, we prefer to specifically use the prior as a natural way to incorporate knowledge we have about a problem. It is also important to note that the influence of the prior on the estimation decreases as the number of observations increases.

5.2 Application to error related potentials

To test the framework, we have applied it in the context of ErrP classification. This choice has been motivated by the fact that they are well known EEG responses. This potential is elicited when a human user monitors an external system upon which he has no control whatsoever. When a subject perceives an error, there is typical EEG response that is distinguishable for the EEG response in case of no error (see Figure 5.2a). The experimental protocol used is described in [91]. In the mentioned study, 6 subjects were recorded and are included in the current study (corresponding to subject 6 to 12). Using the same protocol, we also used data from another study where 6 other subjects were recorded.

EEG signal of 12 subjects has been recorded with a sampling rate of 512 Hz using 64 electrodes according to the standard 10/20 international system. Data was re-referenced to the common average reference (CAR) and a 1-30Hz band-pass filter was applied. We then focus the analysis on the data recorded at Cz electrode from 0 s to 0.8 s after the stimulus onset. In order to keep the approach similar to what has been done in the previous chapter, we use the frequency power density as feature of the classifier. The power density of the 7 Hz (corresponding roughly to the frequency band containing the most energy of the ErrP) has been computed through a complex wavelet transform (using a Morlet mother wavelet) : for each time sample, we have an estimation of the band power of the 7 Hz component. As such, the power or amplitude does not follow a Gaussian distribution. Therefore all the power is log transformed in

5.2. Application to error related potentials




FIGURE 5.2 – Grand average of an ErrP of one subject : in red : the grand average of EEG response of trial containing an error. In blue : grand average of correct trials. In green : the difference between the two conditions. (a) Grand average of the EEG signal in the time domain. (b) Grand average of the data with the feature used in the study (log transform of the band power around 7 Hz). The variances around the mean are plotted in dashed line

order to make their distribution more normally distributed (Figure 5.2b)

5.2.1 Model

We model both informative and non-informative set distributions as Gaussian. λ represents the probability of a sample in a trial to belong to the informative set. In the case of ErrP, it can be expected that the waveform with the prominent negative and positive peak will be recognised as the informative set. Since this waveform is approximately 0.25 s in a trial of 0.8 s, λ can be expected to be estimated around




FIGURE 5.3 – Distribution beta as used for the prior of λ .

Chapitre 5. Bayesian approach

0.3. Accordingly, the prior of λ is modelled as a beta distribution centred on 0.3 (see Figure 5.3). In order to favour a discriminant solution, we model the prior distribution of the absolute value of the difference between the means of the informative sets as a Gaussian ($\mu_d = 1, \sigma_d = 1$).

Under these conditions, all distributions considered in the model belong to the exponential family (see Robert [77] for more explanation). It is then possible to express all the posterior probabilities of each individual parameter to estimate conditionally on all other parameters as a distribution that is possible to simulate (i.e. there exist algorithms to sample the distribution).

We will detail the reasoning of that lead to the expression of the conditional posterior probability for one variable (actually two as we will immediately see), then we will briefly express the conditional posterior of the others and without giving the details of the calculus. The reasoning that lead to the expression of the conditional posteriors is usually the same : the posterior conditional is derived from the full posterior expression and we try to express it as a function of the variable of interest from which the form of the distribution and its parameters are recognised.

To find the expression of the conditional of λ , we introduce latent variables $k_{t,j}$ indicating whether the t -th sample $y_{t,j}$ of the j -th trial is an informative sample ($k_{t,j} = 1$) or not ($k_{t,j} = 0$). We note M the total number of trials in the training set. Additionally, we note N_i and N_n the total number of respectively informative and non informative time samples in the train set. We note collectively "rest" the parameters of the informative and non-informative model distribution.

$$p(\lambda | \mathbf{y}_{j=1:M}, \mathbf{k}_{j=1:M}, rest) \propto p(\mathbf{y}_{j=1:M} | \mathbf{k}_{j=1:M}, \lambda, rest) p(\mathbf{k}_{j=1:M} | \lambda, rest) p(\lambda | rest) \quad (5.8)$$

The introduction of the variable $k_{t,j}$ allows to rewrite the likelihood of the observation $y_{t,j}$. So by noting c_j the class of j -th trial :

$$p(y_{t,j} | c_j, \lambda, k_{t,j}, rest) = \lambda p_i(y_{t,j} | c_j) + (1 - \lambda) p_n(y_{t,j}) \quad (5.9)$$

becomes :

$$p(y_{t,j} | c_j, \lambda, k_{t,j}, rest) = p_i(y_{t,j} | c_j)^{k_{t,j}} p_n(y_{t,j})^{1-k_{t,j}} \quad (5.10)$$

The shows that given $\mathbf{k}_{j=1:M}$ and the rest of the variables, $y_{t,j}$ does not depend on λ . In addition, given λ , $\mathbf{k}_{j=1:M}$ do not depend on the rest of the variable and λ do not

5.2. Application to error related potentials

depend on the rest of the variables. Thus we have :

$$p(\lambda|\mathbf{y}_{j=1:M}, \mathbf{k}_{j=1:M}, rest) \propto \underbrace{p(\mathbf{y}_{j=1:M}|\mathbf{k}_{j=1:M}, rest)}_{\text{independent from } \lambda} p(\mathbf{k}_{j=1:M}|\lambda)p(\lambda) \quad (5.11)$$

$$\propto p(\mathbf{k}_{j=1:M}|\lambda)p(\lambda) \quad (5.12)$$

$$\propto \left(\prod_{\substack{t=1:T \\ j=1:M}} p(k_{t,j}|\lambda) \right) p(\lambda) \quad (5.13)$$

$$\propto \lambda^{N_i} (1 - \lambda)^{N_n} p(\lambda) \quad (5.14)$$

If note α and β the parameters of beta distribution of the prior of λ , we have then :

$$p(\lambda|\mathbf{y}_{j=1:M}, \mathbf{k}_{j=1:M}, rest) \propto \lambda^{N_i} (1 - \lambda)^{N_n} \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \lambda^{\alpha-1} (1 - \lambda)^{\beta-1} \quad (5.15)$$

$$\propto \lambda^{(\alpha+N_i)-1} (1 - \lambda)^{(\beta+N_n)-1} \quad (5.16)$$

We then recognise the form of a beta distribution with parameters $\alpha + N_i$ and $\beta + N_n$. We can notice from the previous equations that if we assume a flat prior for λ , its prosterior conditional would have also a beta distribution but with the parameter $N_i + 1$ and $N_n + 1$.

Using a similar reasoning, we deduce the posterior conditional distribution of the other parameters. To simplify the expressions, we introduce few definition : N_{ic} and N_n will be the number of samples in respectively the informative set of the class c and the non informative set, μ_c and σ_c^2 the mean and variance of the informative set of the class c , μ_n and σ_n^2 those of the non-informative set. \bar{y}_{ic} and \bar{y}_n denotes the mean of the sample recognised respectively as the informative for the class c and non-informative. Finally I_c denotes the set of trials of class c in the training set.

After deriving the expressions, it turns out that the posterior conditional of $k_{t,j}$ is a simple Bernouilli distribution (the most simple discrete probability distribution, which takes value 1 with success probability p or 0 otherwise) with a probability :

$$p = \frac{\lambda \varphi(y_{t,j}, \mu_{c_j}, \sigma_{c_j}^2)}{\lambda \varphi(y_{t,j}, \mu_{c_j}, \sigma_{c_j}^2) + (1 - \lambda) \varphi(y_{t,j}, \mu_n, \sigma_n^2)} \quad (5.17)$$

In the same way, we derive the conditional posterior of the parameters of the informative and the non-informative sets and recognise that they follows these distributions

TABLE 5.1 – Estimated lambda

subject	s1	s2	s3	s4	s5	s6	s7	s8	s9	s10	s11	s12
λ	0.14	0.14	0.12	0.73	0.12	0.75	0.75	0.11	0.24	0.76	0.17	0.09

(assuming a flat prior for each of them) :

$$\mu_c \sim \mathcal{N} \left(\frac{1}{\sigma_d^2 + \sigma_c^2} \left(\sigma_d^2 \bar{y}_{ic} + \frac{\sigma_c^2}{N_{ic}} (\mu_n + \mu_d) \right), \frac{\sigma_c \sigma_d}{\sigma_c + N_{ic} \sigma_d} \right) \quad (5.18)$$

$$\sigma_c^2 \sim \mathcal{IG} \left(\frac{N_{ic}}{2} + 1, \frac{2N_{ic}}{\sum_{j \in I_c, t} (k_{t,j} y_{t,j} - \mu_c)^2} \right) \quad (5.19)$$

$$\mu_n \sim \mathcal{N} \left(\bar{y}_n, \frac{\sigma_n^2}{N_n} \right) \quad (5.20)$$

$$\sigma_c^2 \sim \mathcal{IG} \left(\frac{N_n}{2} + 1, \frac{2N_c}{\sum_{j,t} ((1 - k_{t,j}) y_{t,j} - \mu_n)^2} \right) \quad (5.21)$$

All posterior conditional have a distribution model that can be sampled. Therefore, to infer the estimated value to estimate the parameters λ and the means and variances of the informative and non-informative sets from the posterior distribution, it is possible to the Gibbs sampler which is an efficient MCMC method.

5.2.2 Results

The classification procedure has been tested on approximately 1150 trials per subject containing 20% of trials of the error condition. The first two thirds (chronologically) trials has been used for training and the rest reserved for testing. The results of the estimated λ are provided in Table 5.1. For most of the subjects, λ is estimated around 0.15 which means that on average 0.12s of a trial is recognised as informative. If taken continuous, these length may correspond a little bit more than the length between the two major negative and positive peaks of an ErrP. For 4 subjects, λ has been estimated higher than expected. This is however not an issue since it is possible that for these subjects the trials of the two classes exhibits difference beyond only the expected peak shift (the inversion from the negative peak to the positive one as seen on the grand average in Figure 5.2a) usually located between 0.2s and 0.35s.

Given the estimated parameters, we can infer the probability for each individual sample in a trial belonging to the informative set of each class. For each trials of both classes, the probability of each sample to belong to their respective informative set is computed. By averaging this result accross trials of the same class, we have an approximation of how likely a sample measured at a specific time within the trial belongs to the informative set of its respective class. The results are presented in

5.2. Application to error related potentials




FIGURE 5.4 – Mean probability of a sample of a specific time being recognised as belonging to the informative set (in red : samples of class error, in blue : samples of class correct, y-axis : probability, x-axis : time in seconds).

TABLE 5.2 – Classification results on test set

	Bayesian			LDA
	error acc	correct acc	mcc	mcc
s1	0.73	0.62	0.28	0.647
s2	0.58	0.64	0.18	0.411
s3	0.59	0.69	0.23	0.629
s4	0.88	0.63	0.40	0.418
s5	0.60	0.75	0.30	0.228
s6	0.59	0.73	0.27	0.233
s7	0.62	0.78	0.34	0.293
s8	0.41	0.52	-0.05	0.147
s9	0.45	0.67	0.10	0.226
s10	0.62	0.73	0.30	0.323
s11	0.48	0.81	0.28	0.470
s12	0.58	0.61	0.16	0.329

Figure 5.4. Interestingly, we can see that for all subjects but two (subjects 8 and 9), more samples located around 0.3s are recognised to belong to the informative set for the error class. This indicates clearly that the estimation procedure has correctly recognised the peak shift between 0.2s and 0.35s as being part of the informative set. On the other hand, the localisation of the informative samples of class correct is less obvious than the class error. For some subjects, the distribution is almost flat. This indicates that either their informative samples are less synchronised (spread over the whole trial) than those of class error, or that the estimation procedure had difficulties to identify an informative set different from the non-informative set in the class correct.

Based on the estimated parameters on training set, we can compute for each trial the posterior probability of belonging to a particular class using the equation (5.1). The results are compared with those obtained with a synchronous approach : the LDA is computed on the signal in time domain in a particular time-window on highly downsampled data (64 Hz after applying a 1-10 Hz bandpass).

The results presented in Table 5.2 reports the classification performance using MCC values (see subsection 3.4.1 of Chapter 3). MCC are used because of the number of trials in each are unbalanced. The table shows that the performances of our approach are comparable for many subjects although it performs worse for others than the classical synchronous approach (see the MCC results). However, it should be noted that the classification accuracy for class correct is surprisingly low given the size of this class in the data set (80% of the trials). This means that our classification method has a lot of false positives.




FIGURE 5.5 – Overlapped plots of all trials of class error and class correct of subject 7 (x-axis : time in seconds).

5.3 Discussion

The approach describes in this chapter works under the assumption that the method can extract an informative set whose distribution is specific to each class. But while it is obvious from Figure 5.5 that the trials in class error which exhibits a high peak of amplitude are specific to this class, we can also see many trials in both classes that shows a maximal amplitude that are similar which can lead informative set distributions that overlap too much. From the raw potentials, it seems that shapes of the waveform could be more specific to each class than the amplitude of a frequency component. This is an hypothesis that will be considered in the next chapters.

In the model describes in this chapter, we have assumed the independence of the observation within a trial. In this sense, the method uses here a naive Bayes classifier. This has the advantage to lead to a simple classification rule. However when considering Figure 5.5, the hypothesis of independence is obviously not valid. However, this is not a problem for classification purpose since it makes the correct maximum a posteriori [92]. However it fails to produce a good estimate for the correct class probabilities which will tends to often close to 0 or 1. This might be an issue if the estimated class probabilities must used in a larger decision system.

The method is based on the hypothesis that there exists a set of time samples in the trial whose distribution is specific to each one class and the rest is distributed by a distribution that is common to both classes (the non-informative set). However the way the parameters of these distributions are estimated, nothing enforces that the distribution of the non-informative set differs from the informatives sets. It will be the case only if the training data present this characteristic. Otherwise, if the distribution of the informative set of one class is hardly different from the distribution of the non-

Chapitre 5. Bayesian approach

informative, this will tend to lead to an high estimated λ as we can see in Table 5.1 for four subjects. However this does not necessarily lead to huge performance drop (see subject 4). But this necessarily means that the size of the informative set of the other class will be overestimated.

A way to temper this issue would be to modify the model to allow to estimate λ per class. This is feasible by doing small changes to the signal model. However, in this case, the estimation of the parameters will need to be carefully performed in order to avoid λ per class to be too large which would increase the problem of the naive Bayes hypothesis described earlier. In the rest of the thesis we will modify the method in different way, so we will not cover this aspect.

To conclude, this chapter proposes a framework of asynchronous classification that aims to be more formal compared to the one described in the previous chapter. It is designed to be flexible enough to be able to tackle problems in the same context of classification. It reduces the number of hyper-parameters by incorporating them directly into the model (for example the percentile threshold of the last chapter is turned into the probability λ to be in the informative set). However this incorporation has come with a slight increase of the number of the parameters to estimate. Nevertheless these parameters are also easier to understand from the problem. From the understanding of the ErrP it is natural to provide a prior on the probability λ as we have seen in the chapter.

The results show that we can classify EEG evoked potentials with this approach without relying on strong time cue. However its performance is generally worse than the classical synchronous approach because of high rate of false positive. The shape of the waveform of the ErrP seems more specific to each class. Thus incorporating it in the model of the informative part might lead to better results. This modification will be the topic of the next chapter.

6 Template based realignment

6.1 Introduction

Describing a signal in a time-frequency fashion has its advantage when trying to study an asynchronous phenomenon that has many components in frequency domain. However when its frequency components are roughly defined (because, as instantaneous phenomenon, it does not contain many cycles), its time-frequency description may not be localised enough to be useful for recognition purpose. This is what happens for the ErrP we have seen previously on Figure 5.5 : the relevant pattern seems to be well localised in time-domain but when analysing the 7Hz component (the most salient frequency component) its localisation seems more fuzzy. A better description would likely be obtained using more frequency components around 7 Hz. However when dealing with a new problem, it is not priory known how precise the frequency resolution should be (note that we were using a continuous complex wavelet transform) in order to have a good description of the phenomenon.

So it is proposed in this chapter to use different approach which would analyse the signal in the time domain. Many EEG signal in particular evoked potential have a suitable description in the time-domain we have seen in the subsection 2.2.1 of the chapter 2. It also not excluded that many induced activity may also have a description in the time-domain suitable for classification but which is difficult to detect due to the absence of time-locked nature of the signal. An approach to this problem is presented in this chapter. In the opposite of the previous study the relevant information will not be described in terms of snapshots in time of the signal in the transformed space. Instead, the relevant part is considered as a complete sequence of the trial. As such, the previous approach is turned into a pattern recognition problem.

This approach assumes that a similar signal pattern can be found in the trials of each class, i.e. that particular features in the trial can be found synchronised relatively to each other, although the time onset of that pattern may differ from trial to trial.

As such, this method could be seen as a method of template matching whose the template would have a probabilistic model. An advantage of such approach compared to one proposed in the previous chapter is that it allows to characterise differently the amplitude of the different peaks of an evoked potential when applied on evoked response data.

6.2 Model

In this chapter, we hypothesise that the informative part is a pattern in the time domain that can be found in all trials. It represents a sequence that is present in all trials but whose onset differs in each trial. This informative part can be seen as a template that is searched within the trial. However, it should be noted that this template is not simply described in term of sequence of points, but instead in a probabilistic way.

The non-informative part is also described in a probabilistic way but contrary to the previous chapter, its model (its probabilistic description) is very different for the informative part. In this section we will see in the first two subsections the chosen models for the template (i.e. the informative part) and for the non-informative part. Figure 6.1 provides a graphical model of the trial signal. Given these models, we will then see in the last subsection how to estimate their parameters from the recorded trials using a Bayesian approach.

6.2.1 Template model

This new approach makes this hypothesis that there is a similar pattern of same length in each trial of the asynchronous process. In other words, the informative part is assumed to be a sequence of points that follows a certain probability distribution that needs to be estimated (Figure 6.2).

To keep the model simple, the points of the sequence are assumed to be independent. Although previous model assumed independence of the samples within a trial, the difference here is that the independence is conditioned on the pattern onset : given the onset τ in the trial, the i -th point in the informative part follows a distribution $h(x|\theta_i)$ that depends on its distance from the onset but is independent from the other points in the pattern :

$$X_{i+\tau} \sim h(x|\theta_i) \tag{6.1}$$

A model of distribution of each point within the pattern should be chosen. When considering the case of the ErrP in Figure 5.5, it is remarkable that of one the most salient difference between error and correct class lies in the peaks of the evoked

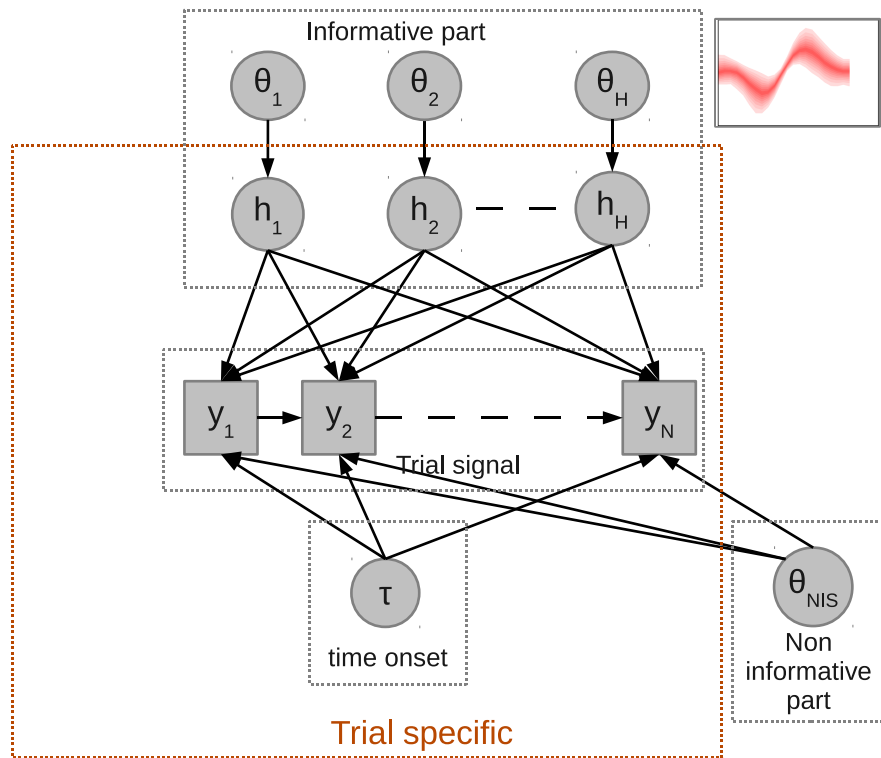


FIGURE 6.1 – Graphical model of the signal distribution of the trial. θ denotes collectively the parameters of the distribution of an element.




