Spatial Patterns for Discriminative Estimation (SP♠DE)[1]

A. Llera[1,2,3], R. Chauvin[1,2,3], P. Mulders[1,2,3,4], J. Naaijen [1,2,3], M. Mennes[1,2,3], C. F. Beckmann[1,2,3,5]

1- Radboud University Nijmegen, Comeniuslaan 4, 6525 HP, Nijmegen, The Netherlands.

2- Donders Institute for Brain, Cognition and Behavior, Centre for Cognitive Neuroimaging, 6500 GL, Nijmegen, The Netherlands.

3- Department of Cognitive Neuroscience, Radboud University Medical Centre, 6525 EN, Nijmegen, The Netherlands.

4- Department of Psychiatry, Radboud University Medical Centre, 6500 HB, Nijmegen, The Netherlands.

5- Oxford Centre for Functional Magnetic Resonance Imaging of the Brain (FMRIB), University of Oxford, Oxford, OX3 9DU, United Kingdom.

Abstract

Functional connectivity between brain regions is modulated by cognitive states or experimental conditions. Such connectivity variations can be appreciated in covariance or correlation matrices obtained from fMRI data. A multivariate methodology that can capture fMRI connectivity maps in light of different experimental conditions would be of primary importance to learn about the specific roles of the different brain areas involved in the observed connectivity variations. Here we introduce Spatial Patterns for Discriminative Estimation (SP♠DE), a supervised dimensionality reduction model that provides a full multivariate characterization by achieving optimal discriminative linear spatial filters in terms of variance and provides interpretable spatial maps directly reflecting the brain areas involved in the accounted connectivity changes. We demonstrate the strength of SPADE using fMRI data from the Human connectome project to show that the model provides close to perfect discrimination between different fMRI tasks at low dimensionality. Additionally, the straightforward interpretability of the model is demonstrated by the obtained linear filters relating to anatomical areas well known to be involved in each considered fMRI task. Further, we also show that such an approach provides an alternative view to traditional task fMRI analyses by looking at changes in the covariance structure as a substitute to changes in the mean signal as the general linear model (GLM) analyses. We use 2-back and 0-back working memory task fMRI data from the Human connectome project to show that SPADE discovers connectivity changes in memory related areas during the 2-back periods, while attention network areas are involved during the 0-back task periods. We conclude that SPADE is a robust tool to investigate brain connectivity alterations across induced cognitive changes and has the potential to be used in pathological or pharmacological cohort studies.

Introduction

Functional magnetic resonance imaging (fMRI) is the primary tool used to investigate how the brain (re)acts under different mental states and the biological underpinnings of brain disorders. Most commonly, these state or group-differences are investigated using generalized linear models (GLM) [1] to detect brain regions that differentiate between experimental conditions. The GLM is for example the standard tool used to analyze task fMRI data incorporating a temporal model for the induced brain activity [2] that allows obtaining a supervised voxel-wise characterization of changes in the blood-oxygen-level dependent (BOLD) response [1]. Yet, while powerful and straightforward to interpret, the GLM has two crucial disadvantages. Foremost, while fMRI data is multivariate in nature, the GLM is applied as a mass univariate approach (i.e., independently to each voxel in the brain) and does not take into account interactions between different regional effects. Second, it ignores the variance in the data by looking at single mean BOLD changes across conditions. To overcome these issues, one can use multivariate models that allow studying the second order interactions between different brain areas [3,4]. Although these models overcome some of the issues with univariate modelling, they are not specifically

[1] Manuscript under review by NeuroImage

tailored for state/group characterization. Further inclusion into models for state/group identification, i.e. regression or classification, makes the interpretation of such results difficult.

Consequently, to improve interpretability of results obtained from analyzing multivariate fMRI data while optimizing discrimination between two conditions or groups of fMRI time-series, it would be beneficial to use a multivariate approach to extract brain patterns that maximize discriminative properties. To that end we introduce a model that extracts Spatial Patterns for Discriminative Estimation (SPADE). SPADE is a covariance based discriminative approach that provides optimal linear discriminative filters in terms of variance [5]. The basics of this approach are commonly used in EEG-based brain-computer interfaces [6,7]. Here, we expand its original formulation with a permutation testing statistical approach for model order selection. To our knowledge this method has not been applied to fMRI data yet as it requires a large number of observations to compute well-posed covariance matrices. However, with recent advances in data acquisition and large-scale data collection it is now possible to collect a large number of observations in relatively short time periods, thus enabling application of this methodology to fMRI data.

Here, we make use of high-quality fMRI data from the Human Connectome Project (HCP) [8,9] to uncover the potential power of our method for discriminative estimation. We consider resting state, motor, and working memory task fMRI data from over 1000 subjects and demonstrate that we are able to learn low-dimensional representations of the data that provide close to perfect discrimination between the different mental states. Furthermore, the method is fast and efficient, and the learned filters are robust, specific, and relate to brain areas well-known to be involved during their respective paradigms [10,11], making it useful to address various research questions.

Methods

Dataset & preprocessing

In this work we use resting state and task fMRI data from the Human Connectome Project (HCP) [8,9]. For each subject, we consider the motor task (MT), the working memory task (WM), and the first of the two available resting state sessions (RS). We use the HCP1200 release which includes data from 1200 subjects and only consider data from subjects for whom all three data modalities (RS, MT and WM) were available resulting in a total of 1063 subjects. In all cases we used the fully preprocessed data in MNI152 space as delivered in the HCP1200 release. For the full preprocessing details we refer the reader to [12].

