



HAL
open science

Multivariate semi-blind deconvolution of fMRI time series

Hamza Cherkaoui, Thomas Moreau, Abderrahim Halimi, Claire Leroy,
Philippe Ciuciu

► **To cite this version:**

Hamza Cherkaoui, Thomas Moreau, Abderrahim Halimi, Claire Leroy, Philippe Ciuciu. Multivariate semi-blind deconvolution of fMRI time series. 2020. hal-03005584v1

HAL Id: hal-03005584

<https://hal.science/hal-03005584v1>

Preprint submitted on 14 Nov 2020 (v1), last revised 2 Sep 2021 (v2)

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.

Multivariate semi-blind deconvolution of fMRI time series

Hamza Cherkaoui^{b,a,c}, Thomas Moreau^c, Abderrahim Halimi^d, Claire Leroy^b, and
Philippe Ciuciu^{a,c,e}

^aCEA/NeuroSpin, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France.

^bUniversité Paris-Saclay, CEA, CNRS, Inserm, BioMaps, Orsay, 91401, France.

^cParietal Team, Université Paris-Saclay, CEA, Inria, Gif-sur-Yvette, 91190, France.

^dSchool of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh UK.

^eCorresponding author.

Abstract

Whole brain estimation of the haemodynamic response function (HRF) in functional magnetic resonance imaging (fMRI) is critical to get insight on the global status of the neurovascular coupling of an individual in healthy or pathological condition. Most of existing approaches in the literature works on task-fMRI data and relies on the experimental paradigm as a surrogate of neural activity, hence remaining inoperative on resting-stage fMRI (rs-fMRI) data. To cope with this issue, recent works have performed either a two-step analysis to detect large neural events and then characterize the HRF shape or a joint estimation of both the neural and haemodynamic components in an univariate fashion. In this work, we express the neural activity signals as a combination of piece-wise constant temporal atoms associated with sparse spatial maps and introduce an haemodynamic parcelation of the brain featuring a temporally dilated version of a given HRF model in each parcel with unknown dilation parameters. We formulate the joint estimation of the HRF shapes and spatio-temporal neural representations as a multivariate *semi-blind deconvolution* problem in a paradigm-free setting and introduce constraints inspired from the dictionary learning literature to ease its identifiability. An efficient alternating minimization algorithm is proposed and validated on both synthetic and real rs-fMRI data at the subject level. To demonstrate its significance at the population level, we apply this new framework to the UK Biobank data set, first for the discrimination of haemodynamic territories between balanced groups ($n = 24$ individuals in each) patients with an history of stroke and healthy controls and second, for the analysis of normal aging on the neurovascular coupling. Overall, we statistically demonstrate that a pathology like stroke or a condition like normal brain aging induce longer haemodynamic delays in certain brain areas (e.g. Willis polygon, occipital, temporal and frontal cortices) and that this haemodynamic feature may be predictive with an accuracy of 74 % of the individual's age in a supervised classification task performed on $n = 459$ subjects.

Keywords: BOLD signal, HRF, sparsity, low-rank decomposition, multivariate modeling, dictionary learning, UK Biobank

1. Introduction

1.1. Context

Functional magnetic resonance imaging (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). The estimation of the response is of a primary interest: a change in the haemodynamic response could be linked to the pharmacological mechanism of a drug (Do et al., 2020), the effect of healthy aging (West et al., 2019) or the consequence of a neuropathological process, for example Alzheimer’s disease (Asemani et al., 2017). Thus, the HRF could be considered as a precious bio-marker to investigate the neurovascular function of the brain in a 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 in the healthy brain or to uncover functional alteration in pathological condition.

1.2. Related works

Several methods have been designed to estimate this haemodynamic response in the case of task-related fMRI (tfMRI). In this setup, the participant is engaged in an experimental paradigm (EP) during the imaging session, which alternates between rest and task periods (Friston et al., 1998; Ciuciu et al., 2003; Lindquist and Wager, 2007; Pedregosa et al., 2015). Commonly, supervised HRF estimation methods fit a model to explain the observed BOLD signal from the EP (Goutte et al., 2000; Ciuciu et al., 2003; Lindquist and Wager, 2007; Vincent et al., 2010; Chaari et al., 2012; Pedregosa et al., 2015). A limitation of these approaches is that the EP is used as a surrogate for the neural activity. Therefore they do not account for possible delays in the subject’s responses compared to the task onsets, thus yielding a biased HRF estimate. Moreover, these methods cannot be used on resting-state fMRI data (rs-fMRI), where the participant is laying still in the MRI scanner and where no EP is available to serve as surrogate for neural activity.

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

36 the fMRI timescale of seconds, see (Glover, 1999) for Wiener filtering and smooth
37 estimation of activity-inducing signals and (Gitelman et al., 2003) for its general-
38 ization to study psychophysiologic interactions at the neuronal level. Importantly,
39 within a temporal fMRI deconvolution framework, most recent works have re-
40 lied on sparse regularization to recover a limited number of spike-like activations
41 directly in the time domain (Hernandez-Garcia and Ulfarsson, 2011; Caballero-
42 Gaudes et al., 2012) or imposing sparsity in the activelet-domain, which is a
43 wavelet basis that is tailored to the haemodynamic properties (Khalidov et al.,
44 2011). Next, a foundational work (Karahanoğlu et al., 2013) has proposed a
45 spatio-temporal model of the underlying activity-inducing signal including both
46 temporal and spatial sparsity-based regularization. By doing so, the recovered
47 neural activity profiles are used to define functional networks, hence converging to
48 the original approach proposed in (Wu et al., 2013) that uncovers functional net-
49 works from deconvolved BOLD signals. However, the total activation work (Kara-
50 hanoğlu et al., 2013) suffers from two main limitations: first, the HRF shape in
51 this modeling was fixed and kept constant throughout the brain, making this ap-
52 proach mainly suitable for normal populations; second, for computational reasons,
53 spatial regularization was introduced within predefined brain regions in a given
54 atlas (Karahanoğlu et al., 2013) instead of impacting whole brain activity.

55 Alternatively some recent works have suggested to estimate both the neural
56 activity and the HRF profile at the same time (Cherkaoui et al., 2019; Farouj
57 et al., 2019). These approaches are often referred to as *semi-blind deconvolution*
58 *schemes* of the BOLD signal as they do not fully estimate an unknown HRF shape
59 from scratch but instead try to uncover a transformation from a pre-existing HRF
60 profile such as the canonical shape. Both (Cherkaoui et al., 2019) and (Farouj et al.,
61 2019) rely on the hypothesis of a block-type property for the neural activation
62 signal as initially proposed in (Karahanoğlu et al., 2013) but are able to infer the
63 haemodynamic profile for instance with a limited parameterization to deal with
64 the magnitude and delay ambiguities between the neural input the haemodynamic
65 filter.

66 1.3. Goals and contributions

67 This paper extends preliminary contributions (Cherkaoui et al., 2019) and of-
68 fers a new algorithm that aims to fit a rich multivariate decomposition of the
69 BOLD data using a semi-blind deconvolution and low-rank sparse decomposition.
70 The model distinguishes two major parts in the BOLD signal: the neurovascu-
71 lar coupling and the neural activity signal. One of its main features is to tackle
72 the intra-subject haemodynamic variability by introducing regionwise HRF pro-
73 files over a brain parcellation, encoding the spatial variations of the neurovascular
74 coupling. Also, in regard to neural activity, we follow ideas developed in the con-
75 volutional dictionary learning literature to develop a low-rank approximation of

76 a signal (Grosse et al., 2007; Dupré La Tour et al., 2018). We thus model the
 77 neural activation signals as a combination of a limited number of piece-wise con-
 78 stant temporal profiles associated with spatial maps that capture the magnitude of
 79 functional networks. As such, the neural input signals are represented with their
 80 own spatial representation that departs from the neurovascular parcellation. We of
 81 course present a scalable optimization algorithm that is able to fit all parameters
 82 of interest in the whole brain in a reasonable computing time.

83 The rest of the paper is organized as follows. Section 2 introduces our modeling
 84 of the BOLD data and presents our semi-blind blind deconvolution algorithm.
 85 Next, our technique is validated on simulated data and on real rs-fMRI data at
 86 the individual level in Section 3. In Section 4 we illustrate the proposed framework
 87 at the population level on the large UK Biobank database. Two main applications
 88 were targeted, namely the neurovascular discrimination of patients with stroke
 89 episodes as compared to healthy controls and the prediction of brain age. Section 5
 90 discusses the potential impact of our method as well as its limitations. Conclusion
 91 and outlook are synthesized in Section 6.

92 2. Multivariate low-rank decomposition of the BOLD signal

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

95 *Notation*

96 In what follows, y_i denotes the i^{th} entry in vector \mathbf{y} . Let $\tilde{T} = T - L + 1$, the
 97 convolution of two signals $\mathbf{a} \in \mathbb{R}^{1 \times \tilde{T}}$ and $\mathbf{v} \in \mathbb{R}^{1 \times L}$ is denoted by $\mathbf{v} * \mathbf{a} \in \mathbb{R}^{1 \times T}$.
 98 For $\mathbf{A} \in \mathbb{R}^{P \times \tilde{T}}$, $\mathbf{v} \bar{*} \mathbf{A} \in \mathbb{R}^{P \times T}$ is obtained by convolving each line of \mathbf{A} with \mathbf{v} .
 99 For $\mathbf{V} \in \mathbb{R}^{P \times L}$, $\mathbf{V} \dot{*} \mathbf{A} \in \mathbb{R}^{P \times T}$ is obtained by convolving each line of \mathbf{A} with the
 100 corresponding line of \mathbf{V} . Let ∇ be the first-order difference operator such that
 101 $\forall \mathbf{x} \in \mathbb{R}^{\tilde{T}}$, $\nabla \mathbf{x} \in \mathbb{R}^{\tilde{T}-1}$ with $(\nabla \mathbf{x})_i = x_i - x_{i-1}$, $\forall i \in \{2, \dots, \tilde{T}\}$.

102 2.1. Linear and time-invariant modeling

103 2.1.1. 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. 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:

$$\mathbf{y}_j = \mathbf{v} * \tilde{\mathbf{a}}_j + \mathbf{e}_j . \quad (1)$$




Figure 1: Illustration of the voxel-wise 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 $\mathbf{a} * \mathbf{v} + \mathbf{e}$.

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

109 Typically, the HRF \mathbf{v} has a restricted support in time of about 20 s. The
 110 challenge with HRF modeling is to find a fair trade-off between a flexible model
 111 that is able to capture the true haemodynamic response in each brain area and a
 112 reliable one that limits overfitting by reducing the number of degrees of freedom
 113 (diminishing the variance). Since our approach estimates the neural activity along
 114 with the HRF, reducing the number of degrees of freedom is critical to avoid the
 115 aforementioned overfitting. In this paper, we assume M different HRF with \mathbf{v}_m
 116 being the HRF corresponding to the m^{th} region Θ_m . Numerous approaches have
 117 been proposed to model this haemodynamic response \mathbf{v}_m . The Finite Impulse Re-
 118 sponse (FIR) (Dale, 1999; Glover, 1999) model does not assume any particular
 119 shape for the HRF which make it very flexible but prone to data overfitting in the
 120 presence of noise. Regularization has thus been introduced to constrain the overall
 121 HRF shape in FIR models and limit their tendency to overfitting, see for instance
 122 penalization over the second-order derivative to end up with physiologically plau-
 123 sible smooth HRF estimates (Ciuciu et al., 2003; Casanova et al., 2008). Alterna-
 124 tively, the HRF has been modeled as a linear decomposition of predefined atoms
 125 such as B-splines (Zhang et al., 2007; Vakorin et al., 2007), wavelets (Khalidov
 126 et al., 2011), a sensitivity-selected set (Woolrich et al., 2004) or more physiologi-
 127 cally informed patterns such as the canonical HRF and its derivatives in time and
 128 with respect to the dispersion parameter (Friston et al., 1998). All these methods
 129 intend to capture fluctuations in haemodynamic delay or shape with the mini-
 130 mum number of parameters. Last, to constrain even more the parameter values
 131 and reduce variance estimates, parametric models such as the *inverse logit trans-*
 132 *form* (Lindquist and Wager, 2007) have been proposed and successfully tested
 133 when benchmarking over multiple fMRI data sets (Lindquist et al., 2009). In this
 134 work, we propose the time dilation HRF model (Cherkaoui et al., 2019), which cap-
 135 tures the haemodynamic delay by dilating a reference HRF pattern: $\mathbf{v} = \mathbf{v}_{ref}(\delta t)$

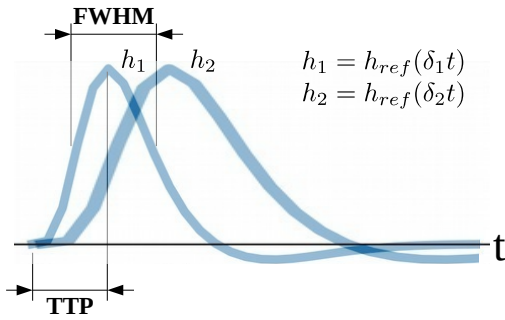


Figure 2: Illustration of two haemodynamic response functions (HRF) denoted h_1, h_2 with the full-width-at-half-maximum (FWHM) and the time-to-peak (TTP) (of h_1) depicted. Here we illustrate these two HRFs with different dilation parameters (δ_1, δ_2) such as $\delta_2 \leq \delta_1$, leading to $\text{TTP}(h_1)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

