



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. *NeuroImage*, 2021, 10.1016/j.neuroimage.2021.118418 . hal-03005584v2

HAL Id: hal-03005584

<https://hal.science/hal-03005584v2>

Submitted on 2 Sep 2021

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

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

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}

^a*CEA/NeuroSpin, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France.*

^b*Université Paris-Saclay, CEA, CNRS, Inserm, BioMaps, Orsay, 91401, France.*

^c*Parietal Team, Université Paris-Saclay, CEA, Inria, Gif-sur-Yvette, 91190, France.*

^d*School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh UK.*

^e*Corresponding author (philippe.ciuciu@cea.fr). Postal Address: CEA Saclay, NeuroSpin,
Bat 145 - PC 156, F-91191 Gif-sur-Yvette cedex, France.*

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. A fast alternating minimization algorithm, along with its efficient implementation, 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 (Asemi 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; Degras and Lindquist, 2014; Rosa et al., 2015; Pedregosa et al., 2015; Eickenberg et al., 2017). 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

34 where no EP is available to serve as surrogate for neural activity. A potential
35 solution would consist in uncovering neural activity by thresholding large values
36 in the BOLD signal and then constructing a sparse binary sequence (Wu et al.,
37 2013) as an artificial EP to be subsequently used in supervised HRF estimation
38 techniques. However, due to this two-step procedure, any error in neural activity
39 detection tends to bias the recovery of HRF shape, especially in regard to the
40 haemodynamic delay.

41 On the other hand, a long-standing literature on fMRI deconvolution methods
42 has emerged since the late 90s to uncover the underlying *activity-inducing sig-*
43 *nal* at the fMRI timescale of seconds, see (Glover, 1999) for Wiener filtering and
44 smooth estimation of activity-inducing signals and (Gitelman et al., 2003) for its
45 generalization to study psychophysiologic interactions at the neuronal level. Im-
46 portantly, within a temporal fMRI deconvolution framework, the neural activity
47 signal was estimated by involving the ℓ_2 -norm regularization (Caballero-Gaudes
48 et al., 2011). Interestingly, most recent works have relied on sparse regulariza-
49 tion to recover a limited number of spike-like activations directly in the time do-
50 main (Hernandez-Garcia and Ulfarsson, 2011; Caballero-Gaudes et al., 2011, 2012,
51 2013, 2018) or imposing sparsity in the activelet-domain, which is a wavelet ba-
52 sis that is tailored to the haemodynamic properties (Khalidov et al., 2007, 2011).
53 Next, a foundational work, called Total Activation (Karahanoğlu et al., 2013),
54 has proposed a spatio-temporal model of the underlying activity-inducing signal
55 including both temporal and spatial sparsity-based regularization. In doing so,
56 the recovered neural activity profiles are used to define functional networks from
57 deconvolved BOLD signals. Such approaches have thus revealed the existence of
58 neurally-related events in the BOLD signal during periods of rest (Petridou et al.,
59 2013; Karahanoğlu and Van De Ville, 2015; Gonzalez-Castillo et al., 2019). How-
60 ever, Total Activation (Karahanoğlu et al., 2013) suffers from a main limitations:
61 The HRF shape in this model was fixed and kept constant throughout the brain,
62 making this approach primarily suitable for normal populations. Without disen-
63 tangling the neural and vascular components, but taking into consideration the
64 spatial variability in the neurovascular coupling, Caballero-Gaudes et al. (2012)
65 have proposed a voxelwise HRF modeling.

66 Alternatively some recent works have suggested either a separate or concurrent
67 estimation of the neural activity and the HRF profile (Wu et al., 2013; Cherkaoui
68 et al., 2019; Farouj et al., 2019). These approaches are often referred to as *semi-*
69 *blind deconvolution schemes* of the BOLD signal as they do not fully estimate an
70 unknown HRF shape from scratch but instead try to uncover a transformation
71 from a pre-existing HRF profile such as the canonical shape. Both (Cherkaoui
72 et al., 2019) and (Farouj et al., 2019) rely on the hypothesis of a block-type prop-
73 erty for the neural activation signal as initially proposed in (Karahanoğlu et al.,

2013) but are able to infer the haemodynamic profile for instance with a limited parameterization to deal with the magnitude and delay ambiguities between the neural input the haemodynamic filter.

1.3. Goals and contributions

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

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

2. Multivariate low-rank decomposition of the BOLD signal

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

Notation

In what follows, y_i denotes the i^{th} entry in vector \mathbf{y} . Let $\tilde{T} = T - L + 1$, the 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 $\mathbf{v} * \mathbf{a} \in \mathbb{R}^{1 \times T}$. 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} . 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




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 white Gaussian noise. The measured fMRI signal is denoted \mathbf{y} and obtained via $\mathbf{a} * \mathbf{v} + \mathbf{e}$.

111 corresponding line of \mathbf{V} . Let ∇ be the first-order difference operator such that
 112 $\forall \mathbf{x} \in \mathbb{R}^{\tilde{T}}, \nabla \mathbf{x} \in \mathbb{R}^{\tilde{T}-1}$ with $(\nabla \mathbf{x})_i = x_i - x_{i-1}, \forall i \in \{2, \dots, \tilde{T}\}$.

113 2.1. Linear and time-invariant modeling

114 2.1.1. Univariate modeling

115 A common model for the multivariate (P voxels, T scans) BOLD data $\mathbf{Y} \in$
 116 $\mathbb{R}^{P \times T}$ with $\mathbf{Y} = (\mathbf{y}_j)_{j=1}^P$ is the linear and time-invariant model (LTI) (Boynton
 117 et al., 1996). This model is illustrated in Fig. 1. For each voxel, the measured
 118 time series, denoted by $\mathbf{y}_j \in \mathbb{R}^{1 \times T}$, is the convolution of a neural activation signal
 119 $\tilde{\mathbf{a}}_j \in \mathbb{R}^{1 \times \tilde{T}}$, 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
 120 Gaussian noise, which leads to:

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

121 Although the noise that contaminates the BOLD effect is serially correlated in
 122 time (Woolrich et al., 2001), here for the sake of simplicity we assume that the
 123 noise corrupting the fMRI time series is white and Gaussian. As shown hereafter
 124 in Section 2.1.2, this assumption is tenable as our model will be flexible enough
 125 to capture and segregate noise-related components. However, for the sake of com-
 126 pleteness, additional justification can be found in the supplementary material (cf.
 127 Section A.8). In the latter, we studied on simulations the impact of a first-order
 128 autoregressive Gaussian noise on the estimation accuracy and showed that for
 129 SNR ($\simeq -8$ dB) and autoregressive parameter (≤ 0.3) values met on real fMRI
 130 data, the reconstruction error due to the model mismatch remains quite low.

131 Typically, the HRF \mathbf{v} has a restricted support in time of about 20 s. The
 132 challenge with HRF modeling is to find a fair trade-off between a flexible model
 133 that is able to capture the true haemodynamic response in each brain area and a
 134 reliable one that limits overfitting by reducing the number of degrees of freedom
 135 (diminishing the variance). Since our approach estimates the neural activity along
 136 with the HRF, reducing the number of degrees of freedom is critical to avoid the
 137 aforementioned overfitting. In this paper, we assume M different HRF with \mathbf{v}_m
 138 being the HRF corresponding to the m^{th} region Θ_m . Numerous approaches have




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 $TTP(h_1) < TTP(h_2)$ and $FWHM(h_1) < FWHM(h_2)$.

139 been proposed to model this haemodynamic response \mathbf{v}_m . The Finite Impulse Re-
 140 sponse (FIR) (Dale, 1999; Glover, 1999) model does not assume any particular
 141 shape for the HRF which make it very flexible but prone to data overfitting in the
 142 presence of noise. Regularization has thus been introduced to constrain the overall
 143 HRF shape in FIR models and limit their tendency to overfitting, see for instance
 144 penalization over the second-order derivative to end up with physiologically plau-
 145 sible smooth HRF estimates (Ciuciu et al., 2003; Casanova et al., 2008). Alterna-
 146 tively, the HRF has been modeled as a linear decomposition of predefined atoms
 147 such as B-splines (Zhang et al., 2007; Vakorin et al., 2007), wavelets (Khalidov
 148 et al., 2011), a sensitivity-selected set (Woolrich et al., 2004) or more physiologi-
 149 cally informed patterns such as the canonical HRF and its derivatives in time and
 150 with respect to the dispersion parameter (Friston et al., 1998). All these methods
 151 intend to capture fluctuations in haemodynamic delay or shape with the mini-
 152 mum number of parameters. Last, to constrain even more the parameter values
 153 and reduce variance estimates, parametric models such as the *inverse logit trans-*
 154 *form* (Lindquist and Wager, 2007) have been proposed and successfully tested
 155 when benchmarking over multiple fMRI datasets (Lindquist et al., 2009). In this
 156 work, we propose the time dilation HRF model (Cherkaoui et al., 2019), which cap-
 157 tures the haemodynamic delay by dilating a reference HRF pattern: $\mathbf{v}_\delta = \mathbf{v}_{ref}(\delta t)$
 158 where \mathbf{v}_{ref} is a reference temporal profile, here the canonical HRF. This approach
 159 is efficient while simple as it encodes delay fluctuations through a single scalar
 160 parameter δ (one degree of freedom). One limitation of this choice is that it leads
 161 to the simultaneous variation of the full width-at-half-maximum (FWHM) of
 162 the HRF and its time-to-peak (TTP) (see also Fig. 2). This coupling between
 163 TTP and FWHM partly alleviates time-shift ambiguities between the HRF and
 164 the neural activity signal during the blind deconvolution process. Although this
 165 model is quite simple, it remains attractive as it provides a low variance HRF esti-
 166 mate while capturing the main effects on the neurovascular coupling as described
 167 in Section 4. Our voxelwise model reads as follows:

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

168 *2.1.2. Multivariate modeling*

169 To better account for the spatial structure of the neurovascular system and
 170 the intrinsic organization of functional networks, we extend this model to the
 171 multivariate setting. Our multivariate model reads as follows:

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

172 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
 173 straightforward approach is that it constrains the haemodynamic response to be
 174 the same across the whole brain. As the HRF shape depends on the neurovascular
 175 coupling, its features vary in space over different brain areas and between individ-
 176 uals (Handwerker et al., 2004; Badillo et al., 2013). This suggests that, for a given
 177 subject, the HRF should be modeled locally in the brain. An appropriate approach
 178 for doing so is to rely on existing brain parcellation (Varoquaux and Craddock,
 179 2013). Ideally to accurately fit the real haemodynamic response function in a sub-
 180 ject, we would favor a large number of regions. However, the larger this number,
 181 the smaller the number of voxels per region, which could impair the stability of
 182 HRF estimation. In practice, using a minimum of 50 voxels per region leads to sta-
 183 ble HRF estimates. In contrast, aggregating too many voxels (e.g. $> 10^3$) results
 184 in identical HRFs over the whole brain. Hence, in all our experiments, we used
 185 parcellations where each region consisted of a few hundred voxels. We refer the
 186 reader to (Vincent et al., 2008) for more information on this point. In this work,
 187 we rely on the Harvard-Oxford probabilistic brain atlas (Desikan et al., 2006). We
 188 threshold the probabilities to obtain a fine brain parcellation that offers enough
 189 flexibility to adapt to the true haemodynamic system. In what follows, we math-
 190 ematically introduce a brain parcellation with M regions with $(\Theta_m)_{m=1}^M \in \{0, 1\}^P$
 191 1 if the i^{th} voxel belongs to the m^{th} region and 0 if not. This allows us to extend
 192 Eq. (3) as follows:

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

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

197 In our work, we rather introduce a low-rank constraint and learn both K temporal
 198 atoms (with $K \ll P$) and corresponding spatial maps. These maps encode various
 199 functional networks, each of them being summarized by specific neural activation
 200 profile. Mathematically, this can be modeled by replacing each vector $\tilde{\mathbf{a}}_j$ in Eq. (4)
 201 with a linear combination of neural activation patterns $\mathbf{Z} = (\mathbf{z}_k)_{k=1}^K \in \mathbb{R}^{K \times \tilde{T}}$, with
 202 $\mathbf{z}_k \in \mathbb{R}^{1 \times \tilde{T}}$, modulated in space by the spatial maps $\mathbf{U} = (\mathbf{u}_k)_{k=1}^K \in \mathbb{R}^{K \times P}$, with




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.