Spatial Patterns for Discriminative Estimation (SP♠DE)

SPADE indexes connectivity changes from two sets of fMRI data using the simultaneous diagonalization of two covariance matrices as originally introduced in (Fukunaga, 1990); full details are provided in the supplementary material. Although ideally one would like to compute full brain spatial covariance matrices, this is not computationally achievable yet due to memory constraints, and a spatial dimensionality reduction must be performed in order to apply the algorithm to fMRI data. For each subject and for each of the three considered fMRI data modalities we performed a spatial dimensionality reduction into 165 regions of interest (ROI) from a functional parcellation [13] by extracting the mean time-series across the voxels in each ROI. Each ROI time-series was then independently demeaned and divided by its standard deviation before further processing.

Then, considering such ROI time-series gathered under two different fMRI modalities (e.g., resting state versus working memory task fMRI), we compute a unique covariance matrix per modality, and use these two covariance matrices to estimate their simultaneous diagonalization [5]. This results in a set of discriminative linear spatial filters that optimally separate the two initial modalities in terms of variance. We address the model order selection, i.e. the number of spatial filters selected to optimize discrimination, using permutation testing. Briefly, we select the filters that explain significantly more variance for a given condition than the variance obtained when randomizing the data across conditions. Full details are provided in the supplementary material.

An important feature of the SPADE model is that connectivity changes can be summarized for each basis vector as a unique spatial map that provides a spatial weight for each ROI. This results in a straightforward interpretation of the relevant changes in connectivity. Since the simultaneous diagonalization assumes no implicit noise model, the estimated basis vectors cannot be directly interpreted back to the brain, and the interpretable associated spatial maps are obtained using structural coefficients [10,11]. Brief details are provided in the supplementary material and for a more extensive explanation we refer the reader to [10].

For the remainder of the paper we will denote the full process involving spatial dimensionality reduction of two fMRI data modalities, simultaneous diagonalization, model order estimation by permutation testing and estimation of the associated interpretable spatial maps as SPADE. A toolbox providing full automatized estimation will be made publicly available upon publication.

SPADE evaluations

We use the resting state (RS), motor task (MT) and working memory task (WM) fMRI from the HCP sample and apply the introduced SPADE methodology to find brain connectivity differences between each pair of fMRI modalities independently: RS vs MT, RS vs WM and MT vs WM. To quantify the quality of the learned filters we perform a discriminative analysis to distinguish between fMRI modalities using a ten-fold cross validation approach for each of the three scenarios (RS vs MT, RS vs WM and MT vs WM) independently. At each fold, SPADE discriminative filters are learned from the covariance matrices of the two selected modalities using data from 90% of the subjects and the number of filters is selected using permutation testing. Then we compute the projection of these subjects' data into the newly learned basis, compute the logarithmic variance of the resulting time-series as features, and use them to train a Linear Discriminant Analysis (LDA) classifier [14] to distinguish between the two included modalities (e.g. RS vs MT). The filters are then used to project the remaining unseen data (~10% subjects) and extract the log-variance of these projections as features to test the classifier quality. Each classifier is evaluated using the accuracy which we define as the percentage of correctly classified samples. For comparison we also evaluate the results obtained using a linear support vector machine (SVM) classifier in the full covariance space, i.e. feature space of dimension (165*166)/2= 13695. Finally, we visually evaluate the spatial extent of the SPADE linear filters using structural coefficients for visualization of the associated spatial maps [10,11]. For simplicity we only report the top four filters at each experiment. The full set of significant filters will be uploaded as supplementary material.

In addition, we perform a secondary evaluation where we apply the SPADE model to a scenario previously studied using a GLM analysis. Using the WM task data from the HCP sample, we compared the 0-back memory task periods to the 2-back memory task periods using the SPADE approach. In this case we report the spatial maps associated to the top four spatial maps for comparison to the results presented in [15]. The full set of spatial filters will also be provided as supplementary material.

Results

To illustrate the discriminative performance of the SPADE methodology, Figure 1 presents two-dimensional representations of fMRI data obtained using the SPADE model. For each comparison between two fMRI modalities, we plotted one random fold from the ten-fold-cross-validation used to evaluate discrimination and divide the data into training and testing datasets. Note that by construction, the first and the last SPADE filters maximize variance for one class while minimizing it for the other one. Consequently, the discriminative power of the algorithm usually comes from the combination of pairs of filters at both extremes of the eigen-spectrum. In the first row we can appreciate, for each of the three comparisons, a very clear two-dimensional separation between the training features obtained by combining the two extreme SPADE filters. The second row clearly shows the generalization ability of the SPADE model on the test-set data. We observed that in all cases the classification using uniquely one pair of filters (two-dimensional feature space) is very high, between 98.6% and 99.76% respectively; exact statistics are provided in the second column of Table 1.