136 where \mathbf{v}_{ref} is a reference temporal profile, here the canonical HRF. This approach
 137 is efficient while simple as it encodes delay fluctuations through a single scalar
 138 parameter δ (one degree of freedom). One limitation of this choice is that it leads
 139 to the simultaneously variation of the full width-at-half-maximum (FWHM) of the
 140 HRF and its time-to-peak (TTP) (see also Fig. 2). Our voxelwise model reads as
 141 follows:

$$\mathbf{y}_j = \mathbf{v}_\delta * \tilde{\mathbf{a}}_j + \mathbf{e}_j . \quad (2)$$

142 2.1.2. 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 multivariate model reads as follows:

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

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. (3) as follows:

$$\mathbf{Y} = \left(\sum_{m=1}^M \Theta_m^\top \mathbf{v}_{\delta_m} \right) \ast \tilde{\mathbf{A}} + \mathbf{E} , \quad (4)$$

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

147 In our work, we rather introduce a low-rank constraint and learn both K temporal
 148 atoms (with $K \ll P$) and corresponding spatial maps. These maps encode various
 149 functional networks, each of them being summarized by specific neural activation
 150 profile. Mathematically, this can be modeled by replacing each vector $\tilde{\mathbf{a}}_j$ in Eq. (4)
 151 with a linear combination of neural activation patterns $\mathbf{Z} = (\mathbf{z}_k)_{k=1}^K \in \mathbb{R}^{K \times \tilde{T}}$, with
 152 $\mathbf{z}_k \in \mathbb{R}^{1 \times \tilde{T}}$, modulated in space by the spatial maps $\mathbf{U} = (\mathbf{u}_k)_{k=1}^K \in \mathbb{R}^{K \times P}$, with
 153 $\mathbf{u}_k \in \mathbb{R}^{1 \times P}$, such that: $\tilde{\mathbf{A}} = \sum_{k=1}^K \mathbf{u}_k^\top \mathbf{z}_k$. In other words, the spatial configuration
 154 \mathbf{u}_k encodes which voxels are linked to a given neural activation profile $\mathbf{z}_k \in \mathbb{R}^{1 \times \tilde{T}}$.
 155 Note that a voxel may belong to different functional networks. This is coherent
 156 with the fact that a voxel contains about one hundred thousands neurons for a
 157 typical spatial resolution (1.5 mm isotropic), and thus the underlying neural acti-
 158 vation signals are combined with possibly different temporal fingerprints. Finally,
 159 our forward model for BOLD fMRI data is given by Eq. 5 (see also Fig. 3):

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

160 2.2. Prior information and regularizing constraints

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

168 The forward model described in Eq. (5) is *trilinear* as it depends on both un-
 169 known spatial (\mathbf{U}) and temporal (\mathbf{Z}) inputs and unknown convolution filters (\mathbf{v}_{δ_m}).
 170 This means that any timing variation in neural activation signals \mathbf{Z} can be sym-
 171 metrically compensated by the opposite variation in the haemodynamic filter \mathbf{v}_{δ_m}

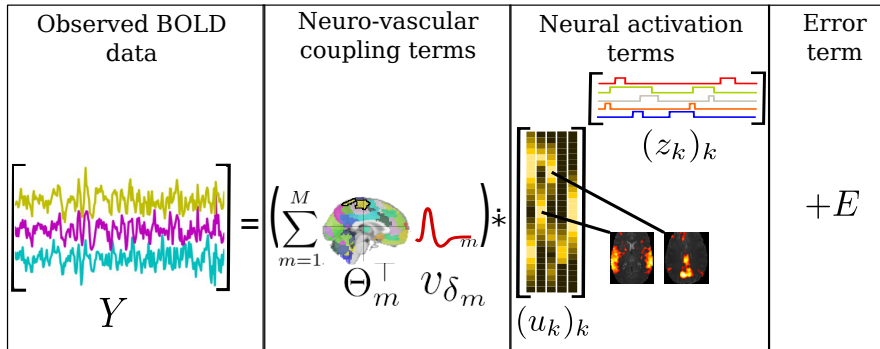


Figure 3: Illustration of the low-rank multivariate BOLD signal model (the colors are here for illustrative purposes). \mathbf{Y} stands for the observed BOLD data, \mathbf{v}_m the m^{th} HRF, Θ_m represents the predefined brain parcellation, $\mathbf{U} = (\mathbf{u}_k)_{k=1}^K$ defines the K spatial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

172 or that any sign or scale variation in \mathbf{U} can be compensated by the inverse vari-
173 ation in \mathbf{Z} or \mathbf{v}_{δ_m} . To overcome these sign, scale and time-shift ambiguities, we
174 introduce some regularizing constraints in the computation of the parameter esti-
175 mates. First, to avoid any sign ambiguity in the convolution between the neural
176 input signals and the haemodynamic filter, the HRF filter has a constant and posi-
177 tive maximal amplitude (see Fig. 2). Second, as there is an interplay between
178 the spatial and temporal components in the input signals, we also impose a non-
179 negativity constraint over the entries of each spatial map \mathbf{u}_k , i.e. $\forall j, k_{kj} \geq 0$.
180 By doing this, we only allow the deactivation of a functional network to be en-
181 coded in the corresponding \mathbf{z}_k . Third, to deal with the scale ambiguity, we impose
182 $\|\mathbf{u}_k\|_1 = \eta, \forall k = 1, \dots, K$, where $\eta \in \mathbb{R}_+$ is a user-defined parameter that sets
183 the magnitude of each spatial map. As our HRF model has a constant maximal
184 amplitude, only the neural activity signals \mathbf{Z} capture the observed BOLD signal
185 fluctuations. Forth, to deal with the time-shift ambiguity, we constrain the dila-
186 tion parameter δ in the HRF model to be in $[0.5, 2.0]$. This should also permit
187 the recovery of physiologically plausible haemodynamic delays. Complementary
188 to these constraints, akin to (Caballero-Gaudes et al., 2012; Karahanoglu et al.,
189 2013; Cherkaoui et al., 2019) we will also assume the neural activation signals \mathbf{Z} to
190 be temporally piecewise constant or equivalently assume their first-order deriva-
191 tive $\nabla \mathbf{z}_k, \forall k$ to be sparse. Practically speaking, we add a total variation (TV)
192 regularization term to our model.

193 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{U} and the neural activation signals \mathbf{Z} from the model depicted in Eq. (5) and the aforementioned regularizing constraints. The

estimates for these parameters can be obtained by solving the following constrained minimization problem:

$$\begin{aligned} & \arg \min_{(\mathbf{U}, \mathbf{Z}, \boldsymbol{\delta})} \frac{1}{2} \left\| \mathbf{Y} - \left(\sum_{m=1}^M \boldsymbol{\Theta}_m^\top \mathbf{v}_{\delta_m} \right) \ast \left(\sum_{k=1}^K \mathbf{u}_k^\top \mathbf{z}_k \right) \right\|_F^2 + \lambda \sum_{k=1}^K \|\nabla \mathbf{z}_k\|_1, \quad (6) \\ & \text{subject to } \forall k, \|\mathbf{u}_k\|_1 = \eta, \quad \forall j, u_{kj} \geq 0, \quad \forall m, \delta_m \in [0.5, 2.0]. \end{aligned}$$

194 The Gaussian noise hypothesis leads to a quadratic data fidelity term, to measure
 195 how well we reconstruct the observed fMRI signals. Moreover, in alignment with
 196 with Caballero-Gaudes et al. (2012); Karahanoglu et al. (2013); Cherkaoui et al.
 197 (2019), we enforce the first-order derivative of the temporal atoms $(\nabla \mathbf{z}_k)_{k=1}^K$ to
 198 be sparse in order to constrain each \mathbf{z}_k to a piecewise constant signal. For that
 199 purpose, we use a TV regularization term, which corresponds to the ℓ_1 norm of the
 200 gradient in time $\nabla \mathbf{z}_k$. Importantly, this modeling of the neural activation signals
 201 allows us to fully adapt to task-fMRI and rs-fMRI experiments and to perform
 202 paradigm-free fMRI data analyses. In the first case, as the task-related BOLD
 203 signal is classically modeled as the convolution of an input block signal, repre-
 204 senting the experimental paradigm (the onsets of the stimulus trials) with a HRF
 205 filter, we can recover neural activation signals close to the experimental paradigm.
 206 The neural activation profiles being inferred from the data, we can estimate both
 207 input signals corresponding to block and event-related designs, the only difference
 208 between the two being the length of activation blocks, which in the latter case can
 209 be restricted to a single time point. More interestingly, the proposed framework
 210 is even more appealing for processing rs-fMRI data and uncover spontaneous and
 211 time-varying fluctuations of brain activity as the the block duration may change
 212 from one instance to the next.

213 Moving to the technical aspects for solving the constrained optimization prob-
 214 lem (6), it is worth mentioning that it is not globally convex. However, when \mathbf{U}
 215 and $\boldsymbol{\delta}$ are fixed, problem (6) becomes convex in \mathbf{Z} and similarly when \mathbf{Z} and $\boldsymbol{\delta}$
 216 are fixed, it becomes convex in \mathbf{U} . Our minimization strategy of Eq. (6) thus re-
 217 lies on a block-coordinate descent algorithm, where we alternate the minimization
 218 between the two convex problems in \mathbf{U} and \mathbf{Z} followed by the non-convex one
 219 involving \mathbf{V} . Also, the non-negativity constraints are activated when solving for
 220 the spatial maps \mathbf{U} and the boundary constraints over $\boldsymbol{\delta}$ are handled in parallel
 221 for each m , i.e. each HRF pattern when solving for \mathbf{V} . Algorithm 1 details these
 222 three main steps.

223

Algorithm 1: Multivariate deconvolution and low-rank decomposition of the BOLD signal.

Input: BOLD signal \mathbf{Y} , ϵ

1 initialization: $\forall k, \mathbf{z}_k^{(0)} = \mathbf{0}_{\bar{T}}, \mathbf{u}_k^{(0)} = \mathbf{u}_k^{(init)}, \boldsymbol{\delta}^{(0)} = \boldsymbol{\delta}^{(init)}, i = 1$;

2 **repeat**

3 Estimate the temporal atoms $\mathbf{Z}^{(i)}$ with fixed $\mathbf{U}^{(i-1)}$ and $\boldsymbol{\delta}^{(i-1)}$:

$$(\mathbf{z}_k^{(i)})_k = \arg \min_{(\mathbf{z}_k)_k} \frac{1}{2} \left\| \mathbf{Y} - \left(\sum_{m=1}^M \boldsymbol{\Theta}_m^\top \mathbf{v}_{\delta_m}^{(i-1)} \right) * \left(\sum_{k=1}^K \mathbf{u}_k^{(i-1)\top} \mathbf{z}_k \right) \right\|_F^2 + \lambda \sum_{k=1}^K \|\nabla \mathbf{z}_k\|_1 .$$

4 Estimate the spatial maps $\mathbf{U}^{(i)}$ with fixed $\mathbf{Z}^{(i)}$ and $\boldsymbol{\delta}^{(i-1)}$:

$$(\mathbf{u}_k^{(i)})_k = \arg \min_{(\mathbf{u}_k)_k} \frac{1}{2} \left\| \mathbf{Y} - \left(\sum_{m=1}^M \boldsymbol{\Theta}_m^\top \mathbf{v}_{\delta_m}^{(i-1)} \right) * \left(\sum_{k=1}^K \mathbf{u}_k^\top \mathbf{z}_k^{(i)} \right) \right\|_F^2 ,$$

subject to $\{\forall k, \|\mathbf{u}_k\|_1 = \eta \text{ and } \forall j, u_{kj} \geq 0\}$.

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

$$(\mathbf{v}_m^{(i)})_m = \arg \min_{(\delta_m)_m} \frac{1}{2} \left\| \mathbf{Y} - \left(\sum_{m=1}^M \boldsymbol{\Theta}_m^\top \mathbf{v}_{\delta_m} \right) * \left(\sum_{k=1}^K \mathbf{u}_k^{(i-1)\top} \mathbf{z}_k^{(i)} \right) \right\|_F^2 ,$$

subject to $\delta_m \in [0.5, 2.0]$.

5 **until** $\frac{J((\mathbf{z}_k^{(i-1)})_k, (\mathbf{u}_k^{(i-1)})_k, (\mathbf{v}_m^{(i-1)})_m) - J((\mathbf{z}_k^{(i)})_k, (\mathbf{u}_k^{(i)})_k, (\mathbf{v}_m^{(i)})_m)}{J((\mathbf{z}_k^{(i-1)})_k, (\mathbf{u}_k^{(i-1)})_k, (\mathbf{v}_m^{(i-1)})_m)} \leq \epsilon$;

224

225 In regard to the $(\mathbf{z}_k)_{k=1}^K$ step, we performed the minimization using an adaptive-
 226 restart accelerated forward-backward algorithm (O’Donoghue and Candes, 2015).
 227 In regards to the $(\mathbf{u}_k)_{k=1}^K$ step, we first benchmark various algorithms in the dic-
 228 tionary learning literature and selected the most efficient, namely the one used to
 229 update the dictionary in Mairal et al. (2009). Last, for the minimization with re-
 230 spect to $(\delta_m)_{m=1}^M$ we used the accelerated forward-backward algorithm (Combettes
 231 and Pesquet, 2009) after checking that it leads to a correct estimation of $\boldsymbol{\delta}$. The
 232 reader can find all details of the gradient computation *w.r.t* $(\mathbf{z}_k)_{k=1}^K$ and $(\mathbf{u}_k)_{k=1}^K$
 233 in Cherkaoui et al. (2019), for the gradient *w.r.t* $\boldsymbol{\delta}$ the computation is reported
 234 in the supplementary material Section 7.1. For each step, we implemented the
 235 corresponding gradient (i.e. forward move) in an efficient manner to limit multiple
 236 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{Z} 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. $g_z(\mathbf{z}_k) = \lambda \|\nabla \mathbf{z}_k\|_1$. 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{U}), we also proceed separately on the K components: the proximity operator of $g_u(\mathbf{u}_k) = \mathbb{1}_{\|\mathbf{u}_k\|_1=\eta} + \mathbb{1}_{u_{kj}\geq 0}$ where $\mathbb{1}$ stands for the indicator function², is given by:

$$\text{prox } g_u(\mathbf{u}_k) = [(u_{kj} - \mu)_+]_{1\leq j\leq P} \quad (7)$$

where μ is defined as $\sum_{j=1}^P \max\{0, u_{kj} - \mu\} = \eta$ 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). For the rest of the paper, 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 1 converges to a local minimizer Eq. (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 $(\mathbf{u}_k)_{k=1}^K$, 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 $(\mathbf{z}_k)_{k=1\dots K}$ to zero and each entry of $(\delta_m)_{m\in\{1..M\}}$ to 1.