203 $\mathbf{u}_k \in \mathbb{R}^{1 \times P}$, such that: $\tilde{\mathbf{A}} = \sum_{k=1}^K \mathbf{u}_k^\top \mathbf{z}_k$. In other words, the spatial config-
 204 \mathbf{u}_k encodes which voxels are linked to a given neural activation profile
 205 $\mathbf{z}_k \in \mathbb{R}^{1 \times T}$. Note that a voxel may belong to different functional networks as over-
 206 lapping can occur (Karahanoğlu et al., 2013; Najafi et al., 2016). Additionally,
 207 the sub-components of a functional network might also not activate simultane-
 208 ously (Allan et al., 2015; Betzel et al., 2016) as they can belong to different spatial
 209 maps \mathbf{u}_k . This is coherent with the fact that a voxel contains about one hundred
 210 thousands neurons for a typical spatial resolution (1.5 mm isotropic), and thus
 211 the underlying neural activation signals are combined with possibly different tem-
 212 poral fingerprints. Finally, our forward model for BOLD fMRI data is given by
 213 Eq. (5) (see also Fig. 3):

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

214 2.2. Prior information and regularizing constraints

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

222 The forward model described in Eq. (5) is *trilinear* as it depends on both un-
 223 known spatial (\mathbf{U}) and temporal (\mathbf{Z}) inputs and unknown convolution filters (\mathbf{v}_{δ_m}).

224 This means that any timing variation in neural activation signals \mathbf{Z} can be sym-
 225 metrically compensated by the opposite variation in the haemodynamic filter \mathbf{v}_{δ_m}
 226 or that any sign or scale variation in \mathbf{U} can be compensated by the inverse vari-
 227 ation in \mathbf{Z} or \mathbf{v}_{δ_m} . To overcome these sign, scale and time-shift ambiguities, we
 228 introduce some regularizing constraints in the computation of the parameter esti-
 229 mates. First, to avoid any sign ambiguity in the convolution between the neural
 230 input signals and the haemodynamic filter, the HRF filter has a constant and pos-
 231 itive maximal amplitude (see Fig. 2). Second, as there is an interplay between
 232 the spatial and temporal components in the input signals, we also impose a non-
 233 negativity constraint over the entries of each spatial map \mathbf{u}_k , i.e. $\forall j, k_{kj} \geq 0$.
 234 By doing this, we only allow the deactivation of a functional network to be en-
 235 coded in the corresponding \mathbf{z}_k . Third, to deal with the scale ambiguity, we impose
 236 $\|\mathbf{u}_k\|_1 = \eta, \forall k = 1, \dots, K$, where $\eta \in \mathbb{R}_+$ is a user-defined parameter that sets
 237 the magnitude of each spatial map. As our HRF model has a constant maximal
 238 amplitude, only the neural activity signals \mathbf{Z} capture the observed BOLD sig-
 239 nal fluctuations. Fourth, to deal with potential time-shift indeterminacy, recover
 240 plausible physiological HRF shape and make the optimization easier, we constrain
 241 the dilation parameters $\boldsymbol{\delta}$ in the HRF model to lie in $[0.5, 2.0]$. Complementary
 242 to these constraints, similar to (Karahanoğlu et al., 2013; Cherkaoui et al., 2019;
 243 Urunuela et al., 2020) we will assume the neural activation signals \mathbf{Z} to be tempo-
 244 rally piecewise constant or equivalently consider their first-order derivative $\nabla \mathbf{z}_k, \forall k$
 245 to be sparse. Practically speaking, we add a total variation (TV) regularization
 246 term to our model.

2.3. Optimization problem

248 In this section, we derive an algorithm for estimating the dilation parameters
 249 $\boldsymbol{\delta} = (\delta_m)_{m=1}^M$, the spatial maps \mathbf{U} and the neural activation signals \mathbf{Z} from the
 250 model depicted in Eq. (5) and the aforementioned regularizing constraints. The
 251 estimates for these parameters can be obtained by solving the following constrained
 252 minimization problem:

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

253 The white Gaussian noise hypothesis leads to a quadratic data fidelity term,
 254 to measure how well we reconstruct the observed fMRI signals. Moreover, in
 255 alignment with with Caballero-Gaudes et al. (2012); Karahanoğlu et al. (2013);
 256 Cherkaoui et al. (2019), we enforce the first-order derivative of the temporal atoms
 257 $(\nabla \mathbf{z}_k)_{k=1}^K$ to be sparse in order to constrain each \mathbf{z}_k to a piecewise constant sig-
 258 nal. For that purpose, we use a TV regularization term, which corresponds to

259 the ℓ_1 norm of the gradient in time $\nabla \mathbf{z}_k$. Interestingly, the sparsity of the first-
 260 order derivative of the temporal components $(\mathbf{D}\mathbf{z}_k)_k$ tends to impose independence
 261 between these temporal components (Daubechies et al., 2009). This property is
 262 preserved by the discrete integration operator. Thus, the complete low-rank neural
 263 components should tend to be de-coupled due to the independence of the tempo-
 264 ral processes $(\mathbf{z}_k)_k$. Importantly, this modeling of the neural activation signals
 265 allows us to fully adapt to task-fMRI and rs-fMRI experiments and to perform
 266 paradigm-free fMRI data analyses (Deneux and Faugeras, 2010; Caballero-Gaudes
 267 et al., 2011). In the first case, as the task-related BOLD signal is classically mod-
 268 eled as the convolution of an input block signal, representing the experimental
 269 paradigm (the onsets of the stimulus trials) with a HRF filter, we can recover neu-
 270 ral activation signals close to the experimental paradigm. The neural activation
 271 profiles by being inferred from the data, allows estimation of input signals cor-
 272 responding either to block or event-related paradigms, the second scenario being
 273 seen as a single time point block. More interestingly, the proposed framework is
 274 even more appealing for processing rs-fMRI data and uncover spontaneous and
 275 time-varying fluctuations of brain activity as the block duration may change from
 276 one instance to the next (see also (Petridou et al., 2013; Karahanoglu and Van
 277 De Ville, 2015)).

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

288

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

289

290 In regards to the $(\mathbf{z}_k)_{k=1}^K$ step, we performed the minimization using an adaptive-
 291 restart accelerated forward-backward algorithm (O’Donoghue and Candes, 2015).
 292 In regards to the $(\mathbf{u}_k)_{k=1}^K$ step, we noticed that it can be cast as a standard
 293 quadratic program of the form $\arg \min_{\text{vec}(\mathbf{U})} \text{vec}(\mathbf{U})' \mathbf{A} \text{vec}(\mathbf{U}) - 2 \text{vec}(\mathbf{U})' \mathbf{b}$ with
 294 linear and non-negativity constraints on $\text{vec}(\mathbf{U})$, with $\text{vec}(\cdot)$ the usual matrix-
 295 to-vector transformation. Although the matrix \mathbf{A} and constraints are sparse, the
 296 problem is high-dimensional (PK variables, K linear constraints) and may not be
 297 efficiently solvable by standard mathematical programming solvers. Consequently,
 298 we benchmarked various algorithms in the dictionary learning literature and se-
 299 lected the most efficient, namely the one used to update the dictionary in Mairal
 300 et al. (2009). Last, for the minimization with respect to $(\delta_m)_{m=1}^M$ we used the
 301 accelerated forward-backward algorithm (Beck and Teboulle, 2009; Bioucas-Dias
 302 and Figueiredo, 2007) after checking that it leads to a correct estimation of $\boldsymbol{\delta}$. The
 303 reader can find all details of the gradient computation w.r.t. $(\mathbf{z}_k)_{k=1}^K$ and $(\mathbf{u}_k)_{k=1}^K$
 304 in Cherkaoui et al. (2019), for the gradient w.r.t. $\boldsymbol{\delta}$ the computation is reported in
 305 the Supplementary Material, see Section A.1. For each step, we implemented the
 306 corresponding gradient (i.e. forward move) in an efficient manner to limit multiple
 307 computations over the iterations.

308 Critical steps for the efficiency of this algorithm are the computation of proxi-

mal 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 use 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 a constant \mathbf{Z} is solution of Eq. (6), we detail the computation of λ_{\max} in Supplementary Material, cf. Appendix A.2. 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 $(\delta_m)_{m \in \{1 \dots M\}}$ to 1. As the current optimisation problem is non-convex, we investigated different initialization strategies to assess how much the recovered parameters depend on the initial values (see Supplementary Material, Section A.4).

For the sake of reproducibility of the current results and others not shown here, we offer an efficient software implementation of our fast algorithm and with numerous code optimizations. The underlying `hemolearn` Python package is open source³ and available to the neuroimaging community.

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

¹<https://pypi.org/project/prox.tv/>

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

³code available at <https://github.com/hemolearn/hemolearn>.

342 shape compared to a single deconvolution scheme. Next, we will demonstrate the
 343 usefulness of the proposed framework on real rs-fMRI data at the individual level.
 344 In particular, we will highlight the impact of the hyper-parameter selection on the
 345 decomposition, describe the interpretation of its component and its stability. The
 346 usefulness of our method in the context of large cohorts will be investigated in
 347 Section 4.

348 3.1. Numerical simulations

349 3.1.1. Synthetic data

350 We generated two temporal Dirac signals of length $\tilde{T} = 500$ with a fixed
 351 sparsity level. The first generated Dirac signal is composed of randomly drawn
 352 signed spikes, with location chosen uniformly in time and intensity drawn from
 353 a Gaussian distribution $\mathcal{N}(0, 1)$ to simulate spontaneous (i.e. resting-state) neu-
 354 ral activity fluctuations. To produce the second temporal atom, we generated 4
 355 spikes (2 with positive amplitudes, 2 with negative amplitudes) to simulate a task-
 356 related response. To produce the corresponding block signals \mathbf{Z} – shown in blue
 357 in Fig. 4[right panels] – we integrated over time these signals and convolved them
 358 with a predefined HRF \mathbf{v}_δ to yield two corresponding pure BOLD time series. The
 359 chosen HRF has length $L = 25$ and is shown in blue in Fig. 4[left panel]. For the
 360 sake of simplicity, we considered a single HRF profile ($M = 1$) in this synthetic
 361 setting, so the haemodynamic properties were supposed constant in space. We
 362 then assigned these BOLD signals to spatial locations. Hence, we defined $K = 2$
 363 corresponding 2D maps $\mathbf{U} = (\mathbf{u}_k)_{k=1}^2$ (20×20 , i.e. $P = 200$). Each spatial map
 364 has a single activating squared region consisting of $4 \times 4 = 16$ pixels. Each active
 365 pixel has a randomly drawn non-negative magnitude, the other ones being set to
 366 zero. Then, we normalized each map by its ℓ_1 -norm. Finally we added white
 367 Gaussian random noise to produce observed, i.e. noisy BOLD signals \mathbf{Y} of length
 368 $T = \tilde{T} + L - 1 = 524$ scans ($TR = 1$ s) with a signal-to-noise-ratio (SNR) of -10 dB.
 369 The mean synthetic BOLD signals are reported in black traces in Fig. 4(a)-(b) for
 370 both activated regions (in bottom panels) while standard deviation across acti-
 371 vated voxels is encoded by transparency around these mean curves. To make a
 372 comparison with respect to the state of the art, we deconvolved the observed signal
 373 using the method proposed in (Wu et al., 2013). This approach first detects the
 374 main neurally-related events by thresholding the observed BOLD signal. A binary
 375 sparse sequence associated with these spikes is thus constructed to be used as a
 376 proxy of a neural input signal. Then, for each voxel, the HRF profile is fitted to
 377 the data using this sequence of neural events. As a result, the BOLD signal can
 378 then be deconvolved voxelwise using the estimated HRF.

379 *3.1.2. Numerical results*

380 In a first step, we only estimated the pair (\mathbf{Z}, \mathbf{U}) from the synthetic fMRI time
381 series \mathbf{Y} and kept the HRF profile \mathbf{v} constant. The results are reported in Fig. 4(a).
382 The HRF shape used in this deconvolution process is shown in green in Fig. 4(a)
383 and actually differs from the true shape used for simulating the data. Because
384 of this discrepancy in terms of haemodynamic delay and peak magnitude, the
385 neural activation signals are not properly recovered (orange traces in Fig. 4(a)).
386 The magnitude of the estimates $\hat{\mathbf{Z}}$ is much larger than the true one. This is
387 partly due to compensate for the smaller magnitude of the HRF (green trace in
388 Fig. 4(a)) used for deconvolving the BOLD signals. Consequently, the residual
389 mean square errors (RMSEs) computed on the neural activation signals are pretty
390 large. However, we noticed that in both spatial maps, the non-negative magnitudes
391 $\hat{\mathbf{U}}$ are very well estimated. This is a direct consequence of using non-overlapping
392 activating regions for the two neural traces. Additionally, we report, using a dashed
393 green line, the case where the HRF profile is set to the true shape. In this case,
394 we observe that the neural temporal components are accurately recovered.

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

404 Finally, for comparison purposes, we report in Fig. 4(c), the estimates com-
405 puted by (Wu et al., 2013). We noticed an average error of 3s on the haemody-
406 namic delay estimate. Moreover, for a selected voxel localized in the first activated
407 region, we display the estimated neural events as well as the deconvolved BOLD
408 signal and compare it to the true first temporal atom, its first-order derivative
409 and the observed noisy BOLD signal. It is worth noticing that Wu et al. (2013)
410 provided a poor alignment of the estimated neural spikes with the true ones as
411 compared to our approach.