FIGURE 6.2 – Examples of trials (synthetic data) : the informative part is highlighted in blue (y-axis : EEG potential, x-axis : time)




FIGURE 6.3 – Distribution of the peaks of ErrP. In the figure (a), the histogram of the negative peak values is displayed in red for class error and in blue for class correct (Electrode Cz, subject 7). In the figure (b), the histogram is plotted in the similar manner as (a). (y-axis : number of elements, x-axis : EEG potential in μV)

potential. Now by analysing the distribution of these peaks (Figure 6.3), it can be that the variability around the mean peak value is asymmetric : the variability is higher for those which are far from zero. In other words, the distribution of the extrema are skewed.

In such a context, choosing a Gaussian model as it may come naturally in mind might not be the ideal choice for the distribution model of the points of the pattern. Instead it is better to use a model of a skew normal distribution which is simply a generalisation of the Gaussian distribution with a skewness parameter [93]. More formally, a univariate random variable X follows a standard skew-normal distribution \mathcal{SN} with skewness parameter α if its density follows :

$$p(x; \alpha) = 2\varphi(x, 0, 1)\Phi(\alpha x, 0, 1) \tag{6.2}$$

where φ is the probability density function of a normal distribution and Φ the cumulative distribution function of a normal distribution. With a trivial change of variable $Y = \xi + \omega X$ with $X \sim \mathcal{SN}(\alpha)$, the distribution can be adjusted easily to any location ξ and scale ω (See Figure 6.4). Since the skew normal distribution is a model generalising of the Gaussian distribution, it is a suitable choice for the model of the point of the pattern : the skewness parameter can indeed always be estimated to null if the method is applied on a EEG signal that does not require a skewed distribution

6.2.2 Non-informative samples

For the non informative part, we need to choose a model that will represent the local evolution of the dynamic of the signal while staying general enough to represent signal




FIGURE 6.4 – Examples of skew normal distributions for different location, scale and skewness.

sequences that may be quite different from a trial to another. Most of the studies of EEG pattern (in particular evoked potential) requires a step of band pass filtering the EEG signal in the preprocessing. Thus trying to represent the evolution of coloured noise could be a good approach to model the non informative part of the trial.

In order to achieve this, a stochastic model of coloured noise is proposed. The signal x is decomposed into two components estimated from the analytic representation x_a (using the Hilbert transform) : its local frequency and its local amplitude. The local amplitude is then assumed to follows a linear stochastic evolution :

$$|x_a|_{t+1} = \lambda |x_a|_t + (1 - \lambda)m + \epsilon \quad \text{with} \quad \epsilon \sim \mathcal{N}(0, s) \quad (6.3)$$

whose parameters m , λ and s are common in all non informative part.

6.2.3 Training the models

Following a Bayesian approach, the estimation of the parameters of the model is inferred from the posterior probability which requires to define the prior we have on each of them. To perform the inference, we intend to use a MCMC method which preferably would be the Gibbs sampler like in the previous chapter. In order to use it, it requires to be able to sample from each posterior conditional distribution of each parameter. The use of the skew normal complicates this task and we will see how to overcome this issue. Finally, we will see the results of the derivation of the conditional posteriors as they are used in the implementation of the Gibbs sampler for the method

described in this chapter.

Priors

Like in the previous chapter, the priors must be chosen so that they will constraint the estimation to sensible values while not impairing the generality of the method. For this purpose, we notice that we can easily provide a prior distribution on two set of variables : the onsets τ_t of the informative pattern in the each trial and the parameters of the non-informative samples.

The prior of the onsets can indeed be set according to the knowledge the experimenter has about the problem he is dealing with. If we take the example of ErrP, we know –from previous studies or more simply by observing the grand average– that the informative pattern will likely be located between 150 and 450 ms after the stimulus onset. Given the length of the informative pattern used for the model, we can turn this into a prior probability distribution for τ_t . More interestingly, as we will see later, the conjugate prior of the onset is a general discrete distribution on the possible values for the pattern onset. In other words, the choice of the model of this prior does not influence the computation performance.

The other set of variables whose prior can be easily chosen are the parameters of the non-informative part of the trial. As we have seen the previous section, its model has been chosen such as it can represent the local evolution of a general EEG signal. By taking the whole EEG signal, no matter whether it corresponds the signal of interest or not, we can estimate the parameters of this model with a "typical" EEG signal. If we assume that the non-informative part is representative of this "typical" signal, the prior can then be immediately chosen.

Sampling the parameters of a skew normal model

Like in the previous chapter, we want to estimate the parameters of the model by sampling the posterior distribution using a MCMC algorithm. However taking into account the skewness in the distribution of the informative pattern has some side effects. Contrary to the normal distribution, there is no generator of the conditional distribution of a skew normal distribution. A first possibility would be to use a Metropolis Hastings algorithm to sample them. However compared to the Gibbs sampler (its special case), this algorithm can quickly slow down the convergence if the instrument law diverges from the target density (See section 2.4.3 for reminder about the Metropolis-Hastings algorithm).

Another hybrid approach could have been to slightly modify the model and make the parameters of the informative pattern directly dependent from the other variable of

the model, thus estimating them at each iteration of the MCMC algorithm. However, this problem has also issues since there is no analytic form of the maximum likelihood estimator for a skew normal distribution and the numerical estimator is quite unstable especially when the skewness parameter is close to 0 [93, 94].

To tackle the problem in an elegant way, it is better to resort to data augmentation which is a method for constructing iterative optimisation or sampling algorithms via the introduction of unobserved data or latent variables [95]. The idea is to find a new parameterisation of the same distribution such as the conditional posteriors of the new parameters to distributions that can be sampled. So it is what is presented here.

Let us consider the skew normal distribution $\mathcal{SN}(\xi, \omega^2, \alpha)$. It can be shown [96, 97] that if :

$$\begin{aligned} Y &= \xi + \omega\delta Z + \omega\sqrt{1 - \delta^2}\varepsilon \\ Z &\sim \mathcal{TN}_{[0,+\infty)}(0, 1) \\ \varepsilon &\sim \mathcal{N}(0, 1) \\ \delta &= \frac{\alpha}{\sqrt{1 + \alpha^2}} \end{aligned}$$

where \mathcal{TN} refers to the truncated normal distribution (see notations in Appendix C), then the random variable Y follows the skew normal distribution $\mathcal{SN}(\xi, \omega^2, \alpha)$

With this, we have then a method to sample any skew normal distribution since it is possible to sample a truncated normal distribution (See Appendix B for the method).

At this stage we can introduce a reparameterisation (ξ, ψ, σ^2) with :

$$\psi = \omega\delta \quad \text{and} \quad \sigma = \omega(1 - \delta^2) \tag{6.4}$$

If the random variables z , ϵ and y are defined such as :

$$\begin{aligned} z &\sim \mathcal{TN}_{[0,+\infty)}(0, 1) \\ \epsilon &\sim \mathcal{N}(0, \sigma^2) \\ y &= \xi + \psi z + \epsilon \end{aligned}$$

Then

$$y \sim \mathcal{SN}(\xi, \omega^2, \alpha)$$

In this new parameterisation, the parameters are independent conditionally to the variable z . Thus, it is now easy to calculate their conditional posterior distribution and run a Gibbs algorithm to sample the density $p(\xi, \psi, \sigma^2 | \mathbf{x})$. All of this has been made possible by introducing the latent variable z which is necessary to run the sampling

Chapitre 6. Template based realignment

algorithm but in the end, the estimation of the variables ξ , ω and α do not depends on it :

$$\alpha = \frac{\psi}{\sigma} \quad \text{and} \quad \omega = \sigma^2 + \psi^2 \quad (6.5)$$

The Gibbs sampler implementation

As described in the previous sections, we have the following model : in each trial t , there is a part that is distributed by the probability model of the informative pattern whose onset starts at the index τ_t . Assuming that the pattern has a length of H points and a trial T points, we have for any j from 1 to T :

$$\begin{cases} x_{t,j} = \xi_i + \psi_i z_{t,i} + \epsilon_{t,i} \text{ with } i = j - \tau_t & \text{if } j \in [\tau_t, H + \tau_t] \text{ (informative)} \\ |x_a|_{t,j+1} = \lambda |x_a|_{t,j} + (1 - \lambda)m + \epsilon_{t,j} & \text{if } j \notin [\tau_t, H + \tau_t] \text{ (non-informative)} \end{cases} \quad (6.6)$$

with

$$\epsilon_{t,i} \sim \mathcal{N}(0, \sigma_i^2), \quad z_{t,i} \sim \mathcal{TN}_{[0,+\infty)}(0, 1), \quad \epsilon_{t,j} \sim \mathcal{N}(0, s)$$

So , each time point in the trial is modelled as a random variable distributed either by the model of the informative set or by the model of the non-informative set. With this model, we can calculate the conditional posterior distribution for each parameter to be estimated. For the sake of simplicity, the details of the calculus are skipped. For the parameters of the informative pattern, we have (by noting collectively by *rest* the observations and the other variables) :

$$z_{t,i}|rest \sim \mathcal{TN}_{[0,+\infty)}\left(\frac{\psi^2}{\psi^2 + \sigma_i^2} (x_{t,i+\tau_t} - \psi_i), \frac{\sigma_i^2}{\psi^2 + \sigma_i^2}\right) \quad (6.7)$$

$$\xi_i|rest \sim \mathcal{N}\left(\sum_{t=1:N} \left(\frac{x_{t,i+\tau_t} - \psi_i z_{t,i}}{N}\right), \frac{\sigma_i^2}{N}\right) \quad (6.8)$$

$$\sigma_i^2|rest \sim \mathcal{IG}\left(\frac{N}{2} + 1, \frac{2}{\sum_{t=1:N} (x_{t,i+\tau_t} - \psi_i z_{t,i} - \xi_i)^2}\right) \quad (6.9)$$

$$\psi_i^2 | rest \sim \mathcal{N} \left(\frac{\sum_{t=1:N} (z_{t,i} (x_{t,i+\tau_t} - \psi_i z_{t,i} - \xi_i))}{\sum_{t=1:N} z_{t,i}}, \frac{\sigma_i^2}{\sum_{t=1:N} z_{t,i}} \right) \quad (6.10)$$

For the time onsets of the informative pattern, we have ($p(\tau_t)$ is the prior of the time onset) :

$$f_1(\mathbf{x}_t, \boldsymbol{\sigma}^2, \boldsymbol{\psi}, \mathbf{z}_t, \boldsymbol{\xi}, \tau) = - \sum_{i=1:H} \frac{(x_{t,\tau+i} - \xi_i - \psi_i z_{t,i})^2}{2\sigma_i^2} \quad (6.11)$$

$$f_2(\mathbf{x}_t, m, s, \lambda, \tau) = - \frac{1}{2s^2} \sum_{i=1:\tau, \tau+H:N} \left(|x_a|_{t,i+1} - \lambda |x_a|_{t,i} - (1-\lambda)m \right) \quad (6.12)$$

$$p(\tau_t | rest) \propto \exp \left(f_1(\mathbf{x}_t, \boldsymbol{\sigma}^2, \boldsymbol{\psi}, \mathbf{z}_t, \boldsymbol{\xi}, \tau_t) f_2(\mathbf{x}_t, m, s, \lambda, \tau_t) \right) p(\tau_t) \quad (6.13)$$

Finally for the parameters of non informative part :

$$m | rest \sim \mathcal{N} \left(\sum_{\substack{t=1:N \\ i \notin [\tau_t, \tau_t+H]}} \frac{|x_a|_{t,i+1} - \lambda |x_a|_{t,i}}{N(T-H)(1-\lambda)}, \frac{\sigma_a^2}{N(T-H)(1-\lambda)} \right) \quad (6.14)$$

$$\sigma_a^2 | rest \sim \text{IG} \left(\frac{N(T-H)}{2} + 1, \frac{2}{\sum_{\substack{t=1:N \\ i \notin [\tau_t, \tau_t+H]}} \left(|x_a|_{t,i+1} - \lambda |x_a|_{t,i} - (1-\lambda)m \right)^2} \right) \quad (6.15)$$

$$\lambda | rest \sim \mathcal{TN}_{[0,1]} \left(\frac{\sum_{\substack{t=1:N \\ i \notin [\tau_t, \tau_t+H]}} (|x_a|_{t,i-1} - m)(|x_a|_{t,i} - m)}{\sum_{\substack{t=1:N \\ i \notin [\tau_t, \tau_t+H]}} \left(|x_a|_{t,i-1} - m \right)^2}, \frac{\sigma_a^2}{\sum_{\substack{t=1:N \\ i \notin [\tau_t, \tau_t+H]}} \left(|x_a|_{t,i-1} - m \right)^2} \right) \quad (6.16)$$

6.3 Result on synthetic data

The inference procedure has been tested on synthetic data [98]. The rationale is to verify that if we have data that follows exactly the assumed models, the inference




FIGURE 6.5 – 100 trials of synthetic data simulating ErrP. The y-axis correspond to the trials, the x-axis the time in the trial and the color the EEG signal. .

methods converge correctly to the expected values. In order to do this, Error related potentials have been simulated using the models described in the previous section.

So 500 trials of 1 second each have been synthesised using a sampling frequency of 128 Hz. The ErrP pattern used is 0.3 s long : ξ_i is generated by a sine wave centred around 7 Hz modulated by a Gaussian function. The α_i parameters have been chosen proportional to ξ_i , $\omega_i = \xi_i + K$ (K constant). With such parameters, the pattern obtained is a sequence that contained two big peaks, a negative and then a positive one whose skewness increases with the peak value. It is similar to what has been seen in Figure 6.3. However, contrary to the ErrP data that have been studied in the previous chapter, the standard deviation of the time onsets has been arbitrarily increased (standard deviation : 0.1 s) to test whether the method can recover from a moderate jitter in the data. To illustrate this, the first 100 trials have been displayed in Figure 6.5.

In addition, no prior has been used. The initial parameters of the pattern distribution (for the Gibbs sampler) have been initialised so that the mean and the skewness of each point in the pattern are zeros and the variance is very high. The time onsets have been initialised such as each informative part is located in the middle of each trial. The parameters of the non-informative part have been initialised to the values estimated by taking the whole synthetic EEG signal and by considering it as a whole as non-informative samples. This correspond to the procedure that would be carried out on real EEG signal.

The test shows that the method converges to a decent estimation. (Figure 6.6). In the figure, we notice that the estimated informative pattern is very similar to the one that has been used to generate the synthetic data and has been correctly localised. This is a general result on all trials : the difference between the estimated onsets and the one that have been used to generate the data is very small (mean : 31.3 ms standard deviation : 4.15×10^{-2} ms). But estimated onsets are consistently shifted compared to the real one. This is due to the fact that the estimated pattern is a slightly shifted version of the original pattern. This is explained by the fact that the most informative part of the pattern is located in its centre and there is no constraint on how the pattern




FIGURE 6.6 – Example of convergence in 3 synthetic trials : the sequence of ξ_i (approximately the mean of the pattern) is displayed at the location of the time-onset : in red the one that have been used to generate the trial, in blue the estimated one.




FIGURE 6.7 – Probability distribution of the prior used for the time-onset τ of the informative part.

should be aligned with it. It is an issue that may reappear with real data whenever the length of the used pattern is bigger that the real informative part of the trials. But this inconvenience is in practice not very important as long as the estimated pattern contains the informative part.

6.4 Test on real ErrP data

6.4.1 Methods

The method has been tested on the same ErrP data set used in the previous chapter. For modelling the informative part, we use a length of informative sequence of 0.17s (in a trial of 0.8s). As described in the subsection 6.2.3, we use a prior on the time-onset of the informative part. We know that the evoked response for an error potential occurs around 300ms after the stimulus. Therefore, the prior is modelled by a cropped Gaussian distribution on $[0s, 0.63s]$ of mean 0.2s and standard deviation 0.1s. This corresponds to the distribution presented in Figure 6.7. The initial values of the time-onset used in the MCMC algorithm are also drawn from this prior distribution.

In order to test the robustness of the estimation algorithm to desynchronised data, the

Chapitre 6. Template based realignment

TABLE 6.1 – Correlation of the grand average of the realigned data with the grand average of the original data

subject	normal trials	jittered trials
s1	0.805	0.762
s2	0.890	0.798
s3	0.902	0.911
s4	0.931	0.951
s5	0.845	0.845
s6	0.927	0.913
s7	0.909	0.931
s8	0.803	0.737
s9	0.571	0.638
s10	0.872	0.876
s11	0.883	0.893
s12	0.816	0.764

method has also been applied on the same ErrP data with a random jitter. To this end, we have drawn for each trial a random time shift from a uniform distribution between -0.2s to 0.2s. We then extract the new data from the original recording by shifting the data trial with these offsets which are stored to evaluate how much this jitter can be recovered. In the rest of the section, these trials will be referred as the jittered data. When referring to the results on this data, it is worth noting that the training has been performed on the jittered trials as well.

6.4.2 Results

Since the original trials are assumed to be reasonably time locked, it is expected that the distribution of the time-onset of the informative part is well centred around 0.2s (because the informative part is modelled as being 0.17s long). This is actually the case for all subjects (see the third column of Figure 6.8 and 6.9) except subject 8 and 9 (see Table 6.1). However, for these subjects, the grand average does not present a shape of a usual grand average of ErrP, so it is possible that the original data of these subjects is not well time-locked.

The second column of Figure 6.8 and 6.9 shows the grand average of the trials when they are realigned based on the time-onset found for the informative part in each trial. We can see that for most subjects the realigned average has a similar shape to the original average (except for subjects 8 and 9). However the realigned average is also often sharper and the positive and negative peaks are bigger than in the original average.

These results can be confirmed by computing an approximation of the signal to noise

6.4. Test on real ErrP data




FIGURE 6.8 – Results of estimation on ErrP (subject 1-6) : From left to the right, The first column shows the grand average, the second the grand average of all trials realigned (based on the detected time onset for each trial), the third column contains the distribution of the detected time onset. The fourth column contains the grand average of the randomly jittered trials. The fifth column contains the grand average of the jittered plus realigned trials. Finally the last column shows the distribution of the difference between the estimated onset and the shift introduced in each trial (recentred on the mean distribution).

