

A. Supplementary material

A.1. Gradient derivation w.r.t the HRF dilation parameter

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

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

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

$$\nabla_{\boldsymbol{\delta}} J(\boldsymbol{\delta}) = \left[\frac{\partial J(\boldsymbol{\delta})}{\partial \delta_1}, \dots, \frac{\partial J(\boldsymbol{\delta})}{\partial \delta_M} \right]^\top \in \mathbb{R}^M \quad (\text{A.1})$$

To this end, we proceed componentwise:

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

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

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

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

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

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

$$v'(t) = \left(\frac{a-1}{t} - 1 \right) \frac{t^{a-1}e^{-\delta t}}{\Gamma(a)} - c \left(\frac{b-1}{t} - 1 \right) \frac{t^{b-1}e^{-\delta t}}{\Gamma(b)} \quad (\text{A.3})$$

$$\frac{\partial}{\partial \delta} v(\delta t) = tv'(\delta t) = \left(\frac{a-1}{\delta} - t \right) \frac{(\delta t)^{a-1}e^{-\delta t}}{\Gamma(a)} - c \left(\frac{b-1}{\delta} - t \right) \frac{(\delta t)^{b-1}e^{-\delta t}}{\Gamma(b)} \quad (\text{A.4})$$

The value of $\frac{\partial v_{\delta_m}}{\partial \delta_m}$ can thus be computed by taking the discrete time points corresponding to the sampling rate of the BOLD signal and the length of the considered HRF. By replacing its value in the computation of $\frac{\partial J(\boldsymbol{\delta})}{\partial \delta_m}$ from Eq. (A.2), we obtain a closed form expression for the gradient of J w.r.t the HRF dilation parameter δ i.e. $\nabla_{\boldsymbol{\delta}} J(\boldsymbol{\delta})$.

A.2. Derivation of λ_{\max}

In this subsection, we derive the computation of λ_{\max} which is the minimal value of λ for which a constant \mathbf{Z}^* is solution of Eq. (6) given in the main text.

In Eq. (A.5), we re-formulate the cost-function J in Eq. (6) in a formulation where unrelated terms are compressed:

$$J(\mathbf{Z}) = \frac{1}{2} \left\| \mathbf{Y} - \sum_{k=1}^K \mathbf{D}_k \star \mathbf{z}_k \right\|_F^2 + \lambda \sum_{k=1}^K \|\nabla \mathbf{z}_k\|_1 + C_{U,\delta} \quad (\text{A.5})$$

where $\mathbf{D}_k = \left(\sum_{m=1}^M \boldsymbol{\Theta}_m^\top \mathbf{v}_{\delta_m} \right) \odot \mathbf{u}_k^\top$ and $C_{U,\delta}$ a constant that does not depend on \mathbf{z}_k for any k .

Then, we compute the optimal solution for which each \mathbf{z}_k is constant, i.e. $\mathbf{z}_k = c_k \mathbf{1}$. We will denote by \mathbf{c} the vector of $(c_k)_{k=1}^K$. For this set of solution, the regularization $\|\nabla \mathbf{z}_k\|_1$ is always 0 and this problem reduces to

$$\min_{\mathbf{c}} \frac{1}{2} \left\| \mathbf{Y} - \sum_{k=1}^K c_k (\mathbf{D}_k \star \mathbf{1}) \right\|_F^2 \quad (\text{A.6})$$

If we denote by \mathbf{y} the flatten vector of all values of \mathbf{Y} and by \mathbb{D} the matrix composed of lines \mathbb{D}_k , which correspond to the flatten vectors $\mathbf{D}_k \star \mathbf{1}$, then the above formulation amounts to solving a simple linear system whose solution is given by $\mathbf{c} = \mathbb{D}^\dagger \mathbf{y}$, where \mathbb{D}^\dagger denotes the pseudo-inverse of matrix \mathbb{D} .

Consequently, we need to find the smallest value of λ such that $\mathbf{z}_k = c_k \mathbf{1}$ is solution of Eq. (A.5). As the sub-gradient of $\|\nabla \mathbf{z}_k\|_1$ is complex, let's use the equivalent synthesis formulation of this problem, as described in (Cherkaoui et al., 2020). Let \mathbf{L} be the discrete integration operator and \mathbf{R} the identity matrix with

a 0 in the first diagonal term. Then, we can use the change of variable $\mathbf{z}_k = \mathbf{L}\mathbf{q}_k$, by noting that \mathbf{L} is invertible and $\nabla \mathbf{z}_k = \mathbf{R}\mathbf{q}_k$. We end up with the following problem

$$\underbrace{\frac{1}{2} \left\| \mathbf{Y} - \sum_{k=1}^K \mathbf{D}_k \star \mathbf{L}\mathbf{q}_k \right\|_F^2}_{f(\mathbf{Q})} + \underbrace{\sum_{k=1}^K \lambda \|\mathbf{R}\mathbf{q}_k\|_1}_{g(\mathbf{Q})} \quad (\text{A.7})$$