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

416 3.2. Single-subject analysis on rs-fMRI data

417 3.2.1. Dataset and parameter setting

418 To illustrate the proposed semi-blind deconvolution algorithm, we analyzed a
419 single subject extracted from the UKBB resting-state fMRI data set. More inves-
420 tigation on a larger cohort of this data set will be presented in Section 4. The
421 rs-fMRI data was 6 min10s long with TR = 0.735s. The first ten seconds were
422 discarded (dummy scans) so that we end up with $T = 490$ scans (6min). The data
423 was collected on a 3T Skyra Siemens MAGNETOM MR system at an isotropic res-
424 olution of $2.4 \times 2.4 \times 2.4 \text{ mm}^3$ using the multi-band GRE sequence (mb = 8).⁴ Stan-
425 dard preprocessing steps were applied: motion correction using MCFLIRT [Jenkinson et al. \(2002\)](#);
426 grand-mean intensity normalisation of the entire 4D data set
427 by a single multiplicative factor; high-pass temporal filtering (Gaussian-weighted
428 least-squares straight line fitting); EPI unwarping; gradient distortion correction
429 unwarping. Structured artefacts are removed by ICA processing, see the documen-
430 tation⁵ for a full description. Finally, we down-sampled the voxels to reduce the
431 spatial dimension of our decomposition to 61,605 voxels to match the resolution
432 of 2 mm isotropic of the Harvard-Oxford probabilistic atlas ([Desikan et al., 2006](#)).

433 In this subsection, we manually set the temporal regularization parameter to
434 $\lambda_f = 0.8$ (remember $\lambda_f \in [0, 1]$). This setting achieves a bias-variance trade-off
435 between two extreme situations, namely data overfitting ($\lambda_f = 0$) on one hand
436 and entirely sparse neural activation signals ($\lambda_f = 1$ as $\lambda = \lambda_{\max}$) on the other
437 hand. The question of the unsupervised tuning of λ_f is critical. It could be driven
438 either from a statistical viewpoint (e.g. using the maximum likelihood criterion)
439 that characterizes how likely the measured time series may be observed or using an
440 external task and its corresponding metric such as classification performance (e.g.
441 accuracy in prediction). Because the former approach does not admit a closed
442 form solution, we explored in the Supplementary Material (see Section A) the
443 impact of changing the temporal regularization (i.e. amount of sparsity in the
444 activation neural signals) on the spatial map of haemodynamic delays. The spa-
445 tial consistency we reported across regularization levels gave us confidence on the
446 haemodynamic parameter estimates to be further used in subsequent classification
447 tasks (see Section 4). In this setting, we implemented a cross-validation step with
448 a leave-one-out loop to tune λ_{\max} in an unsupervised way.

449 In regard to the number of spatio-temporal atoms K , we set it using the ex-
450 plained variance (or R^2 -score averaged over the time series) as target metric in a
451 preliminary study. For this set of parameters, the model estimation took around

⁴Acquisition details can be found at <https://www.fmrib.ox.ac.uk/ukbiobank/protocol/>.

⁵Preprocessing details can be found at https://biobank.ctsu.ox.ac.uk/crystal/crystal/docs/brain_mri.pdf

452 1 minute on a machine with 15 GB of RAM and an Intel processor i7-7600U (2
453 physical cores, 2.80 GHz).

454 3.2.2. Results

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

462 The R^2 -score is defined as follows: $R^2 = 1 - \frac{SS_{\text{res}}}{SS_{\text{tot}}}$ where SS_{tot} quantifies the vari-
463 ance of the data \mathbf{Y} and SS_{res} the variance of the residuals after fitting model (5)
464 by minimizing the cost function described in Eq. (6). The R^2 -score may vary from
465 $-\infty$ in pathological cases to 1 for a perfectly matching model. A good model is
466 normally associated with $R^2 > 0$ and means that the L_2 norm of the residual is
467 lower than the variance of the data. We therefore ran multiple model fitting for
468 K in a range of $\{2, 3, 4, \dots, 10, 15, 20, \dots, 50\}$. The results are shown in Fig. 5(a)
469 and illustrate that the model accuracy first increases as a function of K up to
470 reaching a plateau around $R^2 \simeq 0.55$ for $K = 20$. So adding more spatio-temporal
471 components no longer improves its ability to capture variability in the data while
472 it becomes more complex.

473 The second information measure we used to help us select K was based on the
474 determinant of the correlation matrix $\Sigma_K = (\mathbb{E}[(\mathbf{z}_k - \mathbf{m}_k)(\mathbf{z}_\ell - \mathbf{m}_\ell)^T / \sigma_k^2 \sigma_\ell^2])_{k,\ell}$
475 between the temporal atoms $(\mathbf{z}_k)_{k=1}^K$. The quantities σ_k^2 and σ_ℓ^2 define the variance
476 of the neural activation signals \mathbf{z}_k and \mathbf{z}_ℓ . As Σ_K is semi-positive definite with
477 entries between 0 and 1, its eigenvalues are positive or null and so its determinant
478 varies between 0 and 1: $\det \Sigma_K = 1$ when matrix Σ_K defines a basis, which means
479 that all atoms are orthogonal and decorrelated like in a PCA decomposition. In
480 contrast, $\det \Sigma_K = 0$ when matrix Σ_K is not of full rank so at least one atom
481 could be obtained as a linear combination from the others. Therefore, as before
482 we ran multiple model fitting for K in a range of $\{2, 3, 4, \dots, 10, 15, 20, \dots, 50\}$
483 and we plotted in Fig. 5(b) the evolution of the determinant of Σ_K as a function
484 of K . The results show us that beyond $K \geq 20$, we get a correlation matrix with
485 $\det \Sigma_K \leq 10^{-10}$ which tends to zero. According to this criterion, from the we
486 should therefore not exceed 20 temporal atoms. Thus, from the R^2 -score criterion
487 and this $\det \Sigma_K$ criterion, in our following experiments we will keep $K = 20$.

488 **Analysis of spatial decomposition.** Fig. 6 shows the spatial maps of this
489 spatio-temporal decomposition for this individual and Tab. 1 summarizes the list
490 of functional networks retrieved in this setting. It is worth mentioning that the

Table 1: **Taxonomy of functional networks** involved in the spatio-temporal decomposition (5) with $K = 20$, whose spatial maps are shown in Fig. 6.

Networks	#Comp. id k
Default Mode	6, 11, 18, 20
Attention	5
Central executive	10, 12, 13
Saliency	19
Sensori-motor	3, 7, 9, 15, 17
Auditory-language	2, 8, 14
Visual	1, 4, 16

491 sensory (visual, auditory and motor) and resting-state networks (RSN) are quite
492 well retrieved by multiple components, localized respectively in the occipital (com-
493 ponents 1, 4, 16) temporal (components 2, 8 and 14), fronto-parietal (components
494 5, 10, 11, 13, 17 and 20), frontal/fronto-polar (components 6, 12, 18 and 19) and
495 sensori-motor (components 3, 7, 9 and 15) cortices. The maps associated with
496 the motor network are split (3 and 9 lateralized on the left and right hemispheres,
497 respectively). The different areas of the visual system are split too between the pri-
498 mary visual cortex (components 1 and 4) and the extrastriate cortex (components
499 16). The language system is also spread across several components (see Tab. 1).
500 Part of the well known RNS (Menon, 2015) have been captured: The default
501 mode network (DMN), which deactivates during cognitive tasks is represented in
502 components 6, 11, 18 and 20, while the attention and saliency networks by a sin-
503 gular component (5 and 19, respectively). The central executive network (CEN) is
504 spread over components 10, 12 and 13.

505 **Spatio-temporal decomposition.** To fully illustrate our method on real rs-
506 fMRI data, we show the whole set of output features (neural activation signals,
507 spatial maps, HRF shapes) in Fig. 7. We also depict a voxel-based denoised
508 BOLD signal reconstructed and the convolution between the neural input and the
509 HRF estimate. Fig. 7(a) represents together a neural activation signal in the pri-
510 mary visual cortex and the corresponding spatial map (component 1 in the above
511 mentioned decomposition). The proposed axial views permit the identification
512 of the primary visual cortex and the calcarine fissure. Fig. 7(b) depicts similar
513 features in the DMN (component 11) and the Pearson correlation coefficient with
514 the neural time course in the visual cortex. Its small negative value confirms a
515 slight negative correlation between the task-positive and the DMN network. Both
516 time courses actually present alternating periods of positive and negative activity
517 but they are almost uncorrelated. Fig. 7(c)-(d) illustrate the fastest and slow-
518 est HRF time courses estimated in the regions of interest depicted in red. The

519 fastest haemodynamic response ($\text{FWHM}_f = 5.1\text{s}$) was found in the middle tem-
520 poral gyrus while the slowest ($\text{FWHM}_s = 8.0\text{s}$) is localized in the frontal orbital
521 cortex. Fig. 7(e) finally shows how well our approach is able to fit the rs-fMRI
522 time course measured in voxel marked by the black cross in Fig. 7(c). The neural
523 activation signal is piecewise constant and ahead in time compared to the BOLD
524 time series. Once convolved with the HRF profile, the denoised BOLD signal ap-
525 pears as a smoother version of the measured BOLD time course: its magnitude is
526 smaller and its fluctuations in time are slower. This is a direct consequence of the
527 temporal regularization used to recover sparse input signals.

528 **Analysis of correlation structure.** Next, to go beyond the spatial analysis,
529 Fig. 8 depicts the correlation matrix between the corresponding neural activation
530 signals. It is then insightful to notice that the correlation between the multiple
531 components in a given network are quite strong. For the visual network we ob-
532 served a correlation coefficient varying between 0.31 and 0.48, the largest value
533 being reached for areas located both in the extrastriate cortex. The same conclu-
534 sion holds in the motor network with a correlation level varying between 0.1 and
535 0.36. In regard to the DMN, component 11 plays the role of a hub as it corre-
536 lates with components 10 and 20 pretty strongly⁶ between 0.46 and 0.54. Overall,
537 this analysis shows that the proposed approach does not separate RSN in single
538 components. However, it still achieves a meaningful decomposition. Although we
539 do not report here the spatial decomposition for $K \in \{5, 8, 10, 15, 30, 40, 50\}$, the
540 latter can be easily obtained using the HemoLearn toolbox⁷. Interestingly, in such
541 decomposition some RSNs are either not recovered or mixed together for small
542 $K \leq 15$ whereas for large values of $K \in \{25, 30, \dots, 50\}$ each RSN is split in
543 multiple components. Increasing the number of components extends the range of
544 variation of the correlation coefficients in both positive and negative senses. This
545 confirms that our approach does not have statistical independence guaranties like
546 PCA and ICA do when decomposing the neural activity.

547 3.3. Single-subject analysis on task-fMRI data

548 3.3.1. Data set and parameter setting

549 To further illustrate the proposed semi-blind deconvolution algorithm, we an-
550 alyzed a single subject extracted from the HCP (Van Essen et al., 2013) data set
551 and focused on the motor task in the activation fMRI experiment. The fMRI run
552 was 3min34s long with $\text{TR} = 0.753\text{s}$. The first ten seconds were discarded (dummy
553 scans) so that we ended up with $T = 284$ scans collected by interleaved simulta-
554 neous multislice echo-planar images with a multi-band factor of 8 and a spatial

⁶no statistical test performed at the individual level

⁷<https://github.com/hemolearn/hemolearn>

555 resolution of $2 \times 2 \times 2\text{mm}^3$. The fMRI images were already preprocessed using a
556 classical pipeline including realignment, coregistration, spatial normalization and
557 smoothing (5mm isotropic)⁸. Additionally, we spatially down-sampled the fMRI
558 images to reduce the spatial dimension of our decomposition to 61,605 voxels.
559 The EP was divided in two sets of motor tasks, with 15s fixation blocks at the
560 beginning, in the middle of the acquisition and at the end of the recording. Each
561 task set was composed of 5 blocks of 12 s each, preceded by a 3s cue indicating the
562 task to be performed by the participant. The latter corresponded to moving the
563 tongue, tapping the left or right finger or contracting the left or right toes. In what
564 follows, we only consider one participant even though our results are reproducible
565 across individuals.

566 In this subsection, we manually set the temporal regularization parameter to
567 $\lambda_f = 0.8$ as in the prior rs-fMRI analysis and the number of spatio-temporal atoms
568 to $K = 30$ to offer a richer spatio-temporal decomposition with possibly more
569 focal components, compared to resting-state analysis. We compared our approach
570 to (Pedregosa et al., 2015) in which a voxelwise HRF estimate is yielded using a
571 rank-1 GLM model. This permits a fair comparison of both haemodynamic delays
572 and spatial activation maps corresponding to a given condition from the EP.