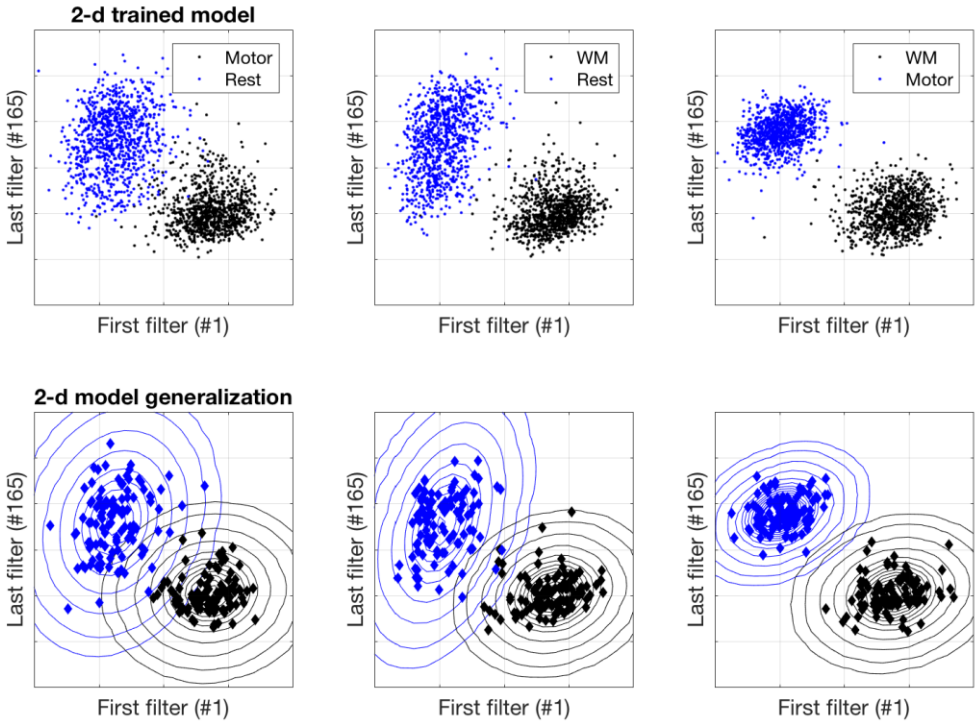

Figure 1: Two-dimensional fMRI data representations obtained using SPADE. The columns represent three different comparisons, MT vs RS, WM vs RS and WM vs MT task fMRI respectively. In the first row we present the logarithmic variance of the training fMRI data projected onto the first and last SPADE filters learned to discriminate each pair of fMRI modalities as indicated in the figure legends. Dots represent subjects and blue and black color encode the two different fMRI tasks considered on each scenario. In all three cases we can see that the training set is correctly separated in two visually distinguishable sets. In the second row we represent the learned two-dimensional training feature distributions presented in the first row as two Gaussian densities, and visualize the logarithmic variance of the testing fMRI data projections into the new basis as color coded diamonds. Appreciate the generalization ability of the SPADE model that provides an optimal two-dimensional classification space.

However, to fully evaluate the discriminative performance of the SPADE model we applied model order selection via permutation testing and selected filters at significance Bonferroni corrected level $p<0.05/165 = 3.3\text{x}10^{-4}$ (see previous section and supplementary material for complete details).

Although in practice the dimensionality was computed on each fold independently, for illustration we provide the exact set of filters selected when using the full sample at each comparison in Table 1 column 'SPADE selected filters'. On each of the three comparisons the classification performance obtained using SPADE at the dimensionality selected using the permutation testing strategy is presented in Table 1 column 'SPADE accuracy'. For comparison we also provide in the last column of Table 1 the classification results obtained using a SVM in the full covariance space. Although the mean improvement with respect to the two-dimensional scenario or the SVM is relatively small, the improvement it is clearly notable in the reduced standard deviation of the results across folds. These results highlight that the top filters are the most relevant for discrimination but using the model order selection generalizes better across folds. Further, the comparison with the SVM provides remarkable results, especially in light of the reduced dimension in which SPADE classification is performed and the linear nature of the SPADE filters that allow for a straightforward spatial interpretation.

Comparison	SPADE-2D accuracy	SPADE Selected filters	SPADE accuracy	SVM accuracy
RT vs MT	98.64 (0.95) %	1-6 and 155-165	99.3 (2*10^-3) %	98.9 (0.37) %
RT vs WM	99.34 (0.33) %	1-2 and 152-165	99.77 (6*10^-5) %	99.8 (0.40) %
WM vs MT	99.76 (0.46) %	1-4 and 153-165	99.63 (2*10^-4) %	99.58 (0.35) %
mean	99.24 (0.58) %		99.56 (1.3*10^-3) %	99.42 (0.37) %

Table 1: Summary results of the SPADE analyses. The second row presents results when comparing resting state (RT) with motor task (MT) fMRI, the third row shows RT with working memory (WM) task fMRI, and the forth row compares WM and MT fMRI. The fifth row shows the mean results across the three comparisons. The second column (SPADE 2-D) presents the classification results when discriminating each pair of fMRI modalities using a SPADE two-dimensional projection of the data. The column 'SPADE selected filters' shows the eigen-spectrum location of the filters selected by the model order selection. The column 'SPADE accuracy' shows the classification results obtained using the model order selection and the column 'SVM accuracy' shows the results obtained using a SVM in full covariance space (i.e. feature space of dimension (165*166)/2= 13695).

To evaluate the spatial distributions of the filters associated to the discriminative results reported, Figure 2 summarizes the spatial maps obtained in all three comparisons; the full set of maps will be uploaded as supplementary material. Since each fold provides slightly different discriminative filters, for visual validation we performed SPADE analyses including all subjects in the learning phase. The obtained spatial filters were then transformed into interpretable spatial maps [10,11].