Then, the Karush–Kuhn–Tucker (KKT) conditions for the system (A.7) in the point \mathbf{Q}_c read:

$$-\nabla f(\mathbf{Q}_c) \in \partial g(\mathbf{Q}_c) = [-\lambda, \lambda] \quad (\text{A.8})$$

Thus, if $\lambda \geq \lambda_{\max} = \|\nabla f(\mathbf{Q}_c)\|_{\infty} = \max_{k_0 \in \{1..K\}, t \in [2, \tilde{T}]} |(\nabla f(\mathbf{Q}_c))_{k_0}(t)|$, the constant signal \mathbf{Q}_c is solution of Eq. (A.7). Computing this value gives:

$$\forall k_0, \quad (\nabla f(\mathbf{Q}_c))_{k_0}(t) = \left((\mathbf{L}^{\top} \mathbf{D}_{k_0}) \star^{\uparrow} \left(\mathbf{Y} - \sum_{k=1}^K c_k (\mathbf{D}_k \star \mathbf{1}) \right) \right) (t) \quad (\text{A.9})$$

with \star^{\uparrow} the time-reversed convolution operator applied line by line. Note that once the optimal constants \mathbf{c} have been computed, this value can be easily obtained by computing the gradient of the ℓ_2 norm and applying the linear operator \mathbf{L}^{\top} which is a reversed cumulative sum.

A.3. Algorithm details for each estimation step

Each estimation step relies on a proximal gradient descent algorithm that involves a clever gradient computation to reduce the computation cost at each iteration. We detail the algorithm for each estimation step: \mathbf{Z} -estimation in Algorithm A.1, \mathbf{U} -estimation in Algorithm A.2 and δ -estimation in Algorithm A.3.

Algorithm A.1: Estimation of the temporal atoms $\mathbf{Z}^{(i)}$ with fixed $\mathbf{U}^{(i-1)}$ and $\delta^{(i-1)}$

- Input:** BOLD signal \mathbf{Y} , ϵ , $\mathbf{U}^{(i-1)}$, $\delta^{(i-1)}$ and $(L_k)_{k=1}^K$
- 1 initialization: $\forall k, \mathbf{z}_k^{(0)} = \mathbf{z}_k^{(i-1)}$, $t = 1$ (iteration counter)
 - 2 pre-computation: $\forall k, \mathbf{B}_k = \mathbf{u}_k^{(i-1)} \star^{\uparrow} \left(\sum_{m=1}^M \Theta_m^{\top} \mathbf{v}_{\delta_m}^{(i-1)} \right) \mathbf{Y}$
 - 3 $\forall k, \mathbf{C}_k = \mathbf{u}_k^{(i-1)} \left(\sum_{m=1}^M \Theta_m^{\top} \mathbf{v}_{\delta_m}^{(i-1)} \right) \star^{\uparrow} \left(\sum_{m=1}^M \Theta_m^{\top} \mathbf{v}_{\delta_m}^{(i-1)} \right) \star (\mathbf{u}_k^{(i-1)})^{\top}$
 - 4 **repeat**
 - 5 **for** $k \leftarrow 1$ **to** K **do**
 - 6 $\mathbf{z}_k^{(t)} = \text{prox}_{\text{TV}} \left(\mathbf{z}_k^{(t-1)} - \frac{1}{L_k} \left(\mathbf{B}_k - \mathbf{C}_k \mathbf{z}_k^{(t-1)} \right) \right)$
 - 7 **until** $\frac{J((\mathbf{z}_k^{(t-1)})_k) - J((\mathbf{z}_k^{(t)})_k)}{J((\mathbf{z}_k^{(t-1)})_k)} \leq \epsilon$;
- with $\text{prox}_{\text{TV}}(\cdot)$ being the proximal operator of the total variation regularization and $\forall k$, L_k the Lipschitz being constant of $\nabla_{\mathbf{z}_k} J(\mathbf{z}_k)$.
-

Algorithm A.2: Estimation of the spatial maps $\mathbf{U}^{(i)}$ with fixed $\mathbf{Z}^{(i)}$ and $\boldsymbol{\delta}^{(i-1)}$

Input: BOLD signal \mathbf{Y} , ϵ , $\mathbf{Z}^{(i)}$, $\boldsymbol{\delta}^{(i-1)}$ and $(L_k)_{k=1}^K$