573 3.3.2. Results

574 For illustrative purposes, we focused on the *left-hand* motor task. Fig. 9(a)
575 represents the 100 most activated voxels estimated by (Pedregosa et al., 2015)
576 in response to this condition. These voxels were obtained by thresholding the
577 regression coefficients of the model displayed in Fig. 9(c). From the $K = 30$
578 spatio-temporal components of our model, we manually chose the 22nd temporal
579 component that offers the most correlated temporal pattern with the experimental
580 block paradigm associated with the *left-hand* motor task. Fig. 9(b) represents the
581 100 most activated voxels, localized in the right motor cortex, obtained by thresh-
582 olding the 22nd spatial map shown in Fig. 9(d). We notice a clear disagreement
583 among the two approaches on the localization of the most activated regions in
584 response to this experimental condition. Indeed, the activation of left part of the
585 cerebellum found by (Pedregosa et al., 2015) is quite suspect as we expect the *left-*
586 *hand* motor task to induce evoked activity in the right motor cortex. This result
587 might be explained by the absence of a clear evoked activity in response to the
588 second block of stimulation as shown in the average BOLD signal over the right
589 motor cortex and the fact that Pedregosa et al. (2015)’s method is *univariate*, see
590 Fig. 9(e). In contrast, the neural activity recovered by the proposed *multivariate*

⁸Acquisition details can be found at https://humanconnectome.org/storage/app/media/documentation/s1200/HCP_S1200_Release_Appendix_VI.pdf

591 method (cf. Fig. 9(e)) and associated with the voxels localized in the right motor
 592 cortex shows strong consistency in time with the onsets of the blocks of the motor
 593 task the participant has to undergo. For ease of legibility, this part of the experi-
 594 mental paradigm is displayed in light gray. Last, we compared the HRF profiles in
 595 the activated regions the two approaches provide (cf. Fig. 9(f)). Again, we notice
 596 a clear disagreement among the two competing methods on the haemodynamic
 597 delays, which is explained by the different voxel localization (shorter TTP for Pe-
 598 dregosa et al. (2015)’s method due to the recovered activity in the visual cortex
 599 vs longer TTP for our method that picks up the activation in the primary right
 600 motor cortex).

601 3.4. haemodynamic estimation stability over time

602 The shape of the HRF is controlled by the neurovascular coupling including
 603 both neural and non neural factors such as glial cell activity, cerebral energy
 604 metabolism, and the cerebral vasculature. Abnormalities in the local vascular sys-
 605 tem or cell communication due to pathological state or changes in cerebral blood
 606 flow upon psychoactive drugs could influence this haemodynamic response. As we
 607 expect the HRF estimate to be stable if none of those events took place, we propose
 608 to study the within-subject stability of HRF estimates over time, namely between
 609 consecutive time periods. For doing so, we compare the within-subject variability
 610 of the HRF whole brain dilation parameter vector δ to the inter-subject variability
 611 of the same quantity. We thus introduce two reference ℓ_2 distances, namely the
 612 within-subject distance $WS(\delta_1^s, \delta_2^s) = \|\delta_1^s - \delta_2^s\|_2^2$ where (δ_1^s, δ_2^s) correspond to the
 613 vectors of spatially aggregated HRF dilation parameters that were estimated over
 614 two periods of time T_1 and T_2 in the same individual s . Similarly, for any pair of
 615 subjects (s_1, s_2) and a given period T , we measure the between-subject distance
 616 between $(\delta_T^{s_1}, \delta_T^{s_2})$ vectors as follows: $BS(\delta_T^{s_1}, \delta_T^{s_2}) = \|\delta_T^{s_1} - \delta_T^{s_2}\|_2^2$. The goal is then
 617 to compare the within- and between-subject distances across individuals and show
 618 that the intra-subject variability is significantly lower than the inter-subject one
 619 over a sufficiently large population.

620 3.4.1. Data set and numerical analysis

621 We selected 100 healthy subjects from the Human Connectome Project (HCP)
 622 data set (Van Essen et al., 2013) at random. We used this data set because
 623 of the availability of a 12-min long rs-fMRI run with a short time of repeti-
 624 tion ($TR=0.753s$), see Glasser et al. (2013) for a full description of the acqui-
 625 sition parameters and the pre-statistics processing steps. In this rs-fMRI run for
 626 each individual, we extracted two segments of 4 minutes each, denoted as T_1 and
 627 T_2 hereafter, the first and last parts of the recording. We then applied the pro-
 628 posed multivariate spatio-temporal decomposition to each segment using $K = 8$
 629 spatio-temporal atoms $(z_k, u_k)_{k=1}^K$ and a brain atlas $\Theta = (\Theta_m)_{m=1}^M$ (Desikan et al.,

2006) composed of $M = 96$ regions of interest (ROIs). This haemodynamic brain parcellation thus yields 96 HRF dilation parameters $\boldsymbol{\delta} = (\delta_m)_{m=1}^M$ for each individual. In practice, in the definition of $\text{WS}(\cdot, \cdot)$ and $\text{BS}(\cdot, \cdot)$, the true vectors $\boldsymbol{\delta}_{T_j}^{s_i}$ ($i = 1, \dots, 100, j = 1, 2$) have been replaced by their estimates $\widehat{\boldsymbol{\delta}}_{T_j}^{s_i}$ computed by solving Eq. (6) for the two 4-min rs fMRI data sets (T_1 and T_2). To make sure that our conclusions hold for a large scale of temporal regularization parameters, we spanned the range $\lambda_f \in [0, 1]$ and repeated the same procedure over 6 discrete values of λ_f within this interval.

3.4.2. Results

In Fig. 10, the box plot in blue shows the within-subject distance $\text{WS}(\widehat{\boldsymbol{\delta}}_{T_1}^{s_i}, \widehat{\boldsymbol{\delta}}_{T_2}^{s_i})$ between the two 4-min rs-fMRI segments for all individuals and across 6 values of λ_f covering the whole interval $[0, 1]$. The orange and green box plots in Fig. 10 depict the between-subject distances computed over the first and second segments 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 the within-subject (i.e. inter-segment) variability is systematically lower than the between-subject variability and that all metrics remain stable across regularization levels. To go further, we performed a statistical analysis (paired t -test) by comparing the mean of the WS and BS distributions and we obtained significant p -values ($p < 10^{-8}$) showing that the within-subject haemodynamic variability is significantly lower than the between-subject fluctuations. In contrast, the statistical inter-individual comparison between the two segments is not significant ($p \simeq 10^{-2}$). These results are valid for all tested regularization levels indicating a minor impact of the regularization parameter onto the haemodynamic parameter estimate. In sum, this analysis demonstrates that the whole brain characterization of the vascular system remains stable in a given individual between two periods shortly spaced in time, compared to the same analysis between individuals and so that the haemodynamic response discriminates each subject from the others.

4. Clinical validation at the population level

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

667 *4.1. Characterization of patients with an history of stroke*

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

676 We considered 24 patients of both genders and various ages who suffered from
 677 a stroke in the past from the UK Biobank database. For comparison purposes, we
 678 selected 24 healthy controls matched in age and gender from the same database.
 679 We applied the same decomposition ($K = 20$, $M = 96$, same λ_f) to each patient
 680 and healthy control. Fig. 11(a) and Fig. 11(b) show respectively the corresponding
 681 normalized maps of haemodynamic dilation parameters $(\delta_m)_{m=1}^{96}$ in a healthy con-
 682 trol and stroke patient, respectively. The normalization has been done by dividing
 683 all dilation parameter values by their within-subject average, namely $\bar{\delta}^{\text{HC}}$ and $\bar{\delta}^{\text{SP}}$
 684 respectively. We first observed that the dilation parameters were larger in average
 685 in the healthy condition compared to stroke ($\bar{\delta}^{\text{HC}} > \bar{\delta}^{\text{SP}}$). This corresponds to
 686 shorter and more homogeneous TTPs in the brain in the healthy condition. The
 687 shortest TTP found in the healthy control was actually located in the primary
 688 visual cortex (axial slice, $z=-2$, left hemisphere), a result consistent with the lit-
 689 erature on fastest haemodynamic responses often detected in visual areas (Handw-
 690 erker et al., 2004; Badillo et al., 2013). In contrast, Fig. 11(b) illustrates that the
 691 haemodynamic dilation parameters δ^{SP} are smaller (so the TTPs longer) in the
 692 stroke patient Also, we found less variability in the healthy subject since the dif-
 693 ference between the maximum and minimum TTPs were smaller ($\Delta_{\text{TTP}}^{\text{HC}} = 1.25$ s)
 694 compared to the stroke patient ($\Delta_{\text{TTP}}^{\text{SP}} = 2.25$ s). Importantly, Fig. 11(a) illustrates
 695 the relative symmetry of haemodynamic territories that exists in normal subjects
 696 between both hemispheres (Raemaekers et al., 2018).

697 On the contrary, Fig. 11(b) reveals a wider asymmetry between the two hemi-
 698 spheres in the stroke patient. Interestingly, in this patient we noticed the presence
 699 of larger TTPs in the middle left precentral gyrus and left motor cortex (resp.
 700 $z=44$ and $z=60$), namely the brain regions supposedly impacted by the stroke
 701 episode.

702 To go one step further, we quantified the spatial asymmetry of the haemo-
 703 dynamic structure within each individual. To this end, we computed the inter-
 704 hemispheric haemodynamic ℓ_2 distance (IHD) between the HRF dilation param-
 705 eter vectors estimated over the left and right hemispheres in laterally matched brain
 706 regions, respectively denoted δ_{L} and δ_{R} . This normalized intra-subject distance is

707 defined as follows:

$$\text{IHD}(\delta_{\text{R}}^s, \delta_{\text{L}}^s) = \frac{\|\delta_{\text{L}}^s - \delta_{\text{R}}^s\|_2}{\|\delta_{\text{L+R}}^s\|_2}, \quad \forall s = 1, \dots, 24. \quad (8)$$

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

711 By pulling down the values of IHD across all individuals within each group (HC
712 vs SP), we estimated the IHD distributions for the two populations of interest, as
713 shown in Fig. 11(c). In the latter graph, we illustrate how different the two cohorts
714 are in terms of neurovascular asymmetry. The group of 24 stroke patients exhibit
715 larger haemodynamic differences between the ischemic and normal hemispheres.
716 We statistically assessed such difference between the two distributions using a two-
717 sample Kolmogorov-Smirnov test and found a significant p-value ($p = 9.1 \cdot 10^{-5}$).
718 This quantification thus confirmed preliminary visual assessment. We report this
719 p-value with a temporal regularization defined such as $\lambda_f = 0.001$. However, we
720 obtain similar p-value results when using the 5 others levels of temporal regular-
721 ization ($\lambda_f \in [0.001, 0.9]$).

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

726 4.2. Middle-age vs elderly subjects classification

727 In the previous part, we performed group-level statistical analysis in the clas-
728 sical way. In this part, we intend to assess the prediction power of the proposed
729 framework in order to classify middle-age vs elderly subjects using standard ma-
730 chine learning tools (Pedregosa et al., 2011). The reason for choosing this classi-
731 fication task between middle-age and elderly individuals lies first in the fact that
732 multiple studies have pointed out the modification of the haemodynamic system
733 with healthy aging (Ances et al., 2009; Li et al., 2018; West et al., 2019) and
734 second in a regain of interest in the literature for brain age analysis using multi-
735 ple neuroimaging techniques (Engemann et al., 2020). We thus intend to assess
736 whether our approach is able to capture the effect of aging and if so, whether the
737 haemodynamic features reflect more brain aging compared to the neural activity
738 signals.

739 For that purpose, we still relied on the UK Biobank database as in the previous
740 experiments as the short TR (TR=0.735 s) in the rs-fMRI data set (6 min long)
741 provides a suitable setting to investigate the HRF evolution with aging. Here, we

742 selected 459 healthy subjects of both genders and divided them in two balanced
 743 groups: the middle-age (MA: 40-44 yo) and elderly (E: 64-70 yo) groups. We
 744 applied the decomposition (6) to each subject using 5 levels of temporal regular-
 745 ization ($\lambda_f \in [0.001, 0.9]$), $K = 20$ temporal components and we used the same
 746 brain parcellation of $M = 96$ ROIs as before to segregate the HRFs in space.

