



HAL
open science

FastSRM: A fast, memory efficient and identifiable implementation of the shared response model

Hugo Richard, Bertrand Thirion

► **To cite this version:**

Hugo Richard, Bertrand Thirion. FastSRM: A fast, memory efficient and identifiable implementation of the shared response model. *Aperture Neuro*, 2023, 10.52294/001c.87954 . hal-04353509

HAL Id: hal-04353509

<https://inria.hal.science/hal-04353509>

Submitted on 19 Dec 2023

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.



Distributed under a Creative Commons Attribution 4.0 International License

FastSRM: A fast, memory efficient and identifiable implementation of the shared response model

Hugo Richard^a, Bertrand Thirion^b

^a *Criteo AI Lab, Paris, France*

^b *MIND, INRIA, CEA, Neurospin, Gif-sur-Yvette, France*

ABSTRACT

The shared response model (SRM) provides a simple but effective framework to analyze fMRI data of subjects exposed to naturalistic stimuli. However, when the number of subjects or runs is large, fitting the model requires a large amount of memory and computational power, which limits its use in practice. Furthermore, SRM is not identifiable, which makes the shared response difficult to interpret.

In this work, we implement an identifiable version of SRM and show on real data that it improves the stability of the recovered shared response. We then introduce FastSRM, which relies on a dimension reduction step and yields the same solution as the original algorithm. We show experimentally using synthetic and real fMRI data that FastSRM is considerably faster and more memory efficient than current implementations.

The experiments performed in this article are *fully* reproducible: our code available at <https://github.com/hugorichard/FastSRM> allows you to download the data, run the experiments and plot the figures.

For a more pleasant reading of the math equations, we advise the reader to read the version of this article formatted with latex: <https://github.com/hugorichard/FastSRM/blob/master/paper/main.pdf>

Corresponding author: Dr. Hugo Richard, ENSAE, CREST, Institut Polytechnique de Paris, France, email: hugo.richard@ensae.fr

Date Received: April 1, 2022

Date Accepted: August 19, 2022

INTRODUCTION

Complex stimulations such as movie watching, story reading or music listening are more and more popular in the neuroscientific community. Indeed such naturalistic paradigms are unconstrained from behavioral manipulations and thus, more ecological with respect to everyday cognitive conditions. If one wants to use data acquired under naturalistic conditions to map functional responses to the brain, that is, in an encoding setting (1), one has to deal with the fact that no design matrix is specified for naturalistic stimuli. Such a design matrix encodes the temporal events that affect brain signals during an experiment; these events typically reflect the occurrence of some features of interest in the stimuli. Although some works have used manual annotations (2) or deep neural networks (3,4), see also <https://neuroscout.org>, to create an estimate of the design matrix of naturalistic

stimuli, it is a hard task. Without a design matrix, models such as the general linear model (5) cannot be used.

Another approach is to learn jointly a set of reference time courses and the spatial maps in an unsupervised way. These reference time courses formally replace the design matrix in the spatiotemporal decomposition of the data, but cannot be considered as a design matrix, as it is data driven. In the case where the data set only contains one subject, independent component analysis (6) (ICA) is the method of choice. ICA assumes that components are independent, a defensible assumption that ensures identifiability up to permutation and scaling of the model and can be fitted efficiently. When the data set contains multiple subjects undergoing the same protocol, a natural assumption is to assume that the temporal reference is shared across subjects. Many different methods can produce shared components from different subjects. Some assume independent components (7–10),

while others are only based on second-order information such as multiset canonical correlation analysis (11) or the shared response model (SRM) (12).

This work focuses on the SRM. Because of its simplicity and its built-in dimension reduction, this model is widely used (13–17), in particular, as a preprocessing step for source identification (7,8), classification (18–20) or as a basis for transfer learning (21).

What makes the analysis of fMRI data particularly difficult is that the number of features (voxels) is usually much larger than the number of available samples (time frames). Yet, SRM has initially been designed to work within regions of interest using only a few subjects. When using full brain data, computational costs become high. In addition, memory requirements are difficult to meet since the full data set needs to hold in memory. Lastly, another issue of SRM is its nonidentifiability (see the proof in (8) and Appendix D, also reported in Appendix B for completeness), which makes it challenging to interpret the extracted shared components as they are defined up to a rotation.

In this work, we introduce FastSRM, an identifiable model that uses an optimal spatial decomposition to speed up the computations with provably no loss of performance. FastSRM gains identifiability by imposing that the covariance of the shared components is diagonal (see (8), Appendix D or Appendix B of this paper for a proof). The gains in speed and memory usage are based on the observation that a low-dimensional, yet the distance-preserving representation of the images yields the same result as the full data. Such a representation can be interpreted as an initial spatial decomposition of the data. We show that FastSRM brings several practical benefits: on real data, it produces more stable estimates of the shared components than its nonidentifiable counterpart, and it is much faster and more memory efficient than other implementations that do not make use of an optimal spatial decomposition.

A Python implementation is available at <https://github.com/hugorichard/FastSRM>. All our experiments are fully reproducible. Scripts to download the data, run the experiments and plot the figures are included. A continuous integration pipeline runs the tests automatically when any change to the core algorithm is made.

BACKGROUND: THE SHARED RESPONSE MODEL

The SRM (12) is a multiview latent factor model. The data $x_1 \dots x_m$ are modeled as random vectors following the model:

$$x_i = A_i s + n_i \quad (1)$$

$$A_i^T A_i = I_p \quad (2)$$

where $x_i \in \mathbb{R}^v$ is the data of subject i , called view i , $A_i \in \mathbb{R}^{p \times v}$ is the mixing matrix of view i , n_i is the noise of view i

and $s \in \mathbb{R}^p$ are the shared components referred to as the *shared response* in fMRI applications. p is the number of time frames, v is the number of voxels and m is the number of subjects.

The mixing matrices A_i are assumed to be orthogonal so that $A_i^T A_i = I_p$. However, in general, the matrix $A_i A_i^T$ is different from identity. The noise n_i is assumed to be Gaussian with covariance Σ_i and independent across views. We assume the number of features v to be much larger than the number of components p : $v \gg p$.

The conceptual Figure 1 illustrates an application of the SRM to fMRI data. The mixing matrices are spatial topographies specific to each subject while the shared components represent the common timecourses. In (12, 22), two versions of the SRM are introduced, namely, the deterministic and probabilistic models.

Deterministic SRM

Let us consider n observations (scans) of x_i and s that we stack into matrices $X_i \in \mathbb{R}^{v \times n}$ and $S \in \mathbb{R}^{p \times n}$. The deterministic SRM considers both the mixing matrices A_i and the n observations of the shared response S as parameters to be estimated. The noise variance is fixed to a multiple of identity: $\forall i, \Sigma_i = \sigma^2 I_v$ where σ is a hyperparameter to choose.

The model is optimized by maximizing the log-likelihood, assuming that the noise is Gaussian distributed. The likelihood is then given by $p(x) = \prod_i N(x_i; A_i s, \sigma^2 I)$ and therefore the empirical negative log-likelihood is given up to a constant independent of A_i and S by

$$L = \frac{1}{n} \sum_i \|A_i S - X_i\|^2 = \frac{1}{n} (\|S\|^2 - 2 \langle A_i S, X_i \rangle + \|X_i\|^2) \quad (3)$$

L is optimized by performing alternate minimization on $(A_1 \dots A_m)$ and S . Note that the hyperparameter σ does not have an influence on the loss and can therefore be ignored.

The gradient with respect to S is given by $\sum_i A_i^T (A_i S - X_i) = \sum_i (S - A_i^T X_i)$ yielding the closed form updates:

$$S \leftarrow \frac{1}{m} \sum_i (A_i^T X_i) \quad (4)$$

From (3), minimizing L with respect to A_i is equivalent to maximizing $\langle A_i, X_i S^T \rangle$ and therefore we have

$$A_i \leftarrow P \left(\frac{1}{n} X_i S^T \right) \quad (5)$$

where P is the projection on the Stiefel manifold: $P(M) = M(M^T M)^{-\frac{1}{2}}$.