- 1 initialization: $\forall k, \mathbf{u}_k^{(0)} = \mathbf{u}_k^{(i-1)}$, $t = 1$ (iteration counter) ;
- 2 pre-computation: $\forall k, \mathbf{B}_k = \mathbf{z}_k^{(i)} \ast \left[\sum_{m=1}^M \boldsymbol{\Theta}_m^\top \mathbf{v}_{\delta_m}^{(i-1)} \right] \mathbf{Y}$
- 3 $\forall k, \mathbf{C}_k = \mathbf{z}_k^{(i)} \left(\sum_{m=1}^M \boldsymbol{\Theta}_m^\top \mathbf{v}_{\delta_m}^{(i-1)} \right) \ast \left[\sum_{m=1}^M \boldsymbol{\Theta}_m^\top \mathbf{v}_{\delta_m}^{(i-1)} \right] \ast (\mathbf{z}_k^{(i)})^\top$
- 4 **repeat**
- 5 | **for** $k \leftarrow 1$ **to** K **do**
- 6 | | $\mathbf{u}_k^{(t)} = \text{proj} \left(\mathbf{u}_k^{(t-1)} - \frac{1}{L_k} \left(\mathbf{B}_k - \mathbf{C}_k \mathbf{u}_k^{(t-1)} \right) \right)$
- 7 **until** $\frac{J((\mathbf{u}_k^{(t-1)})_k) - J((\mathbf{u}_k^{(t)})_k)}{J((\mathbf{u}_k^{(t-1)})_k)} \leq \epsilon$;

with $\text{proj}(\cdot)$ being the projection onto $\{\mathbf{u} \mid \|\mathbf{u}\|_1 = \eta \text{ and } \forall j, u_j > 0\}$ and $\forall k, L_k$ being the Lipschitz constant of $\nabla_{\mathbf{u}_k} J(\mathbf{u}_k)$.

Algorithm A.3: Estimation of the haemodynamic delays $\boldsymbol{\delta}^{(i-1)}$ with fixed $\mathbf{Z}^{(i)}$ and $\mathbf{U}^{(i)}$

Input: BOLD signal \mathbf{Y} , ϵ , $\mathbf{U}^{(i)}$, $\mathbf{Z}^{(i)}$ and L

- 1 initialization: $\boldsymbol{\delta}^{(0)} = \boldsymbol{\delta}^{(i-1)}$, $t = 1$ (iteration counter) ;
- 2 **repeat**
- 3 | $\boldsymbol{\delta}^{(t)} = \text{proj} \left(\boldsymbol{\delta}^{(t-1)} - \frac{1}{L} \nabla_{\boldsymbol{\delta}} J(\boldsymbol{\delta}^{(t-1)}) \right)$
- 4 **until** $\frac{J((\boldsymbol{\delta}^{(t-1)})_k) - J((\boldsymbol{\delta}^{(t)})_k)}{J((\boldsymbol{\delta}^{(t-1)})_k)} \leq \epsilon$;

with $\text{proj}(\cdot)$ being the projection onto $\{\boldsymbol{\delta} \mid \forall m, \delta_m \in [0.5, 2.0]\}$ and L the Lipschitz constant of $\nabla_{\boldsymbol{\delta}} J(\boldsymbol{\delta})$.

A.4. Investigation of Initialization strategies

As the cost function described in Eq. (5) is not convex, the final solution reached during its minimization depends on the initialization, hence we investigate here different strategies to set up the spatial maps \mathbf{U} . We consider three approaches: **[•] Gaussian initialization.** First, the spatial maps $\mathbf{U} \in \mathbb{R}^{K \times P}$ were initialized with K random zero-mean and unitary Gaussian vectors of length P . Then these maps were further projected onto the admissible set of non-negative and sparsity constraints.

[•] Patch-based initialization. Second, we extracted from the BOLD data \mathbf{Y} a sub-matrix $\tilde{\mathbf{Y}}$ of dimensions $P \times L$ (P voxels, L time points) starting at a random initial onset t_0 and computed the matrix-vector product $\tilde{\mathbf{Y}} \mathbf{v}^{(0)}$ with the initialization of the HRF $\mathbf{v}^{(0)}$ to get a first spatial patch $\mathbf{u}^{(0)}$. By repeating this random selection of a given onset K times, we ended up with a full patch-based initialization of K spatial maps in $\mathbf{U}^{(0)}$. Of course, these initial maps are then

projected onto the admissible set of constraints. To some extent, these K maps quantify the presence of the HRF shape (i.e. matched filter approach) in different temporal segments of the BOLD data.

[•] ICA-based initialization. Third, we initialized the K spatial maps using the ICA decomposition of the observed data $\mathbf{Y} = \mathbf{W}^{(0)}\mathbf{U}^{(0)}$ and then projected $\mathbf{U}^{(0)}$ onto the admissible set of constraints.

To determine the impact of each initialization strategy, we used the same simulated experiment as before with two extreme SNR associated with very high and very low levels of white Gaussian noise that were added to the signals of interest (SNR ranging from -20 dB to 40 dB to cover both realistic and noise-free cases). For each noise level, we randomly generated 100 synthetic data sets \mathbf{Y} and estimated the different parameters using these three initialization strategies.