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 4.

3.1. Numerical simulations

3.1.1. 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

¹https://pypi.org/project/prox_tv/

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

264 in blue in Fig. 4[right panels] – we integrated over time these signals and con-
 265 volved them with a predefined HRF \mathbf{v}_δ to yield two corresponding pure BOLD
 266 time series. The chosen HRF has length $L = 25$ and is shown in blue in Fig. 4[left
 267 panel]. For the sake of simplicity, we considered a single HRF profile ($M = 1$)
 268 in this synthetic setting, so the haemodynamic properties were supposed constant
 269 in space. We then assigned these BOLD signals to spatial locations. Hence, we
 270 defined $K = 2$ corresponding 2D maps $\mathbf{U} = (\mathbf{u}_k)_{k=1}^2$ (10×10 , i.e. $P = 100$).
 271 Each spatial map has a single activating region consisting of 4 pixels. Each ac-
 272 tive pixel has a randomly drawn non-negative magnitude, the other ones being
 273 set to zero. Then, we normalized each map by its ℓ_1 -norm. Finally we added
 274 Gaussian random noise to produce observed, i.e. noisy BOLD signals \mathbf{Y} of length
 275 $T = \tilde{T} + L - 1 = 524$ scans ($TR = 1$ s) with a signal-to-noise-ratio (SNR) of -1 dB.
 276 The mean synthetic BOLD signals are reported in black traces in Fig. 4(a)-(b) for
 277 both activated regions (in bottom panels) while standard deviation across acti-
 278 vated voxels is encoded by transparency around these mean curves.

279 3.1.2. Numerical results

280 In a first step, we only estimated the pair (\mathbf{Z}, \mathbf{U}) from the synthetic fMRI time
 281 series \mathbf{Y} and kept the HRF profile \mathbf{v} constant. The results are reported in Fig. 4(a).
 282 The HRF shape used in this deconvolution process is shown in green in Fig. 4(a)
 283 and actually differs from the true shape used for simulating the data. Because
 284 of this discrepancy in terms of haemodynamic delay and peak magnitude, the
 285 neural activation signals are not properly recovered (orange traces in Fig. 4(a)).
 286 The magnitude of the estimates $\hat{\mathbf{Z}}$ is much larger than the true one. This is
 287 partly due to compensate for the smaller magnitude of the HRF (green trace in
 288 Fig. 4(a)) used for deconvolving the BOLD signals. Consequently, the residual
 289 mean square errors (RMSEs) computed on the neural activation signals are pretty
 290 large. However, we noticed that in both spatial maps, the non-negative magnitudes
 291 $\hat{\mathbf{U}}$ are very well estimated. This is a direct consequence of using non-overlapping
 292 activating regions for the two neural traces.

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

302 These results on synthetic data confirmed the good expected behavior of the
 303 proposed method. From a computational viewpoint, the estimation with constant

304 HRF ran in 0.5s while the full estimation took 1s approximately on a machine
305 with 15 GB of RAM and an Intel processor i7-7600U (2 physical cores, 2.80 GHz).

306 3.2. Single-subject analysis on rs-fMRI data

307 3.2.1. Data set and parameter setting

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

321 In this subsection, we manually set the temporal regularization parameter to
322 $\lambda_f = 0.8$ (remember $\lambda_f \in [0, 1]$). This setting achieves a bias-variance trade-off
323 between two extreme situations, namely data overfitting ($\lambda_f = 0$) on one hand
324 and entirely sparse neural activation signals ($\lambda_f = 1$ as $\lambda = \lambda_{\max}$) on the other
325 hand. The question of the unsupervised tuning of λ_f is critical. It could be driven
326 either from a statistical viewpoint (e.g. using the maximum likelihood criterion)
327 that characterizes how likely the measured time series may be observed or using an
328 external task and its corresponding metric such as classification performance (e.g.
329 accuracy in prediction). Because the former approach does not admit a closed
330 form solution, we explored in the Supplementary Material (see Section 7) the
331 impact of changing the temporal regularization (i.e. amount of sparsity in the
332 activation neural signals) on the spatial map of haemodynamic delays. The spa-
333 tial consistency we reported across regularization levels gave us confidence on the
334 haemodynamic parameter estimates to be further used in subsequent classification
335 tasks (see Section 4). In this setting, we implemented a cross-validation step with
336 a leave-one-out loop to tune λ_{\max} in an unsupervised way.

337 In regard to the number of spatio-temporal atoms K , we set it using the ex-
338 plained variance (or R^2 -score) as target metric in a preliminary study. For this set

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

339 of parameters, the model estimation took around 1 minute on a machine with 15
 340 GB of RAM and an Intel processor i7-7600U (2 physical cores, 2.80 GHz).

341 3.2.2. Results

342 **Model selection.** The first question we addressed on real rs-fMRI data was to
 343 optimally set the number of spatio-temporal atoms K and to find the best compro-
 344 mise between model complexity and model accuracy. For this purpose, we looked
 345 at two complementary criteria. The first one is standard and corresponds to the
 346 R^2 score that quantifies the variance explained by model (5) over the total sum
 347 of squares whereas the second one is given by the determinant of the correlation
 348 matrix between the neural activation signals.

349 The R^2 -score is defined as follows: $R^2 = 1 - \frac{SS_{\text{res}}}{SS_{\text{tot}}}$ where SS_{tot} quantifies the vari-
 350 ance of the data \mathbf{Y} and SS_{res} the variance of the residuals after fitting model (5)
 351 by minimizing the cost function described in Eq. (6). The R^2 -score may vary from
 352 $-\infty$ in pathological cases to 1 for a perfectly matching model. A good model is
 353 normally associated with $R^2 > 0$ and means that the L_2 norm of the residual is
 354 lower than the variance of the data. We therefore ran multiple model fitting for
 355 K in a range of $\{2, 3, 4, \dots, 10, 15, 20, \dots, 50\}$. The results are shown in Fig. 5(a)
 356 and illustrate that the model accuracy first increases as a function of K up to
 357 reaching a plateau around $R^2 \simeq 0.55$ for $K = 20$. So adding more spatio-temporal
 358 components no longer improves its ability to capture variability in the data while
 359 it becomes more complex.

360 The second information measure we used to help us select K was based on the
 361 determinant of the correlation matrix $\Sigma_K = (\mathbb{E}[(\mathbf{z}_k - \mathbf{m}_k)(\mathbf{z}_\ell - \mathbf{m}_\ell)^T / \sigma_k^2 \sigma_\ell^2])_{k,\ell}$
 362 between the temporal atoms $(\mathbf{z}_k)_{k=1}^K$. The quantities σ_k^2 and σ_ℓ^2 define the variance
 363 of the neural activation signals \mathbf{z}_k and \mathbf{z}_ℓ . As Σ_K is semi-positive definite with
 364 entries between 0 and 1, its eigenvalues are positive or null and so its determinant
 365 varies between 0 and 1: $\det \Sigma_K = 1$ when matrix Σ_K defines a basis, which means
 366 that all atoms are orthogonal and decorrelated like in a PCA decomposition. In
 367 contrast, $\det \Sigma_K = 0$ when matrix Σ_K is not of full rank so at least one atom
 368 could be obtained as a linear combination from the others. Therefore, as before
 369 we ran multiple model fitting for K in a range of $\{2, 3, 4, \dots, 10, 15, 20, \dots, 50\}$
 370 and we plotted in Fig. 5(b) the evolution of the determinant of Σ_K as a function
 371 of K . The results show us that beyond $K \geq 20$, we get a correlation matrix with
 372 $\det \Sigma_K \leq 10^{-10}$ which tends to zero. According to this criterion, from the we
 373 should therefore not exceed 20 temporal atoms. Thus, from the R^2 -score criterion
 374 and this $\det \Sigma_K$ criterion, in our following experiments we will keep $K = 20$.

375 **Analysis of spatial decomposition.** Fig. 6 shows the spatial maps of this
 376 spatio-temporal decomposition for this individual and Tab. 1 summarizes the list
 377 of main regions and functional networks retrieved in this setting.

378 It is worth mentioning that the sensory networks (visual, auditory and motor)
 379 are quite well retrieved by a single or multiple components, respectively located
 380 respectively in the occipital (components 10, 16 and 19), temporal (component
 381 7) and motor (components 6, 15 and 18) cortices. The maps associated with the
 382 motor network are split and lateralized (6 and 18 on the right hemisphere while
 383 15 in the left). The different areas of the visual system are split too between the
 384 primary visual cortex (component 10) and the extrastriate cortex (components 16,
 385 19). The language system has also been identified by a single component. Part
 386 of the well known intrinsic resting-state networks (RSN; Menon 2015) have been
 387 captured: (i) the right and left-lateralized fronto-parietal resting-state networks
 388 spatially similar to the bilateral dorsal attention network are captured by com-
 389 ponent 1 and 4, respectively while the left and right lateral frontoparietal central
 390 executive networks appear in component 9. The default mode network (DMN),
 391 which deactivates during demanding cognitive tasks is represented in component 2
 392 and 20: the angular gyrus (AG) appear in both components while the posterior cin-
 393 gulate cortex (PCC) is captured only by component 20 and the medial prefrontal
 394 cortex (mPF) by component 2. However, we found that component 9 actually
 395 mixes the left AG in the DMN with a left-lateralized fronto-parietal network that
 396 perfectly matches the CEN (Menon, 2015). In contrast, we did not clearly re-
 397 trieve neither the salience network – usually anchored in anterior insula and dorsal
 398 anterior cingulate cortex – nor the right CEN.

399 **Spatio-temporal decomposition.** To fully illustrate our method on real rs-
 400 fMRI data, we show the whole set of output features (neural activation signals,
 401 spatial maps, HRF shapes) in Fig. 7. We also depict a voxel-based denoised
 402 BOLD signal reconstructed a the convolution between the neural input and the
 403 HRF estimate. Fig. 7(a) represents together a neural activation signal in the
 404 primary visual cortex and the corresponding spatial map (component 10 in the
 405 above mentioned decomposition). The proposed axial views allow us to identify the
 406 primary visual cortex and the calcarine fissure. Fig. 7(b) depicts similar features
 407 in the DMN (component 2) and the Pearson correlation coefficient with the neural
 408 time course in the visual cortex. Its negative value confirms a negative correlation
 409 between the task-positive and the DMN network. Both time courses actually
 410 present alternating periods of positive and negative activity but they are almost
 411 uncorrelated. Fig. 7(c)-(d) illustrate the fastest and slowest HRF time courses
 412 estimated in the regions of interest depicted in red. The fastest haemodynamic
 413 response ($\text{FWHM}_f = 5.1\text{s}$) was found in the middle temporal gyrus while the
 414 slowest ($\text{FWHM}_s = 8.0\text{s}$) is located in the frontal orbital cortex. Fig. 7(e) finally
 415 shows how well our approach is able to fit the rs-fMRI time course measured
 416 in voxel marked by the black cross in Fig. 7(c). The neural activation signal is
 417 piecewise constant and ahead in time compared to the BOLD time series. Once

Table 1: **Taxonomy of brain regions and functional networks** involved in the spatio-temporal decomposition (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 and have been identified from the [AAL template](#).

Network	# Comp.	Brain areas
Visual	10	R calcarine fissure and surrounding cortex
	16	L 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

418 convolved with the HRF profile, the denoised BOLD signal appears as a smoother
 419 version of the measured BOLD time course: its magnitude is smaller and its
 420 fluctuations in time are slower. This is a direct consequence of the temporal
 421 regularization used to recover sparse input signals.

422 **Analysis of correlation structure.** Next, to go beyond the spatial analysis,
 423 Fig. 8 depicts the correlation matrix between the corresponding neural activation
 424 signals. It is then insightful to notice that the correlation between the multiple
 425 components in a given network are quite strong. For the visual network we ob-
 426 served a correlation coefficient varying between 0.35 and 0.69, the largest value
 427 being reached for areas located both in the extrastriate cortex. The same con-
 428 clusion holds in the motor network with a correlation level varying between 0.5
 429 and 0.65. In regard to the DMN, component 2 plays the role of a hub as it
 430 correlates with components 9 and 20 pretty strongly⁵ between 0.27 and 0.4. How-
 431 ever, component 9 is almost decorrelated from component 20 as it mixes regions
 432 in the dAN and the DMN. Overall, this analysis shows that the proposed ap-
 433 proach does not separate RSN in single components. However, it still achieves a
 434 meaningful decomposition. For illustrative purposes, the spatial decomposition for
 435 $K \in \{5, 8, 10, 15, 30, 40, 50\}$ and the corresponding correlation matrices between
 436 the temporal atoms are reported in the Supplementary Material. Interestingly,
 437 these results show that some RSNs are either not recovered or mixed together for
 438 small $K \leq 15$ whereas for large values of $K \in \{25, 30, \dots, 50\}$ each RSN is split in
 439 multiple components. Increasing the number of components extends the range of
 440 variation of the correlation coefficients in both positive and negative senses. This
 441 confirms that our approach does not have statistical independence guaranties like
 442 PCA and ICA do when decomposing the neural activity.

