



HAL
open science

Efficient whole brain estimation of the haemodynamic response function for TV-regularized semi-blind deconvolution of neural activity in fMRI

Hamza Cherkaoui

► **To cite this version:**

Hamza Cherkaoui. Efficient whole brain estimation of the haemodynamic response function for TV-regularized semi-blind deconvolution of neural activity in fMRI. Imaging. Université Paris-Saclay, 2021. English. NNT: 2021UPAST022 . tel-03269435

HAL Id: tel-03269435

<https://theses.hal.science/tel-03269435v1>

Submitted on 24 Jun 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Efficient whole brain estimation
of the haemodynamic response
function for TV-regularized
semi-blind deconvolution of
neural activity in fMRI
*Estimation efficace de la réponse
hémodynamique cerveau-entier par
déconvolution semi-aveugle, via la
régularisation par variation totale, à partir
de données d'IRM fonctionnelle*

Thèse de doctorat de l'Université Paris-Saclay

École doctorale n° 575: Electrical, Optical, Bio-Physics and
Engineering (EOBE)

Spécialité de doctorat: Imagerie et physique médicale

Unité de recherche: Université Paris-Saclay, CEA, CNRS, Inserm,
Laboratoire d'Imagerie Biomédicale Multimodale Paris Saclay, 91401,
Orsay, France

Référent: Faculté des sciences d'Orsay

Thèse présentée et soutenue à Paris-Saclay, le 03 mars 2021, par

Hamza CHERKAOUI

Composition du jury:

Charles SOUSSEN Professeur, Centrale Supélec	Président
Julien MAIRAL Directeur de recherche, Inria Grenoble	Rapporteur & Examineur
Dimitri VAN DE VILLE Professeur, Ecole Polytechnique de Lausanne	Rapporteur & Examineur
Florence FORBES Directrice de recherche, Inria Grenoble	Examinatrice
Myriam EDJLALI-GOUJON Docteur, Centre Hospitalier Saint Anne	Examinatrice

Direction de la thèse:

Philippe CIUCIU Chargé de recherche, HDR, CEA NeuroSpin	Directeur de thèse
Thomas MOREAU Chargé de recherche, Inria Saclay	Coencadrant de thèse & Examineur
Claire LEROY Chargée de recherche, CEA Service Hospitalier Frédéric Joliot	Coencadrante de thèse & Examinatrice

Nosce te ipsum
— Socrates, 470 - 399 B.C.

Acknowledgements

The work presented in this thesis was only possible with the support and the help of many people. I would like to take the opportunity to express my gratitude.

First I would like to express the deepest appreciation to my supervisors Dr. **Philippe Ciuciu**, Dr. **Claire Leroy** and Dr. **Thomas Moreau** for their continuous support during these three years. I really appreciated their patience, the trust and the freedom they gave me to conduct my work.

I would like to thank Dr. **Abderrahim Halimi** for the contribution and precious advice he gave me during this thesis. I would also like to express my gratitude to Dr. **Bertrand Thirion** for his kind and regular advice. I am also really grateful to Dr. **Nicolas Tournier**, Dr. **Michel Bottlaender** and Dr. **Brice Fernandez** for your advice during our insightful meetings. Thank you for your time, for sharing your knowledge with me.

In addition to the people directly involved in my work, I would like to thank my fellows at **Parietal**, who contributed making the time spent here memorable. To Dr. **Pierre Ablin**, **Thomas Bazeille**, **Zaccharie Ramzi**, **Jérôme-Alexis Chevalier** and **Hugo Richard** I could never thank you enough for spending those difficult periods, you were always here and ready to give me some precious advice and laughs. To Dr. **Arthur Mensch** thank you for your help and for all the useful conversations we had on optimization in high dimension. To Dr. **Kamalaker Reddy Dadi** thank you for your help and for all the useful conversations we had on real data validation, your help was precious. To my former colleagues Dr. **Elvis Dohmatob**, Dr. **Mathurin Massias** and Dr. **Jérômes Dockes** you were models. I also want to thank all my **Metric colleagues**, it has been a great pleasure to work with you during my internship at the beginning of my research journey.

Finally I also want to thank my relatives and friends for supporting me during the PhD thesis and in my life in general. To my parents and my sisters, I thank you for always being here for your kindness and support.

Contents

Notation	vii
General Introduction	1
I Context	7
1 Introduction to fMRI	9
1.1 Introduction to neuroscience	10
1.2 Principle of Magnetic Resonance Imaging (MRI)	13
1.3 Introduction fMRI	19
1.4 Neurovascular coupling in fMRI	23
1.5 Chapter conclusion	27
2 How to segregate the vascular and neuronal components in fMRI?	29
2.1 Introduction to the General Linear Model (GLM)	29
2.2 HRF estimation with fixed neural activity	31
2.3 Neural activity estimation with fixed HRF	34
2.4 Joint estimation of the neural activity and the neurovascular coupling from BOLD signal	37
2.5 Conclusion	38
3 TV problems minimization with iterative algorithms	39
3.1 Introduction to regularization in optimization	39
3.2 Solving a TV regularized problem	42
3.3 Performance comparison	47
3.4 Chapter conclusion	50
II Methodological Developments	53
4 TV problems minimization with learned algorithms	55
4.1 Introduction to TV regularization in 1D	56
4.2 Solving a TV regularized problem	57

4.3	Performance comparison	69
4.4	Conclusion	72
5	Contribution to neurovascular disentangling	75
5.1	Sparsity-based blind deconvolution of neural activation signal in fMRI	76
5.2	fMRI BOLD signal decomposition using a multivariate low-rank model	81
5.3	Conclusion	88
6	Multivariate joint estimation neural activity HRF	89
6.1	Introduction	90
6.2	Multivariate low-rank decomposition of the BOLD signal	90
6.3	Model Validation	97
6.4	Clinical validation at the population level	109
6.5	Conclusion	113
7	Clinical application	115
7.1	Analysis of the Synchropioid cohort	115
7.2	Summary	121
	General Conclusion and Perspectives	123
	Contributions	123
	Limitations & Perspectives	125
	Appendices	129
A	Résumé en français (Abstract in French)	131
	Bibliography	139

Notation

In this manuscript, the mathematical notation follows general typographical conventions.

Symbol	Description
\mathbb{N}	The set of Natural numbers.
\mathbb{R}	The set of Real numbers.
α	Lower-case letters denote scalars.
x_i	subscript i denotes the i^{th} entry of vector \mathbf{x} .
\mathbf{x}	Lower-case bold letters denote vectors.
\mathbf{A}	Upper-case bold letters denote matrices.
\mathbf{A}^\top or \mathbf{x}^\top	Matrix (\mathbf{A}) or Vector (\mathbf{x}) transpose is represented by an exponent \top .
\times	Cross product.
$\ \mathbf{A}\ _F$	The Frobenius norm of a matrix defined as $\ \mathbf{A}\ _2 = \sqrt{\sum_{i=1}^m \sum_{j=1}^n A_{i,j} ^2}$.
$\ \mathbf{x}\ _2$	The ℓ_2 -norm of a vector defined as $\ \mathbf{x}\ _2 = \sqrt{\sum_{i=1}^m x_i^2}$.
$\ \mathbf{x}\ _1$	The ℓ_1 -norm of a vector defined as $\ \mathbf{x}\ _1 = \sum_{i=1}^m x_i $.
$\ \mathbf{x}\ _0$	The ℓ_0 pseudo-norm of a vector defined as $\ \mathbf{x}\ _0 = \sum_{i=1}^m 1_{x_i \neq 0}$
$\mathbf{v} * \mathbf{a} \in \mathbb{R}^{1 \times T}$	The convolution of two signals $\mathbf{a} \in \mathbb{R}^{1 \times (T+L-1)}$ and $\mathbf{v} \in \mathbb{R}^{1 \times L}$
$\mathbf{v} \bar{*} \mathbf{A} \in \mathbb{R}^{P \times T}$	The convolution of each line of $\mathbf{A} \in \mathbb{R}^{P \times (T+L-1)}$ with $\mathbf{v} \in \mathbb{R}^{1 \times L}$
$\mathbf{V} \dot{*} \mathbf{A} \in \mathbb{R}^{P \times T}$	The convolution of each line of $\mathbf{A} \in \mathbb{R}^{P \times (T+L-1)}$ with the corresponding line of $\mathbf{V} \in \mathbb{R}^{P \times L}$
$\nabla \mathbf{x} \in \mathbb{R}^{T-1}$	the first-order difference operator such as $(\nabla \mathbf{x})_i = x_i - x_{i-1}, \forall i \in \{2, \dots, T\}$
$\Gamma(x)$	The Gamma function evaluated on $x \in \mathbb{R}$ defined as $\Gamma(x) = \int_0^\infty x^{t-1} e^{-x} dx$.
$\mathcal{N}(\boldsymbol{\mu}, \mathbf{V})$	Multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix \mathbf{V} .
$\Gamma_0(\mathbb{R}^T)$	The set of convex, proper, lower semi-continuous functions on $\mathbb{R}^T \rightarrow \mathbb{R} \cup +\infty$.

General Introduction

Context & Motivations

FUNCTIONAL magnetic resonance imaging (fMRI) non-invasively records brain activity by dynamically measuring the blood oxygenation level-dependent (BOLD) contrast. The latter indirectly measures neural activity through the neurovascular coupling [Ogawa et al., 1992]. This coupling is usually characterized as a linear and time-invariant system and thus summarized by its impulse response, the so-called **haemodynamic response function** (HRF) [Bandettini et al., 1993, Boynton et al., 1996]. The estimation of this response is of primary interest for clinicians: a change in the haemodynamic response could be linked to the pharmacological mechanism of a drug [Do et al., 2020], the effect of normal aging [West et al., 2019] or the consequence of a neuropathological process [Asemani et al., 2017]. Thus, the HRF could be considered as a **precious biomarker** to investigate the neurovascular function of the brain in healthy or pathological condition. Moreover, its estimation also links the observed BOLD signal to the underlying neural activity, which can in turn be used to better understand cognitive processes. Several methods have been designed to estimate the haemodynamic response function in the context of task-related fMRI (tfMRI). In this setup, the participant is engaged in an experimental paradigm (EP) during the imaging session, which typically alternates between rest and task periods [Friston et al., 1998a, Ciuciu et al., 2003, Lindquist and Wager, 2007, Pedregosa et al., 2015] during which the participant is submitted to sensorymotor stimuli or more cognitively demanding tasks. Commonly, supervised HRF estimation methods fit a model to explain the observed BOLD signal from the EP, namely binary input signal corresponding to stimulus onsets [Ciuciu et al., 2003, Lindquist and Wager, 2007, Vincent et al., 2010, Pedregosa et al., 2015]. A limitation of these approaches is that the EP is used as a surrogate for the neural activity. Therefore these methods cannot be used on resting-state fMRI data (rs-fMRI), where no EP is available. On the other hand, a long-standing literature on fMRI deconvolution methods has emerged since the late 90s to uncover the underlying *activity-inducing signal* at the fMRI timescale of seconds [Glover, 1999, Gitelman et al., 2003, Hernandez-Garcia and Ulfarsson, 2011, Khalidov et al., 2011]. Importantly, a foundational work [Karahanoğlu et al., 2013] has proposed a spatio-temporal model of the underlying activity-inducing signal including both temporal and spatial sparsity-based regularization. By doing so, the recovered neural activity profiles are used to characterize functional networks, hence converging to the original approach proposed by Wu et al. [2013] that reveals functional networks from

deconvolved BOLD signals. Alternatively a recent work [Farouj et al., 2019] has suggested to simultaneously estimate both the neural activity and the HRF profile with a limited parameterization. This approach is often referred to as *semi-blind deconvolution schemes* of the BOLD signal. Farouj et al. [2019] relies on the hypothesis of a constant block signal for the neural activity as initially proposed in Karahanoğlu et al. [2013]. However in [Farouj et al., 2019], the authors are able to infer the haemodynamic parameters to deal with the magnitude and delay ambiguities between the neural input and the HRF.

Contributions

IN this thesis we propose a novel approach to disentangle the *activity-inducing signal* (neural activation signal) from the *neurovascular coupling* (HRF) for both task-related and resting-state fMRI (tfMRI and rs-fMRI, respectively). In short, the general idea is to define a model that is able to separate these two components at the scale of the whole brain, along with a suitable optimization algorithm as a means to perform parameter estimation in an efficient way. After an introduction to neuroimaging techniques, in particular to magnetic resonance imaging (MRI), as well as to the mechanisms underlying the BOLD signal measured in fMRI (chapter 1), we present, in chapter 2, the state-of-the-art methods that estimate the HRF or the neural activity signals from fMRI data. In particular, we focus on the seminal work introduced by Karahanoğlu et al. [2013] that models the neural activation signal as a piecewise constant block time series. This hypothesis is especially relevant for tfMRI where the EP has such block structure and potentially the neural signals too, but it also makes sense in the context of rs-fMRI as a means to regularize the spontaneous neural activity.

This piecewise constant hypothesis on the neural inputs is enforced using 1-dimensional (1D) Total Variation (TV) regularization which promotes sparsity of the first-order derivative of the neural signal and thus leads to block-shaped signals. The resolution of such deconvolution problem with TV regularization can be computationally expensive. After a review and a benchmark of classical approaches to solve 1D TV-regularized problems (chapter 3), we theoretically compare the performances of the **analysis vs synthesis** approaches in chapter 4, and show that the analysis algorithms can be much more efficient than their synthesis counterparts for the problems under study. In this work, we also propose a way to perform **differentiable algorithm unrolling** [Gregor and Le Cun, 2010] for proximal gradient descent (PGD) with TV-regularized optimization problems by either directly computing the gradient of the TV proximal operator with an analytic formula (LPGD-Taut) or computing it with a nested unrolled algorithm (LPGD-LISTA). Comparison on real fMRI data shows promising results for learned optimization algorithm based on this unrolling compared to iterative algorithms.

In order to study the haemodynamic response over the whole brain, we first propose to estimate voxelwise HRFs (first part of chapter 5). This univariate method is based on the same block hypothesis for the neural activity and introduces a simplistic but novel modeling for the HRF. The haemodynamic response is modeled as a time dilated version of the canonical **HRF** $v_{\text{ref}}(t)$, whose definition relies on the difference between two Gamma distributions Friston et al. [1998a]. This parameterization allows us to summarize the HRF with a single dilation parameter δ . One limitation of this contribution is its massively univariate aspect. Indeed, as we estimate both the neural input signal and the HRF, the total number of parameters in

this model is similar to the number of scans which could lead to overfitting. To cope with this issue, we propose two contributions [Cherkaoui et al., 2019b, 2020a] that extend the previous model in a multivariate setting in order to capture more efficiently the neurovascular coupling. In the second part of chapter 5, we introduce K **temporal components** and their corresponding **spatial maps** to encode the neural activity signals and respectively localize their contribution to the measured BOLD fMRI data. Moreover, in Cherkaoui et al. [2020a], we extend this model to estimate the neurovascular coupling over the whole brain as a time dilated HRF for each region from a **predefined brain parcellation** (chapter 6). At the subject level, we demonstrate that the estimated **neural activity spatial maps are related to meaningful functional networks**. To assess the significance of this approach at the population level, we statistically demonstrate that a pathology like stroke or a condition like normal brain aging induce longer haemodynamic delays in certain brain areas and that this haemodynamic feature may be predictive of the individual status in a machine learning (classification) task. Moreover, we investigate in chapter 7 the usefulness of our method in a clinical context. We analyze PET/fMRI imaging data recorded simultaneously on a PET/MR system installed at the Service Hospitalier Frederic Joliot (CEA/DRF/Joliot, Orsay) in the context of a protocol called Synchropioid. The latter aims to understand the pharmacological mechanism of the buprenorphine – an analgesic drug – on the brain in order to understand and unveil its side effects like some habituation or the need to rapidly increase the dose to obtain an analgesic effect in some specific patients. Using the PET data, we aim to respectively localize where the buprenorphine is fixated in the brain and in which brain regions it is mostly concentrated. Next, from the concomitantly recorded rs-fMRI data, our objective is to study the effect of the drug on the haemodynamic system, in particular whether it impacts the HRF profile and to what extent the observed effect, if any, is consistent with our findings in PET imaging. In short, our preliminary results on four subjects (two healthy volunteers treated with the analgesic dose vs two healthy controls) indicate that the buprenorphine is mostly concentrated in the Putamen (subcortical nucleus) and in the Insula and that these regions are mostly affected by a deceleration of the neurovascular coupling, observed through longer time-to-peaks in the HRFs.

Thesis Outline

This dissertation is organized as follows:

Chapter 1: *Introduction to fMRI* proposes an introduction to the neuroscience and their most common imaging techniques. We focus more specifically on the MRI, next on functional MRI and finally introduce the neurovascular coupling along with the definition of the HRF.

Chapter 2: *How to segregate the vascular and neuronal components in fMRI?* introduces the state-of-the-art approaches to disentangle the vascular coupling from the neural activity involved in the fMRI signal.

Chapter 3: *TV problems minimization with iterative algorithms* proposes to summarize the methods used to solve the optimization problem that involved the Total Variation regularization, as this regularization is central in our contribution to estimate the HRF.

Chapter 4: *TV problems minimization with learned algorithms* introduces our novel approach to solve more efficiently the optimization problem that involved the Total Variation

regularization based on neural network.

Chapter 5: *Contribution to neurovascular disentangling* proposed to expose our two first contributions done to disentangle the vascular coupling from the neurovascular coupling.

Chapter 6: *Multivariate joint estimation neural activity HRF* exposes our major contribution to disentangle the vascular coupling from the neurovascular coupling in a multivariate fashion for the whole brain.

Chapter 7: *Clinical application* proposes to apply our multivariate disentangling method to the Synchronoid cohort to investigate the pharmacological mechanism of the buprenorphine.

Publications

Articles submitted in Peer-Reviewed Journals

Hamza Cherkaoui, Thomas Moreau, Abderrahim Halimi, Claire Leroy and Philippe Ciuciu. *Multivariate semi-blind deconvolution of fMRI time series*. under review at NeuroImage

International Conferences Paper Presented with Reading Committee and Proceedings

Hamza Cherkaoui, Jeremias Sulam and Thomas Moreau. *Learning to solve TV regularized problems with unrolled algorithms..* Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS), Virtual, 2020. Poster.

Hamza Cherkaoui, Thomas Moreau, Abderrahim Halimi and Philippe Ciuciu. *fMRI BOLD signal decomposition using a multivariate low-rank model..* Proceedings of the 27th European Signal Processing Conference (EUSIPCO), A Coruña, Spain, 2019. Oral.

Hamza Cherkaoui, Thomas Moreau, Abderrahim Halimi and Philippe Ciuciu. *Sparsity-based blind deconvolution of neural activation signal in fMRI*. Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Brighton, United Kingdom, 2019. Poster.

Hamza Cherkaoui, Loubna El Gueddari, Carole Lazarus, Antoine Grigis, Fabrice Poupon, Alexandre Vignaud, Sammuell Farrens, Jean-Luc Starck and Philippe Ciuciu. *Analysis vs synthesis-based regularization for combined compressed sensing and parallel MRI reconstruction at 7 Tesla..* Proceedings of the 26th European Signal Processing Conference (EUSIPCO), Rome, Italy, 2018. Poster.

Abstract in peer-reviewed international conferences without proceedings

Hamza Cherkaoui, Thomas Moreau, Abderrahim Halimi, Claire Leroy and Philippe Ciuciu. *Data-driven haemodynamic response function estimation: a semi-blind multivariate deconvolution of the fMRI signal*. Organization for Human Brain Mapping, Virtual, 2020. Poster.

Part I

Context

Introduction to fMRI

Chapter Outline

1.1	Introduction to neuroscience	10
1.1.1	Description of the brain	10
1.1.2	Neuroscience and the usual neuroimaging modalities.	11
1.2	Principle of Magnetic Resonance Imaging (MRI).	13
1.2.1	Signal-to-noise-ratio (SNR)	14
1.2.2	Polarization using radio-frequency (RF) waves	15
1.2.3	Relaxation	15
1.2.4	Weighting contrast	15
1.2.5	Measuring the FID signal	16
1.2.6	Spatial encoding	17
1.2.7	Image reconstruction	18
1.2.8	MRI in clinical practice	19
1.3	Introduction fMRI.	19
1.3.1	Principle of Functional MRI (fMRI).	19
1.3.2	Classical fMRI preprocessing	21
1.3.3	fMRI acquisition.	22
1.4	Neurovascular coupling in fMRI	23
1.4.1	Introduction to the neurovascular coupling and its modeling	23
1.4.2	Clinical application.	25
1.5	Chapter conclusion	27

IN this chapter we introduce the brain and its structure and give a general introduction to neuroscience before focusing more specifically on functional Magnetic Resonance Imaging (fMRI), the imaging technique dedicated to probe brain function both at rest and during task performance. In regards to the imaging technique, we will first describe the blood oxygenated level dependent (BOLD) effect at the origin of the signal measured in fMRI, and then will cover the way the data is collected from an acquisition perspective up to the classical preprocessing pipeline. Finally, we will focus on the fMRI signal and more specially on the neuro-vascular coupling involved in the signal generation, how to model it and how this coupling can be relevant as a biomarker for further statistical analysis.

1.1 Introduction to neuroscience

1.1.1 Description of the brain

As the remaining of the on manuscript focus on functional Magnetic Resonance Imaging (fMRI), its signal (the BOLD contrast) and more specifically on the neurovascular coupling which is at the core of the fMRI signal generation, we first briefly propose a general introduction to the brain to provide a meaningful context.

The brain structure

The brain is part of the central nervous system and is composed of three major parts: (i) the brainstem that makes the junction with the rest of the nervous system, (ii) the cerebellum that regulates motor movements and coordinates voluntary movements and the neocortex that is notably responsible for the higher cognitive functions such as the memory, the learning or the language for example.

In the brain, we can differentiate two major types of cells: the glial cells and the neurons. Briefly speaking, neurons are responsible for all sensory-motor and cognitive functions while the glial cells are responsible for supporting and protecting the neurons. The brain contains about 86 billion neurons. The neurons are connected one another in complex structure, the neuropil, which defines specific functionality. The population of neurons is densely connected and mainly located in the gray matter. The main second component of the brain tissues is the white matter where lie the myelinated fiber bundles of neuronal axons. Myelin is a lipid-rich substance that insulates axons to allow for a better propagation of electric spikes and thus connects distant gray matter regions across the whole brain.

This mechanism of distant neuronal connections through the white matter enables efficient high cognitive functions in the brain.

Introduction to the neuron

The neuron contains a nucleus inside the cell body. Neurons are connected to each other by extensions at the level of their cell bodies named dendrites and their axon terminals, namely synaptic endings. The axon is surrounded by myelin sheath and schwann cell which are at the base of the white matter, see illustration [Figure 1.1-1 \(b\)](#) to observe a schema of multipolar neuron. There exists a hundred of different types of neurons in the brain with specific message-carrying abilities, [Figure 1.1-1 \(a\)](#) is an illustration of staining of pyramidal neurons [[Del Río and De Felipe, 1994](#)] in the primate neocortex with which we notice the long axon and the dendrite connecting the neurons to each others.

The signal spreading across neuron's membranes from the synapse to the dendrites is partly electrical and partly chemical. When a neuron is excited on its axon terminals from others neurons' dendrites, if the overall voltage exceeds a certain threshold the neuron itself will in turn produce an electrochemical spike, called *action potential*, that permits to propagate the signal to other neurons. By cascading these elementary action potentials across the vastness of neural connections, the information flow may circulate at high speed in the whole brain to synchronize distant brain regions and thus underpin cognition functions.

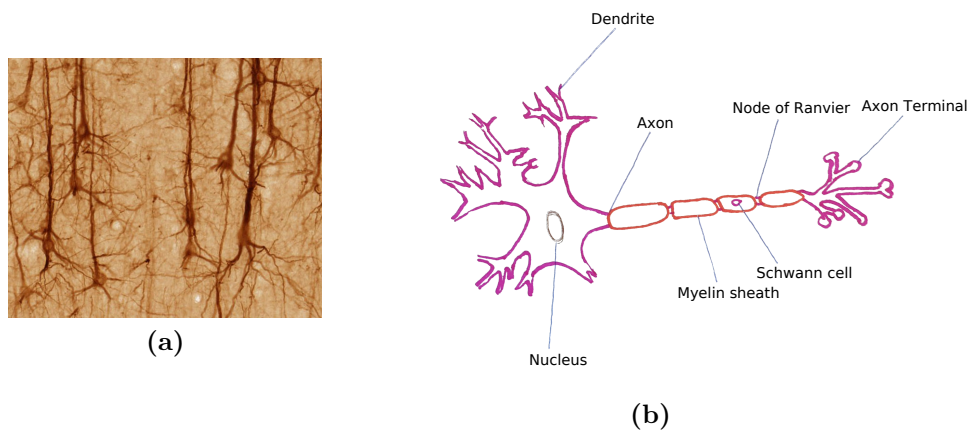


Figure 1.1-1 – (a) **Multipolar neuron photo in the primate neocortex** - (b) **Multipolar neuron illustration**. In (a) we display SMI32-stained pyramidal neurons which allows us to observe the complex connection produce by the axon terminals and dendrites. In (b) we propose a schematic illustrate of the structure of a neuron from the dendrite attached to the cell body with its nucleus at its center, to the axon, envelopped with a myelin sheath produced by the Schwann cell separated by the node of Ranvier, up to the axon terminal.

We briefly focus on the glial cells, as they have a determinant role for the remaining for our work. As explained, those cells support and protect the neurons to ensure their proper working. We can distinguish four different types, see [Figure 1.1-2](#):

- *Ependymal* are responsible for the production of cerebrospinal fluid that surround the brain and protects it from shocks.
- *Astrocytes* have numerous different roles. Among them they provide neurons with nutrients and they serve as intermediaries in neuronal regulation of blood flow.
- *Oligodendrocytes* provide support and insulation to axons by creating the myelin sheath (the white matter).
- *Microglial* are macrophages that actively participate in immune defense of the brain.

In the present work, defining a model to capture the vascular coupling to the neural activation signals is indirectly interested in astrocytic activity. Therefore, the role of astrocytes in modulating the cerebral blood flow (CBF) to the metabolic demand of brain activity is of fundamental importance to provide an appropriate and consistent energy supply to support brain function.

1.1.2 Neuroscience and the usual neuroimaging modalities

Modern neuroscience studies the brain in healthy or pathological condition to understand either the cognitive functions (language, consciousness, memory, time perception, etc.) or diseases' mechanisms in neurogenerative pathology (e.g. Parkinson, Alzheimer). To this end, modern neuroscience heavily relies on neuroimaging, both structural and functional to analyse the brain structure and function. To that aim, several imaging techniques can be employed (see list in [Figure 1.1-3](#)), some of them being more sensitive to the timing of events in the brain with a millisecond temporal resolution (EEG, MEG), while others (i.e. MRI




Figure 1.1-2 – **Illustration of the different types of Glial cells.** The figure displays the four different types of glial cells present in the brain: the pia mater cells, astrocytes (with their endfoot), the oligodendrocytes and the microglial cells. The neurons are represented with their synapses and axons and the capillary in red. Source: Atlas de neurosciences humaines de Netter Felten and Shetty [2011].

and fMRI) have a better ability to probe brain structure and localize which brain regions are engaged during task performance or even at rest due to a millimetric spatial resolution. Last but not least, nuclear imaging techniques like PET may access molecular sensitivity to decipher metabolic activity (e.g. ATP synthesis or glucose consumption) that cannot be reached using non-invasive and non-ionizing techniques such as fMRI.




Figure 1.1-3 – **Imaging modalities represented along the two axes describing the spatial and temporal resolution.** Additionally, the color indicates the degree of invasiveness. The most common imaging techniques in neuroscience: the Electroencephalography (EEG), the Magnetoencephalography (MEG), the structural Magnetic Resonance Imaging (MRI) and the functional Magnetic Resonance Imaging (fMRI) and the nuclear imaging modalities PET and SPECT, namely Positron Emission Tomography and Single Photon Emission Computerized Tomography. Note the of temporal resolution of the fMRI modality while being non-invasive.

Neuroimaging techniques

- *Electroencephalography* (EEG) is a time-resolved modality (a millisecond time resolution) that records in a non-invasive manner the voltage fluctuations, resulting from the neural activity (i.e. local field potentials as a marker of post-synaptic activity) on the scalp surface [Logothetis et al., 2001]. An EEG acquisition is obtained by placing an headset made of 64 to 256 electrodes to record the potential fluctuations in the range of microvolts (μV).
- *Magnetoencephalography* (MEG) is a time-resolved modality (a millisecond time resolu-

tion) that records in a non-invasive manner the fluctuations in magnetic fields generated by the neuronal activity in the cortex, resulting from the neural activity on the scalp surface. A MEG acquisition is obtained by placing the head of the subject inside a recording large and fixed helmet. As the recorded magnetic fields are of very low intensity (in the order of 10^{-13} Tesla), the MEG machine is installed within a Faraday cage to remove the contribution of external magnetic perturbations due for instance to the Earth's magnetic field.

- *Positron Emission Tomography* (PET) is an imaging technique that localizes a radiotracer to quantify various physiological activities such as neurotransmission, metabolic processes or blood flow. The radiotracer is injected at a tracer dose which corresponds to a small concentration that does not noticeably influence the pharmacology or pharmacokinetics of the process being imaged. This radiotracer is marked by an isotope (carbon, fluorine, nitrogen, oxygen, etc.) that emits positrons whose annihilation produces the emission of two photons. The detection of the trajectory of those photons by external detectors positioned at different orientations in the PET camera localizes the emission position and therefore the concentration of the radiotracer in the brain.
- *Structural Magnetic Resonance Imaging* (MRI) is a non-invasive and non-ionizing imaging technique that images the morphological structure of the brain (gray and white matter, cerebrospinal fluid, skull, etc.) by manipulating the magnetization of the hydrogen atoms contained in the water molecules. This process involves a polarization step which consists in deviating the spins of hydrogen atoms from their equilibrium and then measuring how fast they return to this equilibrium state. Because the different tissues present in the brain have different relaxation properties, one may create a contrast between these tissues and create precise images that reflect the anatomy. The principle of MR data acquisition and image reconstruction are given in [section 1.2](#).
- *Functional Magnetic Resonance Imaging* (fMRI) is a non-invasive techniques that records indirectly the neural activity by taking advantage of the magnetic fluctuations induced by the level of oxygenation in the blood. Its temporal resolution ranges around 1 to 2 seconds and its spatial resolution is nowadays around 1.5 mm isotropic. This technique thus captures low frequency activity compared to EEG and MEG. As fMRI data is at the core of this thesis, we will dedicate [section 1.3](#) to its introduction, from the origin of the measured signal to the common preprocessing pipeline and the type of acquisitions one usually meets.

1.2 Principle of Magnetic Resonance Imaging (MRI)

Magnetic Resonance Imaging is a non invasive imaging technique widely used in the medical field, mainly to probe soft tissues in the human body. The principle of MRI is based on the spin property of some molecules, such as hydrogen, phosphorus or sodium, was first described in 1940 [[Bloch et al., 1946](#)], which have a non-zero magnetic moment due to their even number of nucleons (protons and neutrons).

As the water molecules are the most prominent in the human body, one can easily manipulate the spins of hydrogen atoms embedded in water molecules to create MR images. In essence, the vast majority of MRI exams typically produce maps of “water density” in first approximation.

A MR system actually relies on three major components:

- A large static magnetic field, denoted B_0 (e.g. 3 Tesla), which gives the MRI machine its distinctive aspect, cf. [Figure 1.2-4](#). The role of this static B_0 field is to align all hydrogen atoms present in water molecules with its main direction.
- An organ-specific coil that is used to deliver the radio-frequency waves at the Larmor frequency ($\omega_0 = \gamma B_0$ and $\gamma = 42.58\text{MHz/T}$ defines the gyromagnetic ratio for the hydrogen atom) to generate the resonance phenomenon and collect in return the MR signal. In the MR field, one refers to this radio-frequency wave as the B_1^+ (transmit) and B_1^- (receive) fields. For instance, the coil may have a birdcage form in neuroimaging and can comprise a large number of channels to boost the signal-to-noise-ratio (SNR). The coil is also tuned for the type of imaging undertaken. This means that a coil used for proton imaging cannot be used for Sodium imaging as the gyromagnetic ratios γ differ between the two atoms.
- The magnetic field gradients system that allows for the spatial encoding of the MR signal into the k-space or the spatial frequency domain. This encoding can be done in 2D or in 3D leading to 2D and 3D imaging respectively. Most of MR pulse sequences rely on 2D imaging for the sake of speed. Hence, 3D volumes are then reconstructed by stacking multiple 2D slices.



Figure 1.2-4 – **MRI scanner of 3T**. Example of a 3T MRI system installed at Neurospin (CEA Saclay, France) to perform structural or functional neuroimaging.

1.2.1 Signal-to-noise-ratio (SNR)

The image resolution in MRI depends on several factors but the larger the SNR, the higher the resolution one can reach. As B_0 contributes to the gain in SNR more than linearly ($\text{SNR} \propto B_0^{1.65}$), higher image resolution is expected at 7 Tesla compared to 3 and 1.5 Tesla. Currently, the more powerful scanner for clinical usage is at 10.5T in the CMRR center, Minneapolis (MN, USA). However, the 11.7T Iseult scanner installed at NeuroSpin¹ should be available for human exams in 2022 after a 1-year safety check period. As a gain in image resolution requires longer acquisition times, most often the MR exams are performed at millimetric resolution to keep scan time compatible with clinical routine. This allows the physician to maintain a scan time close to a few minutes per imaging contrast. Higher

¹see <https://www.cea.fr/english/Pages/News/Iseult-MRI-Magnet-Record.aspx>

resolution imaging (e.g. 600 μm isotropic) can be performed using dedicated acceleration techniques, as detailed hereafter.

1.2.2 Polarization using radio-frequency (RF) waves

During the acquisition, the aligned spinning molecules of hydrogen are polarized with a RF wave (i.e. B_1^+ magnetic field delivered through a transmit coil) which produces a tilt of the spinning molecules. The frequency of the RF wave must be equal to the Larmor frequency of the hydrogen atoms to tilt them in the plan (say (Oxy)) which is orthogonal to the B_0 field direction (Oz).

1.2.3 Relaxation

Once the RF pulse is stopped, the polarized molecules lose their excess of energy by emitting back a magnetic echo to return to their equilibrium. This process is accompanied with a double relaxation, the longitudinal relaxation $M_z(t)$ that describes the regrowth of the magnetization along (Oz) and the transverse relaxation $M_{xy}(t)$ that describes how fast the magnetization vanishes in (Oxy). The corresponding dynamics are given by the simplest form of the Bloch equations:

$$M_z(t) = M_0(1 - e^{-t/T1}) \quad (1.1)$$

$$M_{xy}(t) = M_0 e^{-t/T2} \quad (1.2)$$

The longitudinal relaxation $M_z(t)$ reflects the spin-lattice interaction while the transverse relaxation $M_{xy}(t)$ is associated with the local spin-spin coupling. The two key parameters that quantify these phenomena are called the $T1$ and $T2$ parameters and basically represent 63 % of recovery of M_0 (the total magnetization at equilibrium) and 63 % of signal decay in the transverse plan. These parameters vary across tissues and are responsible for the contrasts in MR images. In short, $T1$ is typically one order of magnitude larger than $T2$ in viscous liquids that correspond to biological samples and they range as follows: $T1 \in (250, 4000)\text{ms}$ while $T2 \in (70, 2000)\text{ms}$. Note that the upper bounds are reached for water or cerebro-spinal fluid (CSF). The MRI signal, also called the free induction decay (FID), that is actually measured corresponds to the transverse relaxation $M_{xy}(t)$. However, the measuring process, also termed a pulse sequence (e.g. spin or gradient echo), is performed at steady state and thus involves both $T1$ and $T2$ weighting simultaneously (see details hereafter).

1.2.4 Weighting contrast

The MRI scanner is able to produce multiple brain imaging contrasts depending on the setting of acquisition parameters (echo time:TE, time of repetition: TR).

The time of repetition (TR) corresponds to the time interval separating two consecutive applications of RF pulses and thus controls the recovery of $M_z(t)$. As such, it controls the $T1$ -weighting. A long TR value permits the recovery of M_0 in all tissues and is not discriminant in $T1$. In contrast, a short TR value (e.g. $< 500\text{ms}$ in 2D imaging) segregates the tissues according to their $T1$ value. The white matter being associated with a shorter $T1$ than the grey matter, this explains why white matter is bright in $T1$ -weighted ($T1w$) images.

The CSF is dark in T1w images. However, this happens only when the second acquisition parameter TE, i.e. the echo time, is well calibrated.

Physically, the TE value corresponds to the time point at which the echo is centered and thus it matches the middle of the readout in k-space. TE controls the T_2 weighting, i.e. how much the transverse relaxation diminishes and contributes to the measured signal. A short TE value (e.g. 5ms) does not discriminate different tissues while a longer setting ($TE = 30\text{ms}$) permits to disentangle the contribution of the white and the grey matter.

To summarize, a pair of short (TE, TR) values is used for T1w imaging as the T2 contribution is unweighted in the measured FID signal while the T1 one is enforced. On the contrary, a pair of long (TE, TR) parameters is used for T2w imaging as the T1 contribution is unweighted in the measured FID signal while the T2 one is enforced. Another important setting that one can meet in clinical routine is proton density (PD) imaging which corresponds to unweighting the T1 and T2 contributions to better approximate M_0 and as such is used with long TR and short TE.

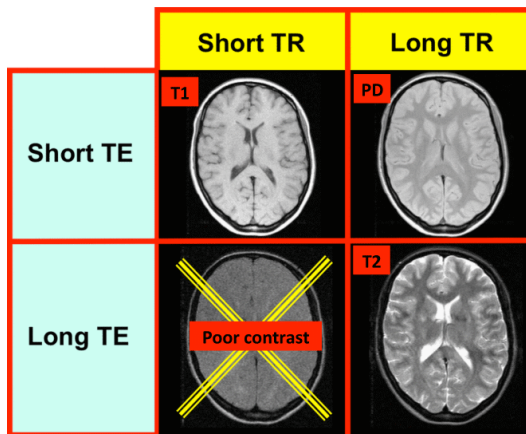


Figure 1.2-5 – MRI T1 contrast, T2 contrast and Proton Density contrast *w.r.t* repetition time (TR) and the echo time (TE). (top-left) the T1 contrast is obtained with a short TR and a short TE, (top-right) the T2 contrast is obtained with a long TR and a long TE and (bottom-right) the Proton Density contrast is obtained with a long TR and a short TE.

1.2.5 Measuring the FID signal

Except for T2w imaging, fast 2D imaging is currently achieved using the gradient echo pulse sequence. It consists in reversing the gradient on the readout axis to form the echo, in contrast to the RF spin echo sequence, which relies on a π -refocusing pulse. The gradient echo sequence is the short name of the the gradient recalled or refocused echo sequence (GRE). In such setting, the flip angle of the B_1^+ RF field, called θ is lower than 90° .

However, the T2 weighting is no longer preserved. Instead it is replaced by a T_2^* weighting, where T_2^* is shorter than T2 as it reflects the local inhomogeneities induced by magnetic field gradient irregularities as follows: $1/T_2^* = 1/T_2 + \Delta B_0$, where the ΔB_0 describes the local fluctuations of the B_0 field induced by the magnetic field gradients. However, the transverse magnetization $M_{xy}(t)$ which is the actual MRI signal has no longer a T2 decay but a T_2^* instead. This sequence can be also easily extended to 3D imaging.

After an initial transient state, the steady state of the dynamic equilibrium can be reached. This means that from TR to TR the value of M_z and M_{xy} remains the same. Typically, the

steady-state MRI signal equation reads as follows:

$$M_{xy}(TE) = \frac{M_0 \sin \theta (1 - E_1)}{1 - \cos \theta E_1} e^{-TE/T2^*} \quad (1.3)$$

where θ is typically set to the Ernst angle:

$$\theta_E = \cos^{-1} E_1 \quad (1.4)$$

$$\text{with } E_1 = e^{-TR/T1}. \quad (1.5)$$

The MRI signal thus depends on the tissue properties and varies between the gray and white matter, the fat and cerebro spinal fluid (CSF). In practice, this signal is measured through a receiver coil (e.g. B_1^- magnetic field). In brain imaging, most often the same coil is used for transmission and reception up to 3 Tesla and is located around the head (birdcage form). At ultra-high magnetic fields ($\geq 7T$), specific designs of transmit and receiver coils with a varying number of channels are available: the larger the number of channels in transmit mode, the more homogeneous the RF field can be delivered whereas in reception the larger the number of channels, the better the SNR although the latter is varying in space with a stronger signal in the outer surface compared to the inner parts of the brain (e.g. subcortical nuclei).

1.2.6 Spatial encoding

In MRI, spatial encoding of the resonating hydrogen atom's spins is performed using gradient magnetic fields (G_x , G_y and G_z), which are applied in the three spatial dimensions to carry out slice selection (e.g. G_z along the head feet axis), phase (e.g. G_y along the posterior anterior axis) and frequency (e.g. G_x In 2D imaging, slice selection is applied during the RF pulse delivery as a means to isolate the resonating spins of a single plan (i.e. $\omega_0(z) = \gamma(B_0 + G_z \cdot z)$). Next, the linear variations of G_x and G_y introduce a one-to-one mapping between the frequency (centered around the Larmor frequency) and phase of resonating spins and their spatial location. In practice, the phase encoding gradient G_y is then applied for a short period of time ΔT as follows: $\omega_0(y) = \gamma(B_0 + G_y \cdot y)$ with $\Delta\varphi(y) = \gamma G_y \cdot y \Delta T$. Last, the frequency encoding gradient G_x is applied during signal acquisition ($\omega_0(x) = \gamma(B_0 + G_x \cdot x)$), i.e. when opening the analog-to-digital converter for accumulating a discrete version of the FID signal. Importantly, the manipulation of gradient profiles (say in 2D G_x and G_y) over time permits to cover a wide range of spatial frequencies ($k_x k_y$) which are defined as follows:

$$k_x(t) = \frac{\gamma}{2\pi} \int_0^t G_x(\tau) d\tau$$

$$k_y(t) = \frac{\gamma}{2\pi} \int_0^t G_y(\tau) d\tau.$$

Hence, the area under the gradient waveforms defines the spatial frequencies and this shows that $\mathbf{k} = (k_x, k_y)$ defines a 2D trajectory in the Fourier domain:

$$\mathbf{k}(t) = \frac{\gamma}{2\pi} \int_0^t \mathbf{G}(\tau) d\tau.$$

The time spent to traverse the k-space depends on the type of imaging. If the trajectory is segmented in multiple shots, the acquisition gets slower as for instance a single line of

k-space is collected per TR. If one wants to get faster such as for fMRI, we can use single shot trajectories like zigzag patterns also called Echo Planar Imaging (EPI). In the latter situation, a single k-space plan can be collected in about 30 ms for conventional image resolution.

As the FID signal is the superimposition of all resonating spins over the whole field of view (FOV), i.e. across spatial positions $\mathbf{r} = (r_x, r_y)$, the measured signal in MRI $s(t)$ – at least without taking all sources of artifacts into account (e.g. off-resonance effects due to B_0 inhomogeneities) is given by the following noisy discrete time series:

$$y(t_i) = s(t_i) + n_i(t) \quad \forall t_i = i \times \text{TR} \quad (1.6)$$

$$s(t) = \int_{\text{FOV}} x(\mathbf{r}) e^{-2i\pi \mathbf{k}(t) \cdot \mathbf{r}} d\mathbf{r} \quad (1.7)$$

where $n(t)$ is a 2D Gaussian white noise in the k-space domain and x defines the image of the organ under investigation. The signal $s(t)$ is nothing but than the Fourier transform of the original 2D image x corresponding to the tilted slice of the 3D volume.

To sum up, the data in 2D MRI is collected in the 2D Fourier space also called the k-space and represents the spatial frequencies in the MR image. In conventional fMRI studies, one usually collects a single volume slice by slice as this allows to speed up the acquisition along the third dimension using multiband multislice imaging techniques [Moeller et al., 2010]. Other imaging sequences (e.g. 3D GRE EPI) are available for instance at 7 Tesla to reach higher spatial resolution. However, in that case, the temporal resolution is slower than 1 s.

1.2.7 Image reconstruction




Figure 1.2-6 – **k-space of a brain.** (right) displays the reconstructed image of a brain seen from the sagittal plane with its associated k-space in (left). Source: mriquestions.com

The simplest way to reconstruct an MR image from k-space data simply consists in computing an inverse fast Fourier transform, see Figure 1.2-6. However, nowadays, in an attempt to speed up MRI exams, one collects less data in k-space either using parallel imaging (SENSE or GRAPPA) [Pruessmann et al., 1999, Griswold et al., 2002] or using compressed sensing techniques [Lustig et al., 2007]. The former consists in regularly under-sampling the k-space phase encoding steps while the latter performs pseudo-random variable density under-sampling to maximize the local incoherence between the acquisition space and the sparsifying transform domain (i.e. spatial wavelets domains in MRI) [Lustig et al., 2007, Lazarus et al., 2019]. Indeed, the MR images can be well represented in a sparse manner using wavelet transforms. In this case, the image reconstruction gets more complicated

as it requires to solve an ill-posed inverse problem using appropriate convex optimization algorithms [Beck and Teboulle, 2009, Chaâri et al., 2011].

1.2.8 MRI in clinical practice

MRI is currently used for clinical diagnosis of brain tumors and other neurological disorders (e.g. multiple sclerosis) or brain pathology. The radiologists examine several 3D MR images of the brain structure using different imaging contrasts (T_1 , T_2 , FLAIR² and T_2^*) as the latter provide complementary information on the brain tissue properties. Additionally, diffusion-weighted MR imaging is used in daily routine as it is very sensitive and specific to Stroke acute episodes in the first hours. Interestingly, the T_2^* -weighted imaging contrast is sensitive to iron concentration. This can be very useful to detect small hemorrhages in tiny vessels by combining the magnitude and phase information as done in susceptibility weighted imaging. As such T_2^* -weighted imaging is sensitive to blood properties. Although its role for structural imaging is not negligible, it actually occupies a privileged role in functional MRI, as explained hereafter.

1.3 Introduction fMRI

In this section, we first describe the fMRI signal and its origin, and then we detail the classical statistical processing steps from the raw fMRI images and the two main kinds of exams that are used with this imaging technique.

1.3.1 Principle of Functional MRI (fMRI)

In this part, we first describe the effect of the oxygenation of the blood onto the MRI signal and see that this effect – termed the BOLD effect – allows to indirectly measure fluctuations in neural activity, which is at the core of fMRI.

BOLD contrast

In Ogawa et al. [1990, 1992], Dr Ogawa demonstrated that the blood volume increased in the human visual cortex during visual stimulation using BOLD contrast. This first contribution was at the origin of the fMRI imaging technique.

The main mechanism that underlies the fMRI technique relies on the fluctuation of oxygenated blood in the brain tissue. When a group of neurons fires together they will consume oxygen and glucose. To provide a continuous supply, the astrocytes, see Figure 1.1-2, will regulate the blood flow to ensure that oxygenated blood irrigates the tissue. This vasodilation mechanism is the ground of the fMRI imaging technique. Upon activation, oxygen is extracted by the cells, thereby increasing the level of deoxyhaemoglobin in the blood. This is compensated for by an increase in blood flow in the vicinity of the active cells, leading to a net increase in oxyhaemoglobin, see Figure 1.3-8. This mechanism induces local B_0 inhomogeneities. As reported in Figure 1.3-7, the oxyhemoglobin (oxygenated blood) has a diamagnetic property which will produce an increase of the magnetic signal during the echo time and hence a longer T_2^* relaxation parameter. On the contrary, the deoxyhemoglobin

²FLuid Attenuated Inversion Recovery pulse sequence that is used to suppress the bright signal in the CSF as this is a T_2 sequence.




Figure 1.3-7 – **Para-magnetic property of the oxygen molecule.** The deoxyhemoglobin (de-oxygenated blood) has a paramagnetic property, in contrary of the oxyhemoglobin. This magnetic feature is at the origin of the BOLD contrast. Source Huettel, Song & McCarthy, 2004, Functional Magnetic Resonance Imaging

has a paramagnetic property, which in return will decrease the magnetic signal and hence shorten the T_2^* effect. This fluctuation of the fMRI signal regarding the oxyhemoglobin concentration is named the Blood Oxygen Level Dependent (BOLD) contrast.




Figure 1.3-8 – **(a) deoxygenated blood at rest - (b) oxygenated blood following a local neural activity.** The neural activity in the brain changes locally the blood flow and the oxygen level in the tissue. Source Huettel, Song & McCarthy, 2004, Functional Magnetic Resonance Imaging

fMRI data

Based on this principle, the MRI scanner records a T_2^* -weighted volume of the brain at each time step (volumic TR), each volume being acquired slicewise along the z -axis (*i.e.* the axial axis), see Figure 1.3-9. This results in a four dimensional data set (three dimensions in space + time) in which each voxel (*i.e.* a three dimensional rectangular cuboid on a regularly spaced three-dimensional grid) fluctuates over time and represents, notably, the local evolution in oxyhemoglobin concentration as described previously. Usually, the time between two consecutive volumes, namely the volumic TR (*i.e.* the slicewise TR multiplied by the number of slices), is in the range of 1-2 second (s). However, as the MRI acquisition techniques are getting faster with the recent advances in multiband pulse sequences [Feinberg

et al., 2010, Smitha et al., 2018], they permit lower TR and thus improved temporal resolution. Recent large fMRI data sets such as the Human Connectome Project (HCP) [Van Essen et al., 2013] exhibit TR lower than a second to explore higher frequencies in the BOLD signal, be less prone to physiological artifacts such as heart beat and breathing and potentially capture oscillatory neural dynamics (in the delta band) which play an important role in the coordination of large-scale brain networks [Lewis et al., 2016]. In regards to the spatial resolution in fMRI, the voxel dimension ranges from 1 to 3mm along each axis. Of course, a gain in spatial resolution may be achieved at the expense of a lower temporal resolution due to the traditional trade-off between the two. Improved spatial resolution may help finely delineate cortical activation as far as the sensitivity of detection is maintained, which in turn requires a good temporal SNR. The temporal SNR (tSNR) is thus a key property in fMRI [Triantafyllou et al., 2005], defined voxelwise as the ratio of the average signal across volumes divided by the corresponding standard deviation.




Figure 1.3-9 – **Illustration of the four dimensions in fMRI.** During an fMRI acquisition, consecutive brain volumes are collected at a predefined pace defined at the volumic TR. Source nilearn.github.io

1.3.2 Classical fMRI preprocessing

In this subsection, we describe the classical pipeline to perform the spatio-temporal preprocessing of fMRI data. Indeed, nuisance signals and artifacts may corrupt the BOLD signal during acquisition. Firstly, the brain scans are not perfectly aligned one another and thus should be spatially re-aligned. Secondly, the subject may move in the scanner during acquisition, which should be compensated for. Thirdly, all slices in a given volume are not collected simultaneously, indeed, each slice takes approximately 30 to 40ms to be acquired depending on the selected spatial encoding scheme (e.g. echo planar imaging combined with parallel imaging). This leads to a significant delay between the first and last slices that must be compensated for to permit further statistical analysis in which all voxels are supposed to be collected at the same time point (see Chapter 2). Most often, to meet the timing constraints of one volume every 1-2s., a high-speed *single shot* Echo Planar Imaging (EPI) readout is used to collect fMRI data slicewise. Consequently, the data are prone to local B_0 inhomogeneities and thus to off-resonance effects, which introduce geometric distortions and signal drop out in fMRI data, especially in voxels located at the interfaces between air and tissues (e.g. nasal cavities, ear canal). To compensate for this kind of artifacts, the most popular strategy in fMRI is called *topup* [Jenkinson et al., 2012] and consists in collecting twice the data corresponding to the same slice, first in the top left/bottom right direction in the 2D k-space plane and second in the opposite direction (bottom right/top left). By doing so, the local B_0 inhomogeneities accumulated along the readout axis (left/right) are canceled out. However, this doubles the acquisition time.

The obtained fMRI data could be then registered either to structural MRI data to perform intra-subject normalization or to a brain template to have a spatially coherent basis

to perform comparison with multiple subjects. Moreover, to maximize region overlap between subjects, fMRI data are usually convolved with a three-dimensional Gaussian, resulting in spatial smoothing of the data. Finally, the nuisance artifacts that could still be present in the data, could be estimated for future removal (confound regressors estimation).

To summarize the classical preprocessing steps are:

- *Alignment* : Adjust for movement between scans.
- *Head-motion estimation*: Correct for movement from the subject.
- *Slice-timing correction*: Correct for delays between slices.
- *Susceptibility distortion*: Correct for the presence of artifact (geometric distortions and signal loss) due to B_0 inhomogeneities.
- *Coregistration* : Overlay structural and functional images.
- *Normalisation* : Wrap images to fit to a standard template brain for statistical group analyses.
- *Smoothing* : Increase signal-to-noise ratio and maximize between-subject overlap of activations.
- *Confound estimation* : Estimate the nuisance regressors for future removal.

1.3.3 fMRI acquisition

In this subsection, we briefly describe the different types of acquisition that are currently used in fMRI to map cognitive functions and investigate ongoing fluctuations of brain activity. To this endeavour, we will introduce first *task-related fMRI* (tfMRI) acquisitions and second *resting-state fMRI* (rs-fMRI) acquisitions and describe the specificity of each imaging session.

Task fMRI acquisition

Task fMRI (tfMRI) is used when the participant is engaged in an experimental paradigm (EP) during the imaging session. This means that the subject is submitted to some specific stimuli or is asked to perform certain tasks while his brain activity is recorded at regular intervals. Resting periods are introduced between stimuli. The stimuli may be grouped in blocks of approximately 30s (block designs) or instead submitted to the subject in a sparse way (event-related designs). In the last scenario, the distance in time between consecutive stimuli determines whether the paradigm is slow (15 to 20s) or fast event-related (3s). Block designs are known to generate maximal contrast-to-noise ratio (CNR) and thus ease the detection of evoked activity. In contrast, event-related designs may be useful to estimate the haemodynamic response function (see chapter 2). Fast event-related designs are particularly appealing when mixing different types of stimulus (i.e. different experimental conditions) as they permit to increase the pace between consecutive stimuli and thus preserve a reasonable CNR compared to their slower counterpart.

The nature of the task submitted to the participant varies depending on the purpose of investigation. A simple example is the characterisation of the motor cortex, the corresponding EP for this imaging session could be the *finger tapping* in which the participant alternates

between rest period and finger tapping. Such acquisition permits to activate the motor cortex and, with an adapted estimation procedure, to determine the precise localization of the subject's brain regions involved in task performance.

Numerous tasks can be submitted to the participant during the imaging session, from listening to specific sound patterns to performing a challenging working memory task (e.g. n-back task). The first main objectives of the tfMRI are to spatially localize the subject's brain activity in response to these external stimuli. This is easily achieved using block designs. However, when we are interested in the rich fluctuations of the BOLD signal to try to uncover some neural signature in response to some specific stimuli (e.g. priming or repetition suppression effect), event-related designs are very useful.

This kind of imaging paradigm allows a better understanding of the brain function. However, as each paradigm is focused on a specific cognitive function and because the time spent in the MRI scanner is limited, the analysis of fMRI data provides a limited view on the global functioning of the brain.

Resting-state fMRI acquisition

Rs-fMRI is the type of acquisition where the participant is not engaged in an experimental paradigm but instead stays still in the scanner, either keeping his eyes open or closed. The general idea is to record the nominal and awaken brain functioning to better understand its global functional structure.

Importantly, rs-fMRI reveals functionally connected regions, see [Figure 1.3-10](#), that are organized into reproducible, well-known and large-scale functional brain resting-state networks (RSNs), for examples : default mode network (DMN), attention network, executive control network (ECN), visual network, motor network. Alterations in these RSNs are thought to precede symptoms and structural changes in neurodegeneration by several years, and might allow the identification and quantification of preclinical disease stages.

1.4 Neurovascular coupling in fMRI

In the previous section we highlight the dependence of the fMRI signal on the oxygenation level of the blood, yielding to an indirect observation of the neural activity. In this section, we focus on the mechanism between neural activation and an observable change in the BOLD signal. Then we will detail how useful the estimation of this coupling can be in a clinical context.

