

fMRI Single Trial Discovery of Spatio-Temporal Brain Activity Patterns

Michele Allegra,^{1*} Shima Seyed-Allaei,^{2,3,4} Fabrizio Pizzagalli,^{1,5}
Fahimeh Baftizadeh,⁶ Marta Maieron,⁷ Carlo Reverberi,^{2,3}
Alessandro Laio,¹ and Daniele Amati¹

¹SISSA-International School for Advanced Studies, Via Bonomea, Trieste 265, Italy

²Psychology Department, University of Milan Bicocca, Milan, Italy

³Milan Center for Neuroscience, Milan, Italy

⁴Department of Electrical and Computer Engineering, University of Tehran, Tehran, Iran

⁵Imaging Genetics Center, Mark and Mary Stevens Neuroimaging & Informatics Institute, Keck School of Medicine, the University of Southern California, Marina del Rey, California

⁶Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts

⁷Medical Physics Department, AOUD S. Maria della Misericordia Hospital, Udine, Italy



Abstract: There is growing interest in the description of short-lived patterns in the spatiotemporal cortical activity monitored via neuroimaging. Most traditional analysis methods, designed to estimate relatively long-term brain dynamics, are not always appropriate to capture these patterns. Here we introduce a novel data-driven approach for detecting short-lived fMRI brain activity patterns. Exploiting Density Peak Clustering (Rodriguez and Laio [2014]), our approach reveals well localized clusters by identifying and grouping together voxels whose time-series are similar, irrespective of their brain location, even when very short time windows (~10 volumes) are used. The method, which we call Coherence Density Peak Clustering (CDPC), is first tested on simulated data and compared with a standard unsupervised approach for fMRI analysis, independent component analysis (ICA). CDPC identifies activated voxels with essentially no false-positives and proves more reliable than ICA, which is troubled by a number of false positives comparable to that of true positives. The reliability of the method is demonstrated on real fMRI data from a simple motor task, containing brief iterations of the same movement. The clusters identified are found in regions expected to be involved in the task, and repeat synchronously with the paradigm. The methodology proposed is especially suitable for the

Additional Supporting Information may be found in the online version of this article.

Conflict of interest: The authors declare no competing financial interests.

Contract grant sponsor: The PRIN grant; Contract grant number: 2010RP5RNM_001 from the Italian Ministry of University (CR and SSA).

*Correspondence to: Michele Allegra; SISSA-International School for Advanced Studies, Via Bonomea, Trieste 265, Italy. E-mail: mallegra@sissa.it

Received for publication 5 April 2016; Revised 30 September 2016; Accepted 1 November 2016.

DOI: 10.1002/hbm.23463

Published online 23 November 2016 in Wiley Online Library (wileyonlinelibrary.com).

study of short-time brain dynamics and single trial experiments, where the event or task of interest cannot be repeated for the same subject, as happens, for instance, in problem-solving, learning and decision-making. A GUI implementation of our method is available for download at <https://github.com/micheleallegra/CDPC>. *Hum Brain Mapp* 38:1421–1437, 2017. © 2016 Wiley Periodicals, Inc.

Key words: unsupervised fMRI analysis; time-dependent connectivity; short-lived brain activity patterns

INTRODUCTION

Several cognitive processes take place over a rapid time scale, are not necessarily synchronous with external events, cannot be repeated in the same subject, and may occur at different times for different subjects. Examples are the change of strategy during the solution of a problem, decision-making in a novel setting, and learning [Bassett et al., 2011; Braun et al., 2015; Schuck et al., 2015]. These processes offer a severe challenge to fMRI inquiry [Baert et al., 2000; Dale, 1999; Henson, 2007; Jezzard et al., 2001], because the transient nature of the features defies established analysis approaches.

Available data analysis tools for fMRI can be classified as supervised or unsupervised. Supervised approaches are based on specific assumptions about the measurable response to external events. For instance, in statistical parametric mapping [Friston et al., 1995] one assumes that the signal associated to an experiment is given by a design matrix representing experimental conditions, convoluted with an appropriate response function. Robust though they may be, supervised approaches are of little use when there is no obvious design matrix. In these conditions, unsupervised approaches seem more appropriate, since they probe brain activity without assuming a relation to predefined events or other predictors. The most commonly used unsupervised approaches are Principal Component Analysis (PCA) [Hansen et al., 1999; Lai and Fang, 1999], Independent Component Analysis (ICA) [Beckmann and Smith, 2004; Calhoun et al., 2001; Daubechies et al., 2009; McKeown et al., 1998], clustering analysis [Craddock et al., 2012; Dimitridou et al., 2004; Yee and Gao, 2002], and paradigm-free mapping [Caballero et al., 2011, 2013]. The latter approach provided a proof-of-principle demonstration of the viability of fMRI analysis on single trials.

The limitations of current unsupervised methods call for the development of new approaches that are sensitive and robust enough to identify an interesting signal in a single subject, for very short time windows, and independently of correlation with external events. In this work, we develop a novel unsupervised fMRI analysis method aimed at meeting these challenges. Our method assumes that voxels participating to the same brain function should behave *coherently*, that is, with similar time development. Therefore, we propose a procedure to identify similarly behaving voxels within a (narrow) time window and then group

them into clusters through the recently introduced Density Peak Clustering (DPC) algorithm [Rodriguez and Laio, 2014]. Our method differs from others in two aspects. First, it automatically embeds signal filters that allow removing artifacts and optimizing the detection of genuine coherent signals. Second, at variance with other clustering-based analysis methods, it identifies the correct number of clusters with no a priori assumptions on their shape in the signal space, location in the brain, and size. As a result, we achieve an automatic and unsupervised procedure that allows identifying and clustering together voxels that behave coherently in a time window as short as approximately 10 volumes, irrespective of their brain location or their relation with task events, at the same time discarding many noise-related artifacts. We call this procedure Coherence Density Peak Clustering (CDPC).