443 3.3. haemodynamic estimation stability over time

444 The shape of the HRF is controlled by the neurovascular coupling including
 445 both neural and non neural factors such as glial cell activity, cerebral energy
 446 metabolism, and the cerebral vasculature. Abnormalities in the local vascular sys-
 447 tem or cell communication due to pathological state or changes in cerebral blood
 448 flow upon psychoactive drugs could influence this haemodynamic response. As we
 449 expect the HRF estimate to be stable if none of those events took place, we propose
 450 to study the intra-subject stability of HRF estimates over time, namely between
 451 consecutive time periods. For doing so, we compare the intra-subject variability of
 452 the HRF whole brain dilation parameter vector δ to the inter-subject variability
 453 of the same quantity. We thus introduce two reference ℓ_2 distances, namely the
 454 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

⁵no statistical test performed at the individual level

455 vectors of spatially aggregated HRF dilation parameters that were estimated over
 456 two periods of time T_1 and T_2 in the same individual s . Similarly, for any pair of
 457 subjects (s_1, s_2) and a given period T , we measure the between-subject distance
 458 between $(\boldsymbol{\delta}_T^{s_i})_i$ vectors as follows: $\text{BS}(\boldsymbol{\delta}_T^{s_1}, \boldsymbol{\delta}_T^{s_2}) = \|\boldsymbol{\delta}_T^{s_1} - \boldsymbol{\delta}_T^{s_2}\|_2^2$. The goal is then to
 459 compare the within- and between-subject distances across individuals and show
 460 that the intra-subject variability is significantly lower than the inter-subject one
 461 over a sufficiently large population.

462 3.3.1. Data set and numerical analysis

463 We selected 100 healthy subjects from the Human Connectome Project (HCP)
 464 data set (Van Essen et al., 2013) at random. We used this data set because
 465 of the availability of a 12-min long rs-fMRI run with a short time of repeti-
 466 tion ($TR=0.72\text{s}$), see Glasser et al. (2013) for a full description of the acqui-
 467 sition parameters and the pre-statistics processing steps. In this rs-fMRI run for
 468 each individual, we extracted two segments of 4 minutes each, denoted as T_1 and
 469 T_2 hereafter, the first and last parts of the recording. We then applied the pro-
 470 posed multivariate spatio-temporal decomposition to each segment using $K = 8$
 471 spatio-temporal atoms $(\mathbf{z}_k, \mathbf{u}_k)_{k=1}^K$ and a brain atlas $\Theta = (\Theta_m)_{m=1}^M$ (Desikan et al.,
 472 2006) composed of $M = 96$ regions of interest (ROIs). This haemodynamic brain
 473 parcellation thus yields 96 HRF dilation parameters $\boldsymbol{\delta} = (\delta_m)_{m=1}^M$ for each in-
 474 dividual. In practice, in the definition of $\text{WS}(\cdot, \cdot)$ and $\text{BS}(\cdot, \cdot)$, the true vectors
 475 $\boldsymbol{\delta}_{T_j}^{s_i}$ ($i = 1, \dots, 100, j = 1, 2$) have been replaced by their estimates $\widehat{\boldsymbol{\delta}}_{T_j}^{s_i}$ computed
 476 by solving Eq. (6) for the two 4-min rs fMRI data sets (T_1 and T_2). To make sure
 477 that our conclusions hold for a large scale of temporal regularization parameters,
 478 we spanned the range $\lambda_f \in [0, 1]$ and repeated the same procedure over 10 discrete
 479 values of λ_f within this interval.

480 3.3.2. Results

481 In Fig. 9, the box plot in blue shows the within-subject distance $\text{WS}(\widehat{\boldsymbol{\delta}}_{T_1}^{s_i}, \widehat{\boldsymbol{\delta}}_{T_2}^{s_i})$
 482 between the two 4-min rs-fMRI segments for all individuals and across 5 values
 483 of λ_f covering the whole interval $[0, 1]$. The orange and green box plots in Fig. 9
 484 depict the between-subject distances computed over the first and second segments
 485 respectively, namely $\text{BS}(\widehat{\boldsymbol{\delta}}_{T_1}^{s_i}, \widehat{\boldsymbol{\delta}}_{T_1}^{s_j})$ and $\text{BS}(\widehat{\boldsymbol{\delta}}_{T_2}^{s_i}, \widehat{\boldsymbol{\delta}}_{T_2}^{s_j})$ with $i \neq j$. We observed that
 486 the within-subject (i.e. inter-segment) variability is systematically lower than the
 487 between-subject variability and that all metrics remain stable across regulariza-
 488 tion levels. To go further, we performed a statistical analysis (paired t -test) by
 489 comparing the mean of the WS and BS distributions and we obtained significant p -
 490 values ($p < 10^{-8}$) showing that the within-subject haemodynamic variability is sig-
 491 nificantly lower than the between-subject fluctuations. In contrast, the statistical
 492 inter-individual comparison between the two segments is not significant ($p \simeq 10^{-2}$).

493 These results are valid for all tested regularization levels indicating a minor im-
 494 pact of the regularization parameter onto the haemodynamic parameter estimate.
 495 In sum, this analysis demonstrates that the whole brain characterization of the
 496 vascular system remains stable in a given individual between two periods shortly
 497 spaced in time, compared to the same analysis between individuals and so that
 498 the haemodynamic response discriminates each subject from the others.

499 4. Clinical validation at the population level

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

509 4.1. Characterization of patients with an history of stroke

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

518 We considered 24 patients of both genders and various ages who suffered from
 519 a stroke in the past from the UK Biobank database. For comparison purposes, we
 520 selected 24 healthy controls matched in age and gender from the same database.
 521 We applied the same decomposition ($K = 20$, $M = 96$, same λ_f) to each patient
 522 and healthy control. Fig. 10(a) and Fig. 10(b) show respectively the corresponding
 523 normalized maps of haemodynamic dilation parameters $(\delta_m)_{m=1}^{96}$ in a healthy con-
 524 trol and stroke patient, respectively. The normalization has been done by dividing
 525 all dilation parameter values by their within-subject average, namely $\bar{\delta}^{\text{HC}}$ and $\bar{\delta}^{\text{SP}}$
 526 respectively. We first observed that the dilation parameters were larger in average
 527 in the healthy condition compared to stroke ($\bar{\delta}^{\text{HC}} > \bar{\delta}^{\text{SP}}$). This corresponds to
 528 shorter and more homogeneous TTPs in the brain in the healthy condition. The
 529 shortest TTP found in the healthy control was actually located in the primary

530 visual cortex (axial slice, $z=-2$, left hemisphere), a result consistent with the liter-
 531 ature on fastest haemodynamic responses often detected in visual areas (Handw-
 532 erker et al., 2004; Badillo et al., 2013). In contrast, Fig. 10(b) illustrates that the
 533 haemodynamic dilation parameters δ^{SP} are smaller (so the TTPs longer) in the
 534 stroke patient. Also, we found less variability in the healthy subject since the dif-
 535 ference between the maximum and minimum TTPs were smaller ($\Delta_{\text{TTP}}^{\text{HC}} = 1.25$ s)
 536 compared to the stroke patient ($\Delta_{\text{TTP}}^{\text{SP}} = 2.25$ s). Importantly, Fig. 10(a) illustrates
 537 the relative symmetry of haemodynamic territories that exists in normal subjects
 538 between both hemispheres (Raemaekers et al., 2018).

539 On the contrary, Fig. 10(b) reveals a wider asymmetry between the two hemi-
 540 spheres in the stroke patient. Interestingly, in this patient we noticed the presence
 541 of larger TTPs in the middle left precentral gyrus and left motor cortex (resp.
 542 $z=44$ and $z=60$), namely the brain regions supposedly impacted by the stroke
 543 episode.

To go one step further, we quantified the spatial asymmetry of the haemo-
 dynamic structure within each individual. To this end, we computed the inter-
 hemispheric haemodynamic ℓ_2 distance (IHD) between the HRF dilation param-
 eter 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:

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

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

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

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

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) 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 t -scores to compare the dilation parameters in each region. The results are presented in Fig. 11(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 log-transformed p-values, i.e. $-\log_{10} p_{\text{val}}$ (shown in Fig. 11(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,

602 (ii) HRF shapes and (iii) haemodynamic dilation parameters. We did not consider
603 the spatial activity maps as input features in these models as they do not permit
604 to perform dimension reduction. We trained these LR models using the Scikit-
605 Learn software (Pedregosa et al., 2011) to predict the age label for each subject
606 (1 for the elderly people, 0 for middle-aged people). A ℓ_2 -norm regularization was
607 used in the estimation of the LR model parameters with an hyper-parameter $\beta >$
608 0. We grid-searched the temporal regularization parameter λ_f and the classifier
609 hyper-parameter β (see Fig. S3 in Supplementary Materials for the stability of the
610 setting). We chose the accuracy as the classification metric and cross-validated the
611 score to provide an estimation of the generalization error with a 10 times repeated
612 stratified 4-fold split.

613 Fig. 12(a) shows that the haemodynamic properties have an improved predic-
614 tion power to discriminate the age compared to the neural activation signals (i.e.
615 temporal components). The latter actually reaches an average accuracy score of
616 0.557, whereas the mean accuracy associated with the HRF shape and haemo-
617 dynamic dilation parameter estimates respectively goes up to 0.741 and 0.743.
618 Also, the distribution of accuracy scores across trials is more concentrated for the
619 HRF dilation parameters compared to the whole HRF shape. This is likely due
620 to the dimension reduction operated to extract this parameter which fluctuates
621 less than the complete profile of the haemodynamic response. This analysis thus
622 demonstrates that our decomposition is able to capture the brain age based on
623 neurovascular information. However, due to the large between-subject variability
624 even within each class of age the neural activation signals do not define a good
625 feature for the brain age prediction. Complementary to that, Fig. 12(b) illustrates
626 the progression of the mean accuracy score with the number of individuals in-
627 volved in the LR model and clearly depicts that a plateau is reached around 459
628 subjects (the total size of the sampled cohort) both for the haemodynamic dila-
629 tion parameter. Also, one can see the rapid progression of the mean accuracy with
630 the number of individuals for the LR model based on haemodynamic properties
631 compared to the one constructed from the neural activation signals.

632 Overall, this experiment has permitted to demonstrate that haemodynamic
633 features are a good biomarker of the normal aging, as already reported in the lit-
634 erature (Grady and Garrett, 2014; West et al., 2019). Moreover, it highlighted that
635 the inter-hemispheric asymmetry in neurovascular coupling brings key information
636 to discriminate middle-age from elderly people.

637 5. Discussion

638 *Separating vascular and neuronal components on fMRI BOLD signals.* Both neu-
639 ral and non neural factors such as glial cell activity, cerebral energy metabolism and
640 the cerebral vasculature contribute independently and synergistically to the fMRI

641 BOLD signal. A mis-estimation of individual and regional HRFs may lead to an in-
642 terpretation that haemodynamic changes as neural variations could have consider-
643 able implications for the interpretability and reliability of findings in fMRI studies.
644 Previous literature supports the notion that HRF variability corrupts fMRI data
645 analysis (Rangaprakash et al., 2017, 2018; Yan et al., 2018). Deconvolution-based
646 approaches such as the one presented here allows discrimination of cerebrovascular
647 components from neural activations and minimizes the confound of HRF variabil-
648 ity in the exploration of brain physiology, functional connectivity and cognitive
649 processes.

650 In this work, we developed a new algorithm that proposed the joint estimation
651 of the HRF and neural activation signal as a semi-blind deconvolution multivariate
652 problem in a paradigm-free setting. Synthetic and real resting-state fMRI data
653 allowed us to demonstrate that this approach is able to faithfully capture the
654 individual’s haemodynamic response function and intrinsic functional networks
655 with low intra-subject variability and relative minimal impact of hyper-parameters
656 on the reliability of HRF estimation. Aims of this paper were also to verify that
657 these methodological developments have practical impacts as related to detection
658 and classification. Further demonstrations using well defined research protocols
659 will be required to refine the tool for use in clinical applications.

660 *Whole brain analysis of haemodynamic properties.* Generally, supervised HRF es-
661 timation methods explain the observed BOLD signal by focusing on a set of brain
662 areas (e.g. visual, auditory and sensorimotor cortices) that are typically involved
663 in a specific activation paradigm (Goutte et al., 2000; Marrelec et al., 2003; Handw-
664 erker et al., 2004; Lindquist and Wager, 2007; Vincent et al., 2010; Pedregosa et al.,
665 2015). To the best of our knowledge, the proposed approach is the first time a
666 particular method is proposed to perform whole brain analysis of haemodynamic
667 properties using resting-state fMRI data in a paradigm-free manner. It should
668 be noted that we used a common brain parcellation – the Havard-Oxford proba-
669 bilistic atlas (Desikan et al., 2006) – across all individuals, to ensure group-level
670 analysis and facilitate between-group comparisons. As such, the major findings we
671 reported on the asymmetries in haemodynamic features between stroke patients
672 and healthy controls, on one hand and the differences related to normal aging on
673 the other hand, are dependent on this atlas and could slightly differ with another
674 parcellation. This question is left for future research.