1.4.1 Introduction to the neurovascular coupling and its modeling

The neurovascular coupling relies on multiple actors, from the astrocyte cells to smooth muscle cells notably. The principle aspect of the response is the vasodilation of the surrounding blood vessels. Once a group of neurons fire together, the glial cells will dilate blood vessels to supply the neurons in oxygen. As explained previously, this re-oxygenation is then captured by the MRI scanner through an increased $T2^*$ decay which produces the BOLD contrast.

A major difficulty related to the BOLD signal is that this oxygenation supply happens a couple of seconds after the neural firing. If we consider a single spike as a proxy for neural

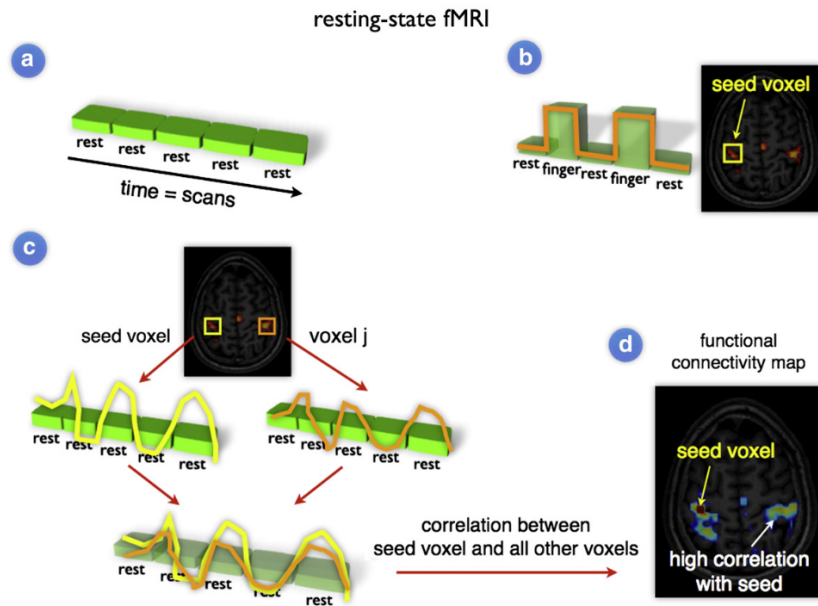


Figure 1.3-10 – **Illustration of the rs-fMRI analysis.** Identification of the voxels timely correlated to estimate the functional networks present in the brain. Source M. P. van den Heuvel and H. E. Hulshoff Pol, Exploring the brain network: A review on resting-state fMRI functional connectivity.

activity, the BOLD signal will peak in average 5s later and then return to its baseline after approximately 20 to 30s seconds.

In order to perform analysis on the BOLD data, this neurovascular coupling was model as a transfer function which takes as input the neural activation signal and output the BOLD signal. This function was described through a differential equation system and can be summarized by its Green function. The latter represents the noise-free response of the neurovascular system to a spike input, which could model a very brief neural activation, see Figure 1.4-11. This response is named the Haemodynamic Response Function (HRF). Numerous models have been proposed to encode the HRF shape in various ways but the most common is the double Gamma distribution model introduced in Friston et al. [1998a], see Equation 1.8 where $c, \alpha_1, \alpha_2, \beta_1, \beta_2$ are fixed scalars defined by the authors as follows:

$$v(t) = \frac{t^{\alpha_1-1} \beta_1^{\alpha_1} e^{-\beta_1 t}}{\Gamma(\alpha_1)} - c \frac{t^{\alpha_2-1} \beta_2^{\alpha_2} e^{-\beta_2 t}}{\Gamma(\alpha_2)} \quad (1.8)$$




Figure 1.4-11 – **Illustration of the Haemodynamic Response Function (HRF).** The yellow line represents the neural activation in time u and the red line represents the BOLD response y to this spike input, namely the HRF. Source Huettel, Song & McCarthy, 2004, Functional Magnetic Resonance Imaging.

Moreover, to obtain the response $y(t)$ from any neural activity signal denoted $u(t)$, the input is convolved with the HRF $v(t)$. Indeed, as described in Boynton et al. [1996], the neurovascular coupling exhibits linear and time invariant (LTI) properties and can then be

faithfully described by a convolution operator:

$$y(t) = \int_0^T v(\tau)u(t-\tau)d\tau = (v \star u)(t). \quad (1.9)$$

Consequently, the following properties hold:

- *Multiplicative scaling*: if the amplitude of the neural activation signal $u(t)$ is multiplied by a factor λ , then the produced BOLD signal is scaled by λ too: $y(\lambda u(t)) = (v \star (\lambda u))(t) = \lambda y(t)$.
- *Additivity*: if the neural activation signal is a sum of multiple spikes $(u_i(t))_i$, then the corresponding BOLD signal is the sum of the responses $y_i(t)$ to each signal $u_i(t)$: $y(\sum_i u_i(t)) = (v \star (\sum_i u_i))(t) = \sum_i y_i(t)$.
- *Time shift invariance*: if the neural activation signal is shifted in time from a specific delay Δt , then the associated BOLD signal $y(t)$ is shifted by Δt too: $y(u(t - \Delta t)) = (v \star u)(t - \Delta t) = y(t - \Delta t)$.

In the discrete setting, the counterpart of Equation 1.9 reads:

$$\mathbf{y} = \mathbf{v} * \mathbf{u} \quad \text{with} \quad y_n = \sum_{i=0}^{K-1} v_i u_{n-i}. \quad (1.10)$$

However, this formulation does not take any noise component into account. A common assumption is to consider the noise ϵ as an auto-regressive noise of order 1 (also denoted as AR(1): $\forall n, \epsilon_n = a\epsilon_{n-1} + \nu_n$ where a stands for the auto-regressive parameter and ν the innovation white Gaussian noise of variance σ^2), see Equation 1.11 and Figure 1.4-12.

$$\mathbf{y} = \mathbf{v} * \mathbf{u} + \epsilon \quad (1.11)$$




Figure 1.4-12 – **Illustration of BOLD signal modeling.** The BOLD signal is obtained by convolving the neural activity signal \mathbf{u} (in red) with the Haemodynamic Response Function \mathbf{v} (in blue) and by adding noise ϵ (in yellow).

1.4.2 Clinical application

In this section, we propose to discuss the opportunity to use the HRF as a biomarker as a means to investigate brain function in a healthy or a pathological condition.

The Haemodynamic Response Function as a biomarker

As described earlier, the HRF models a complex cascade of events produced notably by the glial cells. A modification that will affect those cells would induce a change in the HRF shape. Thus, the latter may provide a valuable biomarker to identify the regions that are impaired by a neurodegenerative pathology (e.g. Parkinson's disease), a neurological disorder (e.g.

epilepsy) or specific drugs (e.g. caffeine). Moreover, depending on its precise representation, the HRF could be used to identify the change occurred in the BOLD response, whether it implies a re-scaling of its magnitude or a longer delay in the vascular coupling.

The Synchronoid project

The analgesic derived from opium (opioid analgesics) are largely used nowadays with notably an ongoing crisis due to their massive usage in the USA³. A specific example of opioid is the buprenorphine. The buprenorphine is an opioid used as analgesic for post-surgical pain or neoplasm pain, it is also used to treat opioid use disorder. This molecule produces a variable analgesic response for each patient along with a tolerance effect which leads to increase the drug's dosage to obtain the same analgesic effect. Thanks to the ceiling effect of its μ -agonist activity, buprenorphine is also used to treat opioid use disorder, with greater safety and less addictive potential than methadone. Such medication induces a patient-dependent analgesic effect and various degrees of addiction. The reason for this variability has remained unknown so far. To uncover the variability sources and potential addiction factors, the Synchronoid project has been set up by the Service Hospitalier Frederic Joliot since 2018 to investigate and characterize the mechanisms of action of the opioid drugs, see Figure 1.4-13. The study was approved by the biomedical research ethics committee in decembre 2017 (Etude CEA 100- 040; EudraCT 2017-001897-41; CPP 2017-12-02-ter).




Figure 1.4-13 – **Illustration of the different steps from the injection to the analgesic effect of the buprenorphine.** The main steps from the injection of the painkiller to the analgesic effect on the patient. The Synchronoid project aims to investigate the receptor occupation step as well as the neural activity step to better characterize the effect of the buprenorphine on the neurovascular coupling in the brain along with the neural activity.

More precisely, the project aims to characterize the action of the buprenorphine spatially and temporally. Buprenorphine is classified as a partial agonist. It has a high affinity, but low efficacy at the μ -receptor where it yields a partial effect upon binding. It also, however, possesses κ -receptor antagonist activity making it useful not only as an analgesic, but also in opioid abuse deterrence, detoxification, and maintenance therapies. The objective is to form a cohort of 60 subjects (healthy volunteers with no history of substance abuse or addiction) who will receive a dose of the pain killer while being scanned in a with a fully integrated 3T PET/MR system (GE SIGNA PET/MR scanner)⁴. This is a hybrid scanner that allows simultaneous PET and MR scans for patients, see the photo Figure 1.4-14. This novel type of machines allows for the precise recording of the localization of the painkiller in the brain using the PET modality while recording the ongoing fluctuations of the BOLD signal

³<https://www.theguardian.com/us-news/opioids>

⁴At the time of this manuscript is written, the cohort is not yet fully recruited and the project will be continuing beyond this PhD thesis.

using the MRI system during resting-state acquisitions. Half of the participants, defined in a randomized controlled trial, will receive a placebo dose (control group) while the other half will receive an analgesic dose (0.2 mg intravenously injected) in order to exhibit the buprenorphine effect by contrasting the two populations and making statistical comparisons, as done in clinical trials. The group-comparison of the HRF estimates will permit to uncover in the brain which regions show analgesic-related changes. The general idea here is to use the HRF as a biomarker to investigate the analgesic effect of opioid and open new avenues of research to better control its effects on patients.



Figure 1.4-14 – **PET MR scanner**. Photo of the PET MR scanner presents at the Service Hospitalier Frederic Joliot involved in the Synchronopoid protocol.

1.5 Chapter conclusion

We have introduced the background in neuroimaging, through the description of the brain structure and function and the various neuroimaging techniques used to investigate them. We have paid attention to the fMRI acquisition setup and described the neurovascular coupling involved in the origin of the BOLD signal. We have presented the main assumptions underlying this coupling, basically summarized by the impulse response of a linear and time invariant system also called the Haemodynamic Response Function (HRF). Additionally, we have emphasized how its estimation is relevant to investigate brain activity in a healthy or disease condition. In the following, we will describe how to estimate the HRF from fMRI data and we will introduce novel and efficient methods to that aim.

* * *
* *
*

How to segregate the vascular and neuronal components in fMRI?

Chapter Outline

2.1	Introduction to the General Linear Model (GLM)	29
2.2	HRF estimation with fixed neural activity	31
2.2.1	HRF modelling	32
2.2.2	Fitting the HRF model to the observed fMRI data	33
2.3	Neural activity estimation with fixed HRF	34
2.4	Joint estimation of the neural activity and the neurovascular coupling from BOLD signal	37
2.5	Conclusion	38

THIS chapter introduces state-of-the-art approaches to disentangle the neurovascular coupling from fMRI data. First, we will detail the General Linear Model (GLM) as it is the most common method used to identify the brain activity that correlates to a specific task during the imaging session. Then, we will present the main approaches that have been proposed in the literature to estimate the HRF from fMRI data collected during task performance, i.e. along the course of an experimental paradigm. In the third section, we will introduce methods that estimate the neural activation signals from the observed signal with a fixed HRF. Finally, we will introduce the most recent methods that allow for a *semi-blind deconvolution* of neural activity from an unknown HRF in a paradigm-free setting, meaning that the knowledge of the experimental paradigm is no longer required to conduct this kind of analysis and hence such approaches may apply to both task-related and resting-state fMRI data.

2.1 Introduction to the General Linear Model (GLM)

In this section, we introduce the General Linear Model (GLM) approach. This model describes the observed task-related fMRI (tfMRI) data with a linear combination of predefined regressors (i.e. temporal atoms) based of the different experimental conditions performed by the participant during the imaging session. The original GLM formulation does not

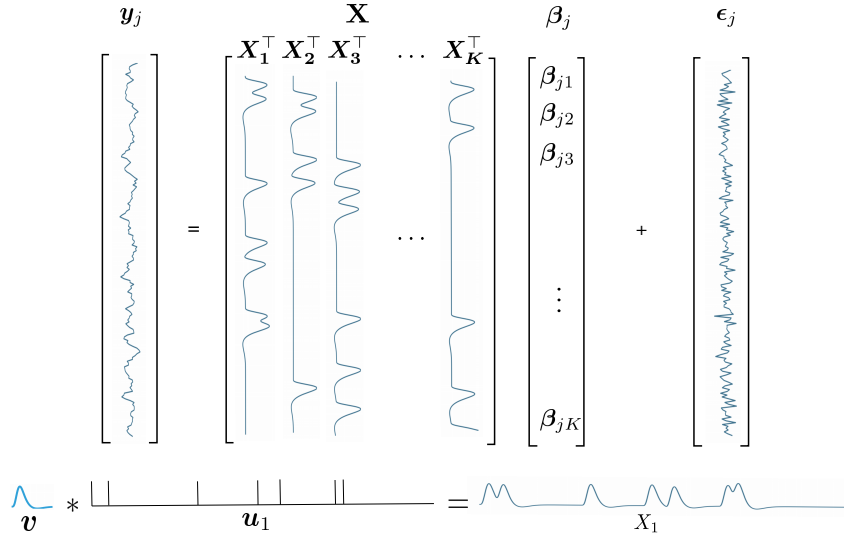


Figure 2.1-1 – **General Linear Model illustration.** The observed BOLD signal \mathbf{y}_j is described as the linear combination such as $\mathbf{X}\boldsymbol{\beta}_j$ with the addition of an auto-correlated noise [Woolrich et al., 2001] denoted $\boldsymbol{\epsilon}_j$, which is usually model as an autoregressive process of order 1. Note that the components of \mathbf{X} follow the LTI model such as $\mathbf{X}_k = \mathbf{v} * \mathbf{u}_k$ with \mathbf{u}_k being the temporal signature of each conditions.

allow for the estimation of either the HRF or the underlying neural activity signal, however as it is the most common method used to perform the analysis of tfMRI, we describe it briefly.

The GLM proposes a massively univariate description of the tfMRI data $\mathbf{Y} \in \mathbb{R}^{T \times P}$. The model makes the hypothesis that voxel's temporal signal can be recomposed as a linear combination of K predefined temporal components of length T . Each component is described by the temporal signature of a specific task (condition) of the experimental paradigm (EP). This signature is convolved with a canonical HRF \mathbf{v} that has a fixed and constant shape for the whole brain to better explained the observed signal. All these temporal signatures form the *design matrix* $\mathbf{X} \in \mathbb{R}^{T \times K}$ of the experiment, such as $\mathbf{X} = (\mathbf{X}_k)_{k=1}^K$ with $\forall k \in [1..K]$ $\mathbf{X}_k = \mathbf{v} * \mathbf{u}_k$ with $\mathbf{u}_k \in \mathbb{R}^{\tilde{T}}$ describing the temporal onsets of the k^{th} condition. The GLM model is summarized in Equation 2.2 (see also Figure 2.1-1).

For one voxel:

$$\mathbf{y}_j = \mathbf{X}\boldsymbol{\beta}_j + \boldsymbol{\epsilon}_j, \quad (2.1)$$

with $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{V})$

For P voxels:

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{E}, \quad (2.2)$$

with $\mathbf{Y} = (\mathbf{y}_j)_{j=1}^P$, $\boldsymbol{\beta} = (\boldsymbol{\beta}_j)_{j=1}^P$ and $\mathbf{E} = (\boldsymbol{\epsilon}_j)_{j=1}^P$

such that $\forall j \in \{1..P\}$ $\boldsymbol{\epsilon}_j \sim \mathcal{N}(0, \mathbf{V})$.

Most of approaches propose to minimize a simple quadratic loss as a cost-function to fit the model parameters to the observed tfMRI data. However, the quadratic loss is well suited only with the assumption that the noise is Gaussian. As shown by Woolrich et al. [2001], the noise in fMRI is best modelled by a first-order autoregressive process. In

order to remove the correlation in the noise, the most common solution is to apply a prior whitening on the observed data \mathbf{Y} . The correlated noise is given by $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{V})$ where $\mathbf{V} \in \mathbb{R}^{T \times T}$ is the symmetric correlation matrix and σ the standard deviation of the noise. The Cholesky decomposition can be used to find a matrix \mathbf{Q} such that $\mathbf{V} = (\mathbf{Q}^\top \mathbf{Q})^{-1}$. One approach to estimate \mathbf{V} is to fit a first GLM on the observed data and consider the residual $\hat{\mathbf{E}} = \mathbf{Y} - \mathbf{X}\boldsymbol{\beta}$. Then, we fit, on $\hat{\mathbf{E}}$, a weighted sum of predefined correlation matrices to this residual to estimate \mathbf{V} [Bollmann et al., 2018]. Thus, the GLM decomposition is used on the prewhitening problem as :

$$\begin{aligned} \mathbf{QY} &= \mathbf{QX}\boldsymbol{\beta} + \mathbf{Q}\boldsymbol{\epsilon}, \\ \text{with } \mathbf{Q}\boldsymbol{\epsilon} &\sim \mathcal{N}(0, \sigma^2 \mathbf{I}). \end{aligned} \tag{2.3}$$

As $\boldsymbol{\beta}$ belongs to $\mathbb{R}^{K \times P}$, each row $\boldsymbol{\beta}_k \in \mathbb{R}^P$ is a spatial map that encodes which voxels has a signal that is mostly correlated with the k^{th} experimental condition. Indeed, the computation of a T-statistics based on the vector $\boldsymbol{\beta}_k$ allows us to identify the most significant entries which in return permit to localize the most activated voxels in response to the stimulation (e.g. listen to sounds, music or read sentences) associated with the k^{th} condition.

The GLM is impeded by several limitations. First, the haemodynamic response is crudely modeled for the whole brain. Indeed, despite enhancements proposed by more elaborated formulations that allow to adapt to some degree the HRF, the GLM assumes a fixed and common neurovascular coupling for the whole brain. This assumption is known to be false and thus makes the standard GLM approach too crude to accurately model fMRI data. Moreover, the LTI model used in the GLM supposes that the conditions $((\mathbf{u}_k)_{k=1}^K)$ are used as a proxy for the neural activity. However, this hypothesis does not take into account the possible delays in the participant's responses during task performance. Second, the simplistic block or Dirac modelling for a stimulus or an action can depart from the true neural activation signal as this one should be continuous, in particular in time. Overall, the oversimplistic description of the neural activity impedes the precise modeling of the BOLD fMRI data.

To cope with these issues, numerous contributions proposed to improve the estimation of the HRF to better explain the observed data. We will detail these contributions in the next section.

2.2 HRF estimation with fixed neural activity

In this section, we detail the most common modelling for the HRF and describe how to estimate the corresponding model parameters from the observed fMRI data. All the contributions discussed here [Friston et al., 1998a, Aguirre et al., 1998, Friston et al., 1999b, Dale, 1999, Glover, 1999, Goutte et al., 2000, Ciuciu et al., 2003, Casanova et al., 2008, Khalidov and Unser, 2006, Lindquist and Wager, 2007, Vakorin et al., 2007, Lindquist et al., 2009, Khalidov et al., 2011, Pedregosa et al., 2015, Zhang et al., 2013] were proposed in task-related fMRI context where an Experimental Paradigm (EP) is available as a proxy to the neural activity signal, following the LTI model described in the previous section.

2.2.1 HRF modelling

Most of those contributions to estimate the HRF rely on a few HRF models. With the same HRF model, the estimation process can vary by making the regularization of the HRF evolve to better capture the true haemodynamic response in a given subject. The main challenge with HRF modeling is to find an efficient trade-off between a flexible model that is able to precisely capture the true haemodynamic response in each brain area (reducing the estimation bias) and a reliable model that limits the overfitting by reducing the number of degrees of freedom (diminishing the estimation variance). In this subsection, we propose to list and detail the main HRF models available in the literature.

The Finite Impulse Response (FIR) model. The FIR model [Dale, 1999, Glover, 1999] does not assume any particular shape for the HRF which makes it very flexible but prone to data overfitting in the presence of noise. Regularization has thus been introduced to constrain the overall HRF shape in FIR models, for instance by penalizing the second-order derivative to end up with physiologically plausible smooth HRF estimates [Goutte et al., 2000, Ciuciu et al., 2003, Casanova et al., 2008].

The linear combination of predefined atoms. Alternatively, to limit the number of parameters in the model, the HRF $\mathbf{v} \in \mathbb{R}^l$ has been set as a linear combination of predefined atoms [Aguirre et al., 1998, Friston et al., 1998a, 1999b], such as $\mathbf{v} = \boldsymbol{\alpha}^\top \mathbf{B}$, with $\boldsymbol{\alpha} \in \mathbb{R}^R$ being the coefficients of decomposition and $\mathbf{B} = (\mathbf{b}_r)_{r=1}^R \in \mathbb{R}^{R \times l}$ the predefined atoms. The atoms can then be set in order to span all plausible HRFs and hence better capture the neurovascular coupling.

Two widespread sets of predefined atoms are:

- *Physiologically informed pattern set:* Friston et al. [1998a] propose to consider the canonical HRF (the double Gamma distribution model) and its derivatives in time and with respect to the dispersion parameter. The two last aforementioned atoms allows to express a potentially delayed (i.e. slower or shorter) and sustained (wider HRF peak) neurovascular coupling that better fits the observed data. Note that, based on this double Gamma distribution function some models are limited to the first two atoms, nevertheless the number of atoms are usually set to three. Note that, each of these atoms can't be expressed as a linear combination of the others, however this set does not strictly define an orthogonal basis.
- *B-splines:* (or basis spline) is a basis defined from spline functions. A spline is a function that has minimal support with respect to a given degree, smoothness, and domain partition. In these approaches [Vakorin et al., 2007, Zhang et al., 2013], the HRF is defined as a linear combination of splines. In Vakorin et al. [2007] the authors propose to penalize the ℓ_2 -norm of the first-order derivative of the splines to obtain a smooth HRF estimate.

Other alternatives exist in the literature, such as Khalidov and Unser [2006] and Khalidov et al. [2011] who use wavelets designed to encode the HRF and estimate the neural activity. This example is detailed in the next subsection as a means to estimate the neural activation signal.

The logit model. To limit model overfitting, parametric HRF models have been proposed as an alternative to complex regularization schemes. Interestingly, the authors of [Lindquist and Wager \[2007\]](#) propose the *inverse logit transform* and demonstrate its efficiency over multiple fMRI data sets [[Lindquist et al., 2009](#)]. The inverse logit function L is defined by $\forall x \in \mathbb{R} \quad L(x) = \frac{1}{1+e^{-x}}$. The proposed model summarizes the HRF with only seven parameters (two parameters are fixed by the constraint of a finite support):

$$v(t) = \alpha_1 L\left(\frac{t-T_1}{D_1}\right) + \alpha_2 L\left(\frac{t-T_2}{D_2}\right) + \alpha_3 L\left(\frac{t-T_3}{D_3}\right) \quad (2.4)$$

2.2.2 Fitting the HRF model to the observed fMRI data

In this sub-section, we will detail how to fit the parameters of the HRF model to the observed BOLD data.

As a first approach, some contributions, as in [Lindquist et al. \[2009\]](#), proposes to consider the neural activity per voxel as known. This hypothesis is plausible if the most correlated voxels to a given condition, are identified in the brain and used to estimate the HRF.

However, in the general case, the neural activity per voxel is unknown. One way to cope with this, is to fit a GLM that embedded the HRF modelling. Considering the fMRI data pre-whiten, we fit the GLM to the observed fMRI data by minimizing a simple quadratic data-fidelity term. Although, estimating both the HRF's parameters and the coefficient β leads to an ill-posed problem. To manage this problem, the authors, in [Pedregosa et al. \[2015\]](#), propose to introduce a low-rank hypothesis to the GLM, by minimizing the cost-function defined in :

$$\begin{aligned} \arg \min_{\alpha_j, \beta_j, c_j} \quad & \frac{1}{2} \|\mathbf{y}_j - \mathbf{X}_B \text{vec}(\alpha_j \beta_j^\top) - \mathbf{X}_c c_j\|_2^2 \\ \text{subject to} \quad & \|\alpha_j^\top \mathbf{B}\|_\infty = 1 \quad \text{and} \quad \langle \alpha_j, \alpha_{\text{ref}} \rangle > 0 \end{aligned} \quad (2.5)$$

Here, $\mathbf{y}_j \in \mathbb{R}^T$ refers to the BOLD signal for *one voxel*. The nuisance phenomenon are modeled with $\mathbf{X}_c c_j$ where \mathbf{X}_c refers to predefined nuisance regressors with their associated coefficients c_j . The signal of interest is captured by the term $\mathbf{X}_B \text{vec}(\alpha_j \beta_j^\top)$. The design matrix $\mathbf{X}_B \in \mathbb{R}^{T \times RK}$ is defined by successively stacking the regressors obtained by convolving each of the basis elements with the stimulus onsets of each condition – similarly to the usual *design matrix* used in the GLM models. α_j refers to the parameters of the HRF, for one voxel, modeled *w.r.t* a set of predefined atoms $\mathbf{B} = (\mathbf{b}_r)_{r=1}^R \in \mathbb{R}^{R \times l}$ such as $\mathbf{v}_j = \alpha_j^\top \mathbf{B}$. These parameters α_j are constrained with the ℓ_∞ -norm to avoid amplitude indeterminacy and constrained to have a positive scalar product with fixed user-defined reference HRF parameters α_{ref} to prevent too large deviations from expected HRF shapes. Finally, the β_j denotes the regressor coefficients for the model.

In this work, the authors rely on the assumption that a linear combination of predefined atoms \mathbf{B} for the HRF, such as $\mathbf{v} = \alpha_j^\top \mathbf{B}$ and enforce a rank-1 constraint on the coefficient α_j for the considered voxel such as $\text{vec}(\alpha_j \beta_j^\top) = [\alpha_j \beta_{j1}, \alpha_j \beta_{j2}, \dots, \alpha_j \beta_{jK}] \in \mathbb{R}^{RK \times 1}$. By doing so, the authors provide a method to estimate one HRF per voxel while identifying the most correlated voxel to each EP in the design matrix.

In this section, we have introduced most common models for the HRF and how their parameters are estimated from the observed BOLD data. In the following section we will see how to estimate the neural activity with a fixed and constant neurovascular coupling for the whole brain.

2.3 Neural activity estimation with fixed HRF

In this section, we give an overview of the different methods proposed to estimate the neural activity signals in the brain from fMRI data. Importantly, we will detail the fundamental work [Karahanoğlu et al., 2013] that proposed a block signal modelling for the neural activity. This hypothesis is critical for our contributions presented in the next chapters.

We introduce the Toeplitz matrix such that $\forall \mathbf{u} \in \mathbb{R}^{T-L+1}$ $\mathbf{H}_v \mathbf{u} = \mathbf{v} * \mathbf{u}$. One can estimate the neural activity signal \mathbf{u} by minimizing a simple quadratic data-fidelity term as :

$$\arg \min_{\mathbf{u} \in \mathbb{R}^{T-L+1}} \frac{1}{2} \|\mathbf{Y} - \mathbf{H}_v \mathbf{u}\|_2^2 \quad (2.6)$$

Since \mathbf{H}_v is not invertible, Equation 2.6 admits multiple minimizer. In order to reduce the variance of estimation, one can introduce modelling hypothesis on \mathbf{u} to make the estimation problem better conditioned. The neural activity signal models the activity of approximately 10^5 neurons over time during a TR. If we normalize the signal to 1, we can interpret its value as the proportion of firing neurons among the population belonging to a voxel. A reasonable hypothesis would be to have a smooth evolution of this percentage over time. However, in a BOLD deconvolution context, a flexible smoothness regularization is not sufficient to achieve robust estimation of the neural activity as the overall problem stays ill-posed.

Agnostic neural activation signals modelling. In Glover [1999], the authors propose to obtain a proxy for the HRF by taking the observed BOLD response of an event-related condition. Indeed, as the HRF is the response of the neurovascular coupling to a Dirac input (impulse event), a direct observation of the BOLD signal after an event will give a noisy observation of the HRF. With this proxy, they propose to estimate the neural activity signals, denoted $u(t)$ using the Wiener filter :

$$u(t) = \mathcal{F}^{-1} \left(\frac{H^*(\omega)M(\omega)}{|H(\omega)|^2 + |N(\omega)|^2} \right) \quad (2.7)$$

with $*$ the complex conjugate, $H(\omega)$ the Fourier transform of the HRF, $M(\omega)$ the Fourier transform of the observed BOLD data and $N(\omega)$ the Fourier transform of the noise in the observed data.

This contribution is of special interest since it does not impose any model on the neural activity signal. However, the described estimation process can only be applied to event-related task fMRI data for which an estimate of the HRF can be derived beforehand. Moreover, to perform Wiener filtering, one needs to estimate the term $|N(\omega)|^2$. These two limitations reduce the applicability of this approach.

Sparse hypothesis for the neural activation signals. An attempt to improve regularization of the neural activity model is to consider that the signal is sparse. In other words, we consider that most of the coefficients of the 1D vector modelling the neural signal are zeros.

Doing so, the neural activity is summarized into a few time-locked temporal activations named *events*.

In Khalidov and Unser [2006], Khalidov et al. [2011] the authors propose to estimate the neural activation signal by computing it as an output of the *Activelet* transform of the observed BOLD data. These wavelets form a dictionary in which the canonical HRF is sparsely represented. The neural activity signal is estimated by enforcing sparsity to find the most parsimonious representation for the observed BOLD signal through the minimization of a LASSO problem.

In Caballero-Gaudes et al. [2011, 2012, 2013] the authors introduce their “paradigm-free mapping” with which they estimate the neural activation signal by regularizing it notably with the ℓ_1 -norm in order to recover a parsimonious activation.

Numerous others approaches propose to estimate the neural activity from the BOLD signal under this sparse hypothesis by introducing a Bayesian framework [Ciuciu et al., 2010, Chaari et al., 2012, Bakhous et al., 2012]. As the remaining of the manuscript focuses on variational approaches, we do not detail those approaches here. Moreover, this sparse hypothesis is known to poorly capture the true neural activity. Representing neural activity as a piecewise constant or block signal has been introduced later on in Karahanoğlu et al. [2013] to improve its representation

Piecewise constant or block signal hypothesis for the neural activation signals.

To the best of our knowledge, the first contribution that has introduced this block signal hypothesis in the estimation of the neural activity signal is Karahanoğlu et al. [2013].

- *Deconvolution with the Total Variation regularization.* In Karahanoğlu et al. [2013] framework, the authors impose sparsity of the first-order derivative of the estimated neural activity signal which estimates a piecewise-constant signal, as described in Equation 2.8 and illustrated in Figure 2.3-2. This regularization involves a Total Variation (TV) term which will be detailed later on in subsection 3.1.2. Note that these authors also impose a spatial smoothness within predefined regions of interest concerning the denoised BOLD signal.

$$\begin{aligned} & \arg \min_{\mathbf{X} \in \mathbb{R}^{P \times T}} \frac{1}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + R_T(\mathbf{X}) + R_S(\mathbf{X}) \quad (2.8) \\ & \text{with } R_T(\mathbf{X}) = \sum_{j=1}^P \lambda_{1j} \sum_{t=1}^T |(\Delta_L(\mathbf{X}_{p,:}))_{:,t}| \\ & \text{and with } R_S(\mathbf{X}) = \sum_{t=1}^T \lambda_{2t} \|\Delta_{\text{Lap}}(\mathbf{X}_{:,t})\|_{2,1} \end{aligned}$$

with, as previously, P the number of voxels, T the number of scans, $\mathbf{Y} \in \mathbb{R}^{P \times T}$ the observed BOLD signal and $\mathbf{X} \in \mathbb{R}^{P \times T}$ the denoised BOLD signal. $\Delta_L = \Delta_D \circ \Delta_h$ is the operator that transforms the BOLD signal into the first order derivative of the neural activity signal, by first applying the HRF deconvolution operator and then the discrete differential operator, and finally Δ_{Lap} is the Laplacian operator in 3D.

The considered optimization problem is convex in \mathbf{X} and minimized with an alternating optimization with a Generalized Forward Backward algorithm [Raguet et al., 2013].

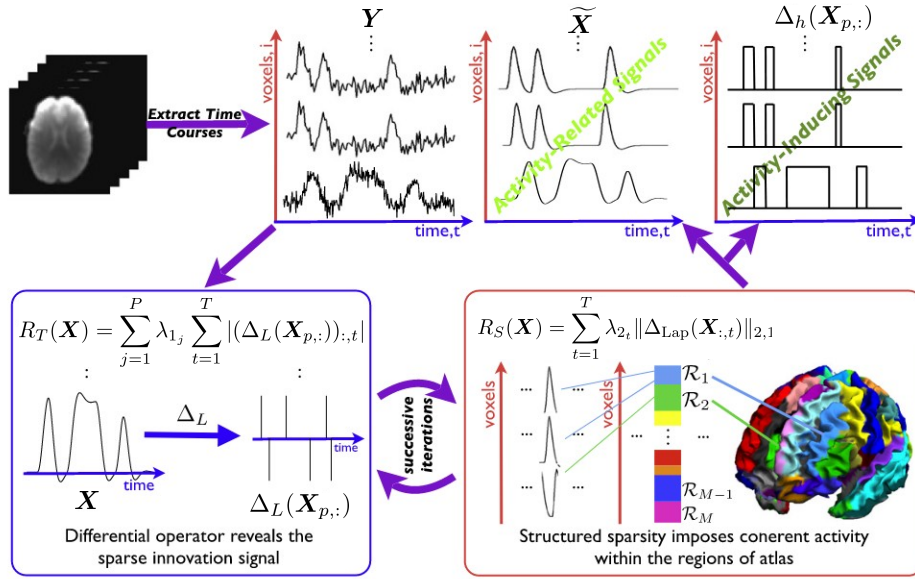


Figure 2.3-2 – Scheme of the alternated minimization procedure described in [Karahanoğlu et al. \[2013\]](#) to estimate the neural activity $\Delta_h(\mathbf{X}_{p,:})$ from the observed BOLD data \mathbf{Y} . The authors propose to estimate it by introducing notably the block signal hypothesis on the neural activity signal $\Delta_h(\mathbf{X}_{p,:})$.

The block signal hypothesis is particularly relevant for task fMRI data collected along an experimental paradigm with a block design structure as the neural activity signals may follow the same pattern. However, this model makes sense as well in the context of resting-state fMRI data as a means to regularize the spontaneous neural activity.

- *Deconvolution with the LASSO.* In [Costantini et al. \[2018\]](#) the authors propose to estimate the neural activation signal by fitting a linear model with a quadratic term and a ℓ_1 -norm regularization term. They minimize the convex cost-function defined by:

$$\arg \min_{z \in \mathbb{R}^T} \frac{1}{2P} \|\mathbf{Y} - \mathbf{H}_v \mathbf{L} z\|_2^2 + \lambda \|z\|_1 \quad (2.9)$$

with P the number of voxels, \mathbf{H}_v the Toeplitz matrix associated with the predefined and fixed HRF, and \mathbf{L} the discrete integration operator. They propose to minimize [Equation 2.9](#) with the Least Angle Regression (LARS) algorithm for fast estimation.

While this method might appear as modeling the neural activity as a sparse signal, it is in fact the synthesis formulation for a block model. Indeed, as z is sparse, the recovered signal $\mathbf{L}z$ obtained with the integration operator \mathbf{L} has a piece-wise constant structure. While this formulation leads to a simpler optimization procedure, it is usually slower as we show in [chapter 4](#).

In this section, we have described the most common neural activity signal modeling along with their estimation method from the observed BOLD signal. However, as these approaches rely on a constant and fixed HRF shape to deconvolve the observed signal, they also introduce a bias *w.r.t* the neurovascular coupling. In the following section, we will present recent contributions that aim to estimate both the HRF and the neural activity in an attempt to remove this bias.

2.4 Joint estimation of the neural activity and the neurovascular coupling from BOLD signal

A natural question would be to know if one can estimate these two components jointly from fMRI data. Since we aim at estimating the neural activity without an explicit knowledge of the HRF, one may refer to this problem as the *blind deconvolution*. In our case, as we constrain the structure of the HRF by using a limited set of parameters, we refer to this problem as the *semi-blind deconvolution*. We next describe some interesting attempts to solve this problem.

As previously introduced, the two main assumptions that dominate the neural activity signal modeling are the sparse signal and the block signal hypothesis. We propose to detail recent contributions in both case in the perspective of blind deconvolution.

Sparse hypothesis for the neural activation signals. In Wu et al. [2013] the authors propose to estimate the most prominent peaks in the observed BOLD signal by retaining only the entries that exceed a certain threshold. With this summary of the fMRI data, they estimate the haemodynamic delay. Then, they fit the HRF parameters (the response height, the Full Width at Half Max (FWHM) of the HRF and its Time to Peak (TP)) to match this delay and in return they deconvolve the observed BOLD signal with this estimated HRF, see Figure 2.4-3 for illustration. Note that this different steps are not repeated, hence, the HRF is not corrected from the estimated neural activity and inversely.

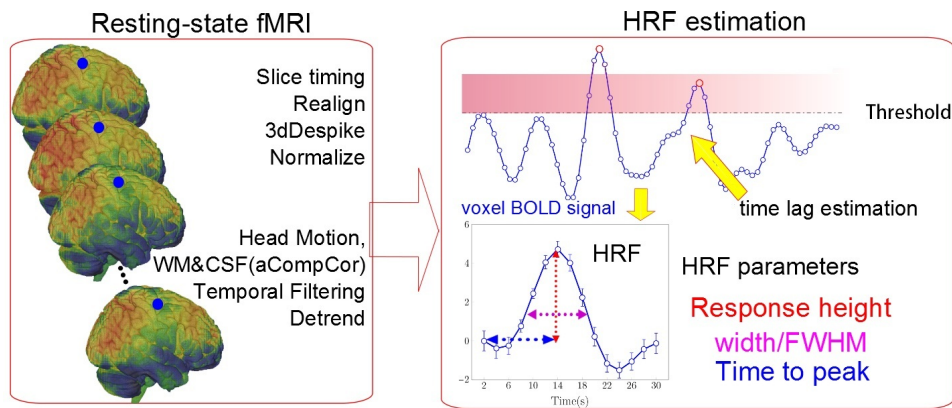


Figure 2.4-3 – Estimation process proposed in Wu et al. [2013]. The BOLD signal is summarized with its most prominent peaks on which the haemodynamic delay is estimated. An HRF is then fitted to correspond to this estimated delay which in return allows to deconvolve the observed BOLD signal.

However, as described in the previous section, the sparse hypothesis poorly captures the neural activity. We once again turn to the block signal hypothesis which allows for a gain in flexibility and thus in estimation performance.

Block signal hypothesis for the neural activation signals. In Farouj et al. [2019] the authors propose to minimize Equation 2.10 to estimate both the HRF coefficients defined as a linear combination of predefined atoms and estimate the neural activation signal similarly to Karahanoğlu et al. [2013]. They consider the cost-function defined by :

$$\arg \min_{\alpha \in \mathbb{R}^2, \mathbf{x}_j \in \mathbb{R}^T} \frac{1}{2} \|\mathbf{y}_j - \mathbf{x}_j\|_F^2 + \lambda_1 \|D\mathbf{H}_\alpha^{-1} \mathbf{x}_j\|_1 + \lambda_2 \|\alpha - \alpha_{\text{ref}}\|_2^2 \quad (2.10)$$

with $\mathbf{y}_j \in \mathbb{R}^T$ the observed BOLD data and $\mathbf{x}_j \in \mathbb{R}^T$ the estimated denoised BOLD data and λ_1 and λ_2 the regularization parameter. The parameters of the model are estimated by minimizing Equation 2.10, using a the generalized forward–backward splitting algorithm, as proposed in Karahanoglu et al. [2013].




Figure 2.4-4 – Semi-blind deconvolution proposed in Farouj et al. [2019]. The HRF modelling is based on the linear combination of predefined atoms and the neural activity model corresponds to the hypothesis of a block signal.

Point process analysis. Interestingly, in Tagliazucchi et al. [2010], the authors proposes . This characterization of the inter-regions interactions leads to recover the HRF for the target region and allows the deconvolution of the observed BOLD with the estimation neurovascular coupling for the region. The noteworthy aspect of this method, is that

2.5 Conclusion

In this chapter, we have progressively introduced methods that estimate (i) the vascular coupling in the context of task fMRI experiments for a given representation of neural activity through the stimulus onsets, then (ii) inversely and finally (iii) both components in a *semi-blind deconvolution* setting. The main limitations of these works have been outlined. More specifically, they proceed voxelwise, i.e. in an univariate manner and do not take the spatial structure into account, whether this structure concerns the vascular coupling or the spread of neural activity throughout the brain. Including this kind of structure in new methods could be a significant step forward to capture both components in a more robust way and make these novel approaches less prone to data overfitting. Extending the methods presented here the better account for the multivariate nature of the signal is one of the core contribution of this thesis, that we describe in chapter 5 and chapter 6.

* * *
* *
*

TV problems minimization with iterative algorithms

Chapter Outline

3.1	Introduction to regularization in optimization	39
3.1.1	General introduction to regularization	39
3.1.2	Introduction to the TV regularization	40
3.2	Solving a TV regularized problem	42
3.2.1	Analysis formulation	42
3.2.2	Synthesis formulation	46
3.3	Performance comparison	47
3.3.1	Material and methods	47
3.3.2	Results	48
3.4	Chapter conclusion	50

IN the previous chapter, we have introduced the Total Variation regularization (TV-reg) as a means to regularize the neural activation signals and make the global estimation better conditioned. The purpose of this chapter is to introduce the TV-reg in 1-dimension (1D) in detail. We define and explore the specific characteristics of this regularization, explain how to minimize a TV regularized problem by introducing several algorithms to that aim and compare their performances experimentally. A theoretical comparison of the different formulations of the problem is given in [chapter 4](#).

3.1 Introduction to regularization in optimization

3.1.1 General introduction to regularization

In machine learning or signal processing ill-posed problems often appear such as image reconstruction, source localization or classification. In order to make the estimation feasible, one can introduce a prior to promote solutions that satisfies certain properties. A usual solution in the Bayesian setting is to provide a specific distribution regarding the parameters to estimate. In a variational context, a solution is to add possibly multiple regularization terms




Figure 3.1-1 – **Cost-function evolution *w.r.t* the regularization parameter λ** : In this simplistic example, we display the cost function level set with different level of regularization $\lambda \in \{0.0, 1.0, 10.0, 1000.0\}$. We can notice that the minimum set evolve from an hyperplan of dimension 1 (a straight line) when $\lambda = 0$ to a singleton for $\lambda > 0$.

in the associated optimization problem. Depending on the nature of the introduced regularization, one can add information on the structure of the estimate, its norm, its smoothness, etc.

To give a simple intuition on how additional terms in a cost-function can improve the conditioning of the problem, let's consider $\mathbf{A} \in \mathbb{R}^{m \times n}$ with $m \leq n$ and $\mathbf{y} \in \mathbb{R}^m$. If we try to find $\mathbf{u} \in \mathbb{R}^n$ such that $\mathbf{y} \approx \mathbf{A}\mathbf{u}$, we might consider the following cost-function F defined by $L(\cdot) = \frac{1}{2} \|\mathbf{y} - \mathbf{A} \cdot\|_2^2$. However, with that simple loss function, we can't select a unique solution. Indeed, for any minimizer \mathbf{u} and for $\mathbf{u}_0 \in \text{Ker}[\mathbf{A}] \setminus \{0\}$, we have $L(\mathbf{u}) = L(\mathbf{u} + \mathbf{u}_0)$. The problem is said to be ill-posed as there are infinitely many equivalent solutions. By adding a regularization term, for example here the ℓ_2 -norm, such that $F(\cdot) = L(\cdot) + \lambda \|\cdot\|_2^2$, we can distinguish these equivalent solutions. Indeed, while the data-fidelity terms are equals for all $\mathbf{u} + \mathbf{u}_0$, if \mathbf{u} is in $\text{Ker}[\mathbf{A}]^\perp$, then \mathbf{u} has the smallest ℓ_2 norm and will be promoted among all the other equivalent solutions. Thus, the regularization allows to select among all solutions with equivalent data-fitting term a solution that verifies a certain property – here a small ℓ_2 norm. The unicity of the solution is guaranteed here since L is strongly convex. This is illustrated on a simplistic example with $\mathbf{y} \in \mathbb{R}^2$ such as $\mathbf{y} = [4 \ 0]^\top$ and $\mathbf{A} \in \mathbb{R}^{2 \times 2}$ with $\text{rank}(\mathbf{A}) = 1$ such as:

$$\mathbf{A} = \begin{bmatrix} 2 & 0 \\ 0 & 0 \end{bmatrix}$$

Figure 3.1-1 shows that, if we set $\lambda = 0$, the set of minimizers of the cost-function F is the one dimensional vector space defined by $\Theta = \{\mathbf{u} \in \mathbb{R}^2 \text{ such as } \mathbf{u} = [2 \ a]^\top \text{ with } a \in \mathbb{R}\}$. However, whenever $\lambda > 0$, the solution set reduces to a singleton. Hence, the estimation problem associated with F is better conditioned. A classical example of regularization is the ℓ_2 -norm, ℓ_1 -norm or the elastic-net [Zou et al., 2004] norm which combines both the previous ones.

3.1.2 Introduction to the TV regularization

The TV-reg promotes piece-wise constant estimates by penalizing the ℓ_1 -norm of the first order derivative of the estimated signal. If we consider a signal $\mathbf{u} \in \mathbb{R}^n$, the associated TV regularization is defined by: $\|\mathbf{u}\|_{TV} = \|\mathbf{D}\mathbf{u}\|_1 = \sum_{i=2}^n |u_i - u_{i-1}|$ with \mathbf{D} being the discrete differential operator such as $\mathbf{D} \in \mathbb{R}^{(n-1) \times n}$ and defined by:

$$D = \begin{bmatrix} -1 & 1 & 0 & \dots & 0 \\ 0 & -1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix}$$

As the ℓ_1 -norm promotes sparsity on the first order derivative, namely $\mathbf{z} = D\mathbf{u}$, this pushes \mathbf{z} to be a Dirac signal and thus \mathbf{u} to be a piece-wise constant signal. Figure 3.1-2 displays two examples where we generate sparse signals $D\mathbf{u}_1$ and $D\mathbf{u}_2$ and integrate them to obtain the piece-wise constant signals \mathbf{u}_1 and \mathbf{u}_2 .




Figure 3.1-2 – Examples of two piece-wise constant signals: \mathbf{u}_1 , \mathbf{u}_2 , along with their corresponding Dirac signals $D\mathbf{u}_1$, $D\mathbf{u}_2$. The sparsity of $D\mathbf{u}_1$ is higher than the one of $D\mathbf{u}_2$, i.e. $\|D\mathbf{u}_1\|_0 \leq \|D\mathbf{u}_2\|_0$

To illustrate the effect of TV-reg, we provide a simple 1D illustration. We consider the BOLD signal deconvolution problem for a synthetic signal Figure 3.1-3. We generate a time series \mathbf{z} with a support $S(\mathbf{z}) = 6$ which we integrate to obtain a block signal \mathbf{u} . This signal is then convolved with the canonical HRF [Friston et al., 1998a] and we add a Gaussian noise. We estimate the synthetic source signal $\hat{\mathbf{u}}$ by minimizing the cost function $P(\mathbf{u}) = \frac{1}{2}\|\mathbf{y} - \mathbf{u} * \mathbf{v}\|_2^2 + \lambda\|\mathbf{u}\|_{TV}$, with λ set manually. In Figure 3.1-3, we can see that we recover most of the variations of the original signal as the minimization of P deconvolves and denoises the observed signal.

Finally, we can mention that the TV regularization is a useful tool for a large number of applications. For 1D signals, we already mentioned that it has been used in neuro-imaging for deconvolution problems [Karahanoğlu et al., 2013]. Recent work also used TV-regularisation as a detrending step to preprocess oculo-motor recordings [Lalanne et al., 2020]. Most application of TV-reg are for 2D image processing problems such as denoising [Rudin et al., 1992, Chambolle, 2004, Darbon and Sigelle, 2006], medical image reconstruction [Tian et al., 2011] along with many of other applications [Rodríguez, 2013]. In this chapter, we will consider only the 1D case to focus on the BOLD deconvolution model that we introduced previously.




Figure 3.1-3 – Synthetic BOLD deconvolution example with the TV regularization: We display the estimation of the source signal in orange from the observed signal in blue which was obtained by convolving the original source signal with the canonical HRF and by adding a Gaussian noise.

In the previous chapter, we consider the cost-function $F(\cdot) = \frac{1}{2}\|\mathbf{v} * \cdot - \mathbf{y}\|_2^2 + \|\mathbf{D} \cdot\|_1$. By introducing the Toeplitz matrix associated to \mathbf{v} , we rewrite F such that $F(\cdot) = \frac{1}{2}\|\mathbf{H}_v \cdot - \mathbf{y}\|_2^2 + \|\mathbf{D} \cdot\|_1$. Moreover to generalize the problem we consider $\mathbf{A} \in \mathbb{R}^{n \times m}$ instead of \mathbf{H}_v . In this manner, we introduce the *TV regularized problem* which consists in finding the minimizer of the function P .

$$\min_{\mathbf{u} \in \mathbb{R}^m} \left[P(\mathbf{u}) = \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{u}\|_2^2 + \lambda\|\mathbf{D}\mathbf{u}\|_1 \right] , \quad (3.1)$$

We can make a couple of remark on P . First, it is a proper and closed function which means that P is a lower semicontinuous function. Moreover, P is coercive which means that P admits a minimizer. However, it is also convex since the data fidelity term is a quadratic term, which means that any local minimum will be a global minimum. However, even though $f(\cdot) = \frac{1}{2}\|\mathbf{y} - \mathbf{A} \cdot\|_2^2$ is a smooth function, the use of the ℓ_1 -norm in the regularization makes $g(\cdot) = \lambda\|\cdot\|_{TV}$ not differentiable. In the following section we will see how to minimize P considering those properties.

3.2 Solving a TV regularized problem

In this section, we detail several approaches to solve the TV-reg problem. The formulation proposed in Equation 3.1 is named the analysis formulation. Other equivalent formulations are also presented as well as the classical algorithms to minimize them efficiently.

3.2.1 Analysis formulation

Solving the analysis formulation by the book

Considering a cost function defined by $F = f + g$ with $f : \mathbb{R}^{\tilde{T}} \rightarrow \mathbb{R}^+$ being convex and smooth and $g : \mathbb{R}^{\tilde{T}} \rightarrow \mathbb{R}^+$ being convex, the most straightforward approach to minimize F is to consider the proximal gradient descent (PGD) [Rockafellar, 1976]. PGD algorithm alternates between two steps during each iteration t . The first is called the gradient step and decreases the cost-function along the descent direction defined by the gradient $\nabla f(\mathbf{u}^{(t)})$, such that $\tilde{\mathbf{u}}^{(t+1)} = \mathbf{u}^{(t)} - \rho \nabla f(\mathbf{u}^{(t)})$. The step-size ρ can be chosen in $[0, \frac{2}{L}[$ with L the Lipschitz constant of the gradient operator to have theoretical guaranty on the convergence. The second step is the proximal step, where the regularization term g in the cost-function is taken into account to provide a feasible point by solving at each iteration the proximal operator of $\text{prox}(g, \rho)(\cdot)$. The proximal operator step consists in solving the problem Equation 3.2 w.r.t the iterate $\tilde{\mathbf{u}}^{(t+1)}$.

$$\mathbf{u}^{(t+1)} = \arg \min_{\mathbf{u} \in \mathbb{R}^m} \frac{1}{2}\|\tilde{\mathbf{u}}^{(t+1)} - \mathbf{u}\|_2^2 + g(\mathbf{u}) , \quad (3.2)$$

In some cases, the proximal operator admits a closed-form solution. In the case of the ℓ_1 -norm, this is the so call soft-thresholding or for a convex set characteristic function, this is the orthogonal projection on the set. In the case of the TV-reg, the proximal operator step in Algorithm 1 does not have a closed-form solution. Hence, one needs to solve Equation 3.2 numerically with one of the algorithms exposed in the next paragraphs. However, in the 1D case, powerful dynamic programming algorithms have been developed to solve it exactly (up to numerical precision), such as the *taut-strings algorithm* [Condat, 2013a]. Interestingly,

this algorithm has a complexity of $O(\tilde{T})$ in most practical cases, which allows to efficiently solve the proximal step in [Algorithm 1](#).

Note that the ℓ_1 -norm can be replaced by a smooth function such as the Huber loss in order to smooth the cost-function. In this case, the function can be minimized with second order approaches [[Byrd et al., 1995](#)].

If we consider our cost-function P in [Equation 3.1](#), here we have f such as $f(\mathbf{u}) = \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{u}\|_2^2$ and $\nabla f(\mathbf{u}^{(t)}) = \mathbf{A}^\top(\mathbf{A}\mathbf{u}^{(t)} - \mathbf{y})$; the regularization term g is the aforementioned TV-reg, such as $g(\mathbf{u}) = \lambda\|\mathbf{D}\mathbf{u}\|_1$. Hence, for a scaling parameter $\frac{1}{\rho}$, the proximal operator $\text{prox}(g, \rho)(u)$ is defined by $\text{prox}(g, \rho)(\cdot) = \arg \min_{\mathbf{u} \in \mathbb{R}^T} \left[\frac{1}{2}\|\mathbf{u} - \cdot\|_2^2 + \frac{\lambda}{\rho}\|\mathbf{D}\mathbf{u}\|_1 \right]$. The corresponding PGD algorithm in the analysis formulation is defined in [Algorithm 1](#).

Algorithm 1: Analysis Proximal gradient descent

Input: BOLD signal \mathbf{y} , ϵ , ρ

- 1 initialization: $\mathbf{u}^{(0)} = \mathbf{0}_{\tilde{T}}$, $t = 0$;
- 2 **repeat**
- 3 $\tilde{\mathbf{u}}^{(t+1)} = \mathbf{u}^{(t)} - \frac{1}{\rho}\mathbf{A}^\top(\mathbf{A}\mathbf{u}^{(t)} - \mathbf{y})$
- 4 $\mathbf{u}^{(t+1)} = \arg \min_{\mathbf{u} \in \mathbb{R}^T} \left[\frac{1}{2}\|\tilde{\mathbf{u}}^{(t+1)} - \mathbf{u}\|_2^2 + \frac{\lambda}{\rho}\|\mathbf{D}\mathbf{u}\|_1 \right]$
- 5 **until** $\|\mathbf{u}^{(t+1)} - \mathbf{u}^{(t)}\|_2 \leq \epsilon$;

[algorithm 1](#) is guaranteed to converge to a solution of [Equation 3.1](#) when the step size is properly set as the previously mentioned – as the inverse of the Lipschitz constant of the gradient. Additionally, one can compute at each iteration an optimal decreasing step ρ^* to have a better descent, with the backtracking line search based on the Armijo-Goldstein condition [[Armijo, 1966](#)] for example. Another way to accelerate the [algorithm 1](#) is to introduce a momentum term [[Nesterov, 2013](#), [Liang and Schönlieb, 2018](#)]. The idea is to consider the previous update direction $\mathbf{u}^{(t)} - \mathbf{u}^{(t-1)}$ to update the next iterate $\mathbf{u}^{(t+1)}$. To go further in the acceleration, a restarting scheme can also be added [[Kim and Fessler, 2018](#)], the main objective of those techniques is to cancel the potential rebound of the cost-function evolution *w.r.t* the iteration. We will not go into detail but the accelerating and the restarting scheme will be implemented and their performances will be compared with the classical PGD algorithm in the next section.

Solving the analysis formulation with primal-dual algorithms

Another type of algorithms to solve cost-functions involving a regularization term g is the *primal-dual* algorithm. Here the term *dual* refers to Lagrangian duality. The idea behinds duality is to consider that an optimization problem can be seen in two different ways: a direct minimization of the cost-function or a maximization of a concave lower bound of this cost-function.