Chapitre 6. Template based realignment

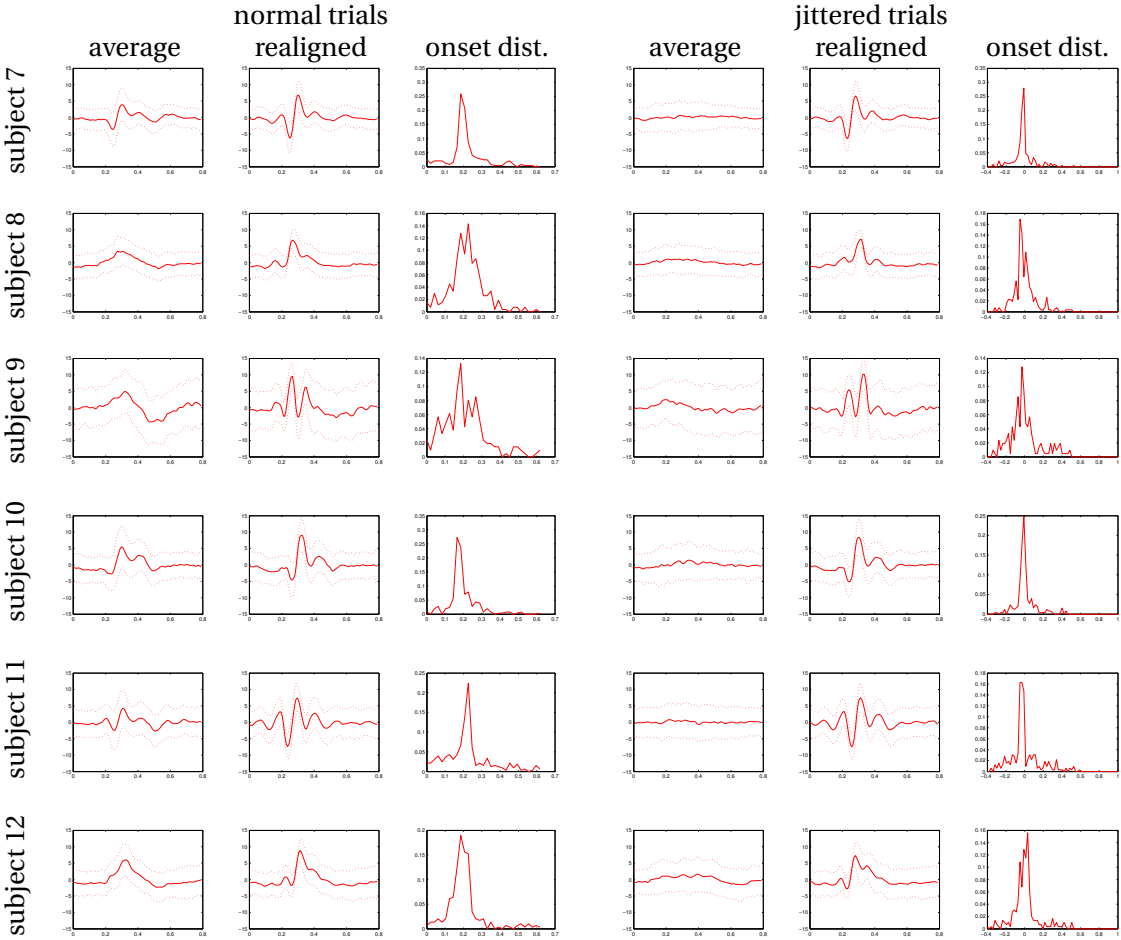


FIGURE 6.9 – Results of estimation ErrP (subject 7-12) : From left to the right, The first column shows the grand average, the second the grand average of all trials realigned (based on the detected time onset for each trial), the third column contains the distribution of the detected time onset. The fourth column contains the grand average of the randomly jittered trials. The fifth column contains the grand average of the jittered plus realigned trials. Finally the last column shows the distribution of the difference between the estimated onset and the shift introduced in each trial (recentred on the mean distribution).

TABLE 6.2 – Signal to noise ratio of the evoked response

subject	SNR original	SNR realigned
s1	2.192	3.505
s2	1.813	3.072
s3	2.232	3.218
s4	1.894	3.792
s5	1.207	2.899
s6	1.531	3.229
s7	1.611	3.647
s8	0.661	2.492
s9	0.668	2.561
s10	1.418	3.240
s11	1.251	3.862
s12	1.535	3.137

ratio around the part of interest, i.e. where the evoked response is expected to be the largest, so in the [0.2s 0.37s] window. To estimate this signal to noise ratio, the average peak to peak difference divided by the mean standard deviation in the window of interest :

$$SNR = \frac{\max_i(\bar{x}_i) - \min_i(\bar{x}_i)}{\frac{1}{H} \sum_{t=1:H} \sqrt{\frac{1}{N} \sum_{t=1:N} (x_{t,i} - \bar{x}_i)^2}} \quad \text{with } i \in [\tau_t, \tau_t + H]$$

The results are presented in Table 6.2. For all the subjects the realignment based on the estimated time-onsets improves the signal to noise ratio. In particular subjects 8 and 9 that showed an unusual grand average shows a SNR of the realigned trials up to 4 times bigger than the original data which reinforced the hypothesis of desynchronised original data. More interestingly, subjects 8, 10 and 12 that do not show in the original data the usual grand average characteristic to an error related potential their average of realigned data is more similar to the usual ErrP.

At this stage, we have seen that the method seems to recover from the small jitter present in the original data. We verify this fact with the artificially jittered data. The three last columns of Figure 6.8 and 6.9 show the results using the jittered trials. The fourth column contains the grand average and confirms that the jitter introduced was sufficient to distort any alignment in the data. The fifth column is the average on the trials realigned on the time-onset found in the jittered trials. It can be seen that for all subjects the average of the realigned after jittered data is very similar to the average of the realigned original trials. This means that the convergence of the model of the informative pattern is robust to jitter.

Moreover, when compared to the artificial jitter introduced in each trial, the estimated

time-onset is very close. Indeed, if we look at the last column of Figure 6.8 and 6.9, we see that the distribution of the difference between the estimated time onset of the jittered data and the introduced jitter is very concentrated around one value for most subjects. The subjects who do not present this concentration are also the same whose the distribution of estimated time-onset where not very concentrated in the original data. This means that the method is actually recovering most of the jitter that has been artificially introduced in the data, thus indicating the robustness of the estimation method to jitter.

6.5 Discussion

We have seen in the subsection 6.2.1 that the informative part was modelled as a time sequence of independent points. This is reasonable assumption to model short phenomenon which can be expected to have the same waveform of over a short period of time. However, the method developed here is unlikely to recognise long pattern since they would likely contain a certain variability in the latency between the peaks of the pattern. A possible approach could be to use a model of two (or more) template whose onsets would be estimated independently and whose the time latency between each of them could be modelled as well. However the complexity added by this goes beyond the scope of this thesis.

It is worth noting that the hypothesis of the signal modelled as a non-informative part and an informative pattern has consequence on the type of preprocessing that the method can handle. The informative pattern is indeed not added to a background activity and any variability is directly model by the distribution of the non-informative part or by the distribution of the pattern. Consequently, the analysed data should not contain any Direct Current (DC) component. As such, the method can be effective only on high passed signal. However, this is a very constraining limitations as the vast majority of the EEG are carried out on band passed signal.

In this chapter, we have also seen the strength of an approach of sampling to perform the inference of the parameters using an MCMC method. If we had to compute maximum a posteriori estimate of the parameter using a classical optimisation algorithm, when would need to estimate at each iteration the full posterior probability. This would require to compute the product of the likelihoods of the product of informative for all trials. The problem would then to represent a value which many order time smaller than the precision of floating number of a computer. Instead the Gibbs sampler provide a way to converge to the estimation by *sampling* the posterior distribution instead *computing* it.

The estimation of the parameters are obtained after 10000 iteration of the Gibbs sampler which represent takes approximately 1800 seconds per subject on common

desktop computer. It is worth noting that the implementation is not optimised at all, and more importantly, a smaller number of iteration (maybe 10 times less) would be sufficient to obtain the same parameter estimation of very close. The issue here is that the Gibbs sampler need a certain number of iterations to converge to the full posterior distribution and a high number of iteration has been chosen to ensure it. Determining an efficient stop criterion is still a research topic [99, 100], but not the one of this thesis.

The results on jittered trials not only the indicate the robustness of the estimation to the jitter but they also suggest potential application of the method. In this case, the informative pattern model has indeed been estimated on highly jittered data which do not show any recognisable evoked potential. This context is often the case of induced EEG activity whose study in the time-domain could be considered with the approach developed in this chapter. Another possible application of this method could be the study of evoked potential recorded in a complex environment : the complexity often increase the variability of the reaction time of the subject.

To conclude, in this chapter, we have developed a method that estimates the common temporal component in many repetition of the same process. It does so by building a probabilistic model of the temporal sequence of interest and of the rest of the samples so called the non-informative part. Although the method assumes that the sequence of interest has always the same length, it can handle data in which the informative part is located differently (in the temporal domain) in each trial and it can even handle a strong jitter. The problem was in this chapter dealing with a problem of estimation of unaligned pattern. The next chapter will describe its adaptation to a problem of classification.

7 Template based classification

In the previous chapter, we have presented a method to estimate the probability distribution of an unknown sequence that is assumed to be present in all trials of the same phenomenon. This sequence is assumed to appear in different time in all trials although the knowledge of the problem to which the method is applied can often easily provide a prior distribution of where the sequence is more likely to appear. In the end, this method estimates the probability distribution of an asynchronous phenomena present in all trials. But it considers the model of only one asynchronous phenomenon, i.e. only one class. In this chapter, we will see how to adapt the previous approach for a classification problem.

7.1 Estimation in a 2-classes problem

7.1.1 Model

A problem of classification is simply a problem a estimation as described in section 2.4 where the parameter to estimate is a discrete variable which corresponds to the class c of the observed trial $x_{1:N}$.

$$p(c|x_{1:N}) \propto p(x_{1:N}|c)p(c)$$

From this formula, the posterior probability of the class given the measure (the trial) is proportional to the product of the prior of c and the likelihood of the trial under the hypothesis of the class c . If there were only one class, this likelihood would be given by the model described in the previous chapter. So to tackle the problem of classification using an approach similar to the previous one, we need to extend the previous model to make suitable for the case of several classes. For the rest of the chapter, we will restrict to the problem of classification to the case of two classes. Although extending to more classes is straightforward.




FIGURE 7.1 – Graphical model of the signal distribution of the trial. θ denotes collectively the parameters of the distribution of an element.

As before, we will approach the problem by making the dichotomy between the informative and non-informative parts of a trial which is central to the whole approach of asynchronous detection/classification described in this thesis. Like in the chapter 5, we hypothesise that the non-informative part have the same probability distribution in both classes while the informative one has a distribution specific to each class. However from the model described in the previous chapter there are now several aspects that can characterise the specificity of each classes : the distribution of the informative pattern and the distribution of the time-onsets of the informative part in each trial. Figure 7.1 presents the graphical model of the trial signal.

Many evoked potential have their time-locked response to the stimulus that differs with the condition (class). For example, this is the case for error related potential and P300 potentials as we have seen in the section 2.2.1 of chapter 2. It is this time-locked response difference that allows the processing used to show these potentials : computing the average response of all trials per class implicitly assumes that timing of the response does not vary per class and the difference must be studied in term of difference of the shape of the evoked responses. As such, we tend to expect that the most salient difference between the classes will be seen in terms of differences of the informative pattern distribution. We also expect the distribution of the time-onsets τ of the informative parts will not be discriminant.

Under these assumptions, the distribution model of the time-onsets stays the same as it was in the previous chapter since it does not need to be parametrised, only the parameters of the informative pattern become class specific, but the whole estimation procedure remains mostly the same and the parameters of the model are estimated from the full posterior distribution sampled by a Gibbs sampler.

For the parameters of the informative pattern, we have the parameters that depends on the class c (t_c refers to the trials of classes c and N_c the number of trial of class c). So the posterior conditional used in the Gibbs sampler are similar to the one presented in the subsection 6.2.3 adapted to the two class problem. So for each class c , we have :

$$z_{t_c,i}|rest \sim \mathcal{TN}_{[0,+\infty[} \left(\frac{\psi_c^2}{\psi_c^2 + \sigma_{c,i}^2} (x_{t_c,i+\tau_{t_c}} - \psi_{c,i}), \frac{\sigma_{c,i}^2}{\psi_c^2 + \sigma_{c,i}^2} \right) \quad (7.1)$$

$$\xi_{c,i}|rest \sim \mathcal{N} \left(\sum_{t_c=1:N_c} \left(\frac{x_{t_c,i+\tau_{t_c}} - \psi_{c,i} z_{t_c,i}}{N_c} \right), \frac{\sigma_{c,i}^2}{N_c} \right) \quad (7.2)$$

$$\sigma_{c,i}^2 | rest \sim \mathcal{IG} \left(\frac{N_c}{2} + 1, \frac{2}{\sum_{t_c=1:N_c} (x_{t_c,i+\tau_{t_c}} - \psi_i z_{t_c,i} - \xi_{c,i})} \right) \quad (7.3)$$

$$\psi_{c,i}^2 | rest \sim \mathcal{N} \left(\frac{\sum_{t_c=1:N_c} (z_{t_c,i} (x_{t_c,i+\tau_{t_c}} - \psi_{c,i} z_{t_c,i} - \xi_{c,i}))}{\sum_{t_c=1:N_c} z_{t_c,i}}, \frac{\sigma_{c,i}^2}{\sum_{t_c=1:N_c} z_{t_c,i}} \right) \quad (7.4)$$

For the time onsets of the informative pattern, we have :

$$f_1(\mathbf{x}_t, \boldsymbol{\sigma}^2, \boldsymbol{\psi}, \mathbf{z}_t, \boldsymbol{\xi}, \tau, c) = - \sum_{i=1:H} \frac{(x_{t,\tau+i} - \xi_{c,i} - \psi_{c,i} z_{t,i})^2}{2\sigma_{c,i}^2} \quad (7.5)$$

$$f_2(\mathbf{x}_t, m, s, \lambda, \tau) = -\frac{1}{2s^2} \sum_{i=1:\tau, \tau+H:N} (|x_a|_{t,i+1} - \lambda |x_a|_{t,i} - (1-\lambda)m) \quad (7.6)$$

$$\pi(\tau_t | \mathbf{x}, rest) \propto \exp(f_1(\mathbf{x}_t, \boldsymbol{\sigma}^2, \boldsymbol{\psi}, \mathbf{z}_t, \boldsymbol{\xi}, \tau_t, c) f_2(\mathbf{x}_t, m, s, \lambda, \tau_t)) \pi(\tau_t) \quad (7.7)$$

Finally, the parameters of the non informative part are the same with N referring to the total number of trial ($N = \sum_{c=1:2} N_c$) :

$$m | rest \sim \mathcal{N} \left(\sum_{\substack{t=1:N \\ i \notin [\tau_t, \tau_t+H]}} \frac{|x_a|_{t,i+1} - \lambda |x_a|_{t,i}}{N(T-H)(1-\lambda)}, \frac{\sigma_a^2}{N(T-H)(1-\lambda)} \right) \quad (7.8)$$

$$\sigma_a^2 | rest \sim \mathcal{IG} \left(\frac{N(T-H)}{2} + 1, \frac{2}{\sum_{\substack{t=1:N \\ i \notin [\tau_t, \tau_t+H]}} (|x_a|_{t,i+1} - \lambda |x_a|_{t,i} - (1-\lambda)m)} \right) \quad (7.9)$$

$$\lambda | rest \sim \mathcal{TN}_{[0,1]} \left(\frac{\sum_{\substack{t=1:N \\ i \notin [\tau_t, \tau_t+H]}} (|x_a|_{t,i-1} - m)(|x_a|_{t,i} - m)}{\sum_{\substack{t=1:N \\ i \notin [\tau_t, \tau_t+H]}} (|x_a|_{t,i-1} - m)^2}, \frac{\sigma_a^2}{\sum_{\substack{t=1:N \\ i \notin [\tau_t, \tau_t+H]}} (|x_a|_{t,i-1} - m)^2} \right) \quad (7.10)$$

7.1.2 Results

The method has been tested on the same Error Related Potential data set as chapter 5 and 6. In order to test the classification method, the data set has divided (per subject) into a training set formed by the first 2/3 (in the chronological order) of the trials which will be used for training the parameters of the model distribution and the last 1/3 forming the testing set is reserved to test the classification procedure. We will first analyse the results of estimation and later in this chapter present the classification results. The results of parameter estimation, presented in Figure 7.2 and 7.3 show, per class, the grand average of the original trials, the grand average of the realigned data, the distribution of the time-onsets of the informative parts and the distribution of the informative pattern.

It is worth noting from the figures that the distribution of the time-onsets of the two classes are quite different. Only for 2 subjects (8 and 9) the distribution do not differ significantly. For almost all of the others (excepting subject 12), the distribution of the onsets of the non-error classes are less concentrated around a particular time than those of the error class are.

When looking at the distribution of the informative pattern, we note that they do not differs between class : for 4 subjects (1, 8, 11 and 12), they estimated distributions are similar in the 2 classes. For some other subject like the 9, the estimated patterns are just a time shifted version of each other. This problem can occur as explained in section 6.3 due to shift uncertainty in the pattern estimation. However, when considering the onset distribution of the same subject, we notice that if the two template distributions were shifted to match, the time-onset distribution difference would be increased.

At this point, we can conclude that the assumption that the onset distributions are not discriminative is not completely accurate. It is rather the combination of the informative pattern distribution with the onset distribution with allows to discriminate the error class against the non-error class. More importantly we have seen that the distribution of the non-error class onset is less concentrated than the error class. It is worth-noting that this might explain the difference of the grand average of the two classes in the original data and might also explain why this difference seems to disappear in the grand average of the realigned data.

7.2 Classification

7.2.1 Model

We have seen in the previous section that a proper classification method based on the previous approach would require using a model of time-onset distribution specific




FIGURE 7.2 – Results of the parameter estimation on ErrP (subject 1-6) : For all figures, the red plot refers to class error and the blue plot refers to the non-error class. From left to the right, the first column shows the grand average per class, the second the grand average of all trials realigned (based on the detected time onset for each trial), the third column contains the distribution of the detected time onsets. The fourth column contains the probability distribution of the informative pattern (the more intense is the color, the more probable is the value).




FIGURE 7.3 – Results of parameter estimation on ErrP (subject 7-12) : For all figures, the red plot refers to class error and the blue plot refers to the non-error class. From left to the right, the first column shows the grand average per class, the second the grand average of all trials realigned (based on the detected time onset for each trial), the third column contains the distribution of the detected time onsets. The fourth column contains the probability distribution of the informative pattern (the more intense is the color, the more probable is the value).




FIGURE 7.4 – Examples of beta distribution with different parameters.

for each class. From is shown in Figure 7.2 and 7.3, the time-onset distribution are unimodal and by construction of the model, bounded by the trial limits (diminished by the length of the informative part). So an immediate candidate for the model distribution would be a beta distribution. The model described in the previous section could be modified to integrate a model of beta distribution for time-onsets per class. However, the beta distribution suffers from discontinuities when the parameters varies (for example when the α parameter crosses 1 : see Figure 7.4). As usual, in a Bayesian approach, the way to mitigate this problem is to provide a prior on the parameters. However in this present case, the choice of prior on these parameters is not necessarily very intuitive. So for the sake of simplicity, we will keep the model as it has been trained previously and the distributions of the onset will be estimated after the model parameters have been trained on the training set.

So the parameters of the onset distributions are estimated for each class by the maximum likelihood estimates. Moreover, in order to favour a distribution centred on one value, i.e. favour a distribution whose probability in the boundaries tends to 0, we estimate the parameters on the 80% onsets closest to the mean onset. By doing so, we simply eliminate the influence of the outliers on the parameter estimation.