The complexity of deterministic SRM is in $\tilde{O}(n_{\text{iter}} m p v n)$, where n_{iter} is the number of iterations. We monitor the convergence by computing the ℓ_∞ norm of the gradient. Note that we can monitor the gradient without any increase in complexity. Indeed, after the updates with respect to each

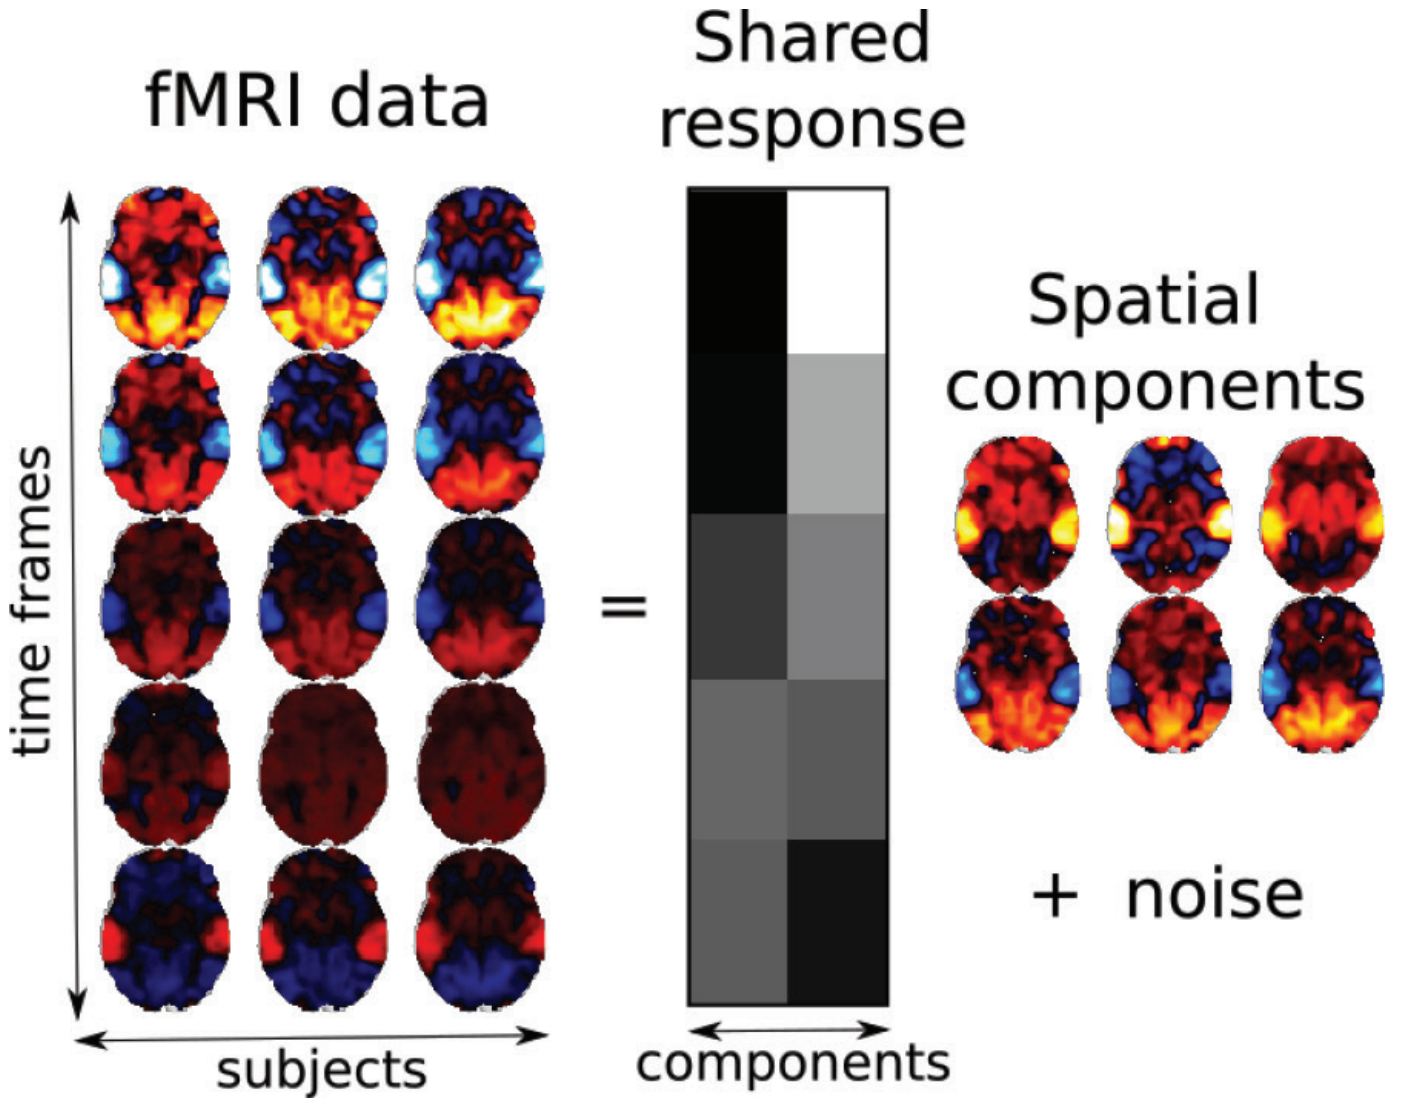


Fig. 1. Shared response model: The raw fMRI data are modeled as a weighted combination of subject-specific spatial components with additive noise. The weights are shared between subjects and constitute the shared response to the stimuli.

mixing matrix have been carried out, the gradient with respect to $(A_i)_{i=1}^m$ is zero and therefore to compute the ℓ_∞ norm of the gradient, we only need the gradient with respect to S : $mS - \sum_i A_i^T \mathbf{x}_i$, where the right-hand side is used in the updates of S . The algorithm is stopped when the gradient falls below a chosen tolerance.

Probabilistic SRM

In probabilistic SRM, $\Sigma_i = \sigma_i^2 I_v$ and the shared components are assumed to be Gaussian $s \sim N(0, \Sigma_s)$.

In (12) and (22), Σ_s is only assumed to be a definite positive. As already highlighted in introduction, this causes the model to be unidentifiable (see (8), Appendix D or Appendix B of this paper for a proof). Enforcing a diagonal Σ_s ensures identifiability, provided that the diagonal values are different. So we assume here that Σ_s is diagonal (and refer the interested reader to (12) and (22) for the original formulation of probabilistic SRM without the

diagonal constraint). The following paragraph gives update rules for λ_i and A_i and Σ_s . We highlight that the derivations used to obtain these updates are not new. They are obtained following the strategy introduced in (12) and (22).

The model is optimized via the expectation-maximization algorithm. We give all details on the formulation and derivation of update rules in Appendix A. Denoting $\mathbb{V}[s | \mathbf{x}] = \left(\sum_i \frac{1}{\sigma_i^2} I + \Sigma_s^{-1} \right)^{-1}$ and $\mathbb{E}[s | \mathbf{x}] = \mathbb{V}[s | \mathbf{x}] \sum_i \frac{1}{\sigma_i^2} A_i^T \mathbf{x}_i$, the updates are given by

$$\sigma_i^2 \leftarrow \frac{1}{v} (\mathbb{E}[\|\mathbf{x}_i - A_i \mathbb{E}[s | \mathbf{x}]\|^2] + \|\text{diag}(\mathbb{V}[s | \mathbf{x}])\|) \quad (6)$$

$$A_i \leftarrow P(\mathbb{E}[\mathbf{x}_i \mathbb{E}[s | \mathbf{x}]^T]) \quad (7)$$

$$\Sigma_s \leftarrow \mathbb{V}[s | \mathbf{x}] + \mathbb{E}[\mathbb{E}[s | \mathbf{x}] \mathbb{E}[s | \mathbf{x}]^T] \quad (8)$$

The complexity of probabilistic SRM is $\tilde{O}(n_{\text{iter}} mpvn)$, the same as in deterministic SRM. We can monitor the

convergence by computing the log-likelihood decrease at each iteration and stop the algorithm when the magnitude of the decrease is below some tolerance. The storage requirements of deterministic or probabilistic SRM are in $\tilde{O}(m\nu n)$, which simply means that the data set needs to hold in memory.

THE FASTSRM ALGORITHM

Reducing the computational burden by the use of spatial decompositions

SRM algorithms use different sets of parameters θ to represent the data. In deterministic SRM, $\theta = ((A_i)_{i=1}^m, S)$, where $(A_i)_{i=1}^m$ are the mixing matrices and S is the shared response, while in probabilistic SRM, $\theta = ((A_i)_{i=1}^m, \Sigma_s, (\sigma_i)_{i=1}^m)$, where $(A_i)_{i=1}^m$ are the mixing matrices, $(\sigma_i)_{i=1}^m$ the noise levels and Σ_s the components variance.