The main benefit of *primal-dual* algorithms is that both the primal variable \mathbf{u} and the dual variable denoted $\boldsymbol{\mu}$, in our case, are updated simultaneously. To that aim, we introduce a variable \mathbf{z} and re-write the regularization term as a constrain: our main problem [Equation 3.1](#) is equivalent to the minimization problem [Equation 3.3](#).

$$\min_{\mathbf{u} \in \mathbb{R}^m} \left[P(\mathbf{u}) = \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{u}\|_2^2 + \lambda \|\mathbf{z}\|_1 \right], \quad (3.3)$$

subject to $\mathbf{D}\mathbf{u} - \mathbf{z} = 0$.

We derive the associated Augmented Lagrangian function \mathcal{L} Equation 3.4 from Equation 3.3 to relax the constraint by introducing the dual variable $\boldsymbol{\mu} \in \mathbb{R}^{m-1}$

$$\mathcal{L}(\mathbf{u}, \mathbf{y}, \boldsymbol{\mu}) = \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{u}\|_2^2 + \lambda \|\mathbf{z}\|_1 + \boldsymbol{\mu}^\top (\mathbf{D}\mathbf{u} - \mathbf{z}) + \frac{\gamma}{2} \|\mathbf{D}\mathbf{u} - \mathbf{z}\|_2^2, \quad (3.4)$$

\mathcal{L} is the Augmented Lagrangian and can be minimized via any primal-dual algorithms. Here we provide the example of the Alternating Direction Method of Multipliers (ADMM) introduced in Boyd et al. [2011], see algorithm 2.

Algorithm 2: ADMM

Input: BOLD signal \mathbf{y} , ϵ , ρ

1 initialization: $\mathbf{u}^{(0)} = \mathbf{0}_{\tilde{T}}$, $\mathbf{z}^{(0)} = \mathbf{0}_{\tilde{T}-1}$, $\boldsymbol{\mu}^{(0)} = \mathbf{0}_{\tilde{T}-1}$, $t = 0$;

2 **repeat**

3 $\mathbf{u}^{(t+1)} = (\mathbf{A}^\top \mathbf{A} + \rho \mathbf{D}^\top \mathbf{A})^\dagger (\mathbf{A}^\top \mathbf{y} - \mathbf{D}^\top \boldsymbol{\mu}^{(t)} + \rho \mathbf{D}^\top \mathbf{z}^{(t)})$

4 $\forall i \in [1..\tilde{T}] \quad z_i^{(t+1)} = \frac{1}{\rho} (\boldsymbol{\mu}^{(t)})_i + (\mathbf{D}\mathbf{u}^{(t+1)})_i + \frac{\lambda}{\gamma} \text{sign}(z_i^{(t)})$

5 $\boldsymbol{\mu}^{(t+1)} = \boldsymbol{\mu}^{(t)} + \gamma (\mathbf{D}\mathbf{u}^{(t+1)} - \mathbf{z}^{(t+1)})$

6 **until** $\|\mathbf{D}\mathbf{u}^{(t)} - \mathbf{z}^{(t)}\|_2 \leq \epsilon$;

At each iteration we minimize *mathcal{L}* w.r.t \mathbf{u} , \mathbf{z} and then we make a ascent gradient step concerning $\boldsymbol{\mu}$ such as $\boldsymbol{\mu}^{(t+1)} = \boldsymbol{\mu}^{(t)} + \gamma (\mathbf{D}\mathbf{u}^{(t+1)} - \mathbf{z}^{(t+1)})$. Note that the update of \mathbf{u} and \mathbf{z} can be distributed.

Condat [2013b] proposes an accelerated primal-dual algorithm, by doing a joint gradient descent in the primal along in the dual. The algorithm and its properties are detailed in Condat [2013b], we simply report the definition of the algorithm here algorithm 3:

Algorithm 3: Condat-Vu

Input: BOLD signal \mathbf{y} , ϵ , τ , σ , ρ

1 initialization: $\mathbf{u}^{(0)} = \mathbf{0}_{\tilde{T}}$, $\mathbf{z}^{(0)} = \mathbf{0}_{\tilde{T}-1}$, $t = 0$;

2 **repeat**

3 $\tilde{\mathbf{u}}^{(t+1)} = \mathbf{u}^{(t)} - \tau \mathbf{A}^\top (\mathbf{A}\mathbf{u}^{(t)} - \mathbf{y}) - \tau \mathbf{D}^\top \mathbf{z}^{(t)}$

4 $\tilde{\mathbf{z}}^{(t+1)} = \mathbf{z}^{(t)} + \sigma \mathbf{D} (2\tilde{\mathbf{u}}^{(t+1)} - \mathbf{u}^{(t)}) - \text{ST}_\sigma(\mathbf{z}^{(t)} + \sigma \mathbf{D} (2\tilde{\mathbf{u}}^{(t+1)} - \mathbf{u}^{(t)}))$

5 $\mathbf{u}^{(t+1)} = \eta \tilde{\mathbf{u}}^{(t+1)} + (1 - \eta) \mathbf{u}^{(t)}$

6 $\mathbf{z}^{(t+1)} = \eta \tilde{\mathbf{z}}^{(t+1)} + (1 - \eta) \mathbf{z}^{(t)}$

7 **until** $\|\mathbf{D}\mathbf{u}^{(t)} - \mathbf{z}^{(t)}\|_2 \leq \epsilon$;

In this section, we have briefly introduce *primal-dual* algorithms which are based on the concept of Lagrangian duality. In the next section, we will see that there is another possible definition for the dual and we introduce the Fenchel's duality.

Solving the dual of analysis formulation

The Fenchel's duality is a generalization of the Lagrangian duality, thus the intuition of the approach is same. To obtain the associated dual problem, we define the Fenchel's transform, defined by:

$$\forall \mathbf{u}^* \in \mathbb{R}^m \quad f^*(\mathbf{u}^*) = - \inf_{\mathbf{u} \in \mathbb{R}^m} \left(f(\mathbf{u}) - \mathbf{u}^{*\top} \mathbf{u} \right) \quad (3.5)$$

Applying this transformation to our cost function gives us the dual formulation of our problem, which leads to a problem with a simpler proximal operator associated to the new regularization term. To prove it, we now derive the dual of Equation 3.1. By defining f and g , such as $f(\cdot) = \frac{1}{2} \|\mathbf{A} \cdot - \mathbf{y}\|_2^2$ and $g(\cdot) = \lambda \|\cdot\|_1$ and by denoting p the minimum of Equation 3.1 w.r.t \mathbf{u} , the problem reads:

$$\min_{\mathbf{u} \in \mathbb{R}^m} f(\mathbf{u}) + g(\mathbf{D}\mathbf{u}) \quad (3.6)$$

With the Fenchel-Rockafellar duality theorem, we derive the dual in Equation 3.7.

$$\min_{\boldsymbol{\mu} \in \mathbb{R}^m} f^*(-\mathbf{D}^\top \boldsymbol{\mu}) + g^*(\boldsymbol{\mu}) \quad (3.7)$$

With $\boldsymbol{\mu} \in \mathbb{R}^m$ being the dual variable, we have $g^*(\boldsymbol{\mu}) = - \min_{\mathbf{u} \in \mathbb{R}^m} g(\mathbf{u}) - \boldsymbol{\mu}^\top \mathbf{u}$. With a component-wise minimization, we obtain $g^*(\boldsymbol{\mu})_i = \delta_{|\mu_i| \leq \lambda}$ with δ being the convex indicator. Thus, we deduce that $g^*(\boldsymbol{\mu}) = \delta_{\|\boldsymbol{\mu}\|_\infty \leq \lambda}$. Then, we have $f^*(\boldsymbol{\mu}) = - \min_{\mathbf{u} \in \mathbb{R}^m} f(\mathbf{u}) - \boldsymbol{\mu}^\top \mathbf{u}$. By cancelling the gradient we obtain: $f^*(\boldsymbol{\mu}) = \frac{1}{2} \|\mathbf{A}^\dagger \boldsymbol{\mu}\|_2^2 + \boldsymbol{\mu}^\top \mathbf{A}^\dagger \mathbf{y}$, with \mathbf{A}^\dagger being the pseudo-inverse of \mathbf{A} . Note that, if we set $\mathbf{A} = \mathbf{I}$, we obtain the same problem as Chambolle [2004], Barbero and Sra [2018].

Finally, we obtain the Fenchel dual of the optimization problem Equation 3.1:

$$\begin{aligned} \min_{\boldsymbol{\mu} \in \mathbb{R}^{m-1}} \quad & \frac{1}{2} \|\mathbf{A}^\dagger \mathbf{D}^\top \boldsymbol{\mu}\|_2^2 - \boldsymbol{\mu}^\top \mathbf{D} \mathbf{A}^\dagger \mathbf{y} \\ \text{s.t.} \quad & \|\boldsymbol{\mu}\|_\infty \leq \lambda \end{aligned} \quad (3.8)$$

With Equation 3.8, we can derive a projected gradient descent algorithm easily. Indeed the proximal operator associated to the constraint term $\boldsymbol{\mu} \rightarrow \delta_{\{\|\boldsymbol{\mu}\|_\infty \leq \lambda\}}$ is simply a projection such as $P_j(\boldsymbol{\mu})_i = \begin{cases} \tilde{\mu}_i & \text{if } |\tilde{\mu}_i| \leq \lambda \\ \text{sign}(\tilde{\mu}_i) \lambda & \text{if } |\tilde{\mu}_i| > \lambda \end{cases}$

Algorithm 4: Dual Projected gradient descent

Input : BOLD signal \mathbf{y} , ϵ

- 1 initialization: $\mathbf{z}^{(0)} = \mathbf{0}_{\tilde{T}}$, $t = 0$;
- 2 repeat
- 3 $\tilde{\boldsymbol{\mu}}^{(t+1)} = \boldsymbol{\mu}^{(t)} - \frac{1}{\rho} ((\mathbf{D}\mathbf{A}^\dagger)(\mathbf{D}\mathbf{A}^\dagger)^\top \boldsymbol{\mu} - \mathbf{y}^\top (\mathbf{D}\mathbf{A}^\dagger))$
- 4 $\boldsymbol{\mu}_i^{(t+1)} = \begin{cases} \tilde{\mu}_i^{(t+1)} & \text{if } |\tilde{\mu}_i^{(t+1)}| \leq \lambda \\ \text{sign}(\tilde{\mu}_i^{(t+1)}) \lambda & \text{if } |\tilde{\mu}_i^{(t+1)}| > \lambda \end{cases}$
- 5 until $\|(\mathbf{D}\mathbf{A}^\dagger)(\mathbf{D}\mathbf{A}^\dagger)^\top \boldsymbol{\mu} - \mathbf{y}^\top (\mathbf{D}\mathbf{A}^\dagger)\|_2 \leq \epsilon$;

As we mention, the proximal operator involved in this algorithm has a closed form solution and thus is easy to implement. However, the gradient term involved \mathbf{A}^\dagger the pseudo-inverse of \mathbf{A} . Depending on the rank of the observation matrix its pseudo-inverse could be ill-conditioned and the Lipschitz constant of the gradient could be too high, thus producing a too small step size to achieve a fast minimization.

3.2.2 Synthesis formulation

Finally, in this subsection we detail that the analysis formulation of Equation 3.1 can be rewritten in the so called synthesis formulation. The general idea is to simplify the regularization term with a change of variable. Suppose that we consider a regularization defined by g such as $\forall \mathbf{u} \in \mathbb{R}^m$ $g(\mathbf{u}) = \|\mathbf{M}\mathbf{u}\|_1$ with $\mathbf{M} \in \mathbb{R}^{m \times m}$ being invertible. Then, we could define $\mathbf{z} = \mathbf{M}\mathbf{u} \in \mathbb{R}^m$ such as $g(\mathbf{u}) = \|\mathbf{z}\|_1$. Note that the proximal operator of the ℓ_1 -norm is the Soft-Thresholding operator. Thus, with this change of variable we obtain a closed-form solution for the proximal operator involved in the PGD algorithm.

However, if we consider our cost-function in Equation 3.1 we have $\mathbf{M} = \mathbf{D}$, with \mathbf{D} being not invertible. Thus, we should slightly adapt the described scheme in our case. We introduce $\mathbf{L} \in \mathbb{R}^{T \times T}$ and $\mathbf{R} \in \mathbb{R}^{T \times T}$, such as:

$$\mathbf{L} = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 1 & 1 & \dots & \vdots \\ \vdots & \vdots & \ddots & 0 \\ 1 & \dots & 1 & 1 \end{bmatrix} \quad \mathbf{R} = \begin{bmatrix} 0 & 0 & \dots & 0 \\ 0 & 1 & \dots & \vdots \\ \vdots & \vdots & \ddots & 0 \\ 0 & \dots & 0 & 1 \end{bmatrix}$$

\mathbf{L} is the discrete integration operator and \mathbf{R} asserts that $\forall \mathbf{z} \in \mathbb{R}^T$ such as $\mathbf{u} = \mathbf{L}\mathbf{z}$ we have $\|\mathbf{D}\mathbf{u}\|_1 = \|\mathbf{R}\mathbf{z}\|_1$. Indeed, with \mathbf{R} , we regularized all the coefficients of \mathbf{z} except the first one, which allow recover \mathbf{u} with the integration, such as $\mathbf{u} = \mathbf{L}\mathbf{z}$, as desired. With this adapted change of variable, we now derive the equivalent synthesis formulation in Equation 3.9:

$$\min_{\mathbf{z} \in \mathbb{R}^T} \left[P_{\mathbf{AL}, \lambda}(z) = \frac{1}{2} \|\mathbf{y} - \mathbf{AL}\mathbf{z}\|_2^2 + \lambda \|\mathbf{R}\mathbf{z}\|_1 \right], \quad (3.9)$$

With Equation 3.9, we recover the Least Absolute Shrinkage and Selection Operator (LASSO) problem formulation, which involves a quadratic fidelity term and a ℓ_1 -norm. We can now derive easily a PGD algorithm to minimize it, see algorithm 5:

Algorithm 5: Synthesis Proximal gradient descent (LASSO)

Input: BOLD signal \mathbf{y} , ϵ , ρ

- 1 initialization: $\mathbf{z}^{(0)} = \mathbf{0}_T$, $t = 0$;
- 2 **repeat**
- 3 $\tilde{\mathbf{z}}^{(t+1)} = \mathbf{z}^{(t)} - \frac{1}{\rho} (\mathbf{AL})^\top (\mathbf{AL}\mathbf{z}^{(t)} - \mathbf{y})$
- 4 $\mathbf{z}^{(t+1)} = \text{ST}_\lambda(\tilde{\mathbf{z}}^{(t+1)})$
- 5 **until** $\|\mathbf{z}^{(t+1)} - \mathbf{z}^{(t)}\|_2^2 \leq \epsilon$;

This algorithm is also known as the Iterative Soft-Thresholding Algorithm (ISTA) proposed by Daubechies et al. [2004]. It can be accelerated using a momentum term, as proposed in Beck and Teboulle [2009] with the algorithm Fast ISTA (FISTA).

Moreover, to minimize the LASSO problem at Equation 3.9, the Least Angle Regression (LARS) algorithm can be used. LARS algorithm includes iteratively the coefficients of \mathbf{z} to produce an estimate with the smallest support $S = \{i | z_i \neq 0\}$. At each iteration k the algorithm selects the line \mathbf{AL}_i most correlated with the residual $\mathbf{y} - \mathbf{AL}\mathbf{z}^{(k)}$ and updates the corresponding coefficient $z_i^{(k)}$ with a gradient step. At iteration l , if another line \mathbf{AL}_j becomes more correlated with the residual, the algorithm includes the corresponding coefficient $z_j^{(l)}$ to the active-set, such as $\hat{S}^{(l)} = \{j\} \cup \hat{S}^{(l-1)} = \{i, j\}$, and updates it (*i.e.* $z_i^{(l)}$ and $z_j^{(l)}$)

with a gradient step. If a coefficient hits 0 the algorithm removes it from the active-set. This approach continues to include coefficients until all the coefficients have been considered.

The LARS algorithm is similar to stepwise regression approaches that include iteratively coefficients and minimize exactly the cost-function *w.r.t* the current active-set. The LARS algorithm was introduced by Efron et al. [2004] and it was first applied to the case of the LASSO by Harchaoui and Lévy-Leduc [2007].

To recap, in this section, we have listed and detailed most of the different approaches to solve the TV-reg problem. We have seen that we could derive multiple equivalent formulation of the same problem which produce different algorithms. A natural question would be to investigate which algorithm proposes the fastest minimization of our original problem, we will see this in the next section with an experiment benchmark on synthesis data.

3.3 Performance comparison

To practically explore which algorithms has the best performance we propose to benchmark each of them in a simple deconvolution problem as exposed in Cherkaoui et al. [2019a].

3.3.1 Material and methods

We consider a source signal $\mathbf{u} \in \mathbb{R}^T$ with $T = 100$ such as $\|\mathbf{D}\mathbf{u}\|_0 = 5$ in a first scenario and $\|\mathbf{D}\mathbf{u}\|_0 = 50$ in a second scenario. We consider an HRF $\mathbf{v} \in \mathbb{R}^l$ with $l = 60$ following the HRF model proposed in Cherkaoui et al. [2019a] with the time dilation parameter $\delta = 1$. We introduce the Toeplitz matrix $\mathbf{H} \in \mathbb{R}^{T+l-1 \times T}$ associated to \mathbf{v} . We generate a single synthetic BOLD time serie defined as $\mathbf{y} = \mathbf{v} * \mathbf{u} + \epsilon$ such as $\mathbf{y} \in \mathbb{R}^{T+l-1}$ and $\epsilon \in \mathbb{R}^{T+l-1}$ being a Gaussian noise such as the observed signal as a signal to noise ratio (SNR) of 0.0 dB.

To estimate the source signal \mathbf{u} from the observed signal \mathbf{y} we propose to solve Equation 3.10, defined as:

$$\min_{\mathbf{u} \in \mathbb{R}^T} \left[P_{\mathbf{H},\lambda}(\mathbf{u}) = \frac{1}{2} \|\mathbf{y} - \mathbf{H}\mathbf{u}\|_2^2 + \lambda \|\mathbf{D}\mathbf{u}\|_1 \right], \quad (3.10)$$

Equation 3.10 is similar to Equation 3.1, the only difference being the observation matrix, since $\mathbf{A} = \mathbf{H}$ which provide here a specific structure in this observation operator, note that in this case the matrix is not invertible. Here we fix arbitrary $\lambda = 1.0$. In general the setting of this hyper-parameter is important since it controls the bias of the estimator defined by $\mathbf{y} \rightarrow \arg \min_{\mathbf{u} \in \mathbb{R}^T} [P_{\mathbf{H},\lambda}](\mathbf{u})$. However, since our objective in this section is simply an optimization performance comparison we will set λ arbitrary.

We choose to benchmark the proximal gradient descent with analysis primal formulation, the Condat-Vu algorithm, the analysis dual formulation, the synthesis primal formulation. We also consider the LARS algorithm with the synthesis primal formulation. We made the choice to not include the ADMM algorithm 2 since its performances are known to be sub-optimal in comparison of the Condat-Vu algorithm 3. Moreover, as it's known that the accelerating and restarting scheme can greatly improve the optimization performance of an

PGD algorithm, we include them in the benchmark.

We fix all the step size to the inverse of the Lipschitz constant of the gradient, except for the Condat-Vu algorithm. For the later, we manually fix $\sigma = 0.5$ after a previous benchmark and set optimally τ *w.r.t* σ according to Condat [2013b].

We propose to introduce the momentum acceleration and the restarting scheme to the analysis primal PDG algorithm and the synthesis primal PDG algorithm. For the analysis formulation, we define the momentum, at iteration $t + 1$, as $\theta^{t+1} = \frac{1}{1 + \sqrt{1 + 4(\theta^t)^2}}$ with $\theta^0 = 1$ such as at each iteration $\mathbf{u}^{(t+1)} = \tilde{\mathbf{u}}^{(t+1)} + \frac{\theta^t - 1}{\theta^{t+1}}(\tilde{\mathbf{u}}^{(t+1)} - \mathbf{u}^{(t)})$ with $\tilde{\mathbf{u}}^{(t)}$ being the output of the proximal operator step of the iteration; we do similarly for the synthesis formulation *w.r.t* $\mathbf{z}^{(t)}$. Moreover, concerning the restarting scheme, we reset the momentum such as $\theta^{t+1} = \theta^t$ when $P_{\mathbf{H},\lambda}(\mathbf{u}^{(t+1)}) > P_{\mathbf{H},\lambda}(\mathbf{u}^{(t)})$; we do similarly for the synthesis formulation.

Each algorithm performs 30000 iterations completely and we store at each iteration the cost-function value $P_{\mathbf{H},\lambda}(\mathbf{u}^{(t)})$ along with the duration of the iteration without the cost-function computation duration.

To unsure, fair comparison we initialize each algorithm to the same variable, we choose $\boldsymbol{\mu}^{(0)} = \mathbf{0}_{T-1}$, by applying the K.K.T conditions we obtain $\mathbf{u} = (\mathbf{A}^\top \mathbf{A})^\dagger (\mathbf{A}^\top \mathbf{y} - \mathbf{D}^\top \boldsymbol{\mu})$ such as $\mathbf{u}^{(0)} = (\mathbf{A}^\top \mathbf{A})^\dagger \mathbf{A}^\top \mathbf{y}$ which correspond to $\mathbf{u}^{(0)} = \mathbf{A}^\dagger \mathbf{y}$ and $\mathbf{z}^{(0)} = \mathbf{D}\mathbf{u}^{(0)}$. Those equivalent initializations assert that the cost-function evolution starts at the same value for each algorithm.

A summary of the parameters of the different algorithms can be found in Table 3.1.

3.3.2 Results

We investigate two scenarios, $\|\mathbf{D}\mathbf{u}\|_0 = 5$ and $\|\mathbf{D}\mathbf{u}\|_0 = 50$ with $\mathbf{D}\mathbf{u} \in \mathbb{R}^{99}$. In Figure 3.3-4 and Figure 3.3-5, we report the evolution of the cost-function $P_{\mathbf{H},\lambda}(\mathbf{x}^{(t)}) - P_{\mathbf{H},\lambda}(\mathbf{u}^*)$ *w.r.t* the iterations t and the time in seconds for each algorithms in log-scale for both scenarios. We exclude the LARS algorithm from the loss comparison but will report later its average convergence time.

In both case, we observe that the order of performance for each algorithms stays the same. We notice that only the synthesis primal PGD and APGD did not reach the minimum of the cost-function whereas all the other algorithms have. We observe that the synthesis formulation algorithms (PGD, APGD, RAPGD) have the worst performance, as a comparison, the fastest synthesis formulated algorithm (RAPGD) takes 100 more iterations to converge to the same cost-function value than the RAPGD analysis primal formulation. The Condat-Vu algorithm has in-between performance, yet it has similar performance of the RAPGD synthesis primal formulation algorithm. Moreover, the PGD analysis dual formulation algorithm has one of the worst performance. In addition, all the analysis primal formulation algorithms (PGD, APGD, RAPGD) have the best minimization performance, the RAPGD scheme being the fastest algorithm to reach the minimum for both scenarios. The good performance of the analysis primal formulation algorithm can be explained by the efficiency of the proximal operator step in comparison of the synthesis formulation.

Table 3.1 – Summary of the algorithm’s hyper-parameters used to benchmark their performance.

Algorithm	Step size	Momentum	Restarting	Initialization
Analysis primal Proximal gradient descent (Analysis primal PGD)	$\rho = \ \mathbf{H}\ _F$	-	-	$\mathbf{u}^{(0)} = \mathbf{A}^\dagger \mathbf{y}$
Accelerated Analysis primal Proximal gradient descent (Analysis primal APGD)	$\rho = \ \mathbf{H}\ _F$	$\theta^{t+1} = \frac{1}{1 + \sqrt{1 + 4(\theta^t)^2}}$ $\mathbf{u}^{(t+1)} = \tilde{\mathbf{u}}^{(t+1)} + \frac{\theta^t - 1}{\theta^t + 1}(\tilde{\mathbf{u}}^{(t+1)} - \mathbf{u}^{(t)})$ ^a	-	$\mathbf{u}^{(0)} = \mathbf{A}^\dagger \mathbf{y}$
Restarting Analysis primal Proximal gradient descent (Analysis primal RAPGD)	$\rho = \ \mathbf{H}\ _F$	$\theta^{t+1} = \frac{1}{1 + \sqrt{1 + 4(\theta^t)^2}}$ $\mathbf{u}^{(t+1)} = \tilde{\mathbf{u}}^{(t+1)} + \frac{\theta^t - 1}{\theta^t + 1}(\tilde{\mathbf{u}}^{(t+1)} - \mathbf{u}^{(t)})$ ^b	if $P_{\mathbf{H},\lambda}(\mathbf{u}^{(t+1)}) > P_{\mathbf{H},\lambda}(\mathbf{u}^{(t)})$ then $\theta^{t+1} = \theta^t$	$\mathbf{u}^{(0)} = \mathbf{A}^\dagger \mathbf{y}$
Condat-Vu	$\frac{1}{\tau} = \frac{\sigma = 0.5}{\frac{\ \mathbf{H}\ _F}{2} + \sigma \ \mathbf{D}\ _F}$ $\eta = 1.0$	-	-	$\mathbf{u}^{(0)} = \mathbf{A}^\dagger \mathbf{y}$ and $\mathbf{z}^{(0)} = \mathbf{D}\mathbf{u}^{(0)}$
Dual Projected gradient descent (Analysis dual PGD)	$\rho = \ \mathbf{H}^\dagger \mathbf{D}\ _F$	-	-	$\mathbf{v}^{(0)} = \mathbf{0}_{T-1}$
Synthesis Proximal gradient descent (Synthesis primal PGD)	$\rho = \ \mathbf{HL}\ _F$	-	-	$\mathbf{z}^{(0)} = \mathbf{D}\mathbf{A}^\dagger \mathbf{y}$
Accelerated Synthesis Proximal gradient descent (Synthesis primal APGD)	$\rho = \ \mathbf{HL}\ _F$	$\theta^{t+1} = \frac{1}{1 + \sqrt{1 + 4(\theta^t)^2}}$ $\mathbf{z}^{(t+1)} = \tilde{\mathbf{z}}^{(t+1)} + \frac{\theta^t - 1}{\theta^t + 1}(\tilde{\mathbf{z}}^{(t+1)} - \mathbf{z}^{(t)})$ ^c	-	$\mathbf{z}^{(0)} = \mathbf{D}\mathbf{A}^\dagger \mathbf{y}$
Restarting Accelerated Synthesis Proximal gradient descent (Synthesis primal RAPGD)	$\rho = \ \mathbf{HL}\ _F$	$\theta^{t+1} = \frac{1}{1 + \sqrt{1 + 4(\theta^t)^2}}$ $\mathbf{z}^{(t+1)} = \tilde{\mathbf{z}}^{(t+1)} + \frac{\theta^t - 1}{\theta^t + 1}(\tilde{\mathbf{z}}^{(t+1)} - \mathbf{z}^{(t)})$ ^d	if $P_{\mathbf{H},\lambda}(\mathbf{z}^{(t+1)}) > P_{\mathbf{H},\lambda}(\mathbf{z}^{(t)})$ then $\theta^{t+1} = \theta^t$	$\mathbf{z}^{(0)} = \mathbf{D}\mathbf{A}^\dagger \mathbf{y}$
LARS algorithm (LARS-LASSO)	-	-	-	$\mathbf{z}^{(0)} = \mathbf{D}\mathbf{A}^\dagger \mathbf{y}$

^a $\tilde{\mathbf{u}}$ refers to the output of the proximal operator step
^b See footnote 3.1
^c $\tilde{\mathbf{z}}$ refers to the output of the proximal operator step
^d See footnote 3.1




Figure 3.3-4 – (left) Performance comparison *w.r.t* the iterations, (right) Performance comparison *w.r.t* the time. We compare the performance of each algorithms in the scenario where $\|\mathbf{D}\mathbf{u}\|_0 = 5$.

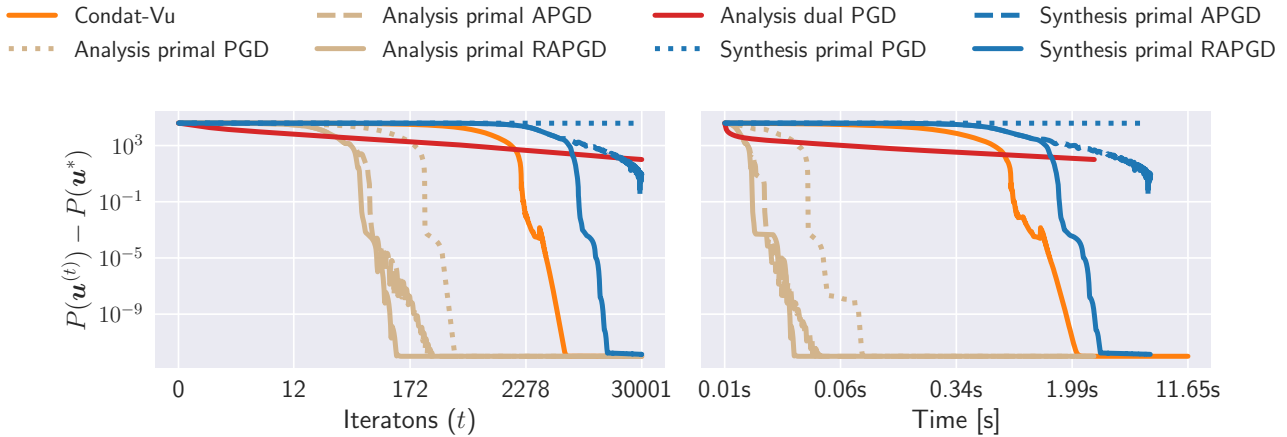


Figure 3.3-5 – (left) **Performance comparison *w.r.t* the iterations**, (right) **Performance comparison *w.r.t* the time**. We compare the performance of each algorithms in the scenario where $\|\mathbf{D}\mathbf{u}\|_0 = 50$.

Additionally, we investigate the time spent for each algorithm (including the LARS algorithm) to reach the minimum of the cost-function $P_{\mathbf{H},\lambda}(\mathbf{u}^*)$ with a precision of 10^{-3} . We report the results for $\lambda \in \{0.001, 1.0, 10.0\}$. We study the scenario where $\|\mathbf{D}\mathbf{u}\|_0 = 5$. Most of the synthesis formulation algorithm (PGD, APGD) and the dual formulation algorithm are not reported since they do not reach this precision sufficiently fast. The time spent is displayed in log-scale for all the scenarios. To achieve fair comparison, we run 10 times each algorithm. In Figure 3.3-6, we display the average time for each method and represent the variance with a black line on top of each bar.

In all cases, we notice that the analysis formulation (PGD, APGD, RAPGD) and the LARS algorithm with the synthesis formulation are the faster than the other methods, the RAPGD scheme being the fastest algorithm. The Condat-Vu algorithm has in-between performance and the RAPGD synthesis formulation algorithm has the worst performance of the reported algorithms. We notice that the higher λ , the faster all the algorithms are (except the LARS algorithm). This can be explained as the support estimation is easier in the case where the solution has more os. Concerning the LARS algorithm, the level of regularization does not affect the convergence time.

3.4 Chapter conclusion

We have defined and described what is the TV regularization, how to solve a problem that involve it and compared experimentally these algorithms. We found that in practice, if one can implement the Taut-string algorithm in an efficient way then the analysis formulation with the prox-TV operator achieves the best performance.

However, in the next chapter, we will see that it is possible to further accelerate the resolution of TV regularized problems, in particular using learned algorithms.

* * *
* *
*

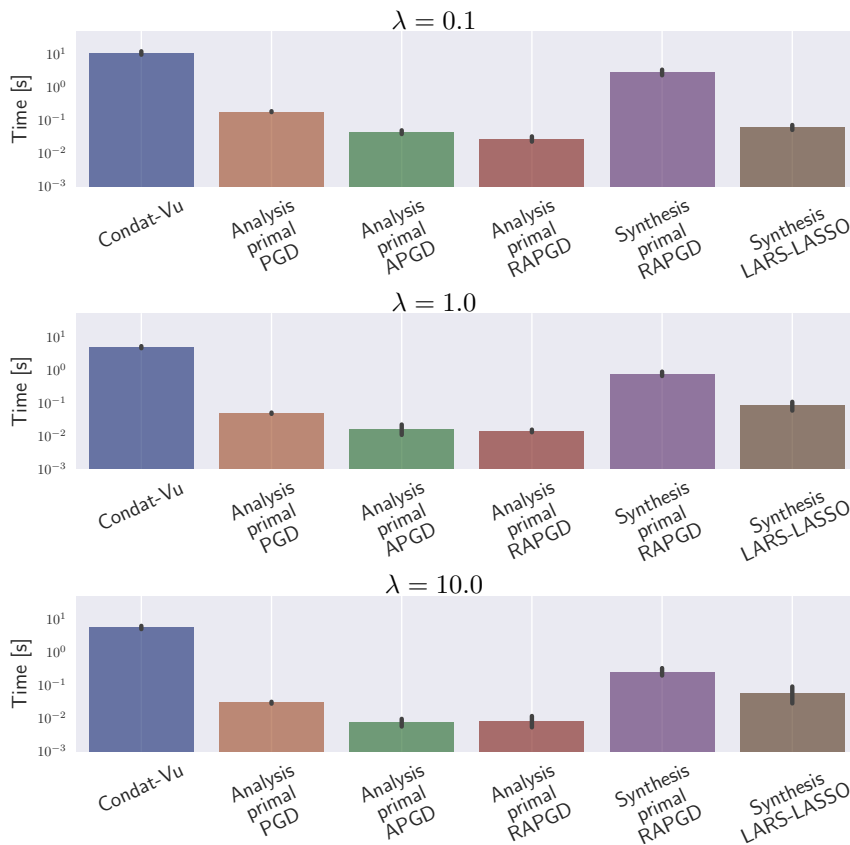


Figure 3.3-6 – **Convergence time comparison.** We report for each algorithm the average time spent to reach a value of the cost function with sub-optimality 10^{-3} . We consider the scenario where $\lambda \in \{0.1, 1.0, 10.0\}$. We notice the higher the TV regularization the faster the algorithms are. Their relative performance is quite stable, with a clear advantage for the Analysis methods compared to the other, due to the efficiency of the prox-TV.

Part II

Methodological Developments

TV problems minimization with learned algorithms

Chapter Outline

4.1	Introduction to TV regularization in 1D	56
4.2	Solving a TV regularized problem	57
4.2.1	Solving TV-regularized problems	57
4.2.2	Back-propagating through TV proximal operator	63
4.2.3	Unrolled prox-TV	66
4.3	Performance comparison	69
4.3.1	Simulation	70
4.3.2	fMRI data deconvolution.	72
4.4	Conclusion	72

This chapter was presented at an International conference as:

H. Cherkaoui, J. Sulam, and T. Moreau. Learning to solve tv regularised problems with unrolled algorithms. In *34th Conference and Workshop on Neural Information Processing Systems (NeurIPS)*, pages 1–21, 2020b.

THE purpose of this chapter is to introduce our novel approach based on neural networks to efficiently minimize a cost-function involving the Total Variation (TV) regularization term nearby a quadratic data consistency term. We propose a theoretical comparison of multiple equivalent formulations of the considered optimization problem, introduce our novel method and provide experimental validations to demonstrate its usefulness notably on real fMRI data.

4.1 Introduction to TV regularization in 1D

As explained, the TV-regularized problems are typically convex, and so a wide variety of algorithms are in principle applicable. Since the ℓ_1 norm in the TV term is non-smooth, Proximal Gradient Descent (PGD) is the most popular choice [Rockafellar, 1976]. Yet, the computation for the corresponding proximal operator (denoted prox-TV) represents a major difficulty in this case as it does not have a closed-form analytic solution. For 1D problems, it is possible to rely on dynamic programming to compute prox-TV, such as the taut string algorithm [Davies and Kovac, 2001, Condat, 2013a]. Another alternative consists in computing the proximal operator with iterative first order algorithm [Chambolle, 2004, Beck and Teboulle, 2009, Boyd et al., 2011, Condat, 2013b]. Other algorithms to solve TV-regularized problems rely on primal dual algorithms [Chambolle and Pock, 2011, Condat, 2013b] or Alternating Direction Method of Multipliers (ADMM) [Boyd et al., 2011]. These algorithms typically use one sequence of estimates for each term in the objective and try to make them as close as possible while minimizing the associated term. While these algorithms are efficient for denoising problems – where one is mainly concerned with good reconstruction – they can result in estimate that are not very well regularized if the two sequences are not close enough.

When on fixed computational budget, iterative optimization methods can become impractical as they often require many iterations to give a satisfactory estimate. To accelerate the resolution of these problems with a finite (and small) number of iterations, one can resort to unrolled and learned optimization algorithms (see Monga et al. 2019 for a review). In their seminal work, Gregor and Le Cun [2010] proposed the Learned ISTA (LISTA), where the parameters of an unfolded Iterative Shrinkage-Thresholding Algorithm (ISTA) are learned with gradient descent and back-propagation. This allows to accelerate the approximate solution of a Lasso problem [Tibshirani, 1996], with a fixed number of iteration, for signals from a certain distribution. The core principle behind the success of this approach is that the network parameters can adaptively leverage the sensing matrix structure [Moreau and Bruna, 2017] as well as the input distribution [Giryas et al., 2018, Ablin et al., 2019]. Many extensions of this original idea have been proposed to learn different algorithms [Sprechmann et al., 2012, 2013, Borgerding et al., 2017] or for different classes of problem [Xin et al., 2016, Giryas et al., 2018, Sulam et al., 2019]. The motif in most of these adaptations is that all operations in the learned algorithms are either linear or separable, thus resulting in sub-differentials that are easy to compute and implement via back-propagation. Algorithm unrolling is also used in the context of bi-level optimization problems such as hyper-parameter selection. Here, the unrolled architecture provides a way to compute the derivative of the inner optimization problem solution compared to another variable such as the regularisation parameter using back-propagation [Bertrand et al., 2020].

The focus of this chapter is to apply algorithm unrolling to TV-regularized problems in the 1D case. While one could indeed apply the LISTA approach directly to the *synthesis* formulation of these problems, we show in this chapter that using such formulation leads to slower iterative or learned algorithms compared to their *analysis* counterparts. The extension of learnable algorithms to the analysis formulation is not trivial, as the inner proximal operator does not have an analytical or separable expression. We propose two architectures that can learn TV-solvers in their analysis form directly based on PGD. The first architecture

uses an exact algorithm to compute the prox-TV and we derive the formulation of its weak Jacobian in order to learn the network’s parameters. Our second method rely on a nested LISTA network in order to approximate the prox-TV itself in a differentiable way. This latter approach can be linked to inexact proximal gradient methods [Schmidt et al., 2011, Machart et al., 2012]. These results are backed with numerical experiments on synthetic and real data. Concurrently to our work, Lecouat et al. [2020] also proposed an approach to differentiate the solution of TV-regularized problems. While their work can be applied in the context of 2D signals, they rely on smoothing the regularization term using Moreau-Yosida regularization, which results in smoother estimates from their learned networks. In contrast, our work allows to compute sharper signals but can only be applied to 1D signals.

4.2 Solving a TV regularized problem

The chapter is organized as follows. In subsection 4.2.1, we describe the different formulations for TV-regularized problems and their complexity. We also recall central ideas of algorithm unfolding. subsection 4.2.2 introduces our two approaches for learnable network architectures based on PGD. Finally, the two proposed methods are evaluated on synthetic and real data in subsection 4.3.1 and subsection 4.3.2.

Notations For a vector $x \in \mathbb{R}^k$, we denote $\|x\|_q$ its ℓ_q -norm. For a matrix $A \in \mathbb{R}^{m \times k}$, we denote $\|A\|_2$ its ℓ_2 -norm, which corresponds to its largest singular value and A^\dagger denotes its pseudo-inverse. For an ordered subset of indices $\mathcal{S} \subset \{1, \dots, k\}$, $x_{\mathcal{S}}$ denote the vector in $\mathbb{R}^{|\mathcal{S}|}$ with element $(x_{\mathcal{S}})_t = x_{i_t}$ for $i_t \in \mathcal{S}$. For a matrix $A \in \mathbb{R}^{m \times k}$, $A_{:, \mathcal{S}}$ denotes the sub-matrix $[A_{:, i_1}, \dots, A_{:, i_{|\mathcal{S}|}}]$ composed with the columns $A_{:, i_t}$ of index $i_t \in \mathcal{S}$ of A . For the rest of the chapter, we refer to the operators $D \in \mathbb{R}^{k-1 \times k}$, $\tilde{D} \in \mathbb{R}^{k \times k}$, $L \in \mathbb{R}^{k \times k}$ and $R \in \mathbb{R}^{k \times k}$ as:

$$D = \begin{bmatrix} -1 & 1 & 0 & \dots & 0 \\ 0 & -1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix} \quad \tilde{D} = \begin{bmatrix} 1 & 0 & \dots & 0 \\ -1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & -1 & 1 & \end{bmatrix} \quad L = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 1 & \dots & 1 & 1 \end{bmatrix} \quad R = \begin{bmatrix} 0 & 0 & \dots & 0 \\ 0 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & 1 \end{bmatrix}$$

4.2.1 Solving TV-regularized problems

We begin by detailing the TV-regularized problem that will be the main focus of our work. Consider a latent vector $u \in \mathbb{R}^k$, a design matrix $A \in \mathbb{R}^{m \times k}$ and the corresponding observation $x \in \mathbb{R}^m$. The original formulation of the TV-regularized regression problem is referred to as the *analysis* formulation [Rudin et al., 1992]. For a given regularization parameter $\lambda > 0$, it reads

$$\min_{u \in \mathbb{R}^k} P(u) = \frac{1}{2} \|x - Au\|_2^2 + \lambda \|u\|_{TV}, \quad (4.1)$$

where $\|u\|_{TV} = \|Du\|_1$, and $D \in \mathbb{R}^{k-1 \times k}$ stands for the first order finite difference operator, as defined above. The problem in Equation 4.1 can be seen as a special case of a Generalized Lasso problem [Tibshirani and Taylor, 2011]; one in which the analysis operator is D . Note that problem P is convex, but the TV -norm is non-smooth. In these cases, a practical alternative is the PGD, which iterates between a gradient descent step and the prox-TV.

This algorithm's iterates read

$$u^{(t+1)} = \text{prox}_{\frac{\lambda}{\rho}\|\cdot\|_{TV}} \left(u^{(t)} - \frac{1}{\rho} A^\top (Au^{(t)} - x) \right), \quad (4.2)$$

where $\rho = \|A\|_2^2$ and the prox-TV is defined as

$$\text{prox}_{\mu\|\cdot\|_{TV}}(y) = \arg \min_{u \in \mathbb{R}^k} F_y(u) = \frac{1}{2} \|y - u\|_2^2 + \mu \|u\|_{TV}. \quad (4.3)$$

Problem [Equation 4.3](#) does not have a closed-form solution, and one needs to resort to iterative techniques to compute it. In our case, as the problem is 1D, the prox-TV problem can be addressed with a dynamic programming approach, such as the taut-string algorithm [[Condat, 2013a](#)]. This scales as $O(k)$ in all practical situations and is thus much more efficient than other optimization based iterative algorithms [[Rockafellar, 1976](#), [Chambolle, 2004](#), [Condat, 2013b](#)] for which each iteration is $O(k^2)$ at best.

With a generic matrix $A \in \mathbb{R}^{m \times k}$, the PGD algorithm is known to have a sublinear convergence rate [[Combettes and Bauschke, 2011](#)]. More precisely, for any initialization $u^{(0)}$ and solution u^* , the iterates satisfy

$$P(u^{(t)}) - P(u^*) \leq \frac{\rho}{2t} \|u^{(0)} - u^*\|_2^2, \quad (4.4)$$

where u^* is a solution of the problem in [Equation 4.1](#). Note that the constant ρ can have a significant effect. Indeed, it is clear from [Equation 4.4](#) that doubling ρ leads to consider doubling the number of iterations.

Synthesis formulation

An alternative formulation for TV-regularized problems relies on removing the analysis operator D from the ℓ_1 -norm and translating it into a synthesis expression [[Elad et al., 2007](#)]. Removing D from the non-smooth term simplifies the expression of the proximal operator by making it separable, as in the Lasso. The operator D is not directly invertible but keeping the first value of the vector u allows for perfect reconstruction. This motivates the definition of the operator $\tilde{D} \in \mathbb{R}^{k \times k}$, and its inverse $L \in \mathbb{R}^{k \times k}$, as defined previously. Naturally, L is the discrete integration operator. Considering the change of variable $z = \tilde{D}u$, and using the operator $R \in \mathbb{R}^{k \times k}$, the problem in [Equation 4.1](#) is equivalent to

$$\min_{z \in \mathbb{R}^k} S(z) = \frac{1}{2} \|x - ALz\|_2^2 + \lambda \|Rz\|_1. \quad (4.5)$$

Note that for any $z \in \mathbb{R}^k$, $S(z) = P(Lz)$. There is thus an exact equivalence between solutions from the synthesis and the analysis formulation, and the solution for the analysis can be obtained with $u^* = Lz^*$. The benefit of this formulation is that the problem above now reduces to a Lasso problem [[Tibshirani, 1996](#)]. In this case, the PGD algorithm is reduced to the ISTA with a closed-form proximal operator (the soft-thresholding). Note that this simple formulation is only possible in 1D where the first order derivative space is unconstrained. In larger dimensions, the derivative must be constrained to verify the Fubini's formula that enforces the symmetry of integration over dimensions. While it is also possible to derive synthesis formulation in higher dimension [[Elad et al., 2007](#)], this does not lead to simplistic proximal operator.

For this synthesis formulation, with a generic matrix $A \in \mathbb{R}^{m \times k}$, the PGD algorithm has also a sublinear convergence rate [Beck and Teboulle, 2009] such that

$$P(u^{(t)}) - P(u^*) \leq \frac{2\tilde{\rho}}{t} \|u^{(0)} - u^*\|_2^2, \quad (4.6)$$

with $\tilde{\rho} = \|AL\|_2^2$.

Proof. The convergence rate of ISTA for the synthesis formulation reads

$$S(z^{(t)}) - S(z^*) \leq \frac{\tilde{\rho}}{2t} \|z^{(0)} - z^*\|_2^2. \quad (4.7)$$

We use $S(z^{(t)}) = P(Lz^{(t)}) = P(u^{(t)})$ to get the correct left-hand side term. For the right hand side, we use $z^{(0)} = \tilde{D}u^{(0)}$, and $z^* = \tilde{D}u^*$, which gives $\|z^{(0)} - z^*\|_2^2 = \|\tilde{D}(u^{(0)} - u^*)\|_2^2 \leq 4\|u^{(0)} - u^*\|_2^2$. The last majoration comes from the fact that $\|\tilde{D}\|_2^2 \leq 4$, as shown per Lemma 4.2.1. This yields

$$P(u^{(t)}) - P(u^*) \leq \frac{2\tilde{\rho}}{t} \|u^{(0)} - u^*\|_2^2. \quad (4.8)$$

□

Lemma 4.2.1. [Singular values of L] The singular values of $L \in \mathbb{R}^{k \times k}$ are given by

$$\sigma_l = \frac{1}{2 \cos\left(\frac{\pi l}{2k+1}\right)}, \quad \forall l \in \{1, \dots, k\}.$$

Thus, $\|L\|_2 = \frac{2k+1}{\pi} + o(1)$.

Proof. As L is invertible, so is $L^\top L$. To compute the singular values σ_l of L , we will compute the eigenvalues μ_l of $(L^\top L)^{-1}$ and use the relation

$$\sigma_l = \frac{1}{\sqrt{\mu_l}} \quad (4.9)$$

With simple computations, we obtain

$$M_k = (L^\top L)^{-1} = L^{-1}(L^\top)^{-1} = \tilde{D}\tilde{D}^\top = \begin{bmatrix} 1 & -1 & 0 & \dots & \\ -1 & 2 & -1 & 0 & \dots \\ & \ddots & \ddots & \ddots & \\ & & 0 & -1 & 2 & -1 \\ & & & & -1 & 2 \end{bmatrix} \quad (4.10)$$

This matrix is tri-diagonal with a quasi-toeplitz structure. Its characteristic polynomial $P_k(\mu)$ is given by:

$$P_k(\mu) = |\mu \text{Id} - M_k| = \begin{vmatrix} \mu - 1 & 1 & 0 & \dots & \\ 1 & \mu - 2 & 1 & 0 & \dots \\ & \ddots & \ddots & \ddots & \\ & & 0 & 1 & \mu - 2 & 1 \\ & & & 0 & 1 & \mu - 2 \end{vmatrix} \quad (4.11)$$

$$= (\mu - 1)Q_{k-1}(\mu) - Q_{k-2}(\mu) \quad (4.12)$$

where Equation 4.12 is obtained by developing the determinant relatively to the first line and $Q_k(\mu)$ is the characteristic polynomial of matrix \widetilde{M}_k equal to M_k except for the the top left coefficient which is replaced by 2

$$\widetilde{M}_k = \begin{bmatrix} 2 & -1 & 0 & \dots & \\ -1 & 2 & -1 & 0 & \dots \\ & \ddots & \ddots & \ddots & \\ & & 0 & -1 & 2 & -1 \\ & & & 0 & -1 & 2 \end{bmatrix} \quad (4.13)$$

Using the same development as Equation 4.12, one can show that Q_k verifies the recurrence relationship

$$Q_k(\mu) = (\mu - 2)Q_{k-1}(\mu) - Q_{k-2}(\mu); \quad Q_1(\mu) = 2 - \mu, \quad Q_0(\mu) = 1 . \quad (4.14)$$

Using this with Equation 4.12 yields

$$P_k(\mu) = Q_k(\mu) + Q_{k-1}(\mu) \quad (4.15)$$

With the change of variable $\nu = \frac{\mu-2}{2}$ and denoting $U_k(\nu) = Q_k(2 + 2\nu)$, the recursion becomes

$$U_k(\nu) = 2\nu U_{k-1}(\nu) - U_{k-2}(\nu); \quad U_1(\nu) = 2\nu, \quad U_0(\mu) = 1 . \quad (4.16)$$

This recursion defines the Chebyshev polynomials of the second kind [Chebyshev, 1853] which verifies the following relation

$$\forall \theta \in [0, 2\pi], \quad U_k(\cos(\theta)) \sin(\theta) = \sin((k+1)\theta) . \quad (4.17)$$

Translating this relationship to Q_k gives

$$\forall \theta \in [0, 2\pi], \quad Q_k(2 + 2 \cos(\theta)) \sin(\theta) = \sin((k+1)\theta) . \quad (4.18)$$

Using this in Equation 4.15 shows that for $\theta \in [0, 2\pi[$ P_k verify

$$P_k(2 + 2 \cos(\theta)) \sin(\theta) = \sin((k+1)\theta) + \sin(k\theta) . \quad (4.19)$$

The equation

$$\sin((k+1)\theta) + \sin(k\theta) = 0 , \quad (4.20)$$

has l solution in $[0, 2\pi[$ that are given by $\theta_l = \frac{2\pi l}{2k+1}$ for $l \in \{1, \dots, n\}$. As for all l , $\sin(\theta_l) \neq 0$, the values $\mu_l = 2 + 2 \cos(\theta_l) = 4 \cos^2(\frac{\pi l}{2k+1})$ are the roots of P_k and therefor the eigenvalues of M_k . Using Equation 4.9 yields the expected value for σ_l .

The singular value of L is thus obtain for $l = k$ and we get

$$\|L\|_2 = \sigma_k = \frac{1}{2 \cos\left(\frac{\pi k}{2k+1}\right)} = \frac{1}{2 \cos\left(\frac{\pi}{2}\left(1 - \frac{1}{2k+1}\right)\right)} , \quad (4.21)$$

$$= \frac{1}{2 \sin\left(\frac{\pi}{2} \frac{1}{2k+1}\right)} = \frac{2k+1}{\pi} + o(1) . \quad (4.22)$$

Where the last approximation comes from $\frac{1}{\sin(x)} = 1/x + o(1)$ when x is close to 0. \square

While the rate of this algorithm is the same as in the analysis formulation – in $O(\frac{1}{t})$ – the constant $\tilde{\rho}$ related to the operator norm differs. We now present two results that will characterize the value of $\tilde{\rho}$.

Proposition 4.2.2. *[Lower bound for the ratio $\frac{\|AL\|_2^2}{\|A\|_2^2}$ expectation] Let A be a random matrix in $\mathbb{R}^{m \times k}$ with i.i.d normally distributed entries. The expectation of $\|AL\|_2^2/\|A\|_2^2$ is asymptotically lower bounded when k tends to ∞ by*

$$\mathbb{E} \left[\frac{\|AL\|_2^2}{\|A\|_2^2} \right] \geq \frac{2k+1}{4\pi^2} + o(1)$$

Proof. Finding the norm of AL can be written as

$$\|AL\|_2^2 = \max_{x \in \mathbb{R}^k} xL^\top A^\top ALx; \quad st \|x\|_2 = 1 \quad (4.23)$$

From [Lemma 4.2.1](#), we can write $L = W^\top \Sigma V$ with V, W two unitary matrices and Σ a diagonal matrix with $\Sigma_{l,l} = \sigma_l$ for all $l \in \{1, \dots, k\}$.

First, we consider the case where $A^\top A$ is a rank one matrix with $A^\top A = \|A\|_2^2 u_1 u_1^\top$, with vector u_1 uniformly sampled from the ℓ_2 -ball and fixed $\|A\|_2$. In this case, as W is unitary, $w_1 = Wu_1$ is also a vector uniformly sampled from the sphere. Also as V is unitary, it is possible to re-parametrize [Equation 4.23](#) by $y = Vx$ such that

$$\max_{y \in \mathbb{R}^k} \|A\|_2^2 y^\top \Sigma w_1 w_1^\top \Sigma y; \quad st \|y\|_2 = 1 \quad (4.24)$$

This problem can be maximized by taking $y = \frac{\Sigma u_1}{\|\Sigma u_1\|_2}$, which gives

$$\|AL\|_2^2 = \|A\|_2^2 \|\Sigma w_1\|_2^2 \quad (4.25)$$

Then, we compute the expectation of $\|\Sigma w_1\|_2^2$ with respect with w_1 , a random vector sampled in the ℓ_2 unit ball,

$$\mathbb{E}_{w_1} [\|\Sigma w_1\|_2^2] = \sum_{l=1}^k \sigma_l^2 \mathbb{E}[u_{1,i}^2] = \sum_{l=1}^k \frac{1}{4 \cos^2 \frac{\pi l}{2k+1}} \frac{1}{k} = \frac{1}{2\pi} \sum_{l=1}^k \frac{\pi}{2k} \frac{1}{\cos^2 \frac{\pi l}{2k+1}}. \quad (4.26)$$

Here, we made use of the fact that for a random vector u_1 on the sphere in dimension k , $\mathbb{E}[u_{1,i}] = \frac{1}{k}$. In the last part of the equation, we recognize a Riemann sum for the interval $[0, \frac{\pi}{2}]$. However, $x \mapsto \frac{1}{\cos^2(x)}$ is not integrable on this interval. As the function is positive and monotone, we can still use the integral to highlight the asymptotic behavior of the series. For k large enough, we consider the integral

$$\int_0^{\frac{\pi}{2} - \frac{\pi}{2k+1}} \frac{1}{\cos^2(x)} dx = \left[\frac{\sin(x)}{\cos(x)} \right]_0^{\frac{\pi}{2} - \frac{\pi}{2k+1}} = \frac{\cos \frac{\pi}{2k+1}}{\sin \frac{\pi}{2k+1}} = \frac{2k+1}{\pi} + o(1) \quad (4.27)$$

Thus, for k large enough, we obtain

$$\mathbb{E}_{w_1} [\|\Sigma w_1\|_2^2] = \frac{1}{2\pi} \left(\frac{2k+1}{\pi} + o(1) \right) \quad (4.28)$$

Thus, we get

$$\mathbb{E} \left[\frac{\|AL\|_2^2}{\|A\|_2^2} \right] = \left(\frac{k + \frac{1}{2}}{\pi^2} + o(1) \right) \quad (4.29)$$

This concludes the case where $A^\top A$ is of rank-1 with uniformly distributed eigenvector.