When looking at differences between resting state and motor task [16], in Figure 2 A) we observe that areas synchronizing to maximize variance for the motor task (left side, filters 1 and 2) involve bilateral supplementary motor cortex, caudate and operculum in the first filter, while primary motor cortex and contralateral cerebellar cortex involvement is present in the second filter. Areas identified by these filters are well known to be related to motor preparation, execution and control [17,18]. On the other side of the spectrum (right side, filters 164 and 165) we find the filters maximizing variance for resting state data while minimizing it for motor task data. The last filter (#165) involves sensory motor cortex, bilateral dorso-lateral PFC (attention network), cerebellum and several temporal areas including the temporal gyrus and temporal pole. On the other side, the penultimate filter (#164) resembles the default mode network (DMN), anterior cingulate cortex and thalamus.

When looking at the differences between resting state and working memory task, in Figure 2 B) we observe that the areas synchronizing to maximize variance for the working memory task in the first filter involve areas well known to working memory performance [19,20] as visual and precuneus area, similar to the dorsal pathway, and temporal-occipital areas, as present in the ventral pathway [21,22]. Filter number two involves areas similar to number one but includes bilateral hippocampus and thalamus and more temporal areas. On the other side of the eigen-spectrum, we find that the last filter (#165) involves the anterior cingulate cortex, cerebellum, frontal pole, thalamus, middle temporal gyrus and temporal pole, frontal opercula. Further, the penultimate filter (#164) involves again the DMN and the thalamus. We observe that the filters relating to increased variance during rest (Figure 2A and 2B) are in both cases informing about a stronger involvement of default mode network (DMN) related areas during the resting state and are extremely similar in both independent analyses (Figure 2 A and B right sides), showing the consistency of the SPADE algorithm.