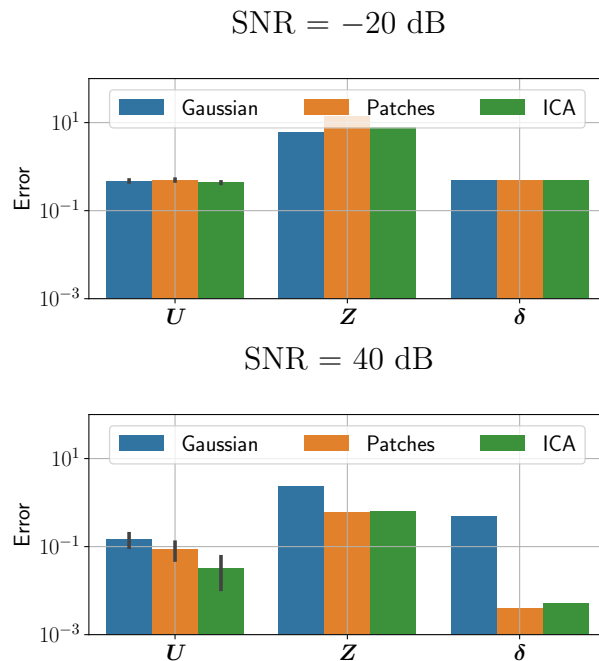


Figure A: **Bar chart comparing the estimation errors in U , Z and δ as a function of the initialization strategy.** As variability across the 100 realizations was only introduced on the initialization of U , the standard deviation of errors is solely reported for U in gray vertical line. In blue, we display the estimation errors associated with the random Gaussian initialization of U . In orange, we depict the estimation errors corresponding to the patch-based initialization. In green, we report the estimation errors that result from using the ICA decomposition to set up the K spatial maps in U . Two SNR levels were used (SNR= -20 dB and SNR= 40 dB) in the same simulated framework as the one used in the core paper.

In Fig. A, for each set up we report the mean estimation errors that were averaged over the 100 realizations. The errors were computed as follows: For the spatial

maps \mathbf{U} we used the following metric: $\text{Error}(\mathbf{U}) = \frac{1}{K} \sum_{k=1}^K (1 - \text{corr}(\hat{\mathbf{u}}_k, \mathbf{u}_k^{\text{True}}))$. For the temporal components \mathbf{Z} and the HRF dilation parameters $\boldsymbol{\delta}$, we considered the normalized residual ℓ_2 -norms: $\text{Error}(\mathbf{Z}) = \frac{1}{K} \sum_{k=1}^K \left(\frac{\|\hat{\mathbf{z}}_k - \mathbf{z}_k^{\text{True}}\|_2}{\|\mathbf{z}_k^{\text{True}}\|_2} \right)$ and $\text{Error}(\boldsymbol{\delta}) = \frac{\|\hat{\boldsymbol{\delta}} - \boldsymbol{\delta}^{\text{True}}\|_2}{\|\boldsymbol{\delta}^{\text{True}}\|_2}$.

As expected we first observed significantly lower errors in the high SNR regime (SNR=40 dB) for all parameters of interest but more importantly for \mathbf{Z} and $\boldsymbol{\delta}$. Additionally, in the low SNR mode we noticed a similar behavior for all three initialization strategies. In contrast, in the high SNR regime, the ICA initialization provides us with better parameter estimates, especially for \mathbf{U} . The patch-based approach achieves a good trade-off between the purely random Gaussian initialization, which performs the worse as it does not take the observed data into account, and the ICA.

A.5. Single-subject analysis on rs-fMRI data

In Fig. B, we complete the estimated spatial maps reported in Fig. 6 with the associated temporal components that capture the activation profiles of the given spatial maps.

A.6. Haemodynamic parameter estimate stability across various levels of temporal regularization