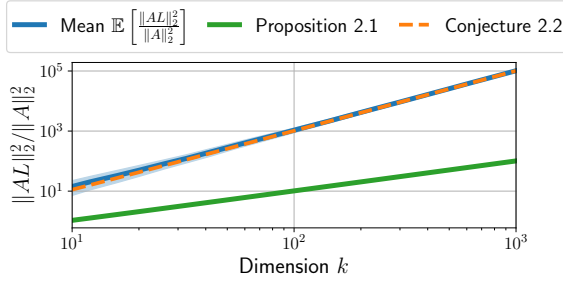


Figure 4.2-1 – Evolution of $\mathbb{E} \left[\frac{\|AL\|_2^2}{\|A\|_2^2} \right]$ w.r.t the dimension k for random matrices A with *i.i.d* normally distributed entries

. In light blue is the confidence interval $[0.1, 0.9]$ computed with the quantiles. We observe that it scales as $O(k^2)$ and that our conjectured bound seems tight.

In the case where $A^\top A$ is larger rank, it is lower bounded by $\|A\|_2^2 u_1 u_1^\top$ where u_1 is its eigenvector associated to its largest eigenvalue, since it is *psd*. Since $A^\top A$ is a Whishart matrix, its eigenvectors are uniformly distributed on the sphere [Silverstein, 1989]. We can thus use the same lower bound as previously for the whole matrix. \square

The full proof can be found in Proposition 4.2.1. The lower bound is constructed by using $A^\top A \succeq \|A\|_2^2 u_1 u_1^\top$ for a unit vector u_1 and computing explicitly the expectation for rank one matrices. To assess the tightness of this bound, we evaluated numerically $\mathbb{E} \left[\frac{\|AL\|_2^2}{\|A\|_2^2} \right]$ on a set of 1000 matrices sampled with *i.i.d* normally distributed entries. The results are displayed w.r.t the dimension k in Figure 4.2-1. It is clear that the lower bound from Proposition 4.2.2 is not tight. This is expected as we consider only the leading eigenvector of A to derive it in the proof. The following conjecture gives a tighter bound.

Conjecture 4.2.3 (Expectation for the ratio $\frac{\|AL\|_2^2}{\|A\|_2^2}$). *Under the same conditions as in Proposition 4.2.2, the expectation of $\|AL\|_2^2 / \|A\|_2^2$ is given by*

$$\mathbb{E} \left[\frac{\|AL\|_2^2}{\|A\|_2^2} \right] = \frac{(2k+1)^2}{16\pi^2} + o(1) .$$

We believe this conjecture can potentially be proven with analogous developments as those in Proposition 4.2.2, but integrating over all dimensions. However, a main difficulty lies in the fact that integration over all eigenvectors have to be carried out jointly as they are not independent. This is subject of current ongoing work.

Finally, we can expect that $\tilde{\rho}/\rho$ scales as $\Theta(k^2)$. This leads to the observation that $\frac{\tilde{\rho}}{\rho} \gg \rho$ in large enough dimension. As a result, the analysis formulation should be much more efficient in terms of iterations than the synthesis formulation – as long as the prox-TV can be dealt with efficiently.

Unrolled iterative algorithms

As shown by Gregor and Le Cun [2010], ISTA is equivalent to a recurrent neural network (RNN) with a particular structure. This observation can be generalized to PGD algorithms for any penalized least squares problem of the form

$$u^*(x) = \arg \min_u \mathcal{L}(x, u) = \frac{1}{2} \|x - Bu\|_2^2 + \lambda g(u) , \quad (4.30)$$

where g is proper and convex, as depicted in Figure 4.2-2. By unrolling this architecture with T layers, we obtain a network $\phi_{\Theta(T)}(x) = u^{(T)}$ – illustrated in Figure 4.2-3 – with

parameters $\Theta^{(T)} = \{W_x^{(t)}, W_u^{(t)}, \mu^{(t)}\}_{t=1}^T$, defined by the following recursion

$$u^{(0)} = B^\dagger x; \quad u^{(t)} = \text{prox}_{\mu^{(t)}g}(W_x^{(t)}x + W_u^{(t)}u^{(t-1)}) . \quad (4.31)$$

As underlined by Equation 4.4, a good estimate $u^{(0)}$ is crucial in order to have a fast convergence toward $u^*(x)$. However, this chosen initialization is mitigated by the first layer of the network which learns to set a good initial guess for $u^{(1)}$. For a network with T layers, one recovers exactly the T -th iteration of PGD if the weights are chosen constant equal to

$$W_x^{(t)} = \frac{1}{\rho}B^\top, \quad W_u^{(t)} = (\text{Id} - \frac{1}{\rho}B^\top B), \quad \mu^{(t)} = \frac{\lambda}{\rho}, \quad \text{with } \rho = \|B\|_2^2 . \quad (4.32)$$

In practice, this choice of parameters are used as initialization for a posterior training stage. In many practical applications, one is interested in minimizing the loss Equation 4.30 for a fixed B and a particular distribution over the space of x , \mathcal{P} . As a result, the goal of this training stage is to find parameters $\Theta^{(T)}$ that minimize the risk, or expected loss, $\mathbb{E}[\mathcal{L}(x, \phi_{\Theta^{(T)}}(x))]$ over \mathcal{P} . Since one does not have access to this distribution, and following an empirical risk minimization approach with a given training set $\{x_1, \dots, x_N\}$ (assumed sampled *i.i.d* from \mathcal{P}), the network is trained by minimizing

$$\min_{\Theta^{(T)}} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(x_i, \phi_{\Theta^{(T)}}(x_i)) . \quad (4.33)$$

Note that when $T \rightarrow +\infty$, the presented initialization in Equation 4.32 gives a global minimizer of the loss for all x_i , as the network converges to exact PGD. When T is fixed, however, the output of the network is not a minimizer of Equation 4.30 in general. Minimizing this empirical risk can therefore find a weight configuration that reduces the sub-optimality of the network relative to Equation 4.30 over the input distribution used to train the network. In such a way, the network learns an algorithm to approximate the solution of Equation 4.30 for a particular class or distributions of signals. It is important to note here that while this procedure can accelerate the resolution the problem, the learned algorithm will only be valid for inputs x_i coming from the same input distribution \mathcal{P} as the training samples. The algorithm might not converge for samples which are too different from the training set, unlike the iterative algorithm which is guaranteed to converge for any sample.

This network architecture design can be directly applied to TV regularized problems if the synthesis formulation Equation 4.5 is used. Indeed, in this case PGD reduces to the ISTA algorithm, with $B = AL$ and $\text{prox}_{\mu g} = \text{st}(\cdot, \mu)$ becomes simply a soft-thresholding operator (which is only applied on the coordinates $\{2, \dots, k\}$, following the definition of R). However, as discussed in Proposition 4.2.2, the conditioning of the synthesis problem makes the estimation of the solution slow, increasing the number of network layers needed to get a good estimate of the solution. In the next section, we will extend these learning-based ideas directly to the analysis formulation by deriving a way to obtain exact and approximate expressions for the sub-differential of the non-separable prox-TV.

4.2.2 Back-propagating through TV proximal operator

Our two approaches to define learnable networks based on PGD for TV-regularized problems in the analysis formulation differ on the computation of the prox-TV and its derivatives. Our first approach consists in directly computing the weak derivatives of the exact proximal operator while the second one uses a differentiable approximation.




Figure 4.2-2 – **PGD** - Recurrent Neural Network




Figure 4.2-3 – **LPGD** - Unfolded network for Learned PGD with $T = 3$

Figure 4.2-4 – **Algorithm Unrolling** - Neural network representation of iterative algorithms. The parameters $\Theta^{(t)} = \{W_x^{(t)}, W_u^{(t)}, \mu^{(t)}\}$ can be learned by minimizing the loss Equation 4.33 to approximate good solution of Equation 4.30 on average.

Derivative of prox-TV

While there is no analytic solution to the prox-TV, it can be computed exactly (numerically) for 1D problems using the taut-string algorithm [Condat, 2013a]. This operator can thus be applied at each layer of the network, reproducing the architecture described in Figure 4.2-3. We define the LPGD-Taut network $\phi_{\Theta^{(T)}}(x)$ with the following recursion formula

$$\phi_{\Theta^{(T)}}(x) = \text{prox}_{\mu^{(T)} \|\cdot\|_{TV}} \left(W_x^{(T)} x + W_u^{(T)} \phi_{\Theta^{(T-1)}}(x) \right) \quad (4.34)$$

To be able to learn the parameters through gradient descent, one needs to compute the derivatives of Equation 4.33 w.r.t the parameters $\Theta^{(T)}$. Denoting $h = W_x^{(t)} x + W_u^{(t)} \phi_{\Theta^{(t-1)}}(x)$ and $u = \text{prox}_{\mu^{(t)} \|\cdot\|_{TV}}(h)$, the application of the chain rule (as implemented efficiently by automatic differentiation) results in

$$\frac{\partial \mathcal{L}}{\partial h} = J_x(h, \mu^{(t)})^\top \frac{\partial \mathcal{L}}{\partial u}, \quad \text{and} \quad \frac{\partial \mathcal{L}}{\partial \mu^{(t)}} = J_\mu(h, \mu^{(t)})^\top \frac{\partial \mathcal{L}}{\partial u}, \quad (4.35)$$

where $J_x(h, \mu) \in \mathbb{R}^{k \times k}$ and $J_\mu(h, \mu) \in \mathbb{R}^{k \times 1}$ denotes the weak Jacobian of the output of the proximal operator u with respect to the first and second input respectively. We now give the analytic formulation of these weak Jacobians in the following proposition.

Proposition 4.2.4. [Weak Jacobian of prox-TV] Let $x \in \mathbb{R}^k$ and $u = \text{prox}_{\mu \|\cdot\|_{TV}}(x)$, and denote by \mathcal{S} the support of $z = \tilde{D}u$. Then, the weak Jacobian J_x and J_μ of the prox-TV relative to x and μ can be computed as

$$J_x(x, \mu) = L_{:, \mathcal{S}} (L_{:, \mathcal{S}}^\top L_{:, \mathcal{S}})^{-1} L_{:, \mathcal{S}}^\top \quad \text{and} \quad J_\mu(x, \mu) = -L_{:, \mathcal{S}} (L_{:, \mathcal{S}}^\top L_{:, \mathcal{S}})^{-1} \text{sign}(Du)_{\mathcal{S}}$$

First, we recall Lemma 4.2.5 to weakly derive the soft-thresholding.

Lemma 4.2.5 (Weak derivative of the soft-thresholding; Deledalle et al. 2014). The soft-thresholding operator $ST : \mathbb{R} \times \mathbb{R}_+ \mapsto \mathbb{R}$ defined by $ST(t, \tau) = \text{sign}(t)(|t| - \tau)_+$ is weakly differentiable with weak derivatives

$$\frac{\partial st}{\partial t}(t, \tau) = \mathbb{1}_{\{|t| > \tau\}}, \quad \text{and} \quad \frac{\partial st}{\partial \tau}(t, \tau) = -\text{sign}(t) \cdot \mathbb{1}_{\{|t| > \tau\}},$$

where

$$\mathbb{1}_{\{|t| > \tau\}} = \begin{cases} 1, & \text{if } |t| > \tau, \\ 0, & \text{otherwise.} \end{cases}$$

A very important remark here is to notice that if one denote $z = st(t, \tau)$, one can rewrite these weak derivatives as

$$\frac{\partial st}{\partial t}(t, \tau) = \mathbb{1}_{\{|z|>0\}} \quad , \quad \text{and} \quad \frac{\partial st}{\partial \tau}(t, \tau) = -\text{sign}(z) \cdot \mathbb{1}_{\{|z|>0\}} \quad . \quad (4.36)$$

Indeed, when $|t| > \tau$, $|z| = |t| - \tau > 0$ and conversely, $|z| = 0$ when $|t| < \tau$. Moreover, when $|t| > \tau$, we have $\text{sign}(t) = \text{sign}(z)$ and thus the two expressions for $\frac{\partial st}{\partial \tau}$ match.

Using this [Lemma 4.2.5](#), we now give the proof of [Proposition 4.2.4](#).

Proof. The proof is inspired from the proof from [Bertrand et al. \[2020, Proposition 1\]](#). We denote $u(x, \mu) = \text{prox}_{\mu\|\cdot\|_{TV}}(x)$, hence $u(x, \mu)$ is defined by

$$u(x, \mu) = \arg \min_{\hat{u}} \frac{1}{2} \|x - \hat{u}\|_2^2 + \mu \|\hat{u}\|_{TV} \quad (4.37)$$

Equivalently, as we have seen previously in [Equation 4.5](#), using the change of variable $\hat{u} = L\hat{z}$ and minimizing over \hat{z} gives

$$\min_{\hat{z}} \frac{1}{2} \|x - L\hat{z}\|_2^2 + \mu \|R\hat{z}\|_1 \quad . \quad (4.38)$$

We denote by $z(x, \mu)$ the minimizer of the previous equation. Thus, the solution $u(x, \mu)$ of the original problem [Equation 4.37](#) can be recovered using $u(x, \mu) = Lz(x, \mu)$. Iterative PGD can be used to solve [Equation 4.38](#) and $z(x, \mu)$ is a fixed point of the iterative procedure. That is to say the solution z verifies

$$\begin{cases} z_1(x, \mu) = z_1(x, \mu) - \frac{1}{\rho} (L^\top (Lz(x, \mu) - x))_1 \quad , \\ z_i(x, \mu) = st \left(z_i(x, \mu) - \frac{1}{\rho} (L^\top (Lz(x, \mu) - x))_i, \frac{\mu}{\rho} \right) \quad \text{for } i = 2 \dots k \quad . \end{cases} \quad (4.39)$$

Using the result from [Lemma 4.2.5](#), we can differentiate [Equation 4.39](#) and obtain the following equation for the weak Jacobian $\hat{J}_x(x, \mu) = \frac{\partial z}{\partial x}(x, \mu)$ of $z(x, \mu)$ relative to x

$$\hat{J}_x(x, \mu) = \begin{pmatrix} 1 \\ \mathbb{1}_{\{|z_2(x, \mu)|>0\}} \\ \vdots \\ \mathbb{1}_{\{|z_k(x, \mu)|>0\}} \end{pmatrix} \odot \left[\left(\text{Id} - \frac{1}{\rho} L^\top L \right) \hat{J}_x(x, \mu) + \frac{1}{\rho} L^\top \text{Id} \right] \quad . \quad (4.40)$$

Identifying the non-zero coefficient in the indicator vectors yields

$$\begin{cases} \hat{J}_{x, \mathcal{S}^c}(x, \mu) = 0 \\ \hat{J}_{x, \mathcal{S}}(x, \mu) = \left(\text{Id} - \frac{1}{\rho} L_{:, \mathcal{S}}^\top L_{:, \mathcal{S}} \right) \hat{J}_{x, \mathcal{S}}(x, \mu) + \frac{1}{\rho} L_{:, \mathcal{S}}^\top \quad . \end{cases} \quad (4.41)$$

As, L is invertible, so is $L_{:, \mathcal{S}}^\top L_{:, \mathcal{S}}$ for any support \mathcal{S} and solving the second equations yields the following

$$\hat{J}_{x, \mathcal{S}} = (L_{:, \mathcal{S}}^\top L_{:, \mathcal{S}})^{-1} L_{:, \mathcal{S}}^\top \quad (4.42)$$

Using $u = Lz$ and the chain rules yields the expecting result for the weak Jacobian relative to x , noticing that as $\hat{J}_{x, \mathcal{S}^c} = 0$, $L\hat{J}_x = L_{:, \mathcal{S}}\hat{J}_{x, \mathcal{S}}$.

Similarly, concerning, $\hat{J}_\mu(x, \mu)$, we use the result from [Lemma 4.2.5](#) an differentiale [Equation 4.39](#) and obtain $\hat{J}_\mu(x, \mu) = \frac{\partial z}{\partial \mu}(x, \mu)$ of $z(x, \mu)$ relative to μ

$$\widehat{J}_\mu(x, \mu) = \begin{pmatrix} 1 \\ \mathbb{1}_{\{|z_2(x, \mu)| > 0\}} \\ \vdots \\ \mathbb{1}_{\{|z_k(x, \mu)| > 0\}} \end{pmatrix} \odot \left[\left(\text{Id} - \frac{1}{\rho} L^\top L \right) \widehat{J}_\mu(x, \mu) \right] + \frac{1}{\rho} \begin{pmatrix} 1 \\ -\text{sign}(z_2(x, \mu)) \mathbb{1}_{\{|z_2(x, \mu)| > 0\}} \\ \vdots \\ -\text{sign}(z_k(x, \mu)) \mathbb{1}_{\{|z_k(x, \mu)| > 0\}} \end{pmatrix}. \quad (4.43)$$

Identifying the non-zero coefficient in the indicator vectors yields

$$\begin{cases} \widehat{J}_{\mu, \mathcal{S}^c}(x, \mu) &= 0 \\ \widehat{J}_{\mu, \mathcal{S}}(x, \mu) &= \widehat{J}_{\mu, \mathcal{S}^c}(x, \mu) - \frac{1}{\rho} L_{:, \mathcal{S}}^\top L_{:, \mathcal{S}} \widehat{J}_{\mu, \mathcal{S}^c}(x, \mu) - \frac{1}{\rho} \text{sign}(z_{\mathcal{S}}(x, \mu)). \end{cases} \quad (4.44)$$

As previous, solving the second equation yields the following

$$\widehat{J}_{\mu, \mathcal{S}} = -(L_{:, \mathcal{S}}^\top L_{:, \mathcal{S}})^{-1} \text{sign}(z_{\mathcal{S}}(x, \mu)) \quad (4.45)$$

Using $u = Lz$ and the chain rules yields the expecting result for the weak Jacobian relative to μ , noticing that as $\widehat{J}_{\mu, \mathcal{S}^c} = 0$.

□

The proof of this proposition can be found in [Equation 4.2.2](#). Note that the dependency in the inputs is only through \mathcal{S} and $\text{sign}(Du)$, where u is a short-hand for $\text{prox}_{\mu \|\cdot\|_{TV}}(x)$. As a result, computing these weak Jacobians can be done efficiently by simply storing $\text{sign}(Du)$ as a mask, as it would be done for a RELU or the soft-thresholding activations, and requiring just $2(k-1)$ bits. With these expressions, it is thus possible to compute gradient relatively to all parameters in the network, and employ them via back-propagation.

4.2.3 Unrolled prox-TV

As an alternative to the previous approach, we propose to use the LISTA network to approximate the prox-TV [Equation 4.3](#). The prox-TV can be reformulated with a synthesis approach resulting in a Lasso *i.e.*

$$z^* = \arg \min_z \frac{1}{2} \|h - Lz\|_2^2 + \mu \|Rz\|_1 \quad (4.46)$$

The proximal operator solution can then be retrieved with $\text{prox}_{\mu \|\cdot\|_{TV}}(h) = Lz^*$. This problem can be solved using ISTA, and approximated efficiently with a LISTA network [Gregor and Le Cun \[2010\]](#). For the resulting architecture – dubbed LPGD-LISTA – $\text{prox}_{\mu \|\cdot\|_{TV}}(h)$ is replaced by a nested LISTA network with a fixed number of layers T_{in} defined recursively with $z^{(0)} = Dh$ and

$$z^{(\ell+1)} = \text{st} \left(W_z^{(\ell, t)} z^{(\ell)} + W_h^{(\ell, t)} \Phi_{\Theta^{(t)}}, \frac{\mu^{(\ell, t)}}{\rho} \right). \quad (4.47)$$

Here, $W_z^{(\ell, t)}, W_h^{(\ell, t)}, \mu^{(\ell, t)}$ are the weights of the nested LISTA network for layer ℓ . They are initialized with weights chosen as in [Equation 4.32](#) to ensure that the initial state approximates the prox-TV. Note that the weights of each of these inner layers are also learned through back-propagation during training.

The choice of this architecture provides a differentiable (approximate) proximal operator. Indeed, the LISTA network is composed only of linear and soft-thresholding layers – standard tools for deep-learning libraries. The gradient of the network’s parameters can thus be computed using classic automatic differentiation. Moreover, if the inner network is not trained, the gradient computed with this method will converge toward the gradient computed using [Proposition 4.2.4](#) as T_{in} goes to ∞ , see the next [Proposition 4.2.6](#).

Proposition 4.2.6. *Linear convergence of the weak Jacobian* We consider the mapping $z^{(T_{in})} : \mu\mathbb{R}^k \times \mathbb{R}_+ \mapsto \mathbb{R}^k$ defined where $z^{(T_{in})}(x)$ is defined by recursion

$$z^{(t)}(x, \mu) = ST(z^{(t-1)}(x, \mu) - \frac{1}{\|L\|_2^2} L^\top (Lz^{(t-1)}(x, \mu) - x), \frac{\mu}{\|L\|_2^2}). \quad (4.48)$$

Then the weak $\mathcal{J}_x = L \frac{\partial z^{(T_{in})}}{\partial x}$ and $\mathcal{J}_\mu = L \frac{\partial z^{(T_{in})}}{\partial \mu}$ of this mapping relative to the inputs x and μ converges linearly toward the weak Jacobian J_x and J_μ of $\text{prox}_{\mu\|\cdot\|_{TV}}(x)$ defined in [Proposition 4.2.4](#).

This mapping defined in [Equation 4.48](#) corresponds to the inner network in LPGD-LISTA when the weights of the network have not been learned.

Proof. As L is invertible, problem [Equation 4.38](#) is strongly convex and have a unique solution. We can thus apply the result from [Bertrand et al. \[2020, Proposition 2\]](#) which shows the linear convergence of the weak Jacobian $\hat{\mathcal{J}}_x = \frac{\partial z^{(T_{in})}}{\partial x}$ and $\hat{\mathcal{J}}_\mu = \frac{\partial z^{(T_{in})}}{\partial \mu}$ for ISTA toward \hat{J}_x and \hat{J}_μ of the synthesis formulation of the prox. Using the linear relationship between the analysis and the synthesis formulations yields the expected result. \square

Thus, in this untrained setting with infinitely many inner layers, the network is equivalent to LPGD-Taut as the output of the layer also converges toward the exact proximal operator.

Connections to inexact PGD A drawback of approximating the prox-TV via an iterative procedure is, precisely, that it is not exact. This optimization error results from a trade-off between computational cost and convergence rate. Using results from [Machart et al. \[2012\]](#), one can compute the scaling of T and T_{in} to reach an error level of δ with an untrained network.

Proposition 4.2.7. [*Scaling of T and T_{in} w.r.t the error level δ*] Let δ the error defined such as $P_x(u^{(T)}) - P_x(u^*) \leq \delta$.

We suppose there exists some constants $C_0 \geq \|u^{(0)} - u^*\|_2$ and $C_1 \geq \max_\ell \|u^{(\ell)} - \text{prox}_{\frac{\mu}{\rho}}(u^{(\ell)})\|_2$. Then, T the number of layers for the global network and T_{in} the inner number of layers for the prox-TV scale are given by

$$T_{in} = \frac{\ln \frac{1}{\delta} + \ln 6\sqrt{2\rho}C_1}{\ln \frac{1}{1-\gamma}} \quad \text{and} \quad T = \frac{2\rho C_0^2}{\delta}$$

with ρ defined as in [Equation 4.2](#)

Proof. As discussed by [Machart et al. \[2012\]](#), the global convergence rate of inexact PGD with T_{in} inner iteration is given by

$$P_x(u^{(T)}) - P_x(u^*) \leq \frac{\rho}{2T} \left(\|u^{(0)} - u^*\|_2 + 3 \sum_{\ell=1}^T \sqrt{\frac{2(1-\gamma)^{T_{in}} \|u^{(\ell-1)} - \text{prox}_{\frac{\mu}{\rho}}(u^{(\ell-1)})\|_2^2}{\rho}} \right)^2, \quad (4.49)$$

where γ is the condition number for L i.e. $\frac{\cos(\frac{\pi}{2k+1})}{\sin(\frac{\pi}{2k+1})}$.

We are looking for minimal parameters T and T_{in} such that the error bound in Equation 4.49 is below a certain error level δ .

We consider the case where there exists some constants $C_0 \geq \|u^{(0)} - u^*\|_2$ and $C_1 \geq \max_\ell \|u^{(\ell)} - \text{prox}_{\frac{\mu}{\rho}}(u^{(\ell)})\|_2$ upper bounding how far the initialization can be compared to the result of the global problem and the sub-problems respectively.

We denote $\alpha_1 = 3\sqrt{\frac{2}{\rho}}C_1$. The right hand side of Equation 4.49 can be upper bounded by as

$$\begin{aligned} \frac{\rho}{2T} \left(\|u^{(0)} - u^*\|_2 + 3 \sum_{\ell=1}^T \sqrt{\frac{2(1-\gamma)^{T_{in}} \|u^{(\ell-1)} - \text{prox}_{\frac{\mu}{\rho}}(u^{(\ell-1)})\|_2^2}{\rho}} \right)^2 \\ \leq \frac{\rho}{2T} \left(C_0 + \alpha_1 T (1-\gamma)^{T_{in}/2} \right)^2 \end{aligned} \quad (4.50)$$

Then, we are looking for T, T_{in} such that this upper bound is lower than δ , i.e.

$$\frac{\rho}{2T} \left(C_0 + \alpha_1 T (1-\gamma)^{T_{in}/2} \right)^2 \leq \delta \quad (4.51)$$

$$\Leftrightarrow \left(C_0 + \alpha_1 T (1-\gamma)^{T_{in}/2} \right)^2 - \frac{2\delta}{\rho} T \leq 0 \quad (4.52)$$

$$\Leftrightarrow \left(C_0 + \alpha_1 T (1-\gamma)^{T_{in}/2} - \sqrt{\frac{2\delta}{\rho}} \sqrt{T} \right) \underbrace{\left(B + \alpha_1 T (1-\gamma)^{T_{in}/2} + \sqrt{\frac{2\delta}{\rho}} \sqrt{T} \right)}_{\geq 0} \leq 0 \quad (4.53)$$

$$\Leftrightarrow C_0 + \alpha_1 T (1-\gamma)^{T_{in}/2} - \sqrt{\frac{2\delta}{\rho}} \sqrt{T} \leq 0 \quad (4.54)$$

$$(4.55)$$

Denoting $\alpha_2 = \sqrt{\frac{2\delta}{\rho}}$ and $X = \sqrt{T}$, we get the following function of X and T_{in}

$$f(X, T_{in}) = \alpha_1 (1-\gamma)^{T_{in}/2} X^2 - \alpha_2 X + C_0 \quad (4.56)$$

The inequality $f(X, T_{in}) \leq 0$ has a solution if and only if $\alpha_2^2 - 4C_0\alpha_1(1-\gamma)^{T_{in}/2} \geq 0$ i.e.

$$T_{in} \geq 2 \frac{\ln \frac{\alpha_2^2}{4\alpha_1 C_0}}{\ln 1-\gamma}$$

Taking the minimal value for T_{in} i.e. $T_{in} = 2 \frac{\ln \frac{\alpha_2^2}{4\alpha_1 C_0}}{\ln 1-\gamma} = \frac{\ln \frac{1}{\delta} + \ln 6\sqrt{2\rho}C_1}{\ln \frac{1}{1-\gamma}}$ yields

$$f(X, T_{in}) = \frac{\alpha_2^2}{4C_0} X^2 - \alpha_2 X + C_0 = \frac{\alpha_2^2}{4C_0} \left(X - \frac{2C_0}{\alpha_2} \right)^2$$

for $X = \frac{2C_0}{\alpha_2} = \frac{\sqrt{2\rho}C_0}{\sqrt{\delta}}$ i.e. $T = \frac{2\rho C_0^2}{\delta}$. \square

Proposition 4.2.7 shows that without learning, T should scale as $O(\frac{1}{\delta})$ and T_{in} should be larger than $O(\ln(\frac{1}{\delta}))$. This scaling gives potential guidelines to set these parameters, as one can expect that learning the parameters of the network would reduce these requirement.

4.3 Performance comparison

All experiments are performed in Python using PyTorch [Paszke et al., 2019]. We used the implementation¹ of Barbero and Sra [2018] to compute TV proximal operator using taut-string algorithm. The code to reproduce the figures is available online².

In all experiments, we initialize $u_0 = A^\dagger x$. Moreover, we employed a normalized λ_{reg} as a penalty parameter: we first compute the value of λ_{max} (which is the minimal value for which $z = 0$ is solution of Equation 4.5) and we refer to λ as the ratio so that $\lambda_{reg} = \lambda \lambda_{max}$, with $\lambda \in [0, 1]$.

Definition of λ_{max} The definition of λ_{max} is the smallest value for the regularisation parameter λ such that the solution of the TV-regularized problem is constant. This corresponds to the definition of λ_{max} in the Lasso, which is the smallest regularisation parameter such that 0 is solution. We here derive its analytic value which is used to rescale all experiments. This is important to define an equiregularisation set for the training and testing samples, to have a coherent and generalizable training.

Proposition 4.3.1. *The value of λ_{max} for the TV-regularized problem is*

$$\lambda_{max} = \|A^\top (Ac\mathbf{1} - x)\|_\infty$$

where $c = \frac{\sum_{i=1}^p S_i x_i}{\sum_{i=1}^p S_i^2}$ and $S_i = \sum_{j=1}^k A_{i,j}$.

Proof. We first derive the constant solution of the ℓ_2 -regression problem associated to Equation 4.1. For $c \in \mathbb{R}$, we consider a constant vector $c\mathbf{1}$. The best constant solution for the ℓ_2 -regression problem is obtained by solving

$$\min_{c \in \mathbb{R}} f_x(c) = \frac{1}{2} \|x - cA\mathbf{1}\|_2^2. \quad (4.57)$$

The first order optimality condition in c reads

$$\nabla f_x(c) = \sum_{i=1}^n \left(\sum_{j=1}^k A_{i,j} \right) (c \sum_{j=1}^k A_{i,j} - x_i) = \sum_{i=1}^n S_i (cS_i - x_i) = 0, \quad (4.58)$$

and thus $c = \frac{\sum_{i=1}^p S_i x_i}{\sum_{i=1}^p S_i^2}$.

Then, we look at the conditions on λ to ensure that the constant solution $c\mathbf{1}$ is solution of the regularized problem. The first order conditions on the regularized problem reads

$$0 \in \partial P_x(c\mathbf{1}) = A^\top (Ac\mathbf{1} - x) + \lambda \partial \|Dc\mathbf{1}\|_1 \quad (4.59)$$

Next, we develop the previous equality:

$$\forall j \in \{2, \dots, k\}, \quad A_j^\top (Ac\mathbf{1} - x) \in \lambda \partial (\|Dc\mathbf{1}\|_1)_j = [-\lambda, \lambda] \quad \text{since } Dc\mathbf{1} = 0 \quad (4.60)$$

Thus, the constrains are all satisfied for $\lambda \geq \lambda_{max}$, with $\lambda_{max} = \|A^\top (Ac\mathbf{1} - x)\|_\infty$ and as c is solution for the unregularized problem reduced to a constant, $c\mathbf{1}$ is solution of the TV-regularized problem for all $\lambda \geq \lambda_{max}$. □

¹Available at <https://github.com/albarji/proxTV>

²Available at <https://github.com/hcherkaoui/carpet>.

As the computational complexity of all compared algorithms is the same except for the proximal operator, we compare them in term of iterations.

Optimization algorithm for training In our experiments, all networks are trained using Gradient Descent (GD) with back-tracking line search. The gradients are computed using automatic differentiation in Pytorch [Paszke et al., 2019] for most layers and the weak Jacobian proposed in section 4.2.2 for the back-propagation through the prox-TV. The learning is stopped once a step-size of $\eta_{limit} = 10^{-20}$ is reached in the back-tracking step. For LPGD-LISTA, the weights of the inner LISTA computing the prox-TV are trained jointly with the parameters of the outer unrolled PGD.

Weight initialization All layers for an unrolled algorithm are initialized using the values of weights in Equation 4.32 that ensure the output of the layer with T layers corresponds to the output of T iterations of the original algorithm. To further stabilize the training, we use a layer-wise approach. When training a network with $T_1 + T_2$ layers after having trained a network with T_1 layers, the first T_1 layers in the new network are initialized with the weights of the one trained previously, and the remaining layers are initialized using weights value from Equation 4.32. This ensures that the initial value of the loss for the new network is smaller than the one from the shallower one if the unrolled algorithm is monotonous (as it is the case for PGD).

4.3.1 Simulation

We generate $n = 2000$ times series and used half for training and other half for testing and comparing the different algorithms. We train all the network’s parameters jointly – those to approximate the gradient for each iteration along with those to define the inner proximal operator. We set the length of the source signals $(u_i)_{i=1}^n \in \mathbb{R}^{n \times k}$ to $k = 8$ with a support of $|S| = 2$ non-zero coefficients (larger dimensions will be showcased in the real data application). We generate $A \in \mathbb{R}^{m \times k}$ as a Gaussian matrix with $m = 5$, obtaining then $(u_i)_{i=1}^n \in \mathbb{R}^{n \times p}$. Moreover, we add Gaussian noise to measurements $x_i = Au_i$ with a signal to noise ratio (SNR) of 1.0.

We compare our proposed methods, LPGD-Taut network and the LPGD-LISTA with $T_{in} = 50$ inner layers to PGD and Accelerated PGD with the analysis formulation. For completeness, we also add the FISTA algorithm for the synthesis formulation in order to illustrate Proposition 4.2.2 along with its learned version.

Figure 4.3-5 presents the risk (or expected function value, P) of each algorithm as a function of the number of layers or, equivalently, iterations. For the learned algorithms, the curves in t display the performances of a network with t layer trained specifically. We observe that all the synthesis formulation algorithms are slower than their analysis counterparts, empirically validating Proposition 4.2.2. Moreover, both of the proposed methods accelerate the resolution of Equation 4.1 in a low iteration regime. However, when the regularization parameter is high ($\lambda = 0.8$), we observe that the performance of the LPGD-LISTA tends to plateau. It is possible that such a high level of sparsity require more than 50 layers for the inner network (which computes the prox-TV). According to section 4.2.3, the error associated with this proximity step hinders the global convergence, making the loss function decrease slowly. Increasing the number of inner layers would alleviate this issue, though at the expense

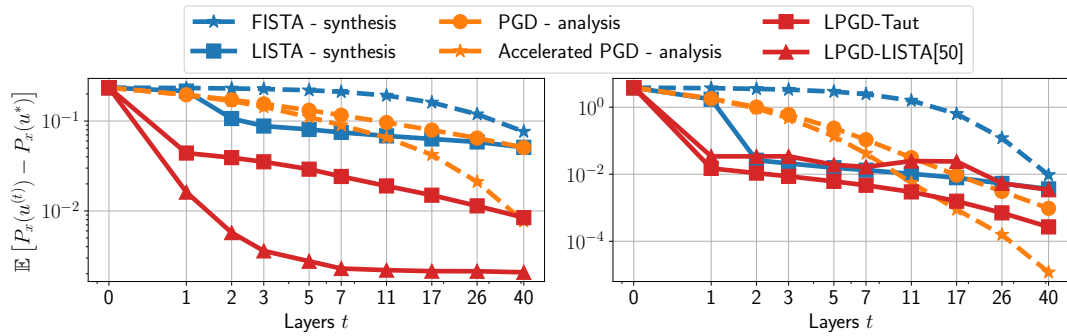


Figure 4.3-5 – **Performance comparison** for different regularisation levels (*left*) $\lambda = 0.1$, (*right*) $\lambda = 0.8$. We see that synthesis formulations are outperformed by the analysis counterpart. Both our methods are able to accelerate the resolution of Equation 4.1, at least in the first iterations.




Figure 4.3-6 – **Proximal operator error comparison** for different regularisation levels (*left*) $\lambda = 0.1$, (*right*) $\lambda = 0.8$. We see that learn the trained unrolled prox-TV barely improve the performance. More interestingly, in a high sparsity context, after a certain point, the error sharply increase.

of increased computational burden for both training and runtime. For LPGD-Taut, while the Taut-string algorithm ensures that the recovered support is exact for the proximal step, the overall support can be badly estimated in the first iterations. This can lead to un-informative gradients as they greatly depend on the support of the solution in this case, and explain the reduced performances of the network in the high sparsity setting.

Inexact prox-TV With the same data $(x_i)_{i=1}^n \in \mathbb{R}^{n \times m}$, we empirically investigate the error of the prox-TV $\epsilon_k^{(t)} = F_{u^{(t)}}(z^{(t)}) - F_{u^{(t)}}(z^*)$ and evaluate it for c with different number of layers ($T \in [20, 50]$). We also investigate the case where the parameter of the nested LISTA in LPGD-LISTA are trained compared to their initialization in untrained version.

Figure 4.3-6 depicts the error ϵ_k for each layer. We see that learning the parameters of the unrolled prox-TV in LPGD-LISTA barely improves the performance. More interestingly, we observe that in a high sparsity setting the error sharply increases after a certain number of layers. This is likely cause by the high sparsity of the estimates, the small numbers of iterations of the inner network (between 20 and 50) are insufficient to obtain an accurate solution to the proximal operator. This is in accordance with inexact PGD theory which predict that such algorithm has no exact convergence guarantees [Schmidt et al., 2011].

4.3.2 fMRI data deconvolution

Functional magnetic resonance imaging (fMRI) is a non-invasive method for recording the brain activity by dynamically measuring blood oxygenation level-dependent (BOLD) contrast, denoted here x . The latter reflects the local changes in the deoxyhemoglobin concentration in the brain [Ogawa et al. \[1992\]](#) and thus indirectly measures neural activity through the neurovascular coupling. This coupling is usually modelled as a linear and time-invariant system and characterized by its impulse response, the so-called haemodynamic response function (HRF), denoted here h . Recent developments propose to estimate either the neural activity signal independently [[Karahanoğlu et al., 2013](#), [Cherkaoui et al., 2019b](#)] or jointly with the HRF [[Cherkaoui et al., 2019a](#), [Farouj et al., 2019](#)]. Estimating the neural activity signal with a fixed HRF is akin to a deconvolution problem regularized with TV-norm,

$$\min_{u \in \mathbb{R}^k} P(u) = \frac{1}{2} \|h * u - x\|_2^2 + \lambda \|u\|_{TV} \quad (4.61)$$

To demonstrate the usefulness of our approach with real data, where the training set has not the exact same distribution than the testing set, we compare the LPGD-Taut to Accelerated PGD for the analysis formulation on this deconvolution problem. We choose two subjects from the UK Bio Bank (UKBB) dataset [[Sudlow et al., 2015](#)], perform the usual fMRI processing and reduce the dimension of the problem to retain only 8000 time-series of 250 time-frames, corresponding to a record of 3 minute 03 seconds. We train the LPGD taut-string network solver on the first subject and [Figure 4.3-9](#) reports the performance of the two algorithms on the second subject for $\lambda = 0.1$.

For this experiment, we investigate the 6 min long rs-fMRI acquisition (TR=0.735 s) from the UK Bio Bank dataset [[Sudlow et al., 2015](#)]. The following pre-processing steps were applied on the images: motion correction, grand-mean intensity normalisation, high-pass temporal filtering, Echo planar imaging unwarping, Gradient Distortion Correction unwarping and structured artefacts removal by Independent Components Analysis. More details on the processing pipeline can found in [Alfaro-Almagro et al. \[2018\]](#).

On top of this preprocessing, we perform a standard fMRI preprocessing proposed in the python package Nilearn³. This standard pipeline includes to detrend the data, standardize it and filter high and low frequencies to reduce the presence of noise.

The performance is reported relatively to the number of iteration as the computational complexity of each iteration or layer for both methods is equivalent. It is clear that LPGD-Taut converges faster than the Accelerated PGD even on real data. In particular, acceleration is higher when the regularization parameter λ is smaller. As mentioned previously, this acceleration is likely to be caused by the better learning capacity of the network in a low sparsity context.

4.4 Conclusion

This chapter studies the optimization of TV-regularized problems via learned PGD. We demonstrated, both analytically and numerically, that it is better to address these problems in their original analysis formulation rather than resorting to the simpler (alas slower) synthesis version. We then proposed two different algorithms that allow for the efficient computation

³<https://nilearn.github.io>

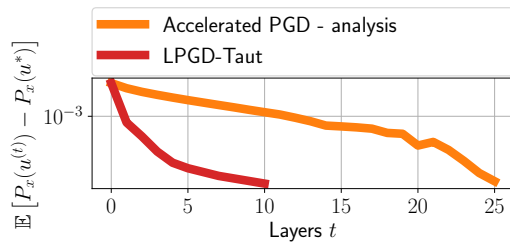
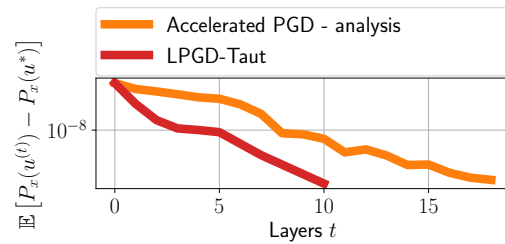
Figure 4.3-7 – $\lambda = 0.1\lambda_{\max}$ Figure 4.3-8 – $\lambda = 0.8\lambda_{\max}$

Figure 4.3-9 – **Performance comparison** ($\lambda = 0.1$ *left* - $\lambda = 0.8$ *right*) between our analytic prox-TV derivative method and the PGD in the analysis formulation for the HRF deconvolution problem with fMRI data. Our proposed method outperform the FISTA algorithm in the analysis formulation. We notice a slight degradation of the acceleration in this high sparsity context.

and derivation of the required prox-TV, exactly or approximately. Our experiments on synthetic and real data demonstrate that our learned networks for prox-TV provide a significant advantage in terms of convergence speed.

Finally, the principles presented in this chapter could be generalized and deployed to other optimization problems, involving not only the TV-norm but more general analysis-type priors. In particular, this chapter only focused on 1D TV problems because the equivalence between Lasso and TV is not exact in higher dimension. In this case, we believe that a dual formulation [Chambolle, 2004] is a key ingredient to tackle this problem and derive similar learnable algorithms.

* * *
* *
*
*

Contribution to neurovascular disentangling

Chapter Outline

5.1	Sparsity-based blind deconvolution of neural activation signal in fMRI . . .	76
5.1.1	Introduction	76
5.1.2	HRF estimation with neural activation sparse model.	77
5.1.3	Numerical Experiments	78
5.1.4	Interim summary	81
5.2	fMRI BOLD signal decomposition using a multivariate low-rank model . . .	81
5.2.1	Introduction	81
5.2.2	Low rank decomposition of the BOLD signal	82
5.2.3	Numerical experiments	85
5.2.4	Summary	87
5.3	Conclusion	88

Parts of this chapter were presented in two International conferences with proceedings:

H. Cherkaoui, T. Moreau, A. Halimi, and P. Ciuciu. Sparsity-based Semi-Blind Deconvolution of Neural Activation Signal in fMRI. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1323–1327, 2019a.

H. Cherkaoui, T. Moreau, A. Halimi, and P. Ciuciu. fMRI BOLD signal decomposition using a multivariate low-rank model. In *European Signal Processing Conference (EUSIPCO)*, pages 1–5, 2019b.

THIS chapter proposes two contributions to the estimation of the HRF from fMRI data. We first propose a univariate *semi-blind deconvolution* approach to estimate a HRF voxelwise along with the associated neural activity. However, as this model is greedy in terms of parameters, we then propose a multivariate approach to deconvolve the neural activity using a fixed HRF shape. This is of course an intermediate step to ease the developments toward a multivariate the semi-blind deconvolution of fMRI data, which will be presented in the next chapter. This intermediate approach makes sense for a regional analysis of task-fMRI data, as illustrated in this chapter.

5.1 Sparsity-based blind deconvolution of neural activation signal in fMRI

5.1.1 Introduction

As explained in the [chapter 1](#), fMRI non-invasively records brain activity by dynamically measuring the blood oxygenation level-dependent (BOLD) contrast. The latter reflects the local changes in the deoxyhemoglobin concentration in the brain [[Ogawa et al., 1992](#)] and thus indirectly measures neural activity through the neurovascular coupling. This coupling is usually characterized as a linear and time-invariant system and thus summarized by its impulse response, the so called haemodynamic response function (HRF) [[Bandettini et al., 1993](#), [Boynton et al., 1996](#)]. Its estimation links the observed signal to the underlying neural activity, which can in turn be used to understand cognitive processes in the healthy brain or to predict neurological diseases. Some recent work proposes to estimate such a surrogate by estimating a block signal using a fixed HRF [[Karahanoğlu et al., 2013](#)]. In doing so, the recovered neural activity signal is used to define functional networks in which the population of neurons have been activated together at the same time. However, as the HRF is not allowed to vary across brain regions, this method potentially produces a biased estimate of the deconvolved neural activity signal.

Following the ideas developed in the dictionary learning literature [[Olshausen and Field, 1997](#)], we propose, for each BOLD time series, namely each voxel, to jointly estimate the neural activation signal and the HRF with properly selected constraints. The resulting optimization problem is non-convex but an approximated solution can be computed using alternate minimization, and we propose efficient procedures to perform each step. This algorithm aims at reducing the bias introduced with arbitrarily fixed HRF or EP, by learning the HRF for each voxel and a neural activity signal that can fluctuate and depart from the EP. Unlike previous contributions [[Sreenivasan et al., 2015](#), [Makni et al., 2005](#)] to make the inversion well-posed, we regularize the neural activity signature with a sparse prior on the first-order derivative and the HRF is parameterized by a single unknown scalar. On real task fMRI datasets, we show that we are able to recover similar effects to state-of-the-art HRF estimation approaches without the knowledge of the EP.

In the following, [subsection 5.1.2](#) introduces our model for the BOLD signal and our algorithm to estimate the HRF. Then, our technique is evaluated against state-of-the-art algorithm in [subsection 5.1.3](#). Conclusions and future work are discussed in [subsection 5.1.4](#).

5.1.2 HRF estimation with neural activation sparse model

In this section, we present our modeling of the BOLD signal and derive an efficient algorithm to estimate its parameters.

$\mathbf{D} \in \mathbb{R}^{\tilde{T} \times \tilde{T}}$ refers to the modified first-order differences operator and $\mathbf{L} \in \mathbb{R}^{\tilde{T} \times \tilde{T}}$ to the discrete integration operator:

$$\mathbf{L} = \begin{bmatrix} 1 & 0 & \dots & & \\ 1 & 1 & 0 & \dots & \\ 1 & 1 & 1 & 0 & \dots \\ \vdots & \ddots & \ddots & \ddots & \ddots \end{bmatrix} \quad \mathbf{D} = \begin{bmatrix} 1 & 0 & \dots & & \\ 1 & -1 & 0 & \dots & \\ 0 & \ddots & \ddots & \ddots & \dots \\ \vdots & \ddots & 0 & 1 & -1 \end{bmatrix}$$

Linear and time-invariant modeling

A common model for the voxelwise BOLD signal $y(t)$ is the linear and time-invariant model (LTI) [Boynton et al., 1996], where the signal is considered as the result of the convolution of a neural activation signal, denoted $u(t)$, with an HRF, here denoted $h(t)$: $y(t) = v(t) * u(t) + \epsilon(t)$ where $\epsilon(t)$ is an additive noise term. Typically, the HRF $v(t)$ has a restricted support in time and quantifies the neurovascular coupling in a specific brain region. The activation signal $u(t)$ captures the periods during which this particular region is involved in task performance.

In practice fMRI data are collected at a discrete sampling rate, called the time of repetition (TR), which typically varies between 1 and 2 s. Vector $\mathbf{y} \in \mathbb{R}^{\tilde{T}}$ thus refers to the BOLD signal measured in each voxel of the brain along \tilde{T} consecutive scans. The discretized LTI model reads: $\mathbf{y} = \mathbf{v} * \mathbf{u} + \boldsymbol{\epsilon}$ with $\mathbf{y}, \boldsymbol{\epsilon} \in \mathbb{R}^{\tilde{T}}$ and $\mathbf{v} \in \mathbb{R}^m$, m being the number of time-points for the HRF, typically smaller than \tilde{T} and spanning over about 20 s.

In task fMRI data, the activation signal is usually represented by the piecewise constant time course associated with the experimental design. A common way to enforce such structure in \mathbf{u} is to consider its first derivative $\mathbf{z} = \mathbf{D}\mathbf{u}$ to be sparse. To make the computations easier, we inject this prior information in the LTI model and re-parameterize it using \mathbf{u} and $\mathbf{u} = \mathbf{L}\mathbf{z}$:

$$\mathbf{y} = \mathbf{v} * \mathbf{L}\mathbf{z} + \boldsymbol{\epsilon}. \quad (5.1)$$

To constrain \mathbf{v} to be physiologically plausible, we choose to restrict our model to parametric HRF shapes \mathbf{v}_α . A classical choice is to select \mathbf{v}_α as a the linear combination of d atoms $\sum_{i=1}^d \alpha_i \mathbf{b}_i$, where $(\mathbf{b}_i)_{i \in [1..d]}$ are some well define HRF atoms [Lindquist and Wager, 2007, Pedregosa et al., 2015, Friston et al., 1998a]. Here instead, we propose to use a reference HRF denoted h_{ref} [Friston et al., 1998a] and dilate the time such as \mathbf{v}_α is the discretization of $v_\alpha(t) = v_{\text{ref}}(\alpha t)$. The main advantage of this choice is to vary the full width at half-maximum (FWHM) of the HRF and its time-to-peak (TP) with only one parameter. The model in Equation 5.1 has an ambiguity in magnitude, as if \mathbf{v} is multiplied by β and \mathbf{u} is scaled down by the same factor, our model remains the same. To fix this scale ambiguity, we set $\|v_{\text{ref}}\|_\infty = 1$.

Semi-blind-deconvolution as a joint optimization problem

If the additive noise in Equation 5.1 is considered to be Gaussian, the parameter of the HRF α and the derivative of the neural activation signal \mathbf{u} can be jointly estimated by solving

$$\begin{aligned} & \arg \min_{\alpha \in \mathbb{R}, \mathbf{z} \in \widetilde{\mathbb{R}}^T} \frac{1}{2} \|\mathbf{v}_\alpha * \mathbf{L}\mathbf{z} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{z}\|_1, \\ & \text{subject to} \quad \alpha_{\min} \leq \alpha \leq \alpha_{\max}. \end{aligned} \quad (5.2)$$

This optimization problem is not jointly convex in α and \mathbf{z} . For a fixed α , it is convex in \mathbf{z} and for a fixed \mathbf{z} , convexity in α is not guaranteed as it depends on the analytical model of \mathbf{v}_{ref} . However, this 1-dimensional optimization problem can be solved easily as α is constrained to lie in $[\alpha_{\min}, \alpha_{\max}]$. We minimize Equation 5.2 using a block-coordinate descent approach, where we alternate the minimization between \mathbf{z} and α . algorithm 6 details the steps of this procedure.

Algorithm 6: Semi-blind deconvolution scheme of the BOLD signal.

Input : BOLD signal \mathbf{y} , stopping rule ν

- 1 initialization: $\alpha^{(0)}, \mathbf{z}^{(0)} = 0, k = 1$;
- 2 **repeat**
- 3 Deconvolution of the BOLD signal for $\mathbf{h}_{\alpha^{(k-1)}}$:

$$\mathbf{z}^{(k)} = \arg \min_{\mathbf{z} \in \widetilde{\mathbb{R}}^T} \frac{1}{2} \|\mathbf{v}_{\alpha^{(k-1)}} * \mathbf{L}\mathbf{z} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{z}\|_1$$
- 4 Estimate the HRF parameter with fixed $\mathbf{u}^{(k)}$:

$$\alpha^{(k)} = \arg \min_{\alpha \in \mathbb{R}} \frac{1}{2} \|\mathbf{v}_\alpha * \mathbf{L}\mathbf{z}^{(k)} - \mathbf{y}\|_2^2$$
 subject to $\alpha_{\min} \leq \alpha \leq \alpha_{\max}$
- 5 **until** $\|\alpha^{(k)} - \alpha^{(k-1)}\|_2 / \|\alpha^{(k)}\|_2 < \nu$;

For the estimation of \mathbf{z} with fixed α , the accelerated proximal gradient descent algorithm [Beck and Teboulle, 2009] was used as it provides fast convergence to the optimal solution. Other algorithms such as coordinate descent methods [Friedman et al., 2007, Moreau et al., 2018], can also be considered. However they do not improve the results as the problem is convex and can only speed up the convergence. For updating α , we resorted to the limited memory BFGS algorithm [Byrd et al., 1995] implemented by Jones et al. [2001-]. We early-stopped the main loop and each sub-problem too once the iterates stabilized themselves. In practice less than 50 iterations of the main loop were needed to converge.

Owing to the global non-convexity, this approach converges to a local minimizer of Eq. (5.2), which may be suboptimal for our semi-blind deconvolution objective. To limit the impact of the initialization selection, we tested multiple random initializations. However, we found experimentally that initializing α to α_{\max} – i.e. initializing \mathbf{v}_α to the HRF with the tighter FWHM – and \mathbf{z} to 0 is enough to ensure the convergence to reliable estimates. Multiple random initializations on α did not improve the quality of the solution.

5.1.3 Numerical Experiments

In this section, we validated our approach both on simulated and real task fMRI data. They were collected during different task performance in order to exhibit a learning effect on the

5.1. Sparsity-based blind deconvolution of neural activation signal in fMRI 79

HRF and compare our method to a state-of-the-art approach [Pedregosa et al., 2015]. All experiments were performed in Python and our implementation along with the code for experimental validation is freely available online¹ in order to support reproducible research.

Results on synthetic data

Artificial time series. We randomly generated 100 neural activation signals \mathbf{u}^* of 5 blocks, with an average duration of 12 s each and a standard deviation of 1 s. We choose a TR of 0.75s and a scan duration of 3 min to mimic the Human Connectome Project (HCP) protocol. We defined a common HRF shape \mathbf{v}^* for all these artificial voxels. Last, we investigated 6 different scenarios with signal-to-noise-ratio (SNR) ranging from 1 to 20 dB.




Figure 5.1-1 – Evolution with respect to the SNR of the ℓ_2 relative error defined as the mean across voxels of $\|\hat{\mathbf{u}} - \mathbf{u}^*\|_2 / \|\mathbf{u}^*\|_2$ for the neural activation signal and as $\|\mathbf{v}_{\hat{\alpha}} - \mathbf{v}^*\|_2 / \|\mathbf{v}^*\|_2$ for the HRF.

Results. We tested our semi-blind deconvolution approach to recover the pair $(\hat{\alpha}, \hat{\mathbf{z}})$ from each measured time series and then deduce the HRF $\mathbf{v}_{\hat{\alpha}}$ and the neural activation signal $\hat{\mathbf{u}} = \mathbf{L}\hat{\mathbf{z}}$. As shown in Figure 5.1-1, in low SNR cases we did not perfectly recover both signals. In contrast, as the SNR increases the error of our estimator is significantly reduced by a factor of 3 on the most challenging problem (estimation of \mathbf{u}). A visual inspection of the different estimates confirms that our approach behaved accurately according to our model.

Results on real data

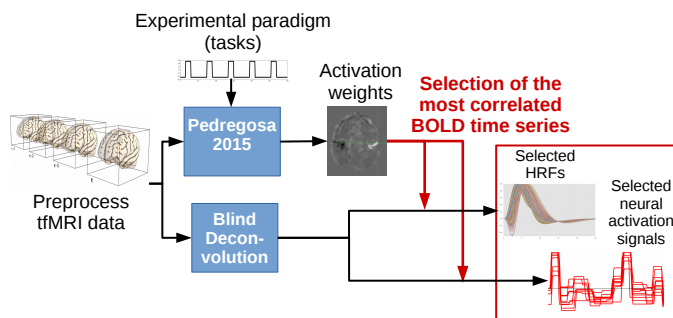


Figure 5.1-2 – Voxel selection procedure for the validation of the semi-blind deconvolution method on HCP data. The voxels are selected based on their correlation level with the EP.