Once estimated on the training set, the onset distribution per class will replace the prior $\pi(\tau)$ used in the Bayesian estimation procedure by $\pi(\tau|c)$ a distribution that depends on the class c . The class of one trial \mathbf{x} of the test set is then easily estimated by doing an inference from the model trained :

$$p(c|\mathbf{x}) \propto \int_{\tau} \prod_{i=1:H} p(x_{i+\tau}|\xi_{c,i}, \omega_{c,i}^2, \alpha_{c,i}) \prod_{j=1:\tau, \tau+H:T} p(x_j|m, s, \lambda, x_{j-1})\pi(\tau|c)d\tau$$

where $p(x|\xi_{c,i}, \omega_{c,i}^2, \alpha_{c,i})$ is the likelihood of the sample x with the i -th point of the informative pattern of the class c (ω^2 and α are calculated back from ψ and σ^2 as

TABLE 7.1 – Classification results on test sets (MCC values)

subjects	original trials			jittered trials		
	LDA	rLDA	Bayes	LDA	rLDA	Bayes
s1	0.647	0.380	0.289	-0.056	0.221	0.266
s2	0.411	0.260	0.226	0.028	0.178	-0.024
s3	0.629	0.422	0.420	-0.110	0.357	0.357
s4	0.418	0.417	0.320	0.035	0.275	0.317
s5	0.228	0.317	0.230	0.120	0.175	0.329
s6	0.233	0.353	0.317	0.150	0.210	0.348
s7	0.293	0.291	0.307	0.051	0.254	0.315
s8	0.147	0.175	0.070	0.013	0.200	0.174
s9	0.226	0.048	0.043	0.123	0.086	-0.069
s10	0.323	0.470	0.419	0.070	0.358	0.374
s11	0.470	0.501	0.453	0.010	0.312	0.475
s12	0.329	0.303	0.299	0.108	0.328	0.221

explained by equation 6.4 in subsection 6.2.3) and $p(x_j|m, s, \lambda, x_{j-1})$ is the likelihood of the j -th sample of the trial to be a non informative sample. Using this equation, we can infer the class of an observed trial and thus test the classification based on the model.

In order to compare the method with the state of the art, each trial has been classified by an LDA classifier using the same data but downsampled to 64Hz. Another classifier has been constructed by realigning each trial based on the maximum peak in the [0.2s, 0.5s] window of each trial and then classify the data using a LDA as before. We will refer this method as the realigned Liner Disriminant Analysis (rLDA) method latter in the rest of the chapter.

7.2.2 Results

The results of classification presented in Table 7.1 correspond to the classification on the test set expressed in MCC values. The results have been presented on both original data and the jittered version as described in section 6.4 : an artificial jitter of +/- 0.2s is applied independently on each trial of the data set and the classification method is trained and tested on this jittered data set.

When considering the results on the original trials, the classification method developed in this chapter performed above 0.2 for all subjects excepting for subjects 8 and 9. It is worth noting that these are the same subjects for which the estimation of the distribution of the time onsets nor the distribution of the informative part do not differ greatly between each class. When compared with the other methods for the same original data set, the LDA method performs the best for all but 3 subjects (s6, s5 and

s10). Concerning the rLDA, although it performs slightly better than our method, the results are comparable.

Concerning the results on jittered data, we observe a drop of performance for all methods which is expected. However the drop is not the same for all of them. Considering the performance of the LDA method, the classification is now at the random level for 8 subjects and it is not very high for the remaining 4. The rLDA method still classifies the data of most of the subjects non randomly, however, when compared to the performance on the original data, the drop is quite high for most subjects.

However when compared the performance of the method developed in this chapter on the original data set with the performance on the jittered data, we observe a relatively small drop of performance. But more importantly, it can be observed that this method outperforms the others. For 9 subjects it performs as good as the rLDA method and for 7 of them, it performs much better.

If we consider our method and the rLDA compared to the classification results of the LDA, it can be concluded that the realignment procedure of the rLDA or the one inherent to our method have positive effect for classifying jittered data. However, it should be noted that contrary to the rLDA, our method does not know in advance how to realign data trials : it learns what the common part in the training set and the realignment is implicitly done on this common part. It is not the case for the rLDA method where it is known in advance that the realignment should be based on the maximum positive peak in the range of [0.2s 0.5s]. It is an interesting property that the method developed in this chapter has and which can be important when applied on a new problem on which little is known.

7.3 Discussion

The method for dealing with asynchronous data introduced in this chapter shows for some subjects a similar pattern of ErrPs while the distributions of the time onsets differ (bigger variability in the time onsets of the non-error class). Such appearance opens a discussion. It might be possible, that the neural circuitry (involving the anterior cingulate cortex, that acts as a comparator of expectancy and outcome [101, 102]) responds in a similar way for both error and non-error conditions but would respond in a more time locked manner when there is an error. This would lead to a clear ErrP on scalp recorded EEG as an ERP, whereas the response for the non-error would not be time locked. It is an hypothesis that could be investigated in future.

The Bayesian method we have developed assumes a skew normal distribution for the points that form the informative pattern. As a generalisation of the Gaussian distribution, it allows to model more different shapes of probability density functions. It

could be argued that the increase of the degree of freedom could lead to an overfitting of the skew parameter of the distribution. However the results of estimation show that it is not the case since for most part of the patterns, the estimated skewness is rather low. The high skewness of mostly found on the peaks of the pattern as it was expected (particularly for subjects 4, 9, 7, 10 and 12). Only subject 5 present an estimated skewness contrary to the expectation (small negative skewness on positive peak) although its pattern model looks globally like an ErrP. So finally, these results indicates that the skew normal distribution model was a suitable choice for an EEG waveform.

The initial approach of not modelling the time-onset distribution have shown not to be optimal. So a class-specific model of beta distribution has been used. To keep the simplicity of the estimation procedure, their estimation have been performed after the estimation of the other parameters. However it is worth noting that this manner of estimating them loose partly the advantages of Bayesian approach : if we had estimated them along the other parameter, it would have been possible to provide a more flexible prior. While we currently can provide a prior of where the time-onsets are located in the trial, estimating the time-onset parameters through a Bayesian approach would have allowed for example to provide a prior on the variability of the time-onsets irrespective to their mean location. This should be considered for the future development of this method.

The results of the classification shows that when applied on jittered ErrP data set, the method outperform the combination of a synchronous classifier (LDA) with a realignment step. Moreover, the realignment step was possible because it is known that the evoked response can be realigned on the maximum peak. In contrast, the method developed in this chapter do not need this information. It was capable to learn a common pattern in all trials of each class and then performed the classification based on these pattern models. This indicates that the method could be used to approach classification problem on data set on which little is known. As such this method could be used as an exploratory tool to study new paradigms.

To conclude, in this chapter, we have developed a method that is capable to classify trials of EEG signal that are not tightly synchronised with the same time information, i.e. whose the relevant information does occurs in the time location in all trials (even of the same class). This is done by assuming that in each trials there is a sequence (and only one) that is specific to each class. This sequence is then modelled in a probabilistic way and a Bayesian method learns the parameters of the distribution of this sequence (for each class) as well as the time of occurrence of this sequence in each trials.

We have also seen that using this approach, we can classify ErrP data with a performance a little bit worse than synchronous classifier like LDA which assumes perfect synchronisation of the trials. However our method performs better when the data a

Chapitre 7. Template based classification

slightly jittered which makes it more suitable for real application when the synchronisation with the stimuli are less controlled.

8 Exploration/exploitation revisited

8.1 Introduction

So far, we have developed a classification method that operates when the outcome of an experiment (a trial) can be expressed as a one dimensional sequence that contains a class-specific pattern whose time-latency can be different in each trial and some parts that is irrelevant to the class. In that context, the method learns the probability distribution of each class-specific template and the probability distribution of the non class-specific part (assuming a model of coloured noise). The method has been tested previously with Error Related Potential in which we artificially increased the jitter.

We now assess its performance on the data recorded with the Exploration/Exploitation paradigm as described in chapter 3. In chapter 4, we have seen that it is actually possible to find EEG correlate of exploratory behaviour and that it is an asynchronous pattern (with respect to the machine selection). The method developed in the previous chapter aims at classifying an asynchronous EEG pattern based on the recognised of an class specific pattern in the recording of one channel. Since it is not known a priori what are the electrodes that could carry such pattern relevant to the exploration/exploitation contrast, we apply the method channel by channel in order to evaluate it.

8.2 Channels selection and prior

To apply the method developed in the previous chapters, we need to know which electrodes will be studied, what is the length of the template that will be used, and ideally, what is the prior we can use for the onset of the informative template in each class. For this, we will study the exploration exploitation contrast in a descriptive manner. We will try to identify the biggest differences in terms of time, frequency, electrodes distribution between the two classes. To this end, we use the same features used in chapter 3 (log-transformed frequency power computed with a continuous wavelet transform).

Instead of classifying the contrast based on a time-locked as done in chapter 3, we are interested in knowing which electrodes and where in time-frequency contributes to the biggest differences between the two classes. For this we have computed a DP value in the following way :

$$DP(t, f, e) = \frac{|\mu_1(t, f, e) - \mu_2(t, f, e)|}{\max(\sigma_1(t, f, e), \sigma_2(t, f, e))}$$

where $\mu_i(t, f, e)$ denotes the mean log-transformed amplitude over all trials of class i at time t , frequency f and electrodes e and $\sigma_i(t, f, e)$ its standard deviation.

Using this measure, we have computed the DP maps for all subjects but the results of only 3 subjects are presented in figures 8.1, 8.2 and 8.1. If taken independently, none of the triplet (time, frequency, electrode) show a significant difference (empirically a value superior to 1) which would have been surprising given the results of chapter 3. However, the DP maps present an evolution over time and location consistent over subjects. For 5 subjects, we notice that, the main difference can be seen around 10 Hz (particularly visible for subjects 2 and 8). This is consistent with the results reported in Table 4.1 of the chapter 4. Moreover, the difference steadily increase over time and are the most important in the parietal areas (bilaterally). Some studies suggest that lateral intraparietal areas play a role remapping abstract valuation to concrete action [103] which could be relevant in the decision to make exploratory or exploitative action. The importance of the 17.5 Hz component in the classification of exploratory behaviour (as shown in Table 4.1) for some subjects (like subject 8) can be seen as well but seems less important compared to the 10 Hz component.

Given this information, we are able to design the asynchronous classification method we are going to apply on the data recorded in the exploration/exploitation paradigm. First we will apply the method mainly on parietal electrodes (CP3, CP4, P3, P4, P5, P6, Pz) and some frontal electrodes (F3,F4,F5,F6,Fz,FCz) since the frontal areas are suspected to play a role in decision-making processes. Secondly, the template length will be set to 0.3s (in a analysis window of 0.9s) to allow the template to contain several cycles of a 10 Hz signal. As we have mentioned in section 6.3, there is less impact on the results by using a template too large rather than too small. Finally, from the deviation maps we have seen, it is a priori expected that the most informative parts to the exploration-exploitation contrast are located at the end of the trial, i.e. close to the machine selection (see Figure 8.4).

8.3 Classification results

The classification method described in the previous chapter has been applied on the data recorded with the 4-slot machines paradigm as described in chapter 3. It uses the same analysis window as before (from 1s to 0.1s prior to the machine selection). The




FIGURE 8.1 – DP : subject 2. The color scale is normalised to the difference between the biggest and smallest value of the all maps of the subject.




FIGURE 8.2 – DP : subject 4. The color scale is normalised to the difference between the biggest and smallest value of the all maps of the subject.




FIGURE 8.3 – DP : subject 10. The color scale is normalised to the difference between the biggest and smallest value of the all maps of the subject.

Chapitre 8. Exploration/exploitation revisited

TABLE 8.1 – Classification results on different electrodes (MCC values) : the values in bold notifies a couple (subject-electrode) that shows a classification that is significantly better than random.

subjects	Fz		
	Bayes	LDA	rLDA
s1	-0.057	0.094	-0.062
s2	0.044	-0.007	-0.000
s3	0.038	-0.032	-0.040
s4	-0.035	0.064	-0.022
s5	0.066	-0.030	-0.018
s6	-0.020	-0.036	-0.081
s7	0.075	-0.045	0.102
s8	0.132	0.082	0.021

subjects	F5			F3			F4			F6		
	Bayes	LDA	rLDA	Bayes	LDA	rLDA	Bayes	LDA	rLDA	Bayes	LDA	rLDA
s1	-0.053	0.005	-0.070	0.071	-0.006	-0.074	-0.049	0.001	-0.030	-0.057	-0.007	-0.099
s2	0.020	0.096	0.048	-0.018	0.059	0.062	0.135	-0.024	-0.047	-0.013	-0.066	-0.045
s3	-0.021	-0.051	0.043	0.120	-0.067	-0.026	0.018	-0.052	0.028	-0.094	-0.016	0.033
s4	-0.042	0.110	-0.034	0.099	0.003	-0.036	0.052	0.052	0.091	-0.085	-0.006	-0.070
s5	0.021	0.021	0.025	0.015	-0.015	0.032	-0.024	-0.003	-0.054	0.132	-0.087	-0.117
s6	0.025	-0.013	-0.001	0.033	-0.123	-0.042	0.113	-0.017	0.000	0.051	0.005	0.046
s7	0.074	-0.003	0.092	0.119	0.013	0.073	0.062	-0.000	0.061	0.023	-0.052	-0.109
s8	0.002	-0.067	-0.047	0.065	-0.098	0.005	0.053	0.090	-0.028	0.098	-0.009	0.005

subjects	FCz		
	Bayes	LDA	rLDA
s1	-0.059	-0.016	-0.061
s2	0.085	-0.037	0.028
s3	-0.079	-0.042	-0.084
s4	0.082	0.062	-0.032
s5	0.053	-0.062	-0.054
s6	-0.047	-0.187	-0.093
s7	0.067	-0.069	0.039
s8	-0.013	0.030	-0.018

subjects	CP3			CP4		
	Bayes	LDA	rLDA	Bayes	LDA	rLDA
s1	0.014	0.091	0.014	0.024	0.071	0.043
s2	0.136	-0.017	0.044	0.149	-0.017	0.053
s3	0.073	-0.024	-0.002	0.059	-0.119	0.014
s4	0.097	0.094	-0.034	-0.072	-0.052	0.160
s5	0.166	0.045	0.020	0.053	0.019	0.060
s6	0.092	0.085	0.016	-0.057	0.010	-0.010
s7	-0.003	0.086	0.122	0.048	0.016	0.010
s8	0.242	-0.103	0.013	0.086	-0.064	-0.019

subjects	Pz		
	Bayes	LDA	rLDA
s1	-0.055	0.036	0.012
s2	-0.012	-0.046	-0.011
s3	0.060	-0.040	-0.070
s4	0.115	0.034	0.152
s5	0.110	-0.083	-0.012
s6	-0.022	0.010	-0.080
s7	0.040	0.038	0.124
s8	0.156	0.142	-0.016

subjects	P5			P3			P4			P6		
	Bayes	LDA	rLDA	Bayes	LDA	rLDA	Bayes	LDA	rLDA	Bayes	LDA	rLDA
s1	-0.053	-0.001	0.013	-0.064	0.113	0.134	-0.071	-0.009	-0.056	0.010	0.040	0.102
s2	0.202	-0.035	0.030	0.095	-0.137	-0.040	0.034	-0.037	-0.033	0.034	0.003	0.066
s3	0.090	0.004	-0.034	0.053	0.071	0.037	0.111	-0.119	0.068	0.045	-0.109	0.063
s4	0.020	0.083	-0.060	0.097	0.066	0.081	-0.010	0.086	0.081	-0.059	0.072	-0.053
s5	-0.005	0.008	-0.019	-0.062	0.008	0.028	0.079	0.033	-0.015	0.032	-0.048	-0.044
s6	0.047	0.064	0.038	0.112	0.053	0.107	-0.035	0.079	0.071	0.104	-0.041	0.023
s7	0.023	-0.057	-0.096	0.002	-0.085	-0.164	-0.002	0.095	-0.095	0.014	0.114	-0.074
s8	0.107	-0.105	0.087	0.306	-0.014	0.044	-0.043	-0.008	-0.020	0.042	-0.020	-0.009




FIGURE 8.4 – Prior on time onset (mean : -0.65s, standard deviation : 0.1).

EEG signal has been downsampled to 128 Hz and filtered with a band pass of [1Hz - 30Hz]. The method has been applied on this signal in the time-domain contrary to the method described in the chapter 4. For the sake of comparison the same LDA and rLDA classifier have been tested on the window [-0.4s -0.1s] for LDA and [-0.6s -0.1s] with a classification window of 0.3s for rLDA. Like in the previous chapter, for each subject, the data has been divided into a training set formed by the first 2/3 of the trials (in the chronological order) and the last third formed the testing set. The results are presented in Table 8.1.

For all subjects many electrodes present a MCC value on testing set lower than 0.1. However, we can find for all subjects but one (subject 1) one or more electrode that show classification results better than random. In addition, the Bayesian method performs globally better than LDA and rLDA : LDA reports only 4 subject-electrode pairs significantly better than random, rLDA reports 7, while the Bayesian method reports 19 pairs. The pairs that have the best classification accuracy with the latter method (above 0.2) are even completely missed by LDA and rLDA.

The electrodes that show the best classification performance are mainly F3, F4, Pz, CP3, P3 and P5. Moreover, the 3 subjects that show the best classification accuracy (subjects 2, 5 and 8) show their best performance in the parietal area (P5 for subjects 2, CP3 for subject 5 and P3 for subject 8) and all of them present non random classification in CP3. This suggests the importance of the parietal area for the contrast between exploration and exploitation. For comparison, the classification results obtained in chapter 4 show the best classification performances for, in the decreasing order, subjects 8, 2, 5 and 3 (results using the combined classifier in Table 4.1). It is worth noting that the subjects that present the best classification accuracy are the same as the ones in the classification results of chapter 4. This is case in spite of the fact the two methods are not applied on the same features (signal in the time-domain here while the first method uses spectral density). The fact that for a single subject, different electrodes present non-random classification result seems to indicate that extending the method to learn informative pattern combining the informative from different channel could be promising.




FIGURE 8.5 – Estimated templates and onset distribution for subject 8 : in red, class exploration, in blue exploitation.




FIGURE 8.6 – Estimated templates and onset distribution for subject 2 : in red, class exploration, in blue exploitation.




FIGURE 8.7 – Estimated templates and onset distribution for subject 1 : in red, class exploration, in blue exploitation. (electrodes CP3 and FCz).

To analyse what has been identified as an informative pattern, we will concentrate on the subjects that present the best classification accuracies (subject 2 and 8) as well the one who presents the worst results (subject 1). The estimated template distributions and onset distributions of the 3 best electrodes in term of classification (CP3, P3, Pz) and one presenting bad classification result (FCz) for subject 8 are presented in Figure 8.5. Those for CP3, P5, CP4 (best) and F5 (bad) for subject 2 are presented in Figure 8.6. The results of CP3 and FCz for subject 1 are presented in Figure 8.7.

We notice from the figures that the best classification accuracies have been obtained when the template for class exploration converges to a clear sine wave pattern. It is worth noting that for both subject, the pattern are very similar and correspond to a sine wave of approximately 10 Hz (corresponding to an alpha rhythm). For the sake of comparison, the best classification performances using the method describes in chapter 4 for subject 2 and 8 were obtained with respectively 10, 23.3 Hz and 17.6, 10, 23.3, 7.5 Hz. It should be noted that the results obtained in the chapter 4 were based on all electrodes together and the classification procedure used a rejection step that could lead to classify some trials as unknown.

It might be argued that the 10 Hz pattern estimated by the Bayesian classifier is simply an artifact of the convergence of the method. After all, the method is trained on a signal that is band passed so it contains inherently a superposition of sine waves contains in the band pass. However, we can rule out this hypothesis easily by looking at the other electrodes. The electrodes that do not show a good classification (FCz for subject 8 and F5 for subject 2) do not present this clear sine wave either : we can notice a lower frequency component (around 6 Hz) for the electrode FCz in Figure 8.5 and more variable. In the case of F5 for subject 2 (Figure 8.6), the template distribution presents a sudden high frequency component (around 25 Hz), which seems shared by the class exploitation. This might explains why it has not been identified as discriminant (random classification). Even worse the estimated template distribution in Figure 8.7 do not present any strong rhythmic component. This seems indicate that the 10 Hz pattern is effectively relevant in the contrast between exploration and exploitation.