747 First, we analyzed the haemodynamic differences between the two populations
 748 by computing a two-sample t -test on the distributions of dilation parameters. We
 749 used the temporal regularization parameter $\lambda_f = 0.675$ which is the one selected
 750 through cross validation for our classification model in the subsequent paragraph.
 751 The results were first quantified with t -scores to compare the dilation paramete-
 752 rs in each region. The results are presented in Fig. 12(a) and illustrate large
 753 differences between the two populations. More specifically, higher dilation param-
 754 eters or shorter TTP were retrieved in middle-age subjects as the t -scores were
 755 mostly positive for the comparison $\bar{\delta}_m^{\text{MA}} > \bar{\delta}_m^{\text{E}}$. This is notably visible in the Willis
 756 Polygon, temporal cortices, angular gyri, the medial prefrontal cortices and the su-
 757 perior frontal cortices. To assess the statistical significance, we also computed the
 758 log-transformed p-values, i.e. $-\log_{10} p_{\text{val}}$ (shown in Fig. 12(b)), after correcting
 759 for multiple comparisons using the Bonferroni correction across the $M = 96$ ROIs.
 760 We noticed first that a large majority of significant brain regions appear bilaterally
 761 indicating larger haemodynamic dilation parameters or shorter TTPs in younger
 762 individuals. Second, the negative t -values reported in the cerebellum are not sta-
 763 tistically significant after correcting for multiple comparisons.

764 Then, in an attempt to be exhaustive we constructed three different logistic
 765 regression (LR) models based either on the individual (i) neural activity signals,
 766 (ii) HRF shapes and (iii) haemodynamic dilation parameters. We did not consider
 767 the spatial activity maps as input features in these models as they do not permit to
 768 perform dimension reduction. We trained these LR models using the Scikit-Learn
 769 software (Pedregosa et al., 2011) to predict the age label for each subject (1 for the
 770 elderly people, 0 for middle-aged people). A ℓ_2 -norm regularization was used in
 771 the estimation of the LR model parameters with an hyper-parameter $\beta > 0$. We
 772 grid-searched the temporal regularization parameter λ_f and the classifier hyper-
 773 parameter β (see Fig. E in Supplementary Materials for the stability of the setting).
 774 We chose the accuracy as the classification metric and cross-validated the score to
 775 provide an estimation of the generalization error with a 10 times repeated stratified
 776 4-fold split.

777 Fig. 13(a) shows that the haemodynamic properties have an improved predic-
 778 tion power to discriminate the age compared to the neural activation signals (i.e.
 779 temporal components). The latter actually reaches an average accuracy score of
 780 0.557, whereas the mean accuracy associated with the HRF shape and haemo-
 781 dynamic dilation parameter estimates respectively goes up to 0.741 and 0.743.

782 Also, the distribution of accuracy scores across trials is more concentrated for the
783 HRF dilation parameters compared to the whole HRF shape. This is likely due
784 to the dimension reduction operated to extract this parameter which fluctuates
785 less than the complete profile of the haemodynamic response. This analysis thus
786 demonstrates that our decomposition is able to capture the brain age based on
787 neurovascular information. However, due to the large between-subject variability
788 even within each class of age the neural activation signals do not define a good
789 feature for the brain age prediction. Complementary to that, Fig. 13(b) illustrates
790 the progression of the mean accuracy score with the number of individuals in-
791 volved in the LR model and clearly depicts that a plateau is reached around 459
792 subjects (the total size of the sampled cohort) both for the haemodynamic dila-
793 tion parameter. Also, one can see the rapid progression of the mean accuracy with
794 the number of individuals for the LR model based on haemodynamic properties
795 compared to the one constructed from the neural activation signals.

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

801 5. Discussion

802 *Separating vascular and neuronal components on fMRI BOLD signals.* Both neu-
803 ral and non neural factors such as glial cell activity, cerebral energy metabolism
804 and the cerebral vasculature contribute independently and synergistically to the
805 fMRI BOLD signal. A mis-estimation of individual and regional HRFs may lead
806 to an interpretation that haemodynamic changes as neural variations could have
807 considerable implications for the interpretability and reliability of findings in fMRI
808 studies. Previous literature supports the notion that HRF variability corrupts
809 fMRI data deconvolution (Gitelman et al., 2003; Roebroeck et al., 2011; Seth et al.,
810 2013). Deconvolution-based approaches such as the one presented here allows dis-
811 crimination of cerebrovascular components from neural activity and minimizes
812 the confound of HRF variability in the exploration of brain physiology, functional
813 connectivity and cognitive processes.

814 In this work, we developed a new algorithm that proposed the joint estimation
815 of the HRF and neural activation signal as a semi-blind deconvolution multivariate
816 problem in a paradigm-free setting. Synthetic and real resting-state fMRI data
817 allowed us to demonstrate that this approach is able to faithfully capture the
818 individual’s haemodynamic response function and intrinsic functional networks
819 with low intra-subject variability and relative minimal impact of hyper-parameters
820 on the reliability of HRF estimation. Aims of this paper were also to verify that

821 these methodological developments have practical impacts as related to detection
822 and classification. Further demonstrations using well defined research protocols
823 will be required to refine the tool for use in clinical applications.

824 *Related works.* Disentangling the neurovascular coupling from the neural activity
825 is well documented in the literature (Goutte et al., 2000; Ciuciu et al., 2003;
826 Handwerker et al., 2004; Lindquist and Wager, 2007; Vincent et al., 2010; Chaari
827 et al., 2012; Pedregosa et al., 2015), however many features present in our model are
828 pivotal to making a significant step forward. First, unlike the previous works our
829 approach is paradigm-free and can be applied to both task-related and resting-state
830 fMRI data. As such, the proposed methodology is not restricted to a few brain
831 regions (e.g. visual, auditory and sensorimotor cortices) that are typically involved
832 in a specific activation. Second, in contrast to (Wu et al., 2013; Farouj et al., 2019;
833 Cherkaoui et al., 2019) we propose a multivariate approach to reliably summarize
834 the neurovascular coupling and the neural activity over the whole brain. Third, to
835 the best of knowledge, the proposed semi-blind deconvolution approach is the first
836 to perform such an extensive experimental validation including i) simulation-based
837 numerical experiments for model validation, ii) performance assessment on task-
838 related and resting-state fMRI data at the individual level, including comparisons
839 to (Wu et al., 2013; Pedregosa et al., 2015), and iii) two meaningful cohort-level
840 evaluations in different clinical settings. It is worth mentioning that we used
841 a common brain parcellation – the Harvard-Oxford probabilistic atlas (Desikan
842 et al., 2006) – across all individuals, to ensure group-level analysis and facilitate
843 between-group comparisons. Hence, our main findings on the asymmetries in
844 haemodynamic features between stroke patients and healthy controls on one hand,
845 and the differences related to normal aging on the other hand, are dependent on
846 this atlas and could slightly differ with another parcellation. This question is left
847 for future research.

848 *Interest in identifying the vascular component for clinical applications.* Non-invasive
849 imaging modalities are of undeniable relevance for the diagnostic and prognostic
850 work-up in patients and this present study constitutes a proof of concept in terms
851 of the interest and feasibility of the proposed approach. This type of analysis
852 was made possible due to the fast convergence of our algorithm and the numer-
853 ous optimizations on the implementation side. By making the HemoLearn Python
854 package open source, we offer a unique tool to the neuroimaging community that
855 will permit complementary investigation in other clinical contexts.

856 The current challenge for diagnostic imaging methods is to find metrics that
857 capture relevant information including biomarkers and the present work on haemo-
858 dynamic blind deconvolution might help uncover these measurable indicators. In
859 that context, in agreement with various imaging studies on aging (Bangen et al.,

2009; Shafto et al., 2014), our algorithm, tested on a large cohort of UK BioBank rs-fMRI data sets (459 subjects) has proven its sensitivity to classify middle-age vs elderly individuals with respect to the estimated neurovascular coupling. Vascular aging, as characterized by progressive deterioration in the cellular structure of the blood vessel wall (Handwerker et al., 2004; Havlicek and Uludağ, 2020), undoubtedly impacts both resting-state cerebral blood flow and the ability to modulate it during neuronal activity. Our study confirmed a regional variability in the HRF features with a significantly slower neurovascular coupling in elderly people. Importantly, we demonstrated that the dilation parameter δ , tightly coupled to the haemodynamic delay, achieves better accuracy scores (0.74) compared to the neural activation signals (0.56, just above chance level), suggesting a causal role of cerebral micro-vascular dysregulation due to cognitive aging.

Based on the proposed haemodynamic asymmetry index, our analysis tested in patients with a history of stroke has also proven its utility to individually detect very slow haemodynamic delays in a restricted brain territory probably related to local ischemic tissue consecutive to stroke. This finding is perfectly consistent with the literature (Altamura et al., 2009) showing that the delay in peak latency that arises as patients advance from the acute to the subacute stroke phase is related to the deterioration of cerebral haemodynamics. Diffusion-weighted MRI and perfusion imaging (e.g. ASL) remain the reference imaging modalities as a noninvasive method to perform the diagnosis in the acute episode of stroke. However, in the post-acute period, rs-fMRI acquisition which does not require the patient’s engagement would be extremely valuable to measure the re-establishment of haemodynamic response function in stroke patients as a prediction of subsequent recovery of cerebral function.

Assessment of cerebrovascular function through approaches that employ haemodynamic deconvolution — as presented in this paper — bring new opportunities for the exploration of brain plasticity and pathogenesis. For example, modeling the haemodynamic response function improves sensitivity of fMRI data to delineate epileptogenic area (Storti et al., 2013). Introducing some flexibility in the HRF shape is of undeniable benefit to consider ‘silent’ vascular risk factors in the pathogenesis or exacerbation of neurological diseases (D’Esposito et al., 2003; Ogaki et al., 2020) or in pharmacological fMRI studies (phMRI) where drugs, used as stimuli, do not always elicit predictable changes in neurovascular coupling (Cherkaoui et al., 2021).

Limitations and extensions. Some limitations of our tool do exist. First, there are free parameters in the proposed modeling (K , λ_f , M) that need to be set in an appropriate manner. We explored two model selection criteria for setting K , namely the R^2 score and the determinant of the correlation matrix between the neural activation signals. Based on these metrics, we found a fair compromise

900 between accuracy and model complexity for $K = 20$. We thus constantly used this
901 value hereafter in the individual decomposition. Of course, other model selection
902 approaches might be envisaged to optimize K and λ_f using for instance a (widely)
903 Bayesian information criterion (Neath and Cavanaugh, 2012; Watanabe, 2013), or
904 the log-likelihood in the standard classical framework. The selected model would
905 thus be the one associated with the lowest BIC value or largest log-likelihood. More
906 recently, the concept of bi-level optimization (Bennett et al., 2006) has emerged to
907 set hyper-parameters. In this case, an upper-level cost function (e.g. a supervised
908 training score on the features of the decomposition) has to be minimized with
909 respect to the unknown hyper-parameters while staying intrinsically connected to
910 the lower-level problem, namely the multivariate decomposition. Because of the
911 extra-computation cost required by these approaches, such aspects are beyond
912 the scope of this paper. Second, the proposed regional analysis is conditioned by
913 the parcellation atlas (and the value of M). It would be interesting to deepen this
914 research by testing the reproducibility of the tool with some atlas variations and the
915 creation of an atlas using subject-specific assessment of the cerebral vasculature.

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

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

938 Sixth, because the proposed HRF model relies solely on a time dilation pa-
939 rameter, its magnitude is fixed and the fluctuations of the BOLD signal across

940 the brain are thus captured through the neural activity atoms $(\mathbf{z}_k)_{k=1}^K$ on one
 941 hand and the spatial maps $(\mathbf{u}_k)_{k=1}^K$ on the other hand. However, the norm
 942 of the spatial maps being constrained, the real BOLD signal amplitude is cap-
 943 tured by the neural activation signals. A recent work [Tsvetanov et al. \(2019\)](#)
 944 has shown that the resting-state fluctuation amplitude is crucial to predict brain
 945 age in healthy subjects. One possible enhancement of the current model would
 946 be to add a magnitude parameter to each HRF. In that case, we should fix the
 947 scale ambiguity issue by setting the amplitude of the temporal atoms $(\mathbf{z}_k)_{k=1}^K$.
 948 This modification would significantly increase the computational complexity due
 949 to the calculation of the proximal operator associated with the new regulariza-
 950 tion term $g_z((\mathbf{z}_k)_k) = \lambda \sum_{k=1}^K (\|\nabla \mathbf{z}_k\|_1 + I_{\|\mathbf{z}_k\|_\infty = \alpha})$. Last, thus far we have used
 951 the canonical HRF as the reference shape in \mathbf{v}_{ref} . This setting could be easily
 952 updated to perform investigations in specific populations (e.g. newborns) where
 953 the true haemodynamic response function is known to deviate from the canonical
 954 shape ([Arichi et al., 2012](#)).