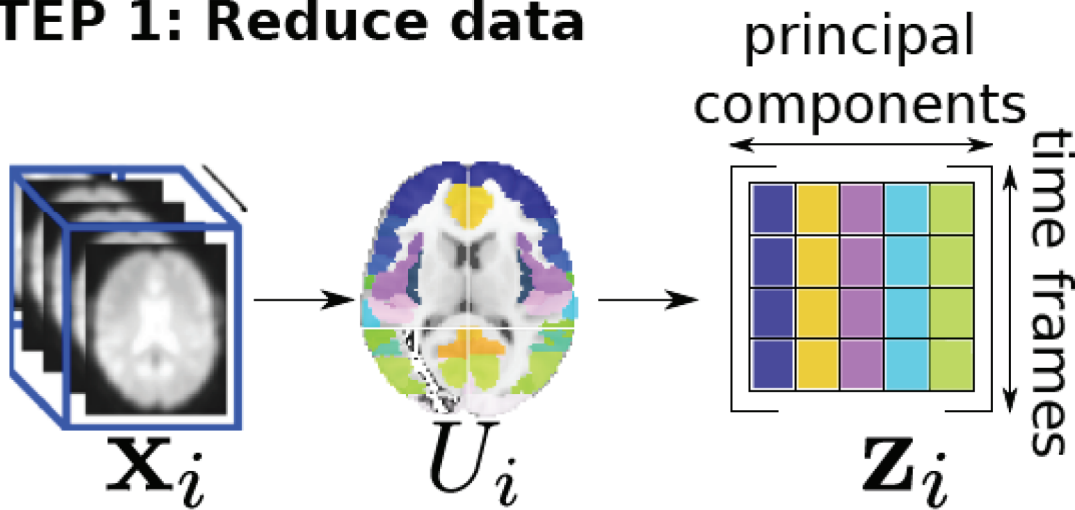
In fMRI, spatial decompositions are often used to reduce the dimensionality of the data. A classical approach

is to apply a deterministic atlas such as (23), which is a parcellation of the brain into r regions. There also exist probabilistic atlases such as (24) that describe each region as a set of weights across the full brain.

Deterministic and probabilistic atlases are spatial decompositions that do not depend on the view at hand. In FastSRM, we consider a set of view-specific spatial decompositions $U_i \in \mathbb{R}^{\nu \times r}$ such that $U_i^T U_i = I_r$, where r is the number of components in the spatial decompositions. Data are reduced using $z_i = U_i^T x_i$ and an SRM algorithm is applied on data z_i yielding parameters θ' . Figure 2 illustrates this process.

Note that the parameters obtained with FastSRM θ' are different from the parameters obtained with the corresponding SRM algorithm θ (the unmixing matrices in θ' do not even have the same shape as the unmixing matrices in θ). However, as we will see in the next section, there exist spatial decompositions such that the models parametrized by θ and θ' are equivalent.

STEP 1: Reduce data



STEP 2: Apply SRM algorithm

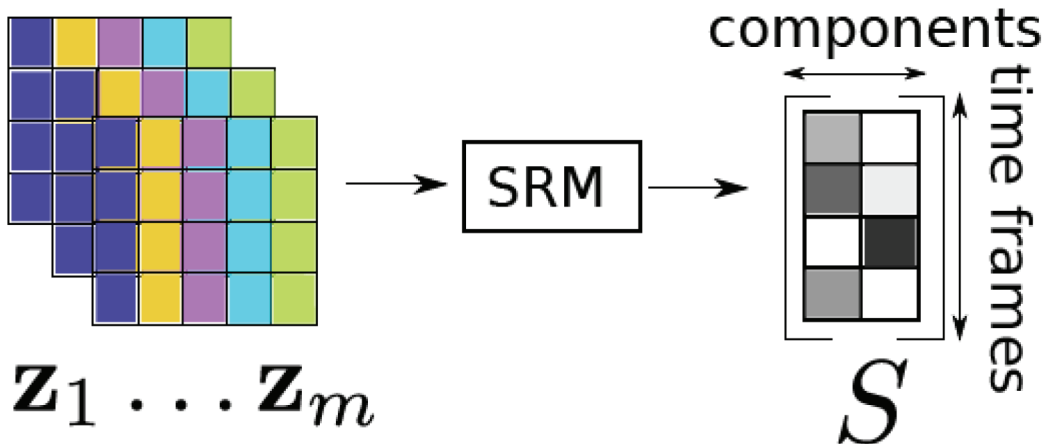


Fig. 2. FastSRM algorithm In step 1, data x_i are projected onto a spatial decomposition U_i that may depend on the subject i (top). In step 2, an SRM algorithm is applied on reduced data to compute the shared response.

From a computational standpoint, dimension reduction provides a large reduction in memory usage. Indeed as the original data are seen only once, it is no longer necessary to keep the full data set in memory (we can load data X_i one after the other and similarly for the spatial decomposition U_i). Therefore, the memory consumption is only in $\tilde{O}(vn)$ (where v is the number of voxels and n is the number of samples), which is lower than that of SRM by a factor of m , the number of subjects. The number of subjects is typically between 10 and 1000. This yields a practical benefit: on fMRI data sets with many subjects, one no longer needs a large cluster to run the SRM but only a modern laptop. Additionally, low memory consumption reduces the risk of thrashing (25), a phenomenon that causes a large increase in computation time when the memory used is close to the total available memory in the hardware.

After preprocessing, the reduced representation z_i is used instead of the original data x_i yielding a time complexity of $\tilde{O}(T_{\text{preprocessing}} + n_{\text{iter}} m p n r)$. Let us highlight that an experiment is often run multiple times, typically when the analysis requires cross-validation procedures. In these cases, the preprocessing is performed only once and the apparent complexity becomes $\tilde{O}(n_{\text{iter}} m p n r)$, which is faster than SRM by a factor of $\frac{v}{r}$. The number of components in large spatial decompositions is about $r = 1000$ and in full brain data, the number of voxels is typically in the 50,000–400,000 range, so that $\frac{v}{r}$ is typically about 50–400. It remains to be shown how to draw a correspondence between FastSRM and SRM, which we do in the following section.

A class of optimal spatial decompositions

In principle, FastSRM can be used with any spatial decomposition. However, in general, working with reduced data induces a loss of information that can be minimized if the spatial decomposition is carefully chosen. In any case, there is little hope to recover the parameters that would have been obtained from SRM from the parameters of FastSRM. Yet, we show that there exists an optimal spatial decomposition in the sense that SRM and FastSRM yield the same results.

Let us consider $x_i = U_{x_i} z_i$ a PCA of x_i with the maximum number of components. As the number of samples n is lower than the number of features, $U_{x_i} \in \mathbb{R}^{v \times n}$ and $z_i \in \mathbb{R}^n$. We also have $U_{x_i}^T U_{x_i} = I$. Therefore, U_{x_i} constitutes a possible choice of subject-specific spatial decomposition. As the next property shows, U_{x_i} is an optimal spatial decomposition for deterministic FastSRM. Note that this optimal spatial decomposition is not unique: any orthogonal matrix with the same span as the columns of X_i would also work.

Proposition 1 (Optimal spatial decomposition for deterministic FastSRM). Let $(A)_i, S$ be the solution obtained by deterministic SRM on data $(X)_i$ and $, S'$ the solution

obtained by deterministic FastSRM on data $(X)_i$ using spatial decompositions $(U_{x_i})_i$, where $X_i = U_{x_i} Z_i$ is a PCA of X_i . Then, $A_i = U_{x_i} A'_i$ and $S = S'$.

The proof is available in Appendix C. In the case of probabilistic SRM, we can obtain very similar results. However, the algorithm applied on reduced data needs to be slightly modified. We have the following result:

Proposition 2 (Optimal spatial decomposition for probabilistic FastSRM). Let $(A)_i, \sigma_i, \Sigma_s$ be the solution obtained by probabilistic SRM on data $x_i \in \mathbb{R}^v$ and $(A'_i)_i, \sigma'_i, \Sigma'_s$ the solution obtained by ProbSRM on data $z_i = U_{x_i}^T x_i, z_i \in \mathbb{R}^n$ but where updates

$$\sigma_i^2 \leftarrow \frac{1}{n} (\mathbb{E}[\|z_i - A_i \mathbb{E}[s|z]\|^2] + \|\text{diag}(\nabla[s|z])\|^2) \quad (9)$$

are replaced by updates

$$\sigma_i^2 \leftarrow \frac{1}{v} (\mathbb{E}[\|z_i - A_i \mathbb{E}[s|z]\|^2] + \|\text{diag}(\nabla[s|z])\|^2) \quad (10)$$

where n is the number of samples and v is the number of voxels.

Then $A_i = U_{x_i} A'_i, \sigma_i = \sigma'_i$ and $\Sigma_s = \Sigma'_s$.