This manuscript is dedicated to introducing and validating the CDPC method for fMRI analysis, providing the ground for future work. Upon introducing CDPC, we first apply it to artificially generated data, simulating a short fMRI recording for a single subject and verifying that it is able to recognize the genuine signal up to a very low signal-to-noise ratio. We then perform a second test with real data, applying CDPC to fMRI images from an experiment involving a simple motor task, showing that it can efficiently detect meaningful signal in the expected activation regions. In both cases, our method is compared with ICA, one of the state-of-the-art approaches for unsupervised analysis of fMRI data. We decided to perform a detailed comparison with ICA because this approach achieves source separation and signal reconstruction on completely different principles and is currently the most widespread technique to investigate spatiotemporal patterns in fMRI. Results of both the simulation and the experiment will show that CDPC is more accurate than ICA, and less prone to detect spurious signal.

MATERIALS AND METHODS

CDPC for fMRI Images

As mentioned in the introduction, the goal of our method (CDPC) is to detect groups of voxels that show coherent patterns of activation. Before describing the method in detail, we illustrate its overall logic with a flow chart (Fig. 1).




Figure 1.

Cartoon representation of the CDPC approach for fMRI analysis. (A) For each pair of voxels, a dissimilarity distance is obtained. (B) According to the number of similar points, voxels are represented on a two-dimensional graph (dubbed the “decision graph”). The outliers in the decision graph (dots of

different colors) are the clusters centers we are looking after. (C) Voxels are assigned to clusters represented in the same colors as their centers in panel B. [Color figure can be viewed at wileyonlinelibrary.com]

As a starting point of our procedure, we define a metric d_{ij} that measures the similarity between the BOLD time series of any pair of voxels i, j in a given time window. We shall call d_{ij} a dissimilarity metric, because d_{ij} is low when the time series of the two voxels is similar (top two traces of Fig. 1A) and high when it is dissimilar (bottom trace in Fig. 1A). After estimating d_{ij} , and applying suitable noise filters, we employ the DPC algorithm [Rodriguez and Laio, 2014] to find the voxels involved in regions behaving coherently. The method identifies the cluster centers by means of a simple graphical representation called decision graph (Fig. 1B). The latter is based on plotting a “density” ρ_i for each voxel (essentially, the number of voxels with a similar BOLD signal) versus a quantity δ_i that measures the distance from the nearest voxel with higher density. As discussed below, δ_i is expected to be high for points that are cluster centers. In the decision graph, cluster centers stand out as isolated points with a large value of ρ_i and δ_i (highlighted in color in Fig. 1B). Other voxels are assigned to each cluster center by a simple iterative procedure. The resulting clusters are not necessarily localized in a single area of the brain but often split in two or more spatially separated regions (Fig. 1C).

CDPC offers some advantages in comparison to other commonly used data-driven fMRI analysis methods. The procedure is fully automatized and unsupervised, and is stable with respect to changes in the free parameters. Second, it yields a visual representation showing the number and ordering of the clusters. Moreover, it includes effective signal filters *before* the clustering step (in the computation of d_{ij} and ρ_i), which reduces the chance of finding artifacts and makes the procedure tailored to conditions where the interesting signal is weak, due to either the paucity of data available (short windows, single subject) or the strong noise.

The (C++) implementation of our method, together with the data used in this work, can be found at <https://>

github.com/micheleallegra/CDPC. Moreover, a GUI implementation of our method as an SPM toolbox is available for download on the same page.

More in detail the CDPC algorithm works as follows. A sequence of fMRI images is represented by the voxel intensities $v_i(t)$ where i labels the voxel ($i = 1, \dots, N$, N being of the order of 100,000 for the whole brain) and t represents the discrete time of successive volumes (scans) or, equivalently, the volume number. We define a metric d_{ij} that measures the dissimilarity between the intensity time-series $v_i(t), v_j(t)$ of any pair of voxels i, j in a given time window T . Our choice of d_{ij} is made in such a way as to properly account for inhomogeneities in the power spectrum of the intensities as well as local differences in the vascularization of brain areas. We first perform the discrete Fourier transform, over the time window including a number T of volumes, of the i th voxel intensity as

$$v_i(\omega) = \sum_{t=0}^{T-1} v_i(t) e^{-i2\pi\omega t} \quad (1)$$

Since $v_i(t)$ is real, the independent components of $v_i(\omega)$ are $T/2$, namely its values for $\omega = 0, \dots, T/2 - 1$. The $\omega = 0$ component of the signal is assumed to be irrelevant, and excluded from any further analysis. From $v_i(\omega)$ we estimate the average power spectrum

$$P(\omega) = \frac{1}{N} \sum_{i=1}^N v_i^*(\omega) v_i(\omega) \quad (2)$$

where the star indicates the complex conjugate. On average, the power is higher for small wave numbers. The dominance of slow components in the signal (i.e., the components in the frequency range 0.0–0.015 Hz) in fMRI time series is well documented in the literature [Bullmore et al., 1996; Genovese et al., 1997; Mattay et al., 1996], and it is normally attributed to cardiac and respiratory aliasing, slow periodic hemodynamic changes [Turner et al., 1998],

heating up of the gradient coils, tissue or fluid shifting, and establishment of thermal equilibrium [Smith et al., 1999]. Usually, the slow components are simply discarded by filtering the signal by a high-pass filter (with a low-frequency cut-off usually set at 0.008 Hz). However, by computing $P(\omega)$ for the seven different subjects in our motor experiment we have verified that the power decays pretty slowly with ω , roughly as $1/\omega$ (Supporting Information Fig. S1). In these conditions, introducing a specific high-pass frequency cutoff can be arbitrary. Instead, we prefer to use a non-linear filter that retains all wave numbers in the analysis, dividing each Fourier component by the average power at the same frequency. In this way we enhance high-frequency components with respect to low-frequency ones, without arbitrarily eliminating the latter. The signal is transformed as follows:

$$\hat{v}_i(\omega) = \frac{v_i(\omega)}{\sqrt{P(\omega)}} \quad (3)$$

We define the dissimilarity d_{ij} between voxel i and j as

$$d_{ij} = \sqrt{\left\| \frac{\hat{v}_i(\omega)}{\hat{v}_i} - \frac{\hat{v}_j(\omega)}{\hat{v}_j} \right\|^2} \quad (4)$$

where the normalization $\hat{v}_i = \max_{\omega} \hat{v}_i(\omega)$ is introduced in order to include in the definition of d_{ij} only genuine variations in the time development, and not differences in signal amplitude that may be trivially due to, for example, differences in the vascularization [Kalcher et al., 2013; Logothetis, 2008]. We remark that d_{ij} is a measure of the dissimilarity between the two-time series, and can thus be considered as a distance in the high-dimensional “signal space” of voxel time-series, with no relation to the spatial distance between the two voxels in the brain.