955 6. Conclusion

956 In this paper, we have presented a semi-blind deconvolution approach to jointly
 957 estimate the haemodynamic response function and the neural activity signals
 958 across the whole brain. As the proposed methodology is paradigm-free, it en-
 959 ables the analysis of resting-state fMRI data in an semi-supervised manner as the
 960 regularization parameters (K, λ_f) may be tuned using a trade-off between model
 961 accuracy and complexity. Beyond the model validation on synthetic and real fMRI
 962 data, we have demonstrated the interest of the proposed approach in two appli-
 963 cations in neuroscience. Both aimed at characterizing cerebral haemodynamic
 964 delays in specific populations, namely stroke patients and elderly people by con-
 965 trasting them with healthy and younger controls. Most importantly, we proposed
 966 an haemodynamic asymmetry index to lateralize the stroke episode while con-
 967 firming the presence of a prolonged haemodynamic delay in these patients. We
 968 also demonstrated that haemodynamic properties are predictable of brain age. Fi-
 969 nally, this new framework opens the door to new research avenues for functional
 970 connectivity analysis based on the neural input signals instead of the BOLD sig-
 971 nal themselves. In contrast to existing techniques ([Wu et al., 2013](#)), our approach
 972 would be less biased by a constant haemodynamic response shape across the whole
 973 brain.

974 Acknowledgment

975 This work was supported by a CEA PhD scholarship, the UK Royal Academy of
 976 Engineering under the RF/201718/17128 grant and the SRPe PECRE 1718/15

977 Award. We would like to thank our colleagues Bertrand Thirion, Kamalaker
978 Reddy Dadi and Thomas Bazeille from Inria for fruitful discussions that helped
979 us investigate the proposed data analyses. Finally we are grateful to Cécilia Gar-
980 rec (CEA/DRF-Dir) for her comments and suggestions that have significantly
981 improved the manuscript.

982 **References**

- 983 S. Ogawa, D. W. Tank, R. Menon, J. M. Ellerman, S. G. Kim, H. Merkle,
984 K. Ugurbil, Intrinsic signal changes accompanying sensory stimulation: func-
985 tional brain mapping with magnetic resonance imaging., *Proceedings of the*
986 *National Academy of Sciences* 89 (1992) 5951–5955.
- 987 P. A. Bandettini, A. Jesmanowicz, E. C. Wong, J. S. Hyde, Processing strategies
988 for time-course data sets in functional MRI of the human brain., *Magnetic*
989 *resonance in medicine* 30 (1993) 161–173.
- 990 G. M. Boynton, S. A. Engel, G. H. Glover, D. J. Heeger, Linear systems analysis of
991 functional magnetic resonance imaging in human V1., *Journal of Neuroscience*
992 16 (1996) 4207–4221.
- 993 C.-T. Do, Z.-M. Manjaly, J. Heinzle, D. Schöbi, L. Kasper, K. P. Pruessmann,
994 K. E. Stephan, S. Frässle, Hemodynamic modeling of aspirin effects on bold
995 responses at 7t., *medRxiv* (2020).
- 996 K. L. West, M. D. Zuppichini, M. P. Turner, D. K. Sivakolundu, Y. Zhao, D. Ab-
997 delkarim, J. S. Spence, B. Rypma, BOLD hemodynamic response function
998 changes significantly with healthy aging., *Neuroimage* 188 (2019) 198–207.
- 999 D. Asemani, H. Morsheddost, M. A. Shalchy, Effects of ageing and alzheimer
1000 disease on haemodynamic response function: a challenge for event-related fMRI.,
1001 *Healthcare Technology Letters* 4 (2017) 109–114.
- 1002 K. J. Friston, P. Fletcher, O. Josephs, A. Holmes, M. D. Rugg, R. Turner, Event-
1003 related fMRI: characterizing differential responses., *Neuroimage* 7 (1998) 30–40.
- 1004 P. Ciuciu, J.-B. Poline, G. Marrelec, J. Idier, C. Pallier, H. Benali, Unsupervised
1005 robust nonparametric estimation of the hemodynamic response function for any
1006 fMRI experiment., *IEEE transactions on Medical Imaging* 22 (2003) 35–51.
- 1007 M. A. Lindquist, T. D. Wager, Validity and power in hemodynamic response
1008 modeling: a comparison study and a new approach., *Human brain mapping* 28
1009 (2007) 764–784.

- 1010 F. Pedregosa, M. Eickenberg, P. Ciuciu, B. Thirion, A. Gramfort, Data-driven
1011 HRF estimation for encoding and decoding models., *NeuroImage* (2015) 209–
1012 220.
- 1013 C. Goutte, F. A. Nielsen, L. K. Hansen, Modeling the haemodynamic response
1014 in fMRI using smooth FIR filters., *IEEE transactions on Medical Imaging* 19
1015 (2000) 1188–1201.
- 1016 T. Vincent, T. Rissern, P. Ciuciu, Spatially adaptive mixture modeling for analysis
1017 of fMRI time series, *IEEE transactions on Medical Imaging* 29 (2010) 59–74.
- 1018 L. Chaari, L. Forbes, T. Vincent, P. Ciuciu, Hemodynamic-informed parcella-
1019 tion of fMRI data in a joint detection estimation framework, *proceedings of*
1020 *International Conference on Medical Image Computing and Computer-Assisted*
1021 *Intervention* 15 (2012) 180–188.
- 1022 D. Degras, M. Lindquist, A hierarchical model for simultaneous detection and
1023 estimation in multi-subject fMRI studies, *NeuroImage* 98 (2014).
- 1024 P. N. Rosa, P. Figueiredo, C. J. Silvestre, On the distinguishability of hrf models
1025 in fMRI, *Frontiers in Computational Neuroscience* 9 (2015) 54.
- 1026 M. Eickenberg, A. Frau-Pascual, A. Hoyos-Idrobo, Gaussian processes for hrf
1027 estimation for bold fmri, *arXiv: Applications* (2017).
- 1028 G. R. Wu, W. Liao, S. Stramaglia, J. R. Ding, H. Chen, D. Marinazzo, A blind
1029 deconvolution approach to recover effective connectivity brain networks from
1030 resting state fMRI data., *Medical Image Analysis* 17 (2013) 365–374.
- 1031 G. H. Glover, Deconvolution of impulse response in event-related bold fMRI.,
1032 *Neuroimage* 9 (1999) 416–429.
- 1033 D. R. Gitelman, W. D. Penny, J. Ashburner, K. J. Friston, Modeling regional
1034 and psychophysiologic interactions in fMRI: the importance of hemodynamic
1035 deconvolution., *Neuroimage* 19 (2003) 200–207.
- 1036 C. Caballero-Gaudes, N. Petridou, I. L. Dryden, L. Bai, S. T. Francis, P. A.
1037 Gowland, Detection and characterization of single-trial fMRI bold responses:
1038 Paradigm free mapping, *Human brain mapping* 32 (2011) 1400–1418.
- 1039 L. Hernandez-Garcia, M. O. Ulfarsson, Neuronal event detection in fMRI time
1040 series using iterative deconvolution techniques., *Magnetic resonance imaging* 29
1041 (2011) 353–364.

- 1042 C. Caballero-Gaudes, D. Van de Ville, N. Petridou, F. Lazeyras, P. Gowland,
1043 Paradigm-free mapping with morphological component analysis: getting most
1044 out of fMRI data, *Optical Engineering + Applications* (2011).
- 1045 C. Caballero-Gaudes, F. I. Karahanoglu, F. Lazeyras, D. Van De Ville, Struc-
1046 tured sparse deconvolution for paradigm free mapping of functional MRI data.,
1047 *International Symposium on Biomedical Imaging* 9 (2012) 322–325.
- 1048 C. Caballero-Gaudes, N. Petridou, S. T. Francis, I. L. Dryden, P. A. Gowland,
1049 Paradigm free mapping with sparse regression automatically detects single-trial
1050 functional magnetic resonance imaging blood oxygenation level dependent re-
1051 sponses, *Human Brain Mapping* 34 (2013) 501–518.
- 1052 C. Caballero-Gaudes, P. Bandettini, J. Gonzalez-Castillo, A temporal deconvo-
1053 lution algorithm for multiecho functional MRI, *International Symposium on*
1054 *Biomedical Imaging* (2018) 608–611.
- 1055 I. Khalidov, D. Van De Ville, J. Fadili, M. Unser, Activelets and sparsity: a new
1056 way to detect brain activation from fmri data., *International Society for Optics*
1057 *and Photonics* 6701 (2007).
- 1058 I. Khalidov, J. Fadili, F. Lazeyras, D. Van De Ville, M. Unser, Activelets: Wavelets
1059 for sparse representation of hemodynamic responses., *Signal processing* 91 (2011)
1060 2810–2821.
- 1061 F. I. Karahanoglu, C. Caballero-Gaudes, F. Lazeyras, D. Van De Ville, Total
1062 activation: fMRI deconvolution through spatio-temporal regularization., *Neu-*
1063 *roImage* 73 (2013) 121–134.
- 1064 N. Petridou, C. Caballero-Gaudes, I. L. Dryden, F. S. T., P. Gowland, Periods of
1065 rest in fMRI contain individual spontaneous events which are related to slowly
1066 fluctuating spontaneous activity., *Hum Brain Mapp* (2013) 1319–1329.
- 1067 F. I. Karahanoglu, D. Van De Ville, Transient brain activity disentangles fmri
1068 resting-state dynamics in terms of spatially and temporally overlapping net-
1069 works., *Nat Commun* 6 (2015) 7751.
- 1070 J. Gonzalez-Castillo, C. Caballero-Gaudes, N. Topolski, D. Handwerker,
1071 F. Pereira, P. Bandettini, Imaging the spontaneous flow of thought: Distinct
1072 periods of cognition contribute to dynamic functional connectivity during rest,
1073 *NeuroImage* 202 (2019).
- 1074 H. Cherkaoui, T. Moreau, A. Halimi, P. Ciuciu, Sparsity-based Semi-Blind Decon-
1075 volution of Neural Activation Signal in fMRI., *IEEE International Conference*
1076 *on Acoustics, Speech and Signal Processing (ICASSP)* (2019) 1323–1327.

- 1077 Y. Farouj, F. I. Karahanoglu, D. Van de Ville, Bold signal deconvolution under
1078 uncertain haemodynamics: A semi-blind approach., International Symposium
1079 on Biomedical Imaging (2019) 1792–1796.
- 1080 H. Cherkaoui, T. Moreau, A. Halimi, P. Ciuciu, fMRI BOLD signal decomposition
1081 using a multivariate low-rank model., European Signal Processing Conference
1082 (EUSIPCO) (2019) 1–5.
- 1083 R. Grosse, R. Raina, H. Kwong, Y. N. Andrew, Shift-invariant sparse coding
1084 for audio classification., proceedings of Uncertainty in Artificial Intelligence 23
1085 (2007) 149–158.
- 1086 T. Dupré La Tour, T. Moreau, M. Jas, A. Gramfort, Multivariate convolutional
1087 sparse coding for electromagnetic brain signals., Advances in Neural Information
1088 Processing System (NeurIPS) (2018) 3292–3302.
- 1089 M. W. Woolrich, B. D. Ripley, M. Brady, S. M. Smith, Temporal autocorrelation
1090 in univariate linear modeling of fMRI data., Neuroimage 14 (2001) 1370–1386.
- 1091 A. M. Dale, Optimal experimental design for event-related fMRI, Human Brain
1092 Mapping 2 (1999) 109–114.
- 1093 R. Casanova, S. Ryali, J. Serences, L. Yang, R. Kraft, P. J. Laurienti, J. A. Mald-
1094 jian, The impact of temporal regularization on estimates of the bold hemo-
1095 dynamic response function: a comparative analysis., NeuroImage 40 (2008)
1096 1606–1618.
- 1097 C. M. Zhang, Y. Jiang, T. Yu, A comparative study of one-level and two-level
1098 semiparametric estimation of hemodynamic response function for fMRI data.,
1099 Statistics in medicine 26 (2007) 3845–3861.
- 1100 V. A. Vakorin, R. Borowsky, G. E. Sarty, Characterizing the functional MRI
1101 response using tikhonov regularization., Statistics in medicine 26 (2007) 3830–
1102 3844.
- 1103 M. W. Woolrich, T. E. J. Behrens, S. M. Smith, Constrained linear basis sets for
1104 hrf modelling using variational bayes., NeuroImage 21 (2004) 1748–1761.
- 1105 M. A. Lindquist, J. Meng Loh, L. Y. Atlas, T. D. Wager, Modeling the hemody-
1106 namic response function in fMRI: efficiency, bias and mis-modeling., NeuroImage
1107 45 (2009) S187–S198.
- 1108 D. A. Handwerker, J. M. Ollinger, M. D’Esposito, Variation of bold hemody-
1109 namic responses across subjects and brain regions and their effects on statistical
1110 analyses., Neuroimage 21 (2004) 1639–1651.