The proof is available in Appendix D. Propositions 1 and 2 show that no information is lost by replacing $x_i \in \mathbb{R}^v$ by its reduced representation $z_i \in \mathbb{R}^n$. A key observation is that Propositions 1 and 2 hold whether or not the model for deterministic (respectively probabilistic) SRM is indeed the generative model of the data.

A complexity analysis shows that finding an optimal spatial decomposition becomes the limiting step of the pipeline. Even with fast implementations, subject-specific PCA is costly. However, FastSRM only works on z_i so we do not need to know the value of U_{x_i} . In practice, we observe data $X_i \in \mathbb{R}^{v \times n}$ and we want to get $Z_i \in \mathbb{R}^{n \times n}$ such that $X_i = U_{x_i} Z_i$. This can be done by performing an SVD of $X_i^T X_i$ yielding $X_i^T X_i = V_i D_i V_i^T$ and setting $Z_i = D_i^{\frac{1}{2}} V_i^T$. Although computing the product $X_i^T X_i$ has time complexity $\tilde{O}(vt^2)$ (the same as a PCA), the constant in the \tilde{O} is exactly one so this operation costs a lot less than a PCA on full data. When estimates of the mixing matrices are needed, they can be obtained by applying equation (26) in the deterministic SRM case and equation (31) in the probabilistic SRM case that only costs $\tilde{O}(mvp^2)$. In practice, the cost of the matrix products $X_i^T X_i$ is often the limiting step of the pipeline (this depends on the number of iterations) but as we show next, it is much more efficient than performing SRM on the full data. Note that if large memory resources are available, these matrix products can be computed in parallel.

RELATED WORK

The implementation of SRM proposed in (12) is quadratic in the number of features, which prevents its application to full brain data. In (22), this issue is addressed by using the inversion lemma to remove the quadratic

dependence. Their implementation is now the most widely used one. In our work, we propose to reduce further the computation time and memory usage compared to the implementation in (22) by the use of optimal spatial decompositions.

Other popular methods used to speed up SRM are searchlight (20) or piecewise (26) approaches. While these methods are efficient, they do not optimize the same objective as the original SRM algorithm and are arguably less principled since the searchlights or pieces are chosen a priori.

We consider SRM as a well principled formulation of the hyperalignment method (27). Many methods exist to perform hyperalignment: deep hyperalignment (28), robust SRM (18), multiview canonical correlation analysis (29), ShICA (7), MultiViewICA (8), optimal transport (30) and many more. SRM has been introduced after hyperalignment in (12) and has been shown to outperform hyperalignment in many different settings. Specifically, SRM is natively a group model, which hyperalignment is not; it is more computationally efficient, with the avoidance of searchlight loops, and sometimes obtains higher accuracy than Procrustes hyperalignment (26). In this article, we do not discuss further the merit of FastSRM with respect to other methods. We simply observe that SRM is widely used and aim to provide a faster, more memory efficient and identifiable implementation. In our view, FastSRM should be used as a plugin replacement for SRM.

EXPERIMENTS

We make several experiments on both synthetic and real data. We used Nilearn (31) (version 0.8.1) for fMRI data preprocessing, Brainiak (32) (version 0.9) for the nonidentifiable version of SRM that implements the work of (12) and (22), Numpy (33) (version 1.19.0) for array processing, Scipy (34) (version 1.5.3) its implementation of the Hungarian algorithm, Matplotlib (35) (version 3.1.2) for plotting, and Sklearn (36) (version 0.23.2) for machine learning pipelines.

Comparing fitting time and performance of FastSRM and SRM on synthetic data

We generate synthetic data x_i according to the model of probabilistic SRM. The parameters of σ_i^2 , A_i and Σ_s are generated randomly. We sample the value of the subject-specific noise level from a normal distribution: $\sigma_i \sim N(0, 0.1)$. The mixing matrices A_i are obtained by sampling their coefficient from a standardized normal distribution. Lastly, the covariance of the shared response Σ_s is diagonal and the diagonal values are obtained by sampling from a Dirichlet distribution with parameter (1...1). We set the number of voxels to $v = 125\,000$, the number of subjects to $m = 10$ and the number of components to $p = 50$. We generate $n = 1000$ samples.

We benchmark deterministic SRM, probabilistic SRM and their FastSRM counterparts in terms of fitting time and performance. Note that in this section, the identifiable implementation of deterministic SRM and probabilistic SRM described in Sections 2.2 and 2.1 is used. Algorithms are designated by the spatial decomposition they use and therefore SRM algorithms are referred to as None because no spatial decomposition is used and FastSRM algorithms will have the label *Optimal*. Note that it would be possible to use FastSRM with suboptimal spatial decompositions (there exists a wide variety of brain spatial decompositions (23,37,38) but using them does not bring any guarantee that the performance is the same as SRM.

We use a number of iterations between 1 and 100 and report the performance, fitting time, and a measure of convergence. In FastSRM, we do not compute the unmixing matrices but only the shared response. We measure the performance of an algorithm by computing the error between the true component $S \in \mathbb{R}^{p \times n}$ and the predicted component $\hat{S} \in \mathbb{R}^{p \times n}$ using the quantity:

$$\text{mse}(\hat{S}, S) = \min_{B \in \mathbb{R}^{p \times p}} \|B\hat{S} - S\|_F^2 = \|S\hat{S}^T\hat{S} - S\|_F^2 \quad (11)$$

This way of measuring errors is insensitive to the indeterminacies in DetSRM. We measure the fitting time in seconds. Lastly, we measure convergence by computing the gradient ℓ_∞ norm in case of DetSRM given by $\max(\text{abs}(G))$, where G is the gradient and use the distance between consecutive values of the loss for ProbSRM. Results are plotted in Figure 3.

We empirically see that the optimal approach is equivalent to using no spatial decomposition in terms of performance. This is predicted by the theory in Section 3.2, where we demonstrate that these two algorithms yield exactly the same output from a given input and initialization. We also see that ProbSRM gives much better results than DetSRM. In terms of fitting time, FastSRM is about a thousand times faster than SRM after 100 iterations. When no spatial decomposition is used, the number of iterations has a very strong impact on performance while it has a small impact when an optimal spatial decomposition is used. Lastly, looking at the convergence curves, we see that even after 100 iterations, algorithms did not fully converge. This means that in practice a larger number of iterations is needed, which would yield an even higher difference in fitting time between methods using no spatial decomposition and methods using an optimal spatial decomposition.

Experiment on fMRI data: identifiability increases stability

We evaluate the performance of the different SRM implementations on the *Sherlock* data sets where fMRI of 17 participants watching “Sherlock” BBC TV show (episode 1)

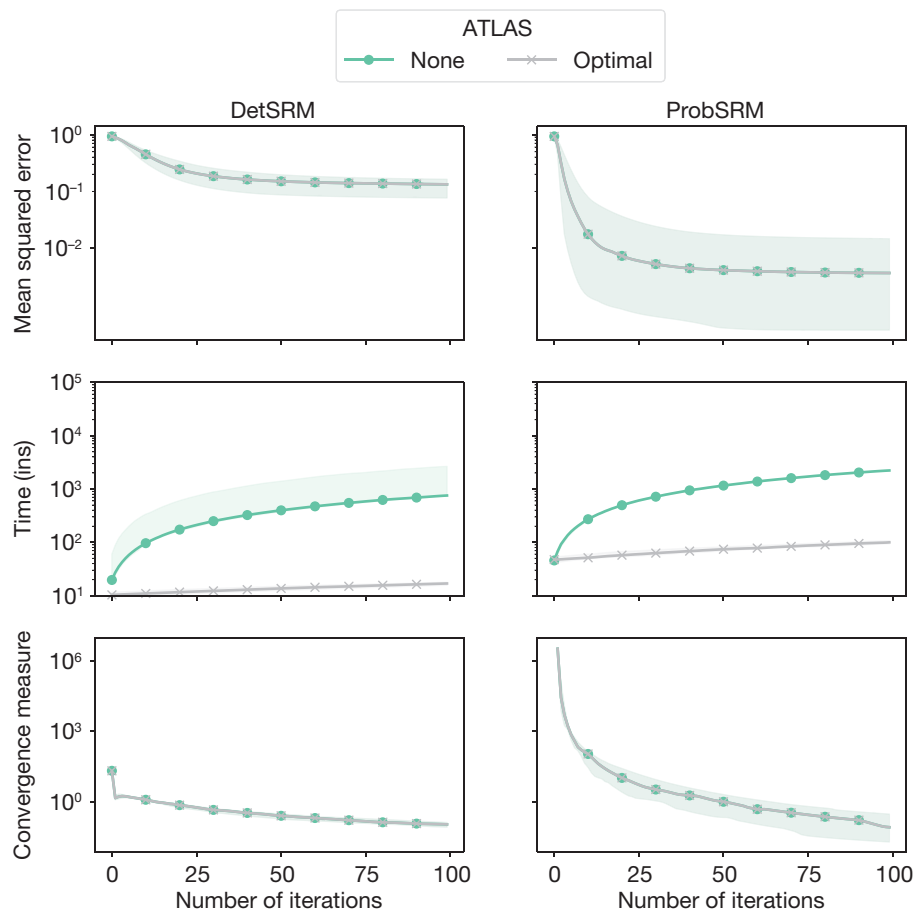


Fig. 3. Benchmark of SRM algorithms on synthetic data: Performance, fitting time and convergence of SRM algorithms in the deterministic (left) or probabilistic (right) case. As expected, when optimal spatial decompositions are used, the performance is the same as if no spatial decomposition is used but the fitting time is much lower. This is even more pronounced when the number of iterations is high (and looking at convergence curves, we see that more iterations could be performed to be closer to a stationary point).