675 *Interest for analyzing normal aging.* Regional variability in the HRF is partly
676 dictated by the size of surrounding blood vessels (Handwerker et al., 2004; Havlicek
677 and Uludağ, 2020). Vascular aging is known to cause progressive deterioration in
678 the cellular structure of the blood vessel wall, with the development of arteriole

679 tortuosity and reduction in capillary density that undoubtedly impact both resting-
680 state cerebral blood flow and the ability to adjust it during neuronal activity. In
681 that context, as concordant with various imaging studies on aging (Bangen et al.,
682 2009), our algorithm has proven its sensitivity to classify middle-age vs elderly
683 subjects on the basis of HRF parameters. Indeed, using recent rs-fMRI data
684 findings obtained on the large-scale CamCAN task-fMRI data set (Shafto et al.,
685 2014) our study confirmed a clear reductive process of the neurovascular coupling
686 in elderly people in multiple brain areas (occipital, temporal and frontal regions).

687 Beyond the age-related statistical comparison, we also validated on a large
688 cohort (459 subjects) the estimated haemodynamic features as potential predictors
689 of brain age in a supervised classification task. Importantly, we demonstrated that
690 the HRF features and particularly the dilation parameter achieves better accuracy
691 scores (0.74) compared to the neural activation signals (0.56, just above chance
692 performance). This confirms that these neural signatures are much more variable
693 across individuals and even groups. This kind of study was made possible due
694 to the fast convergence of our algorithm and its numerous code optimizations.
695 The underlying `hemolearn` Python package is open source⁶ and available to the
696 neuroimaging community for the sake of reproducible science.

697 *Interest for monitoring patients after a stroke episode.* Our approach, using the
698 asymmetry index, has also proven its utility to individually detect, in patients with
699 a history of stroke, very slow haemodynamic delays in a restricted brain territory
700 probably related to local ischemic tissue consecutive to stroke. This finding is per-
701 fectly consistent with the literature (Altamura et al., 2009) showing that the delay
702 in peak latency that arises as patients advance from the acute to the subacute
703 stroke phase is related to the deterioration of cerebral haemodynamics. Conse-
704 quently, remodeling the fMRI haemodynamic response function in stroke patients
705 may optimize the detection of BOLD signal changes. MRI is of course one of the
706 most powerful diagnostic tools in contemporary clinical medicine. However, in the
707 acute episode of stroke, diffusion-weighted MRI and perfusion imaging (e.g. ASL)
708 remain the reference imaging modalities to perform the diagnosis in a noninvasive
709 way. In the post-acute period, rs-fMRI acquisition equipped with the proposed
710 method would be extremely valuable to measure as a prediction of subsequent re-
711 covery of function as it does not require patient’s engagement in an experimental
712 paradigm.

713 *Perspectives for applications.* Such findings bring new opportunities for the ex-
714 ploration of brain plasticity and pathogenesis in humans. In this way, even older

⁶code available at <https://github.com/hcherkaoui/hemolearn>.

715 adults in relatively good health may have undetected, clinically silent vascular
716 pathology and ischemic brain changes such as silent stroke (D’Esposito et al., 2003)
717 . As this could affect neurovascular coupling, it is of critical interest to assess for
718 cerebrovascular function and to consider vascular risk factors in the pathogenesis
719 or exacerbation of age-related degenerative diseases like Alzheimer. The current
720 challenge for diagnostic imaging methods is to find metrics that capture relevant
721 information or biomarkers. Such works on haemodynamic deconvolution, which
722 are not yet used routinely, might help uncover these biomarkers.

723 This present study constitutes a proof of concept in terms of interest and fea-
724 sibility of the proposed approach. However, many other applications on clinically
725 well-characterized populations could be undertaken to improve and demonstrate
726 its robustness. Aside from stroke and neurological vascular diseases, vascular risk
727 factors are associated with an increased risk of epilepsy and could represent a siz-
728 able proportion of cryptogenic cases of epilepsy (Ogaki et al., 2020). Although
729 abnormal neural activities generating interictal epileptiform discharges provoke
730 haemodynamic changes and BOLD activation, (Béнар et al., 2006; Zijlmans et al.,
731 2007) standard MRI scans fail to visualize epileptic source precisely. Some au-
732 thors have shown that standard HRF in the GLM framework can introduce errors
733 on the extension and localization of activating brain areas. However, modeling
734 haemodynamic response function to permit some flexibility in the HRF shape im-
735 proves sensitivity of fMRI data to delineate epileptogenic area (Storti et al., 2013).
736 This improvement is particularly valuable in epileptic patients with drug-resistant
737 focal seizures, where resection of the epileptogenic brain area remains the best
738 therapeutic outcome. In that context, approaches that employ haemodynamic de-
739 convolution – as presented in this paper – promise a more faithful investigation of
740 the cerebral pathology.

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

755 respect to the unknown hyper-parameters while staying intrinsically connected to
 756 the lower-level problem, namely the multivariate decomposition. Because of the
 757 extra-computation cost required by these approaches, such aspects are beyond
 758 the scope of this paper. Second, the proposed regional analysis is conditioned by
 759 the parcellation atlas (and the value of M). It would be interesting to deepen this
 760 research by testing the reproducibility of the tool with some atlas variations and the
 761 creation of an atlas using subject-specific assessment of the cerebral vasculature.

762 Third, to recover more structured spatial maps, an advanced regularization
 763 model based on TV-elastic net (de Pierrefeu et al., 2017) or structured sparsity (Je-
 764 natton et al., 2012; Baldassarre et al., 2012) could be used in space while keeping
 765 the same algorithmic structure. Recent progress in solving the TV proximity op-
 766 erator (Cherkaoui et al., 2020) for instance could also be directly plugged into
 767 the current algorithm. Fourth, as in standard multivariate data-driven methods,
 768 the inter-subject comparison of spatial maps is currently difficult in the proposed
 769 formulation. In the same spirit as group-ICA (Calhoun et al., 2009), canonical
 770 ICA (Varoquaux et al., 2010) or multi-subject dictionary learning (Varoquaux
 771 et al., 2011), the current within-subject decomposition could be extended to the
 772 group-level to become more stable. One possibility would be to impose the same
 773 spatial maps across all individuals like in Calhoun et al. (2009) while another more
 774 flexible approach would permit spatial variations around a group-level spatial tem-
 775 plate Varoquaux et al. (2011). In this context, the neural activation signals could
 776 remain subject-specific with large fluctuations both in timings and magnitudes.
 777 This kind of extension will be investigated in the near future.

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

784 Sixth, because the proposed HRF model relies solely on a time dilation pa-
 785 rameter, its magnitude is fixed and the fluctuations of the BOLD signal across
 786 the brain are thus captured through the neural activity atoms $(\mathbf{z}_k)_{k=1}^K$ on one
 787 hand and the spatial maps $(\mathbf{u}_k)_{k=1}^K$ on the other hand. However, the norm
 788 of the spatial maps being constrained, the real BOLD signal amplitude is cap-
 789 tured by the neural activation signals. A recent work Tsvetanov et al. (2019)
 790 has shown that the resting-state fluctuation amplitude is crucial to predict brain
 791 age in healthy subjects. One possible enhancement of the current model would
 792 be to add a magnitude parameter to each HRF. In that case, we should fix the
 793 scale ambiguity issue by setting the amplitude of the temporal atoms $(\mathbf{z}_k)_{k=1}^K$.
 794 This modification would significantly increase the computational complexity due

795 to the calculation of the proximal operator associated with the new regulariza-
796 tion term $g_z((z_k)_k) = \lambda \sum_{k=1}^K (\|\nabla z_k\|_1 + I_{\|z_k\|_\infty=\alpha})$. Last, thus far we have used
797 the canonical HRF as the reference shape in \mathbf{v}_{ref} . This setting could be easily
798 updated to perform investigations in specific populations (e.g. newborns) where
799 the true haemodynamic response function is known to deviate from the canonical
800 shape (Arichi et al., 2012).

801 6. Conclusion

802 In this paper, we have presented a semi-blind deconvolution approach to jointly
803 estimate the haemodynamic response function and the neural activity signals
804 across the whole brain. As the proposed methodology is paradigm-free, it en-
805 ables the analysis of resting-state fMRI data in an semi-supervised manner as the
806 regularization parameters (K, λ_f) may be tuned using a trade-off between model
807 accuracy and complexity. Beyond the model validation on synthetic and real fMRI
808 data, we have demonstrated the interest of the proposed approach in two appli-
809 cations in neuroscience. Both aimed at characterizing cerebral haemodynamic
810 delays in specific populations, namely stroke patients and elderly people by con-
811 trasting them with healthy and younger controls. Most importantly, we proposed
812 an haemodynamic asymmetry index to lateralize the stroke episode while con-
813 firming the presence of a prolonged haemodynamic delay in these patients. We
814 also demonstrated that haemodynamic properties are predictable of brain age. Fi-
815 nally, this new framework opens the door to new research avenues for functional
816 connectivity analysis based on the neural input signals instead of the BOLD sig-
817 nal themselves. In contrast to existing techniques (Wu et al., 2013), our approach
818 would be less biased by a constant haemodynamic response shape across the whole
819 brain.

820 Acknowledgment

821 This work was supported by a CEA PhD scholarship, the UK Royal Academy
822 of Engineering under the RF/201718/17128 grant and the SRPe PECRE 1718/15
823 Award. We would like to thank our colleagues Bertrand Thirion, Kamalaker Reddy
824 Dadi and Thomas Bazeille from Inria for fruitful discussions that helped us inves-
825 tigate the proposed data analyses.

826 References

827 S. Ogawa, D. W. Tank, R. Menon, J. M. Ellerman, S. G. Kim, H. Merkle, K. Ugur-
828 bil, Intrinsic signal changes accompanying sensory stimulation: functional brain
829 mapping with magnetic resonance imaging., in: Proceedings of the National
830 Academy of Sciences, volume 89, pp. 5951–5955.

- 831 P. A. Bandettini, A. Jesmanowicz, E. C. Wong, J. S. Hyde, Processing strategies
832 for time-course data sets in functional MRI of the human brain., *Magnetic*
833 *resonance in medicine* 30 (1993) 161–173.
- 834 G. M. Boynton, S. A. Engel, G. H. Glover, D. J. Heeger, Linear systems analysis of
835 functional magnetic resonance imaging in human V1., *Journal of Neuroscience*
836 16 (1996) 4207–4221.
- 837 C.-T. Do, Z.-M. Manjaly, J. Heinzle, D. Schöbi, L. Kasper, K. P. Pruessmann,
838 K. E. Stephan, S. Frässle, Hemodynamic modeling of aspirin effects on bold
839 responses at 7t., *medRxiv* (2020).
- 840 K. L. West, M. D. Zuppichini, M. P. Turner, D. K. Sivakolundu, Y. Zhao, D. Ab-
841 delkarim, J. S. Spence, B. Rypma, BOLD hemodynamic response function
842 changes significantly with healthy aging., *Neuroimage* 188 (2019) 198–207.
- 843 D. Asemani, H. Morsheddest, M. A. Shalchy, Effects of ageing and alzheimer
844 disease on haemodynamic response function: a challenge for event-related fMRI.,
845 *Healthcare Technology Letters* 4 (2017) 109–114.
- 846 K. J. Friston, P. Fletcher, O. Josephs, A. Holmes, M. D. Rugg, R. Turner, Event-
847 related fMRI: characterizing differential responses., *Neuroimage* 7 (1998) 30–40.
- 848 P. Ciuciu, J.-B. Poline, G. Marrelec, J. Idier, C. Pallier, H. Benali, Unsupervised
849 robust nonparametric estimation of the hemodynamic response function for any
850 fMRI experiment., *IEEE transactions on Medical Imaging* 22 (2003) 35–51.
- 851 M. A. Lindquist, T. D. Wager, Validity and power in hemodynamic response
852 modeling: a comparison study and a new approach., *Human brain mapping* 28
853 (2007) 764–784.
- 854 F. Pedregosa, M. Eickenberg, P. Ciuciu, B. Thirion, A. Gramfort, Data-driven
855 HRF estimation for encoding and decoding models., *NeuroImage* (2015) 209–
856 220.
- 857 C. Goutte, F. A. Nielsen, L. K. Hansen, Modeling the haemodynamic response
858 in fMRI using smooth FIR filters., *IEEE transactions on Medical Imaging* 19
859 (2000) 1188–1201.
- 860 T. Vincent, T. Rissern, P. Ciuciu, Spatially adaptive mixture modeling for analysis
861 of fMRI time series, *IEEE transactions on Medical Imaging* 29 (2010) 59–74.
- 862 L. Chaari, L. Forbes, T. Vincent, P. Ciuciu, Hemodynamic-informed parcellation
863 of fMRI data in a joint detection estimation framework, in: *proceedings of*

- 864 International Conference on Medical Image Computing and Computer-Assisted
865 Intervention, volume 15, pp. 180–188.
- 866 G. H. Glover, Deconvolution of impulse response in event-related bold fMRI.,
867 Neuroimage 9 (1999) 416–429.
- 868 D. R. Gitelman, W. D. Penny, J. Ashburner, K. J. Friston, Modeling regional
869 and psychophysiologic interactions in fMRI: the importance of hemodynamic
870 deconvolution., Neuroimage 19 (2003) 200–207.
- 871 L. Hernandez-Garcia, M. O. Ulfarsson, Neuronal event detection in fMRI time
872 series using iterative deconvolution techniques., Magnetic resonance imaging 29
873 (2011) 353–364.
- 874 C. Caballero-Gaudes, F. I. Karahanoglu, F. Lazeyras, D. Van De Ville, Structured
875 sparse deconvolution for paradigm free mapping of functional MRI data., in:
876 International Symposium on Biomedical Imaging, volume 9, pp. 322–325.
- 877 I. Khalidov, J. Fadili, F. Lazeyras, D. Van De Ville, M. Unser, Activelets: Wavelets
878 for sparse representation of hemodynamic responses., Signal processing 91 (2011)
879 2810–2821.
- 880 F. I. Karahanoglu, C. Caballero-gaudes, F. Lazeyras, D. Van De Ville, Total activa-
881 tion: fMRI deconvolution through spatio-temporal regularization., NeuroImage
882 73 (2013) 121–134.
- 883 G. R. Wu, W. Liao, S. Stramaglia, J. R. Ding, H. Chen, D. Marinazzo, A blind
884 deconvolution approach to recover effective connectivity brain networks from
885 resting state fMRI data., in: Medical Image Analysis, volume 17, pp. 365–374.
- 886 H. Cherkaoui, T. Moreau, A. Halimi, P. Ciuciu, Sparsity-based Semi-Blind De-
887 convolution of Neural Activation Signal in fMRI., in: IEEE International Con-
888 ference on Acoustics, Speech and Signal Processing (ICASSP), pp. 1323–1327.
- 889 Y. Farouj, F. I. Karahanoglu, D. Van de Ville, Bold signal deconvolution under un-
890 certain haemodynamics: A semi-blind approach., in: International Symposium
891 on Biomedical Imaging, pp. 1792–1796.
- 892 H. Cherkaoui, T. Moreau, A. Halimi, P. Ciuciu, fMRI BOLD signal decompo-
893 sition using a multivariate low-rank model., in: European Signal Processing
894 Conference (EUSIPCO), pp. 1–5.
- 895 R. Grosse, R. Raina, H. Kwong, Y. N. Andrew, Shift-invariant sparse coding for
896 audio classification., in: proceedings of Uncertainty in Artificial Intelligence,
897 volume 23, pp. 149–158.