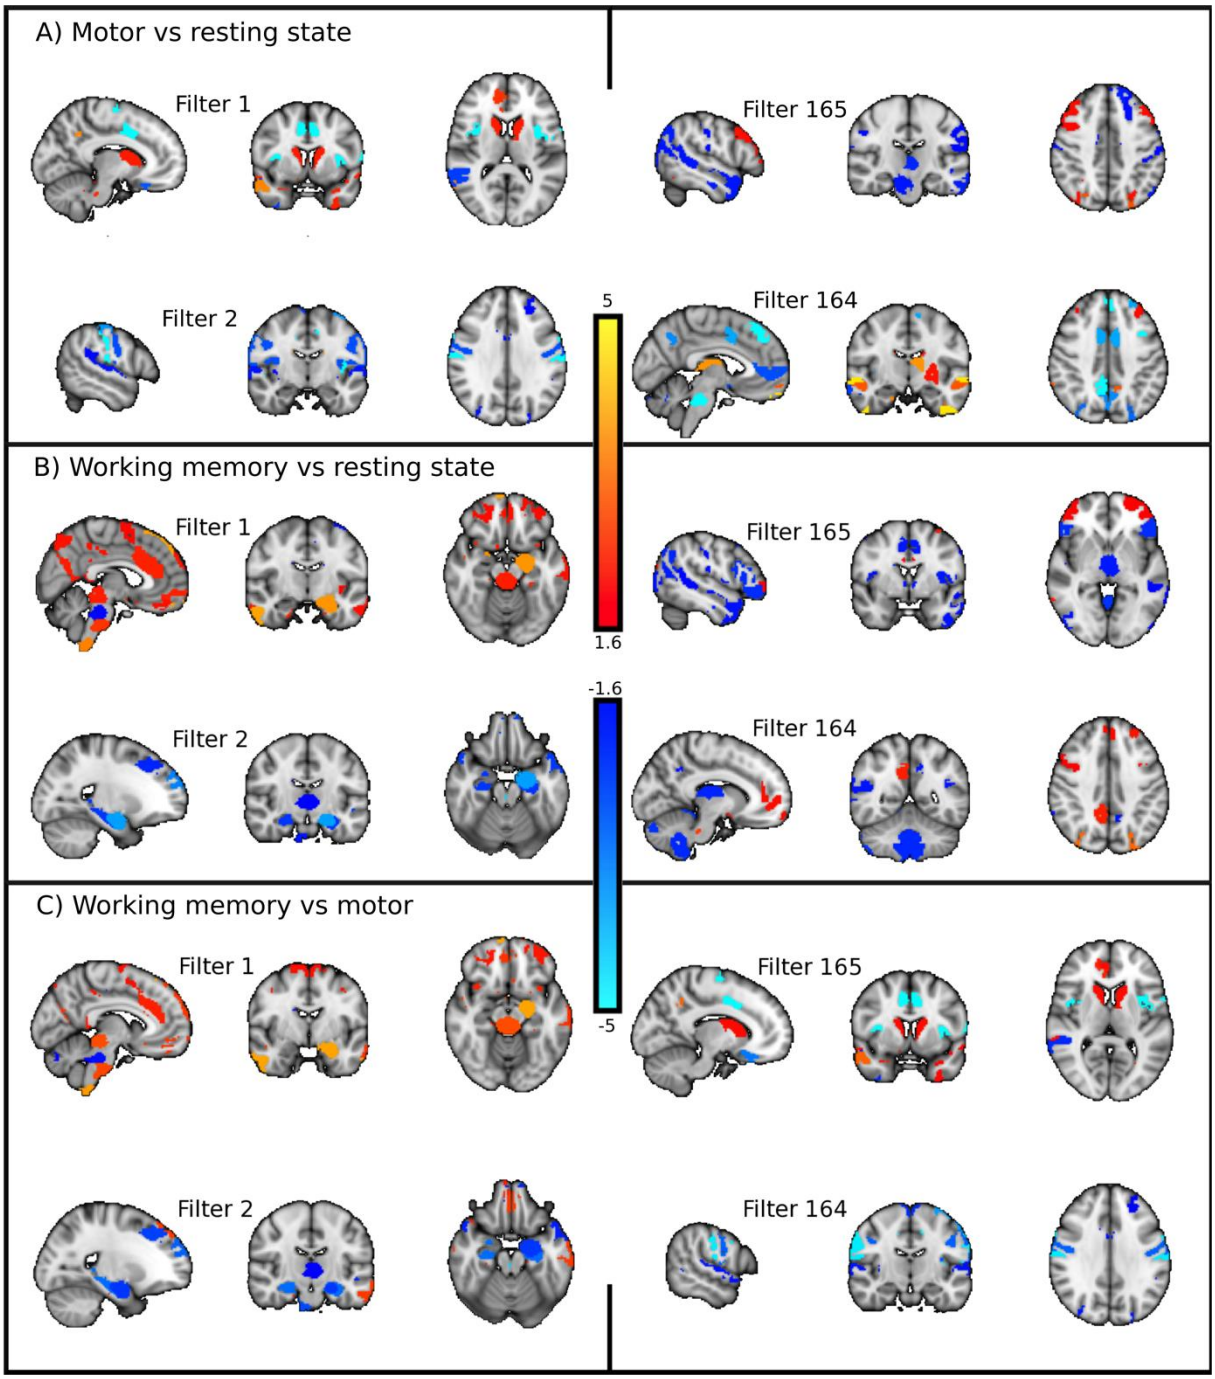

Figure 2: A) Resting state vs. motor task. B) Resting state vs. working memory task. C) Working memory vs. motor task fMRI. For each of the three scenarios we present a summary of the spatial maps associated with the discriminative spatial filters. On the left side we present filters 1 and 2, which maximize variance for A) motor, B) memory and C) memory. On the right side we present filters 165 and 164, which maximize variance for A) resting state, B) resting state and C) motor task. For visualization the value of the maps is converted to pseudo-z-statistics and thresholded at Z=1.6.

When looking at the differences between working memory and motor task (Figure 2 C), we observe that the spatial maps obtained are also very closely related to the ones that differentiated the tasks from the resting state; in particular, filters maximizing for WM (Figure 2 C, filters 1 and 2) resemble the ones maximizing variance for WM with respect to resting state (Figure 2 B filters 1 and 2); and filters maximizing variance for MT (Figure 2 C, filters 164 and 165) exhibit a similar spatial distribution to the ones maximizing variance for MT with respect to resting state (Figure 2 A filters 1 and 2). These results show that the filters extracted for each

modality involve areas well known to be required for the particular task performed and are highly specific since they can be obtained independently of the other covariance structure considered.

In addition to the previously described comparisons, we also performed another validation of the model by considering an analysis where we use the WM task fMRI data and compare the 2-back working memory task to the 0-back working memory task. A common approach to this type of data is a traditional example of task fMRI data analysis where conditions are compared by means of a GLM analysis, and the results of such analyses in this same dataset have been previously reported in [15]. In this study increased activation in the lateral-prefrontal and dorsal parietal cortex was associated with 2-back vs. 0-back working memory task differences. In Figure 3 we present a summary of the spatial maps learned using the SPADE model. The filter providing most variance for the 2-back task, i.e. filter 1, involves left thalamus and hippocampus, bilateral putamen and caudate, temporal pole and language areas. The next filter providing most variance for the 2-back task, i.e. filter 2, involves left thalamus, bilateral hippocampus, brainstem, cingulate and paracingulate gyrus, bilateral inferior frontal gyrus and left inferior temporal gyrus. On the other side of the eigen-spectrum, the filter maximizing variance for the 0-back task, i.e. filter 165, involves brainstem, cingulate gyrus, left thalamus, frontal pole and temporal gyrus and the next strongest filter for the 0-back task, i.e. filter 164, includes bilateral thalamus, brainstem, cingulate gyrus, frontal pole and motor cortex. Summarizing, the networks identified for the 2-back task involve areas purely related to working memory task while the 0-back task involves areas more related to the attention and salience network. This clearly reflects that the 2-back task involves stronger memory processing than the 0-back case which involves sustained attention.