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

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

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

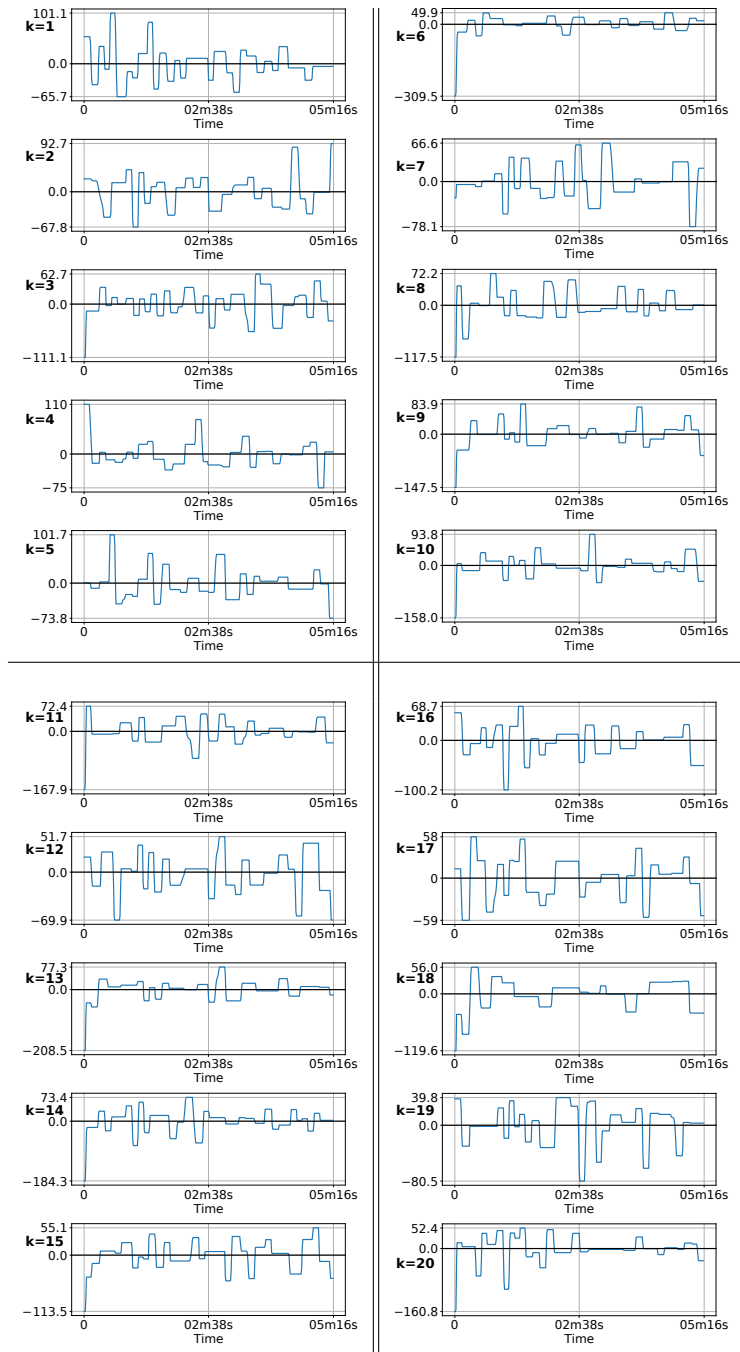


Figure B: **Temporal decomposition of rs-fMRI data for $K = 20$.** From top to bottom and left to right, the twenty labeled temporal components are shown with adapted vertical axis to better observe the temporal pattern of each component. The labeling is arbitrary but matches the order of spatial maps shown in Fig. 6.




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

This previous result entails that the haemodynamic delay estimated by our model (6) may be biased. However, because there is a single temporal regularization parameter, we expect that this bias impacts the whole brain uniformly. To assess this shared effect on the estimated parameter, we observe the relative variations of δ_m . Fig. D displays the value of $\bar{\delta}_m(\lambda_f)$ relative to $\bar{\delta}(\lambda_f)$ for $S = 459$ subjects with three temporal regularization values $\lambda_f \in \{0.001, 0.5, 0.9\}$ on the MNI template. Precisely, for each regularization parameter and for each ROI m , we compute $\bar{\delta}_m/\bar{\delta}$ where $\bar{\delta}_m = \frac{1}{S} \sum_{s=1}^S \hat{\delta}_m^s$ is the average value of the dilation parameter across subjects. While the magnitudes change when the regularization changes, as seen in Fig. C, the spatial structure of dilatation parameters in the brain is globally preserved. Indeed, the normalized maps look very similar for any choice of regularization parameter, showing that the relative variation between each area of the brain are preserve while changing the hyper-parameter. Thus, we can state that the haemodynamic response from the middle temporal gyrus is faster than the response from the frontal orbital cortex, as described in Fig. 7. This means that while the numerical value of the time-to-peak for a given area may not be reflect the actual haemodynamic delay in the brain, the estimated coefficients reflect the spatial variations of the delay between the different areas of the brain. Moreover, these variations are stable with the choice of temporal regularization. Hence, choosing a potentially suboptimal value for λ_f is of limited impact when the primary interest is investigating abnormalities in the neuro-vascular coupling.




Figure D: **Group-level mean of haemodynamic dilation parameter maps normalized by the grand average** $\bar{\delta} = \frac{1}{MS} \sum_{s=1}^S \sum_{m=1}^M \hat{\delta}_m^s$ **as a function of temporal regularization** ($\lambda_f \in [0, 1]$). From top to bottom, axial slices showing the group-level values of the ratio between $\bar{\delta}_m = \frac{1}{S} \sum_{s=1}^S \hat{\delta}_m^s$ and $\bar{\delta}$ in each parcel m for increasing values of $\lambda_f \in \{0.001, 0.5, 0.9\}$. The spatial structure of the maps of haemodynamic dilation parameter remain remarkably stable for various λ_f .