HCP Data. Our validation was performed on the HCP dataset Van Essen et al. [2013] which comprises fMRI recordings of participants performing different motor tasks. The tasks were adapted from the protocol developed in Yeo et al. [2011]. We choose this dataset as it presented both a good temporal and spatial resolution. A short time of repetition (TR=0.720 s) was actually used to collect interleaved simultaneous multislice echo-planar images with a Multi-Band factor of 8 and a spatial resolution of 2x2x2mm. Each fMRI run lasted 3min34s in total during which $\tilde{T} = 284$ scans were acquired.

The fMRI data were already preprocessed using a classical pipeline including realignment, coregistration, spatial normalization and smoothing (5 mm isotropic). The EP was divided

¹<https://github.com/CherkaouiHamza/pybold>




Figure 5.1-3 – Neural activity surrogates normalized by their ℓ_∞ norm. The standard deviation across voxels is encoded by transparency around mean curves for the EP (*green*), the preprocessed BOLD signals \mathbf{y} (*black*) in the most correlated voxels, and the neural activation signals \mathbf{u} (*red*) estimated with our semi-blind deconvolution approach for the same voxels.




Figure 5.1-4 – HRF estimates computed for two different tasks in one participant to the HCP protocol. In (*green*) the canonical SPM HRF, in (*blue*) the reference HRFs estimated using Pedregosa et al. [2015] and in (*red*) the HRFs estimated using the proposed semi-blind deconvolution technique.

in two sets of motor tasks, with 15 s fixation blocks at the beginning, in the middle of the acquisition and at the end of the recording. Each task set was composed of 5 blocks of 12 s each, preceded by a 3 s cue indicating the task to be performed by the participant. The latter corresponded to moving the tongue, tapping the left or right finger or squeezing the left or right toes. In what follows, we only consider one participant even though our results are reproducible across individuals.

Voxel Selection. Each fMRI run comprises a huge data set consisting of 230,314 voxels (i.e. time series) recorded along 284 time points. As our method is so far univariate, it estimates an activation signal and HRF in each voxel independently. Hence, an important aspect in the validation consisted in selecting activated voxels for which these estimates are meaningful. Following the work by Pedregosa et al. [2015], we used a General Linear Model (GLM) that also embeds a supervised voxelwise HRF estimation to regress the convolution of the known EP with the HRF estimate on the measured BOLD signal. From all voxel candidates, we extracted the 100 mostly correlated which are associated with the highest coefficients in the GLM. This process is illustrated in Figure 5.1-2.

Results. Figure 5.1-3 presents the neural activation signals \mathbf{u} estimated with our method for the left hand motor task in one participant. The estimated neural activation signals retrieved the two well defined blocks, suggesting that the model proposes coherent blocks for the neural activation signals with a timing close to that of the EP. Interestingly, one can observe that the measured BOLD signals are postponed in time as compared to the recovered neural activation signals, which is consistent with the sluggishness of the haemodynamic response.

Figure 5.1-4 displays the HRF estimates for two tasks performed by the same participant using the method proposed in Pedregosa et al. [2015] and ours in the semi-blind deconvolution scheme. The HRF estimates were averaged across the 100 selected voxels. For the *visual*

5.2. fMRI BOLD signal decomposition using a multivariate low-rank model⁸¹

fixation task lasting 20 s in total, both methods recover a similar HRF shape. This was expected as the BOLD signal in response to the visual task elicits the strongest activity for both methods. The HRF curves depart from the canonical HRF as all the selected voxels did not confine to the primary visual system but were instead spread between motor and visual regions.

For the *left hand motor* task lasting 24 s in total, the HRF shape recovered by the two approaches differ. The early initial dip found by Pedregosa et al. [2015] is questionable as such depletion may physiologically occur only in the first second after stimulation [Frau-Pascual et al., 2015]. The HRF estimates obtained through semi-blind deconvolution appear more plausible even though the time-to-peak is quite large too. This difference is explained by the capacity of our model to cope with the latencies between the EP and the neural activity signal. Moreover, the HRF estimates obtained using Pedregosa et al. [2015] are close to each other. In contrast, Figure 5.1-4 shows that our HRF estimate significantly differs from the one in response to the *left hand* task. This suggests a task-dependent shape for the HRF as previously demonstrated in the literature [Ciuciu et al., 2003]. Moreover, when we used other tasks available in the HCP dataset we still noticed this coherent task-dependent or learning effect: for instance, the *right* and *left hand* tasks provide similar HRFs to the *right* and *left foot* tasks, respectively.

5.1.4 Interim summary

In this section, we recovered coherent HRF estimates in voxels correlated with a specific task without the explicit knowledge of the experimental paradigm. To do so, we simultaneously estimated a neural activation signal, which may depart from the EP in terms of timing. We observed a clear dependence between the HRF shape and the different tasks involved in the EP. The ℓ_1 norm regularization parameter λ gathers the statistical relevance of our model and is set by hand so far. Future developments, such as a more robust deconvolution technique with a concomitant Lasso or the squared root Lasso could be explored. The dependence on the reference HRF should be investigated too.

However, the described method remains purely univariate and could be extended to a multivariate setting. To limit the number of unknown neural activation signals and take advantage of the spatial structure of the signal, we propose in the next section to aggregate the data with rank-1 constraints [Dupré La Tour et al., 2018].

5.2 fMRI BOLD signal decomposition using a multivariate low-rank model

5.2.1 Introduction

The classical data analysis approach proposes to decompose the BOLD signal using multiple predefined regressors. Each regressor is a time series that models the given temporal signature of an experimental stimulus or task, named condition, convolved with a canonical HRF [Friston et al., 1999a]. Those time-courses are concatenated into a so-called design matrix, and fitted to the observed BOLD data. The estimated coefficients provide the encoding localization of each condition in the brain [Friston et al., 1999a]. The main limitation of this massively univariate approach is twofold: first, it treats one voxel at a time

using the same model; second, it requires the prior knowledge of the experimental paradigm. For these reasons, unsupervised multivariate methods have been introduced in the literature to deal with paradigm-free fMRI datasets such as resting-state recordings. The most famous are likely the principal component analysis (PCA) [Viviani et al., 2005] and the independent component analysis (ICA; Beckmann and Smith 2004, Varoquaux et al. 2010). However, these techniques directly work on the measured BOLD time series and do not deconvolve them to highlight neural activities. An alternative consists in disentangling the neurovascular coupling by deconvolving the BOLD signal using a well chosen HRF [Karahanoğlu et al., 2013, Caballero-Gaudes et al., 2012, Cherkaoui et al., 2019a] and thus recovering voxel-wise neural activation signals. Those approaches provide as many components as the number of voxels. Those components are then used to explore the underlying structure in the data by quantifying either how they cluster together or their functional connectivity.

Goals and contributions

This section presents a new algorithm that aims to offer a rich decomposition of the BOLD signal using low-rank sparse decomposition. Following the ideas developed in the dictionary learning literature [Olshausen and Field, 1997, Moreau et al., 2018], our approach consists in modeling the observed BOLD signal as a linear combination of a limited number of temporal atoms whose first-order derivative is sparse. In that purpose, we introduce spatio-temporal maps which take the neurovascular coupling (temporal aspect) and the localization of activations (spatialization) into account. Then, we jointly estimate those temporal atoms and the associated maps with properly selected constraints. The resulting optimization problem is non-convex but an approximated solution can be computed using an alternate minimization algorithm with an efficient procedure to be performed at each step. Section 5.2.2 introduces our modeling of the BOLD signal and presents our estimation algorithm. Next, our technique is evaluated against state-of-the-art algorithm in Section 5.2.3. Conclusions and future work are discussed in Section 5.2.4.

5.2.2 Low rank decomposition of the BOLD signal

In this section, we present our modeling of the BOLD signal and derive an efficient algorithm to estimate its parameters.

Linear and time-invariant modeling

A common model for the multivariate (P voxels, T scans) BOLD data $\mathbf{Y} \in \mathbb{R}^{P \times T}$ with $\mathbf{Y} = (\mathbf{y}_j)_{j \in \{1..P\}}$ is the linear and time-invariant model (LTI) Boynton et al. [1996], where for each voxel, the measured time series, denoted $\mathbf{y}_j \in \mathbb{R}^{1 \times T}$, is the convolution of a neural activation signal, denoted $\tilde{\mathbf{u}}_j \in \mathbb{R}^{1 \times \tilde{T}}$ with a given HRF, here denoted $\mathbf{v} \in \mathbb{R}^{1 \times L}$, such that $\mathbf{y}_j = \mathbf{v} * \tilde{\mathbf{u}}_j + \mathbf{e}_j$ where $\mathbf{e}_j \in \mathbb{R}^{1 \times T}$ refers to an additive white Gaussian noise Ciuciu et al. [2003]. Typically, the HRF \mathbf{v} has a restricted support in time and quantifies the neurovascular coupling in a specific region of the brain. For the sake of simplicity, the same HRF shape is usually considered for the whole brain and we choose the canonical SPMs double gamma function HRF, as mention in Lindquist et al. [2009]. This model extends as

$$\begin{bmatrix} \mathbf{Y} \\ (P \times T) \end{bmatrix} = \sum_{k=1}^K \begin{bmatrix} \mathbf{u}_k \\ (1 \times T) \end{bmatrix} \star \begin{bmatrix} \mathbf{w}_k \mathbf{v}^T \\ (P \times 1) (1 \times L) \end{bmatrix} + \begin{bmatrix} \mathbf{E} \\ (P \times T) \end{bmatrix}$$

Figure 5.2-5 – Illustration of the low-rank BOLD signal model (the colors are there for illustrative purposes).

follows:

$$\mathbf{Y} = \mathbf{v} \star \tilde{\mathbf{U}} + \mathbf{E} \quad (5.3)$$

with $\mathbf{E} = (\mathbf{e}_j)_{j \in \{1..P\}} \in \mathbb{R}^{P \times T}$ and $\tilde{\mathbf{A}} = (\tilde{\mathbf{u}}_j)_{j \in \{1..P\}} \in \mathbb{R}^{P \times T}$. The activation signals $\tilde{\mathbf{U}}$ capture, in an univariate manner, the periods of time during which some voxels are involved in task performance (or in spontaneous BOLD signal fluctuations). In this univariate model, P independent neural activation signals $(\tilde{\mathbf{u}}_j)_{j \in \{1..P\}}$ are learned, one for each voxel. In our work, we propose to learn K temporal activations $(\mathbf{u}_k)_{k \in \{1..K\}}$ and their associated spatial maps $\mathbf{u}_k \in \mathbb{R}^{P \times 1}$, as we aim to recover K distinctive functional networks with a specific temporal fingerprint. This can be modeled by replacing each vector $\tilde{\mathbf{u}}_j$ in Eq. (5.3) with a linear combination of the activations $(\mathbf{u}_k)_{k \in \{1..K\}}$. A classical assumption for these temporal activation signals is to consider them piecewise constant as in Caballero-Gaudes et al. [2012], Karahanoğlu et al. [2013], Cherkaoui et al. [2019a]. To that aim, we model them as $\mathbf{u}_k = \mathbf{Lz}_k$, where \mathbf{z}_k is sparse. The spatial configuration $\mathbf{w}_k \in \mathbb{R}^{P \times 1}$ encodes which voxels are linked to a given temporal activation $\mathbf{Lz}_k \in \mathbb{R}^{1 \times T}$. In our work, we propose a fixed HRF \mathbf{v} and define the rank-1 spatio-temporal maps $\mathbf{w}_k \mathbf{v}^T \in \mathbb{R}^{P \times V}$ as the convolution kernel with the neural activity, as depicted in Fig. 5.2-5. Learning the HRF will be deferred to future work. Our forward model for BOLD fMRI data thus reads:

$$\mathbf{Y} = \sum_{k=1}^K (\mathbf{Lz}_k) \star (\mathbf{w}_k \mathbf{v}^T) + \mathbf{E}. \quad (5.4)$$

Optimization problem

The spatial maps $(\mathbf{w}_k)_{k \in \{1..K\}}$ and the neural activation signal $(\mathbf{z}_k)_{k \in \{1..K\}}$, from Eq. (5.4), can be jointly estimated by solving the following constrained minimization problem:

$$\begin{aligned}
 J((\mathbf{u}_k)_k, (\mathbf{z}_k)_k) &= \frac{1}{2} \left\| \mathbf{Y} - \sum_{k=1}^K (\mathbf{Lz}_k) \star (\mathbf{w}_k \mathbf{v}^T) \right\|_F^2 + \lambda \sum_{k=1}^K \|\mathbf{z}_k\|_1 \\
 \text{subject to } &\|\mathbf{w}_k\|_1 = \eta \quad \text{and} \quad w_{kj} \geq 0
 \end{aligned} \quad (5.5)$$

To be consistent with Caballero-Gaudes et al. [2012], Karahanoğlu et al. [2013], Cherkaoui et al. [2019a], we enforce the temporal atoms $(\mathbf{z}_k)_{k \in \{1..K\}}$ to be sparse in order to constrain $(\mathbf{Lz}_k)_{k \in \{1..K\}}$ to be piecewise constant signals. Indeed, most of experimental paradims in fMRI propose to model task-related evoked activity using block signals. Moreover, to limit

the indetermination in the convolution we impose the non-negativity in the entries of the spatial maps $(\mathbf{w}_k)_{k \in \{1..K\}}$. Last, to deal with the scale ambiguity – the fact that any solution pair $(\widehat{\mathbf{w}}_k, \widehat{\mathbf{z}}_k)$ is known up to a multiplicative constant – we set $\forall k \in \{1..K\} \quad \|\mathbf{w}_k\|_1 = \eta$, with $\eta \in \mathbb{R}_+$ being a parameter of our method that fixes the magnitude of each spatial maps. This optimization problem is biconvex in $(\mathbf{w}_k)_{k \in \{1..K\}}$ and $(\mathbf{z}_k)_{k \in \{1..K\}}$, meaning that it is convex in each variable but not jointly convex. We minimize Eq. (5.5) using a block-coordinate descent algorithm, where we alternate the minimization between the two convex problems in $(\mathbf{w}_k)_{k \in \{1..K\}}$ and in $(\mathbf{z}_k)_{k \in \{1..K\}}$. Algorithm 7 details these two steps.

Algorithm 7: Low rank decomposition of the BOLD signal.

Input: BOLD signal \mathbf{Y} , ϵ

- 1 initialization: $\mathbf{z}_k^{(0)} = \mathbf{0}_{\widetilde{T}}$, $\mathbf{w}_k^{(0)} = \mathbf{w}_k^{(init)}$, $i = 1$;
- 2 **repeat**
- 3 Estimate the temporal atoms $\mathbf{z}_k^{(i)}$ with fixed $\mathbf{w}_k^{(i-1)}$:

$$\arg \min_{(\mathbf{z}_k)_k} \frac{1}{2} \left\| \mathbf{Y} - \sum_{k=1}^K (\mathbf{L}\mathbf{z}_k) \star (\mathbf{w}_k^{(i-1)} \mathbf{v}^\top) \right\|_F^2 + \lambda \sum_{k=1}^K \|\mathbf{z}_k\|_1$$
- 4 Estimate the spatial maps $\mathbf{w}_k^{(i)}$ with fixed $\mathbf{z}_k^{(i)}$:

$$\arg \min_{(\mathbf{w}_k)_k} \frac{1}{2} \left\| \mathbf{Y} - \sum_{k=1}^K (\mathbf{L}\mathbf{z}_k^{(i)}) \star (\mathbf{w}_k \mathbf{v}^\top) \right\|_F^2$$
 subject to $\|\mathbf{w}_k\|_1 = \eta$ and $w_{kj} \geq 0$
- 5 **until** $\frac{J((\mathbf{z}_k^{(i-1)})_k, (\mathbf{w}_k^{(i-1)})_k) - J((\mathbf{z}_k^{(i)})_k, (\mathbf{w}_k^{(i)})_k)}{J((\mathbf{z}_k^{(i-1)})_k, (\mathbf{w}_k^{(i-1)})_k)} \leq \epsilon$;

We minimize each step with an accelerated forward-backward algorithm [Beck and Teboulle \[2009\]](#) with Armijo backtracking line search [Armijo \[1966\]](#). Recall that for $\mathbf{v} \in \mathbb{R}^{1 \times L}$, $\mathbf{a} \in \mathbb{R}^{1 \times \widetilde{T}}$ and $\mathbf{x} \in \mathbb{R}^{1 \times T}$ $\nabla_{\mathbf{z}} (\frac{1}{2} \|\mathbf{y} - \mathbf{v} \star \mathbf{a}\|_2^2) = -\mathbf{v}^\dagger \star (\mathbf{y} - \mathbf{v} \star \mathbf{a})$ with the time flipped HRF $v_j^\dagger = v_{\widetilde{T}-j}$, thus our gradient steps read:

$$\nabla_{\mathbf{z}_\ell} F_{\mathbf{w}_k}(\mathbf{z}_k) = -\mathbf{L}^\top \left((\mathbf{w}_\ell \mathbf{v}^\top)^\dagger \star \left(\mathbf{Y} - \sum_{k=1}^K (\mathbf{L}\mathbf{z}_k) \star (\mathbf{w}_k \mathbf{v}^\top) \right) \right),$$

$$\nabla_{\mathbf{w}_\ell} F_{\mathbf{z}_k}(\mathbf{w}_k) = -\mathbf{v} \left((\mathbf{L}\mathbf{z}_\ell)^\dagger \star \left(\mathbf{Y} - \sum_{k=1}^K (\mathbf{L}\mathbf{z}_k) \star (\mathbf{w}_k \mathbf{v}^\top) \right) \right).$$

The computation of $\nabla_{\mathbf{z}_\ell} F_{\mathbf{w}_k}(\mathbf{z}_k)$ is optimized by pre-computing $-\mathbf{L}^\top (\mathbf{w}_\ell \mathbf{v}^\top)^\dagger \star \mathbf{Y}$ and $\mathbf{L}^\top (\mathbf{w}_\ell \mathbf{v}^\top)^\dagger \star (\mathbf{w}_k \mathbf{v}^\top) \mathbf{L}$ while that of $\nabla_{\mathbf{w}_\ell} F_{\mathbf{z}_k}(\mathbf{w}_k)$ is accelerated by pre-computing $-\mathbf{v} (\mathbf{L}\mathbf{z}_\ell)^\dagger \star \mathbf{Y}$ and $\mathbf{v} (\mathbf{L}\mathbf{z}_\ell)^\dagger \star \mathbf{v}^\top (\mathbf{L}\mathbf{z}_k)$, as those quantities remain constant during these respective steps.

The proximal operator of $g_z((\mathbf{z}_k)_k) = \lambda \sum_{k=1}^K \|\mathbf{z}_k\|_1$ is the soft-thresholding defined coordinate-wise as $\text{sign}(z)(|z| - \lambda)_+$. For the constraint $g_w((\mathbf{w}_k)_k) = I_{\|\mathbf{w}_k\|_1 = \eta} + I_{w_{kj} \geq 0}$ on the spatial maps, the corresponding proximal operator is given by $\text{prox } g_w(\mathbf{w}_k) = [(\mathbf{w}_{kj} - \mu)_+]_{1 \leq j \leq P}$ with μ is defined as $\sum_{j=1}^P \max\{0, w_{kj} - \mu\} = \eta$ and an efficient implementation can be found in [Condat \[2016\]](#).

We early-stopped the main loop when each main iteration does not decrease sufficiently the cost function. In practice less than 50 iterations of the main loop were needed to converge.

Owing to the global non-convexity, this approach converges to a local minimizer of Eq. (5.5), which may be sub-optimal for our estimation objective. To initialize the spatial




Figure 5.2-6 – (a) In black the observed BOLD signal within the associated activation region normalized by their ℓ_∞ norm, in blue the true temporal atoms, in orange the recovered temporal atoms. (b) The yellow-purple maps define the spatial ground truth and estimates. The standard deviation across voxels is encoded by transparency around mean curves.

maps $(\mathbf{w}_k)_{k=1\dots K}$, we draw each entry as a centered Gaussian variable with variance 1. To limit the impact of the initialization selection, we run multiple times the minimization.

5.2.3 Numerical experiments

In this section, we validate our approach on simulation and illustrate its application to real fMRI data. All experiments were performed in Python and our implementation, as well as the scripts for experimental validation², are freely available online.

Results on synthetic data

Artificial BOLD time series. We randomly generated $P = 100$ BOLD signals \mathbf{Y} . Each time series \mathbf{y}_j was defined as the linear combination of two temporal atoms $(\mathbf{z}_1, \mathbf{z}_2)$ comprising two blocks each whose duration was fixed to 10 s and the magnitude was randomly drawn from a Gaussian distribution centered on 1.0. The weights are defined in two spatial maps $(\mathbf{w}_1, \mathbf{w}_2)$ with a single non-zero pixel in each map. To simulate a realistic scenario, we chose a TR of 1.0 s and a total scan duration of 1min40s ($T = 100$ scans). We added a centered Gaussian noise such that the generated synthetic data has a signal-to-noise ratio of 1.0 dB, defined by

$$\text{SNR} = 10 \log_{10} \left(\frac{\left\| \sum_{k=1}^K (\mathbf{L}\mathbf{z}_k) \star (\mathbf{w}_k \mathbf{v}^\top) \right\|_2^2}{\|\mathbf{E}\|_2^2} \right).$$

Results. For this experiment, we chose $K = 2$ and set the regularization parameter λ to $0.4\lambda_{\max}$, where λ_{\max} is the minimal value for which $\mathbf{0}$ is solution of Eq. (5.5) and we fix the ℓ_1 -norm for each map to be equal to $\eta = 10.0$. Fig. 5.2-6(a) displays the estimated temporal atoms $(\mathbf{L}\hat{\mathbf{z}}_1, \mathbf{L}\hat{\mathbf{z}}_2)$ in orange, along with the observed BOLD signals in black and the true signals in blue. Fig. 5.2-6(b) shows the corresponding spatial map estimates $(\hat{\mathbf{w}}_1, \hat{\mathbf{w}}_2)$ besides the ground truth. The temporal atom estimates recovered well the true signals. The block offsets are correctly temporally aligned but the rising and falling slopes of each block are not perfectly vertical. This indicates that their temporal derivative $(\hat{\mathbf{z}}_1, \hat{\mathbf{z}}_2)$ are not perfectly sparse. The spatial maps are adequately recovered as the activated regions are well localized and the map estimates are sparse.




Figure 5.2-7 – Localization in the *right Precentral Gyrus* region of the two voxels (red crosses) chosen to illustrate the comparison between TA and LRD methods.

Results on real fMRI data

Comparison to Total Activation approach. We qualitatively compare our BOLD low-rank decomposition (LRD) to the state-of-the-art method, called Total Activation (TA) Karahanoglu et al. [2013]. In this approach, the authors propose to deconvolve the BOLD signal by minimizing a convex cost function that involved a sparse temporal constraint and a total variation spatial constraint. Their univariate approach allows to recover, a voxel-specific piecewise constant signal that models the neural activation signal as in our method. The main difference between the two methods is that our multivariate technique allows to recover a much easier-to-interpret decomposition of this neural activation signal. To reduce the computational cost of TA for this experiment, we only considered the temporal regularization.

HCP task fMRI data. Our validation was performed on the Human Connectome Project (HCP) dataset Van Essen et al. [2013] which comprises fMRI recordings of participants performing different motor tasks. The tasks were adapted from the protocol developed in Yeo et al. [2011]. We chose this dataset as it presented both a good temporal and spatial resolution. A short time of repetition ($TR=720$ ms) was actually used to collect interleaved simultaneous multislice echo-planar images with a Multi-Band factor of 8 and a spatial resolution of $2 \times 2 \times 2$ mm. Each fMRI run lasted 3min34s in total during which $T = 284$ scans were acquired. The fMRI data were already preprocessed using a classical pipeline including realignment, coregistration, spatial normalization and smoothing (5 mm isotropic). The experimental paradigm (EP) was divided in two sets of motor tasks, with 15 s fixation blocks at the beginning, in the middle and at the end of the recording. Each set was composed of 5 conditions, each modeled by a blocks of 12 s, preceded by a 3 s cue indicating the task to be performed by the participant. The former corresponded to moving the tongue, tapping the left or right finger or squeezing the left or right toes. In what follows, we only consider one participant even though our results are reproducible across individuals.

Voxel selection. We aim to qualitatively compare the recovery of the neural activation signals for these two techniques (LRD and TA) in each voxel. Each fMRI run comprises a huge data set consisting of 230,314 voxels (time-courses). Thus, we only display results for a specific region the *right Precentral Gyrus*, corresponding to a subsample of 960 voxels in which we chose to display two voxels (see Fig. 5.2-7) illustrating the two methodologies.

Results. For this experiment, we chose $K = 8$ as 8 experimental conditions were involved in the paradigm. As those approach are unsupervised models with no ground truth, we set the regularization parameter for LRD and TA by hand such that $\lambda_{TA} = 0.02\lambda_{\max}$ and $\lambda_{LRD} = 0.07\lambda_{\max}$. Last, we set the ℓ_1 -norm for each map of LRD to be equal to $\eta = 10.0$. Fig. 5.2-8 illustrates the behavior of the LRD and TA deconvolution methods in these two voxels. In **voxel-1**, the low-rank neural activation signal (shown in blue) appears similar to

²<https://github.com/CherkaouiHamza/seven>

5.2. fMRI BOLD signal decomposition using a multivariate low-rank model⁸⁷

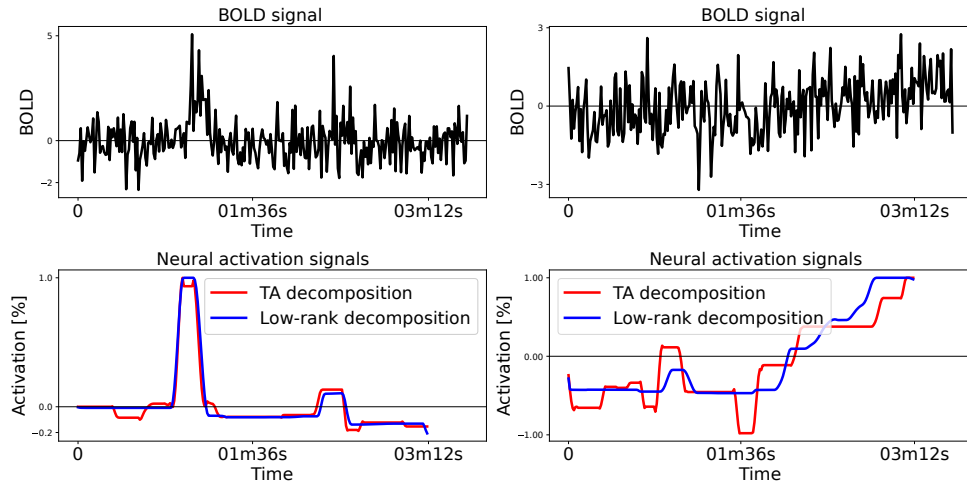


Figure 5.2-8 – In *black* the BOLD signal, in *red* the neural activation signal obtained using TA and in *blue* the neural activation signal obtained using our LRD approach.




Figure 5.2-9 – The different conditions from the experimental paradigm are represented in background: *blue* for *left hand*, *red* for *right hand*, *green* for *left foot* and *yellow* for *right foot*. Each condition lasts 12 s. On the foreground, the estimated temporal atom.

the TA one. Both approaches mainly capture the same dynamics in the measured BOLD signal in this voxel. However, in **voxel-2**, some high frequency components (short-duration activity) that are retrieved in the TA neural activation signal are not captured by our LRD method. This suggests that our model is less sensitive than TA in this voxel. As our temporal atoms are learned across voxels, this is a direct consequence of reducing the number of degrees of freedom in the temporal domain in LRD in contrast to TA. Fig. 5.2-9 and Fig. 5.2-10 depict respectively the temporal activities and the spatial maps associated to 2 of the 8 temporal atoms estimated with our LRD approach. Atom #7 is mainly composed of two blocks locked to the offsets of the condition *left hand* and its spatial map is sparse, with a very well localized region of activation. This suggests that our model has learned the experimental condition that elicits brain activity in this region. In contrast, atom #8 embodies a slightly rising slope between two constant periods, which illustrates its link to the low frequency fluctuations in the fMRI data. The second map displays smoother and wider activation areas in the *right Precentral Gyrus*, suggesting that this model is also capable of modeling trend effects, not related to the conditions.

5.2.4 Summary

This section presents a new low-rank decomposition modeling of the BOLD signal and a corresponding algorithm to perform both the deconvolution in time and the mapping in space. Although such low-rank modeling was already introduced in the dictionary learning literature [Dupré La Tour et al., 2018], we adapted its formulation to fMRI data following




Figure 5.2-10 – The spatial maps associated with each atom. We have limited our analysis to the *right Precentral Gyrus* region (delimited in black). Values of estimated maps are color coded.

ideas from the TA approach. In the validation on real fMRI data, we showed that our method provides a similar decomposition than TA. However, our multivariate model exhibited meaningful components that compose the BOLD signal along with their corresponding spatial maps. Nevertheless, the proposed algorithm relies on a fixed and constant HRF shape, in the next chapter, we will investigate how to estimate the HRF with this framework.

5.3 Conclusion

In this chapter, we have introduced two novel approaches for BOLD signal processing. The first one is dedicated to disentangle the neurovascular coupling voxelwise by introducing a simplistic but efficient single parameter HRF model and by constraining the neural activation to be a block signal. Since this model is univariate, it is greedy in the number of parameters required to fit the data and thus prone to overfitting. To cope with this issue, we have proposed a second approach that captures the neural activity with K temporal components and their associated spatial maps. This model relies on a fixed and constant HRF shape over the whole brain. In the next chapter, we merge these two models to estimate multiple HRFs across the whole brain using a multivariate *semi-blind deconvolution* approach.

* * *
* *
*

Multivariate joint estimation neural activity HRF

Chapter Outline

6.1	Introduction	90
6.2	Multivariate low-rank decomposition of the BOLD signal.	90
6.2.1	Linear and time-invariant modeling	90
6.2.2	Prior information and regularizing constraints	93
6.2.3	Optimization problem.	94
6.3	Model Validation	97
6.3.1	Numerical simulations.	97
6.3.2	Single-subject analysis on rs-fMRI data	98
6.3.3	Haemodynamic estimation stability over time	106
6.4	Clinical validation at the population level	109
6.4.1	Characterization of patients with an history of stroke	109
6.4.2	Middle-age vs elderly subjects classification	110
6.5	Conclusion	113

This chapter was submitted as a Peer-Reviewed Journal as:

H. Cherkaoui, T. Moreau, A. Halimi, C. Leroy, and P. Ciuciu. Multivariate semi-blind deconvolution of fMRI time series. *submitted to NeuroImage, 2020a*.

THIS chapter presents our multivariate approach to disentangle the vascular coupling and neural activity from fMRI data. This approach allows for regionwise HRF estimation in a multivariate manner and summarize neural activity over the whole brain using spatio-temporal atoms. In what follows, we will detail the joint modeling and the associated algorithm to perform parameter estimation from the observed BOLD data. Next, we will show numerical results in a single subject before running group-level analyses to demonstrate the relevance of our approach.

6.1 Introduction

This chapter extends preliminary contributions [Cherkaoui et al., 2019b] and offers a new algorithm that aims to fit a rich multivariate decomposition of the BOLD data using a semi-blind deconvolution and low-rank sparse decomposition. The model distinguishes two major parts in the BOLD signal: the neurovascular coupling and the neural activity signal. One of its main features is to tackle the intra-subject haemodynamic variability by introducing regionwise HRF profiles over a brain parcellation, encoding the spatial variations of the neurovascular coupling. Also, in regard to neural activity, we follow ideas developed in the convolutional dictionary learning literature to develop a low-rank approximation of a signal [Grosse et al., 2007, Dupré La Tour et al., 2018]. We thus model the neural activation signals as a combination of a limited number of piece-wise constant temporal profiles associated with spatial maps that capture the magnitude of functional networks. As such, the neural input signals are represented with their own spatial representation that departs from the neurovascular parcellation. We of course present a scalable optimization algorithm that is able to fit all parameters of interest in the whole brain in a reasonable computing time.

The rest of the chapter is organized as follows. Section 6.2 introduces our modeling of the BOLD data and presents our semi-blind blind deconvolution algorithm. Next, our technique is validated on simulated data and on real rs-fMRI data at the individual level in Section 6.3. In Section 6.4 we illustrate the proposed framework at the population level on the large UK Biobank database. Two main applications were targeted, namely the neurovascular discrimination of patients with stroke episodes as compared to healthy controls and the prediction of brain age. Conclusion and outlook are synthesized in Section 6.5.

6.2 Multivariate low-rank decomposition of the BOLD signal

In this section, we present our modeling of the BOLD signal and derive an efficient algorithm to estimate its parameters.

6.2.1 Linear and time-invariant modeling

Univariate modeling

A common model for the multivariate (P voxels, T scans) BOLD data $\mathbf{Y} \in \mathbb{R}^{P \times T}$ with $\mathbf{Y} = (\mathbf{y}_j)_{j=1}^P$ is the linear and time-invariant model (LTI) [Boynton et al., 1996]. This model is illustrated in Fig. 6.2-1. For each voxel, the measured time series, denoted by $\mathbf{y}_j \in \mathbb{R}^{1 \times T}$, is the convolution of a neural activation signal $\tilde{\mathbf{a}}_j \in \mathbb{R}^{1 \times \tilde{T}}$, with a given HRF, $\mathbf{v} \in \mathbb{R}^{1 \times L}$ and $\mathbf{e}_j \in \mathbb{R}^{1 \times T}$ refers to an additive white Gaussian noise, which leads to:

$$\mathbf{y}_j = \mathbf{v} * \tilde{\mathbf{a}}_j + \mathbf{e}_j . \quad (6.1)$$

Although the noise that contaminates the BOLD effect is serially correlated in time [Woolrich et al., 2001], we assume here that the fMRI data has been pre-whitened. If not, an auto-regressive (AR) modeling for the noise is possible [Penny et al., 2003, Makni et al., 2008] at the expense of an increased computational burden, necessary for identifying the AR parameters.




Figure 6.2-1 – Illustration of the voxelwise BOLD signal modeling: \mathbf{a} is a time series encoding the neural activation signal, \mathbf{v} being the haemodynamic response function (HRF) and \mathbf{e} the additive Gaussian noise. The measured fMRI signal is denoted \mathbf{y} and obtained via $\mathbf{a} * \mathbf{v} + \mathbf{e}$.


Typically, the HRF \mathbf{v} has a restricted support in time of about 20 s. The challenge with HRF modeling is to find a fair trade-off between a flexible model that is able to capture the true haemodynamic response in each brain area and a reliable one that limits overfitting by reducing the number of degrees of freedom (diminishing the variance). Since our approach estimates the neural activity along with the HRF, reducing the number of degrees of freedom is critical to avoid the aforementioned overfitting. In this chapter, we assume M different HRF with \mathbf{v}_m being the HRF corresponding to the m^{th} region Θ_m . Numerous approaches have been proposed to model this haemodynamic response \mathbf{v}_m . The Finite Impulse Response (FIR) [Dale, 1999, Glover, 1999] model does not assume any particular shape for the HRF which make it very flexible but prone to data overfitting in the presence of noise. Regularization has thus been introduced to constrain the overall HRF shape in FIR models and limit their tendency to overfitting, see for instance penalization over the second-order derivative to end up with physiologically plausible smooth HRF estimates [Ciuciu et al., 2003, Casanova et al., 2008]. Alternatively, the HRF has been modeled as a linear decomposition of predefined atoms such as B-splines [Zhang et al., 2007, Vakorin et al., 2007], wavelets [Khalidov et al., 2011], a sensitivity-selected set [Woolrich et al., 2004] or more physiologically informed patterns such as the canonical HRF and its derivatives in time and with respect to the dispersion parameter [Friston et al., 1998a]. All these methods intend to capture fluctuations in haemodynamic delay or shape with the minimum number of parameters. Last, to constrain even more the parameter values and reduce variance estimates, parametric models such as the *inverse logit transform* [Lindquist and Wager, 2007] have been proposed and successfully tested when benchmarking over multiple fMRI data sets [Lindquist et al., 2009]. In this work, we propose the time dilation HRF model [Cherkaoui et al., 2019a], which captures the haemodynamic delay by dilating a reference HRF pattern: $\mathbf{v} = \mathbf{v}_{ref}(\delta t)$ where \mathbf{v}_{ref} is a reference temporal profile, here the canonical HRF. This approach is efficient while simple as it encodes delay fluctuations through a single scalar parameter δ (one degree of freedom). One limitation of this choice is that it leads to the simultaneously variation of the full width-at-half-maximum (FWHM) of the HRF and its time-to-peak (TTP) (see also Fig. 6.2-2). Our voxelwise model reads as follows:

$$\mathbf{y}_j = \mathbf{v}_\delta * \tilde{\mathbf{a}}_j + \mathbf{e}_j . \quad (6.2)$$

Multivariate modeling

To better account for the spatial structure of the neurovascular system and the intrinsic organization of functional networks, we extend this model to the multivariate setting. Our

Figure 6.2-2 – Illustration of two haemodynamic response functions (HRF) denoted h_1, h_2 with the full-width-at-half-maximum (FWHM) and the time-to-peak (TTP) (of h_1) depicted. Here we illustrate these two HRFs with different dilation parameters (δ_1, δ_2) such as $\delta_2 \leq \delta_1$, leading to $\text{TTP}(h_1) < \text{TTP}(h_2)$ and $\text{FWHM}(h_1) < \text{FWHM}(h_2)$.



multivariate model reads as follows:

$$\mathbf{Y} = \mathbf{v}_\delta \bar{*} \tilde{\mathbf{A}} + \mathbf{E} , \quad (6.3)$$

where $\tilde{\mathbf{A}} = (\tilde{\mathbf{a}}_j)_{j=1}^P \in \mathbb{R}^{P \times \tilde{T}}$ and $\mathbf{E} = (\mathbf{e}_j)_{j=1}^P \in \mathbb{R}^{P \times T}$. One limitation of this straightforward approach is that it constrains the haemodynamic response to be the same across the whole brain. As the HRF shape depends on the neurovascular coupling, its features vary in space over different brain areas and between individuals [Handwerker et al., 2004, Badillo et al., 2013]. This suggests that, for a given subject, the HRF should be modeled locally in the brain. An appropriate approach for doing so is to rely on existing brain parcellation [Varoquaux and Craddock, 2013]. Ideally to accurately fit the real haemodynamic response function in a subject, we would favor a large number of regions. However, the larger this number, the smaller the number of voxels per region, which could impair the stability of HRF estimation. For that reason, each region should at least consist of a few hundred voxels. In this work, we rely on the Harvard-Oxford probabilistic brain atlas [Desikan et al., 2006]. We threshold the probabilities to obtain a fine brain parcellation that offers enough flexibility to adapt to the true haemodynamic system. In what follows, we mathematically introduce a brain parcellation with M regions with $(\Theta_m)_{m=1}^M \in \{0, 1\}^P$ 1 if the i^{th} voxel belongs to the m^{th} region and 0 if not. This allows us to extend Eq. (6.3) as follows:

$$\mathbf{Y} = \left(\sum_{m=1}^M \Theta_m^\top \mathbf{v}_{\delta_m} \right) \bar{*} \tilde{\mathbf{A}} + \mathbf{E} , \quad (6.4)$$

The activation signals $\tilde{\mathbf{A}}$ capture for each voxel the periods of time during which any voxel is involved either in task performance or in spontaneous BOLD signal fluctuations. This model remains univariate as P independent neural activation signals $(\tilde{\mathbf{a}}_j)_{j=1}^P \in \mathbb{R}^{P \times \tilde{T}}$ are estimated.

In our work, we rather introduce a low-rank constraint and learn both K temporal atoms (with $K \ll P$) and corresponding spatial maps. These maps encode various functional networks, each of them being summarized by specific neural activation profile. Mathematically, this can be modeled by replacing each vector $\tilde{\mathbf{a}}_j$ in Eq. (6.4) with a linear combination of neural activation patterns $\mathbf{U} = (\mathbf{u}_k)_{k=1}^K \in \mathbb{R}^{K \times \tilde{T}}$, with $\mathbf{u}_k \in \mathbb{R}^{1 \times \tilde{T}}$, modulated in space by the spatial maps $\mathbf{W} = (\mathbf{w}_k)_{k=1}^K \in \mathbb{R}^{K \times P}$, with $\mathbf{w}_k \in \mathbb{R}^{1 \times P}$, such that: $\tilde{\mathbf{A}} = \sum_{k=1}^K \mathbf{w}_k^\top \mathbf{u}_k$. In other words, the spatial configuration \mathbf{w}_k encodes which voxels are linked to a given neural activation profile $\mathbf{u}_k \in \mathbb{R}^{1 \times \tilde{T}}$. Note that a voxel may belong to different functional networks. This is coherent with the fact that a voxel contains about one hundred thousands neurons for a typical spatial resolution (1.5 mm isotropic), and thus the underlying neural activation signals are combined with possibly different temporal fingerprints. Finally, our forward model for BOLD fMRI data is given by Eq. 6.5 (see also Fig. 6.2-3):

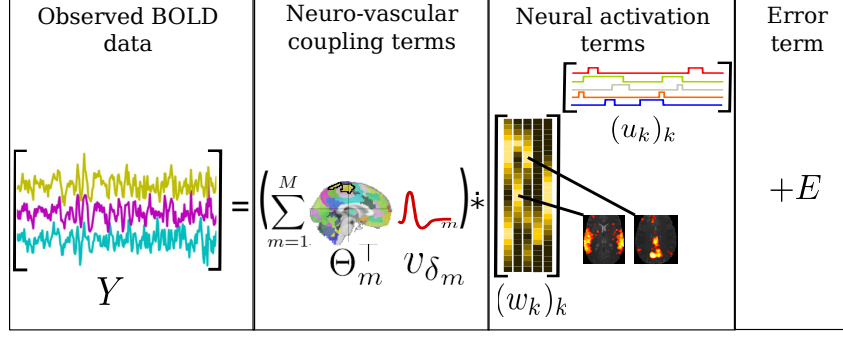


Figure 6.2-3 – Illustration of the low-rank multivariate BOLD signal model (the colors are here for illustrative purposes). \mathbf{Y} stands for the observed BOLD data, \mathbf{v}_m the m^{th} HRF, Θ_m represents the predefined brain parcellation, $\mathbf{W} = (\mathbf{w}_k)_{k=1}^K$ defines the K spatial components and $\mathbf{U} = (\mathbf{u}_k)_{k=1}^K$ the corresponding “neural” activation signals. Last, \mathbf{E} represents the additive white Gaussian noise.

$$\mathbf{Y} = \left(\sum_{m=1}^M \Theta_m^\top \mathbf{v}_{\delta_m} \right) \ast \left(\sum_{k=1}^K \mathbf{w}_k^\top \mathbf{u}_k \right) + \mathbf{E} . \quad (6.5)$$

6.2.2 Prior information and regularizing constraints

The number of unknown parameters $\mathbf{U} \in \mathbb{R}^{K \times \tilde{T}}$, $\mathbf{W} \in \mathbb{R}^{K \times P}$ and $\delta \in \mathbb{R}^{1 \times M}$ is lower than the number of available data \mathbf{Y} . Indeed, the observed fMRI data has P voxels of T scans while the proposed model has to infer K temporal components of \tilde{T} entries, K spatial components of P voxels each, and M HRF dilation parameters. Hence, the number of unknown parameters to be set is $M + (K \times P) + (K \times \tilde{T})$. Importantly, if the number of scans T is large enough and if we adequately choose M and K , then we get $P \times T \gg M + (K \times P) + (K \times \tilde{T})$.

The forward model described in Eq. (6.5) is *trilinear* as it depends on both unknown spatial (\mathbf{W}) and temporal (\mathbf{U}) inputs and unknown convolution filters (\mathbf{v}_{δ_m}). This means that any timing variation in neural activation signals \mathbf{U} can be symmetrically compensated by the opposite variation in the haemodynamic filter \mathbf{v}_{δ_m} or that any sign or scale variation in \mathbf{W} can be compensated by the inverse variation in \mathbf{U} or \mathbf{v}_{δ_m} . To overcome these sign, scale and time-shift ambiguities, we introduce some regularizing constraints in the computation of the parameter estimates. First, to avoid any sign ambiguity in the convolution between the neural input signals and the haemodynamic filter, the HRF filter has a constant and positive maximal amplitude (see Fig. 6.2-2). Second, as there is an interplay between the spatial and temporal components in the input signals, we also impose a non-negativity constraint over the entries of each spatial map \mathbf{w}_k , i.e. $\forall j, k \mathbf{w}_{kj} \geq 0$. By doing this, we only allow the deactivation of a functional network to be encoded in the corresponding \mathbf{w}_k . Third, to deal with the scale ambiguity, we impose $\|\mathbf{w}_k\|_1 = \eta, \forall k = 1, \dots, K$, where $\eta \in \mathbb{R}_+$ is a user-defined parameter that sets the magnitude of each spatial map. As our HRF model has a constant maximal amplitude, only the neural activity signals \mathbf{U} capture the observed BOLD signal fluctuations. Forth, to deal with the time-shift ambiguity, we constrain the dilation parameter δ in the HRF model to be in $[0.5, 2.0]$. This should also permit the recovery of physiologically plausible haemodynamic delays. Complementary to these constraints, akin to [Caballero-Gaudes et al., 2012, Karahanoglu et al., 2013, Cherkaoui et al., 2019b] we

will also assume the neural activation signals \mathbf{U} to be temporally piecewise constant or equivalently assume their first-order derivative $\nabla \mathbf{u}_k, \forall k$ to be sparse. Practically speaking, we add a total variation (TV) regularization term to our model.

6.2.3 Optimization problem

In this section, we derive an algorithm for estimating the dilation parameters $\boldsymbol{\delta} = (\delta_m)_{m=1}^M$, the spatial maps \mathbf{W} and the neural activation signals \mathbf{U} from the model depicted in Eq. (6.5) and the aforementioned regularizing constraints. The estimates for these parameters can be obtained by solving the following constrained minimization problem:

$$\begin{aligned} \arg \min_{(\mathbf{W}, \mathbf{U}, \boldsymbol{\delta})} & \frac{1}{2} \left\| \mathbf{Y} - \left(\sum_{m=1}^M \boldsymbol{\Theta}_m^\top \mathbf{v}_{\delta_m} \right) \ast \left(\sum_{k=1}^K \mathbf{w}_k^\top \mathbf{u}_k \right) \right\|_F^2 + \lambda \sum_{k=1}^K \|\nabla \mathbf{u}_k\|_1, \quad (6.6) \\ \text{subject to} & \quad \forall k, \|\mathbf{w}_k\|_1 = \eta, \quad \forall j, u_{kj} \geq 0, \quad \forall m, \delta_m \in [0.5, 2.0]. \end{aligned}$$

The Gaussian noise hypothesis leads to a quadratic data fidelity term, to measure how well we reconstruct the observed fMRI signals. Moreover, in alignment with Caballero-Gaudes et al. [2012], Karahanoğlu et al. [2013], Cherkaoui et al. [2019a], we enforce the first-order derivative of the temporal atoms $(\nabla \mathbf{u}_k)_{k=1}^K$ to be sparse in order to constrain each \mathbf{u}_k to a piecewise constant signal. For that purpose, we use a TV regularization term, which corresponds to the ℓ_1 norm of the gradient in time $\nabla \mathbf{u}_k$. Importantly, this modeling of the neural activation signals allows us to fully adapt to task-fMRI and rs-fMRI experiments and to perform paradigm-free fMRI data analyses. In the first case, as the task-related BOLD signal is classically modeled as the convolution of an input block signal, representing the experimental paradigm (the onsets of the stimulus trials) with a HRF filter, we can recover neural activation signals close to the experimental paradigm. The neural activation profiles being inferred from the data, we can estimate both input signals corresponding to block and event-related designs, the only difference between the two being the length of activation blocks, which in the latter case can be restricted to a single time point. More interestingly, the proposed framework is even more appealing for processing rs-fMRI data and uncover spontaneous and time-varying fluctuations of brain activity as the the block duration may change from one instance to the next.

Moving to the technical aspects for solving the constrained optimization problem (6.6), it is worth mentioning that it is not globally convex. However, when \mathbf{W} and $\boldsymbol{\delta}$ are fixed, problem (6.6) becomes convex in \mathbf{U} and similarly when \mathbf{U} and $\boldsymbol{\delta}$ are fixed, it becomes convex in \mathbf{W} . Our minimization strategy of Eq. (6.6) thus relies on a block-coordinate descent algorithm, where we alternate the minimization between the two convex problems in \mathbf{W} and \mathbf{U} followed by the non-convex one involving \mathbf{V} . Also, the non-negativity constraints are activated when solving for the spatial maps \mathbf{W} and the boundary constraints over $\boldsymbol{\delta}$ are handled in parallel for each m , i.e. each HRF pattern when solving for \mathbf{V} . Algorithm 8 details these three main steps.

Algorithm 8: Multivariate deconvolution and low-rank decomposition of the BOLD signal.

Input : BOLD signal \mathbf{Y} , ϵ

- 1 initialization: $\forall k, \mathbf{u}_k^{(0)} = \mathbf{0}_{\tilde{T}}, \mathbf{w}_k^{(0)} = \mathbf{w}_k^{(init)}, \delta^{(0)} = \delta^{(init)}, i = 1$;
- 2 **repeat**
- 3 Estimate the temporal atoms $\mathbf{U}^{(i)}$ with fixed $\mathbf{W}^{(i-1)}$ and $\delta^{(i-1)}$:

$$(\mathbf{u}_k^{(i)})_k = \arg \min_{(\mathbf{u}_k)_k} \frac{1}{2} \left\| \mathbf{Y} - \left(\sum_{m=1}^M \Theta_m^\top \mathbf{v}_{\delta_m}^{(i-1)} \right) * \left(\sum_{k=1}^K \mathbf{w}_k^{(i-1)\top} \mathbf{u}_k \right) \right\|_F^2 + \lambda \sum_{k=1}^K \|\nabla \mathbf{u}_k\|_1 .$$
- 4 Estimate the spatial maps $\mathbf{W}^{(i)}$ with fixed $\mathbf{U}^{(i)}$ and $\delta^{(i-1)}$:

$$(\mathbf{w}_k^{(i)})_k = \arg \min_{(\mathbf{w}_k)_k} \frac{1}{2} \left\| \mathbf{Y} - \left(\sum_{m=1}^M \Theta_m^\top \mathbf{v}_{\delta_m}^{(i-1)} \right) * \left(\sum_{k=1}^K \mathbf{w}_k^\top \mathbf{u}_k^{(i)} \right) \right\|_F^2 ,$$
 subject to $\{ \forall k, \|\mathbf{w}_k\|_1 = \eta \text{ and } \forall j, w_{kj} \geq 0 \}$.

Estimate the HRFs $\delta^{(i)}$ with fixed $\mathbf{W}^{(i)}$ and $\mathbf{U}^{(i)}$:

$$(\mathbf{v}_m^{(i)})_m = \arg \min_{(\delta_m)_m} \frac{1}{2} \left\| \mathbf{Y} - \left(\sum_{m=1}^M \Theta_m^\top \mathbf{v}_{\delta_m} \right) * \left(\sum_{k=1}^K \mathbf{w}_k^{(i-1)\top} \mathbf{u}_k^{(i)} \right) \right\|_F^2 ,$$
 subject to $\delta_m \in [0.5, 2.0]$.
- 5 **until** $\frac{J((\mathbf{u}_k^{(i-1)})_k, (\mathbf{w}_k^{(i-1)})_k, (\mathbf{v}_m^{(i-1)})_m) - J((\mathbf{u}_k^{(i)})_k, (\mathbf{w}_k^{(i)})_k, (\mathbf{v}_m^{(i)})_m)}{J((\mathbf{u}_k^{(i-1)})_k, (\mathbf{w}_k^{(i-1)})_k, (\mathbf{v}_m^{(i-1)})_m)} \leq \epsilon$;

In regard to the $(\mathbf{u}_k)_{k=1}^K$ step, we performed the minimization using an adaptive-restart accelerated forward-backward algorithm [O’Donoghue and Candes, 2015]. In regards to the $(\mathbf{w}_k)_{k=1}^K$ step, we first benchmark various algorithms in the dictionary learning literature and selected the most efficient, namely the one used to update the dictionary in Mairal et al. [2009]. Last, for the minimization with respect to $(\delta_m)_{m=1}^M$ we used the accelerated forward-backward algorithm [Combettes and Pesquet, 2009] after checking that it leads to a correct estimation of δ . The reader can find all details of the gradient computation *w.r.t* $(\mathbf{u}_k)_{k=1}^K$ and $(\mathbf{w}_k)_{k=1}^K$ in Cherkaoui et al. [2019b], for the gradient *w.r.t* δ , we detail the gradient derivation of our cost-function from Eq. (6.6) – denoted J hereafter – *w.r.t* δ . Let us define $\tilde{\mathbf{A}} = (\tilde{\mathbf{a}}_j)_{j=1}^P \in \mathbb{R}^{P \times \tilde{T}}$ such as $\tilde{\mathbf{A}} = \sum_{k=1}^K \mathbf{w}_k^\top \mathbf{u}_k$. Moreover, we introduce θ_m the set of indices of voxels belonging to the m^{th} region of the brain parcellation.

$$J(\delta) = \sum_{m=1}^M \sum_{j \in \theta_m} \frac{1}{2} \|\mathbf{v}_{\delta_m} * \mathbf{a}_j - \mathbf{y}_j\|_2^2 + C_{\mathbf{W}, \mathbf{U}}$$

with $C_{\mathbf{W}, \mathbf{U}}$ a constant that does not depend on δ . We aim to compute the gradient of J relative to the value of the parameters δ :

$$\nabla_{\delta} J(\delta) = \left[\frac{\partial J(\delta)}{\partial \delta_1}, \dots, \frac{\partial J(\delta)}{\partial \delta_M} \right]^\top \in \mathbb{R}^M \quad (6.7)$$

To this end, we proceed componentwise:

$$\begin{aligned}
\frac{\partial J(\boldsymbol{\delta})}{\partial \delta_m} &= \frac{1}{2} \sum_{j \in \theta_m} \frac{\partial \|\mathbf{v}_{\delta_m} * \mathbf{a}_j - \mathbf{y}_j\|_2^2}{\partial \delta_m} \\
&= \sum_{j \in \theta_m} \left(\frac{\partial (\mathbf{v}_{\delta_m} * \mathbf{a}_j)}{\partial \delta_m} \right)^\top (\mathbf{v}_{\delta_m} * \mathbf{a}_j - \mathbf{y}_j) \\
&= \left(\frac{\partial \mathbf{v}_{\delta_m}}{\partial \delta_m} \right)^\top \left(\sum_{j \in \theta_m} \mathbf{a}_j^\top * (\mathbf{v}_{\delta_m} * \mathbf{a}_j - \mathbf{y}_j) \right) \\
&= \left(\frac{\partial \mathbf{v}_{\delta_m}}{\partial \delta_m} \right)^\top \underbrace{\left(\mathbf{v}_{\delta_m} * \sum_{j \in \theta_m} \mathbf{a}_j^\top * \mathbf{a}_j - \sum_{j \in \theta_m} \mathbf{a}_j^\top * \mathbf{y}_j \right)}_{\nabla_{\mathbf{v}_{\delta_m}} J}. \tag{6.8}
\end{aligned}$$

Note that $\sum_{j \in \theta_m} \mathbf{a}_j^\top * \mathbf{a}_j$ and $\sum_{j \in \theta_m} \mathbf{a}_j^\top * \mathbf{y}_j$ do not depend on δ_m , thus they can be pre-computed beforehand. The remaining step is to compute $\frac{\partial \mathbf{v}_{\delta_m}}{\partial \delta_m}$. We remind here that \mathbf{v}_{δ_m} is the discretization of the continuous function $\forall t \in \mathbb{R}^+, v_{\delta_m}(t) = v(\delta_m t)$. Thus:

$$\forall t \in \mathbb{R}^+, \quad \frac{\partial}{\partial \delta} v(\delta t) = t v'(\delta t)$$

with \mathbf{v}' the first-order derivative of function \mathbf{v} . Now, taking the definition of $v(\cdot)$ from [Friston et al., 1998b], we get:

$$\forall t \in \mathbb{R}^+, \quad v(t) = \frac{t^{a-1} e^{-t}}{\Gamma(a)} - c \frac{t^{b-1} e^{-t}}{\Gamma(b)}$$

where a, b and c are constants which are given in [Friston et al., 1998b]. A straightforward computation gives us for $t \in \mathbb{R}^+$:

$$v'(t) = \left(\frac{a-1}{t} - 1 \right) \frac{t^{a-1} e^{-t}}{\Gamma(a)} - c \left(\frac{b-1}{t} - 1 \right) \frac{t^{b-1} e^{-t}}{\Gamma(b)} \tag{6.9}$$

$$\frac{\partial}{\partial \delta} v(\delta t) = t v'(\delta t) = \left(\frac{a-1}{\delta} - t \right) \frac{(\delta t)^{a-1} e^{-\delta t}}{\Gamma(a)} - c \left(\frac{b-1}{\delta} - t \right) \frac{(\delta t)^{b-1} e^{-\delta t}}{\Gamma(b)} \tag{6.10}$$

The value of $\frac{\partial \mathbf{v}_{\delta_m}}{\partial \delta_m}$ can thus be computed by taking the discrete time points corresponding to the sampling rate of the BOLD signal and the length of the considered HRF. By replacing its value in the computation of $\frac{\partial J(\boldsymbol{\delta})}{\partial \delta_m}$ from Eq. (6.8), we obtain a closed form expression for the gradient of J w.r.t the HRF dilation parameter δ i.e. $\nabla_{\boldsymbol{\delta}} J(\boldsymbol{\delta})$.