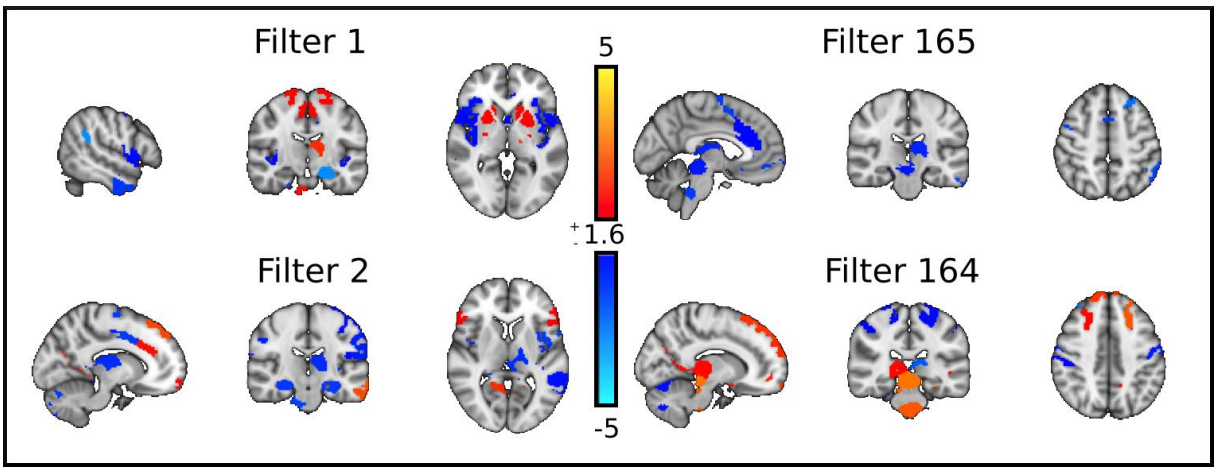

Figure 3: Spatial maps obtained using SPADE to compare the 2-back against 0-back working memory fMRI task data. Left side: filters 1 and 2 show the filters maximizing variance for 2-back memory task. Right side: filters 164 and 165 show the filters maximizing variance for 0-back memory task.

In general our results show some overlap with [15]. More specifically, some areas of filters 1 and 2 are present in both analysis. However, SPADE also reveals additional discriminative areas not found by the direct univariate analysis, including for example the striatum or the superior frontal gyrus.

Discussion

In this work we showed that we can achieve optimal linear filtering of fMRI data by applying an initial spatial dimensionality reduction followed by simultaneous diagonalization of covariance matrices [5,6]. We also introduce a permutation strategy for model order selection, i.e. how many spatial filters optimize discrimination, and we denote the fully developed model as SPADE (Spatial Patterns for Discriminative Estimation). We validated the use of this strategy by using the high-quality HCP data from 1063 subjects. The presented analyses illustrate the power of SPADE filters for discrimination of fMRI data, and although accurate classification can be achieved also using (for example) an SVM in a high-dimensional space (full covariance space), the power of the presented model is the intuitive interpretation of the results that are summarized as a small set of spatial maps that can be used to further develop our understanding of brain function. Contrary to the SVM, SPADE provides simple features that discriminate between the compared conditions and which we

can easily visualize back to brain space. This is certainly a big advantage with respect to discrimination in full covariance, correlation/partial correlation space where the interacting factors between different spatial nodes make the interpretation more difficult.

Our analyses revealed a strong specificity of the spatial distribution of different fMRI modalities across the three considered modalities (rest, motor and working memory). Each of the discriminative spatial maps obtained through SPADE reflects brain areas well-known to be relevant for a particular modality, reflecting for example that a stronger engagement of the DMN is recorded during rest than during task fMRI. In the working memory tasks, filters reveal independent effects for the ventral-dorsal pathways; the sensory integration is represented in one filter and the higher-order memory storage and utilization in another filter. Indeed, these two pathways start by sending sensory information to the parietal and temporal areas, characterizing spatial versus categorical features related to the items to memorize. Both connect to the temporo-medial structure (TMS) via specific connectivity to the parahippocampus, enthorinal cortex, before converging on the hippocampus [23]. The TMS then connects through the temporal pole to the (prefrontal, pole) frontal and (anterior) cingulate to store and use the information. Regarding working memory, we observe the phonological loop with both Wernicke and Broca areas [24–26]. SPADE is extracting differentiating filters out of these networks, demonstrating the existence of independent dynamics that can be the subject of further studies.

Ultimately, the SPADE model can extract regions of connectivity variation, obtained across different scenarios using fMRI data. We have shown its ability when using big data by demonstrating that it is able to pick up on areas relevant to each fMRI scenario but also by showing its close to optimal discrimination between representative fMRI experiments. We consider that the model order estimation, as well as the classification generalization of the model, can be used as a proxy of the specificity of the resulting spatial maps. As a consequence, the quality of the spatial maps reported here for interpretation is very high. Clearly, this occurs due to the strong differences in paradigms that we differentiate; for example, resting state from motor task fMRI.

To show the performance of SPADE to discriminate even between relatively similar paradigms, we applied it to compare 2-back working memory from 0-back working memory task fMRI data. We showed that the 2-back task involves areas well known to be involved in working memory while the 0-back did not involve for example the hippocampus, but instead, the identified networks involved areas related to the salient and attention networks, as the frontal pole or the dorsal anterior cingulate gyrus. In comparing our results with those presented in [15], we note that our analysis point to connectivity changes between 2 and 0-back task at areas known to be involved in working memory tasks that were not clearly identified using the GLM approach in [15]. This is not a surprising result since both analyses are looking at different statistics in the data; while [15] reports changes in the mean fMRI signal across conditions through a GLM analyses, the findings we report using SPADE refer to changes in covariation between brain areas. Consequently, although the results are different, the analysis we present here and that in [15] are not exclusive but rather complementary, with both having their advantages and limitations: while for memory reasons we cannot yet compute the SPADE solutions at the full brain, with a voxel-wise resolution, the computational load of performing the SPADE decomposition itself is much lower than that of a typical GLM. As an example of the computational advantage of the SPADE model we observe that after computing the ROI time-series for the >1000 subjects included in this study the computation of the full SPADE spatial maps took less than two minutes using a standard laptop.