- 898 T. Dupré La Tour, T. Moreau, M. Jas, A. Gramfort, Multivariate convolutional
899 sparse coding for electromagnetic brain signals., in: *Advances in Neural Infor-*
900 *mation Processing System (NeurIPS)*, pp. 3292–3302.
- 901 M. W. Woolrich, B. D. Ripley, M. Brady, S. M. Smith, Temporal autocorrelation
902 in univariate linear modeling of fMRI data., *Neuroimage* 14 (2001) 1370–1386.
- 903 W. Penny, S. Kiebel, K. Friston, Variational bayesian inference for fMRI time
904 series., *NeuroImage* 19 (2003) 727–741.
- 905 S. Makni, J. Idier, T. Vincent, B. Thirion, G. Dehaene-Lambertz, P. Ciuciu, A fully
906 bayesian approach to the parcel-based detection-estimation of brain activity in
907 fMRI., *Neuroimage* 41 (2008) 941–969.
- 908 A. M. Dale, Optimal experimental design for event-related fMRI, *Human Brain*
909 *Mapping* 2 (1999) 109–114.
- 910 R. Casanova, S. Ryali, J. Serences, L. Yang, R. Kraft, P. J. Laurienti, J. A. Mald-
911 jian, The impact of temporal regularization on estimates of the bold hemo-
912 dynamic response function: a comparative analysis., *NeuroImage* 40 (2008)
913 1606–1618.
- 914 C. M. Zhang, Y. Jiang, T. Yu, A comparative study of one-level and two-level
915 semiparametric estimation of hemodynamic response function for fMRI data.,
916 *Statistics in medicine* 26 (2007) 3845–3861.
- 917 V. A. Vakorin, R. Borowsky, G. E. Sarty, Characterizing the functional MRI
918 response using tikhonov regularization., *Statistics in medicine* 26 (2007) 3830–
919 3844.
- 920 M. W. Woolrich, T. E. J. Behrens, S. M. Smith, Constrained linear basis sets for
921 hrf modelling using variational bayes., *NeuroImage* 21 (2004) 1748–1761.
- 922 M. A. Lindquist, J. Meng Loh, L. Y. Atlas, T. D. Wager, Modeling the hemody-
923 namic response function in fMRI: efficiency, bias and mis-modeling., *NeuroImage*
924 45 (2009) S187–S198.
- 925 D. A. Handwerker, J. M. Ollinger, M. D’Esposito, Variation of bold hemody-
926 namic responses across subjects and brain regions and their effects on statistical
927 analyses., *Neuroimage* 21 (2004) 1639–1651.
- 928 S. Badillo, T. Vincent, P. Ciuciu, Group-level impacts of within-and between-
929 subject hemodynamic variability in fMRI., *Neuroimage* 82 (2013) 433–448.

- 930 G. Varoquaux, R. C. Craddock, Learning and comparing functional connectomes
931 across subjects., *NeuroImage* 80 (2013) 405–415.
- 932 R. S. Desikan, F. Segonne, B. Fischl, B. T. Quinn, B. C. Dickerson, D. Blacker,
933 R. L. Buckner, A. M. Dale, R. P. Maguire, B. T. Hyman, M. S. Albert, R. J.
934 Killiany, An automated labeling system for subdividing the human cerebral
935 cortex on MRI scans into gyral based regions of interest., *Neuroimage* 31 (2006)
936 968–980.
- 937 B. O’Donoghue, E. Candes, Adaptive restart for accelerated gradient schemes.,
938 *Foundations of Computational Mathematics* 15 (2015) 715–732.
- 939 J. Mairal, F. Bach, J. Ponce, G. Sapiro, Online dictionary learning for sparse
940 coding., in: *International Conference on Machine Learning*, volume 26, pp.
941 689–696.
- 942 P. L. Combettes, J. C. Pesquet, Proximal splitting methods in signal processing.,
943 *Fixed-Point Algorithms for Inverse Problems in Science and Engineering* (2009).
- 944 A. Chambolle, An Algorithm for Total Variation Minimization and Applications.,
945 *Journal of Mathematical Imaging and Vision* 20 (2004) 89–97.
- 946 A. Barbero, S. Sra, Modular proximal optimization for multidimensional total-
947 variation regularization., in: *Journal of Machine Learning Research*, volume 19,
948 pp. 1–89.
- 949 L. Condat, Fast projection onto the simplex and the l_1 ball., *Mathematical*
950 *Programming Series A* 158 (2016) 575–585.
- 951 F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel,
952 M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos,
953 D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine
954 learning in Python., *Journal of Machine Learning Research* 12 (2011) 2825–2830.
- 955 M. Jenkinson, P. Bannister, M. Brady, S. Smith, Improved optimization for the
956 robust and accurate linear registration and motion correction of brain images.,
957 *NeuroImage* 17 (2002) 825–41.
- 958 V. Menon, Large-scale functional brain organization., *Brain mapping: An ency-*
959 *clopedic reference* 2 (2015) 449–459.
- 960 D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. J. Behrens, E. Yacoub, K. Ugur-
961 bil, The wu-minn Human Connectome Project: An overview., *NeuroImage* 80
962 (2013) 62–79.

- 963 M. Glasser, S. Sotiropoulos, J. Wilson, T. Coalson, B. Fischl, J. Andersson, J. Xu,
964 S. Jbabdi, M. Webster, J. Polimeni, V. DC, M. Jenkinson, The minimal pre-
965 processing pipelines for the human connectome project, *NeuroImage* 80 (2013)
966 105.
- 967 P. H. England, Briefing document: First incidence of stroke estimates for england
968 2007 to 2016., Public Health England publication (2018).
- 969 S. N. Min, S. J. Park, D. J. Kim, M. Subramaniam, K. S. Lee, Development of
970 an algorithm for stroke prediction: A national health insurance database study
971 in korea., *European Neurology* 79 (2018) 214–220.
- 972 M. Raemaekers, W. Schellekens, N. Petridou, N. F. Ramsey, Knowing left from
973 right: asymmetric functional connectivity during resting state., *Brain Structure
974 and Function* 223 (2018) 1909–1922.
- 975 B. M. Ances, C. L. Liang, O. Leontiev, J. E. Perthen, A. S. Fleisher, A. E. Lansing,
976 R. B. Buxton, Effects of aging on cerebral blood flow, oxygen metabolism, and
977 blood oxygenation level dependent responses to visual information stimulation.,
978 *Human Brain Mapping* 30 (2009) 1120–1132.
- 979 Y. Li, W. J. Choi, W. Wei, S. Song, Q. Zhang, J. Liu, , R. K. Wang, Aging-
980 associated changes in cerebral vasculature and blood flow as determined by
981 quantitative optical coherence tomography angiography., *Neurobiol Aging* 70
982 (2018) 148–159.
- 983 D. A. Engemann, O. Kozynets, D. Sabbagh, G. Lemaître, G. Varoquaux, F. Liem,
984 A. Gramfort, Combining magnetoencephalography with magnetic resonance
985 imaging enhances learning of surrogate-biomarkers., *Elife* 9 (2020).
- 986 C. L. Grady, D. D. Garrett, Understanding variability in the bold signal and why
987 it matters for aging., *Brain Imaging Behaviour* 8 (2014) 274–283.
- 988 D. Rangaprakash, M. N. Dretsch, W. Yan, J. S. Katz, T. S. Denney Jr, G. Deshpande,
989 Hemodynamic response function parameters obtained from resting-state
990 functional MRI data in soldiers with trauma., *Data in brief* 14 (2017) 558–562.
- 991 D. Rangaprakash, G.-R. Wu, D. Marinazzo, X. Hu, G. Deshpande, Hemody-
992 namic response function (HRF) variability confounds resting-state fMRI func-
993 tional connectivity., *Magnetic resonance in medicine* 80 (2018) 1697–1713.
- 994 W. Yan, D. Rangaprakash, G. Deshpande, Aberrant hemodynamic responses in
995 autism: Implications for resting state fMRI functional connectivity studies.,
996 *NeuroImage: Clinical* 19 (2018) 320–330.

- 997 G. Marrelec, H. Benali, P. Ciuciu, M. Péligrini-Issac, J.-B. Poline, Robust
998 Bayesian estimation of the hemodynamic response function in event-related
999 BOLD MRI using basic physiological information., *hbm* 19 (2003) 1–17.
- 1000 M. Havlicek, K. Uludağ, A dynamical model of the laminar BOLD response.,
1001 *NeuroImage* 204 (2020) 116–209.
- 1002 K. J. Bangen, K. Restom, T. T. Liu, A. J. Jak, C. E. Wierenga, D. P. Salmon,
1003 M. W. Bondi, Differential age effects on cerebral blood flow and BOLD response
1004 to encoding: associations with cognition and stroke risk., *Neurobiology of aging*
1005 30 (2009) 1276–1287.
- 1006 M. A. Shafto, L. K. Tyler, M. Dixon, J. R. Taylor, J. B. Rowe, R. Cusack, A. J.
1007 Calder, W. D. Marslen-Wilson, J. Duncan, T. Dalgleish, et al., The cambridge
1008 centre for ageing and neuroscience (Cam-CAN) study protocol: a cross-sectional,
1009 lifespan, multidisciplinary examination of healthy cognitive ageing., *BMC neu-*
1010 *rology* 14 (2014) 204.
- 1011 C. Altamura, M. Reinhard, M.-S. Vry, C. P. Kaller, F. Hamzei, F. Vernieri, P. M.
1012 Rossini, A. Hetzel, C. Weiller, D. Saur, The longitudinal changes of BOLD
1013 response and cerebral hemodynamics from acute to subacute stroke. A fMRI
1014 and TCD study., *BMC neuroscience* 10 (2009) 151.
- 1015 M. D’Esposito, L. Y. Deouell, A. Gazzaley, Alterations in the BOLD fMRI sig-
1016 nal with ageing and disease: a challenge for neuroimaging., *Nature Reviews*
1017 *Neuroscience* 4 (2003) 863–872.
- 1018 A. Ogaki, Y. Ikegaya, R. Koyama, Vascular abnormalities and the role of vascular
1019 endothelial growth factor in the epileptic brain, *Frontiers in Pharmacology* 11
1020 (2020).
- 1021 C.-G. Bénar, C. Grova, E. Kobayashi, A. P. Bagshaw, Y. Aghakhani, F. Dubeau,
1022 J. Gotman, EEG–fMRI of epileptic spikes: concordance with EEG source local-
1023 ization and intracranial EEG., *Neuroimage* 30 (2006) 1161–1170.
- 1024 M. Zijlmans, G. Huiskamp, M. Hersevoort, J.-H. Seppenwoolde, A. C. van Huf-
1025 felen, F. S. S. Leijten, EEG–fMRI in the preoperative work-up for epilepsy
1026 surgery., *Brain* 130 (2007) 2343–2353.
- 1027 S. F. Storti, E. Formaggio, A. Bertoldo, P. Manganotti, A. Fiaschi, G. M. Toffolo,
1028 Modelling hemodynamic response function in epilepsy., *Clinical Neurophysiol-*
1029 *ogy* 124 (2013) 2108–2118.