- 1111 S. Badillo, T. Vincent, P. Ciuciu, Group-level impacts of within-and between-
1112 subject hemodynamic variability in fMRI., *Neuroimage* 82 (2013) 433–448.
- 1113 G. Varoquaux, R. C. Craddock, Learning and comparing functional connectomes
1114 across subjects., *NeuroImage* 80 (2013) 405–415.
- 1115 T. Vincent, P. Ciuciu, B. Thirion, Sensitivity analysis of parcellation in the joint
1116 detection-estimation of brain activity in fMRI, *International Symposium on*
1117 *Biomedical Imaging* (2008) 568–571.
- 1118 R. S. Desikan, F. Segonne, B. Fischl, B. T. Quinn, B. C. Dickerson, D. Blacker,
1119 R. L. Buckner, A. M. Dale, R. P. Maguire, B. T. Hyman, M. S. Albert, R. J.
1120 Killiany, An automated labeling system for subdividing the human cerebral
1121 cortex on MRI scans into gyral based regions of interest., *Neuroimage* 31 (2006)
1122 968–980.
- 1123 M. Najafi, B. W. McMenamin, J. Z. Simon, L. Pessoa, Overlapping communities
1124 reveal rich structure in large-scale brain networks during rest and task condi-
1125 tions., *Neuroimage* 135 (2016) 92–106.
- 1126 T. W. Allan, S. T. Francis, C. Caballero-Gaudes, P. G. Morris, E. B. Liddle,
1127 P. F. Liddle, ..., P. A. Gowland, Functional connectivity in MRI is driven by
1128 spontaneous BOLD events., *PloS one* 10 (2015).
- 1129 R. F. Betzel, M. Fukushima, Y. He, X. N. Zuo, O. Sporns, Dynamic fluctuations
1130 coincide with periods of high and low modularity in resting-state functional
1131 brain networks., *NeuroImage* 127 (2016) 287–297.
- 1132 E. Urunuela, S. Jones, A. Crawford, W. Shin, S. Oh, M. Lowe, C. Caballero-
1133 Gaudes, Stability-based sparse paradigm free mapping algorithm for deconvolu-
1134 tion of functional MRI data., *Eng Med Biol Soc* (2020) 1092–1095.
- 1135 I. Daubechies, E. Roussos, S. Takerkart, M. Benharrosh, C. Golden, K. D’Ardenne,
1136 W. Richter, J. D. Cohen, J. Haxby, Independent component analysis for brain
1137 fMRI does not select for independence, *Proceedings of the National Academy*
1138 *of Sciences* 106 (2009) 10415–10422.
- 1139 T. Deneux, O. Faugeras, EEG-fMRI fusion of paradigm-free activity using kalman
1140 filtering, *Neural Computation* 22 (2010) 906–948.
- 1141 B. O’Donoghue, E. Candes, Adaptive restart for accelerated gradient schemes.,
1142 *Foundations of Computational Mathematics* 15 (2015) 715–732.
- 1143 J. Mairal, F. Bach, J. Ponce, G. Sapiro, Online dictionary learning for sparse
1144 coding., *International Conference on Machine Learning* 26 (2009) 689–696.

- 1145 A. Beck, M. Teboulle, A Fast Iterative Shrinkage-Thresholding Algorithm for
1146 Linear Inverse Problems., *SIAM Journal on Imaging Sciences* 2 (2009) 183–202.
- 1147 J. M. Bioucas-Dias, M. A. Figueiredo, A new TwIST: Two-step iterative shrink-
1148 age/thresholding algorithms for image restoration, *IEEE Transactions on Image*
1149 *processing* 16 (2007) 2992–3004.
- 1150 A. Chambolle, An Algorithm for Total Variation Minimization and Applications.,
1151 *Journal of Mathematical Imaging and Vision* 20 (2004) 89–97.
- 1152 A. Barbero, S. Sra, Modular proximal optimization for multidimensional total-
1153 variation regularization., *Journal of Machine Learning Research* 19 (2018) 1–89.
- 1154 L. Condat, Fast projection onto the simplex and the l_1 ball., *Mathematical*
1155 *Programming Series A* 158 (2016) 575–585.
- 1156 F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel,
1157 M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos,
1158 D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine
1159 learning in Python., *Journal of Machine Learning Research* 12 (2011) 2825–2830.
- 1160 M. Jenkinson, P. Bannister, M. Brady, S. Smith, Improved optimization for the
1161 robust and accurate linear registration and motion correction of brain images.,
1162 *NeuroImage* 17 (2002) 825–41.
- 1163 V. Menon, Large-scale functional brain organization., *Brain mapping: An ency-*
1164 *clopedic reference* 2 (2015) 449–459.
- 1165 D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. J. Behrens, E. Yacoub, K. Ugur-
1166 bil, The wu-minn Human Connectome Project: An overview., *NeuroImage* 80
1167 (2013) 62–79.
- 1168 M. Glasser, S. Sotiropoulos, J. Wilson, T. Coalson, B. Fischl, J. Andersson, J. Xu,
1169 S. Jbabdi, M. Webster, J. Polimeni, V. DC, M. Jenkinson, The minimal pre-
1170 processing pipelines for the human connectome project, *NeuroImage* 80 (2013)
1171 105.
- 1172 P. H. England, Briefing document: First incidence of stroke estimates for england
1173 2007 to 2016., *Public Health England publication* (2018).
- 1174 S. N. Min, S. J. Park, D. J. Kim, M. Subramaniam, K. S. Lee, Development of
1175 an algorithm for stroke prediction: A national health insurance database study
1176 in korea., *European Neurology* 79 (2018) 214–220.

- 1177 M. Raemaekers, W. Schellekens, N. Petridou, N. F. Ramsey, Knowing left from
1178 right: asymmetric functional connectivity during resting state., *Brain Structure*
1179 *and Function* 223 (2018) 1909–1922.
- 1180 B. M. Ances, C. L. Liang, O. Leontiev, J. E. Perthen, A. S. Fleisher, A. E. Lansing,
1181 R. B. Buxton, Effects of aging on cerebral blood flow, oxygen metabolism, and
1182 blood oxygenation level dependent responses to visual information stimulation.,
1183 *Human Brain Mapping* 30 (2009) 1120–1132.
- 1184 Y. Li, W. J. Choi, W. Wei, S. Song, Q. Zhang, J. Liu, , R. K. Wang, Aging-
1185 associated changes in cerebral vasculature and blood flow as determined by
1186 quantitative optical coherence tomography angiography., *Neurobiol Aging* 70
1187 (2018) 148–159.
- 1188 D. A. Engemann, O. Kozynets, D. Sabbagh, G. Lemaître, G. Varoquaux, F. Liem,
1189 A. Gramfort, Combining magnetoencephalography with magnetic resonance
1190 imaging enhances learning of surrogate-biomarkers., *Elife* 9 (2020).
- 1191 C. L. Grady, D. D. Garrett, Understanding variability in the bold signal and why
1192 it matters for aging., *Brain Imaging Behaviour* 8 (2014) 274–283.
- 1193 A. Roebroeck, E. Formisano, R. Goebel, The identification of interacting net-
1194 works in the brain using fMRI: model selection, causality and deconvolution.,
1195 *Neuroimage* 58 (2011) 296–302.
- 1196 A. K. Seth, P. Chorley, L. C. Barnett, Granger causality analysis of fMRI BOLD
1197 signals is invariant to hemodynamic convolution but not downsampling., *Neu-
1198 roimage* 65 (2013) 540–555.
- 1199 K. J. Bangen, K. Restom, T. T. Liu, A. J. Jak, C. E. Wierenga, D. P. Salmon,
1200 M. W. Bondi, Differential age effects on cerebral blood flow and BOLD response
1201 to encoding: associations with cognition and stroke risk., *Neurobiology of aging*
1202 30 (2009) 1276–1287.
- 1203 M. A. Shafto, L. K. Tyler, M. Dixon, J. R. Taylor, J. B. Rowe, R. Cusack, A. J.
1204 Calder, W. D. Marslen-Wilson, J. Duncan, T. Dalgleish, et al., The cambridge
1205 centre for ageing and neuroscience (Cam-CAN) study protocol: a cross-sectional,
1206 lifespan, multidisciplinary examination of healthy cognitive ageing., *BMC neu-
1207 rology* 14 (2014) 204.
- 1208 M. Havlicek, K. Uludağ, A dynamical model of the laminar BOLD response.,
1209 *NeuroImage* 204 (2020) 116–209.

- 1210 C. Altamura, M. Reinhard, M.-S. Vry, C. P. Kaller, F. Hamzei, F. Vernieri, P. M.
1211 Rossini, A. Hetzel, C. Weiller, D. Saur, The longitudinal changes of BOLD
1212 response and cerebral hemodynamics from acute to subacute stroke. A fMRI
1213 and TCD study., *BMC neuroscience* 10 (2009) 151.
- 1214 S. F. Storti, E. Formaggio, A. Bertoldo, P. Manganotti, A. Fiaschi, G. M. Toffolo,
1215 Modelling hemodynamic response function in epilepsy., *Clinical Neurophysiol-*
1216 *ogy* 124 (2013) 2108–2118.
- 1217 M. D’Esposito, L. Y. Deouell, A. Gazzaley, Alterations in the BOLD fMRI sig-
1218 nal with ageing and disease: a challenge for neuroimaging., *Nature Reviews*
1219 *Neuroscience* 4 (2003) 863–872.
- 1220 A. Ogaki, Y. Ikegaya, R. Koyama, Vascular abnormalities and the role of vascular
1221 endothelial growth factor in the epileptic brain, *Frontiers in Pharmacology* 11
1222 (2020).
- 1223 H. Cherkaoui, T. Moreau, P. Ciuciu, B. Fernandez, M. Bottlaender, N. Tournier,
1224 C. Leroy, Characterization of the haemodynamic response function after a
1225 buprenorphine challenge study in human healthy volunteer, in: 27th, *Virtual*.
- 1226 A. A. Neath, J. E. Cavanaugh, The Bayesian information criterion: background,
1227 derivation, and applications., *Wiley Interdisciplinary Reviews: Computational*
1228 *Statistics* 4 (2012) 199–203.
- 1229 S. Watanabe, A widely applicable Bayesian information criterion., *Journal of*
1230 *Machine Learning Research* 14 (2013) 867–897.
- 1231 K. P. Bennett, J. Hu, X. Ji, G. Kunapuli, J.-S. Pang, Model selection via bilevel
1232 optimization., *The 2006 IEEE International Joint Conference on Neural Network*
1233 *Proceedings* (2006) 1922–1929.
- 1234 A. de Pierrefeu, T. Löfstedt, F. Hadj-Selem, M. Dubois, R. Jardri, T. Fovet,
1235 P. Ciuciu, V. Frouin, E. Duchesnay, Structured sparse principal components
1236 analysis with the TV-elastic net penalty., *IEEE transactions on medical imaging*
1237 *37* (2017) 396–407.
- 1238 R. Jenatton, A. Gramfort, V. Michel, G. Obozinski, E. Eger, F. Bach, B. Thirion,
1239 Multiscale mining of fMRI data with hierarchical structured sparsity., *SIAM*
1240 *Journal on Imaging Sciences* 5 (2012) 835–856.
- 1241 L. Baldassarre, J. Mourao-Miranda, M. Pontil, Structured sparsity models for
1242 brain decoding from fMRI data., *2012 Second International Workshop on Pat-*
1243 *tern Recognition in NeuroImaging* (2012) 5–8.

- 1244 H. Cherkaoui, J. Sulam, T. Moreau, Learning to solve tv regularised problems with
1245 unrolled algorithms., 34th Conference and Workshop on Neural Information
1246 Processing Systems (NeurIPS) (2020) 1–21.
- 1247 V. D. Calhoun, J. Liu, T. Adalı, A review of group ICA for fMRI data and ICA
1248 for joint inference of imaging, genetic, and erp data., *Neuroimage* 45 (2009)
1249 S163–S172.
- 1250 G. Varoquaux, S. Sadaghiani, P. Pinel, A. Kleinschmidt, J. B. Poline, B. Thirion,
1251 A group model for stable multi-subject ICA on fMRI datasets., *NeuroImage* 51
1252 (2010) 288–299.
- 1253 G. Varoquaux, A. Gramfort, F. Pedregosa, V. Michel, B. Thirion, Multi-subject
1254 dictionary learning to segment an atlas of brain spontaneous activity., *Biennial
1255 International Conference on information processing in medical imaging* (2011)
1256 562–573.
- 1257 M. Milham, D. Fair, M. Mennes, S. Mostofsky, The adhd-200 consortium: a model
1258 to advance the translational potential of neuroimaging in clinical neuroscience.,
1259 *Frontiers in Systems Neuroscience* 6 (2012) 62.
- 1260 D. A. Feinberg, K. Setsompop, Ultra-fast MRI of the human brain with simulta-
1261 neous multi-slice imaging., *Journal of magnetic resonance* 229 (2013) 90–100.
- 1262 J. Hesamoddin, S. Holdsworth, T. Christen, H. Wu, K. Zhu, A. B. Kerr, M. J.
1263 Middione, R. F. Dougherty, M. Moseley, G. Zaharchuk, Advantages of short
1264 repetition time resting-state functional MRI enabled by simultaneous multi-slice
1265 imaging., *Journal of Neuroscience Methods* 311 (2019) 122–132.
- 1266 K. Tsvetanov, R. N. A. Henson, P. S. Jones, H. J. Mutsaerts, D. Fuhrmann, L. K.
1267 Tyler, , J. B. Rowe, The effects of age on resting-state bold signal variability is
1268 explained by cardiovascular and neurovascular factors., preprint bioRxiv (2019).
- 1269 T. Arichi, G. Fagiolo, M. Varela, A. Melendez-Calderon, A. Allievi, N. Merchant,
1270 N. Tumor, S. J. Counsell, E. Burdet, C. F. Beckmann, et al., Development
1271 of BOLD signal hemodynamic responses in the human brain., *Neuroimage* 63
1272 (2012) 663–673.