Although working full brain voxel-wise would be optimal, it is to date computationally intractable and our choice of working in a parcellated full brain [13] could be replaced by working in a resting state network space or even voxel-wise when restricted to an anatomical ROI that might be relevant for a particular investigation. Here we considered a parcellated full brain for ease of results interpretation and visual validation. Further, since the simultaneous diagonalization introduced by [5] can be used for any pair of symmetric positive definite matrices, and given the growing neuroimaging interest in working with partial correlation matrices, we considered the extension of the SPADE algorithm to partial correlation matrices and observe that it did not improve in discrimination with respect to the original formulation. Nevertheless we believe that such an approach warrant more thorough investigation in cases where the initial spatial dimensionality reduction results in more highly correlated data, for example when using an ICA based strategy [3] or a small spatially continuous ROI; note here

that in these two scenarios, although the methodology is identical, the interpretation changes drastically since the filters will relate to resting state networks directly or voxels being involved, respectively. An extra flexibility of the model presented is the possibility of extension to multi-class versions by using combinations of binary problems or with approximations to multiple simultaneous diagonalizations [27]. Another research direction we are investigating for this model includes the study of subject-wise deviation from group functional connectivity as a strategy to address normative modelling [28] on functional connectomes.

To conclude, the SPADE model is able to capture and summarize the brain nodes involved in network covariation changes induced by fMRI experimental manipulation. As introduced here, the SPADE model can be used to address questions of potential interest to diverse scientific and industrial communities.

Acknowledgements

The research leading to these results has received funding from the developing Human Connectome Project through a Synergy Grant by the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013), ERC Grant Agreement no. 319456. We further gratefully acknowledge support from the Netherlands Organization for Scientific Research (NWO) through VIDI grant to CFB (864.12.003) and we also gratefully acknowledge funding from the Wellcome Trust UK Strategic Award (098369/Z/12/Z).

Data were provided [in part] by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

SUPLEMENTARY MATERIAL (SM)

Discriminative simultaneous covariance diagonalization.

Consider two sets of time-series, X and Y, measured at the same spatial locations during two different states of the system being measured, and define the spatial covariance matrices of each of these series as C_X and C_Y. Usually, the unsupervised simultaneous diagonalization of both covariances is achieved through a generalized eigenvalue decomposition, involving a whitening transformation of one of the covariance structures (e.g C_X), combined with a further rotation to also diagonalize the other covariance (C_Y). A different approach was introduced in [5] where the whitening transformation is performed with respect to the sum of both covariances $C := C_X + C_Y$; such whitening can be summarized by two matrices, a rotation W and a scaling diagonal matrix P, such that $I=P^TW^TCWP$. Individually applying these transformations to C_X and C_Y, one obtains two (non-diagonal) matrices $K_X= P^TW^TC_XWP$ and $K_Y= P^TW^TC_YWP$. Performing now the eigenvalue decomposition of (for example) K_Y we obtain matrices Z_Y and diagonal D_Y such that $Z_YK_YZ_Y^T=D_Y$. Defining $V:=WPZ_Y$, we have that $V^TC_XV=I-D_Y$ and $V^TC_YV=D_Y$. This means that V defines a basis for the data that diagonalizes simultaneously C_X and C_Y and, and by construction, the sum of the variances of the projections of data X and Y into each of the new basis vectors adds to one; consequently, for basis directions where the projection of X has large variance (i.e. close to one), the projection of Y must have low variance (i.e. close to zero), and vice versa. The algorithmic imposition to maximize variance for one class and minimize it for the other one, makes of this process a supervised multivariate dimensionality reduction model, and the new learned basis is formed of spatial filters that provide optimal linear discrimination in terms of variance of the projected data (Blankertz et al., 2011). Note that for implementation purposes it all reduces to solving a generalized eigenvalue problem which can be easily accomplished using most software platforms, for example using the matlab command eig (C_x, C) or the python command scipy.linalg.eigh(C_x,C,eigvals_only=False)). In the recent literature, the process described in [5] is commonly used as a feature extractor for EEG based brain computer interfaces [6,7] and the features extracted for discrimination are usually the logarithmic of the variance of the projected data; by construction of the algorithm the most discriminative dimensions are the associated with the extreme eigenvectors since towards the center of the eigen-spectrum the variances for both conditions tend to be close to ½. In practice, for discrimination purposes the dimension is usually reduced to a few pairs of extreme eigenvectors and a low dimensional linear classifier suffices [7].

Interpretation of the linear filters

Since the simultaneous diagonalization assumes no implicit noise model, the spatial filters learned, i.e. the columns of V (V_k) cannot be directly interpreted and the interpretable spatial maps (A_k) are obtained from $A=C_{X|Y}VC_S^{-1}$ [10,11], where X|Y denotes the spatially concatenated X and Y data, and S the data projected to the new basis V, i.e. $S= V^T(X|Y)$. For completeness we note here that X|Y = AS + e, where e denotes the residual of the approximation X|Y = AS and for a more detailed description we refer the reader to [10,11].