A.7. Middle-age vs elderly subjects classification

Finally, we assess the impact of the choice of λ_f on the prediction results from Section 4.2. Fig. E reports the accuracy score for the logistic regression relatively to the choice of regularization parameter β for the classification model and the temporal regularization parameter λ_f for our deconvolution model. The accuracy is almost not impacted by the choice of parameter λ_f , for any value of β . This observation confirms that the choice of λ_f is not critical when studying the relative spatial structure of the haemodynamic delay and that our model can be used in practical cases to evaluate abnormalities in the haemodynamic response.

A.8. Impact of serially correlated Gaussian noise on parameter estimation

Although the noise that contaminates the BOLD effect on fMRI data is serially correlated in time (Woolrich et al., 2001), the proposed approach in the core manuscript neglects this property as we consider a white Gaussian process in Eq. (2). In this section we investigate on simulations the impact of data contaminated by a serially-correlated Gaussian noise on the estimation quality of neural activity and more globally on the parameter estimation in terms of reconstruction error (R2-score). Here, we used a first-order autoregressive noise model (AR(1)), i.e. $e_j(t_n) = \rho e_j(t_{n-1}) + \epsilon_j(t_n)$ where ϵ_j is an innovation, to corrupt synthetic fMRI time series.

In a first simulation, we stick to the synthetic scenario investigated in Section 3.1 except that here the Gaussian noise is auto-correlated in time with an auto-regressive parameter value of $\rho = 0.3$. Note that this value is compatible

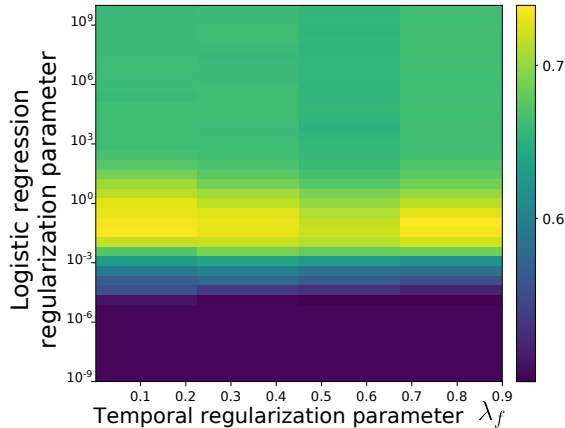


Figure E: Evolution of the accuracy score w.r.t. the logistic regression regularization parameter β and the temporal regularization parameter λ_f . The accuracy score is not impacted by the hyper-parameter λ_f , as moving this parameter mainly impact the magnitude of the estimated delays and not its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

with actual autocorrelation estimates on real fMRI data in the white and gray matter (Penny et al., 2003, Fig. 11). In both cases the signal-to-noise-ratio (SNR) was maintained to -10 dB, a relevant scenario met on real fMRI data.

Fig. F displays the parameter estimates (and the values of residual mean square error, i.e. RMSE, for each atom) along with the true signals that were used for generating the synthetic data. We notice the accurate recovery of the temporal profiles, the spatial maps and the HRF shape (panels (a), (b) and (c), respectively) in both noise case scenarios with a slightly higher RMSE in the situation where the data was corrupted by an AR(1) Gaussian noise.

Next, to deeper investigate the impact of the autocorrelation noise structure on the model mismatch and the results over a wide range of simulations, we analyzed the evolution of the R2-score for multiple SNR levels when injecting an autoregressive Gaussian noise, still with the same AR parameter value ($\rho = 0.3$). To be quite exhaustive, the SNR level was varied between -15 dB and -5 dB and 30 noise realizations were simulated for each SNR level. The motivation for quantifying the performances through the R2-score ($R_2 = 1 - \frac{\|\mathbf{Y} - \hat{\mathbf{Y}}\|_2^2}{\|\mathbf{Y}\|_2^2}$) instead of using solely the RMSE on the temporal atoms was to summarize the overall reconstruction error and the residual in ℓ_2 -norm between the predicted data $\hat{\mathbf{Y}}$, computed from the parameter estimates $(\hat{\mathbf{U}}, \hat{\mathbf{Z}}, \hat{\boldsymbol{\delta}})$, and the multivariate measurements \mathbf{Y} . In Fig. F, we got $R_2 = 0.502$ for the white Gaussian noise (more than 50 % of explained variance by our model) and $R_2 = -0.804$ for the AR(1) Gaussian noise, indicating a loss in explanatory power of our model.

In Fig. G we plot the reconstruction error (i.e. R2-score) of our approach as a function of the SNR level in the white and AR(1) ($\rho = 0.3$) Gaussian noise

scenarios. We show that the R2-scores are positive (good model prediction) and close to each other under the two noise models for SNR values ranging from -5 dB down to -8 dB. Then the performances of our approach degrade quite fast on the data corrupted with the AR(1) noise compared to the white noise contamination in the more noisy scenarios (SNR from -10 dB down to -15 dB).