When observing the estimated onsets for the electrodes that have the best classification, we notice that, although there is a slight bias toward the time of machine selection, their distribution is spread over the whole analysis window. This indicates that the prior used to estimate the pattern do not prevent an informative pattern to be estimated in the least likely regions according to the prior.

8.4 Discussion

In the chapter 4 we have seen that frequencies higher than alpha band playing a role in the classification of the exploration-exploitation contrast. However we have not seen

them in any of the electrodes showing the best classification accuracies. We might have discovered one the limitations of the method here. The Bayesian method developed here is based on the search of a particular pattern found in most of the trials of the same class. But it cannot be excluded that relevant pattern is composed of two frequency components whose one has an amplitude and variability bigger than the other. In this case, unless the components are perfectly synchronised, it is likely that the method will converge to the most prominent and will "integrate" the smaller component in the variance of the probability distribution. In such particular case, a way to deal with the problem could be to decompose the input signal into different frequency component through a filter bank. The choice of the frequency band would need to be carefully chosen in order to avoid to be centred on the relevant components which is not known in advance.

In order to have a fair comparison of the classification results in this chapter with those obtained in Chapter 4, we would need to estimate the informative pattern using the data from several electrodes at the same time. However, since the model in this chapter assumes the independence of the points of the informative pattern, the likelihood of the observation against the pattern model is represented by small values. Using several electrodes would lead to numerical problems in the current state of the method. Another alternative could be to do fusion of classifiers using a voting schema. However, this is neither feasible because of the number of classifiers. The suitable approach would be to change the pattern model to drop the independence hypothesis. This is a development of the method that is considered for the future.

Another study have recently analysed the EEG correlates of the exploration/exploitation contrast using the same experimental protocol [76]. They reported for trials leading to exploration, there was significantly higher activity in dorsolateral prefrontal cortex and the right supra-marginal gyrus. The classification results in this chapter are consistent with frontal activation reported by them. But our classification results emphasises more on the involvement of the left parietal areas. However, it is worth noting that their analysis window was performed post feedback display, while ours focused on the 1s window prior to the machine selection (see Figure 3.1 in the chapter 3 an illustration of the protocol). But the two analysis windows overlap so this fact does not explain all the discrepancy of the results. A possible explanation could be that the difference they found is mostly due to beta or gamma activity which may be masked in our method by lower rhythm (like alpha) if they can be found consistently in trials of the same class.

As a conclusion, we have seen in this chapter that the Bayesian classification method for asynchronous EEG patterns developed in the previous chapter can be applied on the data recorded in the 4-slot machine paradigm in order to classify the exploration/exploitation contrast. Although the performance is not very high, we can find for most subject several channels that classify an exploratory decision better than random. Moreover the method outperforms more classical methods like LDA or LDA

with a realignment procedure. In addition, while able to classify, the method has allowed to identify areas that are likely to play an important role in the exploratory decision-making : the left parietal areas is to provide the best classification performance. Furthermore, for these channels, the informative template has been estimated as a salient sine wave of 10 Hz, indicating the role of alpha rhythm for this process.

We have also seen that, since the best performances per subject, are scattered over several channels, it is expected that a method adapted from the current one but analysing several channels and forming an informative pattern based on these channels, will provide better results. This line of work should be considered later. In addition, the study in this chapter and the comparison with the results of chapter 4 has allowed to spot a limitation of the method : it converges toward the most prominent informative pattern it can find and thus may ignore other components. It is an issue that should be tackle in the future.

9 Conclusion and future work

This thesis has presented methods to analyse EEG signal, that are not time-locked to observable events. This thesis proposes an approach to handle non-time-locked EEG patterns by making the hypothesis that, in each trial, only a part of the signal contains the relevant pattern of interest. This relevant part can appear at any time in the analysis window and differently for each trial. The rest of the trial corresponds to a non-informative part irrelevant to the targeted cognitive task.

9.1 Discussion

First we showed the need of an asynchronous method for classification of decision-making process. One of such cognitive process that is likely to have non-time locked brain activity is the switch exploratory and exploitative behaviour. The N-slot machines is a classical paradigm used in reinforcement learning theory to exemplify the conflict between exploration and exploitation. To this end, we reproduced a experimental protocol (described in chapter 3) that was used in a study using fMRI to find the neural substrates of the exploratory decisions in humans [7]. Decision making process are not easy to study given the fact that observations that the experimenter make do not access directly to the belief-system on which the subjects base their decision. Thus it often required to build a behavioural model of the subject whose parameters are estimated from the observable action of the subjects. This is the case for the exploratory behaviour and studying it required us to build such model. Once built, it allowed us to label each observed machine selection as an exploratory or exploitative decision and then to study their EEG correlates. Using a time-locked approach, the results showed the impossibility to find classification performances significantly better than random. We then concluded in chapter 3 that the EEG correlates (if they existed) of the exploratory decision were not time-clocked to the machine selection, thus justifying the use of asynchronous classification method.

Chapitre 9. Conclusion and future work

In order to deal with the asynchronicity of the data, we propose a method (see chapter 4) which studies the patterns of the EEG signal over the scalp by considering the distributions of each class (exploration and exploitation) as a “bag of time-samples”, meaning that any information about the time location of the sample within the trial is discarded. The rationale of this approach is to detect regions in the feature space that correspond to the most discriminant samples irrespective of the appearance time within each trial. For this we computed the canonical space of the data discarding the time information and select the tails of the distributions to form the informative set for each class. The classification of a trial then used a voting scheme based on the number of identified informative sample of each class allowing unknown trial if needed : in a given trial, the definition of the informative sample do not assure that any of them can be found or that there will not be the same number for each class. In such cases, the trial was classified as unknown. This classification method was applied per frequency on the spectral density of the EEG signal. With this approach, we could show that it is possible to classify the EEG correlates of the exploratory behaviour. In addition, we identified frequency bands (from 10 Hz to 30 Hz for most of the subjects) that provide a classification performance better than random. Furthermore, we saw that by combining the output of the classifiers that we could achieve better result (mostly reducing the unknown) thus indicating that the correlates of exploratory behaviour were not restricted to one particular band.

However this method was prone to over-fitting, as shown by the drop of performance between the train and the test set, and the choice of the limit of the percentile that defined informative set of each class is difficult to determine. To overcome these issues, we propose in chapter 5 to turn the previous method into a generative model. The different aspects of the method (distribution of the informative and non informative sample, probability that an informative sample appear) were modelled, resulting into a decision rule that had a better formalism than the voting scheme of the previous approach. To test this new approach, we used the data of an evoked potential whose waveform is well known : the Error Related Potential (ErrP). Using ErrP data present the advantage that the time of appearance of the evoked response are well known and thus allows us to validate the method.

The use of a generative model slightly increases the number of parameters to estimated compared to the previous approach. However these parameters are also easier to understand from the problem and prior could be provided to help the convergence to the correct parameters. This is an advantage of the Bayesian approach. The results showed that we could classify EEG evoked potentials with this approach without relying on strong time cues. However its performance is generally worse than the classical synchronous approach because of high rate of false positives. But a new improvement to the model by modelling the shape of the waveform could lead to better results.

In order to take advantage of shape of the waveform, we proposed a new model in

chapter 6 keeping the dichotomy between informative and non-informative part in a trial. Instead of taking the informative time samples independently in the trial, it assumes they form a particular pattern in time that the method tries to identify while allowing this pattern to appear anywhere in the trial, while the non informative sample can be distributed by a different model. This method could be seen as a Bayesian asynchronous template matching. At this stage, we first develop the method to handle only one class. After testing on synthetic data to ensure the proper convergence of the approach, we tested it on the ErrP data. By initialising the method to a flat pattern, it converged correctly to the expected waveform of an ErrP. After realignment each trial on the estimated time onset of the informative pattern, we were able to increase the SNR around the waveform of the ErrP by up to 4 times. More interestingly, when a different random time jitter was applied to each trial, the method was able to converge to the correct waveform, thus recovering most of the jitter artificially introduced in each trial.

We then adapted this method to a problem of classification in chapter 7. We built the new method by assuming a different template distribution for each class, while assuming the distribution of the non-informative part will be the same for all classes. After observing that the time-onset distributions were class dependent, we learnt a model for each one based on the training set. We then tested the method again with ErrP data. When applied on usual conditional, we showed that the performance were slightly lower than classical method which assumes perfect synchronisation of the trials. However when introducing jitter of [$\pm 0.2s$], the performance were better than traditional methods. Interestingly, the difference of classification performance of our method on the jittered data compared to the normal one was small.

In the chapter 5, we have assumed the independence of observation given the class of the trial, or in chapters 6 and 7, the independence of the point of the informative. This is often referred as the naive Bayes hypothesis in the literature [104]. This hypothesis simplify the inference since the estimation of the parameters can then be done independently. However, this hypothesis is often a simplification of the reality : when processing band filtered signal, it unlikely that two consecutive time-sample value differ highly. This may have the side effect to underestimate the likelihood of the observations against the model. This usually has little influence on the classification result because the underestimation is consistent over all the possible values. However this may result in a bad estimation of the posterior probability, which for the case of classification, tends to estimated probability close to 0 or 1. This undermine one of key advantages of using a generative Bayesian model. If its hypotheses are verified, the use of a generative model for classification gives indeed a true estimation of reliability of the classifier output which could be reused in a larger system. In order to use the full strength of a Bayesian approach, the hypothesis of independence for be dropped whenever it is not necessary. As consequence, it is highly desirable to improve the model of the informative pattern to take into account the local dependence of the

time-sample on their neighbours.

Applying the method on multichannel data could provide interesting information about the brain dynamics while performing a certain cognitive task. While it is currently limited by the aforementioned issue, solving it could widen the scope of possible analysis that the method could perform. Applying the method on multiple channels could indeed reveal the phase synchrony between brain areas relevant to a specific cognitive process. Alternatively, modelling the latencies between the informative pattern in each channel could be considered.

We finally tested this method of Bayesian asynchronous pattern classification on the data of the exploration/exploitation paradigm (chapter 8). By applying the method to electrode at a time, we were able to identify for 7 subjects out of 8 several channels that classify an exploratory decision better than random. The method allowed to identify areas that are likely to play an important role in the exploratory decision-making : the left parietal areas is to provide the best classification performance. Furthermore, for these channels, the informative template has been estimated as a salient sine wave of 10 Hz, indicating the role of alpha rhythm for this process. This study allowed us to also identify some limitations of the method that must be addressed in future.

More generally, the method in its current state could be used as an exploratory tool to study EEG activity in new experimental paradigm. The method become more useful the less are known about the targeted brain signal. Thus, it could be used for neurophysiological research on single trial. By applied it on band passed EEG signal (not necessarily very narrow), it should be able to spot the most relevant burst of activity that are the most relevant for task. As such it could used in the study of the EEG correlates of perception. For example it could help to study the induced oscillatory activity of face recognition.

It is worth noting that the methods developed in this thesis are not limited to the EEG signal. Although we have tested it only on EEG, the model assumes only certain signal dynamics that could be shared by different bio-electrical signals. That is why, it can be considered to be applied on ECoG, LFP, MEG or electrocardiography.

Finally, the methods developed in this thesis are particularly suitable for BCI applications. In complex environment, the computer system, with which the subject interacts, needs to analyse the scene to relate the measured brain activity with the current context. Depending on the complexity on the scene, this may results into variable latency that might affect the interpretation of the brain activity. The methods developed in this thesis provides an approach to overcome these issues.

9.2 Perspectives

When compared with other classification methods applied on Exploration-Exploitation data like [76], the methods developed in this thesis mainly differ in how they search for the relevant EEG correlates of the task. The former tries to characterise the correlates in terms of scalp distributions of the EEG signal. On the other side, the method described in the thesis aims at characterising them in terms of temporal signatures. Both approaches represent the state of the art and, up to now, there is not ground to prefer one over the other. Concerning the method developed in this thesis, we have already seen in the discussion of chapter 8 that it provides a theoretical framework allowing the analysis using both characterisations (temporal or scalp distribution). However, the assumption of independence of the points of the template prevents, from a numerical point of view, to scale up to many electrodes. This is an issue that needs to be tackled in future.

In a broader perspective, one might wonder what our method brings to the field of EEG pattern analysis and classification. We have seen that it handles asynchronous patterns better than classical approaches. But, in its current state, the method does not seem to provide better performance when analysing synchronous patterns. Nevertheless, our approach may still be beneficial. By its Bayesian nature, it provides a natural framework to incorporate prior knowledge. Of course, depending on the complexity of the type of prior, the model may need to be modified. Easy and principled, this should be compared to the case of classification methods based on discriminative approaches like LDA, QDA or SVM, which offer no means to incorporate priors except for the class probability. Even more complex methods based on neural networks like in [4], leave the experimenter helpless for such an integration. When considering the fact that EEG recording is usually expensive in the sense that it generates a large amount data (dimensionality) compared to the number trials, providing a way to incorporate priors is definitely valuable.

Another interesting aspect of the method is the fact that it is a generative approach. As such, the probability distribution of each individual part of the EEG signal model of a trial is estimated. Thanks to this modular estimation, and its combination with the Bayesian approach, it is possible to reuse parts of the estimated model and integrate them for the estimation of the parameters of a new and more complex model. This is valid if the new model does not depart radically from the initial model. In such a case, what has been learnt (estimated) once is not lost and can form the prior for a more refined model.

Modelling the non-informative is necessary for building a generative approach. Under a more classical approach (often discriminative), the model of the informative part would simply be tested along the whole trial signal to determine the time-onset of the informative part in the trial. In our approach, the model of the non-informative

part is used as an additional test of the context in which the informative part has been found, i.e. if the rest of the trial *does* resemble the usual EEG signal. This measure of the likelihood of the non-informative part combined with the model of the informative part could be used to indicate whether the overall estimated model still explains the data and to determine when the model is not valid any longer. This is crucial for online adaptation to the brain patterns we are looking for [105] and may open new research avenues in this field.

We can conclude that Bayesian generative approaches require an extensive modelling of the different parts of the analysed signal. This task is usually not trivial which makes these approaches less appealing for the experimenter. However, we have seen that they also offer new advantages for EEG processing. In this respect, this thesis has contributed to provide the foundations of a model-based approach for the analysis of brain signals (EEG and other) that can be used and extended to build more refined and better performing methods in the future.

9.3 Future work

The work of this thesis makes essential contributions to build a generative model of classification of asynchronous EEG patterns. Building upon the results presented in this thesis, future work can be devoted in different areas.

Independence of the points in the pattern

In its current state, the model of the informative pattern assumes that the points of the sequence are independently distributed. Although it is an assumption that simplifies the calculus of the posterior conditional probability of the template, it is not exactly the case in reality. We indeed often apply the method on signal that are low pass filtered. This implies that given the value of one time-sample of the EEG signal, it is likely that the next time-sample will be close it. This has undesirable effect : the likelihood probability of a pattern is often an underestimation of what it should be compared when using a model taking into account the dependency on the neighbourhood. To make the best use of a Bayesian model, one must use the most accurate model as possible (this does not prevent it to be highly parameterisable). A possible model for the template could be a mix between of its current state with a stochastic model of EEG signal.

Using multiple channels

From the modelling point of view, extending the current model to multiple channel is very easy. We could even additionally add mixing a factor that could represent the

relevance of each channel in the identification of the informative part. However the problem do not lie in the theory neither the calculus of the conditional posterior probability but in the computational aspect. As mentioned in the previous point, in absence of a model that do not take into account the dependence of points of the informative pattern on each other, the likelihood of the template can vary very much (more than it should) towards very small values which can lead to some round-off problem. It should be possible to integrate few channels together but care must be taken to never reach situation where irrespective to the time onset considered, the likelihood of the template will be 0 (from a computer perspective). Solving the previous point will provide the solution to this one.

Multicomponent pattern

As mentioned in the chapter 8, the method is based on the research of a particular pattern that found in most of the trial of the same class. If the prominent pattern is found in most of them and is significantly different from the model of the non-informative sample, it may "mask" other pattern less prominent (for example higher frequency). To work around the problem, it is possible to analyse the signal through different bandpass as mentioned earlier. But it may suffer from the same problem as using multiple channels.

Time-onsets distribution

We have mentioned in chapter 7 the need of modelling the time onset distribution. We opted for a simple approach by fitting the beta distribution per class on the estimated time-onset (after training) It is better to integrate the time-onset distribution in the Bayesian model and to pose a prior of the parameters of the beta distribution. Using such model of time-onset could be beneficial as it might favour identification of informative part that close. A possible way to tackle this issue would be to use a beta distribution model for the onset per class and to carefully choose the prior on their parameters.

Model of non-informative part

As it is now, the model used for the non-informative sample could maybe improved. The current model use a stochastic Gaussian process to model the local evolution of the amplitude of the EEG signal. However if not constrained, the stochastic Gaussian process allows negative values. This shows that it is not the best fit for this type of signal. However, for the moment no replacement can be proposed.

Enforcing discrimination

There is a great freedom on what type of informative template can be learnt. Currently the method only tries to find a template model that is different from what the distribution non-informative set. Since the non-informative samples of both class share the same distribution, the distribution of the templates of the two classes will converge differently assuming of course that the data of the two classes are different. However when the signal is very rich, there is a possibility that the template converge to inappropriate pattern. In this case, we could improve the method by adding a function of discrimination between the two template. Concretely, it would be a probability density function depending to template parameters of all class that would be unlikely when the parameters are too similar. However the precise model of such function still have to be found. An alternative could be to incorporate of a function of the performance on the training set, at the expense of increasing the computational cost of the estimation.

In conclusion, this work of this thesis has the prospect to be harnessed for developing asynchronous single trial analysis methods for EEG signal, which are highly demanded for complex BCI or for the future analysis of the brain dynamic of cognitive processes.

A From Markov chains to MCMC methods

This chapter presents the mathematical ground on which the MCMC methods are based. We will start by presenting the important results concerning the Markov chains and then, based on them, the MCMC algorithms. The proofs of the important theorems are provided since they are not always found in the text books.

A.1 Markov Chains

Definition A.1. A Markov chain is a sequence of random variables $\Theta[0 : n] = (\Theta[n])_{n \geq 0}$ such as the distribution of $\Theta[n + 1]$ conditionally to the past $(\Theta[\tau])_{\tau \leq n}$ depends only on $\Theta[n]$,

For any $\theta[0 : n + 1] = (\theta[\tau])_{0 \leq \tau \leq n+1}$

$$p_{\Theta[n+1]|\Theta[0:n]}(\theta[n + 1]|\theta[0 : n]) = p_{\Theta[n+1]|\Theta[n]}(\theta[n + 1]|\theta[n])$$

$p_{\Theta[n+1]|\Theta[n]}$ is called the density of the transition kernel.