The next step is grouping voxels in sets that behave in a similar manner, that is, that have a low dissimilarity among them. These sets of voxels are called clusters and they are identified through the DPC method, recently introduced by Rodriguez and Laio [2014] (in the remainder of this section, we will refer to this article as RL2014). The DPC procedure requires first computing the density as $\rho_i = \sum_j \theta(d_c - d_{ij})$ where $\theta(x)$ is the Heaviside step function ($\theta(x) = 1$ if $x \geq 0$ and $\theta(x) = 0$ otherwise) and d_c is a dissimilarity cutoff. This quantity essentially counts the number of voxels that are similar to i , and is called “density” as it represents a simple measure of the density of points surrounding i in signal space. However, we found that this simple definition of density (originally devised for data with low noise) is not optimal for a clustering procedure applied to data affected by strong noise like in the present case. In fact, the density as defined above is highly sensitive to spurious coherence between voxels that can statistically arise even in situations where no genuine coherent signal is actually present. Indeed, upon analyzing a series of fMRI scans on imaging phantoms we obtained values for the density that were

significantly different from zero, despite the clear absence of any functional source of coherence. Therefore, a more advanced definition of density is needed. Analyzing the same phantom data, we found that spurious coherence arising from noise is not spatially localized, that is, it does not involve neighboring voxels. Instead, a truly coherent signal is unlikely to involve spatially isolated voxels: it should rather involve connected regions, so that the voxels involved should behave coherently with at least some of the neighboring voxels. In other words, if all voxels similar to i (i.e., with $d_{ij} < d_c$) are spatially remote from voxel i , it is likely that this similarity is spurious (the result of coherent noise fluctuations): in contrast, if among the voxels similar to i there are several spatially close voxels, we are probably observing a true coherent signal. Thus, if we define the “number of coherent neighbors” as $n_i = \sum_{j \in NN_i} \theta(d_c - d_{ij})$ where the set NN_i includes the voxels inside a small sphere centered on the voxel i , an effective way to distinguish genuine coherence from artifactual coherence is implementing an additional noise filter, eliminating from the computation of the density all voxels with $n_i < n_0$ where n_0 represent a “coherent neighbors” cutoff. Notice that such a cutoff does not eliminate long-range coherence: it just requires that long-range coherence happens between distant, possibly small, connected regions, not distant isolated voxels.

We are thus ready to define the density ρ_i in terms of two parameters, the similarity cutoff d_c and the noise cutoff n_0 . Their optimal values may depend on the total number of voxels and the voxel sizes, so they may need to be adjusted for different experiments. After concluding the presentation of our CDPC method, we will discuss the choice of these two parameters for both simulated and fMRI data in our test and the robustness of the results with respect to their value. In terms of d_c and n_0 , we define the corrected density as

$$\hat{\rho}_i = \sum_j \theta(d_c - d_{ij}) \theta(n_0 - n_i) \theta(n_0 - n_j) \quad (5)$$

We finally normalize the density dividing by its maximum value, $C = \max_i \hat{\rho}_i = \max_i \sum_j \theta(d_c - d_{ij}) \theta(n_0 - n_i) \theta(n_0 - n_j)$, obtaining our final expression for the density:

$$\rho_i = 1/C \sum_j \theta(d_c - d_{ij}) \theta(n_0 - n_i) \theta(n_0 - n_j) \quad (6)$$

In this way, we always have $0 \leq \rho_i \leq 1$. This global normalization (which does not affect the subsequent clustering results) is only introduced to facilitate comparison across different time windows, subjects, and experiments where the density range may slightly vary.

Once defined the density, the number of clusters will be equal to the number of local maxima in the density distribution. Following RL2014, we compute $\delta_i = \min_{j: \rho_j > \rho_i} d_{ij}$, which is the dissimilarity between a voxel and the most similar voxel having a higher density. Cluster centers stand out as isolated points with a large value ρ_i and δ_i . In order to make the approach fully automatic, we do not

visually inspect the decision graph as suggested in RL2014, but we rank the voxels according to their value of δ_i and we consider as putative cluster centers the first $k_{\max} = 10$ voxels (k_{\max} should always be larger than the number of populated clusters obtained at the end of the procedure—see below; if this does not happen, one should repeat the procedure with a larger k_{\max}). After the cluster centers have been chosen in this manner, all voxels with $\rho_i > 0$ are assigned to a cluster. This is done by the following recursive procedure: each voxel is assigned to the same cluster of the most similar voxel having a higher density; if the latter voxel is not yet assigned, one looks for the voxel most similar to it having a higher density, and so forth until either an already assigned voxel or one of the cluster centers is reached. Voxels with $\rho_i = 0$ are not assigned to any cluster, and they are discarded as noise. At the end of this procedure, only some clusters will contain a sizable (>50) number of voxels. In our data, we always find less than eight sizable clusters, which justifies a posteriori our initial choice of k_{\max} . In general, one should select k_{\max} based on some rough expectation on the maximum number of clusters, and verify that the number of sizable clusters is always self-consistently lower than k_{\max} (otherwise, k_{\max} should be increased). For each cluster, we compute the average density of the voxels assigned to it, and we rank the clusters based on this number. Thus, cluster #1 will be the cluster with highest average density.

Parameters Setting

We discuss here the criteria for setting the parameters d_c and n_0 . In the “Results” section we will show the robustness of the results with respect to their variation.

We fix the dissimilarity cutoff d_c to a value for which the average number of voxels within d_c is $M_c = 200$, corresponding to 0.15% of the total number of voxels (see RL2014). As discussed in RL2014, the DPC approach is sensitive only to the relative magnitude of the density, not to its absolute value. Indeed, the relation $\rho_i > \rho_j$ is likely to remain valid if ρ is changed, especially when the difference in ρ is large. This a priori expectation will be checked with our data.

For what regards the noise cut off n_0 , when using phantom data recorded with the same resolution of real data, only a negligible fraction of voxels have more than five spuriously coherent voxels within a sphere of radius 6 mm. Therefore, we decided to implement the additional noise filter by effectively removing from the analysis all voxels with $n_0 < 5$ and defining the set NN_i of neighboring voxels as all voxels within a sphere of 6 mm around the given voxel. We remark that the optimal value of this parameter can depend significantly on the spatial resolution of the image.