is performed. These data are downloaded from <http://arks.princeton.edu/ark:/88435/dsp01nz8062179>. Data were acquired using a 3T scanner with an isotropic spatial resolution of 3 mm. More information including the preprocessing pipeline is available in (39). Subject 5 is removed because of missing data, leaving us with 16 participants. Although SHERLOCK data originally contain one run only, we split them into four blocks of 395 time frames and one block of 396 time frames for the needs of our experiments.

We first show that identifiability is a desirable property as it increases stability. Then, we show that our FastSRM algorithm (that works with optimal spatial decompositions) matches the performance of SRM (that works on full data) but uses much less computational resources.

Identifiability increases stability

Assuming that the data follow the model, identifiability ensures that the problem is well posed since the solutions are well characterized. More precisely, if the data are generated according to an identifiable model, in the

limit of infinite samples, a perfect optimization algorithm is guaranteed to recover the parameters used to generate the data. In contrast, in an unidentifiable model, there are more than one set of parameters that can generate the same data making the search for the good parameters ill defined. In practice, unidentifiable models have a more complex optimization landscape with many more local minima than identifiable models (any local minima in the identifiable model produces as many local minimas as they are indeterminacies in the corresponding unidentifiable model). This is why the unidentifiable models are therefore expected to be more sensitive to initialization and less robust to small changes in the data than their identifiable counterparts. Note however that SRM remains nonconvex so there is no guarantee that the optimal solution is found.

In real fMRI data, the model cannot be expected to hold perfectly, but we can hope for greater stability in the parameters recovered than if a nonidentifiable model is used.

To measure the stability of the shared components obtained from the Sherlock data set, we divide the subjects

into two roughly equal groups and extract the common components in each group. The components are then matched using the Hungarian algorithm (40) and the stability measure is obtained from the average correlation of the matched components. The procedure is repeated nine times using the Brainiak implementation (that is not identifiable since the shared components covariance is unconstrained) and our FastSRM implementation. We plot the results as a scatter plot in Figure 4, where, for each repetition, the x-axis indicates the stability measure obtained with Brainiak's implementation and the y-axis the stability measure obtained with FastSRM (identifiable since the shared components covariance is diagonal). We see that introducing a diagonal source covariance improves stability.

Comparing fitting time, memory usage and performance on a time-segment matching experiment

Time-segment matching has first been introduced in (12). In a nutshell, the time-segment matching accuracy measures the similarity between two multivariate time-series by trying to localize a time-segment in one time-series by correlation with the other. In the context of movie watching, this measure is quite meaningful: if we split the movies in scenes and compute a representation per scene and per subject, it can be assumed that different subjects watching the movie would still have closer representation

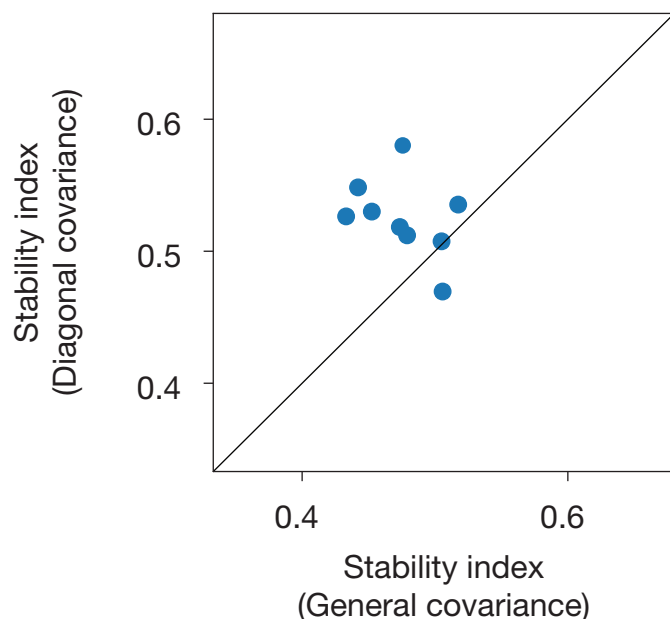


Fig. 4. Identifiability increases stability: We first divide the subjects into two groups and extract the common components in each group. The components of the two groups are then matched using the Hungarian algorithm, and the stability index is determined by the average correlation of the matched components. The procedure is repeated nine times with the Brainiak implementation (not identifiable since the shared components covariance is unconstrained) and our FastSRM implementation (identifiable since the shared components covariance is diagonal).

of the same scenes than of different scenes. This explains why time-segment matching is a standard evaluation of SRM-like methods also used in (41,42) or (20).

We now describe more precisely the experimental design. We split the runs into a train and test set. After fitting the model on the training set, we apply the unmixing matrices $W_i = A_i^{-1}$ of each subject on the test set yielding individual component matrices. We estimate the shared responses by averaging the individual components of each subject but one. We select a target time-segment (nine consecutive time frames) in the shared responses and try to localize the corresponding time-segment in the components of the left-out subject using a maximum-correlation classifier.

The time-segment is said to be correctly classified if the correlation between the sample and target time-segment is higher than with any other time-segment (partially overlapping time windows are excluded).

We use fivefold cross-validation across runs: the training set contains 80% of the runs and the test set 20%, and repeat the experiment using all possible choices for left-out subjects. The mean accuracy is reported in Figure 5 (top). When an optimal spatial decomposition is used, the accuracy is the same as when no spatial decomposition is used but the fitting time is reduced by a factor of 10–100 and so is the memory usage (see Figure 5, bottom).

IMPLEMENTATION AND REPRODUCIBILITY

Our work is fully reproducible, and our code is tested and well documented. All material is available in the <https://github.com/hugorichard/FastSRM> Github repository.

The API is reminiscent of the Scikit-learn one where the model object is instantiated with its parameters and then fitted on data via a method `fit()`. After the model is fit, the learned parameters can be accessed. Then, a `transform()` method computes the shared response from the learned parameters and the data passed as the argument of the `transform()` method.

While Brainiak and Nilearn are useful to reproduce some of our experiments, they are not strong dependencies and FastSRM only needs Numpy (≥ 1.12), Scipy ($\geq 0.18.0$), Matplotlib ($\geq 2.0.0$), Scikit-learn (≥ 0.23), Pytest ($\geq 6.2.5$) and Joblib ($\geq 1.1.0$).

The package comes with a documentation and a tutorial that reproduces, at a smaller scale, the experiment available in Figure 3. These are found at <https://hugorichard.github.io/FastSRM/index.html>. The instructions to fully reproduce the experiments presented in the papers are available in the README: <https://github.com/hugorichard/FastSRM/blob/master/README.md>.

A set of tests reduces the risk of introducing a bug. These tests are run any time a pull request or a merged is performed on the main codebase.