For each step, we implemented the corresponding gradient (i.e. forward move) in an efficient manner to limit multiple computations over the iterations.

Critical steps for the efficiency of this algorithm are the computation of proximal operators for the non-smooth regularizers. In regards to the neural activation patterns (\mathbf{U} step), as the minimization is sequentially performed over the K components, we only need to compute the proximal operator of the TV norm, i.e. $g_u(\mathbf{u}_k) = \lambda \|\nabla \mathbf{u}_k\|_1$. This remains a challenging issue as this operator is not closed form. A seminal contribution has been done in the literature for TV minimization [Chambolle, 2004]. Here, we rather used the Taut-String algorithm proposed by [Barbero and Sra, 2018] for which we use an efficient Python implementation available in an open source package¹. In regards to the constraints on the spatial maps (\mathbf{W}), we also proceed

¹https://pypi.org/project/prox_tv/

separately on the K components: the proximity operator of $g_w(\mathbf{w}_k) = \mathbb{1}_{\|\mathbf{w}_k\|_1=\eta} + \mathbb{1}_{w_{kj}\geq 0}$ where $\mathbb{1}$ stands for the indicator function², is given by:

$$\text{prox } g_w(\mathbf{w}_k) = [(w_{kj} - \mu)_+]_{1\leq j\leq P} \quad (6.11)$$

where μ is defined as $\sum_{j=1}^P \max\{0, w_{kj} - \mu\} = \eta$ and an efficient implementation has been proposed by [Condat \[2016\]](#). We propose to set λ as a fraction of λ_{\max} which is the minimal value of λ for which $\mathbf{0}$ is solution of Eq. (6.6). For the rest of the chapter, we will refer to λ as the fraction of λ_{\max} , such as $\lambda = \lambda_f \lambda_{\max}$, with $\lambda_f \in [0, 1]$.

Algorithm 8 converges to a local minimizer Eq. (6.6) when each main iteration does not decrease sufficiently the cost function. In practice less than 50 iterations of the main loop are needed to converge. To initialize the spatial maps $(\mathbf{w}_k)_{k=1}^K$, we apply an Independent Component Analysis (ICA), implemented in `scikit-learn` [[Pedregosa et al., 2011](#)], on the BOLD signals \mathbf{Y} and retain the produced spatial maps, we initialize the $(\mathbf{u}_k)_{k=1\dots K}$ to zero and each entry of $(\delta_m)_{m\in\{1..M\}}$ to 1.

6.3 Model Validation

We first validate the proposed approach on numerical simulations to illustrate the gain achieved by jointly estimating the neural activity profile and the HRF shape compared to a single deconvolution scheme. Next, we will demonstrate the usefulness of the proposed framework on real rs-fMRI data at the individual level. In particular, we will highlight the impact of the hyper-parameter selection on the decomposition, describe the interpretation of its component and its stability. The usefulness of our method in the context of large cohorts will be investigated in Section 6.4.

6.3.1 Numerical simulations

Synthetic data

We generated two temporal Dirac signals of length $\tilde{T} = 500$ with a fixed sparsity level. Each generated Dirac signal is composed of randomly drawn signed spikes, with location chosen uniformly in time and intensity drawn from a Gaussian distribution $\mathcal{N}(0, 1)$. To produce the corresponding block signals \mathbf{Z} – shown in blue in Fig. 6.3-4[right panels] – we integrated over time these signals and convolved them with a predefined HRF \mathbf{v}_δ to yield two corresponding pure BOLD time series. The chosen HRF has length $L = 25$ and is shown in blue in Fig. 6.3-4[left panel]. For the sake of simplicity, we considered a single HRF profile ($M = 1$) in this synthetic setting, so the haemodynamic properties were supposed constant in space. We then assigned these BOLD signals to spatial locations. Hence, we defined $K = 2$ corresponding 2D maps $\mathbf{W} = (\mathbf{w}_k)_{k=1}^2$ (10×10 , i.e. $P = 100$). Each spatial map has a single activating region consisting of 4 pixels. Each active pixel has a randomly drawn non-negative magnitude, the other ones being set to zero. Then, we normalized each map by its ℓ_1 -norm. Finally we added Gaussian random noise to produce observed, i.e. noisy BOLD signals \mathbf{Y} of length $T = \tilde{T} + L - 1 = 524$ scans ($TR = 1$ s) with a signal-to-noise-ratio (SNR) of -1 dB. The mean synthetic BOLD signals are reported in black traces in Fig. 6.3-4(a)-(b) for both activated

²This function is zero-valued inside the constraint set and equals infinity elsewhere.

regions (in bottom panels) while standard deviation across activated voxels is encoded by transparency around these mean curves.

Numerical results

In a first step, we only estimated the pair (\mathbf{U}, \mathbf{W}) from the synthetic fMRI time series \mathbf{Y} and kept the HRF profile \mathbf{v} constant. The results are reported in Fig. 6.3-4(a). The HRF shape used in this deconvolution process is shown in green in Fig. 6.3-4(a) and actually differs from the true shape used for simulating the data. Because of this discrepancy in terms of haemodynamic delay and peak magnitude, the neural activation signals are not properly recovered (orange traces in Fig. 6.3-4(a)). The magnitude of the estimates $\hat{\mathbf{U}}$ is much larger than the true one. This is partly due to compensate for the smaller magnitude of the HRF (green trace in Fig. 6.3-4(a)) used for deconvolving the BOLD signals. Consequently, the residual mean square errors (RMSEs) computed on the neural activation signals are pretty large. However, we noticed that in both spatial maps, the non-negative magnitudes $\hat{\mathbf{W}}$ are very well estimated. This is a direct consequence of using non-overlapping activating regions for the two neural traces.

In a second step, on the same data set \mathbf{Y} we jointly estimated (\mathbf{U}, \mathbf{W}) and \mathbf{v} using our full semi-blind deconvolution scheme. We kept the same initialization for the HRF shape for the sake of consistency. The results are reported in Fig. 6.3-4(b). The HRF estimate $\hat{\mathbf{v}}$ is shown in orange (dashed line) and actually matches the true curve. Consequently, the neural activation signals $\hat{\mathbf{U}}$ are properly estimated both in time and in magnitude and the corresponding RMSEs are one order of magnitude smaller than those reported in the previous simulated results (see Fig. 6.3-4(a)). This second synthetic setting did not impact the spatial maps, which are still well estimated.

These results on synthetic data confirmed the good expected behavior of the proposed method. From a computational viewpoint, the estimation with constant HRF ran in 0.5s while the full estimation took 1s approximately on a machine with 15 GB of RAM and an Intel processor i7-7600U (2 physical cores, 2.80 GHz).

6.3.2 Single-subject analysis on rs-fMRI data

Data set and parameter setting

To illustrate the proposed semi-blind deconvolution algorithm, we analyzed a single subject extracted from the UKBB resting-state fMRI data set. More investigation on a larger cohort of this data set will be presented in Section 6.4. The rs-fMRI data was 6 min10s long with $\text{TR} = 0.735\text{s}$. The first ten seconds were discarded (dummy scans) so that we end up with $T = 490$ scans (6min). The data was collected on a 3T Skyra Siemens MAGNETOM MR system at an isotropic resolution of $2.4 \times 2.4 \times 2.4 \text{ mm}^3$ using the multi-band GRE sequence ($\text{mb} = 8$).³ Standard pre-statistics processing steps were applied: motion correction using MCFLIRT Jenkinson et al. [2002]; grand-mean intensity normalisation of the entire 4D data set by a single multiplicative factor; high-pass temporal filtering (Gaussian-weighted least-squares straight line fitting); EPI unwarping; gradient distortion correction unwarping. Finally, structured artefacts are removed by ICA processing, see the documentation⁴ for a

³Acquisition details can be found at <https://www.fmrib.ox.ac.uk/ukbiobank/protocol/>.

⁴Preprocessing details can be found at https://biobank.ctsu.ox.ac.uk/crystal/crystal/docs/brain_mri.pdf

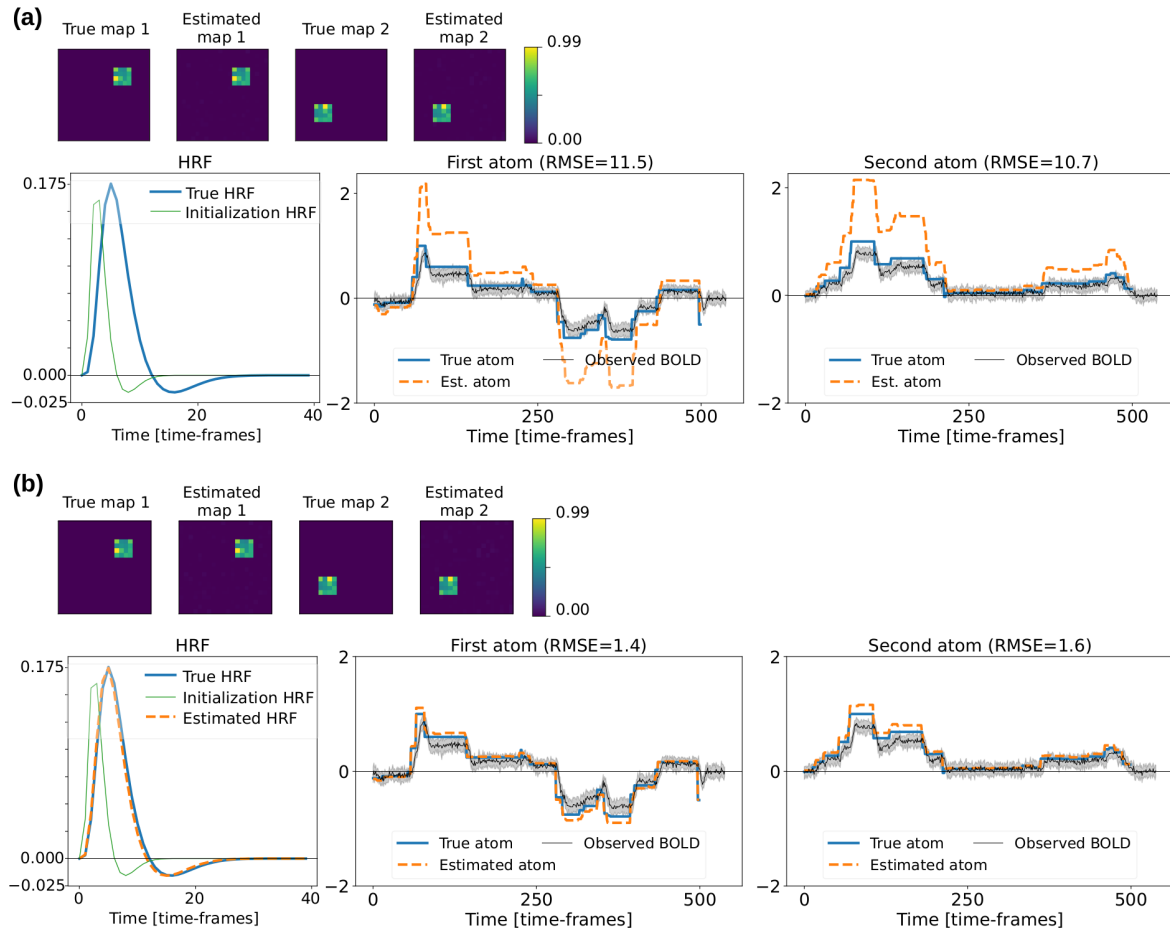


Figure 6.3-4 – **Top (a): deconvolution with fixed HRF.** The top row shows the two true spatial maps and their accurate estimates. The bottom row shows on the left the true HRF shape and the filter used for deconvolution. In the middle and on the right hand side, the true and estimated neural temporal profiles are depicted in blue and orange, respectively for the two activating regions (first atom for map 1, second atom for map 2). The average BOLD time series over the four activating pixels is shown in black and the gray shading is used to report on the variability across activating pixels. **Bottom (b): Semi-blind deconvolution with learned HRF.** The top row shows the two true spatial maps and their accurate estimates. The bottom row shows on the left the true HRF shape, the initial filter used and the final HRF estimate for semi-blind deconvolution. In the middle and on the right hand side, the true and estimated neural temporal profiles are depicted in blue and orange, respectively for the two activating regions (first atom for map 1, second atom for map 2). The average BOLD time series over the four activating pixels is shown in black and the gray shading is used to report on the variability of BOLD signals across activating pixels.

full description.

In this subsection, we manually set the temporal regularization parameter to $\lambda_f = 0.8$ (remember $\lambda_f \in [0, 1]$). This setting achieves a bias-variance trade-off between two extreme situations, namely data overfitting ($\lambda_f = 0$) on one hand and entirely sparse neural activation signals ($\lambda_f = 1$ as $\lambda = \lambda_{\max}$) on the other hand. The question of the unsupervised tuning of λ_f is critical. It could be driven either from a statistical viewpoint (e.g. using the maximum likelihood criterion) that characterizes how likely the measured time series may be observed or using an external task and its corresponding metric such as classification performance (e.g. accuracy in prediction). Because the former approach does not admit a

closed form solution, we explored in the next subsection the impact of changing the temporal regularization (i.e. amount of sparsity in the activation neural signals) on the spatial map of haemodynamic delays. The spatial consistency we reported across regularization levels gave us confidence on the haemodynamic parameter estimates to be further used in subsequent classification tasks (see Section 6.4). In this setting, we implemented a cross-validation step with a leave-one-out loop to tune λ_{\max} in an unsupervised way.

In regard to the number of spatio-temporal atoms K , we set it using the explained variance (or R^2 -score) as target metric in a preliminary study. For this set of parameters, the model estimation took around 1 minute on a machine with 15 GB of RAM and an Intel processor i7-7600U (2 physical cores, 2.80 GHz).

Results

Model selection - setting the temporal regularization parameter λ_f : A well known limitation of regularization methods based on the l_1 -norm such as TV is that large coefficients – here in $(\mathbf{u}_k)_{k=1}^K$ – are shrunken toward zero [Tibshirani, 1996]. Thus, the magnitude of the estimated neural activation signals $(\mathbf{u}_k)_{k=1}^K$ is biased. Moreover, this bias is tightly linked to the choice of the regularization parameter λ_f . Indeed, the larger this parameter is, the more $(\mathbf{u}_k)_{k=1}^K$ are shrunken toward zero. To quantify this effect on our model, we applied the spatio-temporal decomposition with $M = 96$ ROI and $K = 20$ and various temporal regularization level λ_f on the cohort of $S = 459$ subjects sampled from the UK Biobank resting-stage fMRI dataset used in Section 6.4.2. Fig. 6.3-5 reports the grand average of the dilatation parameters

$$\bar{\delta} = \frac{1}{MS} \sum_{s=1}^S \sum_{m=1}^M \hat{\delta}_m^s$$

and its variance with respect to the regularization parameter λ_f . We observed that the HRF dilation parameters decrease with the temporal regularization level – and thus the corresponding time-to-peaks increase with λ_f . This results from the fact that the model with large regularization parameters only accounts for sharp transition in the BOLD signal mean value, which are well approximated with fast HRF.

This previous result entails that the haemodynamic delay estimated by our model (6.6) may be biased. However, because there is a single temporal regularization parameter, we expect that this bias impacts the whole brain uniformly. To assess this shared effect on the estimated parameter, we observe the relative variations of δ_m . Fig. 6.3-6 displays the value of $\bar{\delta}_m(\lambda_f)$ relative to $\bar{\delta}(\lambda_f)$ for $S = 459$ subjects with three temporal regularization values $\lambda_f \in \{0.001, 0.5, 0.9\}$ on the MNI template. Precisely, for each regularization parameter and for each ROI m , we compute $\bar{\delta}_m/\bar{\delta}$ where $\bar{\delta}_m = \frac{1}{S} \sum_{s=1}^S \hat{\delta}_m^s$ is the average value of the dilation parameter across subjects. While the magnitudes change when the regularization changes, as seen in Fig. 6.3-5, the spatial structure of dilatation parameters in the brain is globally preserved. Indeed, the normalized maps look very similar for any choice of regularization parameter, showing that the relative variation between each area of the brain are preserve while changing the hyper-parameter. Thus, we can state that the haemodynamic response from the middle temporal gyrus is faster than the response from the frontal orbital cortex, as described in Fig. 6.3-9. This means that while the numerical value of the time-to-peak for a given area may not be reflect the actual haemodynamic delay in the brain, the estimated




Figure 6.3-5 – **Evolution of the group-level grand average haemodynamic dilation parameter $\bar{\delta}$ as a function of the temporal regularization level $\lambda_f \in [0, 1]$.** The solid blue line reflects the decreasing evolution of $\bar{\delta}$ when $\lambda_f \rightarrow 1$, where the value of $\bar{\delta}$ was spatially averaged over the $M = 96$ parcels and across $S = 459$ subjects from the UK Biobank database. The transparent blue shadow represents the standard deviation around the mean parameter $\bar{\delta}$. In short, the larger λ_f , the smaller $\bar{\delta}$ and thus the larger the mean TTP.




Figure 6.3-6 – **Group-level mean of haemodynamic dilation parameter maps normalized by the grand average $\bar{\delta} = \frac{1}{MS} \sum_{s=1}^S \sum_{m=1}^M \hat{\delta}_m^s$ as a function of temporal regularization ($\lambda_f \in [0, 1]$).** From top to bottom, axial slices showing the group-level values of the ratio between $\bar{\delta}_m = \frac{1}{S} \sum_{s=1}^S \hat{\delta}_m^s$ and $\bar{\delta}$ in each parcel m for increasing values of $\lambda_f \in \{0.001, 0.45, 0.9\}$. The spatial structure of the maps of haemodynamic dilation parameter remain remarkably stable for various λ_f .

coefficients reflect the spatial variations of the delay between the different areas of the brain. Moreover, these variations are stable with the choice of temporal regularization. Hence, choosing a potentially suboptimal value for λ_f is of limited impact when the primary interest is investigating abnormalities in the neuro-vascular coupling.

Model selection - setting the number of spatio-temporal components K : The first question we addressed on real rs-fMRI data was to optimally set the number of spatio-temporal atoms K and to find the best compromise between model complexity and model accuracy. For this purpose, we looked at two complementary criteria. The first one is standard and corresponds to the R^2 score that quantifies the variance explained by model (6.5) over the total sum of squares whereas the second one is given by the determinant of the correlation matrix between the neural activation signals.




Figure 6.3-7 – **Model selection: compromise between model complexity and accuracy.** (a): Evolution of the R^2 score as a function of the number of spatio-temporal atoms K in model (6.5) ranging from 2 to 50. (b): Evolution of the determinant of the correlation matrix Σ_K between neural activation signals as a function of K ranging in the same interval as mentioned earlier.

The R^2 -score is defined as follows: $R^2 = 1 - \frac{SS_{\text{res}}}{SS_{\text{tot}}}$ where SS_{tot} quantifies the variance of the data \mathbf{Y} and SS_{res} the variance of the residuals after fitting model (6.5) by minimizing the cost function described in Eq. (6.6). The R^2 -score may vary from $-\infty$ in pathological cases to 1 for a perfectly matching model. A good model is normally associated with $R^2 > 0$ and means that the L_2 norm of the residual is lower than the variance of the data. We therefore ran multiple model fitting for K in a range of $\{2, 3, 4, \dots, 10, 15, 20, \dots, 50\}$. The results are shown in Fig. 6.3-7(a) and illustrate that the model accuracy first increases as a function of K up to reaching a plateau around $R^2 \simeq 0.55$ for $K = 20$. So adding more spatio-temporal components no longer improves its ability to capture variability in the data while it becomes more complex.

The second information measure we used to help us select K was based on the determinant of the correlation matrix $\Sigma_K = (\mathbb{E}[(\mathbf{u}_k - \mathbf{m}_k)(\mathbf{u}_\ell - \mathbf{m}_\ell)^T / \sigma_k^2 \sigma_\ell^2])_{k,\ell}$ between the temporal atoms $(\mathbf{u}_k)_{k=1}^K$. The quantities σ_k^2 and σ_ℓ^2 define the variance of the neural activation signals \mathbf{u}_k and \mathbf{u}_ℓ . As Σ_K is semi-positive definite with entries between 0 and 1, its eigenvalues are positive or null and so its determinant varies between 0 and 1: $\det \Sigma_K = 1$ when matrix Σ_K defines a basis, which means that all atoms are orthogonal and decorrelated like in a PCA decomposition. In contrast, $\det \Sigma_K = 0$ when matrix Σ_K is not of full rank so at least one atom could be obtained as a linear combination from the others. Therefore, as before we ran multiple model fitting for K in a range of $\{2, 3, 4, \dots, 10, 15, 20, \dots, 50\}$ and we plotted in Fig. 6.3-7(b) the evolution of the determinant of Σ_K as a function of K . The results show us that beyond $K \geq 20$, we get a correlation matrix with $\det \Sigma_K \leq 10^{-10}$ which tends to zero. According to this criterion, from now we should therefore not exceed 20 temporal atoms. Thus, from the R^2 -score criterion and this $\det \Sigma_K$ criterion, in our following experiments we will keep $K = 20$.

Analysis of spatial decomposition Fig. 6.3-8 shows the spatial maps of this spatio-temporal decomposition for this individual and Tab. 6.1 summarizes the list of main regions and functional networks retrieved in this setting.

It is worth mentioning that the sensory networks (visual, auditory and motor) are quite well retrieved by a single or multiple components, respectively located respectively in the occipital (components 10, 16 and 19), temporal (component 7) and motor (components 6, 15 and 18) cortices. The maps associated with the motor network are split and lateralized (6 and 18 on the right hemisphere while 15 in the left). The different areas of the visual system are split too between the primary visual cortex (component 10) and the extrastriate cortex (components 16, 19). The language system has also been identified by a single component. Part of the well known intrinsic resting-state networks (RSN; Menon 2015) have been captured: (i) the right and left-lateralized fronto-parietal resting-state networks spatially similar to the bilateral dorsal attention network are captured by component 1 and 4, respectively while the left and right lateral frontoparietal central executive networks appear in component 9. The default mode network (DMN), which deactivates during demanding cognitive tasks is represented in component 2 and 20: the angular gyrus (AG) appear in both components while the posterior cingulate cortex (PCC) is captured only by component 20 and the medial prefrontal cortex (mPF) by component 2. However, we found that component 9 actually mixes the left AG in the DMN with a left-lateralized fronto-parietal network that perfectly matches the CEN [Menon, 2015]. In contrast, we did not clearly retrieve neither the salience network – usually anchored in anterior insula and dorsal anterior cingulate cortex – nor the right CEN.

Spatio-temporal decomposition To fully illustrate our method on real rs-fMRI data, we show the whole set of output features (neural activation signals, spatial maps, HRF shapes) in Fig. 6.3-9. We also depict a voxel-based denoised BOLD signal reconstructed as the convolution between the neural input and the HRF estimate. Fig. 6.3-9(a) represents together a neural activation signal in the primary visual cortex and the corresponding spatial map (component 10 in the above mentioned decomposition). The proposed axial views allow us to identify the primary visual cortex and the calcarine fissure. Fig. 6.3-9(b) depicts similar features in the DMN (component 2) and the Pearson correlation coefficient with the neural time course in the visual cortex. Its negative value confirms a negative correlation between the task-positive and the DMN network. Both time courses actually present alternating periods of positive and negative activity but they are almost uncorrelated. Fig. 6.3-9(c)-(d) illustrate the fastest and slowest HRF time courses estimated in the regions of interest depicted in red. The fastest haemodynamic response ($\text{FWHM}_f = 5.1\text{s}$) was found in the middle temporal gyrus while the slowest ($\text{FWHM}_s = 8.0\text{s}$) is located in the frontal orbital cortex. Fig. 6.3-9(e) finally shows how well our approach is able to fit the rs-fMRI time course measured in voxel marked by the black cross in Fig. 6.3-9(c). The neural activation signal is piecewise constant and ahead in time compared to the BOLD time series. Once convolved with the HRF profile, the denoised BOLD signal appears as a smoother version of the measured BOLD time course: its magnitude is smaller and its fluctuations in time are slower. This is a direct consequence of the temporal regularization used to recover sparse input signals.

Analysis of correlation structure Next, to go beyond the spatial analysis, Fig. 6.3-10 depicts the correlation matrix between the corresponding neural activation signals. It is then insightful to notice that the correlation between the multiple components in a given network

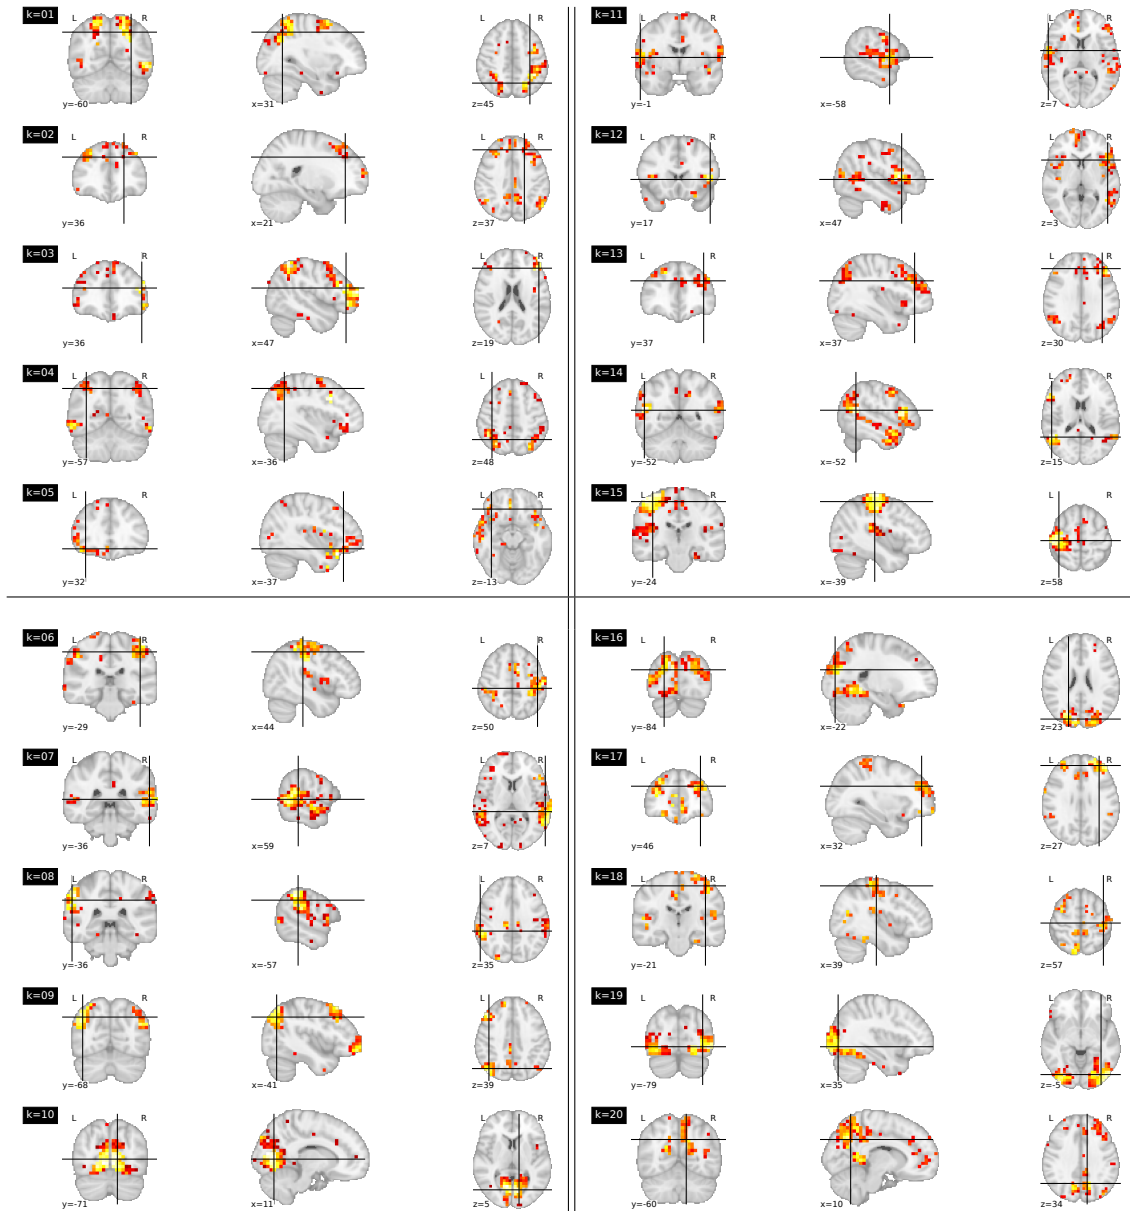


Figure 6.3-8 – **Spatial decomposition of rs-fMRI data for $K = 20$.** From top to bottom and left to right, the twenty labeled spatial maps are shown using the three orthogonal views (coronal on the left, sagittal in the middle and axial on the right). The labeling is arbitrary and the coordinates are given in the MNI space.

Table 6.1 – **Taxonomy of brain regions and functional networks** involved in the spatio-temporal decomposition (6.5) with $K = 20$. We only refer to the main regions in each component. dAN: dorsal Attention Network; DMN=Default Mode Network; IPS=Intra-Parietal Sulcus; FEF: Frontal Eye Fields; CEN: Central Executive Network; R and L stand for left and right hemispheres. The region in bold font matches the location of the cross in Fig. 6.3-8 and have been identified from the AAL template.

Network	# Comp.	Brain areas
Visual	10	R calcarine fissure and surrounding cortex
	16	superior occipital gyrus
	19	Inferior occipital gyrus
Auditory	7	R superior temporal gyrus
Motor	6	R post-central
	15	L precentral gyrus
	18	R precentral gyrus
DMN	2	R superior frontal gyrus, dorsolateral , mPF, AG
	9	L angular gyrus
	20	R precuneus
dAN	1	R inferior parietal , L inferior parietal, R FEF, R inferior frontal gyrus
	4	L inferior parietal , L FEF, L inferior frontal gyrus
Language	14	L middle temporal gyrus , Broca's area
CEN	9	left fronto-parietal
Unclassified	3	R middle frontal gyrus
	5	L inferior frontal gyrus, orbital part
	8	L supramarginal gyrus , R supramarginal gyrus
	11	L Rolandic Operculum
	12	R inferior frontal gyrus , triangular part
	13	R middle frontal gyrus
	17	R middle frontal gyrus , L middle frontal gyrus

are quite strong. For the visual network we observed a correlation coefficient varying between 0.35 and 0.69, the largest value being reached for areas located both in the extrastriate cortex. The same conclusion holds in the motor network with a correlation level varying between 0.5 and 0.65. In regard to the DMN, component 2 plays the role of a hub as it correlates with components 9 and 20 pretty strongly⁵ between 0.27 and 0.4. However, component 9 is almost decorrelated from component 20 as it mixes regions in the dAN and the DMN. Overall, this analysis shows that the proposed approach does not separate RSN in single components. However, it still achieves a meaningful decomposition. For illustrative purposes, the spatial decomposition for $K \in \{5, 8, 10, 15, 30, 40, 50\}$ and the corresponding correlation matrices between the temporal atoms are reported in the Supplementary Material. Interestingly, these results show that some RSNs are either not recovered or mixed together for small $K \leq 15$ whereas for large values of $K \in \{25, 30, \dots, 50\}$ each RSN is split in multiple components. Increasing the number of components extends the range of variation of the correlation coefficients in both positive and negative senses. This confirms that our approach does not have statistical independence guaranties like PCA and ICA do when decomposing the neural activity.

⁵no statistical test performed at the individual level

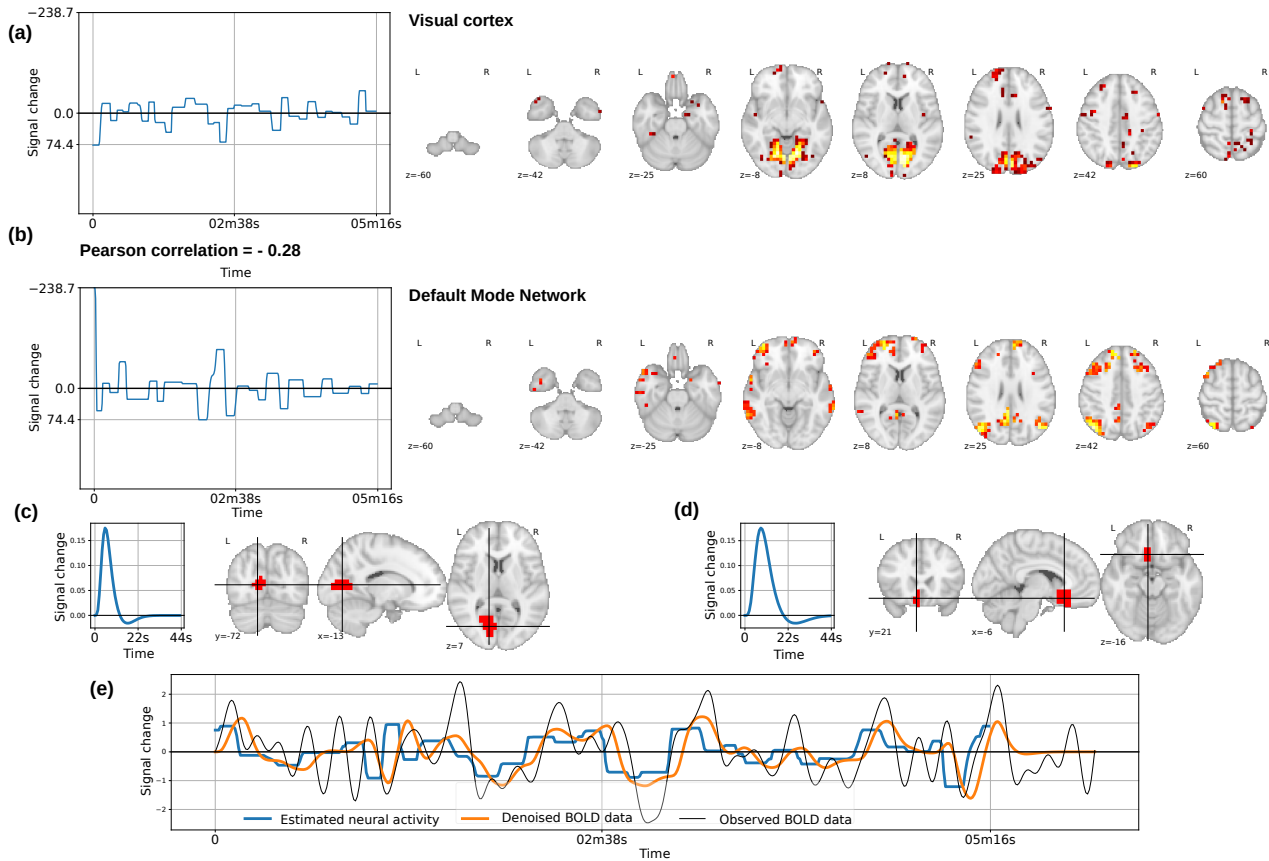


Figure 6.3-9 – **Single subject results from rs-fMRI semi-blind deconvolution analysis.** **Top Row (a):** Neural activation signal \hat{u}_{10} (left) and corresponding spatial map (axial views) \hat{w}_{10} (right), mostly involving activated voxels in the visual cortex. **Second row (b):** Neural activation signal \hat{u}_2 (left) and corresponding spatial map (axial views) \hat{w}_2 (right), mostly involving activated voxels in the default mode network (DMN). **Third row, left (c): Fastest haemodynamic region.** Fastest HRF estimate \hat{v}_{δ_f} (left) located in the middle temporal gyrus as shown on the parcel mask Θ_f (right). **Third row, right (d): Slowest haemodynamic region.** Slowest HRF estimate \hat{v}_{δ_s} (left) located in the frontal orbital cortex as shown on the parcel mask Θ_s (right). **(Bottom row (e): Voxelwise time courses.** Estimate of the neural activation signal (in blue), superimposed on the denoised BOLD signal (in orange) computed as the convolution with the local HRF estimate. The observed BOLD time course is shown in black.

6.3.3 Haemodynamic estimation stability over time

The shape of the HRF is controlled by the neurovascular coupling including both neural and non neural factors such as glial cell activity, cerebral energy metabolism, and the cerebral vasculature. Abnormalities in the local vascular system or cell communication due to pathological state or changes in cerebral blood flow upon psychoactive drugs could influence this haemodynamic response. As we expect the HRF estimate to be stable if none of those events took place, we propose to study the intra-subject stability of HRF estimates over time, namely between consecutive time periods. For doing so, we compare the intra-subject variability of the HRF whole brain dilation parameter vector δ to the inter-subject variability of the same quantity. We thus introduce two reference ℓ_2 distances, namely the within-subject distance $WS(\delta_1^s, \delta_2^s) = \|\delta_1^s - \delta_2^s\|_2^2$ where $(\delta_i^s)_i$ correspond to the vectors of spatially aggregated HRF dilation parameters that were estimated over two periods




Figure 6.3-10 – Correlation matrix between neural activation signals for $K = 20$. Triangular inferior view of the semi-definite positive matrix Σ_K for $K = 20$. All entries vary between -1 and $+1$ as they reflect correlation coefficients.

of time T_1 and T_2 in the same individual s . Similarly, for any pair of subjects (s_1, s_2) and a given period T , we measure the between-subject distance between $(\delta_T^{s_i})_i$ vectors as follows: $\text{BS}(\delta_T^{s_1}, \delta_T^{s_2}) = \|\delta_T^{s_1} - \delta_T^{s_2}\|_2^2$. The goal is then to compare the within- and between-subject distances across individuals and show that the intra-subject variability is significantly lower than the inter-subject one over a sufficiently large population.

Data set and numerical analysis

We selected 100 healthy subjects from the Human Connectome Project (HCP) data set [Van Essen et al., 2013] at random. We used this data set because of the availability of a 12 min long rs-fMRI run with a short time of repetition ($TR = 0.72$ s), see Glasser et al. [2013] for a full description of the acquisition parameters and the pre-statistics processing steps. In this rs-fMRI run for each individual, we extracted two segments of 4 minutes each, denoted as T_1 and T_2 hereafter, the first and last parts of the recording. We then applied the proposed multivariate spatio-temporal decomposition to each segment using $K = 8$ spatio-temporal atoms $(\mathbf{u}_k, \mathbf{w}_k)_{k=1}^K$ and a brain atlas $\Theta = (\Theta_m)_{m=1}^M$ [Desikan et al., 2006] composed of $M = 96$ regions of interest (ROIs). This haemodynamic brain parcellation thus yields 96 HRF dilation parameters $\delta = (\delta_m)_{m=1}^M$ for each individual. In practice, in the definition of $\text{WS}(\cdot, \cdot)$ and $\text{BS}(\cdot, \cdot)$, the true vectors $\delta_{T_j}^{s_i}$ ($i = 1, \dots, 100$, $j = 1, 2$) have been replaced by their estimates $\hat{\delta}_{T_j}^{s_i}$ computed by solving Eq. (6.6) for the two 4-min rs fMRI data sets (T_1

and T_2). To make sure that our conclusions hold for a large scale of temporal regularization parameters, we spanned the range $\lambda_f \in [0, 1]$ and repeated the same procedure over 10 discrete values of λ_f within this interval.

Results

In Fig. 6.3-11, the box plot in blue shows the within-subject distance $WS(\hat{\delta}_{T_1}^{s_i}, \hat{\delta}_{T_2}^{s_i})$ between the two 4 min rs-fMRI segments for all individuals and across 5 values of λ_f covering the whole interval $[0, 1]$. The orange and green box plots in Fig. 6.3-11 depict the between-subject distances computed over the first and second segments respectively, namely $BS(\hat{\delta}_{T_1}^{s_i}, \hat{\delta}_{T_1}^{s_j})$ and $BS(\hat{\delta}_{T_2}^{s_i}, \hat{\delta}_{T_2}^{s_j})$ with $i \neq j$. We observed that the within-subject (i.e. inter-segment) variability is systematically lower than the between-subject variability and that all metrics remain stable across regularization levels. To go further, we performed a statistical analysis (paired t -test) by comparing the mean of the WS and BS distributions and we obtained significant p -values ($p < 10^{-8}$) showing that the within-subject haemodynamic variability is significantly lower than the between-subject fluctuations. In contrast, the statistical inter-individual comparison between the two segments is not significant ($p \simeq 10^{-2}$). These results are valid for all tested regularization levels indicating a minor impact of the regularization parameter onto the haemodynamic parameter estimate. In sum, this analysis demonstrates that the whole brain characterization of the vascular system remains stable in a given individual between two periods shortly spaced in time, compared to the same analysis between individuals and so that the haemodynamic response discriminates each subject from the others.




Figure 6.3-11 – **Within-subject vs between-subject analysis of the haemodynamic variability.** The box plots show respectively in blue, orange and green the distribution of $WS(\hat{\delta}_{T_1}^{s_i}, \hat{\delta}_{T_2}^{s_i})$ for all subjects ($i = 1, \dots, 100$), $BS(\hat{\delta}_{T_1}^{s_i}, \hat{\delta}_{T_1}^{s_j})$ and $BS(\hat{\delta}_{T_2}^{s_i}, \hat{\delta}_{T_2}^{s_j})$ with $i \neq j$. These distributions are assessed for 5 levels of temporal regularization ($\lambda_f \in \{10^{-3}, 0.22, 0.45, 0.67, 0.9\}$) and remain stable. Statistical analysis (paired t -test) was conducted to assess the significance of the difference between the mean of the within- and between-subject ℓ_2^2 norm distributions. Significant differences are marked with a *.

6.4 Clinical validation at the population level

In the previous section, the numerical experiments were devoted to demonstrate the meaningfulness and reliability of the proposed multivariate spatio-temporal within-subject decomposition of fMRI data, especially in resting-state experiments. In this section, our main objective is to showcase the application of this approach to clinical diagnosis. For this purpose, we leverage the functional features (haemodynamic delays, neural activation signals, etc.) output by our approach to first characterize patients with history of stroke compared to healthy controls and then to discriminate middle-age vs elderly subjects. In both analyses, we again used the 6-min long rs-fMRI data from the UK Biobank database.

6.4.1 Characterization of patients with an history of stroke

Stroke is a medical condition in which the blood supply to is interrupted or reduced in a brain area, resulting in ischemic brain tissue and neuronal damage. This pathology is considered as a major health issue nowadays [England, 2018]. In this field, multiple studies [Min et al., 2018] have proposed approaches to better estimate the stroke risk for patients. However, in these attempts a major issue is the precise estimation of the brain damage that occurs in the neurovascular system during and after a stroke episode. To that purpose, we tested our approach to characterize the effect of stroke on the haemodynamic response in the brain.

We considered 24 patients of both genders and various ages who suffered from a stroke in the past from the UK Biobank database. For comparison purposes, we selected 24 healthy controls matched in age and gender from the same database. We applied the same decomposition ($K = 20$, $M = 96$, same λ_f) to each patient and healthy control. Fig. 6.4-12(a) and Fig. 6.4-12(b) show respectively the corresponding normalized maps of haemodynamic dilation parameters (δ_m) $_{m=1}^{96}$ in a healthy control and stroke patient, respectively. The normalization has been done by dividing all dilation parameter values by their within-subject average, namely $\bar{\delta}^{\text{HC}}$ and $\bar{\delta}^{\text{SP}}$ respectively. We first observed that the dilation parameters were larger in average in the healthy condition compared to stroke ($\bar{\delta}^{\text{HC}} > \bar{\delta}^{\text{SP}}$). This corresponds to shorter and more homogeneous TTPs in the brain in the healthy condition. The shortest TTP found in the healthy control was actually located in the primary visual cortex (axial slice, $z=-2$, left hemisphere), a result consistent with the literature on fastest haemodynamic responses often detected in visual areas [Handwerker et al., 2004, Badillo et al., 2013]. In contrast, Fig. 6.4-12(b) illustrates that the haemodynamic dilation parameters δ^{SP} are smaller (so the TTPs longer) in the stroke patient. Also, we found less variability in the healthy subject since the difference between the maximum and minimum TTPs were smaller ($\Delta_{\text{TTP}}^{\text{HC}} = 1.25$ s) compared to the stroke patient ($\Delta_{\text{TTP}}^{\text{SP}} = 2.25$ s). Importantly, Fig. 6.4-12(a) illustrates the relative symmetry of haemodynamic territories that exists in normal subjects between both hemispheres [Raemaekers et al., 2018].

On the contrary, Fig. 6.4-12(b) reveals a wider asymmetry between the two hemispheres in the stroke patient. Interestingly, in this patient we noticed the presence of larger TTPs in the middle left precentral gyrus and left motor cortex (resp. $z=44$ and $z=60$), namely the brain regions supposedly impacted by the stroke episode.

To go one step further, we quantified the spatial asymmetry of the haemodynamic structure within each individual. To this end, we computed the inter-hemispheric haemodynamic ℓ_2 distance (IHD) between the HRF dilation parameter vectors estimated over the left and

right hemispheres in laterally matched brain regions, respectively denoted δ_L and δ_R . This intra-subject distance is defined as follows:

$$\text{IHD}(\delta_R^s, \delta_L^s) = \|\delta_L^s - \delta_R^s\|_2, \quad \forall s = 1, \dots, 24.$$

A zero-valued distance thus reflects a perfect symmetry of the estimated haemodynamic responses. In contrast, we expect to uncover asymmetry between haemodynamic territories respectively located in the ischemic and normal hemispheres.

By pulling down the values of IHD across all individuals within each group (HC vs SP), we estimated the IHD distributions for the two populations of interest, as shown in Fig. 6.4-12(c). In the latter graph, we illustrate how different the two cohorts are in terms of neurovascular asymmetry. The group of 24 stroke patients exhibit larger haemodynamic differences between the ischemic and normal hemispheres. We statistically assessed such difference between the two distributions using a two-sample Kolmogorov-Smirnov test and found a significant p-value ($p = 3.8 \cdot 10^{-4}$). This quantification thus confirmed preliminary visual assessment. We report this p-value with a temporal regularization defined such as $\lambda_f = 0.001$. However, we obtain similar p-value results when using the 5 others levels of temporal regularization ($\lambda_f \in [0.001, 0.9]$).

In summary, this analysis has shown that the proposed framework is instrumental in discriminating healthy subjects from stroke patients, both at the individual and group-levels, using haemodynamic features and an neurovascular asymmetry index, which allowed us to localize pathological haemodynamic delays.

6.4.2 Middle-age vs elderly subjects classification

In the previous part, we performed group-level statistical analysis in the classical way. In this part, we intend to assess the prediction power of the proposed framework in order to classify middle-age vs elderly subjects using standard machine learning tools [Pedregosa et al., 2011]. The reason for choosing this classification task between middle-age and elderly individuals lies first in the fact that multiple studies have pointed out the modification of the haemodynamic system with healthy aging [Ances et al., 2009, Li et al., 2018, West et al., 2019] and second in a regain of interest in the literature for brain age analysis using multiple neuroimaging techniques [Engemann et al., 2020]. We thus intend to assess whether our approach is able to capture the effect of aging and if so, whether the haemodynamic features reflect more brain aging compared to the neural activity signals.

For that purpose, we still relied on the UK Biobank database as in the previous experiments as the short TR (TR=0.735 s) in the rs-fMRI data set (6 min long) provides a suitable setting to investigate the HRF evolution with aging. Here, we selected 459 healthy subjects of both genders and divided them in two balanced groups: the middle-age (MA: 40 – 44 yo) and elderly (E: 64 – 70 yo) groups. We applied the decomposition (6.6) to each subject using 5 levels of temporal regularization ($\lambda_f \in [0.001, 0.9]$), $K = 20$ temporal components and we used the same brain parcellation of $M = 96$ ROIs as before to segregate the HRFs in space.

First, we analyzed the haemodynamic differences between the two populations by computing a two-sample t -test on the distributions of dilation parameters. We used the temporal regularization parameter $\lambda_f = 0.675$ which is the one selected through cross validation for our classification model in the subsequent paragraph. The results were first quantified with

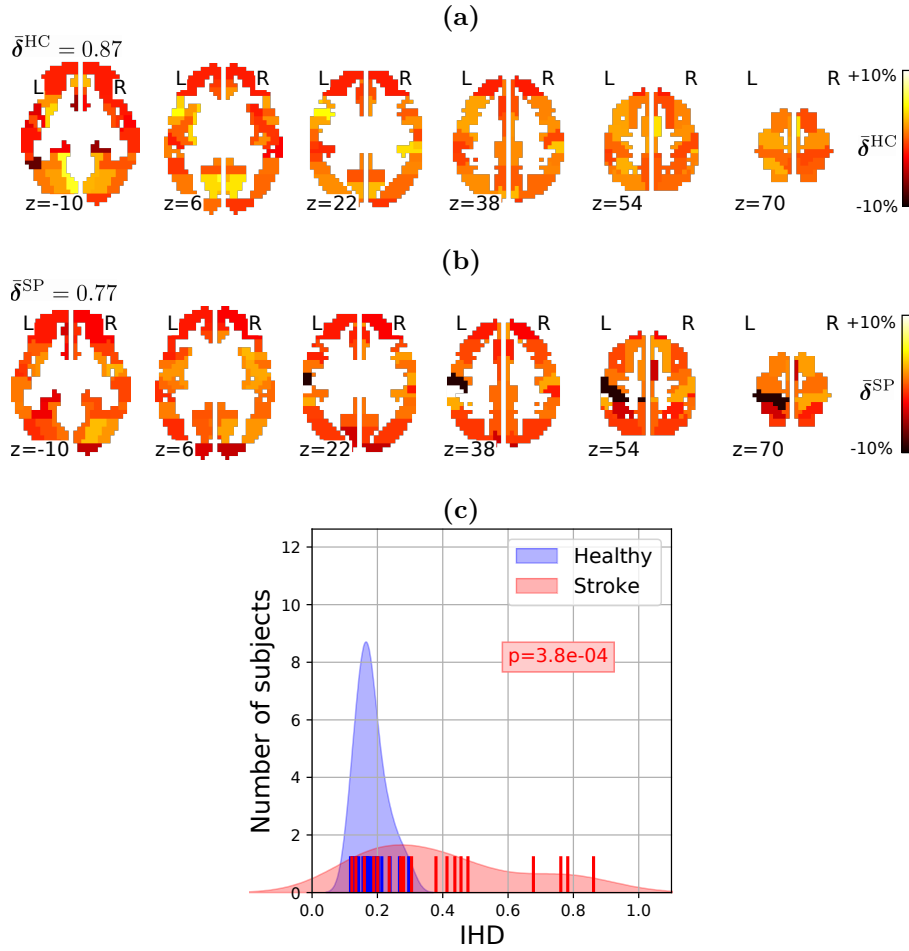


Figure 6.4-12 – **Haemodynamic discrimination between stroke patients (SP) and Healthy Controls (HC).** **Top (a)- Middle (b):** Normalized haemodynamic dilation parameter maps in a healthy control and stroke patient (St), respectively. The maps have been respectively normalized by the within-subject mean value ($\bar{\delta}^s = \frac{1}{M} \sum_{m=1}^{96} \hat{\delta}_m^s$) computed for each subject $s = \text{HC, St}$. Shortest TTP is reached in the visual cortex. Larger haemodynamic dilation parameters maps and thus shorter TTPs are retrieved in healthy condition ($\bar{\delta}^{\text{HC}} = 0.87 > \bar{\delta}^{\text{SP}} = 0.77$). Stronger fluctuations around the mean are observed in the pathological condition as we reported a larger difference between the maximum and the minimum TTP ($\Delta_{\text{TTP}}^{\text{SP}} = 2.25$ s) for the stroke patient than for the Healthy Controls ($\Delta_{\text{TTP}}^{\text{HC}} = 1.25$ s). **Bottom (c):** Histograms of the normalized inter-hemispheric haemodynamic distance (IHD) between dilation parameters computed over the left and right hemispheres (i.e. $\hat{\delta}_L$ and $\hat{\delta}_R$, respectively) in HC (blue) and SP (red), respectively. The significant reported p-value ($p = 3.8 \times 10^{-4}$), which is associated with a two-sample Kolmogorov-Smirnov test between the two distributions, demonstrates that the neurovascular asymmetry in SP is significantly different and actually more spread compared to HC.

t -scores to compare the dilation parameters in each region. The results are presented in Fig. 6.4-13(a) and illustrate large differences between the two populations. More specifically, higher dilation parameters or shorter TTP were retrieved in middle-age subjects as the t -scores were mostly positive for the comparison $\bar{\delta}_m^{\text{MA}} > \bar{\delta}_m^{\text{E}}$. This is notably visible in the Willis Polygon, temporal cortices, angular gyri, the medial prefrontal cortices and the superior frontal cortices. To assess the statistical significance, we also computed the

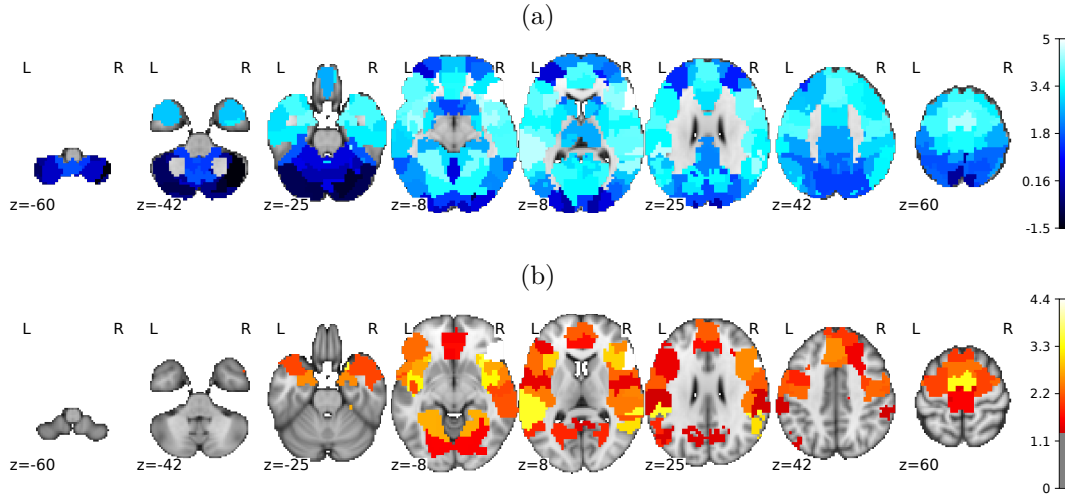


Figure 6.4-13 – **Statistical analysis of the haemodynamic differences between middle-age (MA) and elderly (E) subjects.** (a): T-scores associated with the two-sample t -test between the distributions of haemodynamic dilation parameters in middle-age (MA) and elderly (E) subjects (null hypothesis $H_0 : \bar{\delta}_m^{\text{MA}} = \bar{\delta}_m^{\text{E}}, \forall m = 1, \dots, M$). Note that most of the T-values are positive meaning that $\bar{\delta}_m^{\text{MA}} > \bar{\delta}_m^{\text{E}}$ most often. (b): Thresholded statistical map ($-\log_{10} p_{\text{val}}$) associated with a two-sample t -test performed to assess the mean difference in terms of haemodynamic dilation parameter between the middle-age and elderly subjects. The p-values were Bonferroni corrected for multiple comparisons performed across all ROIs ($M = 96$). The map (axial slices) was thresholded at a significance level of $\alpha = 0.05$ corresponding to $p_{\text{val}} = 1.65$ on the color bar.

log-transformed p-values, i.e. $-\log_{10} p_{\text{val}}$ (shown in Fig. 6.4-13(b)), after correcting for multiple comparisons using the Bonferroni correction across the $M = 96$ ROIs. We noticed first that a large majority of significant brain regions appear bilaterally indicating larger haemodynamic dilation parameters or shorter TTPs in younger individuals. Second, the negative t -values reported in the cerebellum are not statistically significant after correcting for multiple comparisons.