Simulation

As a first test of our new analysis procedure, CDPC, we applied it to simulated fMRI time series. The simulation

was devised so that it mirrored the design of our test experiment (see below). For generating the simulated data we used the neuRosim package [Welvaert et al., 2011], which allows to specify the stimuli, the activated regions, and the properties of noise. The stimuli are modeled as step functions, convoluted with a model hemodynamic response function.

In our simulation, we produce 12 images, analogous to the number of volumes we considered in the analysis of the real data (see below). The simulated TR is 2.5, which gives a total simulation time of 30 s. The number of voxels ($128 \times 34 \times 34$) and voxel dimensions are the same as images from the motor experiment. The fMRI intensity time series at all voxels is obtained by adding the signals on a constant “base” image, that we take to be the mean image of one subject from our motor experiment (s_1), masked with the same brain mask.

We consider three activated regions. Two regions are generated as two spheres with 10 voxel radius and centers at native space coordinates 49,42,24 and 108,89,24. These regions respond to an identical stimulus, a step function that is 1 in the first half of the scans (scans 1–6) and 0 in the second half (scans 7–12). A further, third, region is generated as a sphere with 10 voxel radius and center at coordinates 86,91,18). This third region responds to a different stimulus, modeled as a step function that is 0 in the first half of the scans and 1 in the second half. The signal in the spheres is obtained by convolving the relevant stimulus with a standard hemodynamic response function (double-gamma function). The amplitude of the signal is maximal at the centers of the spheres and slowly fades (with a Gaussian decay coefficient, $0.5 + 0.5e^{-\frac{2}{l^2}}$ with $l^2=10$) as one goes far from the centers. All voxels (both in and out of the regions) are affected by random noise. The magnitude of the noise is fixed by the magnitude of the signal, and by the signal/noise ratio. Given a certain value of signal-to-noise (S/N), the average magnitude of the signal in the activated regions, \bar{s} , is computed, and then the standard deviation of noise over all voxels is fixed as $\sigma = \bar{s}/(S/N)$. We considered four different values of signal-to-noise (S/N) ratio: 20, 5, 3, and 2. The statistical properties of the noise can be adjusted by varying the weights of six standard noise types that model noise from different sources including white noise, physiological noise and spatially correlated noise. The weights of the standard noise types available in neuRosim were fixed as follows: 0.1,0.1,0,0.2,0.2,0.4. These parameters were chosen so as to match as closely as possible the properties of noise found in real images from the motor experiment.

Motor Experiment

Participants

Seven healthy subjects participated in the experiment. They had a mean age of 33.5 years (range 24–46). Three

were males and four females. All participants gave written informed consent. They were right-handed and had normal or corrected to normal vision, no neurological or psychiatric history, and no structural brain abnormalities. The study was approved by the local ethics committee.

Image acquisition

MRI data acquisition was conducted on a whole-body 3 Tesla Philips Achieva (Best, Netherlands) MRI scanner equipped with a SENSE-Head-8 channel coil. Functional runs were acquired using a T2* BOLD-sensitive gradient-recalled EPI sequence, imaging parameters were as follows: TR = 2,500 ms; TE = 35 ms; 90° flip angle; SENSE reduction factor in phase encoding direction = 2; FOV = 23 × 23 cm; 128 × 128 image matrix, yielding an in-plane voxel size of 1.8 × 1.8 mm; 34 axial slices, slice thickness = 3 mm with no gap. Head motion was reduced by a foam custom-built head cushion around the subject's head. The first four volumes in each scan series, collected before equilibrium magnetization was reached, were discarded to allow the MR scanner to reach a steady state. Anatomical T1-weighted images MPRAGE were also acquired (190 sagittal slices; TR = 8.1007; TE = 3.707 ms; flip angle 8°; FOV = 24 cm; voxel size 1 × 1 × 1 mm) to obtain structural three-dimensional (3D) volume.

MATERIALS AND PROCEDURE

Subjects were visually instructed to perform a right or left-hand clenching movement in a blocked design. Each block was composed by 15 s (i.e., 6 volumes) of left/right clenching followed by 15 s of rest. Experimental blocks were repeated eight times. Visual instructions remained visible for the full duration of each task phase, clenching or baseline. The eight blocks were followed by a final baseline acquisition lasting 15 s. The total duration of the fMRI acquisition is, therefore, 102 volumes, corresponding to 255 s. Instructions about the beginning, the end, and the side of the movement were visually cued during the fMRI acquisition. Participants were instructed to relax and remain still, to keep their arms aligned with the sides of the body, and to breathe normally. Stimuli and instructions were presented through a VisuaStim Goggles system (NordicNeuroLab, Bergen, Norway) equipped with the Presentation software (Version 9.9, Neurobehavioral Systems Inc., CA).

Standard Data Analysis

Image preprocessing was performed using SPM12 (Wellcome Trust Centre for Neuroimaging, London) running under Matlab 7.4 (R2007a) (Mathworks, Sherborn, MA). The performed preprocessing steps were: slice timing correction, realignment to correct for motion, co-registration of anatomical images with functional images, and tissue

segmentation to build a brain mask. The same preprocessing steps were performed to all functional images previous to application of all analysis methods (CDPC, ICA). Spatial normalization and spatial smoothing were performed only after ICA and CDPC were applied. In particular, spatial smoothing would interfere with the CDPC, as it enhances correlations between neighboring voxels, creating a spurious source of coherence that impairs the discrimination of genuine signal from noise. Spatial normalization was avoided to prevent the introduction of interpolation biases before results were computed [Coulon et al., 2011; Pizzagalli et al. 2013; Thirion et al., 2006]. Spatial normalization and smoothing were only used to facilitate the comparison of results across subjects. The result maps from ICA and CDPC analyses were spatially normalized to the standard MNI template and spatially smoothed by using a Gaussian kernel with 9 mm FWHM. For visualization purposes, we used the MRICron software (www.mricron.com).