Dataset: SHERLOCK

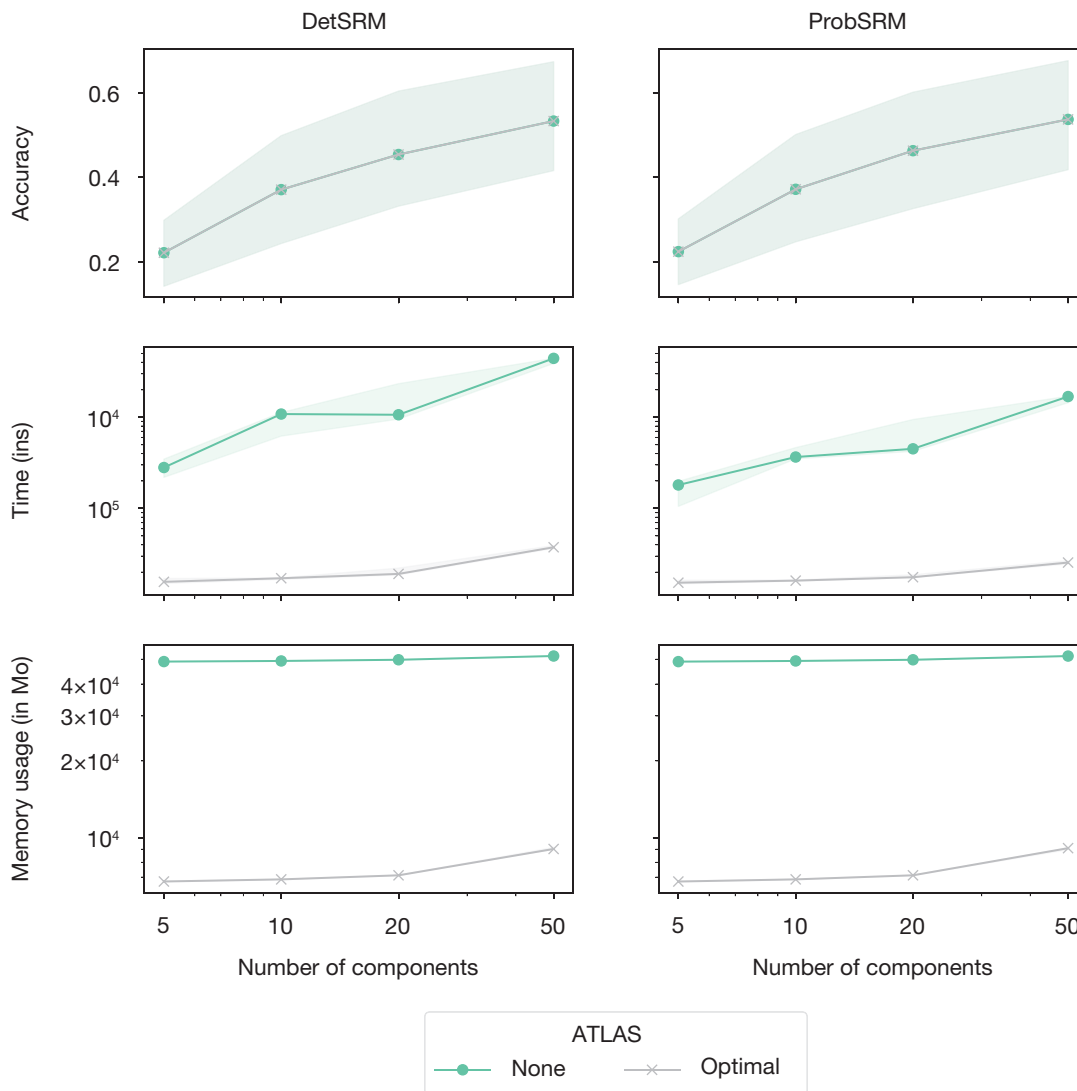


Fig. 5. Benchmark of SRM algorithms on fMRI data: Top: Time-segment matching accuracy. Middle: Fitting time. Bottom: Memory usage. When the optimal spatial decomposition is used, the accuracy is the same as when no spatial decomposition is used but the fitting time is reduced by a factor of 10–100 and so is the memory usage.

CONCLUSION

As studies using naturalistic stimuli will tend to become more common within and across subjects, we need scalable models in terms of computation time and memory usage. This is particularly related to the development of deep phenotyping models such as the Courtois project on Neural Modeling (see <https://docs.cneuromod.ca>). This is precisely what FastSRM provides.

FastSRM is an efficient implementation of SRM that uses optimal spatial decompositions to speed up computations and reduce memory requirements with provably no loss of performance.

We show on synthetic and real data that FastSRM is much faster and more memory efficient than implementations not using an optimal spatial decomposition. Furthermore, FastSRM is identifiable. On real data, we

show that compared to Brainiak’s implementation, which is not identifiable, FastSRM provides more stable estimates of the shared components.

FastSRM inherits all the weaknesses of SRM including the fact that mixing matrices are assumed to be orthogonal, which is debatable in practice. Another desirable property of such decompositions is the spatial smoothness of the components. While it is often observed that SRM components are smooth, nothing imposes such a constraint in the model. We highlight that the optimal spatial decomposition is only optimal with respect to the SRM model. More precisely, it gives a reduced representation such that applying the SRM model on the reduced representation or on the full data gives the same result. Whether this representation is interpretable or not is irrelevant here, since we only focus on accelerating the SRM

algorithm. However, as this representation is not unique, it may be interesting to find the “most interpretable” optimal spatial decomposition. Note that this would not change the end result of the SRM algorithm but could be of interest for interpretation purposes. It is an open question whether dimension reduction and learning of the shared components could be done jointly and efficiently without assuming orthogonal mixing matrices. Our approach based on optimal spatial decomposition shows that SRM-like methods do not need the full data to provide accurate estimates of shared components. We believe that such insights may guide the design of future methods.

Last, FastSRM is useful when the number of features is much larger than the number of samples. This setting is common in fMRI but less common in MEG/EEG where the number of samples is usually much larger than the number of features. Even in this case, it is possible to considerably speed up SRM. Indeed, SRM only depends on second-order statistics. These statistics can be pre-computed making the rest of the pipeline insensitive to the number of samples in the data set. An implementation of this simple idea would make SRM easily applicable to other modalities. We leave it to future work.