In the following, the density of the transition kernel will be noted κ_n and the marginal density ρ_n :

For any (θ, θ') ,

$$\begin{aligned}\kappa_n(\theta|\theta') &= p_{\Theta[n+1]|\Theta[n]}(\theta|\theta') \\ \rho_n(\theta) &= p_{\Theta[n]}(\theta)\end{aligned}$$

It can be easily verified that the sequences $(\kappa_n)_{n \geq 0}$ and $(\rho_n)_{n \geq 0}$ are bound together with the following relation :

Annexe A. From Markov chains to MCMC methods

For any θ' ,

$$\rho_{n+1}(\theta') = \int \kappa_n(\theta'|\theta)\rho_n(\theta)d\theta \quad (\text{A.1})$$

Definition A.2 (homogeneity). *A Markov chain is said homogeneous if the transition kernel is invariant :*

$$\forall n \quad \kappa_n = \kappa$$

Definition A.3 (Invariant distribution). *A distribution density ρ is invariant for the Markov chain with kernel κ_n if, for any n , $\Theta[n]$ follows the law ρ results in $\Theta[n+1]$ following the law ρ .*

Proposition A.1. *If ρ is invariant, it results immediatly from the formula (A.1) that ρ and κ_n are bound by the relation :*

For any θ' and any n ,

$$\rho(\theta') = \int \kappa_n(\theta'|\theta)\rho(\theta)d\theta \quad (\text{A.2})$$

Property A.1. *If κ_n and ρ are linked by the following property of reversibility :*

For any (θ, θ') ,

$$\kappa_n(\theta'|\theta)\rho(\theta) = \kappa_n(\theta|\theta')\rho(\theta') \quad (\text{A.3})$$

then the relation (A.2) is verified.

Démonstration.

$$\begin{aligned} \int \kappa_n(\theta'|\theta)\rho(\theta)d\theta &= \int \kappa_n(\theta|\theta')\rho(\theta')d\theta \\ &= \rho(\theta') \int \kappa_n(\theta|\theta')d\theta \\ &= \rho(\theta') \end{aligned} \quad \square$$

Definition A.4. *A Markov chain with a kernel κ_n is said ergodic if the marginal law ρ_n*

converges towards a limit ρ , independently from the initial law ρ_0 :

$$\rho = \lim_{n \rightarrow +\infty} \rho_n$$

ρ is then the unique invariant distribution of this Markov chain.

Theorem A.1. *If a homogenous Markov chain with the kernel κ has an invariant distribution ρ and if $\kappa(\cdot|\theta)$ is non-zero on the support of ρ for any θ , then the chain is ergodic and the marginal law converges to ρ .*

In addition, let ν be the scalar defined by :

$$\nu = \min_{\{(\theta, \theta') | \rho(\theta') > 0\}} \frac{\kappa(\theta'|\theta)}{\rho(\theta')}$$

$\nu > 0$ if $\kappa(\cdot|\theta)$ is non-zero on the support of ρ and $\nu \leq 1$ because it is a ratio of probability density. Then for any real-valued bounded function f , there is :

$$|\varepsilon\{f(\Theta[+\infty])\} - \varepsilon\{f(\Theta[n])\}| \leq (1 - \nu)^n \max_{(\theta, \theta')} |f(\theta) - f(\theta')|$$

In particular, if f is the indicator function of a measurable set A , we deduce immediatly that :

$$|\text{Prob}(\Theta[+\infty] \in A) - \text{Prob}(\Theta[n] \in A)| \leq (1 - \nu)^n$$

Démonstration. We are going to show recursively that the law ρ_n is necessarily written as :

$$\rho_n(\theta) = [1 - (1 - \nu)^n] \rho(\theta) + (1 - \nu)^n r_n(\theta)$$

where r_n is a density of probability. The property is true for $n = 0$ by taking $r_0 = \rho_0$. Let

Annexe A. From Markov chains to MCMC methods

us assume it is true for n and let us calculate ρ_{n+1} :

$$\begin{aligned}
 \rho_{n+1}(\theta') &= \int \kappa(\theta'|\theta)\rho_n(\theta)d\theta \\
 &= [1 - (1 - \nu)^n] \underbrace{\int \kappa(\theta'|\theta)\rho(\theta)d\theta}_{\rho(\theta')} + (1 - \nu)^n \int \kappa(\theta'|\theta)r_n(\theta)d\theta \\
 &= [1 - (1 - \nu)^n]\rho(\theta') + \\
 &\quad (1 - \nu)^n \int [\kappa(\theta'|\theta) - \nu\rho(\theta') + \nu\rho(\theta')]r_n(\theta)d\theta \\
 &= [1 - (1 - \nu)^{n+1}]\rho(\theta') + \\
 &\quad (1 - \nu)^{n+1} \underbrace{\int \frac{[\kappa(\theta'|\theta) - \nu\rho(\theta')]r_n(\theta)}{1 - \nu}d\theta}_{r_{n+1}(\theta')}
 \end{aligned}$$

It is easily verified that such defined r_{n+1} is effectively a density of probability. So we have indeed $\lim_{n \rightarrow +\infty} \rho_n(\theta) = \rho$ independently from the chosen invariant law.

So there is a unique invariant law. Moreover :

$$\begin{aligned}
 & \left| \varepsilon\{f(\Theta[+\infty])\} - \varepsilon\{f(\Theta[n])\} \right| \\
 &= \left| \int f(\theta)\rho(\theta)d\theta - \int f(\theta)\rho_n(\theta)d\theta \right| \\
 &= (1 - \nu)^n \left| \int f(\theta)\rho(\theta)d\theta - \int f(\theta)r_n(\theta)d\theta \right| \\
 &\leq (1 - \nu)^n \left(\max_{\theta} f(\theta) - \min_{\theta} f(\theta) \right) \\
 &\leq (1 - \nu)^n \max_{(\theta, \theta')} |f(\theta) - f(\theta')| \quad \square
 \end{aligned}$$

A.2 MCMC

In the following, we will consider the Markov chains with the following notations :

Markov chains $\{\Theta[n]\}_{n \in \mathbb{N}}$:

state :	$\Theta[n]$
transition kernel :	$p_{\Theta[n+1] \Theta[n]}$ noted κ_n
marginal law :	$p_{\Theta[n]}$ noted ρ_n

Our goal is to find an algorithm that allows one to sample an arbitrary probability density function ρ . In order to solve this problem, we will find a Markov chain whose transition kernel will converge to the target density ρ .

Let us consider the homogenous Markov chain with an invariant transition kernel κ verifying :

$$\forall(\theta, \theta') \quad \kappa_n(\theta'|\theta)\rho(\theta) = \kappa_n(\theta|\theta')\rho(\theta')$$

In this context, we deduce from the property A.1 that ρ is the invariant density for the aforementioned Markov chain. We will refer later this condition as the *reversibility condition*.

Let us suppose we have a Markov chain with the kernel κ verifying this condition. According to the theorem A.1, the chain is ergodic and the marginal laws ρ_n converge to the target density ρ . We then deduce the following elementary algorithm :

- Draw $\theta[0]$ randomly (such as $\rho(\theta[0]) > 0$)
- For any $n > 0$, generate $\theta[n + 1]$ following $\kappa(\cdot|\theta[n])$

This algorithm assumes that we have a generator of the law $\kappa(\cdot|\theta[n])$ at our disposal. The Metropolis-Hastings algorithm allows to solve the problem starting from an arbitrary generator. In this way, one can employ usual generator (uniform law, normal law...) as well as all that derives from them.

A.2.1 The Metropolis-Hastings algorithm

We are first going to generalize the property A.1.

Property A.2. Let $f : (\theta', \theta) \mapsto f(\theta', \theta)$ be a function :

$$\left\{ \begin{array}{l} \text{non negative} \\ \text{with an integrate} \leq 1 (\theta') \end{array} \right. \quad \begin{array}{l} f(\theta', \theta) \geq 0 \\ r(\theta) = 1 - \int f(\theta', \theta)d\theta' \geq 0 \end{array} \quad (\text{A.4})$$

which verifies the following property of reversibility :

$$\forall(\theta, \theta') \quad f(\theta'|\theta)\rho(\theta) = f(\theta|\theta')\rho(\theta') \quad (\text{A.5})$$

Then the Markov chain defined by the following transition kernel (with δ referring to the Dirac function) :

$$\kappa(\theta'|\theta) = f(\theta', \theta) + r(\theta)\delta(\theta' - \theta)$$

accepts ρ as invariant density.

Annexe A. From Markov chains to MCMC methods

This is equivalent to say that if $\Theta[n] = \theta$, then $\Theta[n + 1] = \theta$ with the probability $r(\theta)$, and that $\Theta[n + 1]$ is drawn from the law $\frac{f(\cdot, \theta)}{1-r(\theta)}$ with the probability $1 - r(\theta)$.

Démonstration. Let us prove that the density ρ verifies the relation of invariance $\rho = \int \kappa(\cdot | \theta) \rho(\theta) d\theta$:

$$\begin{aligned} \int \kappa(\theta' | \theta) \rho(\theta) d\theta &= \int f(\theta', \theta) \rho(\theta) + \int r(\theta) \delta(\theta' - \theta) \rho(\theta) \\ &= \int f(\theta, \theta') \rho(\theta') d\theta + r(\theta') \rho(\theta') \\ &= (1 - r(\theta)) \rho(\theta') + r(\theta') \rho(\theta') \\ &= \rho(\theta') \end{aligned} \quad \square$$

Property A.3. Let q be the density of an instrumental law, non null on the support of ρ :

$$\forall(\theta, \theta') \quad \begin{cases} q(\theta' | \theta) \geq 0 \\ \int q(\theta' | \theta) d\theta = 1 \\ \rho(\theta') \neq 0 \rightarrow q(\theta' | \theta) \neq 0 \end{cases}$$

Then the function f defined by :

$$f(\theta', \theta) = q(\theta' | \theta) \min \left(\frac{q(\theta | \theta') \rho(\theta')}{q(\theta' | \theta) \rho(\theta)}, 1 \right)$$

verifies by construction the formula (A.4), and verifies the property of reversibility (A.5).

Démonstration.

$$\text{Let us pose } \alpha(\theta', \theta) = \min \left(\frac{q(\theta | \theta') \rho(\theta')}{q(\theta' | \theta) \rho(\theta)}, 1 \right).$$

- If $q(\theta' | \theta) = 0$, then $\rho(\theta') = 0$ and $f(\theta', \theta) = 0$, the property of reversibility is properly verified.
- If $\frac{q(\theta | \theta') \rho(\theta')}{q(\theta' | \theta) \rho(\theta)} \geq 1$, then $\alpha(\theta', \theta) = 1$ and $\alpha(\theta, \theta') = \frac{q(\theta' | \theta) \rho(\theta)}{q(\theta' | \theta') \rho(\theta')}$. So we have :

$$\begin{aligned} f(\theta', \theta) \rho(\theta) &= q(\theta' | \theta) \alpha(\theta', \theta) \rho(\theta) \\ &= q(\theta' | \theta) \rho(\theta) \\ &= q(\theta | \theta') \rho(\theta') \frac{q(\theta' | \theta) \rho(\theta)}{q(\theta' | \theta') \rho(\theta')} \\ &= q(\theta | \theta') \rho(\theta') \alpha(\theta, \theta') \\ &= f(\theta, \theta') \rho(\theta') \end{aligned}$$

– If $\frac{q(\theta|\theta')\rho(\theta')}{q(\theta'|\theta)\rho(\theta)} \leq 1$, then $\alpha(\theta', \theta) = \frac{q(\theta|\theta')\rho(\theta')}{q(\theta'|\theta)\rho(\theta)}$ and $\alpha(\theta, \theta') = 1$. So we have :

$$\begin{aligned} f(\theta', \theta)\rho(\theta) &= q(\theta'|\theta)\alpha(\theta', \theta)\rho(\theta) \\ &= q(\theta'|\theta)\frac{q(\theta|\theta')\rho(\theta')}{q(\theta'|\theta)\rho(\theta)}\rho(\theta) \\ &= q(\theta|\theta')\rho(\theta') \\ &= q(\theta|\theta')\alpha(\theta, \theta')\rho(\theta') \\ &= f(\theta, \theta')\rho(\theta') \end{aligned}$$

□

By this property, we realise that given a target law ρ and the instrumental law q , the Markov chain with the transition kernel κ defined by :

$$\kappa(\theta'|\theta) = \alpha(\theta', \theta)q(\theta'|\theta) + r(\theta)\delta(\theta' - \theta)$$

$$\text{where } \begin{cases} \alpha(\theta', \theta) = \min\left(\frac{q(\theta|\theta')\rho(\theta')}{q(\theta'|\theta)\rho(\theta)}, 1\right) \\ r(\theta) = 1 - \int q(\theta'|\theta)\alpha(\theta', \theta)d\theta' \end{cases}$$

accepts ρ as invariant density.

We have then all the elements for the Metropolis-Hastings algorithm :

Metropolis-Hastings

Given

- target density $\rho : \theta \mapsto \rho(\theta)$;
- instrumental density $q : (\theta', \theta) \mapsto q(\theta'|\theta)$;
- initialisation $\theta[0]$.

Recurrence (over $n \geq 0$)

- draw $\theta^*[n+1]$ following $q(\cdot|\theta[n])$;
- draw $u[n]$ following a uniform law between 0 and 1 ;
- calculate $\alpha[n] = \min\left(\frac{q(\theta[n]|\theta^*[n+1])\rho(\theta^*[n+1])}{q(\theta^*[n+1]|\theta[n])\rho(\theta[n])}, 1\right)$;
- $\theta[n+1] = \begin{cases} \theta^*[n+1] & \text{if } \alpha[n] > u[n] \\ \theta[n] & \text{otherwise} \end{cases}$

This indeed corresponds to the previously defined Markov chain :

Démonstration. Let us calculate first the conditional density $p_{\Theta[n+1]|\Theta^*[n+1], \Theta[n]}$:

$$\begin{aligned} p_{\Theta[n+1]|\Theta^*[n+1], \Theta[n]}(\theta'|\theta, \theta^*) &= \Pr(U[n] < \alpha(\theta^*, \theta))\delta(\theta' - \theta^*) + \\ &\quad \Pr(U[n] < \alpha(\theta^*, \theta))\delta(\theta' - \theta) \\ &= \alpha(\theta^*, \theta)\delta(\theta' - \theta^*) + (1 - \alpha(\theta^*, \theta))\delta(\theta' - \theta) \end{aligned}$$

We deduce :

$$\begin{aligned}
 p_{\Theta[n+1]|\Theta[n]}(\theta'|\theta) &= \int p_{\Theta[n+1]|\Theta^*[n+1],\Theta[n]}(\theta'|\theta, \theta^*) \times \underbrace{p_{\Theta^*[n+1]|\Theta[n]}(\theta^*|\theta)}_{q(\theta^*|\theta)} d\theta^* \\
 &= \int \alpha(\theta^*, \theta) \delta(\theta' - \theta^*) q(\theta^*|\theta) d\theta^* + \\
 &\quad \int (1 - \alpha(\theta^*, \theta)) \delta(\theta' - \theta) q(\theta^*|\theta) d\theta^* \\
 &= \alpha(\theta', \theta) q(\theta'|\theta) + \\
 &\quad \left(1 - \int \alpha(\theta^*, \theta) q(\theta^*|\theta) d\theta^* \right) \delta(\theta' - \theta) \\
 &= \alpha(\theta', \theta) q(\theta'|\theta) + r(\theta) \delta(\theta' - \theta) \quad \square
 \end{aligned}$$

A.2.2 Variants and Gibbs sampler

If the Markov chain $(\Theta[n])$ can be decomposed in the following way :

$$\Theta[n] = \begin{bmatrix} \Theta_1[n] \\ \Theta_2[n] \\ \vdots \\ \Theta_d[n] \end{bmatrix}$$

so that it is easier to draw randomly $\Theta_i[n]$ for any $i \in \{1, \dots, d\}$, we can then use two variants of the Metropolis-Hastings algorithm :

- the Metropolis-Hastings algorithm with random update
- the Metropolis-Hastings algorithm with shifting update

Random choice of the variable to be updated

With the previous decomposition, the transition kernel κ can be written :

$$\kappa = \sum_{i=1}^d \omega_i \kappa^{(i)}$$

with $\sum_{i=1}^d \omega_i = 1$ (we can take $\omega_i = \frac{1}{d}$), and where all kernel $\kappa^{(i)}$ verify the condition of invariance. The kernel $\kappa^{(i)}$ then corresponds to a random update of the i -th variable

only :

$$\begin{aligned}\kappa^{(i)}(\theta'|\theta) &= \tilde{\kappa}^{(i)}(\theta'_i|\theta) \prod_{\substack{k=1 \\ k \neq i}}^d \delta(\theta'_k - \theta_k) \\ &= \tilde{\kappa}^{(i)}(\theta'_i|\theta) \delta(\theta'_{1:i-1, i+1:d} - \theta_{1:i-1, i+1:d})\end{aligned}\tag{A.6}$$

The condition of invariance is then written after integration over the $d - 1$ variables unchanged by the transition :

$$\rho(\theta_{1:i-1}, \theta'_i, \theta_{i+1:d}) = \int \tilde{\kappa}^{(i)}(\theta'_i|\theta_{1:i-1}, \theta_i, \theta_{i+1:d}) \rho(\theta_{1:i-1}, \theta_i, \theta_{i+1:d})\tag{A.7}$$

This condition of invariance is verified if each kernel verifies the condition of reversibility :

$$\tilde{\kappa}^{(i)}(\theta'_i|\theta_{1:i-1}, \theta_i, \theta_{i+1:d}) \rho(\theta_{1:i-1}, \theta_i, \theta_{i+1:d}) = \tilde{\kappa}^{(i)}(\theta_i|\theta_{1:i-1}, \theta'_i, \theta_{i+1:d}) \rho(\theta_{1:i-1}, \theta'_i, \theta_{i+1:d})$$

Given an instrumental law $\tilde{q}^{(i)}$, the kernel $\tilde{\kappa}^{(i)}$ defined for any (θ, θ') such as $\theta_k = \theta'_k$ if $k \neq i$ by :

$$\tilde{\kappa}^{(i)}(\theta'_i|\theta) = \tilde{q}^{(i)}(\theta'_i|\theta) \underbrace{\min\left(\frac{\tilde{q}^{(i)}(\theta_i|\theta') \rho(\theta')}{\tilde{q}^{(i)}(\theta'_i|\theta) \rho(\theta)}, 1\right)}_{\alpha_i(\theta'_i, \theta)} + r_i(\theta) \delta(\theta'_i - \theta_i)$$

$$\text{with } r_i(\theta) = 1 - \int \tilde{q}^{(i)}(\theta'_i|\theta) \alpha_i(\theta'_i, \theta) d\theta'_i$$

verifies the condition of reversibility.

We have then all the element for the following Metropolis-Hastings algorithm, initialised by $\theta[0]$:

Metropolis-Hastings (random update)

Given

- target density $\rho : \theta \mapsto \rho(\theta)$;
- instruments $\tilde{q}^{(i)} : (\theta'_i, \theta) \mapsto \tilde{q}^{(i)}(\theta'_i|\theta)$ for $i \in \{1, \dots, d\}$;
- discrete law $(\omega_i)_{1 \leq i \leq d}$;
- initialisation $\theta[0]$.

Recurrence (over $n \geq 0$)

- draw $i[n]$ following the discrete law $(\omega_i)_{1 \leq i \leq d}$;
- draw $u[n]$ following a uniform law between 0 et 1 ;
- draw $\theta_{i[n]}^*[n+1]$ following $\tilde{q}^{(i[n])}(\cdot|\theta[n])$;
- pose $\theta^*[n+1] = (\theta_{1:i[n]-1}[n], \theta_{i[n]}^*[n+1], \theta_{i[n]+1:d}[n])$
- calculate $\alpha[n] = \min \left(\frac{\tilde{q}^{(i[n])}(\theta_{i[n]}[n]|\theta^*[n+1])\rho(\theta^*[n+1])}{\tilde{q}^{(i[n])}(\theta^*[n+1]|\theta[n])\rho(\theta[n])}, 1 \right)$;
- $\theta[n+1] = \begin{cases} \theta^*[n+1] & \text{if } \alpha[n] > u[n] \\ \theta[n] & \text{otherwise} \end{cases}$

Shift of the variable to be updated

In practice, the variable to be updated is not chosen randomly. We propose to draw sequentially the d kernels $\kappa^{(i)}$, for i varying from 1 to d . Let $\theta = \theta^{(0)}$ be the initial point, $\theta' = \theta^{(d)}$ the final point :

$$\begin{aligned} \kappa(\theta'|\theta) &= \kappa(\theta^{(d)}|\theta^{(0)}) \\ &= \int \kappa^{(d)}(\theta^{(d)}|\theta^{(d-1)}) \dots \kappa^{(1)}(\theta^{(1)}|\theta^{(0)}) d\theta^{(1)} \dots d\theta^{(d-1)} \end{aligned}$$

So if the kernel $\kappa^{(i)}$ corresponds to the update of the i -th variable only (formula A.6), by integrating, we get :

$$\kappa(\theta'|\theta) = \prod_{i=1}^d \tilde{\kappa}^{(i)}(\theta'_i|\theta'_{1:i-1}, \theta_{i:d})$$

It is easily verified that if all $\tilde{\kappa}^{(i)}$ verify the condition of invariance (formula A.7), the kernel κ verifies the condition of invariance (A.2).