- 1030 A. A. Neath, J. E. Cavanaugh, The Bayesian information criterion: background,
1031 derivation, and applications., Wiley Interdisciplinary Reviews: Computational
1032 Statistics 4 (2012) 199–203.
- 1033 S. Watanabe, A widely applicable Bayesian information criterion., Journal of
1034 Machine Learning Research 14 (2013) 867–897.
- 1035 K. P. Bennett, J. Hu, X. Ji, G. Kunapuli, J.-S. Pang, Model selection via bilevel
1036 optimization., in: The 2006 IEEE International Joint Conference on Neural
1037 Network Proceedings, pp. 1922–1929.
- 1038 A. de Pierrefeu, T. Löfstedt, F. Hadj-Selem, M. Dubois, R. Jardri, T. Fovet,
1039 P. Ciuciu, V. Frouin, E. Duchesnay, Structured sparse principal components
1040 analysis with the TV-elastic net penalty., IEEE transactions on medical imaging
1041 37 (2017) 396–407.
- 1042 R. Jenatton, A. Gramfort, V. Michel, G. Obozinski, E. Eger, F. Bach, B. Thirion,
1043 Multiscale mining of fMRI data with hierarchical structured sparsity., SIAM
1044 Journal on Imaging Sciences 5 (2012) 835–856.
- 1045 L. Baldassarre, J. Mourao-Miranda, M. Pontil, Structured sparsity models for
1046 brain decoding from fMRI data., in: 2012 Second International Workshop on
1047 Pattern Recognition in NeuroImaging, pp. 5–8.
- 1048 H. Cherkaoui, J. Sulam, T. Moreau, Learning to solve tv regularised problems with
1049 unrolled algorithms., in: 34th Conference and Workshop on Neural Information
1050 Processing Systems (NeurIPS), pp. 1–21.
- 1051 V. D. Calhoun, J. Liu, T. Adalı, A review of group ICA for fMRI data and ICA
1052 for joint inference of imaging, genetic, and erp data., Neuroimage 45 (2009)
1053 S163–S172.
- 1054 G. Varoquaux, S. Sadaghiani, P. Pinel, A. Kleinschmidt, J. B. Poline, B. Thirion,
1055 A group model for stable multi-subject ICA on fMRI datasets., NeuroImage 51
1056 (2010) 288–299.
- 1057 G. Varoquaux, A. Gramfort, F. Pedregosa, V. Michel, B. Thirion, Multi-subject
1058 dictionary learning to segment an atlas of brain spontaneous activity., in: Bi-
1059 ennial International Conference on information processing in medical imaging,
1060 pp. 562–573.
- 1061 M. Milham, D. Fair, M. Mennes, S. Mostofsky, The adhd-200 consortium: a model
1062 to advance the translational potential of neuroimaging in clinical neuroscience.,
1063 Frontiers in Systems Neuroscience 6 (2012) 62.

- 1064 D. A. Feinberg, K. Setsompop, Ultra-fast MRI of the human brain with simulta-
1065 neous multi-slice imaging., *Journal of magnetic resonance* 229 (2013) 90–100.
- 1066 J. Hesamoddin, S. Holdsworth, T. Christen, H. Wu, K. Zhu, A. B. Kerr, M. J.
1067 Middione, R. F. Dougherty, M. Moseley, G. Zaharchuk, Advantages of short
1068 repetition time resting-state functional MRI enabled by simultaneous multi-slice
1069 imaging., *Journal of Neuroscience Methods* 311 (2019) 122–132.
- 1070 K. Tsvetanov, R. N. A. Henson, P. S. Jones, H. J. Mutsaerts, D. Fuhrmann, L. K.
1071 Tyler, , J. B. Rowe, The effects of age on resting-state bold signal variability is
1072 explained by cardiovascular and neurovascular factors., preprint bioRxiv (2019).
- 1073 T. Arichi, G. Fagiolo, M. Varela, A. Melendez-Calderon, A. Allievi, N. Merchant,
1074 N. Tumor, S. J. Counsell, E. Burdet, C. F. Beckmann, et al., Development
1075 of BOLD signal hemodynamic responses in the human brain., *Neuroimage* 63
1076 (2012) 663–673.
- 1077 K. J. Friston, O. Josephs, G. Rees, R. Turner, Nonlinear event-related responses
1078 in fMRI., *Magnetic Resonance in Medicine* 39 (1998) 41–52.
- 1079 R. Tibshirani, Regression Shrinkage and Selection via the Lasso., *Journal of the*
1080 *Royal Statistical Society: Series B (statistical methodology)* 58 (1996) 267–288.

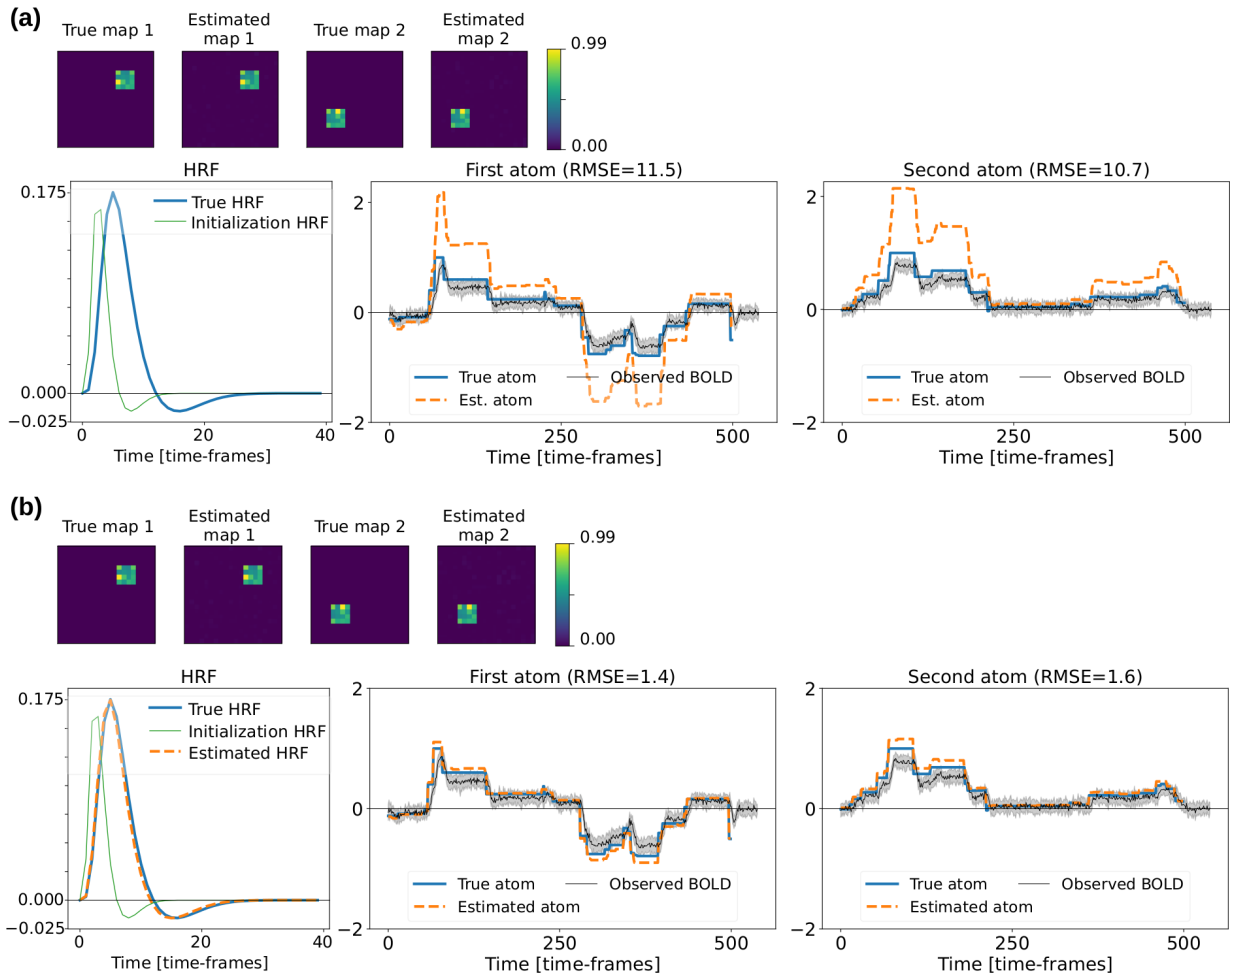


Figure 4: **Top (a): deconvolution with *fixed* HRF.** The top row shows the two true spatial maps and their accurate estimates. The bottom row shows on the left the true HRF shape and the filter used for deconvolution. In the middle and on the right hand side, the true and estimated neural temporal 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.

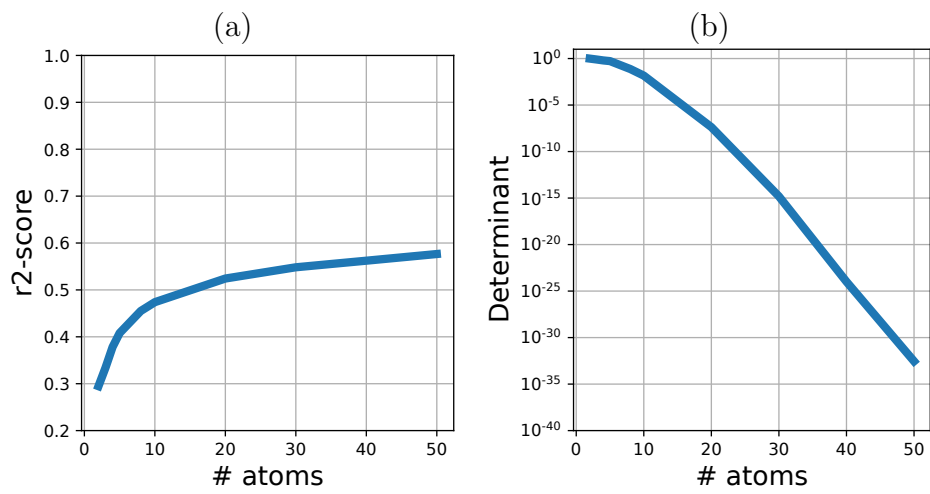


Figure 5: **Model selection: compromise between model complexity and accuracy.** (a): Evolution of the R^2 score as a function of the number of spatio-temporal atoms K in model (5) ranging from 2 to 50. (b): Evolution of the determinant of the correlation matrix Σ_K between neural activation signals as a function of K ranging in the same interval as mentioned earlier.




Figure 6: **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.




Figure 7: **Single subject results from rs-fMRI semi-blind deconvolution analysis.** **Top Row (a):** Neural activation signal \hat{z}_{10} (left) and corresponding spatial map (axial views) \hat{u}_{10} (right), mostly involving activated voxels in the visual cortex. **Second row (b):** Neural activation signal \hat{z}_2 (left) and corresponding spatial map (axial views) \hat{u}_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.

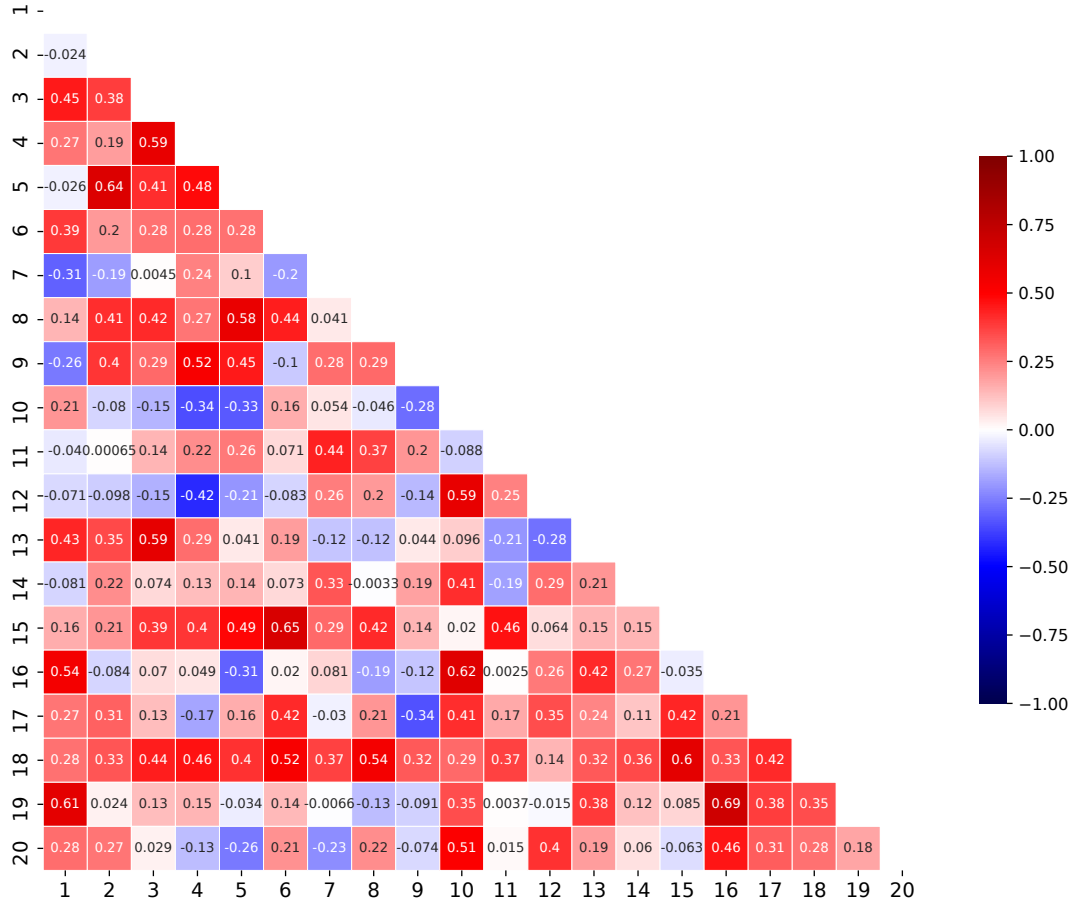


Figure 8: **Correlation matrix between neural activation signals for $K = 20$.** Triangular inferior view of the semi-definite positive matrix Σ_K for $K = 20$. All entries vary between -1 and +1 as they reflect correlation coefficients.