REFERENCES

1. Thomas Naselaris, Kendrick N Kay, Shinji Nishimoto, and Jack L Gallant. Encoding and decoding in fMRI. *Neuroimage*, 56(2):400–410, 2011.
2. Alexander G Huth, Shinji Nishimoto, An T Vu, and Jack L Gallant. A continuous semantic space describes the representation of thousands of object and action categories across the human brain. *Neuron*, 76(6):1210–1224, 2012.
3. Umut Güçlü and Marcel AJ van Gerven. Increasingly complex representations of natural movies across the dorsal stream are shared between subjects. *NeuroImage*, 145:329–336, 2017.
4. Hugo Richard, Ana Pinho, Bertrand Thirion, and Guillaume Charpiat. Optimizing deep video representation to match brain activity. *Computational Cognitive Neuroscience*, 2018.
5. Jean-Baptiste Poline and Matthew Brett. The general linear model and fMRI: Does love last forever? *Neuroimage*, 62(2):871–880, 2012.
6. Christian Jutten and Jeanny Hérault. Blind separation of sources, part I: An adaptive algorithm based on neuromimetic architecture. *Signal Processing*, 24(1):1–10, 1991.
7. H Richard, P Ablin, B Thirion, A Gramfort, and A Hyvarinen. Shared independent component analysis for multi-subject neuroimaging. In *Advances in Neural Information Processing Systems* 34, December 2021.
8. H Richard, L Gresele, A Hyvarinen, B Thirion, A Gramfort, and P Ablin. Modeling shared responses in neuroimaging studies through multiview ICA. In *Advances in Neural Information Processing Systems* 33, December 2020.
9. Gaël Varoquaux, Sepideh Sadaghiani, Jean-Baptiste Poline, and Bertrand Thirion. CanICA: Model-based extraction of reproducible group-level ica patterns from fMRI time series. *arXiv preprint arXiv:0911.4650*, 2009.
10. Vince D Calhoun, Tülay Adalı, Godfrey D Pearson, and James J Pekar. A method for making group inferences from functional MRI data using independent component analysis. *Human Brain Mapping*, 14(3):140–151, 2001.
11. Javier Via, Matthew Anderson, Xi-Lin Li, and Tülay Adalı. Joint blind source separation from second-order statistics: Necessary and sufficient identifiability conditions. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2520–2523. IEEE, 2011.
12. Po-Hsuan Chen, Janice Chen, Yaara Yeshurun, Uri Hasson, James Haxby, and Peter J Ramadge. A reduced-dimension fMRI shared response model. In *Advances in Neural Information Processing Systems*, pages 460–468, 2015.
13. Christopher Baldassano, Janice Chen, Asieh Zadbood, Jonathan W Pillow, Uri Hasson, and Kenneth A Norman. Discovering event structure in continuous narrative perception and memory. *Neuron*, 95(3):709–721, 2017.
14. Jonathan D Cohen, Nathaniel Daw, Barbara Engelhardt, Uri Hasson, Kai Li, Yael Niv, Kenneth A Norman, Jonathan Pillow, Peter J Ramadge, Nicholas B Turk-Browne, et al. Computational approaches to fMRI analysis. *Nature Neuroscience*, 20(3):304–313, 2017.
15. Christopher Baldassano, Uri Hasson, and Kenneth A Norman. Representation of real-world event schemas during narrative perception. *Journal of Neuroscience*, 38(45):9689–9699, 2018.
16. Eshin Jolly, Sushmita Sadhukha, and Luke J Chang. Custom-molded head-cases have limited efficacy in reducing head motion during naturalistic fMRI experiments. *NeuroImage*, 222:117207, 2020.
17. Caroline S Lee, Mariam Aly, and Christopher Baldassano. Anticipation of temporally structured events in the brain. *Elife*, 10:e64972, 2021.
18. Javier S Turek, Cameron T Ellis, Lena J Skalaban, Nicholas B Turk-Browne, and Theodore L Wilke. Capturing shared and individual information in fMRI data. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 826–830. IEEE, 2018.
19. Janice Chen, Yuan C Leong, Christopher J Honey, Chung H Yong, Kenneth A Norman, and Uri Hasson. Shared memories reveal shared structure in neural activity across individuals. *Nature Neuroscience*, 20(1):115–125, 2017.
20. Hejia Zhang, Po-Hsuan Chen, Janice Chen, Xia Zhu, Javier S Turek, Theodore L Wilke, Uri Hasson, and Peter J Ramadge. A searchlight factor model approach for locating shared information in multi-subject fMRI analysis. *arXiv preprint arXiv:1609.09f32*, 2016.
21. Hejia Zhang, Po-Hsuan Chen, and Peter Ramadge. Transfer learning on fMRI datasets. In *International Conference on Artificial Intelligence and Statistics*, pages 595–603, 2018.
22. Michael J Anderson, Mihai Capota, Javier S Turek, Xia Zhu, Theodore L Wilke, Yida Wang, Po-Hsuan Chen, Jeremy R Manning, Peter J Ramadge, and Kenneth A Norman. Enabling factor analysis on thousand-subject neuroimaging datasets. In *2016 IEEE International Conference on Big Data (Big Data)*, pages 1151–1160. IEEE, 2016.
23. Pierre Bellec, Pedro Rosa-Neto, Oliver C Lyttelton, Habib Benali, and Alan C Evans. Multi-level bootstrap analysis of stable clusters in resting-state fMRI. *NeuroImage*, 51(3):1126–1139, 2010.
24. Kamalaker Dadi, Gaël Varoquaux, Antonia Machlouzarides-Shalit, Krzysztof J Gorgolewski, Demian Wassermann, Bertrand Thirion, and Arthur Mensch. Fine-grain atlases of functional modes for fMRI analysis. *NeuroImage*, 221:117126, 2020.
25. Peter J Denning. Thrashing: Its causes and prevention. In *Proceedings of the December 9–11, 1968, Fall Joint Computer Conference, Part I*, pages 915–922, 1968.
26. Thomas Bazeille, Elizabeth Dupre, Hugo Richard, Jean-Baptiste Poline, and Bertrand Thirion. An empirical evaluation of functional alignment using inter-subject decoding. *NeuroImage*, 118:683, 2021.
27. James V Haxby, J Swaroop Guntupalli, Andrew C Connolly, Yaroslav O Halchenko, Bryan R Conroy, M Ida Gobbi, Michael Hanke, and Peter J Ramadge. A common, high-dimensional model of the representational space in human ventral temporal cortex. *Neuron*, 72(2):404–416, 2011.
28. Muhammad Yousefnezhad and Daoqiang Zhang. Deep hyperalignment. *arXiv preprint arXiv:1710.03923*, 2017.
29. Yi-Ou Li, Tülay Adalı, Wei Wang, and Vince D Calhoun. Joint blind source separation by multiset canonical correlation analysis. *IEEE Transactions on Signal Processing*, 57(10):3918–3929, 2009.
30. Thomas Bazeille, Hugo Richard, Hicham Janati, and Bertrand Thirion. Local optimal transport for functional brain template estimation. In *International Conference on Information Processing in Medical Imaging*, pages 237–248. Springer, 2019.
31. Alexandre Abraham, Fabian Pedregosa, Michael Eickenberg, Philippe Gervais, Andreas Mueller, Jean Kossaifi, Alexandre Gramfort, Bertrand Thirion, and Gaël Varoquaux. Machine learning for neuroimaging with scikit-learn. *Frontiers in Neuroinformatics*, 8:14, 2014.
32. Manoj Kumar, Michael J Anderson, James W Antony, Christopher Baldassano, Paula P Brooks, Ming Bo Cai, Po-Hsuan Cameron Chen, Cameron T Ellis, Gregory Henselman-Petrusek, David Huberdeau, et al. Brainiak: The brain imaging analysis kit. *Aperture*, 2020.
33. Charles R Harris, K Jarrod Millman, St’efan J van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al.. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020.
34. Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17:261–272, 2020.
35. John D Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.

36. Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
37. Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, XiNian Zuo, Avram J Holmes, Simon B Eickhoff, and B T Thomas Yeo. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cerebral Cortex*, 28(9):3095–3114, 2017.
38. Arthur Mensch, Julien Mairal, Bertrand Thirion, and Gaël Varoquaux. Extracting universal representations of cognition across brain-imaging studies. *arXiv preprint arXiv:1809.06035*, 2018.
39. J Chen, Y C Leong, K A Norman, and U Hasson. Shared experience, shared memory: A common structure for brain activity during naturalistic recall. *bioRxiv*, 2016.
40. Harold W Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1–2):83–97, 1955.
41. J Swaroop Guntupalli, Ma Feilong, and James V Haxby. A computational model of shared fine-scale structure in the human connectome. *PLoS Computational Biology*, 14(4):e1006120, 2018.
42. Samuel A Nastase, Yun-Fei Liu, Hanna Hillman, Kenneth A Norman, and Uri Hasson. Leveraging shared connectivity to aggregate heterogeneous datasets into a common response space. *bioRxiv*, 2019.

APPENDIX A: DETAILED DERIVATION OF THE PROBABILISTIC SRM ALGORITHM

Here, we describe in detail the log-likelihood underlying the probabilistic SRM and derive the update rules. We highlight that the derivations used to obtain these updates are not new. They are obtained following the strategy introduced in (12) and (22).

Denoting $\mathbb{V}[s | x] = (\sum_i \frac{1}{\sigma_i^2} I + \sum_s^{-1})^{-1}$ and $\mathbb{E}[s | x] = \mathbb{V}[s | x]$.

$\sum_i \frac{1}{\sigma_i^2} A_i^T X_i$, we have

$$p(x, s) = \prod_i \frac{\exp\left(-\frac{\|x_i - A_i s\|^2}{2\sigma_i^2}\right) \exp\left(-\frac{1}{2}\langle s, \sum_s^{-1} s \rangle\right)}{(2\pi\sigma_i^{2\nu})^{\frac{1}{2}} (2\pi|\sum_s|)^{\frac{1}{2}}} \quad (12)$$

$$= c_1 \exp\left[-\frac{1}{2}\left(\sum_i \frac{1}{\sigma_i^2} \|x_i\|^2 - 2\langle \sum_i \frac{1}{\sigma_i^2} A_i^T x_i, s \rangle\right)\right] \quad (13)$$

$$+ \sum_i \frac{1}{\sigma_i^2} \|s\|^2 + \langle s, \sum_s^{-1} s \rangle \quad (14)$$

$$= c_2(x) \exp\left(-\frac{1}{2}\langle (s - \mathbb{E}[s | x], \mathbb{V}[s | x]^{-1}(s - \mathbb{E}[s | x])) \rangle\right) \quad (15)$$

where

$$c_1 = \frac{1}{(2\pi\sigma_i^{2\nu})^{\frac{1}{2}}} \frac{1}{(2\pi|\sum_s|)^{\frac{1}{2}}} \quad (16)$$

and

$$c_2(x) = c_1 \exp\left[-\frac{1}{2}\left(\sum_i \frac{1}{\sigma_i^2} \|x_i\|^2 - \langle \mathbb{E}[s | x], \mathbb{V}[s | x]^{-1} \mathbb{E}[s | x] \rangle\right)\right] \quad (17)$$

are independent from s . Therefore,

$$s | x \sim N(\mathbb{E}[s | x], \mathbb{V}[s | x]) \quad (18)$$

The negative expected completed log-likelihood is given by

$$L = \sum_i \frac{1}{2} \nu \log(\sigma_i^2) + \frac{1}{2\sigma_i^2} \mathbb{E}[\|x_i - A_i s\|^2] \quad (19)$$