Figure 4: **Comparison to (Wu et al., 2013) on synthetic data:** (a): Firstly, we consider the case when the HRF is set to a different shape than the one used for simulating the data. (b): Secondly, we learn the HRF along with the spatio-temporal components (a/b)-top: The yellow-purple maps define the spatial activation patterns. (a/b)-bottom-left: The blue curve represents the true HRF, the green trace corresponds to the initialization and the dashed orange shape corresponds to the HRF used to deconvolve the observed BOLD signal. (a/b)-bottom-right: The observed BOLD signal in the associated activation region is depicted in black, whereas the true and recovered temporal atoms are plotted in blue and orange, respectively. (a)-bottom-right: We add, in dashed green, the estimate of the temporal atoms if we choose the true HRF as the initialization. The standard deviation across voxels is encoded by transparency around mean curves. (c): We report the result of the blind-deconvolution of (Wu et al., 2013) for comparison. (c)-bottom-left: We display in green the histogram of the haemodynamic delays estimated with this concurrent approach for the activated voxels and reports the associated average delay using a green vertical line. The true haemodynamic delay is plotted in blue while the one estimated by our approach is shown in orange. (c)-bottom-right: For the voxel localized in the yellow-purple map (bottom - far right) we display the true neural activation signal in blue and the one estimated by our approach in orange along with the deconvolved BOLD signal obtained with (Wu et al., 2013) (bottom). On top of that panel, in order to quantify how accurate in time the localization of neural events is, we plot the first-order derivative of the true neural activation signal and our estimate with the same color code along with the detected events from Wu et al. (2013)'s method in green.




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.

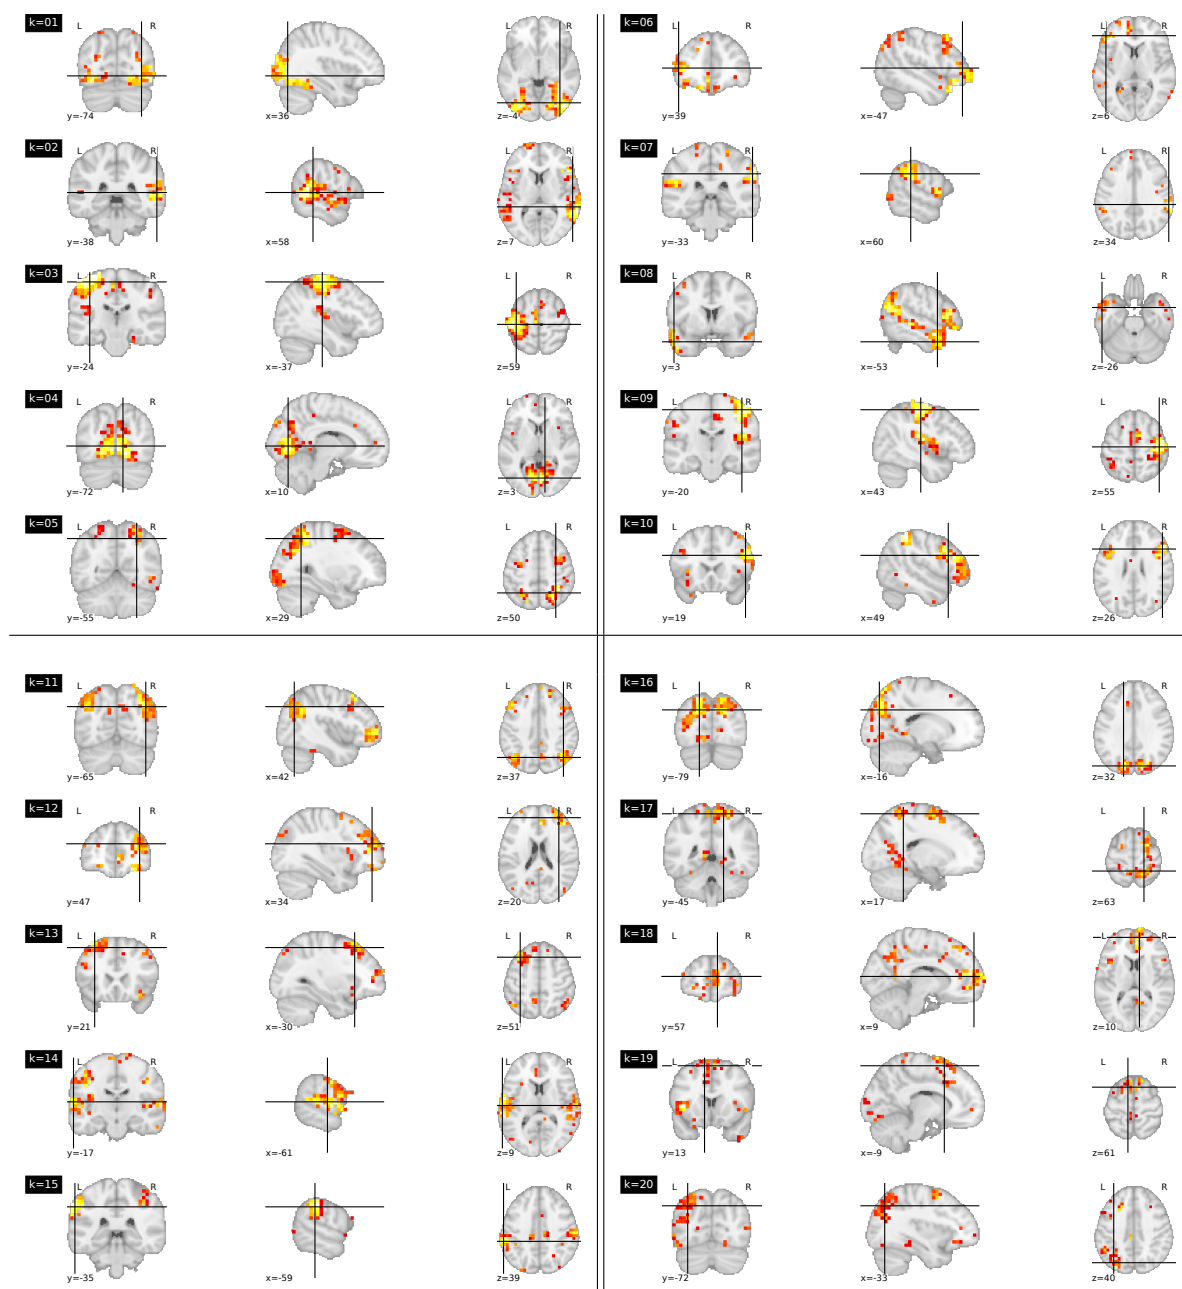


Figure 6: **Spatial decomposition of rs-fMRI data for $K = 20$.** From top to bottom and left to right, the twenty labeled spatial maps are shown using the three orthogonal views (coronal on the left, sagittal in the middle and axial on the right). The labeling is arbitrary and the coordinates are given in the MNI space.




Figure 7: **Single subject results from rs-fMRI semi-blind deconvolution analysis.** **Top Row (a):** Neural activation signal \hat{z}_1 (left) and corresponding spatial map (axial view) \hat{u}_1 (right), mostly involving activated voxels in the visual cortex. **Second row (b):** Neural activation signal \hat{z}_{11} (left) and corresponding spatial map (axial view) \hat{u}_{11} (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) localized 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) localized 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 in 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: **Comparison to (Pedregosa et al., 2015) on task-fMRI data:** We consider the experimental condition of the *left-hand* motor action. **Top (a)-(b):** We localize the 100 most activated voxels w.r.t. the condition. **(a)** corresponds to the activated voxels estimated from the thresholded regression coefficients as done in (Pedregosa et al., 2015) and **(b)** corresponds to the activated voxels in our approach obtained by thresholding the spatial map and keeping the largest values. **Center (c)-(d):** We depict the axial slices of the corresponding activation maps for each approach ((c) refers to our competitor, (d) to our method). We observe that Pedregosa et al. (2015) mainly identify the activation in the cerebellum whereas our approach localizes the activation in the right motor cortex. **Bottom (e):** We display the neural activation time courses estimated by our approach in orange for the selected voxels. The standard deviation across voxels is encoded by transparency around the mean curve. Similarly, we report the observed BOLD signal in blue and display the temporal profile of the motor task in light gray. **Bottom (f):** We display the HRF estimates for each method. The green HRF corresponds to (Pedregosa et al., 2015) with the between-voxel standard deviation encoded by transparency around the mean curve. Similarly, we depict our HRF estimate in orange. We notice larger haemodynamic delays estimated by our method.




Figure 10: **Within- 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 6 levels of temporal regularization ($\lambda_f \in \{10^{-6}, 10^{-3}, 10^{-2}, 10^{-1}, 5 \cdot 10^{-1}, 1 - 10^{-6}\}$) 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 11: **Haemodynamic discrimination between stroke patients (SP) and healthy controls (HC).** **Top (a)- Middle (b):** Normalized haemodynamic dilation parameter maps in a healthy control (top row) and stroke patient (bottom row), 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 $s \in \{\text{HC}, \text{SP}\}$. Larger haemodynamic dilation parameters and thus shorter TTPs are retrieved in healthy condition ($\bar{\delta}^{\text{HC}} = 0.87 > \bar{\delta}^{\text{SP}} = 0.77$). **Bottom (c):** Histograms of the normalized inter-hemispheric haemodynamic distance (NIHD) 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 ($p = 9.1 \cdot 10^{-5}$), which is associated with a two-sample Kolmogorov-Smirnov test between the two distributions, demonstrates that the neurovascular asymmetry in SP, quantified by NIHD, is significantly different and actually larger than the one observed in HC.

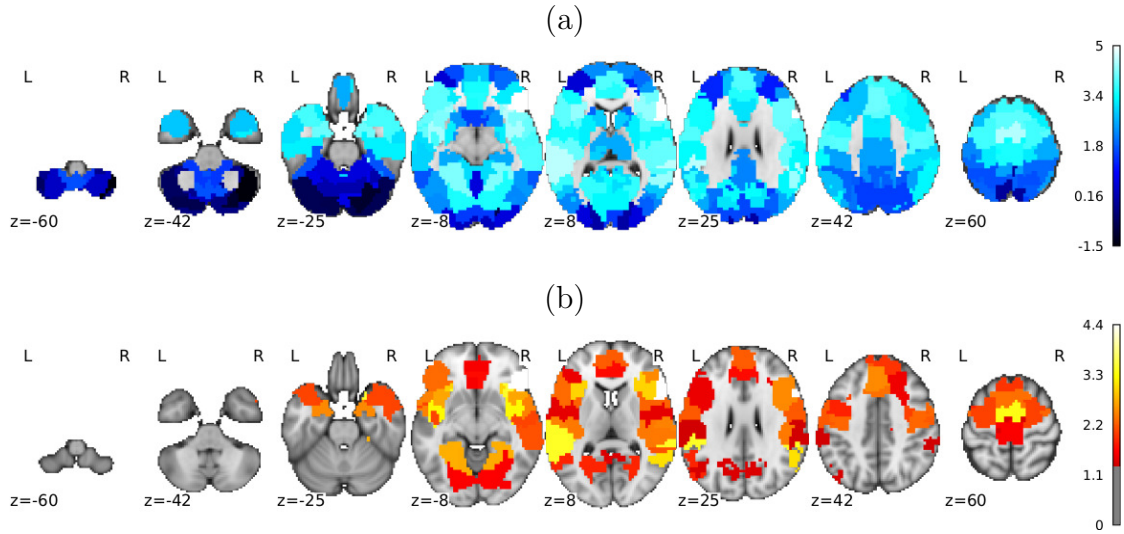


Figure 12: **Statistical analysis of the haemodynamic differences between middle-age (MA) and elderly (E) subjects.** (a): T-scores associated with the two-sample t -test between the distributions of haemodynamic dilation parameters in middle-age (MA) and elderly (E) subjects (null hypothesis $H_0 : \bar{\delta}_m^{\text{MA}} = \bar{\delta}_m^{\text{E}}, \forall m = 1, \dots, M$). 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 13: **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 $(\hat{\delta}_m)_{m=1}^{96}$. 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).