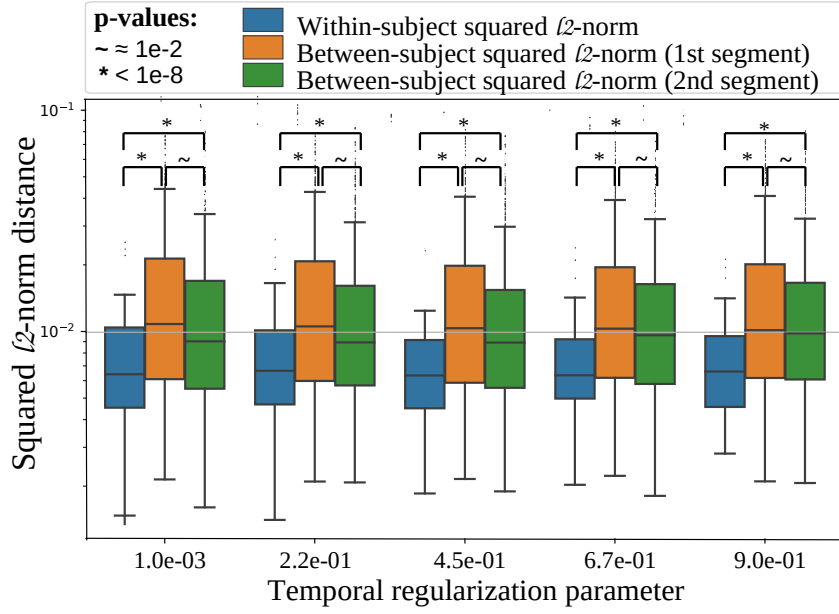


Figure 9: **Within-subject vs between-subject analysis of the haemodynamic variability.** The box plots show respectively in blue, orange and green the distribution of $WS(\widehat{\delta}_{T_1}^{s_i}, \widehat{\delta}_{T_2}^{s_i})$ for all subjects ($i = 1, \dots, 100$), $BS(\widehat{\delta}_{T_1}^{s_i}, \widehat{\delta}_{T_1}^{s_j})$ and $BS(\widehat{\delta}_{T_2}^{s_i}, \widehat{\delta}_{T_2}^{s_j})$ 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 $*$.




Figure 10: **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 $s = \text{HC}, \text{St}$. 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 ($\Delta_{\text{TTP}}^{\text{SP}} = 2.25 \text{ s}$) for the stroke patient than for the Healthy Controls ($\Delta_{\text{TTP}}^{\text{HC}} = 1.25 \text{ s}$). **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.




Figure 11: **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.

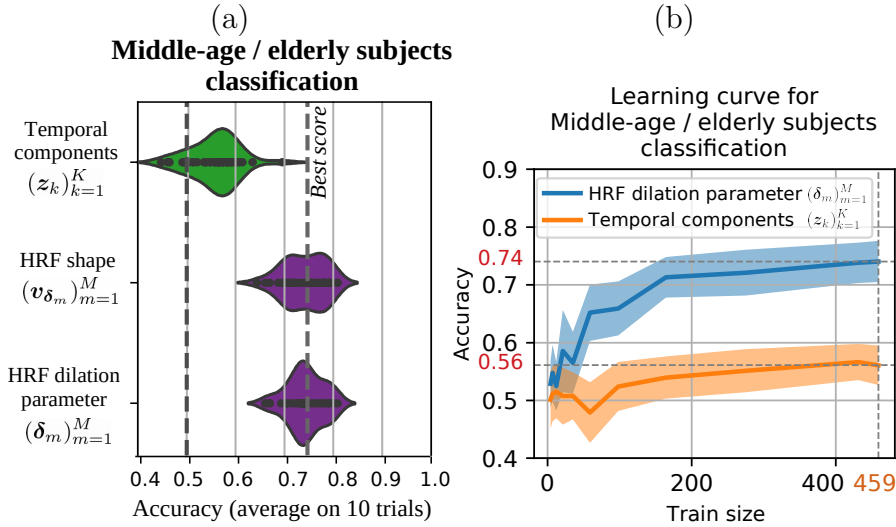


Figure 12: **Accuracy score for classifying middle-age vs. elderly subjects (459 individuals sampled from the UK Biobank database).** **(a):** The prediction was performed by pulling individual features either based on (i) the estimated neural activation signals $(\hat{z}_k)_{k=1}^{20}$, (ii) HRF shape estimates $(\hat{v}_{\delta_m})_{m=1}^{96}$ 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).

1081 **7. Supplementary material**

1082 *7.1. Gradient derivation w.r.t the HRF dilation parameter*

1083 In this subsection, we detail the gradient derivation of our cost-function from
 1084 Eq. (6) – denoted J hereafter – w.r.t $\boldsymbol{\delta}$. Let us define $\tilde{\mathbf{A}} = (\tilde{\mathbf{a}}_j)_{j=1}^P \in \mathbb{R}^{P \times \tilde{T}}$ such as
 1085 $\tilde{\mathbf{A}} = \sum_{k=1}^K \mathbf{u}_k^\top \mathbf{z}_k$. Moreover, we introduce θ_m the set of indices of voxels belonging
 1086 to the m^{th} region of the brain parcellation.

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

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

$$\nabla_{\boldsymbol{\delta}} J(\boldsymbol{\delta}) = \left[\frac{\partial J(\boldsymbol{\delta})}{\partial \delta_1}, \dots, \frac{\partial J(\boldsymbol{\delta})}{\partial \delta_M} \right]^\top \in \mathbb{R}^M \quad (8)$$

To this end, we proceed componentwise:

$$\begin{aligned} \frac{\partial J(\boldsymbol{\delta})}{\partial \delta_m} &= \frac{1}{2} \sum_{j \in \theta_m} \frac{\partial \|\mathbf{v}_{\delta_m} * \mathbf{a}_j - \mathbf{y}_j\|_2^2}{\partial \delta_m} \\ &= \sum_{j \in \theta_m} \left(\frac{\partial (\mathbf{v}_{\delta_m} * \mathbf{a}_j)}{\partial \delta_m} \right)^\top (\mathbf{v}_{\delta_m} * \mathbf{a}_j - \mathbf{y}_j) \\ &= \left(\frac{\partial \mathbf{v}_{\delta_m}}{\partial \delta_m} \right)^\top \left(\sum_{j \in \theta_m} \mathbf{a}_j^\top * (\mathbf{v}_{\delta_m} * \mathbf{a}_j - \mathbf{y}_j) \right) \\ &= \left(\frac{\partial \mathbf{v}_{\delta_m}}{\partial \delta_m} \right)^\top \underbrace{\left(\mathbf{v}_{\delta_m} * \sum_{j \in \theta_m} \mathbf{a}_j^\top * \mathbf{a}_j - \sum_{j \in \theta_m} \mathbf{a}_j^\top * \mathbf{y}_j \right)}_{\nabla_{\mathbf{v}_{\delta_m}} J}. \end{aligned} \quad (9)$$

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 v' the first-order derivative of function v . Now, taking the definition of $v(\cdot)$ from (Friston et al., 1998), 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., 1998). A straightforward computation gives us for $t \in \mathbb{R}^+$:

$$v'(t) = \left(\frac{a-1}{t} - 1 \right) \frac{t^{a-1} e^{-\delta t}}{\Gamma(a)} - c \left(\frac{b-1}{t} - 1 \right) \frac{t^{b-1} e^{-\delta t}}{\Gamma(b)} \quad (10)$$

$$\frac{\partial}{\partial \delta} v(\delta t) = t v'(\delta t) = \left(\frac{a-1}{\delta} - t \right) \frac{(\delta t)^{a-1} e^{-\delta t}}{\Gamma(a)} - c \left(\frac{b-1}{\delta} - t \right) \frac{(\delta t)^{b-1} e^{-\delta t}}{\Gamma(b)} \quad (11)$$

1087 The value of $\frac{\partial v_{\delta_m}}{\partial \delta_m}$ can thus be computed by taking the discrete time points
 1088 corresponding to the sampling rate of the BOLD signal and the length of the
 1089 considered HRF. By replacing its value in the computation of $\frac{\partial J(\boldsymbol{\delta})}{\partial \delta_m}$ from Eq. (9),
 1090 we obtain a closed form expression for the gradient of J w.r.t the HRF dilation
 1091 parameter δ i.e. $\nabla_{\boldsymbol{\delta}} J(\boldsymbol{\delta})$.

1092 7.2. Haemodynamic parameter estimate stability across various levels of temporal 1093 regularization

1094 A well known limitation of regularization methods based on the l_1 -norm such
 1095 as TV is that large coefficients – here in $(\mathbf{z}_{\mathbf{k}})_{k=1}^K$ – are shrunk toward zero (Tib-
 1096 shirani, 1996). Thus, the magnitude of the estimated neural activation signals
 1097 $(\mathbf{z}_{\mathbf{k}})_{k=1}^K$ is biased. Moreover, this bias is tightly linked to the choice of the reg-
 1098 ularization parameter λ_f . Indeed, the larger this parameter is, the more $(\mathbf{z}_{\mathbf{k}})_{k=1}^K$
 1099 are shrunk toward zero. To quantify this effect on our model, we applied the
 1100 spatio-temporal decomposition with $M = 96$ ROI and $K = 20$ and various tem-
 1101 poral regularization level λ_f on the cohort of $S = 459$ subjects sampled from the
 1102 UK Biobank resting-stage fMRI dataset used in Section 4.2. Fig. S1 reports the
 1103 grand average of the dilatation parameters

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

1104 and its variance with respect to the regularization parameter λ_f . We observed
 1105 that the HRF dilation parameters decrease with the temporal regularization level
 1106 – and thus the corresponding time-to-peaks increase with λ_f . This results from the
 1107 fact that the model with large regularization parameters only accounts for sharp
 1108 transition in the BOLD signal mean value, which are well approximated with fast
 1109 HRF.

1110 This previous result entails that the haemodynamic delay estimated by our
 1111 model (6) may be biased. However, because there is a single temporal regular-
 1112 ization parameter, we expect that this bias impacts the whole brain uniformly.
 1113 To assess this shared effect on the estimated parameter, we observe the relative
 1114 variations of δ_m . Fig. S2 displays the value of $\bar{\delta}_m(\lambda_f)$ relative to $\bar{\delta}(\lambda_f)$ for $S = 459$

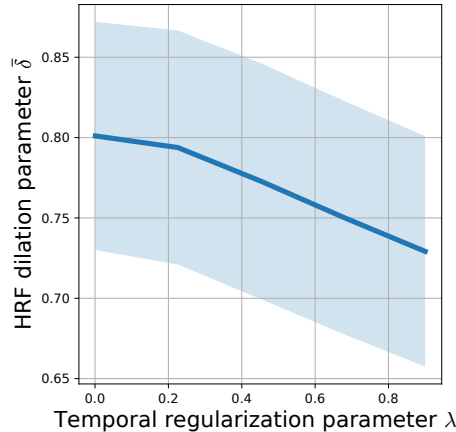


Figure S1: **Evolution of the group-level grand average haemodynamic dilation parameter $\bar{\delta}$ as a function of the temporal regularization level $\lambda_f \in [0, 1]$.** The solid blue line reflects the decreasing evolution of $\bar{\delta}$ when $\lambda_f \rightarrow 1$, where the value of $\bar{\delta}$ 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.

1115 subjects with three temporal regularization values $\lambda_f \in \{0.001, 0.5, 0.9\}$ on the
 1116 MNI template. Precisely, for each regularization parameter and for each ROI m ,
 1117 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 pa-
 1118 rameter across subjects. While the magnitudes change when the regularization
 1119 changes, as seen in Fig. S1, the spatial structure of dilatation parameters in the
 1120 brain is globally preserved. Indeed, the normalized maps look very similar for any
 1121 choice of regularization parameter, showing that the relative variation between
 1122 each area of the brain are preserve while changing the hyper-parameter. Thus,
 1123 we can state that the haemodynamic response from the middle temporal gyrus is
 1124 faster than the response from the frontal orbital cortex, as described in Fig. 7. This
 1125 means that while the numerical value of the time-to-peak for a given area may not
 1126 be reflect the actual haemodynamic delay in the brain, the estimated coefficients
 1127 reflect the spatial variations of the delay between the different areas of the brain.
 1128 Moreover, these variations are stable with the choice of temporal regularization.
 1129 Hence, choosing a potentially suboptimal value for λ_f is of limited impact when
 1130 the primary interest is investigating abnormalities in the neuro-vascular coupling.

1131 Finally, we assess the impact of the choice of λ_f on the prediction results from
 1132 Section 4.2. Fig. S3 reports the accuracy score for the logistic regression relatively
 1133 to the choice of regularization parameter β for the classification model and the
 1134 temporal regularization parameter λ_f for our deconvolution model. The accuracy
 1135 is almost not impacted by the choice of parameter λ_f , for any value of β . This
 1136 observation confirms that the choice of λ_f is not critical when studying the relative




Figure S2: **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 .

1137 spatial structure of the haemodynamic delay and that our model can be used in
 1138 practical cases to evaluate abnormalities in the haemodynamic response.

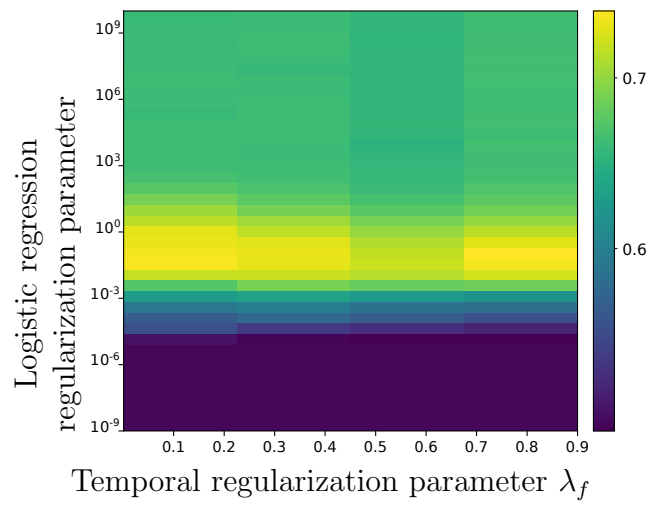


Figure S3: Evolution of the accuracy score *w.r.t* the logistic regression regularization parameter β and the temporal regularization parameter λ_f . The accuracy score is not impacted by the hyper-parameter λ_f , as moving this parameter mainly impact the magnitude of the estimated delays and not its spatial structure.