Independent component analysis (ICA) was carried out using the GIFT toolbox (MIA lab, Mind Research Network, Albuquerque, New Mexico). For the motor experiment data, ICA analysis was performed independently on each subject. The number of components is estimated using Minimum Description Length criteria. Principal component analysis is used to reduce data dimensionality before applying ICA, which is performed by means of the Infomax algorithm. As Infomax is an optimization algorithm, it must be iterated several times to ensure consistency of the optimization results. We used the ICASSO toolbox [Himberg and Hyvärinen, 2003] implemented in GIFT to determine the consistent clusters of the algorithm. The optimization was performed 20 times (in agreement with common practice, for example, in Himberg and Hyvärinen, [2003] the authors performed 15 iterations) and the requirement for consistency was that clusters appear in at least 80% of the optimization iterations (the 80% threshold is recommended by GIFT toolbox).

RESULTS

Test of the Method with Simulated Data

As a first test the CDPC method was applied to simulated data, where the “ground truth” is known. In particular, we are interested in exploring up to which signal/noise ratio CDPC analysis can recognize and discriminate coherent activity in a single subject. Furthermore, we compared CDPC performance with an established method (ICA). It is important to remark the short-lived character of the patterns we consider in this example.

We generated simulated fMRI data covering a time-window of 12 volumes, corresponding to 30 s with a TR = 2.5 s. On a brain mask taken from a real subject we define three activated regions. The first two regions respond to a common stimulus, and thus they should be clustered together. The third region responds to a different stimulus and thus it should appear as a second cluster.

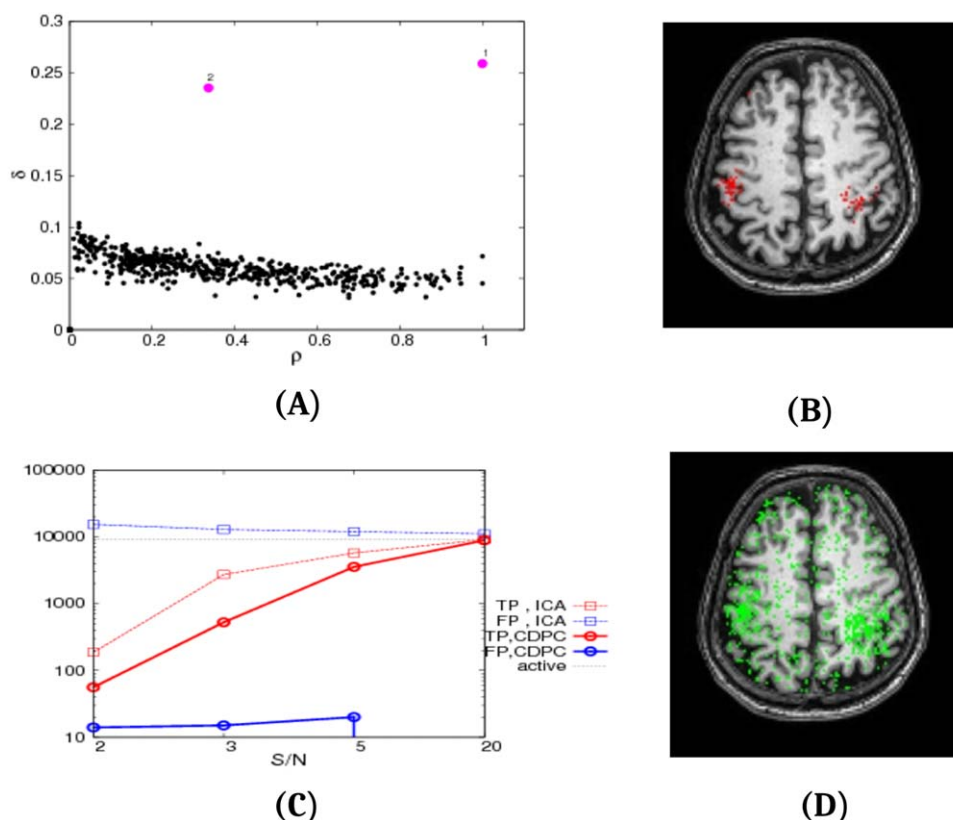


Figure 2.

(A) Decision graph for $S/N = 3$. The purple dots are the density peaks of clusters 1 and 2. (B) Rendering of the voxels assigned to a cluster by CDPC for $S/N = 3$. The true positives concentrate in the centers of the spheres, because the S/N of the simulated signal is fading from the center to the periphery of the spheres. (C) Number of voxels belonging to the activated regions and correctly assigned to a cluster by CDPC (true positives, TP, red) and number of voxels incorrectly assigned to a cluster (false positives, FP, blue) as a function of the S/N ratio.

For comparison, we also show the corresponding number of TP and FP for ICA, where we consider as significant a voxel whose z score is higher than 2.3. For reference, a horizontal line is drawn representing the number of voxels belonging to the activated regions (“active”) (D) Rendering of the voxels assigned to a cluster by ICA at $S/N = 3$. In addition to true positives, one can clearly notice the very large number of false positive voxels generated by ICA—contrary to what happens with CDPC. [Color figure can be viewed at wileyonlinelibrary.com]

All voxels (both inside and outside these regions) are affected by random noise. We considered four different values of signal-to-noise (S/N) ratio: 20, 5, 3, and 2. For each value of S/N , we applied CDPC. For $S/N = 20, 5, 3$ the CDPC analysis correctly finds two clusters. The three activated spheres are correctly identified as two separated clusters, with the first two spheres in the first cluster and the third sphere in the second cluster. In Figure 2A, we show the decision graph in which each voxel is represented with its value of δ (dissimilarity from a point with higher ρ) as function of ρ (i.e., number of voxels coherent with the target voxel) for $S/N = 3$. The two clusters are clearly identified as two well isolated points (depicted in purple in Fig. 2A). In Figure 2B, we show a rendering of the voxels found by the CDPC (voxels assigned to a

cluster) for $S/N = 3$. The activated spheres are clearly visible. The decision graphs for $S/N = 20, S/N = 5$ are very similar, and so are the renderings (except that the spheres are more clearly identified at $S/N = 20, S/N = 5$). For $S/N = 2$ the first two spheres are still identified as the first cluster, but the third sphere is recognized only as the third cluster, since a “spurious” cluster appears as the second cluster. For each value of S/N ratio, the performance of the method was assessed by counting the number of true positives TP (activated voxels that are assigned to the clusters) and false positives FP (nonactivated voxels that are mistakenly assigned). The number of TP and FP as a function of the S/N ratio is presented in Figure 2C. For any value of S/N the number of TP is always higher than the FP. As the S/N ratio is decreased, the number of TP

significantly decreases, but the number of false positives remains very low in comparison, with the ratio between FP and TP reaching at most 0.25 for $S/N = 2$.