Model order estimation

Although for classification purposes the CSP model order selection is not critical due to the latter integration of a classifier, for interpretation purposes we are required to select which filters provide significant connectivity covariation information. To that end we address for the first time the model order estimation on this model by permutation testing to find a robust estimation of which filters explain more variance for a modality than for the other one with respect to a null distribution obtained by permutation. To that end a p-value is computed for each basis vector (column of V, V_k) by comparing the absolute difference between the variances of the data projections, Abs(Var(V'_k*X)- Var(V'_k*Y)), with a distribution of values obtained when the simultaneous diagonalization new basis (V) is computed using randomized groups of covariances. In the dataset set used in the main text the randomization is achieved by permuting half of the subject's covariance matrices across different fMRI modalities.

References

1. Friston, K., Holmes, A. & Worsley, K. Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* (1994).

2. Aguirre, G. K. & Detre, J. A. Experimental Design and the Relative Sensitivity of BOLD and Perfusion fMRI. **500,** 488–500 (2002).

3. Beckmann, C. F. & Smith, S. M. Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Trans Med Imaging* **23(2),** 137–152 (2004).

4. Beckmann, C. F. & Smith, S. M. Tensorial extensions of independent component analysis for multisubject FMRI analysis. **25,** 294–311 (2005).

5. Fukunaga, K. *Introduction to Statistical Pattern Recognition*. *Pattern Recognition* **22,** (1990).

6. Ramoser, H., Müller-Gerking, J. & Pfurtscheller, G. Optimal spatial filtering of single trial EEG during imagined hand movement. *IEEE Trans. Rehabil. Eng.* **8,** 441–446 (2000).

7. Blankertz, B., Lemm, S., Treder, M., Haufe, S. & Müller, K. R. Single-trial analysis and classification of ERP components - A tutorial. *Neuroimage* **56,** 814–825 (2011).

8. Van Essen, D. C. *et al.* The WU-Minn Human Connectome Project: An overview. *Neuroimage* **80,** 62–79 (2013).

9. Elam, J. S. & Van Essen, P. D. Human Connectome Project. in *Encyclopedia of Computational Neuroscience* 1–4 (2013).

10. Haufe, S. *et al.* On the interpretation of weight vectors of linear models in multivariate neuroimaging. *Neuroimage* **87,** 96–110 (2014).

11. Parra, L. C., Spence, C. D., Gerson, A. D. & Sajda, P. Recipes for the linear analysis of EEG. *Neuroimage* **28,** 326–341 (2005).

12. Van Essen, D. C. *et al.* The Human Connectome Project: A data acquisition perspective. *NeuroImage* **62,** 2222–2231 (2012).

13. van Oort, E. S. B. *et al.* Functional parcellation using time courses of instantaneous connectivity. *Neuroimage* **170,** 31–40 (2018).

14. Bishop, C. M. *Pattern Recognition and Machine Learning*. *Pattern Recognition* **4,** (2006).

15. Barch, D. M. *et al.* NeuroImage Function in the human connectome : Task-fMRI and individual differences in behavior. *Neuroimage* **80,** 169–189 (2013).

16. Cole, M. W., Ito, T., Bassett, D. S. & Schultz, D. H. Activity flow over resting-state networks shapes cognitive task activations. *Nat. Neurosci.* **19,** 1718–1726 (2016).

17. Sahyoun, C., Floyer-Lea, A., Johansen-Berg, H. & Matthews, P. . Towards an understanding of gait control: brain activation during the anticipation, preparation and execution of foot movements. *Neuroimage* **21,** 568–575 (2004).

18. Smith, S. M. *et al.* Correspondence of the brain's functional architecture during activation and rest. *Proc. Natl. Acad. Sci.* **106,** 13040–13045 (2009).

19. Owen, A. M., McMillan, K. M., Laird, A. R. & Bullmore, E. N-back working memory paradigm: A meta-analysis of normative functional neuroimaging studies. *Hum. Brain Mapp.* **25,** 46–59 (2005).

20. Soto, D., Rotshtein, P. & Kanai, R. Parietal structure and function explain human variation in working memory biases of visual attention. *Neuroimage* **89,** 289–296 (2014).

21. Takahashi, E., Ohki, K. & Kim, D. S. Dissociation and convergence of the dorsal and ventral visual working memory streams in the human prefrontal cortex. *Neuroimage* **65,** 488–498 (2013).

22. Haak, K. V. & Beckmann, C. F. Objective analysis of the topological organization of the human cortical visual connectome suggests three visual pathways. *Cortex* **98,** 73–83 (2018).

23. Fernandez-Ruiz, A. & Oliva, A. Distributed Representation of 'What' and 'Where' Information in the Parahippocampal Region. *J. Neurosci.* **36,** 8286–8288 (2016).

24. Müller, N. G. & Knight, R. T. The functional neuroanatomy of working memory: Contributions of human brain lesion studies. *Neuroscience* **139,** 51–58 (2006).

25. Baddeley, A. Working memory: Looking back and looking forward. *Nat. Rev. Neurosci.* **4,** 829–839 (2003).

26. Baddeley, A. The episodic buffer: A new component of working memory? *Trends Cogn. Sci.* **4,** 417–423 (2000).

27. Grosse-wentrup, M., Member, S. & Buss, M. Multi-class Common Spatial Patterns and. **55,** 1–10 (2008).

28. Marquand, A. F., Rezek, I., Buitelaar, J. & Beckmann, C. F. Archival Report Understanding Heterogeneity in Clinical Cohorts Using Normative Models : Beyond Case-Control Studies. *Biol. Psychiatry* **80,** 552–561 (2016).