Additionally, to go one step further, we analyzed the evolution of the R2-score for a value of ρ that varies between 0.1 and 0.9. The SNR was set to -10 dB across all these experiments comprising 30 noise realizations for each ρ value.

In Fig. H we plot the reconstruction error (i.e. R2-score) of our approach as a function of the auto-regressive parameter ρ . We show that the R2-score progressively degrades and systematically becomes negative across all experiments (not only in average) as ρ tends to 1, a situation that is rarely met on real fMRI data (Penny et al., 2003; Woolrich et al., 2004).

In conclusion, our model that relies on a white noise assumption seems tenable for AR(1) parameter values ranging from 0 up to 0.3 and for SNR levels larger than -8 dB. These constraints remain quite realistic on actual fMRI data.

References

- K. J. Friston, O. Josephs, G. Rees, R. Turner, Nonlinear event-related responses in fMRI., *Magnetic Resonance in Medicine* 39 (1998) 41–52.
- H. Cherkaoui, J. Sulam, T. Moreau, Learning to solve tv regularised problems with unrolled algorithms., 34th Conference and Workshop on Neural Information Processing Systems (NeurIPS) (2020) 1–21.
- R. Tibshirani, Regression Shrinkage and Selection via the Lasso., *Journal of the Royal Statistical Society: Series B (statistical methodology)* 58 (1996) 267–288.
- M. W. Woolrich, B. D. Ripley, M. Brady, S. M. Smith, Temporal autocorrelation in univariate linear modeling of fMRI data., *Neuroimage* 14 (2001) 1370–1386.
- W. Penny, S. Kiebel, K. Friston, Variational bayesian inference for fMRI time series., *NeuroImage* 19 (2003) 727–741.
- M. W. Woolrich, M. Jenkinson, J. M. Brady, S. M. Smith, Fully Bayesian spatio-temporal modeling of fMRI data, *IEEE Transactions on Medical Imaging* 23 (2004) 213–231.