Comparison with ICA

For comparison, we then applied ICA, one of the state-of-the-art approaches for unsupervised analysis of fMRI data, to the same simulated set. ICA tries to decompose the BOLD signal at each voxel into a linear combination of signals (“timecourses”), in such a way that the spatial maps given by the linear coefficients of each timecourse at all voxels are maximally independent. These maps are called “component maps.” The component maps can be spatially overlapping, at variance with CDPC where each voxel can be assigned to only one cluster. The participation of a given voxel to each component is expressed in terms of a z -score.

For all values of S/N , ICA finds five independent components. In ICA, all three activated regions are associated with the 1st component. The first two regions can be resolved from the last one by looking at the sign of z : the first two regions correspond to positive z values and the third region to negative z values. This is not surprising, given that the two signals in the three regions are anti-correlated. The other components are spurious, noisy components.

In Figure 2B we plot the number of TPs and FPs as a function of the S/N ratio, considering a threshold of significance of $z = 2.3$, which corresponds to a significance value of $P < 0.01$. With this threshold, the number of FP is always higher than the number of TP. Further analysis at various z levels (Supporting Information Fig. S2) indicates that the best z threshold for discriminating between FP and TP depends on the S/N ratio, but is between 3 and 4. However, even for threshold of $z > 3$, corresponding to a very restrictive statistical criterion ($P < 0.001$), the clustering results are much cleaner than those of ICA. Moreover, for several choices of the z -threshold the number of FPs exceeds the number of TPs. This is clearly at odds with the results obtained with our approach, where all voxels assigned to cluster 1 or 2 mostly correspond to TPs regardless of their density ρ , without need of any arbitrary cutoff threshold in ρ . From this analysis, it is clear that the reliability of the CDPC approach surpasses that of ICA. In Figure 2C we show a rendering of the voxels found by the CDPC (voxels assigned to a cluster) and by ICA (voxels with $z = 2.3$) for the different S/N ratios considered. One can clearly see the large number of false positive voxels (i.e., green points outside the spheres) generated by ICA, which is in contrast with the cleaner performance of the CDPC results.

We also benchmarked the performance of another clustering method previously applied to fMRI, fuzzy c -means clustering [Baumgartner et al., 2000], on the same simulated data discussed in this section. The results are shown in

Supporting Information. We obtain a significantly worse performance than CDPC, in agreement with the general findings of Rodriguez and Laio, [2014], where a systematic comparison between the Density Peak algorithm and other clustering algorithms was performed.

Test of CDPC Analysis on a Motor Task

As a next step, we validated our method on real data. We chose to analyze a simple motor task, in which there are clear expectations on the brain regions that should be involved. Functional MRI data were recorded while seven subjects were visually instructed to alternate right-hand or left-hand clenching and rest (see “Materials and Methods” section). Our main goal is to show that, also when considering real data, our method is able to extract meaningful signal even from a short-time window and in a single subject.

Detailed Analysis of a Single Subject

We first consider a single subject (“ s_1 ”) and start by analyzing a single time window of 12 volumes. We applied CDPC to volumes 19–30, the first-time window in which the subject was instructed to move his left hand. In Figure 3A, we show the decision graph. Among the putative cluster centers (the 10 points with the highest δ , see “Materials and Methods” section), only six, highlighted in purple in the graph, yield sizable clusters.

The first cluster includes 2,833 voxels, and corresponds to the maximum density peak, which clearly appears as an isolated point in the decision graph. The other five clusters include 741, 446, 327, 298, and 126 voxels, respectively. This indicates that in this specific time window the activation is well described by a single relatively large cluster of voxels behaving coherently. In Figure 3C we present a surface rendering of the dominant cluster, where voxels are in different colors according to their density. The dominant cluster has a strong density peak in the right primary motor cortex. In addition, voxels with lower density are present in the occipital cortex and in the supplementary motor area. The density ρ of a voxel represents the number of other voxels (far or near) whose signal is similar. Thus, a voxel with higher density has a low dissimilarity with most of the voxels within a cluster. In other words, the density of a voxel measures how much its time series is representative of the dominant time series in a cluster. In Figure 3F, we report the BOLD signals in window 19–30 of all voxels inside five largest connected regions in which cluster #1 is split. The BOLD signals in the five regions exhibit strong coherence, confirming that our clustering analysis is able to spot groups of voxels that behave coherently irrespective of their brain location. For comparison, we also report the BOLD signals in window 19–30 of voxels inside three connected regions of a different cluster (cluster #5) that has a peak in the parietal region. In