Updates are therefore given by

$$\sigma_i^2 \leftarrow \frac{1}{\nu} (\mathbb{E}[\|x_i - A_i \mathbb{E}[s | x]\|^2] + \|\text{diag}(\mathbb{V}[s | x])\|) \quad (20)$$

$$A_i \leftarrow P(\mathbb{E}[x_i \mathbb{E}[s | x]^T]) \quad (21)$$

$$\sum_s \leftarrow \mathbb{V}[s | x] + \mathbb{E}[\mathbb{E}[s | x] \mathbb{E}[s | x]^T] \quad (22)$$

It is useful to access the log-likelihood to check the implementation of the algorithm and monitor the convergence. From equation (15), the likelihood is given by

$$p(x) = c_2(x) \int_s \exp\left(-\frac{1}{2}\langle (s - \mathbb{E}[s | x], \mathbb{V}[s | x]^{-1}(s - \mathbb{E}[s | x])) \rangle\right) ds \quad (23)$$

$$= c_2(x) (2\pi |\mathbb{V}[s | x]|)^{\frac{1}{2}} \quad (24)$$

Replacing $c_2(x)$ by its expression and taking the log, the expected negative log-likelihood is (up to constants) given by

$$\begin{aligned} \mathbb{E}[-\log(p(x))] &= \sum_i \frac{\nu}{2} \log(\sigma_i^2) + \frac{1}{2} \log(|\sum_s|) - \frac{1}{2} \log(|\mathbb{V}[s | x]|) \\ &+ \sum_i \frac{1}{2} \frac{1}{\sigma_i^2} \mathbb{E}[\|x_i\|^2] - \frac{1}{2} \mathbb{E}[\langle \mathbb{E}[s | x], \mathbb{V}[s | x]^{-1} \mathbb{E}[s | x] \rangle] \end{aligned} \quad (25)$$

APPENDIX B: IDENTIFIABILITY RESULTS ON PROBABILISTIC SRM AVAILABLE IN (8) AND APPENDIX D

Proposition 3. Probabilistic SRM with positive definite \sum_s is not identifiable

Proof. Let us assume the data x_i , $i = 1, \dots, m$ follow the probabilistic SRM with parameters \sum_s, A_i, σ_i^2 , $i = 1, \dots, m$ and \sum_s definite positive.

Let us consider an orthogonal matrix $O \in O_k$. We call $A'_i = A_i O$ and $\sum'_s = O^T \sum_s O$. \sum'_s is trivially symmetric positive definite.

Then, the data also follow the SRM with different parameters $\sum'_s, A'_i, \sigma_i^2$, $i = 1, \dots, m$

Proposition 4. We consider the probabilistic SRM where \sum_s is a positive diagonal matrix. We further assume that the values in \sum are all distinct and ranked in ascending order. Then, the probabilistic SRM is identifiable up

to sign indeterminacies on the columns of $\begin{bmatrix} A^1 \\ \vdots \\ A^m \end{bmatrix}$.

Proof. The probabilistic SRM can be written

$$x_i \sim N(0, A_i \sum_s A_i^T + \sigma_i^2 I_\nu) \text{ with } A_i^T A_i = I_p$$

where \sum_s is a positive diagonal matrix with distinct values ranked in ascending order.

Let us assume the data x_i , $i = 1, \dots, m$ follow the probabilistic SRM with parameters \sum_s, A_i, σ_i^2 , $i = 1, \dots, m$. Let us further assume that the data x_i , $i = 1, \dots, m$ follow the probabilistic SRM with another set of parameters $\sum'_s, A'_i, \sigma_i^2$, $i = 1, \dots, m$.

Since the model is Gaussian, we look at the covariances. We have for $i \neq j$

$$\mathbb{E}[x_i (x_j)^T] = A_i \sum_s A_j^T = A'_i \sum'_s A_j'^T,$$

The singular value decomposition is unique up to sign flips and permutation. Since eigenvalues are positive and ranked, the only indeterminacies left are on the eigenvectors. For each eigenvalue, a sign flip can occur simultaneously on the corresponding left and right eigenvector.

APPENDIX D: PROOFS OF PROPOSITION 2

Therefore, we have $\Sigma'_s = \Sigma_s$, $A_i = A'_i D_{ij}$ and $A_j = A'_j D_{ij}$, where $D_{ij} \in \mathbb{R}^{k,k}$ is a diagonal matrix with values in $\{-1, 1\}$. This analysis holds for every $j \neq i$ and therefore $D_{ij} = D$ is the same for all subjects.

We also have for all i

$$\mathbb{E}[X_i(x_j)^T] = A_i \Sigma_s A_i^T + \sigma_i^2 I_v = A'_i \Sigma'_s A_j^T + \sigma_i^2 I_v$$

We therefore conclude $\sigma_i^2 = \sigma_i'^2, i = 1 \dots m$.

Note that if the diagonal subject-specific noise covariance $\sigma_i^2 I_v$ is replaced by any positive definite matrix, the model still enjoys identifiability.

APPENDIX C: PROOFS OF PROPOSITION 1

Proof. Updates of the mixing matrices A_i in deterministic SRM equation (5) can be written as follows:

$$A_i \leftarrow P \left(\frac{1}{n} X_i S^T \right) = U_{x_i} P \left(\frac{1}{n} Z_i S^T \right) \quad (26)$$

where P is the projection on the Stiefel manifold:

$$P(M) = M(M^T M)^{-\frac{1}{2}}$$

Therefore, we can look for A_i as $A_i = U_{x_i} \tilde{A}_i$. We can show that \tilde{A}_i is orthogonal. Indeed,

$$A_i^T A_i = I_p \quad (27)$$

$$\Rightarrow \tilde{A}_i^T U_{x_i}^T U_{x_i} \tilde{A}_i = I_p \quad (28)$$

$$\Rightarrow \tilde{A}_i^T \tilde{A}_i = I_p \quad (29)$$

Then, we use the fact that

$$\|X_i - A_i S\|^2 = \|U_{x_i} Z_i - U_{x_i} \tilde{A}_i S\|^2 = \|Z_i - \tilde{A}_i S\|^2 \quad (30)$$

so that $A'_i = \tilde{A}_i$.

Therefore, the solution of deterministic SRM on data $(z_i)_{i=1}^m$ and $(x_i)_{i=1}^m$ is linked by the change of parameters $\tilde{A}_i = A'_i$ and $S = S'$. This concludes the proof.

Proof. Updates of the mixing matrices A_i in probabilistic SRM equation (21) can be written as follows:

$$A_i \leftarrow U_{x_i} (\mathbb{E}[z_i \mathbb{E}[s | x]^T]) \quad (31)$$

so we can look for A_i as $A_i = U_{x_i} \tilde{A}_i$ and, as in the deterministic case, \tilde{A}_i is orthogonal. Therefore, equality (30) holds.

Then, we consider the expected negative log-likelihood of probabilistic SRM:

$$L = \sum_i \frac{1}{2} v \log(\sigma_i^2) + \frac{1}{2} \log(|\Sigma_s|) + \mathbb{E}[\int_s \sum_i \frac{1}{2\sigma_i^2} \|x_i - A_i s\|^2 + \frac{1}{2} \langle s, \Sigma_s^{-1} s \rangle ds] \quad (32)$$

$$= \sum_i \frac{1}{2} v \log(\sigma_i^2) + \frac{1}{2} \log(|\Sigma_s|) + \mathbb{E}[\int_s \sum_i \frac{1}{2\sigma_i^2} \|z_i - \tilde{A}_i s\|^2 + \frac{1}{2} \langle s, \Sigma_s^{-1} s \rangle ds] \quad (33)$$

where we use equality (30). Optimizing the log-likelihood via expectation maximization yields the exact same updates as probabilistic SRM on data z_i except that updates

$$\sigma_i^2 \leftarrow \frac{1}{n} (\mathbb{E}[\|z_i - \tilde{A}_i \mathbb{E}[s | z]\|^2] + \|\text{diag}(\mathbb{V}[s | z])\|^2) \quad (34)$$

are replaced by updates

$$\sigma_i^2 \leftarrow \frac{1}{v} (\mathbb{E}[\|z_i - \tilde{A}_i \mathbb{E}[s | z]\|^2] + \|\text{diag}(\mathbb{V}[s | z])\|^2) \quad (35)$$

so that $\tilde{A}_i = A'_i$. Therefore, the updates in both algorithms are linked by $A_i = U_{x_i} A'_i$, $\sigma_i^2 = \sigma_i$ and $\Sigma'_s = \Sigma_s$. This concludes the proof.