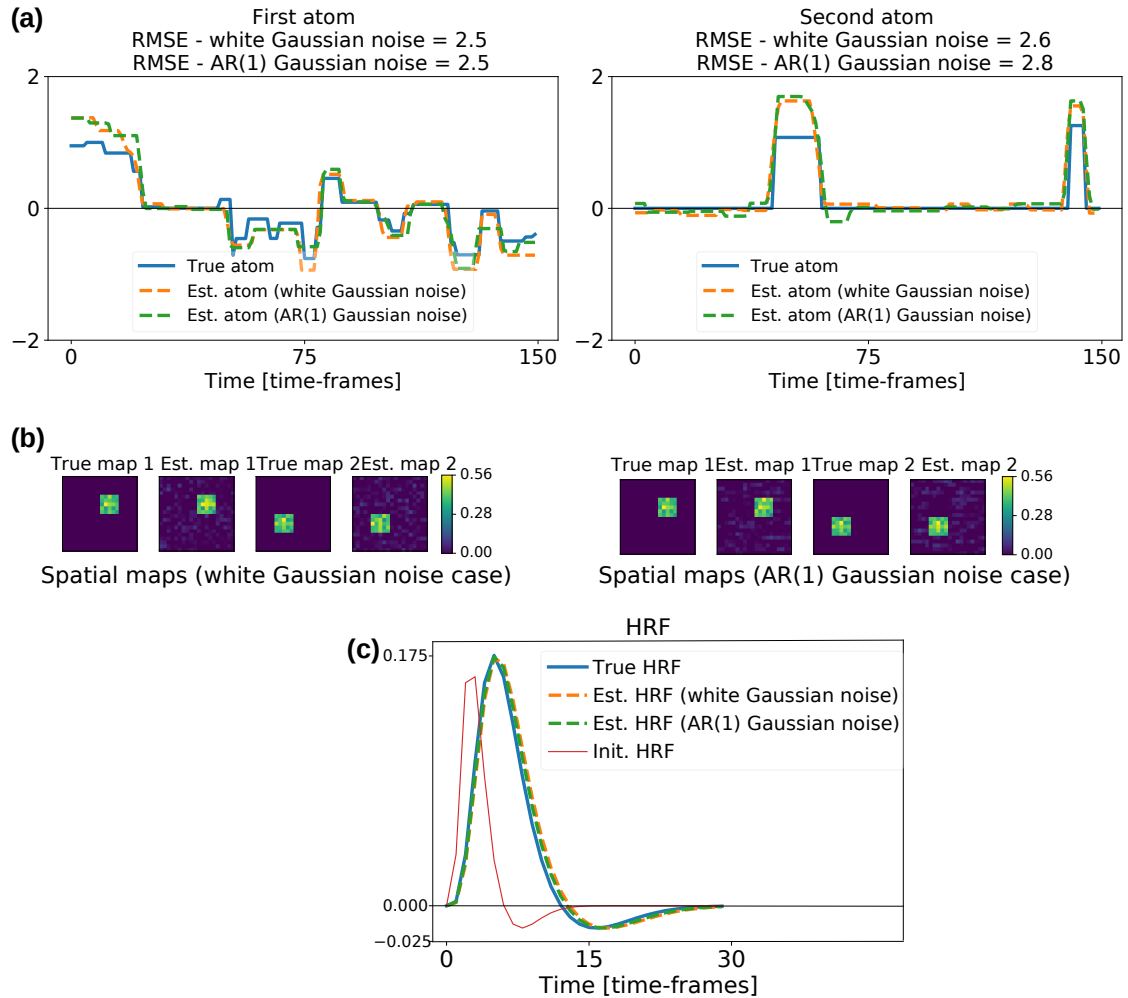


Figure F: **Top (a) the temporal activation signals:** on both sides, the true temporal profiles are depicted in solid blue line whereas the estimated ones are shown in orange and green dashed lines for the white and AR(1) Gaussian noise cases, respectively for the two activating regions (first atom for map 1, second atom for map 2). **Center (b) the spatial maps:** the true and estimated spatial maps for the two noise model scenarios displaying the two activating regions. **Bottom (c) the haemodynamic response function shape:** the true HRF shape is depicted in solid blue line while the HRF estimates are plotted in orange and green dashed lines for the white and AR(1) Gaussian noise cases, respectively. The initialization used in both cases is shown in red.

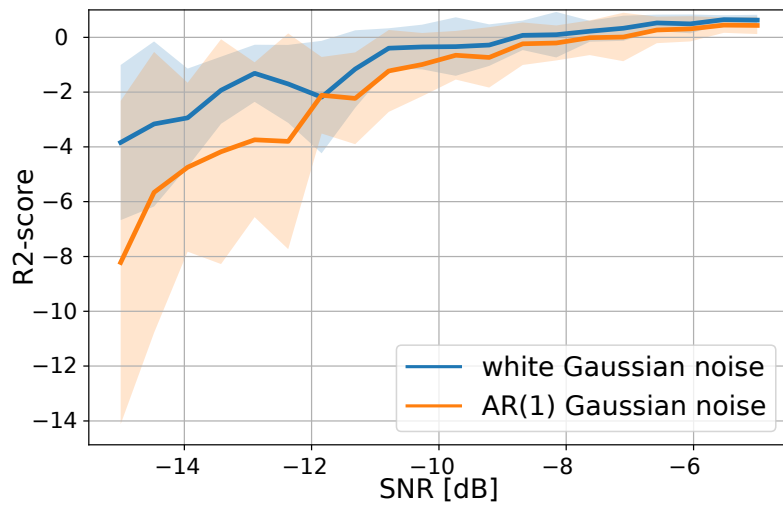


Figure G: **Evolution of the R2-score** with respect to the SNR ranging from -15 dB to -5 dB. The average R2-score is depicted in solid blue line for the white Gaussian noise and in solid orange line for the AR(1) ($\rho = 0.3$) Gaussian noise. The surrounding shading is used to report the variability across the different runs (noise realizations) of the experiment.

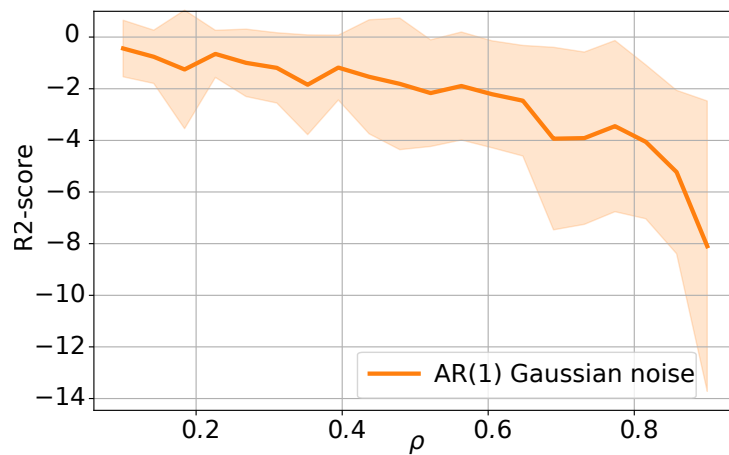


Figure H: **Evolution of the R2-score** with respect to the auto-regressive parameter ρ ranging from 0.1 to 0.9. The average R2-score is depicted in solid orange line for the AR(1) Gaussian noise (SNR= -10 dB). The surrounding shading is used to report the variability across the different runs (noise realizations) of the experiment.