Then, in an attempt to be exhaustive we constructed three different logistic regression (LR) models based either on the individual (i) neural activity signals, (ii) HRF shapes and (iii) haemodynamic dilation parameters. We did not consider the spatial activity maps as input features in these models as they do not permit to perform dimension reduction. We trained these LR models using the Scikit-Learn software [Pedregosa et al., 2011] to predict the age label for each subject (1 for the elderly people, 0 for middle-aged people). A ℓ_2 -norm regularization was used in the estimation of the LR model parameters with an hyper-parameter $\beta > 0$. We grid-searched the temporal regularization parameter λ_f and the classifier hyper-parameter β , see Fig. 6.4-14, for the stability of the setting. We assess the impact of the choice of λ_f on the prediction results from Section 6.4.2. Fig. 6.4-14 reports the accuracy score for the logistic regression relatively to the choice of regularization parameter β for the classification model and the temporal regularization parameter λ_f for our deconvolution model. The accuracy is almost not impacted by the choice of parameter λ_f , for any value of β . This observation confirms that the choice of λ_f is not critical when studying the relative spatial structure of the haemodynamic delay and that our model can be used in practical cases to evaluate abnormalities in the haemodynamic response.

We chose the accuracy as the classification metric and cross-validated the score to provide




Figure 6.4-14 – Evolution of the accuracy score *w.r.t* the logistic regression regularization parameter β and the temporal regularization parameter λ_f . The accuracy score is not impacted by the hyper-parameter λ_f , as moving this parameter mainly impact the magnitude of the estimated delays and not its spatial structure.

an estimation of the generalization error with a 10 times repeated stratified 4-fold split.

Fig. 6.4-15(a) shows that the haemodynamic properties have an improved prediction power to discriminate the age compared to the neural activation signals (i.e. temporal components). The latter actually reaches an average accuracy score of 0.557, whereas the mean accuracy associated with the HRF shape and haemodynamic dilation parameter estimates respectively goes up to 0.741 and 0.743. Also, the distribution of accuracy scores across trials is more concentrated for the HRF dilation parameters compared to the whole HRF shape. This is likely due to the dimension reduction operated to extract this parameter which fluctuates less than the complete profile of the haemodynamic response. This analysis thus demonstrates that our decomposition is able to capture the brain age based on neurovascular information. However, due to the large between-subject variability even within each class of age the neural activation signals do not define a good feature for the brain age prediction. Complementary to that, Fig. 6.4-15(b) illustrates the progression of the mean accuracy score with the number of individuals involved in the LR model and clearly depicts that a plateau is reached around 459 subjects (the total size of the sampled cohort) both for the haemodynamic dilation parameter. Also, one can see the rapid progression of the mean accuracy with the number of individuals for the LR model based on haemodynamic properties compared to the one constructed from the neural activation signals.

Overall, this experiment has permitted to demonstrate that haemodynamic features are a good biomarker of the normal aging, as already reported in the literature [Grady and Garrett, 2014, West et al., 2019]. Moreover, it highlighted that the inter-hemispheric asymmetry in neurovascular coupling brings key information to discriminate middle-age from elderly people.

6.5 Conclusion

In this chapter, we have presented a semi-blind deconvolution approach to jointly estimate the haemodynamic response function and the neural activity signals across the whole brain.

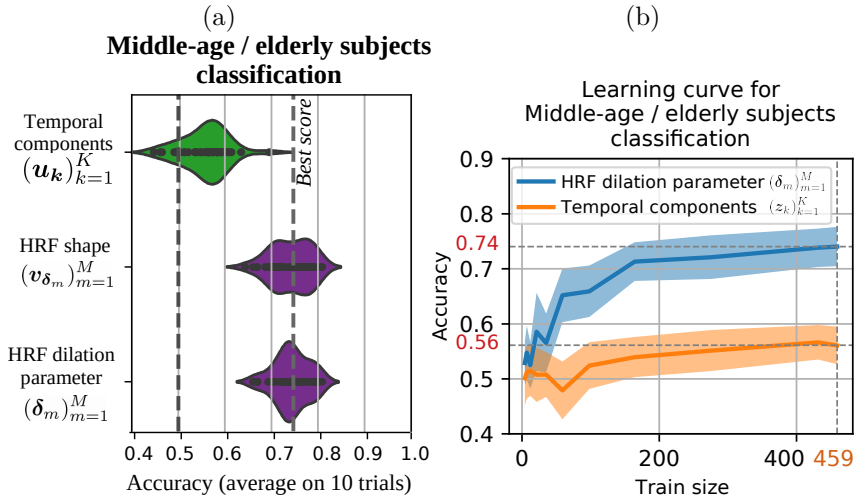


Figure 6.4-15 – **Accuracy score for classifying middle-age vs. elderly subjects (459 individuals sampled from the UK Biobank database).** (a): The prediction was performed by pulling individual features either based on (i) the estimated neural activation signals $(\hat{\mathbf{u}}_k)_{k=1}^{20}$, (ii) HRF shape estimates $(\hat{\mathbf{v}}_{\delta_m})_{m=1}^{96}$ or (iii) the haemodynamic dilation parameters $(\hat{\delta})_{m=1}^{96}$. From top to bottom, the distribution of the classification scores is shown from for the predictive features (i)-(iii), respectively. The best accuracy scores (average 0.74) are reached using the haemodynamic parameters and the smallest variability in the prediction using specifically the dilation parameter estimates. (b): Learning curve of accuracy scores as a function of the number of individuals (middle-age vs elderly subjects) used for the training stage both for the haemodynamic dilation parameter (blue curve) and the temporal components (orange curve). As a plateau is reached for 459 people, we presented the corresponding performances in panel (a).

As the proposed methodology is paradigm-free, it enables the analysis of resting-state fMRI data in an semi-supervised manner as the regularization parameters (K, λ_f) may be tuned using a trade-off between model accuracy and complexity. Beyond the model validation on synthetic and real fMRI data, we have demonstrated the interest of the proposed approach in two applications in neuroscience. Both aimed at characterizing cerebral haemodynamic delays in specific populations, namely stroke patients and elderly people by contrasting them with healthy and younger controls. Most importantly, we proposed an haemodynamic asymmetry index to lateralize the stroke episode while confirming the presence of a prolonged haemodynamic delay in these patients. We also demonstrated that haemodynamic properties are predictable of brain age. Finally, this new framework opens the door to new research avenues for functional connectivity analysis based on the neural input signals instead of the BOLD signal themselves. In contrast to existing methods [Wu et al., 2013], our approach would be less biased by a constant haemodynamic response shape across the whole brain.

* * *
 * *
 *

Clinical application

Chapter Outline

7.1	Analysis of the Synchronpioid cohort	115
7.1.1	Synchronpioid experimental protocol	115
7.1.2	Description of the statistical analysis	116
7.1.3	Preliminary results	117
7.2	Summary	121

THIS chapter proposes to apply our neurovascular decoupling method to the Synchronpioid cohort as a means to investigate the pharmacological mechanism associated with the buprenorphine. We will detail here the acquisition protocol running for the Synchronpioid project, the statistical analysis we deployed and finally preliminary results.

7.1 Analysis of the Synchronpioid cohort

As described in [section 1.4.2](#), to understand and better cope with the side-effects of the buprenorphine in a patient, the Synchronpioid project investigates the pharmacological mechanism associated with this molecule in healthy volunteers. In this chapter, we propose to assess whether the method developed in [chapter 6](#) allows us to uncover the effect of the buprenorphine on the neurovascular coupling and in particular on the HRF properties.

7.1.1 Synchronpioid experimental protocol

The participants received a [C^{11}]buprenorphine of 10.0 μ g injection during the acquisition while laying still in the 3T PET/MR scanner. The control subjects received an additional dose of NaCl at 0.9% while the other volunteers received an intravenous dose of buprenorphine of 0.2 mg. Up to the present time, 16 healthy volunteers participated in this experimental protocol, 10 participants had gone through a PET imaging session and only 4 subjects benefited from a high temporal resolution rs-fMRI acquisition. Within those four subjects, the control group has an average age of 25.5 years with a standard deviation (std) of 2.0 years and the group with an analgesic dose has an average age of 22.5 years with a std of 5.0 years. For these four participants, the overall acquisition contains five rs-fMRI imaging runs. The

third and fifth runs are of particular interest as they have been collected approximately 40min and 75min after injection (AI) using an multi-band acquisition sequence [Moeller et al., 2010], allowing us to get high temporally resolved fMRI data ($TR=0.8s$), see Figure 7.1-1. These fMRI acquisitions last 14 min during which a total of 1065 $T2^*$ -weighted scans was accumulated.

Moreover, the TEP protocol ran for 90min of acquisition along which arterial sampling was performed whenever possible as a means to measure the pharmacokinetic properties of the buprenorphine. The $[C^{11}]$ buprenorphine has a radioactivity of 4MBq/Kg for a total of 400MBq at maximum. In practice, the four participants of interest received in average 276.6MBq during the PET acquisition.




Figure 7.1-1 – **Description of the Synchronioid acquisition setup:** We display the steps of interest in the protocol. We have simultaneously collected PET and rs-fMRI data. The PET imaging session lasts 90min while the fMRI imaging session features multiple consecutive runs. We have a structural MRI and a B_0 field map acquisition along with the *re-blips* acquisitions that allow us to correct for the distortion due to the B_0 inhomogeneities. In addition, we collect two temporally resolved rs-fMRI acquisitions (the third and the fifth runs referenced to as rs-fMRI N°1 and N°2 hereafter), respectively at approximately 40min and 75min after injection.

7.1.2 Description of the statistical analysis

Here we focus on the analysis of rs-fMRI data to investigate the effect of the buprenorphine on the neurovascular coupling. We rely on the fMRIPrep [Esteban et al., 2019] tool to perform the classical preprocessing steps (see subsection 1.3.2 for more details). We analyse the four available subjects (two controls and two participants with an analgesic dose). We applied the multivariate decomposition introduced in chapter 6 to each subject and more specifically to the two rs-fMRI imaging sessions outlined in Figure 7.1-1. We used three levels of temporal regularization $\lambda_f \in \{0.001, 0.1, 0.9\}$ and $K = 20$. For visualization purpose, we back-projected the HRF dilation parameter estimates δ associated with each regularization level onto the MNI template and reported as well the histogram of δ in the each subject. As described in chapter 6, our approach introduces an estimation bias but accurately captures the spatial contrast of the haemodynamic response. Thus, after inspecting the three decomposition results (not shown here), we confirm that they all draw to the same findings and conclusions. Consequently, in what follows, we only report the results for $\lambda = 0.001$.

To localize the distribution of $[C^{11}]$ -buprenorphine binding in the brain from PET data, we computed the Standardized Uptake Value (SUV) normalized in Regions of Interest (ROI) between the 60 and 90min AI. The SUV spatially summarizes the activity concentration of $[C^{11}]$ buprenorphine in tissue normalized by the total injected radioactivity. We consider this timing window to take the temporal stability of the pharmacokinetic at stake into account.

The objective of this preliminary analysis was first to localize which brain regions elicit the most concentrated fixation of buprenorphine and second to characterize potential effects of buprenorphine on the haemodynamic response function.

7.1.3 Preliminary results

In order to compare our results obtained in rs-fMRI with the regional distribution of $[C^{11}]$ buprenorphine uptake which mainly reflects the distribution of μ -opioid receptors, we analyzed the PET data collected in four volunteers who received a placebo dose. Those participants did not benefit from the fine temporal resolution for the rs-fMRI imaging session, hence only their PET data have been processed. The SUV were calculated for all volumes on interest in each volunteer. To compare the intersubject spatial distribution of the radiotracer, SUVr were thereafter calculated by normalizing SUV values by the radioactivity in the Occipital cortex, considered as a suitable reference region for μ -opioid receptors [Hiller and Fan, 1996]. In Figure 7.1-2, we show that the strongest fixation appears in the insula $(-40/40, 16, -2)$, the putamen $(-30/30, 0, 0)$, the thalamus $(-10/10, -18, 4)$, the frontal cortex and the cingulate cortex. As buprenorphine is a μ -receptor agonist and κ -receptor antagonist with a highly stronger affinity for mu opioid, findings of these uptakes in cortical and sub-cortical regions are consistent with previous literature on the distribution of mu opioid receptor in the human brain [Henriksen and Willoch, 2008, Lee et al., 2013, Kantonen et al., 2019].




Figure 7.1-2 – **Spatial distribution of $[C^{11}]$ buprenorphine uptake in a brain (axial slices) of a representative volunteer from the placebo group:** we display the SUVr, *i.e.* the SUV normalized by the radioactivity in the occipital cortex, to locate the fixation of the buprenorphine in the brain. We observe that the fixation is maximal in the insula $(-40/40, 16, -2)$, the putamen (subcortical nucleus) $(-30/30, 0, 0)$, the thalamus (subcortical nucleus) $(-10/10, -18, 4)$, the frontal cortex and the cingulate cortex.

Moreover, to quantify the buprenorphine fixation in some Regions of Interest (ROI), we report in Figure 7.1-3 the SUVr in the regions that features visually a high level of concentration of buprenorphine in Figure 7.1-2. We notice an increase of approximately 10% to 50%, the highest increase ($\approx 50\%$) being located in the Putamen followed by the Insula that features an increase of $\approx 40\%$.

In Figure 7.1-4, for each subject we spatially display the HRF dilation parameters δ back-projected onto the MNI template along with its histogram (rightmost column) for rs-fMRI N°1 that intervenes approximately 40min After Injection (AI). Similar results obtained for rs-fMRI N°2 that intervenes approximately 75min after AI are shown in Figure 7.1-5. In the histograms, we report in black dashed line the grand mean dilation parameter, averaged over the whole brain and across all participants for the two imaging sessions (rs-fMRI N°1 – 2). We also show in orange solid line each subject-specific mean dilation parameter averaged over

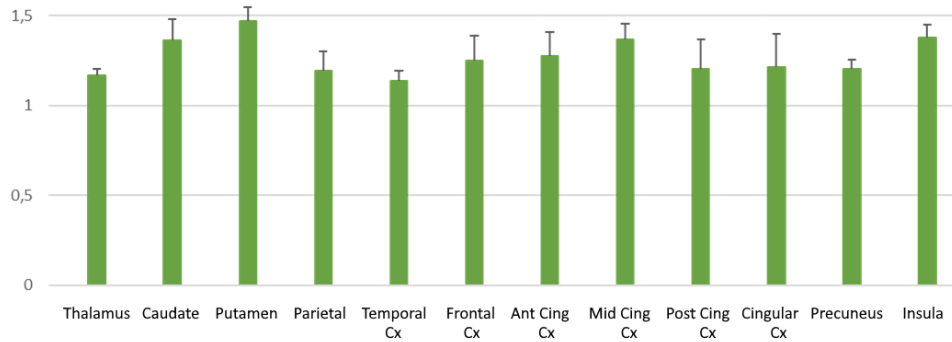


Figure 7.1-3 – **buprenorphine fixation within ROI**: we show the SUVr in the ROIs that feature a high level of buprenorphine concentration in Figure 7.1-2. We notice the highest fixation in the putamen with an increase of $\approx 50\%$, followed by the insula that features an increase of $\approx 40\%$.

the whole brain. The two control subjects (sub-02 and sub-03) are depicted with blue boxes while the participants (sub-01 and sub-04) that received an analgesic dose are highlighted with orange boxes.

Although the number of individuals is very limited, we distinguish a clear separation between the effects of placebo and the analgesic buprenorphine dose on the haemodynamic responses. We report a delayed haemodynamic response function (i.e. a larger haemodynamic delay corresponding to a smaller dilation parameter δ) in participants who received an analgesic dose as compared to the control subjects. This finding is consistent with the existing literature that assessed the effect of opiate substances in cerebral blood flow [Adriaens et al., 2014]. Also, the spatial maps of haemodynamic delays globally depict rather symmetrical haemodynamic responses. However, the strongest effects (i.e. slowest haemodynamic responses in dark blue) are observed in the brain regions (insula and putamen) that correspond to the highest SUVr values in PET imaging, as shown in Figure 7.1-2. Although preliminary, this novel result outlines the spatially localized effect of the buprenorphine on the haemodynamic response thanks to the combination of simultaneous PET and fMRI imaging.

In Figure 7.1-5, we noticed that the spatial symmetry of haemodynamic responses is preserved during the second imaging session (rs-fMRI N°2). Consistently with the results obtained during rs-fMRI N°1, we again observed the largest slowing down effect of the haemodynamic response in the insula. However, in sub-01 we observed the transient effect of buprenorphine as he almost recovered similar haemodynamic responses (i.e. larger dilation parameters δ) in comparison to the control subjects, especially sub-03. In contrast, sub-04 still shows lower dilation parameters, so slower haemodynamic profiles. Therefore, one can anticipate a certain vascular variability among the participants that received an analgesic dose. This is probably due to the pharmacokinetic of the molecule that may fluctuate from one subject to another and should be studied in-depth by cross-correlating these findings with the analysis of PET data using arterial samples.

Last, Figure 7.1-6 depicts the evolution of the haemodynamic dilation parameters δ for each participant between the two imaging sessions (rs-fMRI N°1 and rs-fMRI N°2). This graph permits to realize that a given rs-fMRI run did not happen at the same time points across individuals. Next, we noticed that, in average, the control subjects show

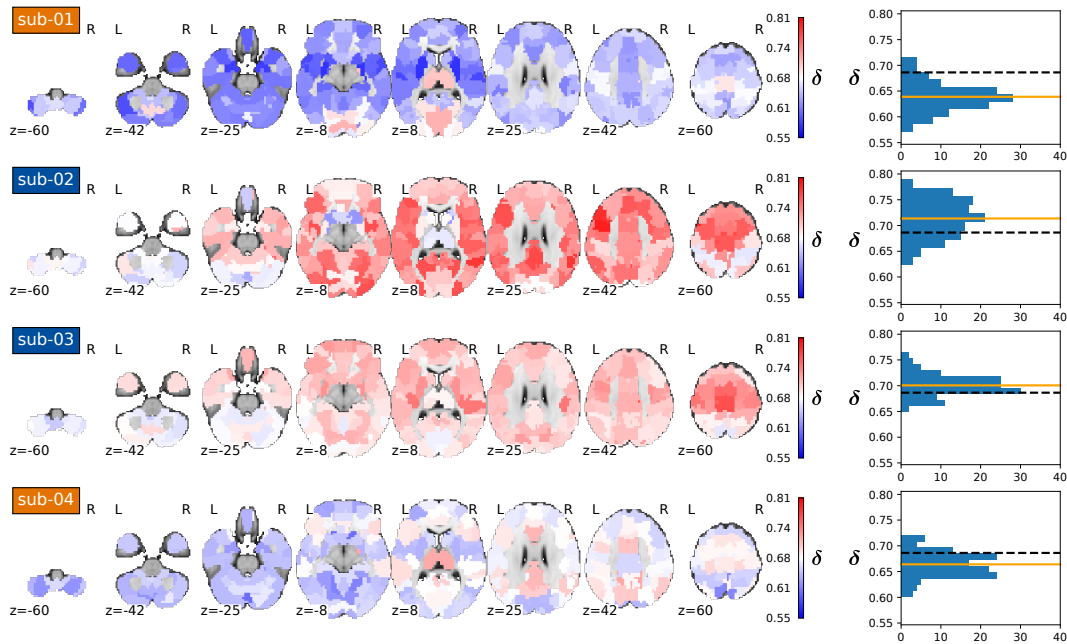


Figure 7.1-4 – **Maps of HRF dilation parameters for rs-fMRI N°1.** (approximately 40min AI): From top to bottom, we show for each subject, from feet to head (i.e. from left to right column), axial slices of HRF dilation parameters δ along with its histogram (rightmost column). On the latter, we report in black dashed line the grand average dilation parameter, computed over the whole brain and across the four participants for the two imaging sessions. We also display subject-specific whole brain mean of dilation parameters in orange solid line for this first imaging session. The control subjects are named with blue-colored boxes (sub-02 and sub-03) while the participants that received an analgesic dose of buprenorphine are named with orange-colored boxes (sub-01 and sub-04). The two groups are clearly separated one another based on their haemodynamic response spatial maps. The volunteers with an analgesic dose of buprenorphine show in average slower haemodynamic responses (i.e. smaller δ values in dark blue) as compared to the control subjects.

faster haemodynamic responses. In the control group, we found pretty consistent values of haemodynamic dilation parameters with a slight vertical shift of the distribution (i.e. larger dilation parameters measured) for sub-02 during the second run as compared to sub-03, which remained very stable across the two imaging sessions. In the group submitted to the analgesic dose of buprenorphine, we observed a discrepancy between the two individuals: Sub-01 recovered faster haemodynamic responses during the second run, getting even closer to sub-03. In contrast, sub-04 shows consistent and low values of haemodynamic dilation parameters, reflecting a slower vascular coupling. This observed variability in time and across individuals, illustrates the complexity of the mechanism that underlies the pharmacological effect of buprenorphine on the neurovascular coupling. In echo to this result, it would be interesting to cross-correlate individually this HRFs estimated with blood and/or PET data such as the blood concentration of buprenorphine, the rate constant K_1 – that represents the unidirectional transport of the tracer from blood to brain – or the rate of $[C^{11}]$ buprenorphine uptake that are indicative to the opioid receptor availability.

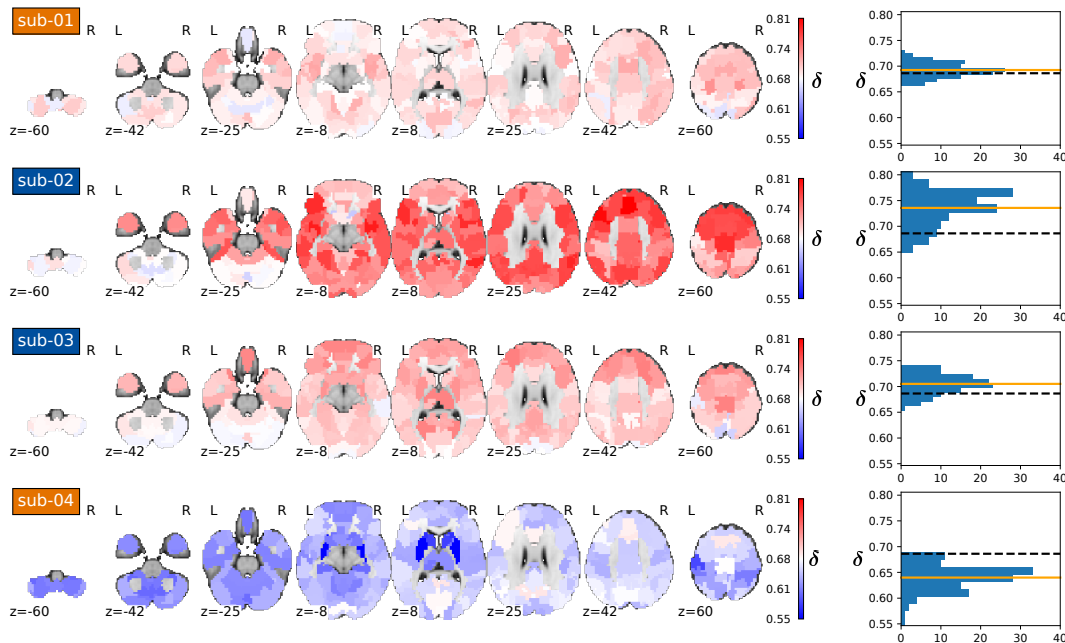


Figure 7.1-5 – **Maps of HRF dilation parameters for rs-fMRI N°2.** (approximately 75min AI): From top to bottom, we show for each subject, from feet to head (i.e. from left to right column), axial slices of HRF dilation parameters δ along with its histogram (rightmost column). On the latter, we report in black dashed line the grand average dilation parameter, computed over the whole brain and across the four participants for the two imaging sessions. We also display subject-specific whole brain mean of dilation parameters in orange solid line for this second imaging session. The control subjects are named with blue-colored boxes (sub-02 and sub-03), the others with orange-colored boxes (sub-01 and sub-04). We notice that during rs-fMRI N°2 sub-01 recovers similar neurovascular responses as compared to the control group whereas sub-04 who even slowed down from his first session.

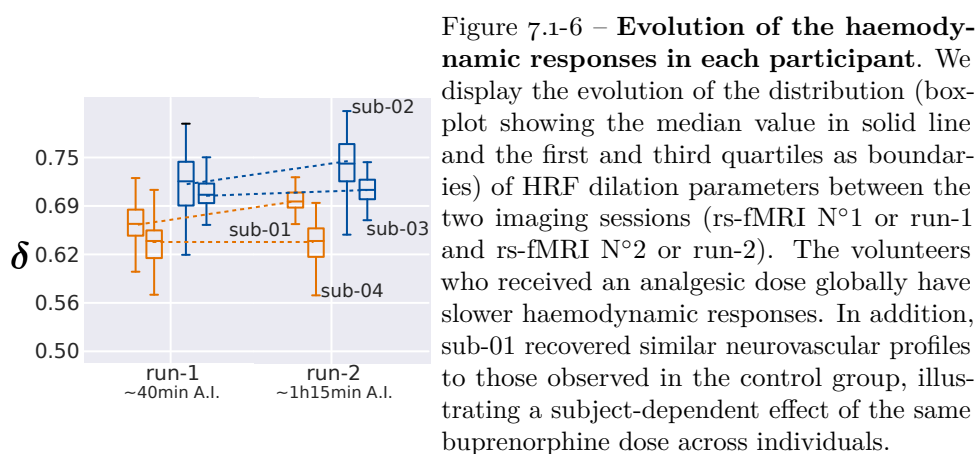


Figure 7.1-6 – **Evolution of the haemodynamic responses in each participant.** We display the evolution of the distribution (box-plot showing the median value in solid line and the first and third quartiles as boundaries) of HRF dilation parameters between the two imaging sessions (rs-fMRI N°1 or run-1 and rs-fMRI N°2 or run-2). The volunteers who received an analgesic dose globally have slower haemodynamic responses. In addition, sub-01 recovered similar neurovascular profiles to those observed in the control group, illustrating a subject-dependent effect of the same buprenorphine dose across individuals.

7.2 Summary

In this chapter, we analyzed the preliminary conjoint PET/fMRI data set collect in the course of the Synchronoid project at the Service Hospitalier Frederic Joliot. In particular, we applied our multivariate semi-blind deconvolution algorithm to temporally resolved fMRI data collected along two consecutive imaging runs in four individuals and we paid attention to the analysis of haemodynamic features. The analysis of neural signatures is left to future work. Importantly, we found out clear haemodynamic differences between the control and test groups, the former received a placebo while the latter the analgesic dose of buprenorphine. On these preliminary data, we showed that the buprenorphine has a global slowing down effect on haemodynamic features, most significantly in the insula where the fixation of the molecule is the highest. However, as the first rs-fMRI imaging session we analysed took place approximately 40min AI, we do not have a baseline BOLD activity for each subject. Getting this baseline could help us to better characterize the effect of the molecule over time. Moreover, this temporal evolution could be set in perspective with the pharmacokinetic of the drug to better understand the inter-subject variability. Of course, these preliminary results need to be confirmed on a larger cohort which is currently under construction. Last but not least, more classical fMRI data analysis will be performed (e.g. functional connectivity analysis) to measure complementary effects of the buprenorphine on brain function.

* * *
* *
*

General Conclusions and Perspectives

Contributions

In the thesis, we have developed new methods and tools to efficiently disentangle the neuronal and vascular activity components in the BOLD fMRI data, both in the context of resting-state or task related experiments. In particular, we have put strong efforts in the scalability of the proposed methods as a means to run whole brain analysis in a few minutes per individual.

Separating vascular and neuronal components in fMRI BOLD signals. Both neural and non-neural factors such as glial cell activity, cerebral energy metabolism and the cerebral vasculature contribute independently and synergistically to the fMRI BOLD signal. A mis-estimation of individual and regional HRFs may lead to interpret haemodynamic changes as neural variations, which could have considerable implications for the interpretability and reliability of findings in fMRI studies. Previous literature supports the hypothesis that HRF variability introduces some bias in fMRI data analysis [Rangaprakash et al., 2017, 2018, Yan et al., 2018]. Deconvolution-based approaches such as the one presented here allow the discrimination of cerebrovascular components from neural activity and minimize the confound of HRF variability in the exploration of brain physiology, functional connectivity and cognitive processes. In this work, we have developed a new algorithm that performs the joint estimation of the HRF and neural activation signal as a semi-blind deconvolution multivariate problem in a paradigm-free setting. Synthetic and real resting-state fMRI data allowed us to demonstrate that this approach is able to faithfully capture the individual's haemodynamic response function and intrinsic functional networks with low intra-subject variability and relative minimal impact of hyper-parameters on the reliability of HRF estimation. To go one step further in terms of validation, we spent a significant amount of time in assessing to what extent these methodological developments have a practical impact in terms of detection of pathology or classification of individuals (e.g. middle aged vs elderly). Further demonstrations using well defined research protocols could be envisaged to refine the tool for its use in clinical applications, especially if our intention is also to use it in the context of neurological disorders (e.g. epilepsy) where abnormal neural activity manifests through interictal spikes between consecutive seizures [Béнар et al., 2006, Zijlmans et al., 2007].

Fast minimization of Total Variation regularized problem. Our approach takes advantage of the Total Variation (TV) regularization to constrain the estimation of the K neural activation temporal atoms. To achieve a fast minimization, we have proposed a novel approach based on a neural network. We unrolled the Proximal Gradient Descent in the analysis formulation of our problem and specified two solutions to compute the proximal operator of the TV penalty term: (i) we defined a sub-network to approximate the operator or (ii) we solved it numerically and provided the formulation for its gradient. We demonstrated the higher performance of our approach on simulations and real fMRI data in the context of the BOLD deconvolution.

Whole brain analysis of haemodynamic properties. Generally, supervised HRF estimation methods explain the observed BOLD signal by focusing on a set of brain areas (e.g. visual, auditory and sensorimotor cortices) that are typically involved in a specific activation paradigm [Goutte et al., 2000, Marrelec et al., 2003, Handwerker et al., 2004, Lindquist and Wager, 2007, Vincent et al., 2010, Pedregosa et al., 2015]. To the best of our knowledge, the proposed approach is the first being able to perform whole brain analysis of haemodynamic properties using resting-state fMRI data in a paradigm-free manner. It is worth mentioning that we used a common brain parcellation – the Harvard-Oxford probabilistic atlas [Desikan et al., 2006] – across all individuals, to ensure the feasibility of group-level analyses and facilitate between-group comparisons. As such, the major findings we reported on the asymmetries in haemodynamic features between stroke patients and healthy controls, on one hand and the differences related to normal aging on the other hand, are dependent on this atlas and could slightly differ with another parcellation. This question is left for future research.

Interest for analyzing normal aging. Regional variability in the HRF is partly dictated by the size of surrounding blood vessels [Handwerker et al., 2004, Havlicek and Uludağ, 2020]. Vascular aging is known to cause progressive deterioration in the cellular structure of the blood vessel wall, with the development of arteriole tortuosity and reduction in capillary density that undoubtedly impact both resting-state cerebral blood flow and the ability to adjust it during neuronal activity. In that context, as concordant with various imaging studies on aging [Bangen et al., 2009], our algorithm has proven its sensitivity to classify middle-age vs elderly subjects on the basis of HRF parameters. Indeed, using recent rs-fMRI data findings obtained on the large-scale CamCAN task-fMRI data set [Shafto et al., 2014] our study confirmed a clear reductive process of the neurovascular coupling in elderly people in multiple brain areas (occipital, temporal and frontal regions). Beyond the age-related statistical comparison, we also validated on a large cohort (459 subjects) the estimated haemodynamic features as potential predictors of brain age in a supervised classification task. This confirms that these neural signatures are much more variable across individuals and even groups.

Interest for monitoring patients after a stroke episode. Our approach, using the proposed haemodynamic asymmetry index, has also proven its utility to detect at the individual level on a single with a history of stroke, very slow haemodynamic delays in a

restricted brain territory probably related to local ischemic tissue consecutive to stroke. This finding is perfectly consistent with the literature [Altamura et al., 2009] showing that the delay in peak latency that arises as patients advance from the acute to the subacute stroke phase is related to the deterioration of cerebral haemodynamics. Consequently, remodeling the fMRI haemodynamic response function in stroke patients may optimize the detection of BOLD signal changes. MRI is of course one of the most powerful diagnostic tools in contemporary clinical medicine. However, in the acute episode of stroke, diffusion-weighted MRI and perfusion imaging (e.g. ASL) remain the reference imaging modalities to perform the diagnosis in a noninvasive way. In the post-acute period, rs-fMRI acquisition equipped with the proposed method would be extremely valuable to predict the potential recovery of brain function as it no longer requires the patient involvement in an experimental paradigm.

Interest for analyzing the effect of a drug on patients. Glial cell activity, cerebral energy metabolism and the cerebral vasculature affect the neurovascular coupling. The presence of a specific drug in the brain would certainly interfere with their activity. Hence the HRF shape of the HRF would be also impacted and this change could be captured by the proposed methodology in this thesis. To this end, we preliminary demonstrated using PET imaging that the buprenorphine, an opioid used as analgesic, is detected and fixated in the brain (most significantly in the Insula) in subjects who received an analgesic dose. Concomitantly, based on the rs-fMRI data collected during the PET imaging protocol we provided evidence for a local increase in the haemodynamic time-to-peak both in the Putamen (subcortical nucleus) and the Insula in the same individuals. We thus exhibited that the buprenorphine slows down the haemodynamic system in these participants. However, we also pointed out a certain inter-subject variability that should be further studied in the future and correlated both with the pharmacokinetic of the drug through the analysis of arterial blood samples and the subject behavior/response to painful stimuli (keep your hand in an ice cube bag).

Limitations & Perspectives

Limitations. Some limitations of our tool do exist. First, there are free parameters in the proposed modeling (K , λ_f , M) that need to be set in an appropriate manner. We explored two model selection criteria for setting K , namely the R^2 score and the determinant of the correlation matrix between the neural activation signals. Based on these metrics, we found a fair compromise between accuracy and model complexity for $K = 20$. We thus constantly used this value hereafter in the individual decomposition. Of course, other model selection approaches might be envisaged to optimize K and λ_f using for instance a (widely) Bayesian information criterion (BIC) [Neath and Cavanaugh, 2012, Watanabe, 2013], or the log-likelihood in the standard classical framework. The selected model would thus be the one associated with the lowest BIC value or largest log-likelihood. More recently, the concept of bi-level optimization [Bennett et al., 2006] has emerged to set hyper-parameters. In this case, an upper-level cost function (e.g. a supervised training score on the features of the decomposition) has to be minimized with respect to the unknown hyper-parameters while staying intrinsically connected to the lower-level problem, namely the multivariate

decomposition. Because of the extra-computation cost required by these approaches, such aspects were left aside but future work could be devoted to address these key points to automate the processing.

Second, the proposed regional analysis is conditioned by the parcellation atlas (and the value of M). It would be interesting to deepen this research by testing the reproducibility of the tool with some atlas variations and the creation of an atlas using subject-specific assessment of the cerebral vasculature.

Third, to recover more structured spatial maps, an advanced regularization model based on TV-elastic net [de Pierrefeu et al., 2017] or structured sparsity [Jenatton et al., 2012, Baldassarre et al., 2012] could be used in space while keeping the same algorithmic structure. Our other contribution for solving the TV proximity operator [Cherkaoui et al., 2020b] could be combined and directly plugged into the current multivariate decomposition algorithm.

Fourth, as in standard multivariate data-driven methods, the inter-subject comparison of spatial maps is currently difficult in the proposed formulation. In the same spirit as group-ICA [Calhoun et al., 2009], canonical ICA [Varoquaux et al., 2010] or multi-subject dictionary learning [Varoquaux et al., 2011], the current within-subject decomposition could be extended to the group-level to become more stable. One possibility would be to impose the same spatial maps across all individuals like in Calhoun et al. [2009] while another more flexible approach would permit spatial variations around a group-level spatial template [Varoquaux et al., 2011]. In this context, the neural activation signals could remain subject-specific with large fluctuations both in timings and magnitudes. This kind of extension would permit to easily perform group-level analyses not only on the HRF parameters but also on the spatio-temporal decomposition.

Fifth, we experimentally observed both on numerical simulations and real fMRI data (ADHD cohort [Milham et al., 2012]) that a TR larger than 1s may be detrimental to a precise estimation of the haemodynamic dilation parameter. For that reason, all analyses were performed on fMRI acquisitions with short TR. This type of data is usually collected using simultaneous multi-slice imaging [Feinberg and Setsompop, 2013, Hesamoddin et al., 2019] to keep this parameter below 1s.

Sixth, because the proposed HRF model relies solely on a time dilation parameter, its magnitude is fixed and the fluctuations of the BOLD signal across the brain are thus captured through the neural activity atoms $(\mathbf{u}_k)_{k=1}^K$ on one hand and the spatial maps $(\mathbf{w}_k)_{k=1}^K$ on the other hand. However, the norm of the spatial maps being constrained, the real BOLD signal amplitude is captured by the neural activation signals. A recent work Tsvetanov et al. [2019] has shown that the amplitude of resting-state activity fluctuations is crucial to predict brain age in healthy subjects. One possible enhancement of the current model would be to add a magnitude parameter to each HRF. In that case, we should fix the scale ambiguity issue by setting the amplitude of the temporal atoms $(\mathbf{z}_k)_{k=1}^K$. This modification would significantly increase the computational complexity due to the calculation of the proximal operator associated with the new regularization term $g_u((\mathbf{u}_k)_k) = \lambda \sum_{k=1}^K (\|\nabla \mathbf{u}_k\|_1 + I_{\|\mathbf{u}_k\|_\infty = \alpha})$.

Last, thus far we have used the canonical HRF as the reference shape in \mathbf{v}_{ref} . This setting could be easily updated to perform investigations in specific populations (e.g. newborns) where the true haemodynamic response function is known to deviate from the canonical shape [Arichi et al., 2012].

Perspectives for applications. Such findings bring new opportunities for the exploration of brain plasticity and pathogenesis in humans. In this way, even older adults in relatively good health may have undetected, clinically silent vascular pathology and ischemic brain changes such as silent stroke [D’Esposito et al., 2003]. As this could affect the neurovascular coupling, it is of critical interest to assess the cerebrovascular function and to consider vascular risk factors in the pathogenesis or exacerbation of age-related degenerative diseases like Alzheimer. The current challenge for diagnostic imaging techniques is to find metrics that capture relevant information or biomarkers. Such works on haemodynamic deconvolution, which are not yet used routinely, might help uncover these biomarkers.

This present study constitutes a proof of concept in terms of interest and feasibility of the proposed approach. However, many other applications on clinically well-characterized populations could be undertaken to improve and demonstrate its robustness. Aside from stroke and neurological vascular diseases, vascular risk factors are associated with an increased risk of epilepsy and could represent a sizable proportion of cryptogenic cases of epilepsy [Ogaki et al., 2020]. Although abnormal neural activities generating interictal epileptiform discharges provoke haemodynamic changes and BOLD activation [Béнар et al., 2006, Zijlmans et al., 2007], standard MRI scans fail to visualize epileptic source precisely. Some authors have shown that standard HRF in the general linear model (GLM) framework can introduce errors on the spatial extent and localization of activated brain regions. Modeling the haemodynamic response function should provide some flexibility in the GLM definition and improve the sensitivity of fMRI data to epileptogenic areas [Storti et al., 2013]. This improvement is particularly valuable in epileptic patients with drug-resistant focal seizures, where resection of the epileptogenic brain area remains the best therapeutic outcome. In that context, approaches that employ haemodynamic deconvolution – as presented in this thesis – promise a more faithful investigation of the cerebral pathology as far as we are able to collect data very quickly (short TR). This demonstrates the need once again of technological breakthroughs in fMRI data acquisition that should come up soon [Chaithya et al., 2020] if one wants to preserve good spatial resolution with short TR.

* * *
* *
*

Appendices

Résumé en français

Abstract in French

Sujet : Estimation efficace cerveau entier de la réponse hémodynamique et déconvolution semi-aveugle de la réponse neurale via la régularisation par variation totale, à partir de donnée d'IRM fonctionnelle.

L'Imagerie par Résonance Magnétique fonctionnelle (IRMf) enregistre de manière non-invasive l'activité cérébrale en mesurant l'évolution du niveau d'oxygénation du sang. C'est le contrast BOLD. Ce dernier mesure de manière indirecte l'activité cérébrale au travers du couplage neurovasculaire [Ogawa et al., 1992]. Ce couplage est caractérisé par un système linéaire et invariant dans le temps, résumé par sa réponse impulsionnelle dénommée la Fonction de Réponse Hémodynamique (FRH) [Bandettini et al., 1993, Boynton et al., 1996]. L'estimation de cette fonction est particulièrement intéressante : un changement de cette FRH peut être lié au mécanisme pharmacologique d'une drogue, l'effet de l'âge sur l'organisme ou les conséquences d'une pathologie cérébrale. De plus, son estimation permet aussi de relier le signal BOLD observé à l'activité cérébrale (i.e. neurale) sous-jacente qui permet ainsi de mieux comprendre les mécanismes cognitifs en jeu.

De nombreuses méthodes ont proposé d'estimer cette fonction de réponse dans le cadre de données IRMf acquises au cours d'un protocole d'activation cérébrale (IRMf-t, i.e. de tâche). Dans ce contexte, le sujet volontaire est soumis durant l'acquisition à un Paradigme Expérimental (PE) qui alterne entre des phases de repos et de stimulation sensorielle, motrice ou cognitive [Friston et al., 1998a, Ciuciu et al., 2003, Lindquist and Wager, 2007, Pedregosa et al., 2015]. Généralement, les méthodes supervisées d'estimation de la FRH proposent d'adapter les paramètres d'un modèle prédéfini afin d'expliquer le signal BOLD observé associé au PE [Ciuciu et al., 2003, Lindquist and Wager, 2007, Vincent et al., 2010, Pedregosa et al., 2015]. Plusieurs limitations sont à souligner parmi les méthodes existantes. Premièrement, le PE est ici utilisé comme un substitut à l'activité cérébrale (i.e. neurale). Un biais peut ainsi être introduit dû, par exemple, à un potentiel retard entre la réponse neurale du sujet et la description temporelle des tâches dans le PE. De plus, ce type d'approches ne peut être utilisé dans le cas d'IRMf-r où aucun PE n'est défini *a priori*. Il s'agit en effet dans ce type d'acquisitions de suivre les fluctuations spontanées de l'activité cérébrale. De manière complémentaire, on retrouve depuis les années 90 une littérature riche concernant les

techniques de déconvolution du signal BOLD afin d’estimer le signal d’activation cérébrale sous-jacent à une échelle temporelle proche de la seconde [Glover, 1999, Gitelman et al., 2003, Hernandez-Garcia and Ulfarsson, 2011, Khalidov et al., 2011]. Une contribution clef, qui est à la base de ces travaux de thèse, est le travail de Karahanoğlu et al. [2013]. Ce dernier propose une modélisation spatio-temporelle de l’activité cérébrale en introduisant une régularisation qui promeut la parcimonie aussi bien temporelle que spatiale sur ce signal estimé. En procédant ainsi, les signaux d’activité cérébrale obtenus permettent de définir des réseaux fonctionnels, et en ce sens, se rapprochent de la méthode proposée par Wu et al. [2013]. Récemment, des travaux du même groupe sont venus prolonger ces précédentes contributions, tel que [Farouj et al., 2019] qui estime conjointement le profil de la FRH et l’activité cérébrale à l’aide d’une paramétrisation limitée. Ce type d’approche est souvent appelé *schéma de déconvolution semi-aveugle* du signal BOLD. Farouj et al. [2019] se base sur l’hypothèse d’un signal d’activité cérébrale constant par morceau ou dit “par bloc”, hypothèse préalablement introduite dans Karahanoğlu et al. [2013] mais est aussi capable d’estimer le paramètre du profil hémodynamique afin de lever l’ambiguïté d’amplitude et d’invariance temporelle entre le signal d’activation neurale et le filtre hémodynamique (i.e. la FRH).

En s’inspirant nous aussi des travaux réalisés par Karahanoğlu et al. [2013], nous proposons une nouvelle approche pour distinguer le signal d’activité neurale du couplage neurovasculaire (FRH) aussi bien pour l’IRMf de repos (IRMf-r) que celui dit de tâche (IRMf-t). En somme, l’objectif principal est de définir un modèle permettant de *capturer le couplage neurovasculaire et l’activité cérébrale à l’échelle du cerveau tout entier*, et de produire un algorithme d’optimisation rapide permettant l’estimation des différents paramètres en jeu sur d’importants volumes de données.

Notre modélisation de l’activité cérébrale se base aussi sur un signal constant par morceau ou dit “par bloc”, noté $\mathbf{u} \in \mathbb{R}^{\tilde{T}}$. De plus, nous notons par $\mathbf{v} \in \mathbb{R}^L$ la FRH. Avec le modèle linéaire et invariant dans le temps du couplage neurovasculaire adopté [Bandettini et al., 1993, Boynton et al., 1996], nous modélisons le signal BOLD observé $\mathbf{y} \in \mathbb{R}^T$ comme suit :

$$\mathbf{y} = \mathbf{v} * \mathbf{u} + \epsilon , \quad (\text{A.1})$$

où $\epsilon \in \mathbb{R}^T$ correspond au bruit observé au sein du signal BOLD. Ce modèle est illustré par Figure A-1.

Il est à noter qu’avec Équation A.1, on ne considère le signal BOLD que dans un voxel, on parle alors de modèle univarié. Ainsi, dans le cas d’une analyse cerveau entier, avec P voxels, il faudrait considérer $\mathbf{Y} \in \mathbb{R}^{P \times T}$ tel que $\mathbf{Y} = (\mathbf{y}_j)_{j=1}^P$. Une des contributions majeures de cette thèse, que nous décrivons plus loin, se propose d’étendre ce modèle de manière multivariée afin d’étudier le cerveau entier mais en utilisant un modèle dont le nombre de paramètres reste raisonnable, i.e. ne croît pas linéairement avec le nombre de voxels.




FIGURE A-1 – **Illustration de la modélisation du signal BOLD.** Le signal BOLD est obtenu par convolution du signal neural (*en rouge*) avec la Fonction de Réponse Hémodynamique (*en bleue*) et en rajoutant du bruit (*en jaune*).

L'hypothèse d'un signal constant par morceau ou dit "par bloc" pour l'activité neurale \mathbf{u} est classiquement imposée grâce à la régularisation par Variation Totale (TV) en 1-dimension (1D). Cette régularisation promeut la parcimonie de la dérivée première du signal $\mathbf{z} = \mathbf{D}\mathbf{u}$, tel que $\|\mathbf{u}\|_{TV} = \|\mathbf{D}\mathbf{u}\|_1 = \sum_{i=2}^n |u_i - u_{i-1}|$, ce qui force le signal \mathbf{u} à être constant par morceaux. Ici, \mathbf{D} est l'opérateur de différentiation discret, défini comme $\mathbf{D} \in \mathbb{R}^{(n-1) \times n}$ avec :

$$\mathbf{D} = \begin{bmatrix} -1 & 1 & 0 & \dots & 0 \\ 0 & -1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix}$$

$$\arg \min_{\mathbf{u} \in \tilde{\mathbb{R}}^T} \frac{1}{2} \|\mathbf{y} - \mathbf{v} * \mathbf{u}\|_2^2 + \lambda \|\mathbf{D}\mathbf{u}\|_1, \quad (\text{A.2})$$

avec $\lambda \in \mathbb{R}^+$ le paramètre de régularisation TV.

La résolution d'un problème de déconvolution avec une régularisation TV (Équation A.2) peut être coûteuse en temps de calcul. Ainsi, nous avons comparé théoriquement les différentes formulations associées au problème Équation A.2 et démontré que l'algorithme issu de la formulation à l'analyse est plus efficace que la formulation à la synthèse. En effet, une comparaison des taux de convergence des différentes formulations, $\tilde{\rho}$ pour la formulation à la synthèse et ρ pour la formulation à l'analyse, démontre que le ratio $\frac{\tilde{\rho}}{\rho}$ évolue au moins en $\Theta(k)$. De plus, nous proposons la conjecture qu'elle évolue en $\Theta(k^2)$ suite aux simulations expérimentales réalisées. Ainsi, en retenant la formulation à l'analyse, nous proposons de "dérouler" de manière différentiable les algorithmes de descente de gradient proximal (PGD) associés au problème d'optimisation avec régularisation TV. "Dérouler" un algorithme signifie définir un réseau de neurones artificiels dont chaque couche est équivalente à une itération de l'algorithme sélectionné. La différentiation à travers l'opérateur proximal TV peut être obtenue soit en calculant directement le gradient via une formule analytique, soit en le calculant via un second algorithme interne "déroulé". Des expérimentations sur des simulations numérique et une comparaison sur des données réelles d'IRMf (voir Figure A-4) donnent des résultats prometteurs concernant les algorithmes "déroulés" en comparaison aux algorithmes itératifs classiques. Ces travaux sur les algorithmes déroulés sont décrits dans le chapitre 4 et ont fait l'objet d'une publication [Cherkaoui et al., 2020b].

Afin d'étudier la réponse hémodynamique dans tout le cerveau, nous proposons, en plus de l'hypothèse d'un signal "bloc" concernant l'activité neurale, une modélisation originale et synthétique de la FRH \mathbf{v} . La réponse hémodynamique est définie par dilatation de la FRH canonique \mathbf{v}_{ref} , tel que $\mathbf{v}_\delta(t) = \mathbf{v}_{\text{ref}}(\delta t)$. Cette paramétrisation permet de résumer la FRH avec un seul paramètre scalaire. Dans la section 5.1 du chapitre 5 nous détaillons notre première contribution mettant en œuvre pour modéliser l'activation neurale l'hypothèse du signal "bloc", ainsi que le modèle de FRH par dilatation, afin d'estimer de façon univariée chacune de ces composantes à partir du signal BOLD. Cette contribution a fait l'objet d'une publication [Cherkaoui et al., 2019a]. Afin de réduire la variance d'estimation de ce modèle, nous proposons de considérer le signal BOLD cerveau entier de manière multivariée, nous obtenons alors :

$$\mathbf{Y} = \mathbf{v}_\delta * \mathbf{U} + \mathbf{E}, \quad (\text{A.3})$$



FIGURE A-2 – $\lambda_f = 0.1\lambda_{\max}$ FIGURE A-3 – $\lambda_f = 0.8\lambda_{\max}$

FIGURE A-4 – **Comparaison des performances pour** ($\lambda_f = 0.1$ gauche - $\lambda_f = 0.8$ droite) entre l’algorithme “déroulé” avec notre formule analytique de la dérivé du prox-TV et l’algorithme itératif de PGD pour la formulation à l’analyse dans le cas du problème de déconvolution du signal BOLD à FRH fixée et connue. Notre méthode présente de meilleures performances que l’algorithme itératif. On constate une légère dégradation des performances dans un contexte où la parcimonie est plus élevée.

avec $\mathbf{U} = (\mathbf{u}_j)_{j=1}^P \in \mathbb{R}^{P \times \tilde{T}}$, $\mathbf{E} = (\epsilon_j)_{j=1}^P \in \mathbb{R}^{P \times T}$ et $\bar{*}$ la convolution ligne à ligne telle que $\mathbf{v}_\delta \bar{*} \mathbf{U} = [\mathbf{v}_\delta * \mathbf{u}_1, \dots, \mathbf{v}_\delta * \mathbf{u}_P]^\top$. Une limitation directe de ce modèle est qu’il ne considère qu’une seule FRH pour tout le cerveau. Or, les caractéristiques du couplage hémodynamique varient suivant la région du cerveau ou l’individu considéré [Handwerker et al., 2004, Badillo et al., 2013]. Cet aspect pousse à estimer localement la FRH dans le cerveau. Une approche adaptée est de se baser sur une parcellisation pré-existante du cerveau en régions fonctionnellement homogènes [Varoquaux and Craddock, 2013]. Idéalement, le nombre de régions définies devraient être grand, cependant, plus on augmente le nombre de régions, plus le nombre de voxels par région est faible, ce qui peut fragiliser la stabilité de l’estimation de la FRH. Pour cette raison, nous considérons des parcellisations d’une centaine de régions telle que l’atlas probabiliste d’Harvard-Oxford [Desikan et al., 2006] avec lequel nous seuillons les probabilités liées aux régions pour obtenir une parcellisation. Nous introduisons donc mathématiquement la parcellisation $(\Theta_m)_{m=1}^M$ à M région telle que $\Theta_m \in \{0, 1\}^P$ avec $\Theta_{mi} = 1$ si le i -ième voxel appartient à la m -ième région et 0 sinon. En introduisant $(\Theta_m)_{m=1}^M$ dans Équation A.3, on obtient :

$$\mathbf{Y} = \left(\sum_{m=1}^M \Theta_m^\top \mathbf{v}_{\delta_m} \right) \bar{*} \mathbf{U} + \mathbf{E} , \quad (\text{A.4})$$

Ici, $\mathbf{U} = (\mathbf{u}_j)_{j=1}^P$ capture de manière univariée l’activité cérébrale pour chaque voxel j . Afin de réduire le nombre de paramètres du modèle et de rendre leur estimation plus robuste, nous introduisons une contrainte de rang faible en introduisant K cartes spatio-temporelles avec $K \ll P$. Chaque carte encode la signature temporelle $\mathbf{u}_k \in \mathbb{R}^{1 \times \tilde{T}}$ d’un réseau fonctionnel (ou l’activité neurale associée) ainsi que sa localisation dans le cerveau $\mathbf{w}_k \in \mathbb{R}^{1 \times P}$ de telle sorte que $\mathbf{U} = \sum_{k=1}^K \mathbf{w}_k^\top \mathbf{u}_k$. Finalement notre modèle est défini par :

$$\mathbf{Y} = \left(\sum_{m=1}^M \Theta_m^\top \mathbf{v}_{\delta_m} \right) \bar{*} \left(\sum_{k=1}^K \mathbf{w}_k^\top \mathbf{u}_k \right) + \mathbf{E} . \quad (\text{A.5})$$

Ce modèle multivarié a d’abord fait l’objet d’un développement avec une FRH fixée pour le cerveau entier qui est décrit dans la section 5.2 du chapitre 5 et a fait l’objet d’une

publication [Cherkaoui et al., 2019b]. Puis, l'extension comprenant l'estimation de plusieurs FRHs pour le cerveau entier, permettant ainsi la déconvolution semi-aveugle du signal BOLD est établi et décrit dans chapitre 6 et a fait l'objet d'une publication [Cherkaoui et al., 2020a]. Afin de démontrer la pertinence de notre approche en terme d'analyse intra-sujet, nous décomposons les données IRMf-r d'un sujet issu de la cohorte UK BioBank [Sudlow et al., 2015]. Figure A-5 permet d'identifier deux réseaux fonctionnels connus (le cortex visuel (a)-droite ainsi que le réseau du mode par défaut (b)-droite). Nous observons aussi une certaine variabilité concernant la réponse vasculaire au sein de ce sujet comme en témoigne la différence entre la FRH la plus rapide (c) localisée dans le gyrus temporal médian et celle la plus lente (d) localisée dans le cortex orbitofrontal.