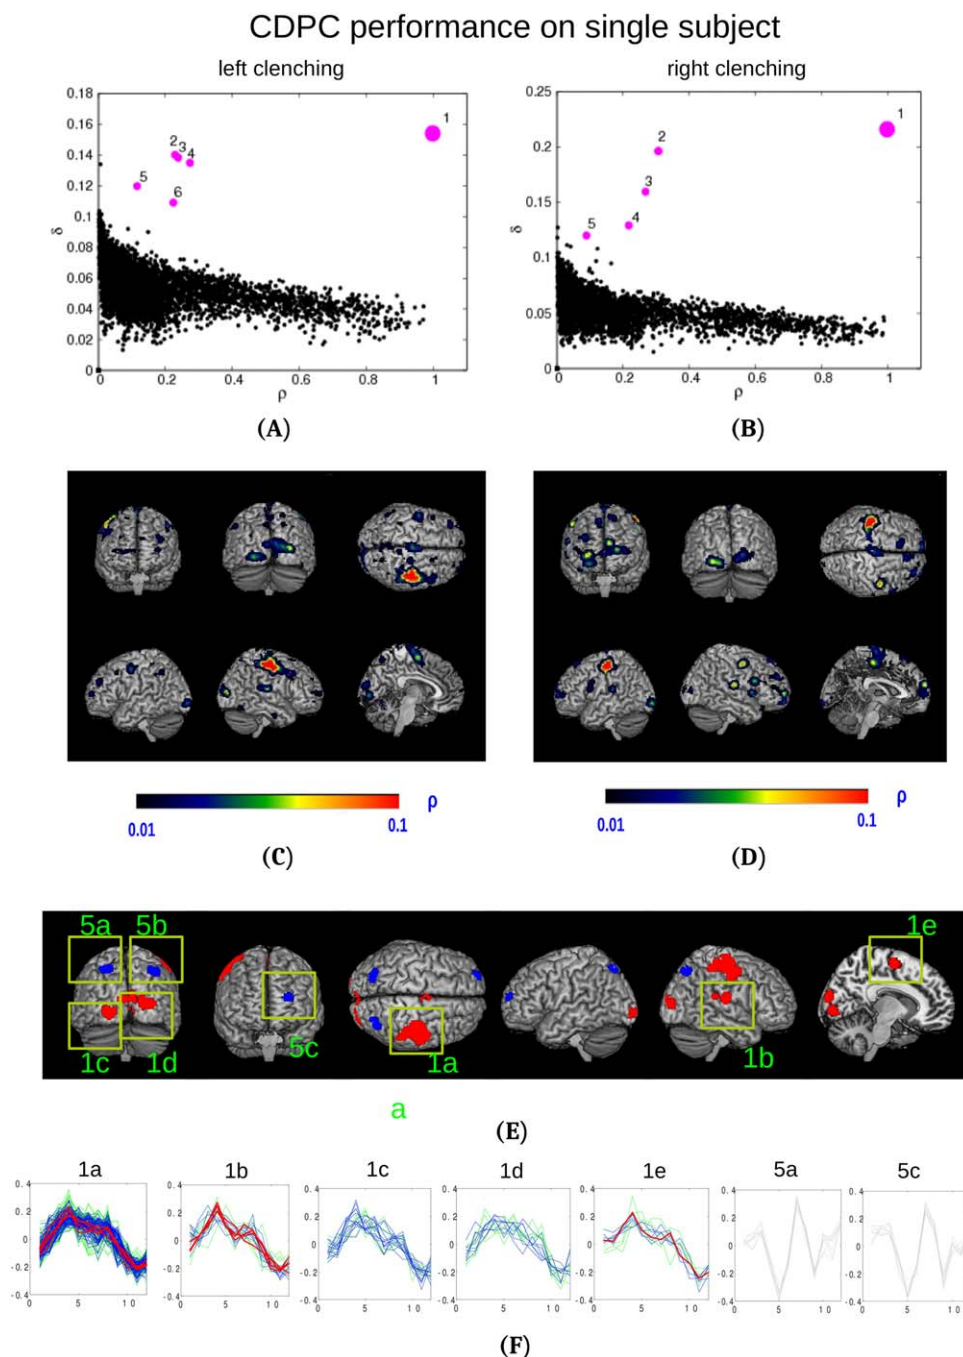


Figure 3.

(A,B): Decision graphs for windows 19–30 (left-hand movement, A) and windows 43–54 (right-hand movement, panel B) for single subject s_1 . We highlight in purple the peaks corresponding to the sizable clusters identified with our method. (C,D): Rendering of the voxels belonging to the first cluster for window 19–30 (left-hand clenching, C) and window 42–53 (right-hand clenching, D). The color of the voxels corresponds to their density (i.e., the horizontal axis of the decision graphs in panels A and B). The density maps have been thresholded at

a value of $\rho = 0.3$. (E) Connected regions of the first cluster with $\rho > 0.5$ (red) and connected regions of the fifth cluster with $\rho > 0.05$ (blue). (F) BOLD signal in window 19–30 for all voxels in the regions in panel E. The signal of each voxel is highlighted in different colors depending on the density of the voxels: (red) $\rho > 0.9$, (blue) $0.9 > \rho > 0.6$, (green) $0.6 > \rho > 0.3$, (gray) $\rho > 0.3$. [Color figure can be viewed at wileyonlinelibrary.com]