The Metropolis-Hastings algorithm initialised by $\theta[n] = \theta_{1:d}[n]$ can then be written :

Metropolis-Hastings (shifting update)

Given

- Target density $\rho : \theta \mapsto \rho(\theta)$;
- instruments $\tilde{q}^{(i)} : (\theta'_i, \theta) \mapsto \tilde{q}^{(i)}(\theta'_i|\theta)$ for $i \in \{1, \dots, d\}$;
- initialisation $\theta[0]$.

Recurrence (over $n \geq 0$)

- For i varying from 1 to d :
 - Pose $\theta^{(i)}[n] = (\theta_{1:i-1}[n+1], \theta_{i:d}[n])$;
 - draw $u_i[n]$ following a uniform law on $[0, 1]$;
 - draw $\theta_i^*[n+1]$ following $\tilde{q}^{(i)}(\cdot|\theta^{(i)}[n])$;
 - pose $\theta^{(i)*}[n] = (\theta_{1:i-1}[n], \theta_i^*[n+1], \theta_{i+1:d}[n])$
 - calculate $\alpha_i[n] = \min \left(\frac{\tilde{q}^{(i)}(\theta_i[n]|\theta^{(i)*}[n+1])\rho(\theta^{(i)*}[n+1])}{\tilde{q}^{(i)}(\theta_i^*[n+1]|\theta^{(i)}[n])\rho(\theta^{(i)}[n])}, 1 \right)$;
 - $\theta_i[n+1] = \begin{cases} \theta_i^*[n+1] & \text{if } \alpha_i[n] > u_i[n] \\ \theta_i[n] & \text{otherwise} \end{cases}$

Gibbs sampler

We use one of the two previous algorithm, with taking $\tilde{q}^{(i)}$ from the target law ρ as the $\Theta_i[n]$ conditionally to the remaining variables $\Theta_{1:i-1, i+1:d}[n]$.

For any (θ, θ') such as $\theta_k = \theta'_k$ if $k \neq i$:

$$\begin{aligned} \tilde{q}^{(i)}(\theta'_i|\theta) &= \tilde{q}^{(i)}(\theta'_i|\theta_{1:d}) \\ &= \frac{\rho(\theta_{1:i-1}, \theta'_i, \theta_{i+1:d})}{c_i(\theta_{1:i-1}, \theta_{i+1:d})} \\ &= \frac{\rho(\theta')}{c_i(\theta_{1:i-1}, \theta_{i+1:d})} \end{aligned}$$

where c_i is the marginal law $\Theta_{1:i-1, i+1:d}[n]$:

$$c_i(\theta_{1:i-1}, \theta_{i+1:d}) = \int \rho(\theta_{1:i-1}, y_i, \theta_{i+1:d}) dy_i$$

Then let us calculate $\alpha_i(\theta'_i, \theta)$:

$$\begin{aligned} \alpha_i(\theta'_i, \theta) &= \min \left(\frac{\tilde{q}^{(i)}(\theta_i|\theta')\rho(\theta')}{\tilde{q}^{(i)}(\theta'_i|\theta)\rho(\theta)}, 1 \right) \\ &= \min \left(\frac{\frac{\rho(\theta)}{c_i(\theta_{1:i-1}, \theta_{i+1:d})}\rho(\theta')}{\frac{\rho(\theta')}{c_i(\theta_{1:i-1}, \theta_{i+1:d})}\rho(\theta)}, 1 \right) \\ &= 1 \end{aligned}$$

Annexe A. From Markov chains to MCMC methods

The algorithm becomes simply :

Gibbs (shifting update)

Given

- target density $\rho : \theta \mapsto \rho(\theta)$;
- initialisation $\theta_{2:d}[0]$.

Recurrence (over $n \geq 0$)

- for i varying from 1 to d ;
 - draw $\theta_i[n + 1]$ following the law :
 $q^{(i)}(\cdot | \theta_{1:i-1}[n + 1], \theta_{i+1:d}[n]) \propto \rho(\theta_{1:i-1}[n + 1], \cdot, \theta_{i+1:d}[n])$

B Method for sampling a truncated normal distribution

In this appendix, we will see how to sample from an univariate truncated normal distribution $\mathcal{TN}_{[\mu^-, +\infty]}(\mu, \sigma)$. We will restrict ourselves to the case of one-sided left truncation. The right truncation can be trivially deduced from the left truncation and the two-sided truncation is not used in the frame of this thesis. The interested reader can refer to [106] for an explanation on the two-sided case as well as the multivariate case. Moreover, for the explanations in this chapter, we will consider the case of $\mu = 0$ and $\sigma = 1$ because :

$$\left. \begin{array}{l} x \sim \mathcal{TN}_{[\frac{\mu^- - \mu}{\sigma}, +\infty]}(0, 1) \\ X = \sigma x + \mu \end{array} \right\} \Rightarrow X \sim \mathcal{TN}_{[\mu^-, +\infty]}(\mu, \sigma)$$

Sampling a truncated normal distribution is more complicated than it seems at first glance. Although there is an obvious algorithm :

repetition algorithm

```
- loop
  - Generate  $z \sim \mathcal{N}(0, 1)$ 
  until  $z > \mu^-$ 
```

its rejection rate makes it useless for practical use when μ^- is becoming much bigger than μ . It is unacceptable when used in a MCMC algorithm which samples repetitively a big number of times the same distribution model (varying the parameters at each iteration).

Another approach could be to use the classical cumulative density function inversion technique :

Annexe B. Method for sampling a truncated normal distribution

inversion algorithm

- Generate $u \sim \mathcal{U}_{[0,1]}$
- Compute $z = \Phi^{-1}(\Phi(\mu^-) + u(1 - \Phi(\mu^-)))$

But due to the precision of the usual estimation of Φ and Φ^{-1} , the algorithm is not very reliable when μ^- is large.

Robert [106] proposed an *accept-reject* algorithm which is more efficient than the other two for $\mu^- > 0$:

accept-reject algorithm

- Compute $\alpha = \frac{\mu^- + \sqrt{(\mu^-)^2 + 4}}{2}$
- loop
 - Generate $z \sim \mathcal{Exp}(\alpha, \mu^-)$
 - Compute $\varrho(z) = \exp(-\frac{(z-\alpha)^2}{2})$
 - Generate $u \sim \mathcal{U}_{[0,1]}$
- until $u < \varrho(z)$

We can then conclude that an efficient sampling of a truncated normal distribution

$\mathcal{TN}_{[\mu^-, +\infty[}(0, 1)$ depends on the value of μ^- :

- if $\mu^- > 0$ use the accept-reject algorithm
- if $\mu^- \in [\alpha, 0]$ use the inversion algorithm
- if $\mu^- < \alpha$ use the repetition algorithm

The choice of α depends roughly on the computational cost of generating a sample from a Gaussian distribution and the probability of rejection compared to the cost of generating a sample from uniform distribution plus the cost of computing $\Phi^{-1}(\Phi(\mu^-) + u(1 - \Phi(\mu^-)))$. This depends of course on the different implementations operation (and on the platform). However with the implementations of MATLAB, this threshold has been estimated empirically on the platform used for the computations in these studies to the value $\alpha = -2.5$.

C notations

C.1 Mathematics

$\mathbf{x} = (x_1, \dots, x_n)$	boldface signifies a vector
$\mathbb{I}_A(t)$	indicator function (1 if $t \in A$, 0 otherwise)
$f(t) \propto h(t)$	functions f and h are proportional
$\Gamma(x)$	gamma function ($x > 0$)
\bar{x}	empirical mean of the sequence $(x_i)_{i \in A}$

C.2 Probability

$f(x \theta)$	density of X , conditional on parameter θ
$X \sim f(x \theta)$	X is distributed with density $f(x \theta)$
$\varphi(x, \mu, \sigma^2)$	probability density function of the normal distribution $\mathcal{N}(\mu, \sigma^2)$
$\varphi(x)$	probability density function of the normal distribution $\mathcal{N}(0, 1)$
$\Phi(x, \mu, \sigma^2)$	cumulative density function of the normal distribution $\mathcal{N}(\mu, \sigma^2)$
$\Phi(x)$	cumulative density function of the normal distribution $\mathcal{N}(0, 1)$
$\pi(\theta)$	generic prior for θ
$\pi(\theta x)$	generic posterior for θ

C.3 Distributions

\mathcal{U}_A	Uniform distribution on the support A
$\mathcal{N}(\mu, \sigma^2)$	normal distribution with mean μ and variance σ^2
$\mathcal{TN}_A(\mu, \sigma^2)$	truncated normal distribution on support A
$\mathcal{SN}(\xi, \omega^2, \alpha)$	skew-normal distribution with location ξ , scale ω and skew parameter α
$\mathcal{Exp}(\alpha, \mu)$	exponential distribution
$\mathcal{Be}(\alpha, \beta)$	beta distribution
$\mathcal{IG}(\alpha, \beta)$	inverse gamma distribution

Bibliographie

- [1] G. Pfurtscheller, C. Neuper, C. Guger, W. Harkam, H. Ramoser, A. Schlögl, B. Obermaier, and M. Pregenzer. Current trends in Graz brain-computer interface (BCI) research. *IEEE Transactions on Rehabilitation Engineering*, 8 :216–219, 2000.
- [2] J. d. R. Millan, M. Franze, J. Mourino, F. Cincotti, and F. Babiloni. Relevant EEG features for the classification of spontaneous motor-related tasks. *Biol Cybern*, 86 :89–95, 2002.
- [3] C. Tallon-Baudry and O. Bertrand. Oscillatory gamma activity in humans and its role in object representation. *Trends Cogn Sci*, 3(4) :151–162, 1999.
- [4] F. Vialatte, C. Martin, R. Dubois, J. Haddad, B. Quenet, R. Gervais, and G. Dreyfus. A machine learning approach to the analysis of time-frequency maps, and its application to neural dynamics. *Neural Netw*, 20 :194–209, 2007.
- [5] W. Yoshida and S. Ishii. Resolution of uncertainty in prefrontal cortex. *Neuron*, 50(5) :781–789, Jun 2006.
- [6] M. F. S. Rushworth. Intention, Choice, and the Medial Frontal Cortex. *Ann N Y Acad Sci*, 1124 :181–207, Mar 2008.
- [7] N. D. Daw, J. P. O’Doherty, P. Dayan, B. Seymour, and R. J. Dolan. Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095) :876–9, Jun 15 2006.
- [8] J. P. Donoghue. Bridging the brain to the world : a perspective on neural interface systems. *Neuron*, 60(3) :511–521, Nov 2008.
- [9] D. R. Kipke, W. Shain, G. Buzsáki, E. Fetz, J. M. Henderson, J. F. Hetke, and G. Schalk. Advanced neurotechnologies for chronic neural interfaces : new horizons and clinical opportunities. *J Neurosci*, 28(46) :11830–11838, Nov 2008.
- [10] M. A. Lebedev and M. A. L. Nicolelis. Brain-machine interfaces : Past, present and future. *Trends Neurosci*, 29(9) :536–546, Sep 2006.

Bibliographie

- [11] J. d. R. Millan and J. Carmena. Invasive or Noninvasive : Understanding Brain-Machine Interface Technology [Conversations in BME]. *Engineering in Medicine and Biology Magazine, IEEE*, 29 :16–22, 2010.
- [12] J. J. Daly and J. R. Wolpaw. Brain-computer interfaces in neurological rehabilitation. *The Lancet Neurology*, 7 :1032–1043, 2008.
- [13] A. Kübler and N. Birbaumer. Brain-computer interfaces and communication in paralysis : Extinction of goal directed thinking in completely paralysed patients? *Clin Neurophysiol*, 119(11) :2658–2666, Nov 2008.
- [14] J. d. R. Millan, R. Rupp, G. Müller-Putz, R. Murray-Smith, C. Giugliemma, M. Tangermann, C. Vidaurre, F. Cincotti, A. Kübler, R. Leeb, C. Neuper, K. Müller, and D. Mattia. Combining Brain-Computer Interfaces and Assistive Technologies : State-of-the-Art and Challenges. *Frontiers in Neuroscience*, 4 :161, 2010.
- [15] J. Wolpaw. Brain-computer interfaces as new brain output pathways. *J Physiol*, 579(Pt 3) :613–619, Mar 2007.
- [16] J. Wolpaw, N. Birbaumer, D. McFarland, G. Pfurtscheller, and T. Vaughan. Brain-computer interfaces for communication and control. *Clin Neurophysiol*, 113(6) :767–791, Jun 2002.
- [17] F. Karmali, M. Polak, and A. Kostov. Environmental control by a brain-computer interface. In *Engineering in Medicine and Biology Society, 2000. Proceedings of the 22nd Annual International Conference of the IEEE*, volume 4, pages 2990–2992 vol.4. 2000.
- [18] X. Gao, D. Xu, M. Cheng, and S. Gao. A BCI-based environmental controller for the motion-disabled. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 11 :137–140, 2003.
- [19] L. A. Farwell and E. Donchin. Talking off the top of your head : toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalography and Clinical Neurophysiology*, 70 :510–523, 1988.
- [20] J. M. Carmena, M. A. Lebedev, R. E. Crist, J. E. O’Doherty, D. M. Santucci, D. F. Dimitrov, P. G. Patil, C. S. Henriquez, and M. A. L. Nicolelis. Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol*, 1(2) :E42, Nov 2003.
- [21] J. d. R. Millan, F. Renkens, J. Mouriño, and W. Gerstner. Brain-actuated interaction. *Artif. Intell.*, 159 :241–259, 2004.
- [22] T. M. Vaughan, W. J. Heetderks, L. J. Trejo, W. Z. Rymer, M. Weinrich, M. M. Moore, A. Kübler, B. H. Dobkin, N. Birbaumer, E. Donchin, E. W. Wolpaw, and J. R. Wolpaw. Brain-computer interface technology : a review of the Second

- International Meeting. *IEEE Trans Neural Syst Rehabil Eng*, 11(2) :94–109, Jun 2003.
- [23] J. R. Wolpaw, N. Birbaumer, W. J. Heetderks, D. J. McFarland, P. H. Peckham, G. Schalk, E. Donchin, L. A. Quatrano, C. J. Robinson, and T. M. Vaughan. Brain-computer interface technology : a review of the first international meeting. *IEEE Transactions on Rehabilitation Engineering*, 8 :164–173, 2000.
- [24] E. E. Fetz. Operant conditioning of cortical unit activity. *Science*, 163 :955–8, 1969.
- [25] A. Sharma, L. Rieth, P. Tathireddy, R. Harrison, H. Oppermann, M. Klein, M. Töpper, E. Jung, R. Normann, G. Clark, and F. Solzbacher. Long term in vitro functional stability and recording longevity of fully integrated wireless neural interfaces based on the Utah slant electrode array. *Journal of Neural Engineering*, 8 :045004, 2011.
- [26] M. Bares and I. Rektor. Basal ganglia involvement in sensory and cognitive processing. A depth electrode CNV study in human subjects. *Clin Neurophysiol*, 112(11) :2022–30, November 2001.
- [27] E. A. Felton, J. A. Wilson, J. C. Williams, and P. C. Garell. Electrocorticographically controlled brain-computer interfaces using motor and sensory imagery in patients with temporary subdural electrode implants. Report of four cases. *Journal of Neurosurgery*, 106 :495–500, 2007.
- [28] U. Mitzdorf. Current source-density method and application in cat cerebral cortex : investigation of evoked potentials and EEG phenomena. *Physiol Rev*, 65(1) :37–100, Jan 1985.
- [29] P. R. Kennedy and R. A. Bakay. Restoration of neural output from a paralysed patient by a direct brain connection. *Neuroreport*, 9 :1707–1711, 1998.
- [30] P. G. Patil, J. M. Carmena, M. A. L. Nicolelis, and D. A. Turner. Ensemble recordings of human subcortical neurons as a source of motor control signals for a brain-machine interface. *Neurosurgery*, 55(1) :27–35 ; discussion 35–8, Jul 2004.
- [31] S. Acharya, M. S. Fifer, H. L. Benz, N. E. Crone, and N. V. Thakor. Electrocorticographic amplitude predicts finger positions during slow grasping motions of the hand. *Journal of Neural Engineering*, 7, 2010.
- [32] E. C. Leuthardt, G. Schalk, J. R. Wolpaw, J. G. Ojemann, and D. W. Moran. A brain-computer interface using electrocorticographic signals in humans. *Journal of Neural Engineering*, 1 :63–71, 2004.
- [33] S. Ray, C. C. Jouny, N. E. Crone, D. Boatman, N. V. Thakor, and P. J. Franaszczuk. Human ECoG analysis during speech perception using matching pursuit : A

Bibliographie

- comparison between stochastic and dyadic dictionaries. *IEEE Transactions on Biomedical Engineering*, 50 :1371–3, 2003.
- [34] K. J. Miller, M. denNijs, P. Shenoy, J. W. Miller, R. P. N. Rao, and J. G. Ojemann. Real-time functional brain mapping using electrocorticography. *Neuroimage*, 37 :504–7, 2007.
- [35] A. Kaplan, A. Fingelkurts, A. Fingelkurts, S. Borisov, and B. Darkhovsky. Nonstationary nature of the brain activity as revealed by EEG/MEG : Methodological, practical and conceptual challenges. *Signal Processing*, 85(11) :2190–2212, 2005.
- [36] A. Bashashati, M. Fatourehchi, R. K. Ward, and G. E. Birch. A survey of signal processing algorithms in brain-computer interfaces based on electrical brain signals. *J Neural Eng*, 4(2) :R32–R57, Jun 2007.
- [37] F. Lotte, M. Congedo, A. Lécuyer, F. Lamarche, and B. Arnaldi. A review of classification algorithms for EEG-based brain-computer interfaces. *J Neural Eng*, 4(2) :R1–R13, Jun 2007.
- [38] S. S. Dalal, J. M. Zumer, A. G. Guggisberg, M. Trumpis, D. D. E. Wong, K. Sekihara, and S. S. Nagarajan. MEG/EEG source reconstruction, statistical evaluation, and visualization with NUTMEG. *Computational Intelligence and Neuroscience*, 2011 :758973, 2011.
- [39] J.-H. Lee, J. Ryu, F. A. Jolesz, Z.-H. Cho, and S.-S. Yoo. Brain-machine interface via real-time fMRI : preliminary study on thought-controlled robotic arm. *Neuroscience Letters*, 450(1) :1–6, Jan 2008.
- [40] R. Sitaram, H. Zhang, C. Guan, M. Thulasidas, Y. Hoshi, A. Ishikawa, K. Shimizu, and N. Birbaumer. Temporal classification of multichannel near-infrared spectroscopy signals of motor imagery for developing a brain-computer interface. *Neuroimage*, 34(4) :1416–1427, Feb 2006.
- [41] G. Strangman, J. P. Culver, J. H. Thompson, and D. A. Boas. A quantitative comparison of simultaneous BOLD fMRI and NIRS recordings during functional brain activation. *NeuroImage*, 17 :719–731, 2002.
- [42] H. Berger. Über das Elektrenkephalogramm des Menschen. *Archiv für Psychiatrie und Nervenkrankheiten*, 87 :527–570, 1929.
- [43] C. Hruska, K. Alter, K. Steinhauer, and A. Steube. Misleading dialogues : Human's brain reaction to prosodic information. In ORAGE-conference. 2001.
- [44] E. Donchin, K. M. Spencer, and R. Wijesinghe. The mental prosthesis : assessing the speed of a P300-based brain-computer interface. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 8 :174–179, 2000.