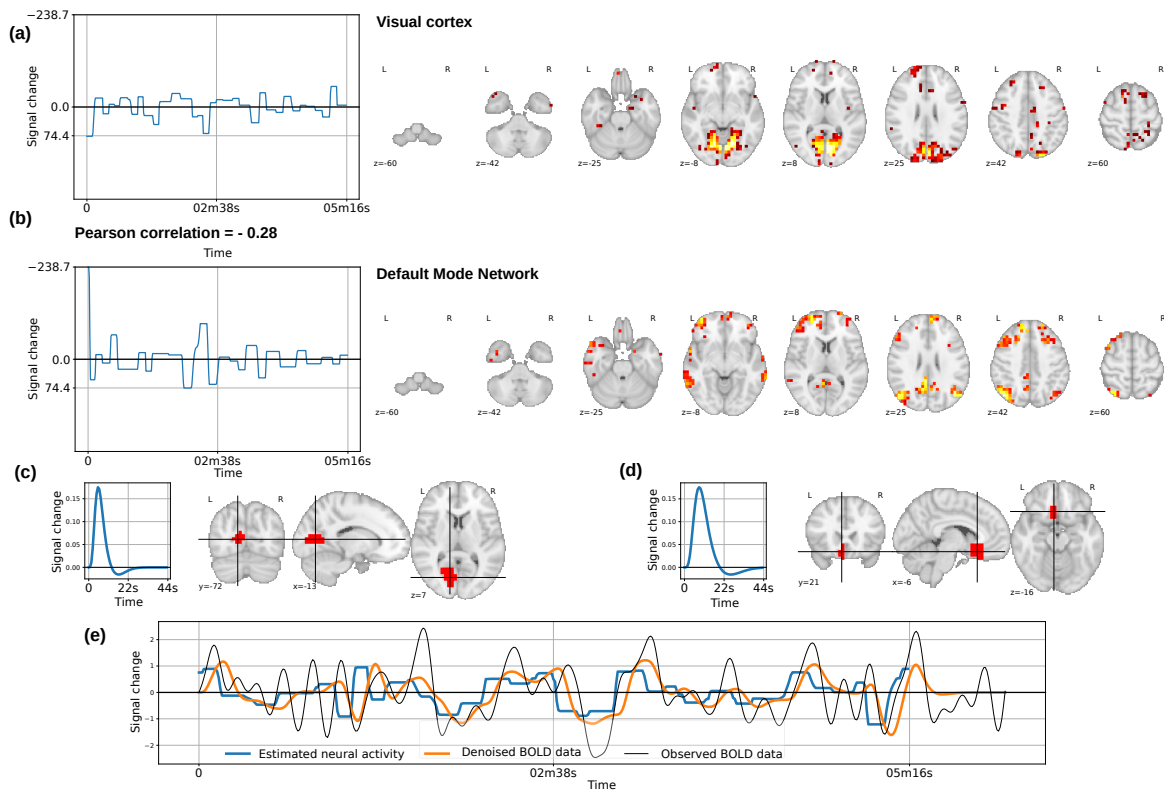


FIGURE A-5 – Résultat de la décomposition par déconvolution semi-aveugle proposée à partir des données d'IRMf-r d'un seul sujet. Rangée supérieure (a) : Activation neurale \hat{u}_{10} (gauche) ainsi que la carte spatiale correspondante (vue axiale) \hat{w}_{10} (droite), qui implique principalement une activation dans le cortex visuel. Rangée centrale (b) : Activation neurale \hat{u}_2 (gauche) ainsi que la carte spatiale correspondante (vue axiale) \hat{w}_2 (droite), qui implique principalement une activation du réseau du mode par défaut (DMN). Troisième rangée, gauche (c) : Région à la réponse hémodynamique la plus rapide. FRH la plus rapide estimée \hat{v}_{δ_f} (gauche) localisée dans le gyrus temporal médian comme illustré par le masque de la région associée Θ_f (droite). Troisième rangée, droite (d) : Région à la réponse hémodynamique la plus lente. FRH la plus lente estimée \hat{v}_{δ_s} (gauche) localisée dans le cortex orbitofrontal comme illustré par le masque de la région associée Θ_s (droite). (Rangée inférieure (e) : Séries temporelles associées au voxel. Signal d'activation neurale estimé (en bleu), signal BOLD débruité (en orange) calculé comme la convolution avec la FRH de la région concernée. Le signal BOLD observé est illustré en noir.

De plus nous illustrons la pertinence de notre décomposition en terme d'analyse de groupe en démontrons statistiquement qu'une pathologie telle que l'accident vasculaire

cérébral (AVC) et la dégradation vasculaire due à l'âge induisent un retard hémodynamique plus important au sein de certaines régions cérébrales et que ces caractéristiques peuvent être utilisés pour prédire l'état de santé de l'individu au moyen de l'apprentissage automatique (tâche de classification). En effet, comme le montre [Figure A-6](#) nous parvenons à différencier les sujets dit d'âge moyen de ceux dit âgés seulement à partir du coefficient de dilatation des FRHs δ associées à chaque individu.

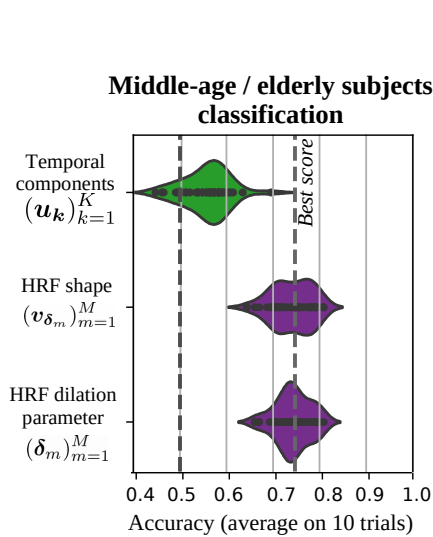



FIGURE A-6 – Score de classification des sujets d'âge moyen et ceux dit âgés (459 individus tirés aléatoirement à partir de la base de données UK Biobank). (a) :

La prédiction est réalisée en tirant de la décomposition par déconvolution semi-aveugle les caractéristiques suivantes (i) les K signatures temporelles estimées $(\hat{\mathbf{u}}_k)_{k=1}^{20}$, (ii) la forme des FRH estimées $(\hat{\mathbf{v}}_{\delta_m})_{m=1}^{96}$ ou (iii) les M coefficients de dilatation hémodynamique $(\hat{\delta})_{m=1}^{96}$. De haut en bas, les distributions des scores de classifications sont montrées pour les caractéristiques (i)-(iii), respectivement. Les meilleurs scores de précision (accuracy) (en moyenne 0.74) sont atteints en utilisant les caractéristiques hémodynamiques et la variance de prédiction la plus faible est atteint lorsqu'on utilise les paramètres de dilatation.

De plus, afin de démontrer l'utilité clinique d'un tel outil nous nous sommes intéressés au projet Synchropioïd. Ce projet vise à étudier les mécanismes pharmacologiques des opiacés (buprenorphine) afin de mieux comprendre et réduire les actuels effets secondaires notamment liés à l'habitué à ce type de molécules. Le protocole expérimental vise à définir une cohorte de 60 sujets dont 30 contrôles. Les sujets contrôles reçoivent une dose traceuse de la buprenorphine afin que sa fixation dans le cerveau puisse être observée grâce au scanner PET. Les 30 autres sujets reçoivent en plus une dose analgésique de buprenorphine. Conjointement et simultanément au scanner PET les sujets bénéficie d'un IRMf de repos afin d'identifier les potentielles modifications de leur activité cérébrale, i.e. hémodynamique. Ainsi, nous avons estimé les FRH de quatre sujets de la cohorte (deux contrôles et deux sujets ayant reçu une dose analgésique) et reportons dans la [Figure A-7](#) l'évolutions des distributions du paramètre de dilatation des FRH (δ) pour chaque sujet. Nous notons que les sujets ayant reçu une dose analgésique ont une réponse hémodynamique plus lente. De plus, une inspection plus approfondie de la localisation de la fixation de buprenorphine et des ralentissements les plus prononcés en terme de réponse hémodynamique montre qu'elles sont cohérentes spatialement et localisées au niveau du Putamen et de l'Insula.

Ainsi, grâce à cette thèse, nous avons proposé une avancée significative concernant l'estimation de la FRH dans tout le cerveau à partir de données d'IRMf de repos ou d'activation et ouvrons ainsi la voie à de nouvelles opportunités en recherche clinique, qui



pourrait bénéficier ainsi d'une estimation plus précise de la réponse hémodynamique.

* * *
* *
*

Bibliography

- P. Ablin, T. Moreau, M. Massias, and A. Gramfort. Learning step sizes for unfolded sparse coding. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 13100–13110, 2019. (Cited on page 56.)
- A. Adriaens, K. Peremans, T. Waelbers, E. Vandermeulen, S. Croubels, L. Duchateau, A. Dobbeleir, K. Audenaert, J. Eersels, S. Vermeire, B. De Spiegeleer, and I. Polis. The effect of morphine on regional cerebral blood flow measured by 99mTc-ECD SPECT in dogs. *PLoS One*, 8, 2014. (Cited on page 118.)
- G. K. Aguirre, E. Zarahn, and M. D’Esposito. The variability of human, bold hemodynamic responses. *NeuroImage*, 8:360–369, 1998. (Cited on pages 31 and 32.)
- F. Alfaro-Almagro, M. Jenkinson, N. K. Bangerter, J. L. R. Andersson, L. Griffanti, G. Douaud, S. N. Sotiropoulos, S. Jbabdi, M. Hernandez-Fernandez, D. Vidaurre, M. Webster, P. McCarthy, C. Rorden, A. Daducci, D. C. Alexander, H. Zhang, I. Dragonu, P. M. Matthews, K. L. Miller, and S. M. Smith. Image Processing and Quality Control for the first 10,000 Brain Imaging Datasets from UK Biobank. *NeuroImage*, 166:400–424, 2018. (Cited on page 72.)
- C. Altamura, M. Reinhard, M.-S. Vry, C. P. Kaller, F. Hamzei, F. Vernieri, P. M. Rossini, A. Hetzel, C. Weiller, and D. Saur. The longitudinal changes of BOLD response and cerebral hemodynamics from acute to subacute stroke. A fMRI and TCD study. *BMC neuroscience*, 10(1):151, 2009. (Cited on page 125.)
- B. M. Ances, C. L. Liang, O. Leontiev, J. E. Perthen, A. S. Fleisher, A. E. Lansing, and R. B. Buxton. Effects of aging on cerebral blood flow, oxygen metabolism, and blood oxygenation level dependent responses to visual information stimulation. *Human Brain Mapping*, 30:1120–1132, 2009. (Cited on page 110.)
- T. Arichi, G. Fagiolo, M. Varela, A. Melendez-Calderon, A. Allievi, N. Merchant, N. Tumor, S. J. Counsell, E. Burdet, C. F. Beckmann, et al. Development of BOLD signal hemodynamic responses in the human brain. *Neuroimage*, 63(2):663–673, 2012. (Cited on page 126.)
- L. Armijo. Minimization of functions having Lipschitz continuous first partial derivatives. *Pacific J. Math.*, 16(1):1–3, 1966. (Cited on pages 43 and 84.)

- D. Asemani, H. Morsheddest, and M. A. Shalchy. Effects of ageing and alzheimer disease on haemodynamic response function: a challenge for event-related fMRI. *Healthcare Technology Letters*, 4:109–114, 2017. (Cited on page 1.)
- S. Badillo, T. Vincent, and P. Ciuciu. Group-level impacts of within-and between-subject hemodynamic variability in fMRI. *Neuroimage*, 82:433–448, 2013. (Cited on pages 92, 109, and 134.)
- C. Bakhous, F. Forbes, T. Vincent, L. Chaari, M. Dojat, and P. Ciuciu. Adaptive experimental condition selection in event-related fMRI. In *International Symposium on Biomedical Imaging*, pages 1755–1758, 2012. (Cited on page 35.)
- L. Baldassarre, J. Mourao-Miranda, and M. Pontil. Structured sparsity models for brain decoding from fMRI data. In *2012 Second International Workshop on Pattern Recognition in NeuroImaging*, pages 5–8, 2012. (Cited on page 126.)
- P. A. Bandettini, A. Jesmanowicz, E. C. Wong, and J. S. Hyde. Processing strategies for time-course data sets in functional MRI of the human brain. *Magnetic resonance in medicine*, 30(2):161–173, 1993. (Cited on pages 1, 76, 131, and 132.)
- K. J. Bangen, K. Restom, T. T. Liu, A. J. Jak, C. E. Wierenga, D. P. Salmon, and M. W. Bondi. Differential age effects on cerebral blood flow and BOLD response to encoding: associations with cognition and stroke risk. *Neurobiology of aging*, 30(8):1276–1287, 2009. (Cited on page 124.)
- À. Barbero and S. Sra. Modular proximal optimization for multidimensional total-variation regularization. *The Journal of Machine Learning Research*, 19(1):2232–2313, 2018. (Cited on pages 45, 69, and 96.)
- A. Beck and M. Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009. (Cited on pages 19, 46, 56, 59, 78, and 84.)
- C. F. Beckmann and S. M. Smith. Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE transactions on medical imaging*, 23(2):137–152, 2004. (Cited on page 82.)
- C.-G. Bénar, C. Grova, E. Kobayashi, A. P. Bagshaw, Y. Aghakhani, F. Dubeau, and J. Gotman. EEG–fMRI of epileptic spikes: concordance with EEG source localization and intracranial EEG. *Neuroimage*, 30(4):1161–1170, 2006. (Cited on pages 123 and 127.)
- K. P. Bennett, J. Hu, X. Ji, G. Kunapuli, and J.-S. Pang. Model selection via bilevel optimization. In *The 2006 IEEE International Joint Conference on Neural Network Proceedings*, pages 1922–1929, 2006. (Cited on page 125.)
- Q. Bertrand, Q. Klopfenstein, M. Blondel, S. Vaiter, A. Gramfort, and J. Salmon. Implicit differentiation of Lasso-type models for hyperparameter optimization. In *International Conference on Machine Learning (ICML)*, pages 3199–3210, 2020. (Cited on pages 56, 65, and 67.)

- F. Bloch, W. Hansen, and M. Packard. Nuclear induction. *Physical Review Journals*, pages 69–127, 1946. (Cited on page 13.)
- S. Bollmann, A. M. Puckett, R. Cunnington, and M. Barth. Serial correlations in single-subject fmri with sub-second TR. *NeuroImage*, 166:152–166, 2018. (Cited on page 31.)
- M. Borgerding, P. Schniter, and S. Rangan. AMP-Inspired Deep Networks for Sparse Linear Inverse Problems. *IEEE Transactions on Signal Processing*, 65(16):4293–4308, 2017. (Cited on page 56.)
- S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011. (Cited on pages 44 and 56.)
- G. M. Boynton, S. A. Engel, G. H. Glover, and D. J. Heeger. Linear systems analysis of functional magnetic resonance imaging in human V1. *Journal of Neuroscience*, 16(13):4207–4221, 1996. (Cited on pages 1, 24, 76, 77, 82, 90, 131, and 132.)
- R. H. Byrd, P. Lu, and J. Nocedal. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific and Statistical Computing*, 16:1190–1208, 1995. (Cited on pages 43 and 78.)
- C. Caballero-Gaudes, D. Van de Ville, N. Petridou, F. Lazeyras, and P. Gowland. Paradigm-free mapping with morphological component analysis: getting most out of fMRI data. *Optical Engineering + Applications*, 2011. (Cited on page 35.)
- C. Caballero-Gaudes, F. I. Karahanoglu, F. Lazeyras, and D. Van De Ville. Structured sparse deconvolution for paradigm free mapping of functional MRI data. In *International Symposium on Biomedical Imaging*, volume 9, pages 322–325, 2012. (Cited on pages 35, 82, 83, 93, and 94.)
- C. Caballero-Gaudes, N. Petridou, S. T. Francis, I. L. Dryden, and P. A. Gowland. Paradigm free mapping with sparse regression automatically detects single-trial functional magnetic resonance imaging blood oxygenation level dependent responses. *Human Brain Mapping*, 34:501–518, 2013. (Cited on page 35.)
- V. D. Calhoun, J. Liu, and T. Adali. A review of group ICA for fMRI data and ICA for joint inference of imaging, genetic, and erp data. *Neuroimage*, 45(1):S163–S172, 2009. (Cited on page 126.)
- R. Casanova, S. Ryali, J. Serences, L. Yang, R. Kraft, P. J. Laurienti, and J. A. Maldjian. The impact of temporal regularization on estimates of the bold hemodynamic response function: a comparative analysis. *NeuroImage*, 40:1606–1618, 2008. (Cited on pages 31, 32, and 91.)
- L. Chaari, L. Forbes, T. Vincent, and P. Ciuciu. Hemodynamic-informed parcellation of fMRI data in a joint detection estimation framework. In *proceedings of International Conference on Medical Image Computing and Computer-Assisted Intervention*, volume 15, pages 180–188, 2012. (Cited on page 35.)

- G.R. Chaithya, P. Weiss, A. Massire, A. Vignaud, and P. Ciuciu. Globally optimized 3D SPARKLING trajectories for high-resolution T2*-weighted Magnetic Resonance Imaging. 2020. (Cited on page 127.)
- A. Chambolle. An Algorithm for Total Variation Minimization and Applications. *Journal of Mathematical Imaging and Vision*, 20(1/2):89–97, 2004. (Cited on pages 41, 45, 56, 58, 73, and 96.)
- A. Chambolle and T. Pock. A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011. (Cited on page 56.)
- L. Chaâri, J.-C. Pesquet, A. Benazza-Benyahia, and P. Ciuciu. A wavelet-based regularized reconstruction algorithm for SENSE parallel MRI with applications to neuroimaging. *Medical Image Analysis*, 15:185–201, 2011. (Cited on page 19.)
- P. L. Chebyshev. Théorie des mécanismes connus sous le nom de parallélogrammes. 1853. (Cited on page 60.)
- H. Cherkaoui, T. Moreau, A. Halimi, and P. Ciuciu. Sparsity-based Semi-Blind Deconvolution of Neural Activation Signal in fMRI. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1323–1327, 2019a. (Cited on pages 47, 72, 75, 82, 83, 91, 94, and 133.)
- H. Cherkaoui, T. Moreau, A. Halimi, and P. Ciuciu. fMRI BOLD signal decomposition using a multivariate low-rank model. In *European Signal Processing Conference (EUSIPCO)*, pages 1–5, 2019b. (Cited on pages 3, 72, 75, 90, 93, 95, and 135.)
- H. Cherkaoui, T. Moreau, A. Halimi, C. Leroy, and P. Ciuciu. Multivariate semi-blind deconvolution of fMRI time series. *submitted to NeuroImage*, 2020a. (Cited on pages 3, 89, and 135.)
- H. Cherkaoui, J. Sulam, and T. Moreau. Learning to solve tv regularised problems with unrolled algorithms. In *34th Conference and Workshop on Neural Information Processing Systems (NeurIPS)*, pages 1–21, 2020b. (Cited on pages 55, 126, and 133.)
- P. Ciuciu, J.-B. Poline, G. Marrelec, J. Idier, C. Pallier, and H. Benali. Unsupervised robust nonparametric estimation of the hemodynamic response function for any fMRI experiment. *IEEE transactions on Medical Imaging*, 22:35–51, 2003. (Cited on pages 1, 31, 32, 81, 82, 91, and 131.)
- P. Ciuciu, T. Vincent, L. Risser, and S. Donnet. A joint detection-estimation framework for analysing within-subject fMRI data. *Journal de la Société Française de Statistique*, 151: 58–89, 2010. (Cited on page 35.)
- P. L. Combettes and H. H. Bauschke. Convex Analysis and Monotone Operator Theory in Hilbert Spaces. *Springer Science & Business Media*, 2011. (Cited on page 58.)
- P. L. Combettes and J. C. Pesquet. Proximal splitting methods in signal processing. *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, 2009. (Cited on page 95.)

- L. Condat. A Direct Algorithm for 1D Total Variation Denoising. *IEEE Signal Processing Letters*, 20(11):1054–1057, 2013a. (Cited on pages 42, 56, 58, and 64.)
- L. Condat. A Primal–Dual Splitting Method for Convex Optimization Involving Lipschitzian, Proximable and Linear Composite Terms. *Journal of Optimization Theory and Applications*, 158(2):460–479, 2013b. (Cited on pages 44, 48, 56, and 58.)
- L. Condat. Fast projection onto the simplex and the l_1 ball. *Mathematical Programming Series A*, 158:575–585, 2016. (Cited on pages 84 and 97.)
- I. Costantini, P. Filipiak, K. Maksymenko, S. Deslauriers-Gauthier, and R. Deriche. fmri deconvolution via temporal regularization using a lasso model and the lars algorithm. 2018. (Cited on page 36.)
- A. M. Dale. Optimal experimental design for event-related fMRI. *Human Brain Mapping*, 2: 109–114, 1999. (Cited on pages 31, 32, and 91.)
- J. Darbon and M. Sigelle. Image Restoration with Discrete Constrained Total Variation Part I: Fast and Exact Optimization. *Journal of Mathematical Imaging and Vision*, 26(3): 261–276, 2006. (Cited on page 41.)
- I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57(11):1413–1457, 2004. (Cited on page 46.)
- P. L. Davies and A. Kovac. Local Extremes, Runs, Strings and Multiresolution. *The Annals of Statistics*, 29(1):1–65, 2001. (Cited on page 56.)
- A. de Pierrefeu, T. Löfstedt, F. Hadj-Selem, M. Dubois, R. Jardri, T. Fovet, P. Ciuciu, V. Frouin, and E. Duchesnay. Structured sparse principal components analysis with the TV-elastic net penalty. *IEEE transactions on medical imaging*, 37(2):396–407, 2017. (Cited on page 126.)
- M. R. Del Río and J. De Felipe. A study of smi 32-stained pyramidal cells, parvalbumin-immunoreactive chandelier cells, and presumptive thalamocortical axons in the human temporal neocortex. *The Journal of comparative neurology*, 342(3):389–408, 1994. (Cited on page 10.)
- C. A. Deledalle, S. Vaiter, J. Fadili, and G. Peyré. Stein Unbiased Gradient estimator of the Risk (SUGAR) for multiple parameter selection. *SIAM Journal on Imaging Sciences*, 7(4): 2448–2487, 2014. (Cited on page 64.)
- R. S. Desikan, F. Segonne, B. Fischl, B. T. Quinn, B. C. Dickerson, D. Blacker, R. L. Buckner, A. M. Dale, R. P. Maguire, B. T. Hyman, M. S. Albert, and R. J. Killiany. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage*, 31:968–980, 2006. (Cited on pages 92, 107, 124, and 134.)
- M. D’Esposito, L. Y. Deouell, and A. Gazzaley. Alterations in the BOLD fMRI signal with ageing and disease: a challenge for neuroimaging. *Nature Reviews Neuroscience*, 4(11): 863–872, 2003. (Cited on page 127.)

- C.-T. Do, Z.-M. Manjaly, J. Heinzle, D. Schöbi, L. Kasper, K. P. Pruessmann, K. Enno Stephan, and S. Frässle. Hemodynamic modeling of aspirin effects on BOLD responses at 7t. *medRxiv*, 2020. (Cited on page 1.)
- T. Dupré La Tour, T. Moreau, M. Jas, and A. Gramfort. Multivariate convolutional sparse coding for electromagnetic brain signals. In *Advances in Neural Information Processing System (NeurIPS)*, pages 3292–3302, 2018. (Cited on pages 81, 87, and 90.)
- Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least angle regression. *The Annals of Statistics*, 32, 2004. (Cited on page 47.)
- M. Elad, P. Milanfar, and R. Rubinstein. Analysis versus synthesis in signal priors. *Inverse Problems*, 23(3):947–968, 2007. (Cited on page 58.)
- D. A. Engemann, O. Kozynets, D. Sabbagh, G. Lemaître, G. Varoquaux, F. Liem, and A. Gramfort. Combining magnetoencephalography with magnetic resonance imaging enhances learning of surrogate-biomarkers. *Elife*, 9, 2020. (Cited on page 110.)
- Public Health England. Briefing document: First incidence of stroke estimates for england 2007 to 2016. *Public Health England publication*, 2018. (Cited on page 109.)
- O. Esteban, C. J. Markiewicz, R. W. Blair, C. A. Moodie, A. I. Isik, A. Erramuzpe, J. D. Kent, M. Goncalves, E. DuPre, M. Snyder, H. Oya, S. S. Ghosh, J. Wright, J. Durnez, R. A. Poldrack, and K. J. Gorgolewski. fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nat Methods*, 16:111–116, 2019. (Cited on page 116.)
- Y. Farouj, F. I. Karahanoğlu, and D. Van De Ville. Bold Signal Deconvolution Under Uncertain HÆModynamics: A Semi-Blind Approach. In *IEEE 16th International Symposium on Biomedical Imaging (ISBI)*, pages 1792–1796, 2019. (Cited on pages 2, 37, 38, 72, and 132.)
- D. A. Feinberg and K. Setsompop. Ultra-fast MRI of the human brain with simultaneous multi-slice imaging. *Journal of magnetic resonance*, 229:90–100, 2013. (Cited on page 126.)
- D. A. Feinberg, S. Moeller, S. M. Smith, E. Auerbach, S. Ramanna, M. F. Glasser, K. L. Miller, K. Ugurbil, and E. Yacoub. Multiplexed echo planar imaging for sub-second whole brain fmri and fast diffusion imaging. *PLOS ONE*, 5:1–11, 2010. (Cited on page 20.)
- D. L. Felten and A. N. Shetty. Atlas de neurosciences humaines de netter. *Elsevier Masson*, 2011. (Cited on page 12.)
- A. Frau-Pascual, F. Forbes, and P. Ciuciu. Physiological models comparison for the analysis of asl fmri data. In *International Symposium on Biomedical Imaging*, volume 12, pages 1348–1351, 2015. (Cited on page 81.)
- J. Friedman, T. Hastie, H. Höfling, and R. Tibshirani. Pathwise coordinate optimization. *The Annals of Applied Statistics*, 1(2):302–332, 2007. (Cited on page 78.)
- K. J. Friston, P. Fletcher, O. Josephs, A. Holmes, M. D. Rugg, and R. Turner. Event-related fMRI: characterizing differential responses. *Neuroimage*, 7:30–40, 1998a. (Cited on pages 1, 2, 24, 31, 32, 41, 77, 91, and 131.)

- K. J. Friston, O. Josephs, G. Rees, and R. Turner. Nonlinear event-related responses in fMRI. *Magnetic Resonance in Medicine*, 39:41–52, 1998b. (Cited on page 96.)
- K. J. Friston, A. P. Holmes, C. J. Price, C. Buchel, and K. J. Worsley. Multisubject fMRI studies and conjunction analyses. *NeuroImage*, 10:385–396, 1999a. (Cited on page 81.)
- K. J. Friston, E. Zarahn, R. N. A. Josephs, O. Henson, and Dale A. M. Stochastic designs in event-related fMRI. *NeuroImage*, 10:607–619, 1999b. (Cited on pages 31 and 32.)
- R. Giryes, Y. C. Eldar, A. M. Bronstein, and G. Sapiro. Tradeoffs between Convergence Speed and Reconstruction Accuracy in Inverse Problems. *IEEE Transaction on Signal Processing*, 66(7):1676–1690, 2018. (Cited on page 56.)
- D. R. Gitelman, W. D. Penny, J. Ashburner, and K. J. Friston. Modeling regional and psychophysiologic interactions in fMRI: the importance of hemodynamic deconvolution. *Neuroimage*, 19(1):200–207, 2003. (Cited on pages 1 and 132.)
- M. Glasser, S. Sotiropoulos, J. Wilson, T. Coalson, B. Fischl, J. Andersson, J. Xu, S. Jbabdi, M. Webster, J. Polimeni, V. DC, and M. Jenkinson. The minimal preprocessing pipelines for the human connectome project. *NeuroImage*, 80:105, 2013. (Cited on page 107.)
- G. H. Glover. Deconvolution of impulse response in event-related bold fMRI. *NeuroImage*, 9(4):416–429, 1999. (Cited on pages 1, 31, 32, 34, 91, and 132.)
- C. Goutte, F. A. Nielsen, and L. K. Hansen. Modeling the haemodynamic response in fMRI using smooth FIR filters. *IEEE transactions on Medical Imaging*, 19:1188–1201, 2000. (Cited on pages 31, 32, and 124.)
- C. L. Grady and D. D. Garrett. Understanding variability in the bold signal and why it matters for aging. *Brain Imaging Behaviour*, 8:274–283, 2014. (Cited on page 113.)
- K. Gregor and Y. Le Cun. Learning Fast Approximations of Sparse Coding. In *International Conference on Machine Learning (ICML)*, pages 399–406, 2010. (Cited on pages 2, 56, 62, and 66.)
- M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic Resonance in Medicine*, 47:1202–1210, 2002. (Cited on page 18.)
- R. Grosse, R. Raina, H. Kwong, and Y. N. Andrew. Shift-invariant sparse coding for audio classification. In *proceedings of Uncertainty in Artificial Intelligence*, volume 23, pages 149–158, 2007. (Cited on page 90.)
- D. A. Handwerker, J. M. Ollinger, and M. D’Esposito. Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *Neuroimage*, 21(4):1639–1651, 2004. (Cited on pages 92, 109, 124, and 134.)
- Zaïd Harchaoui and Céline Lévy-Leduc. Catching change-points with lasso. volume 22, 2007. (Cited on page 47.)
- M. Havlicek and K. Uludağ. A dynamical model of the laminar BOLD response. *NeuroImage*, 204:116–209, 2020. (Cited on page 124.)

- G. Henriksen and F. Willoch. Imaging of opioid receptors in the central nervous system. *Brain*, 131:1171–1196, 2008. (Cited on page 117.)
- L. Hernandez-Garcia and M. O. Ulfarsson. Neuronal event detection in fMRI time series using iterative deconvolution techniques. *Magnetic resonance imaging*, 29:353–364, 2011. (Cited on pages 1 and 132.)
- J. Hesamoddin, S. Holdsworth, T. Christen, H. Wu, K. Zhu, A. B. Kerr, M. J. Middione, R. F. Dougherty, M. Moseley, and G. Zaharchuk. Advantages of short repetition time resting-state functional MRI enabled by simultaneous multi-slice imaging. *Journal of Neuroscience Methods*, 311:122–132, 2019. (Cited on page 126.)
- J. M. Hiller and L. Q. Fan. Laminar distribution of the multiple opioid receptors in the human cerebral cortex. *Neurochem Res.*, 21:1333–1345, 1996. (Cited on page 117.)
- R. Jenatton, A. Gramfort, V. Michel, G. Obozinski, E. Eger, F. Bach, and B. Thirion. Multiscale mining of fMRI data with hierarchical structured sparsity. *SIAM Journal on Imaging Sciences*, 5:835–856, 2012. (Cited on page 126.)
- M. Jenkinson, P. Bannister, M. Brady, and S. Smith. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17: 825–41, 2002. (Cited on page 98.)
- M. Jenkinson, C. F. Beckmann, T. E. J. Behrens, M. W. Woolrich, and S. M. Smith. Fsl. *NeuroImage*, 62:782–790, 2012. (Cited on page 21.)
- E. Jones, T. Oliphant, P. Peterson, et al. Scipy: Open source scientific tools for python, 2001–. (Cited on page 78.)
- T. Kantonen, T. Karjalainen, J. Isojärvi, P. Nuutila, J. Tuisku, J. Rinne, J. Hietala, V. Kaasinen, K. Kalliokoski, H. Scheinin, J. Hirvonen, A. Vehtari, and L. Nummenmaa. Interindividual variability and lateralization of μ -opioid receptors in the human brain. 2019. (Cited on page 117.)
- F. I. Karahanoglu, C. Caballero-gaudes, F. Lazeyras, and D. Van De Ville. Total activation: fMRI deconvolution through spatio-temporal regularization. *NeuroImage*, 73:121–134, 2013. (Cited on pages 1, 2, 34, 35, 36, 37, 38, 41, 72, 76, 82, 83, 86, 93, 94, and 132.)
- I. Khalidov and M. Unser. From differential equations to the construction of new wavelet-like bases. *Signal Processing, IEEE Transactions on*, 54:1256–1267, 2006. (Cited on pages 31, 32, and 35.)
- I. Khalidov, J. Fadili, F. Lazeyras, D. Van De Ville, and M. Unser. Activelets: Wavelets for sparse representation of hemodynamic responses. *Signal processing*, 91(12):2810–2821, 2011. (Cited on pages 1, 31, 32, 35, 91, and 132.)
- D. Kim and J. A. Fessler. Adaptive restart of the optimized gradient method for convex optimization. *Journal of Optimization Theory and Applications*, 178:240–263, 2018. (Cited on page 43.)

- C. Lalanne, M. Rateaux, L. Oudre, M. P. Robert, and T. Moreau. Extraction of Nystagmus Patterns from Eye-Tracker Data with Convolutional Sparse Coding. In *IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 928–931, 2020. (Cited on page 41.)
- C. Lazarus, P. Weiss, N. Chauffert, F. Mauconduit, L. El Gueddari, C. Destrieux, I. Zemmoura, A. Vignaud, and P. Ciuciu. SPARKLING: variable-density k-space filling curves for accelerated t2*-weighted MRI. *Magnetic Resonance in Medicine*, 81:3643–3661, 2019. (Cited on page 18.)
- B. Lecouat, J. Ponce, and J. Mairal. Designing and Learning Trainable Priors with Non-Cooperative Games. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. (Cited on page 57.)
- M. Lee, V. Wanigasekera, and I. Tracey. Imaging opioid analgesia in the human brain and its potential relevance for understanding opioid use in chronic pain. *Neuropharmacology*, 84, 2013. (Cited on page 117.)
- L.D. Lewis, K. Setsompop, B.R. Rosen, , and J.R. Polimeni. Aging-associated changes in cerebral vasculature and blood flow as determined by quantitative optical coherence tomography angiography. *PNAS*, 113:E6679–E6685, 2016. (Cited on page 21.)
- Y. Li, W. J. Choi, W. Wei, S. Song, Q. Zhang, J. Liu, , and R. K. Wang. Aging-associated changes in cerebral vasculature and blood flow as determined by quantitative optical coherence tomography angiography. *Neurobiol Aging*, 70:148–159, 2018. (Cited on page 110.)
- J. Liang and C.-B. Schönlieb. Improving FISTA: Faster, smarter and greedier. *preprint*, 2018. (Cited on page 43.)
- M. A. Lindquist and T. D. Wager. Validity and power in hemodynamic response modeling: a comparison study and a new approach. *Human brain mapping*, 28:764–784, 2007. (Cited on pages 1, 31, 33, 77, 91, 124, and 131.)
- M. A. Lindquist, J. Meng Loh, L. Y. Atlas, and T. D. Wager. Modeling the hemodynamic response function in fMRI: efficiency, bias and mis-modeling. *NeuroImage*, 45:S187–S198, 2009. (Cited on pages 31, 33, 82, and 91.)
- N. K. Logothetis, J. Pauls, M. Augath, T. Trinath, and A. Oeltermann. Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 12:150–157, 2001. (Cited on page 12.)
- M. Lustig, D. L. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, 58:1182–1195, 2007. (Cited on page 18.)
- P. Machart, S. Anthoine, and L. Baldassarre. Optimal Computational Trade-Off of Inexact Proximal Methods. *preprint ArXiv*, 2012. (Cited on pages 57 and 67.)
- J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online dictionary learning for sparse coding. In *International Conference on Machine Learning*, volume 26, pages 689–696, 2009. (Cited on page 95.)

- S. Makni, P. Ciuciu, J. Idier, and J.-B. Poline. Joint detection-estimation of brain activity in functional MRI: a multichannel deconvolution solution. *IEEE Transactions on Signal Processing*, 53:3488–3502, 2005. (Cited on page 76.)
- S. Makni, J. Idier, T. Vincent, B. Thirion, G. Dehaene-Lambertz, and P. Ciuciu. A fully bayesian approach to the parcel-based detection-estimation of brain activity in fMRI. *Neuroimage*, 41:941–969, 2008. (Cited on page 90.)
- G. Marrelec, H. Benali, P. Ciuciu, M. Péligrini-Issac, and J.-B. Poline. Robust Bayesian estimation of the hemodynamic response function in event-related BOLD MRI using basic physiological information. *hbm*, 19:1–17, 2003. (Cited on page 124.)
- V. Menon. Large-scale functional brain organization. *Brain mapping: An encyclopedic reference*, 2:449–459, 2015. (Cited on page 103.)
- M. Milham, D. Fair, M. Mennes, and S. Mostofsky. The adhd-200 consortium: a model to advance the translational potential of neuroimaging in clinical neuroscience. *Frontiers in Systems Neuroscience*, 6:62, 2012. (Cited on page 126.)
- S. N. Min, S. J. Park, D. J. Kim, M. Subramaniam, and K. S. Lee. Development of an algorithm for stroke prediction: A national health insurance database study in korea. *European Neurology*, 79:214–220, 2018. (Cited on page 109.)
- S. Moeller, E. Yacoub, C. A. Olman, E. Auerbach, J. Strupp, N. Harel, and K. Uğurbil. Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. *Magnetic resonance in Medicine*, 63:1144–1153, 2010. (Cited on pages 18 and 116.)
- V. Monga, Y. Li, and Y. C. Eldar. Algorithm Unrolling: Interpretable, Efficient Deep Learning for Signal and Image Processing. *preprint ArXiv*, 2019. (Cited on page 56.)
- T. Moreau and J. Bruna. Understanding Neural Sparse Coding with Matrix Factorization. In *International Conference on Learning Representation (ICLR)*, 2017. (Cited on page 56.)
- T. Moreau, L. Oudre, and N. Vayatis. DICOD: Distributed Convolutional Sparse Coding. In *proceedings of International Conference on Machine Learning*, 2018. (Cited on pages 78 and 82.)
- A. A. Neath and J. E. Cavanaugh. The Bayesian information criterion: background, derivation, and applications. *Wiley Interdisciplinary Reviews: Computational Statistics*, 4:199–203, 2012. (Cited on page 125.)
- Y. Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140:125–161, 2013. (Cited on page 43.)
- B. O’Donoghue and E. Candes. Adaptive restart for accelerated gradient schemes. *Foundations of Computational Mathematics*, 15:715–732, 2015. (Cited on page 95.)
- A. Ogaki, Y. Ikegaya, and R. Koyama. Vascular abnormalities and the role of vascular endothelial growth factor in the epileptic brain. *Frontiers in Pharmacology*, 11, 2020. (Cited on page 127.)

- S. Ogawa, T. M. Lee, A. R. Kay, and D. W. Tank. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences of the United States of America*, 87:9868–9872, 1990. (Cited on page 19.)
- S. Ogawa, D. W. Tank, R. Menon, J. M. Ellermann, S. G. Kim, H. Merkle, and K. Ugurbil. Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Sciences*, 89: 5951–5955, 1992. (Cited on pages 1, 19, 72, 76, and 131.)
- B. A. Olshausen and D. J. Field. Sparse coding with an incomplete basis set: a strategy employed by V1. *Vision Research*, 37:3311–3325, 1997. (Cited on pages 76 and 82.)
- A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems (NeurIPS)*, page 12, 2019. (Cited on pages 69 and 70.)
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. (Cited on pages 97, 110, and 112.)
- F. Pedregosa, M. Eickenberg, P. Ciuciu, B. Thirion, and A. Gramfort. Data-driven HRF estimation for encoding and decoding models. *NeuroImage*, pages 209–220, 2015. (Cited on pages 1, 31, 33, 77, 79, 80, 81, 124, and 131.)
- W. Penny, S. Kiebel, and K. Friston. Variational bayesian inference for fMRI time series. *NeuroImage*, 19:727–741, 2003. (Cited on page 90.)
- K. P. Pruessmann, P. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, 42:952–962, 1999. (Cited on page 18.)
- M. Raemaekers, W. Schellekens, N. Petridou, and N. F. Ramsey. Knowing left from right: asymmetric functional connectivity during resting state. *Brain Structure and Function*, 223:1909–1922, 2018. (Cited on page 109.)
- H. Raguét, J. Fadili, and G. Peyré. Generalized forward-backward splitting. *SIAM Journal on Imaging Sciences*, 6:1199–1226, 2013. (Cited on page 35.)
- D. Rangaprakash, M. N. Dretsch, W. Yan, J. S. Katz, T. S. Denney Jr, and G. Deshpande. Hemodynamic response function parameters obtained from resting-state functional MRI data in soldiers with trauma. *Data in brief*, 14:558–562, 2017. (Cited on page 123.)
- D. Rangaprakash, G.-R. Wu, D. Marinazzo, X. Hu, and G. Deshpande. Hemodynamic response function (HRF) variability confounds resting-state fMRI functional connectivity. *Magnetic resonance in medicine*, 80:1697–1713, 2018. (Cited on page 123.)

- R. T. Rockafellar. Monotone Operators and the Proximal Point Algorithm. *SIAM Journal on Control and Optimization*, 14(5):877–898, 1976. (Cited on pages 42, 56, and 58.)
- P. Rodríguez. Total Variation Regularization Algorithms for Images Corrupted with Different Noise Models: A Review. *Journal of Electrical and Computer Engineering*, pages 1–18, 2013. (Cited on page 41.)
- L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60:259–268, 1992. (Cited on pages 41 and 57.)
- M. Schmidt, N. Le Roux, and F. R. Bach. Convergence Rates of Inexact Proximal-Gradient Methods for Convex Optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1458–1466, 2011. (Cited on pages 57 and 71.)
- M. A. Shafto, L. K. Tyler, M. Dixon, J. R. Taylor, J. B. Rowe, R. Cusack, A. J. Calder, W. D. Marslen-Wilson, J. Duncan, T. Dalgleish, et al. The cambridge centre for ageing and neuroscience (Cam-CAN) study protocol: a cross-sectional, lifespan, multidisciplinary examination of healthy cognitive ageing. *BMC neurology*, 14:204, 2014. (Cited on page 124.)
- J. W. Silverstein. On the eigenvectors of large dimensional sample covariance matrices. *Journal of Multivariate Analysis*, 30:1–16, 1989. (Cited on page 62.)
- K. A. Smitha, K. M. Arun, P. G. Rajesh, S. E. Joel, R. Venkatesan, B. Thomas, and C. Kesavadas. Multiband fmri as a plausible, time-saving technique for resting-state data acquisition: Study on functional connectivity mapping using graph theoretical measures. *Magn Reson Imaging*, 53:1–6, 2018. (Cited on page 21.)
- P. Sprechmann, A. M. Bronstein, and G. Sapiro. Learning Efficient Structured Sparse Models. In *International Conference on Machine Learning (ICML)*, pages 615–622, 2012. (Cited on page 56.)
- P. Sprechmann, R. Litman, and T. B. Yakar. Efficient Supervised Sparse Analysis and Synthesis Operators. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 908–916, 2013. (Cited on page 56.)
- K. R. Sreenivasan, M. Havlicek, and G. Deshpande. Nonparametric hemodynamic deconvolution of fmri using homomorphic filtering. In *IEEE Trans Med Imaging.*, volume 34, pages 1155–1163, 2015. (Cited on page 76.)
- S. F. Storti, E. Formaggio, A. Bertoldo, P. Manganotti, A. Fiaschi, and G. M. Toffolo. Modelling hemodynamic response function in epilepsy. *Clinical Neurophysiology*, 124: 2108–2118, 2013. (Cited on page 127.)
- C. Sudlow, J. Gallacher, N. Allen, V. Beral, P. Burton, J. Danesh, P. Downey, P. Elliott, J. Green, M. Landray, B. Liu, P. Matthews, G. Ong, J. Pell, A. Silman, A. Young, T. Sprosen, T. Peakman, and R. Collins. UK Biobank: An Open Access Resource for Identifying the Causes of a Wide Range of Complex Diseases of Middle and Old Age. *PLOS Medicine*, 12(3), 2015. (Cited on pages 72 and 135.)
- J. Sulam, A. Aberdam, A. Beck, and M. Elad. On Multi-Layer Basis Pursuit, Efficient Algorithms and Convolutional Neural Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2019. (Cited on page 56.)

- Enzo Tagliazucchi, Pablo Balenzuela, Daniel Fraiman, Pedro Montoya, and Dante Chialvo. Spontaneous BOLD event triggered averages for estimating functional connectivity at resting state. *Neuroscience letters*, 488:158–63, 2010. (Cited on page 38.)
- A. Tian, X. Jia, K. Yuan, T. Pan, and S. B. Jiang. Low-dose ct reconstruction via edge-preserving total variation regularization. *Physics in Medicine & Biology*, 56:5949, 2011. (Cited on page 41.)
- R. Tibshirani. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society: Series B (statistical methodology)*, 58:267–288, 1996. (Cited on pages 56, 58, and 100.)
- R. J. Tibshirani and J. Taylor. The solution path of the generalized lasso. *The Annals of Statistics*, 39:1335–1371, 2011. (Cited on page 57.)
- C. Triantafyllou, R. Hoge, G. Krueger, C. J. Wiggins, A. Potthast, G. C. Wiggins, and L. Wald. Comparison of physiological noise at 1.5 t, 3 t and 7 t and optimization of fMRI acquisition parameters. *NeuroImage*, 26:243–50, 2005. (Cited on page 21.)
- K. A. Tsvetanov, R. N. A. Henson, P. S. Jones, H. J. Mutsaerts, D. Fuhrmann, L. K. Tyler, , and J. B. Rowe. The effects of age on resting-state bold signal variability is explained by cardiovascular and neurovascular factors. *preprint bioRxiv*, 2019. (Cited on page 126.)
- V. A. Vakorin, R. Borowsky, and G. E. Sarty. Characterizing the functional MRI response using tikhonov regularization. *Statistics in medicine*, 26(21):3830–3844, 2007. (Cited on pages 31, 32, and 91.)
- D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. J. Behrens, E. Yacoub, and K. Ugurbil. The wu-minn Human Connectome Project: An overview. *NeuroImage*, 80:62–79, 2013. (Cited on pages 21, 79, 86, and 107.)
- G. Varoquaux and R. C. Craddock. Learning and comparing functional connectomes across subjects. *NeuroImage*, 80:405–415, 2013. (Cited on pages 92 and 134.)
- G. Varoquaux, S. Sadaghiani, P. Pinel, A. Kleinschmidt, J. B. Poline, and B. Thirion. A group model for stable multi-subject ICA on fMRI datasets. *NeuroImage*, 51:288–299, 2010. (Cited on pages 82 and 126.)
- G. Varoquaux, A. Gramfort, F. Pedregosa, V. Michel, and B. Thirion. Multi-subject dictionary learning to segment an atlas of brain spontaneous activity. In *Biennial International Conference on information processing in medical imaging*, pages 562–573, 2011. (Cited on page 126.)
- T. Vincent, T. Rissern, and P. Ciuciu. Spatially adaptive mixture modeling for analysis of fMRI time series. *IEEE transactions on Medical Imaging*, 29:59–74, 2010. (Cited on pages 1, 124, and 131.)
- R. Viviani, G Grön, and M. Spitzer. Functional principal component analysis of fMRI data. In *Human Brain Mapping*, volume 24, pages 109–29, 2005. (Cited on page 82.)
- S. Watanabe. A widely applicable Bayesian information criterion. *Journal of Machine Learning Research*, 14:867–897, 2013. (Cited on page 125.)

- K. L. West, M. D. Zuppichini, M. P. Turner, D. K. Sivakolundu, Y. Zhao, D. Abdelkarim, J. S. Spence, and B. Rypma. BOLD hemodynamic response function changes significantly with healthy aging. *Neuroimage*, 188:198–207, 2019. (Cited on pages 1, 110, and 113.)
- M. W. Woolrich, B. D. Ripley, M. Brady, and S. M. Smith. Temporal autocorrelation in univariate linear modeling of fMRI data. *Neuroimage*, 14:1370–1386, 2001. (Cited on pages 30 and 90.)
- M. W. Woolrich, T. E. J. Behrens, and S. M. Smith. Constrained linear basis sets for hrf modelling using variational bayes. *NeuroImage*, 21:1748–1761, 2004. (Cited on page 91.)
- G. R. Wu, W. Liao, S. Stramaglia, J. R. Ding, H. Chen, and D. Marinazzo. A blind deconvolution approach to recover effective connectivity brain networks from resting state fMRI data. In *Medical Image Analysis*, volume 17, pages 365–374, 2013. (Cited on pages 1, 37, 114, and 132.)
- B. Xin, Y. Wang, W. Gao, and D. Wipf. Maximal Sparsity with Deep Networks? In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 4340–4348, 2016. (Cited on page 56.)
- W. Yan, D. Rangaprakash, and G. Deshpande. Aberrant hemodynamic responses in autism: Implications for resting state fMRI functional connectivity studies. *NeuroImage: Clinical*, 19:320–330, 2018. (Cited on page 123.)
- B. T. Yeo, F. M. Krienen, J. Sepulcre, M. R. Sabuncu, D. Lashkari, M. Hollinshead, J. L. Roffman, J. W. Smoller, L. Zollei, J. R. Polimeni, B. Fischl, H. Liu, and R. L. Buckner. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Neurophysiol*, 106:1125–1165, 2011. (Cited on pages 79 and 86.)
- C. M. Zhang, Y. Jiang, and T. Yu. A comparative study of one-level and two-level semiparametric estimation of hemodynamic response function for fMRI data. *Statistics in medicine*, 26:3845–3861, 2007. (Cited on page 91.)
- T Zhang, F. Li, L. Beckes, and J. Coan. A semi-parametric model of the hemodynamic response for multi-subject fmri data. *NeuroImage*, 75, 2013. (Cited on pages 31 and 32.)
- M. Zijlmans, G. Huiskamp, M. Hersevoort, J.-H. Seppenwoolde, A. C. van Huffelen, and F. S. S. Leijten. EEG–fMRI in the preoperative work-up for epilepsy surgery. *Brain*, 130:2343–2353, 2007. (Cited on pages 123 and 127.)
- H. Zou, , and T. Hastie. Regularization and variable selection via the elastic net. *J. R. Statist. Soc. B*, 67:301–320, 2004. (Cited on page 40.)

Titre: Estimation efficace de la réponse hémodynamique cerveau-entier par déconvolution semi-aveugle, via la régularisation par variation totale, à partir de données d'IRM fonctionnelle

Mots clés: FRH, IRMf, modèles multivariés, estimation, apprentissage de dictionnaire, pathologies cérébrales

Résumé: L'Imagerie par Résonance Magnétique fonctionnelle (IRMf) enregistre de manière non-invasive l'activité cérébrale en mesurant l'évolution du niveau d'oxygénation du sang (contrast BOLD). Cette dernière mesure de manière indirecte l'activité cérébrale au travers du couplage neurovasculaire. Ce couplage est caractérisé par la Fonction de Réponse Hémodynamique (FRH). L'estimation de cette réponse est d'un intérêt particulier: un changement de cette réponse hémodynamique peut être lié au mécanisme pharmacologique d'une drogue, l'effet de l'âge sur l'organisme ou les conséquences d'une pathologie cérébrale. De plus, son estimation permet aussi de relier le signal BOLD observé à l'activité cérébrale, i.e. neurale sous-jacente ce qui permet ainsi de mieux comprendre les mécanismes cognitifs en jeu. Nous proposons une nouvelle approche pour distinguer le signal d'activité neurale du couplage neurovasculaire (FRH) aussi bien pour des données d'IRMf de repos (IRMf-r) que d'IRMf de tâche (IRMf-t). L'activité neurale est modélisée par un signal constant par morceau ou dit "par bloc". Cette hypothèse est imposée grâce à la régularisation par variation totale (TV) en 1-dimension (1D). La résolution d'un tel problème de déconvolution par régularisation TV peut être coûteux en temps de calcul. Ainsi nous proposons de comparer théoriquement les différentes formulations associées au problème et démontrons que l'algorithme issu de la formulation à l'analyse est plus efficace que celui associé à la formulation à la synthèse. Dans ce travail, nous proposons aussi de dérouler de manière différentiable les algorithmes de descente de gradient proximal associés au problème d'optimisation avec régulari-

sation TV, en calculant directement le gradient de l'opérateur proximal TV via une formule analytique ou en le calculant via un algorithme interne déroulé. Une comparaison sur données réelles d'IRMf donne des résultats prometteurs concernant les algorithmes déroulés en comparaison aux algorithmes itératifs. Afin d'étudier la réponse hémodynamique sur l'ensemble du cerveau, nous proposons, en plus de l'hypothèse d'un signal bloc concernant l'activité neurale, une modélisation originale et synthétique de la FRH. La réponse hémodynamique est définie par dilatation de la FRH canonique. Cette paramétrisation permet de résumer la FRH avec un seul paramètre scalaire par région appartenant à une parcellisation prédéfinie. De plus, nous introduisons des composantes temporelles et leur carte spatiale correspondante pour encoder l'activité cérébrale et localiser la contribution de chaque composante dans le cerveau. Afin d'illustrer la pertinence de notre approche en terme d'analyse de cohorte, nous démontrons statistiquement qu'une pathologie telle que l'accident vasculaire cérébrale et la dégradation vasculaire due à l'âge induisent un retard hémodynamique au sein de certaines régions cérébrales et que ces caractéristiques peuvent être utilisées pour prédire l'état de santé d'un individu au moyen de méthodes d'apprentissage automatique (tâche de classification). Ainsi, grâce à cette thèse, nous proposons une avancée significative concernant l'estimation de la FRH sur l'ensemble du cerveau qui s'applique à la fois aux données d'IRMf de repos et de tâche de par la déconvolution semi-aveugle de l'activité neurale. Nous ouvrons ainsi de nouvelles perspectives pour l'usage de l'IRMf en recherche clinique au travers de cette caractérisation hémodynamique fiable.

Title: Efficient whole brain estimation of the haemodynamic response function for TV-regularized semi-blind deconvolution of neural activity in fMRI

Keywords: HRF, fMRI, multivariate modeling, estimation, dictionary learning, cerebral pathology

Abstract: Functional magnetic resonance imaging (fMRI) non-invasively records brain activity by dynamically measuring the blood oxygenation level-dependent (BOLD) contrast. The latter indirectly measures neural activity through the neurovascular coupling. This coupling is characterized by the so-called haemodynamic response function (HRF). The estimation of the latter is of a primary interest: a change in the haemodynamic response could be linked to the pharmacological mechanism of a drug, the effect of normal aging or the consequence of a neuropathological process. Moreover, its estimation also links the observed BOLD signal to the underlying neural activity, which can in turn be used to better understand the cognitive processes involved in task performance. We propose a novel approach to disentangle the neural activation signal from the neurovascular coupling for both task-related and resting-state fMRI (tfMRI and rs-fMRI, respectively). We model the neural activation as a piecewise constant block signal, this hypothesis is enforced using 1-dimensional (1D) Total Variation (TV) regularization which promotes sparsity of the first-order derivative of neural activity signals and thus leads to block-shaped signals. The resolution of such deconvolution problem with TV regularization can be computationally expensive. We propose to theoretically compare the performances of the analysis vs synthesis formulations and show that the analysis one can be much more efficient than its synthesis counterpart. In this work, we also propose a mean to perform differentiable algorithm unrolling for proximal gradient descent (PGD) with TV-

regularized optimization problems by either directly computing the TV proximal operator gradient with an analytic formula or computing it with a nested unrolled algorithm. Comparison on real fMRI data shows promising results for learned optimization algorithm based on this unrolling mechanism as compared to iterative algorithms. In order to study the haemodynamic response over the whole brain, we stick to the same block structure hypothesis for the neural activity signal and introduce a simplistic but novel modeling for the HRF. The latter is now modeled as a dilated version of the canonical HRF. This parameterization allows us to capture regionwise variations through a scalar dilation parameter. We introduce temporal components and their corresponding spatial map to encode the neural activity and localize their respective contribution in the brain. Moreover, we introduce the neurovascular coupling over the whole brain by estimating a temporally dilated HRF for each region from a predefined brain parcellation. To illustrate the significance of this approach at the population level, we statistically demonstrate that a pathology like stroke or a condition like normal aging induce longer haemodynamic delays in some brain areas and that this haemodynamic feature may be predictive of the individual status in a machine learning (classification) task. In this thesis we thus made a significant step forward towards efficient whole brain estimation of the HRF in both the rs- and tfMRI contexts and opened a new research avenue in clinical neuroscience through this reliable haemodynamic characterization.