- [45] D. J. Krusienski, E. W. Sellers, F. Cabestaing, S. Bayouth, D. J. McFarland, T. M. Vaughan, and J. R. Wolpaw. A comparison of classification techniques for the P300 Speller. *J Neural Eng*, 3(4) :299–305, Dec 2006.
- [46] J. D. Bayliss. Use of the evoked potential P3 component for control in a virtual apartment. *IEEE Trans Neural Syst Rehabil Eng*, 11(2) :113–116, June 2003.
- [47] I. Iturrate, J. Antelis, A. Kübler, and J. Minguez. A Noninvasive Brain-Actuated Wheelchair Based on a P300 Neurophysiological Protocol and Automated Navigation. *IEEE Transactions on Robotics*, 25(3) :614–627, June 2009.
- [48] B. Z. Allison, D. J. McFarland, G. Schalk, S. D. Zheng, M. M. Jackson, and J. R. Wolpaw. Towards an independent brain-computer interface using steady state visual evoked potentials. *Clin Neurophysiol*, 119(2) :399–408, Feb 2008.
- [49] T. Liu, L. Goldberg, S. Gao, and B. Hong. An online brain-computer interface using non-flashing visual evoked potentials. *J Neural Eng*, 7(3) :036003, Jun 2010.
- [50] G. R. Müller-Putz, R. Scherer, C. Brauneis, and G. Pfurtscheller. Steady-state visual evoked potential (SSVEP)-based communication : impact of harmonic frequency components. *Journal of Neural Engineering*, 2 :1–8, 2005.
- [51] J. Hohnsbein, M. Falkenstein, J. Hoormann, and L. Blanke. Effects of crossmodal divided attention on late ERP components. I. Simple and choice reaction tasks. *Electroencephalogr Clin Neurophysiol*, 78(6) :438–446, Jun 1991.
- [52] W. H. R. Miltner, C. H. Braun, and M. G. H. Coles. Event-Related Brain Potentials Following Incorrect Feedback in a Time-Estimation Task : Evidence for a “Generic” Neural System for Error Detection. *Journal of Cognitive Neuroscience*, 9(6) :788–798, 1997.
- [53] U. Toepel, A. Pannekamp, and K. Alter. Catching the news : Processing strategies in listening to dialogs as measured by ERPs. *Behav Brain Funct*, 3 :53, 2007.
- [54] F. J. Varela, J.-P. Lachaux, E. Rodriguez, and J. Martinerie. The brainweb : phase synchronization and large-scale integration. *Nature Reviews Neuroscience*, 2 :229–239, 2001.
- [55] J.-P. Lachaux, N. George, C. Tallon-Baudry, J. Martinerie, L. Hugueville, L. Minotti, P. Kahane, and B. Renault. The many faces of the gamma band response to complex visual stimuli. *Neuroimage*, 25(2) :491–501, Apr 2005.
- [56] L. Melloni, C. Molina, M. Pena, D. Torres, W. Singer, and E. Rodriguez. Synchronization of neural activity across cortical areas correlates with conscious perception. *The Journal of Neuroscience*, 27 :2858–2865, 2007.
- [57] J.-P. Lachaux, E. Rodriguez, J. Martinerie, and F. J. Varela. Measuring phase synchrony in brain signals. *Human Brain Mapping*, 8 :194–208, 1999.

Bibliographie

- [58] M. L. V. Quyen, J. Foucher, J.-P. Lachaux, E. Rodriguez, A. Lutz, J. Martinerie, and F. J. Varela. Comparison of Hilbert transform and wavelet methods for the analysis of neuronal synchrony. *Journal of Neuroscience Methods*, 111 :83–98, 2001.
- [59] D. Rudrauf, A. Douiri, C. Kovach, J.-P. Lachaux, D. Cosmelli, M. Chavez, C. Adam, B. Renault, J. Martinerie, and M. Le Van Quyen. Frequency flows and the time-frequency dynamics of multivariate phase synchronization in brain signals. *NeuroImage*, 31 :209–227, 2006.
- [60] R. Sutton and A. G. Barto. Reinforcement Learning - An Introduction. MIT Press, Cambridge, Ma, USA, 1998.
- [61] S. Ishii, W. Yoshida, and J. Yoshimoto. Control of exploitation-exploration meta-parameter in reinforcement learning. *Neural Netw*, 15(4-6) :665–687, Jun-Jul 2002.
- [62] S. Lemm, B. Blankertz, T. Dickhaus, and K.-R. Müller. Introduction to machine learning for brain imaging. *Neuroimage*, 56(2) :387–399, May 2011.
- [63] G. Dornhege, J. Millán, T. Hinterberger, D. McFarland, and K. Müller. Towards Brain-Computing Interfacing. Cambridge, MA : MIT Press, 2007.
- [64] G. Gangadhar, R. Chavarriaga, and J. d. R. Millán. Recognition of anticipatory behavior from human EEG. In Proc 4th Intl. Brain-Computer Interface Workshop and Training Course. Graz, Austria, 2008.
- [65] B. Blankertz, S. Lemm, M. Treder, S. Haufe, and K.R.Müller. Single-trial analysis and classification of ERP components - A tutorial. *Neuroimage*, 56(2) :814—825, 2011.
- [66] K. P. Bennett and C. Campbell. Support vector machines : hype or hallelujah ? *SIGKDD Explor. Newsl.*, 2(2) :1–13, December 2000.
- [67] M. Thulasidas, C. Guan, and J. Wu. Robust classification of EEG signal for brain-computer interface. *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, 14(1) :24–29, March 2006.
- [68] M. Zhong, F. Lotte, M. Girolami, and A. Lécuyer. Classifying EEG for brain computer interfaces using Gaussian processes. *Pattern Recognition Letters*, in press :354–359, 2007.
- [69] A. Jordan. On discriminative vs. generative classifiers : A comparison of logistic regression and naive bayes. *Advances in neural information processing systems*, 14 :841, 2002.

- [70] R. Vocat, G. Pourtois, and P. Vuilleumier. Parametric modulation of error-related ERP components by the magnitude of visuo-motor mismatch. *Neuropsychologia*, 49(3) :360–367, Feb 2011.
- [71] T. W. Picton, S. Bentin, P. Berg, E. Donchin, S. A. Hillyard, R. Johnson, G. A. Miller, W. Ritter, D. S. Ruchkin, M. D. Rugg, and M. J. Taylor. Guidelines for using human event-related potentials to study cognition : Recording standards and publication criteria. *Psychophysiology*, 37(2) :127–52, March 2000.
- [72] R. Leeb, D. Friedman, G. R. Müller-Putz, R. Scherer, M. Slater, and G. Pfurtscheller. Self-paced (asynchronous) BCI control of a wheelchair in virtual environments : a case study with a tetraplegics. *Computational Intelligence and Neuroscience*, 2007(Article ID 79642) :1–8, 2007.
- [73] J. d. R. Millan, F. Renkens, J. Mouriño, and W. Gerstner. Noninvasive brain-actuated control of a mobile robot by human EEG. *IEEE Trans Biomed Eng*, 51(6) :1026–1033, 2004.
- [74] S. Haufe, M. S. Treder, M. F. Gugler, M. Sagebaum, G. Curio, and B. Blankertz. EEG potentials predict upcoming emergency brakings during simulated driving. *J Neural Eng*, 8(5) :056001, Jul 2011.
- [75] S. Perdikis, H. Bayati, R. Leeb, and J. d. R. Millan. Evidence Accumulation in asynchronous BCI. *International Journal of Bioelectromagnetism*, 13(3) :131–132, 2011.
- [76] A. Tzovara, M. M. Murray, N. Bourdaud, R. Chavarriaga, J. D. R. Millán, and M. D. Lucia. The timing of exploratory decision-making revealed by single-trial topographic EEG analyses. *Neuroimage*, 60 :1959–1969, Feb 2012.
- [77] C. P. Robert. *The Bayesian Choice : From Decision-Theoretic Foundations to Computational Implementation*. Springer Verlag, New York, 2nd edition, June 2007. ISBN 0387715983.
- [78] W. Gilks, S. Richardson, and D. Spiegelhalter. *Markov Chain Monte Carlo in Practice : Interdisciplinary Statistics (Chapman & Hall/CRC Interdisciplinary Statistics)*. Chapman and Hall/CRC, 1 edition, December 1995. ISBN 0412055511.
- [79] J. Tanabe, L. Thompson, E. Claus, M. Dalwani, K. Hutchison, and M. T. Banich. Prefrontal cortex activity is reduced in gambling and nongambling substance users during decision-making. *Hum Brain Mapp*, 28(12) :1276–1286, Dec 2007.
- [80] M. G. Philiastides and P. Sajda. EEG-informed fMRI reveals spatiotemporal characteristics of perceptual decision making. *J Neurosci*, 27(48) :13082–13091, Nov 2007.


Bibliographie

- [81] W. Yoshida and S. Ishii. Model-based reinforcement learning : A computational model and an fMRI study. *Neurocomputing*, 63 :253–269, 2005. 11th European Symposium on Artificial Neural Networks.
- [82] G. Corrado and K. Doya. Understanding neural coding through the model-based analysis of decision making. *J Neurosci*, 27(31) :8178–8180, Aug 2007.
- [83] F. Galán, J. Palix, R. Chavarriaga, P. W. Ferrez, E. Lew, C. Hauert, and J. Millán. Visuo-Spatial Attention Frame Recognition for Brain-Computer Interfaces. In Int Conf Cognitive Neurodynamics. Shanghai, China, 2007.
- [84] N. Bourdaud, R. Chavarriaga, F. Galán, and J. del R. Millán. Characterizing the EEG Correlates of Exploratory Behavior. *IEEE Trans Neural Sys Rehab Eng*, 16(6) :549–556, 2008.
- [85] W. J. Krzanowski. Principles of multivariate analysis. Oxford University Press, Oxford, 1998.
- [86] F. Galán, P. W. Ferrez, F. Oliva, J. Guàrdia, and J. Millán. Feature Extraction for Multi-class BCI using Canonical Variates Analysis. In IEEE Int Symp Intelligent Signal Processing. 2007.
- [87] W. J. Freeman. Origin, structure, and role of background EEG activity. Part 4 : Neural frame simulation. *Clin Neurophysiol*, 117(3) :572–589, 2006.
- [88] E. M. Whitham, K. J. Pope, S. P. Fitzgibbon, T. Lewis, C. R. Clark, S. Loveless, M. Broberg, A. Wallace, D. DeLosAngeles, P. Lillie, A. Hardy, R. Fronsco, A. Pulbrook, and J. O. Willoughby. Scalp electrical recording during paralysis : Quantitative evidence that EEG frequencies above 20 Hz are contaminated by EMG. *Clin Neurophysiol*, 118(8) :1877–1888, Aug 2007.
- [89] E. M. Whitham, T. Lewis, K. J. Pope, S. P. Fitzgibbon, C. R. Clark, S. Loveless, D. Delosangeles, A. K. Wallace, M. Broberg, and J. O. Willoughby. Thinking activates EMG in scalp electrical recordings. *Clin Neurophysiol*, 119 :1166–1175, Mar 2008.
- [90] M. Falkenstein, J. Hoormann, S. Christ, and J. Hohnsbein. ERP components on reaction errors and their functional significance : A tutorial. *Biol Psychol*, 51(2-3) :87–107, Jan 2000.
- [91] R. Chavarriaga, P. W. Ferrez, and J. Millán. To Err Is Human : Learning from Error Potentials in Brain-Computer Interfaces. In R.Wang, F. Gu, and E. Shen, editors, Int Conf Cognitive Neurodynamics, pages 777–782. Shanghai, China, 2007.
- [92] H. Zhang. The optimality of naive Bayes. *AA*, 1(2) :3, 2004.
- [93] A. Azzalini. A Class of Distributions Which Includes the Normal Ones. *Scandinavian Journal of Statistics*, 12(2) :pp. 171–178, 1985.

-
- [94] A. Azzalini and A. Capitanio. Statistical applications of the multivariate skew normal distribution. *Journal of the Royal Statistical Society : Series B (Statistical Methodology)*, 61(3) :579–602, 1999.
- [95] D. A. Van Dyk and X.-L. Meng. The Art of Data Augmentation. *Journal Of Computational And Graphical Statistics*, 10(1) :1–50, 2001.
- [96] N. Henze. A Probabilistic Representation of the 'Skew-Normal' Distribution. *Scandinavian Journal of Statistics*, 13(4) :pp. 271–275, 1986.
- [97] S. Frühwirth-Schnatter and S. Pyne. Bayesian inference for finite mixtures of univariate and multivariate skew-normal and skew-t distributions. *Biostatistics*, 11(2) :317–336, Apr 2010.
- [98] N. Bourdaud, R. Chavarriaga, and J. Millán del R. Bayesian detection of asynchronous EEG patterns. *International Journal of Bioelectromagnetism*, 13(2) :106–107, 2011.
- [99] M. Cowles and B. Carlin. Markov chain Monte Carlo convergence diagnostics : a comparative review. *Journal of the American Statistical Association*, 5 :883–904, 1996.
- [100] N. Bhatnagar, A. Bogdanov, and E. Mossel. The computational complexity of estimating MCMC convergence time. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 424–435, 2011.
- [101] M. M. Botvinick, J. D. Cohen, and C. S. Carter. Conflict monitoring and anterior cingulate cortex : An update. *Trends in Cognitive Sciences*, 8(12) :539–546, 2004.
- [102] R. G. O'Connell, P. M. Dockree, M. A. Bellgrove, S. P. Kelly, R. Hester, H. Garavan, I. H. Robertson, and J. J. Foxe. The role of cingulate cortex in the detection of errors with and without awareness : a high-density electrical mapping study. *Eur J Neurosci*, 25(8) :2571–2579, Apr 2007.
- [103] L. P. Sugrue, G. S. Corrado, and W. T. Newsome. Matching behavior and the representation of value in the parietal cortex. *Science*, 304(5678) :1782–1787, Jun 2004.
- [104] D. Lewis. Naive (Bayes) at forty : The independence assumption in information retrieval. *Machine Learning : ECML-98*, pages 4–15, 1998.
- [105] J. d. R. Millan. On the Need for On-Line Learning in Brain-Computer Interfaces. In *Proceedings of the International Joint Conference on Neural Networks*. Budapest, Hungary, July 2004. IDIAP-RR 03-30.
- [106] C. P. Robert. Simulation of truncated normal variables. *Statistics and Computing*, 5 :121–125, 1995. 10.1007/BF00143942.

Nicolas Bourdaud

93 Avenue de Choisy
75013 Paris FRANCE
+336 69 73 13 44
nicolas.bourdaud@gmail.com



Education

- 2006–2012 **PhD, EPFL**, Lausanne.
Machine learning, Bayesian estimation, Advanced signal processing: wavelets and application, Brain-computer interfaces
- 2004–2005 **Cognitive sciences Master's degree, ENS**, Paris.
Motor action, Visual perception, Genetic and neurosciences of cognitive development, Concepts and methods in cerebral imaging (EEG/MEG, fMRI, PET...)
- 2001–2004 **Engineering Master's degree, Ecole Centrale de Nantes (French « Grande Ecole »)**.
Signal processing, Robotics, Modelling, Automation, Programming
- 1998–2001 **Preparatory classes, Lycée Buffon**, Paris.
Post-High School intensive education in highly selective preparatory classes leading to competitive entrance examination to the graduate schools of engineering.
- 1998 **Baccalauréat série S**.
Equivalent to A Levels in Sciences.

Work experience

- 2009–2012 **Research assistant, EPFL**, Lausanne.
Detection of cognitive processes using EEG
- 2006–2008 **Research assistant, IDIAP**, Martigny.
Detection of cognitive processes using EEG
- 2005–2006 **Research engineer, Collège de France**, Paris.
Study of the human motor. I provided design of experiments for others and carried out myself studies.
- 2005
(6 months) **Trainee researcher, Collège de France**, Paris.
Study of the human motor control under constrained motion; control of forces exerted upon a surface. Development of a VR system (C++) (OpenGL) using an experimental force-feedback joystick
- 2004
(5 months) **Trainee researcher, IRCCyN (Nantes) / Univerty of Maribor**, France / Slovenia.
Decomposition of surface EMG signals by Bayesian approach; Given a prior statistical model of the firing pattern of motoneurons and a measure of surface EMG, instants of discharge of neurons have to be inferred.
- 2003
(3 months) **Trainee engineer, NEXYAD**, La Garenne-Colombes–France.
Estimation of dynamic parameters of a model using fuzzy logic; estimation of the adherence of a car while driving (C++)
- 2002
(1 month) **Trainee hand worker, PSA**, Poissy–France.
Worker on the assembly lines

Extra curricular activities

- 2002–2003
(10 months) **President of an association**, *École Centrale de Nantes*, Nantes – France.
President of the association organising the party of the Ecole Centrale de Nantes. The party occurs once a year in the buildings of the school and includes several concerts of many artists. 1500 people came to the party whose budget was nearly 100 000 €.
- 2002
(3 months) **Treasurer of an association**, *École Centrale de Nantes*, Nantes – France.
Treasurer of a list for the election of the BDE (Bureau Des Eleves) the association which organizes the associative life in the school.

Languages

- French **Native**
English **Proficient**
German **Basic**

Computer skills

- OS Linux (system administrator and Maintainer of packages for Debian), Windows
programming C (excellent), C++(good), 3D programming (OpenGL), Linux kernel, GTK+, PHP, notion of python, developer of a set of C libraries for BCI (<http://cnbi.epfl.ch/software/index.html>)
interest Realtime signal processing, Standard (POSIX, ISO C99, ISO C11), Cross-platform
software Matlab, L^AT_EX, Microsoft Visual Studio, GNU development toolchain

Personal interests

- Sport I like practicing adventure sports like snowboarding, rock climbing and scuba diving as well jogging
Activities I enjoy playing dramatics as an amateur, travelling on a motorbike and computer programming.

Publications

- N. Bourdaud, R. Chavarriaga, F. Galán, and J. del R. Millán. Characterizing the EEG correlates of exploratory behavior. *IEEE Trans Neural Sys Rehab Eng*, 16(6):549–556, 2008.
- N. Bourdaud, R. Chavarriaga, F. Galán, and J. d. R. Millán. Characterizing EEG correlates of exploratory behavior. In *6th Forum of European Neurosciences*, 2008.
- N. Bourdaud, R. Chavarriaga, and J. d. R. Millán. Template-based bayesian classification of asynchronous EEG patterns. in preparation.
- N. Bourdaud, R. Chavarriaga, and J. Millán del R. Bayesian detection of asynchronous EEG patterns. *International Journal of Bioelectromagnetism*, 13(2):106–107, 2011.
- N. Bourdaud, R. Chavarriaga, and J. d. R. Millán. Recognition of asynchronous cognitive events from EEG. In *4th International Conference on Cognitive Systems*, 2010.
- A. Tzovara, M. M. Murray, N. Bourdaud, R. Chavarriaga, J. D. R. Millán, and M. D. Lucia. The timing of exploratory decision-making revealed by single-trial topographic EEG analyses. *Neuroimage*, 60:1959–1969, Feb 2012.