DPC performance on many subjects

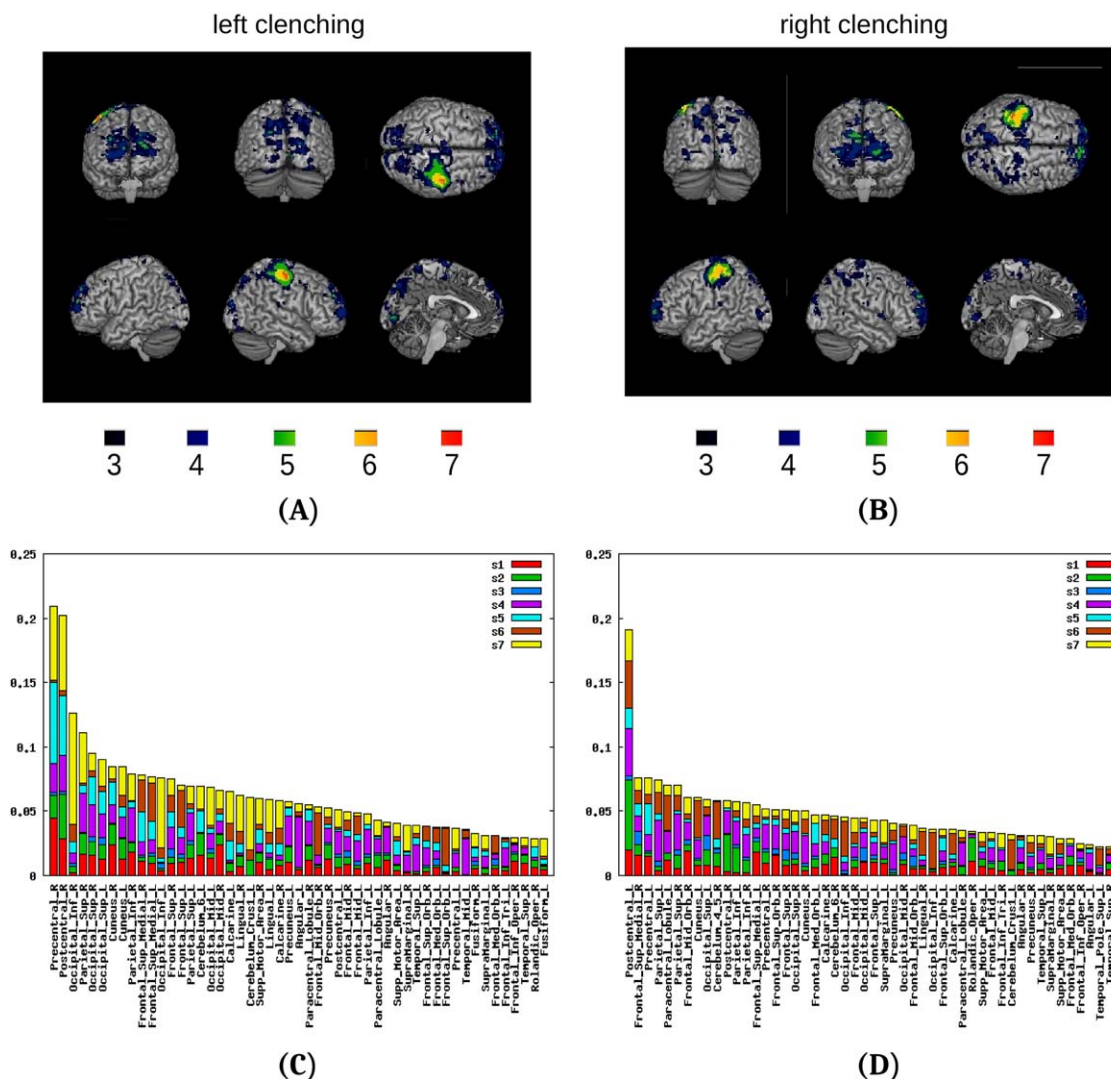


Figure 4.

(A,B) Location of the clusters for the seven subjects for window 19–30 (left-hand clenching, panel A) and 43–54 (right-hand clenching, panel B). In different colors, we show the voxels that pass a density threshold for a different number of subjects: (blue) 3 subjects, (dark green) 4 subjects, (light green) 5 subjects, (orange) 6 subjects, (red) 7 subjects. The density map of

each subject has been normalized to MNI space, smoothed with 9 mm Gaussian smoothing and thresholded at a value of 0.1. (C,D): histogram of the average density of AAL regions for the seven subjects for window 19–30 (left-hand clenching, panel C) and 43–54 (right-hand clenching, panel D). [Color figure can be viewed at wileyonlinelibrary.com]

agreement with expectation, the signals of voxels in cluster #5 are also coherent but very different from signals of cluster #1.

The CDPC procedure was repeated for the volumes 43–54, where the subject was instructed to move his right hand. Qualitatively the results are similar. In Figure 3B, we show the corresponding decision graph. The clusters are in this case five, including 2476, 727, 469, 225, and 172 voxels, respectively. The first cluster again emerges as the dominant

cluster. The dominant cluster (Fig. 3D) has a density peak in the left primary motor cortex. Voxels with lower density are found in occipital cortex, supplementary motor area, pre-frontal cortex, and the right primary motor cortex.

In summary, even when applied on a data set including only 12 volumes and one subject, our method is able to detect a clear signal in several brain regions that are expected to be involved in a visually instructed simple motor task, namely motor cortex and visual cortex.




Figure 5.

Regions of high time-averaged density for subject s_1 , representing the voxels that belong to clusters in many time windows over the experiment. The map is obtained by averaging the density of each voxel over the 91 windows of the experiment. We show the voxels whose time-averaged density is higher than 0.25. We grouped them into seven regions, shown in different colors: 1: primary motor area; 2: inferior occipital; 3: medial

frontal lobe; 4: inferior parietal; 5: superior frontal; 6: supplementary motor area; 7: right motor area. The average density map has been normalized and smoothed with 9-mm Gaussian smoothing (1R and 7 look nearly connected as an effect of normalization and smoothing but they are not in the native space). [Color figure can be viewed at wileyonlinelibrary.com]

Robustness of Results W.R.T. Changes in the Parameters

The results obtained are robust with respect to changes in free parameters entering the CDPC procedure, as well as the preprocessing steps performed. Comparing the clustered voxels obtained with different values of the dissimilarity cutoff d_c (and correspondingly of the coherent neighbor cutoff n_0), we obtain overlaps of the order of 80% or larger (Supporting Information Fig. S3). We find average overlaps of the order of 60% or larger (Supporting Information Fig. S4) for data that have a different preprocessing pipeline (e.g., regression of motion parameters).

Results for All Subjects

We repeated the analysis on the same time windows (volumes 19–30 and volumes 43–54) for the other six participants (Supporting Information Fig. S5). In most cases the first cluster has a peak in the left motor cortex (for right-hand movement) or right motor cortex (for left-hand movement). Voxels that pass a density threshold for at least three subjects are shown in Figure 4A,B for window 19–30 and 43–54, respectively. To provide a quantitative comparison between all subjects, we normalized the density map of each subject and overlapped it with the AAL atlas [Tzourio-Mazoyer et al., 2002], computing the average density in each AAL region. In Figure 4C,D we present cumulative histograms of this quantity for all subjects and for windows 19–30 and 43–54. From this figure, we can clearly see that the highest density concentrates in

specific regions. In particular, in window 19–30 (left-hand clenching) the two regions with the highest density are the right precentral and right postcentral that include the right motor cortex. Symmetrically, in window 43–54 (right-hand clenching) the two regions with highest density are the left precentral and left postcentral. In addition, other areas have consistently high density for several subjects in both windows, including: visual cortex, somatosensory cortex, prefrontal cortex, and cerebellum.

In short, the results of the CDPC method are consistent and reproducible between different subjects. The primary motor areas are identified in every subject even by applying the procedure only on a very short-time window (12 volumes). Other areas plausibly related to the task (e.g., visual cortex) also show a coherent signal for most of the subjects.

Consistency of CDPC Results Across Time Windows

Since during the fMRI acquisition the subject was instructed to repeatedly move the left or the right hand, one expects the same clusters to appear again and again during the experiment. To assess the reliability of our method we analyzed the entire run of subject s_1 in consecutive windows of 12 volumes, yielding 91 cluster analyses covering the 102 volumes in the run. A movie showing the clusters as a function of time is available at <https://github.com/micheleallegra/CDPC>. To provide evidence for the expected regularity, we averaged the density maps

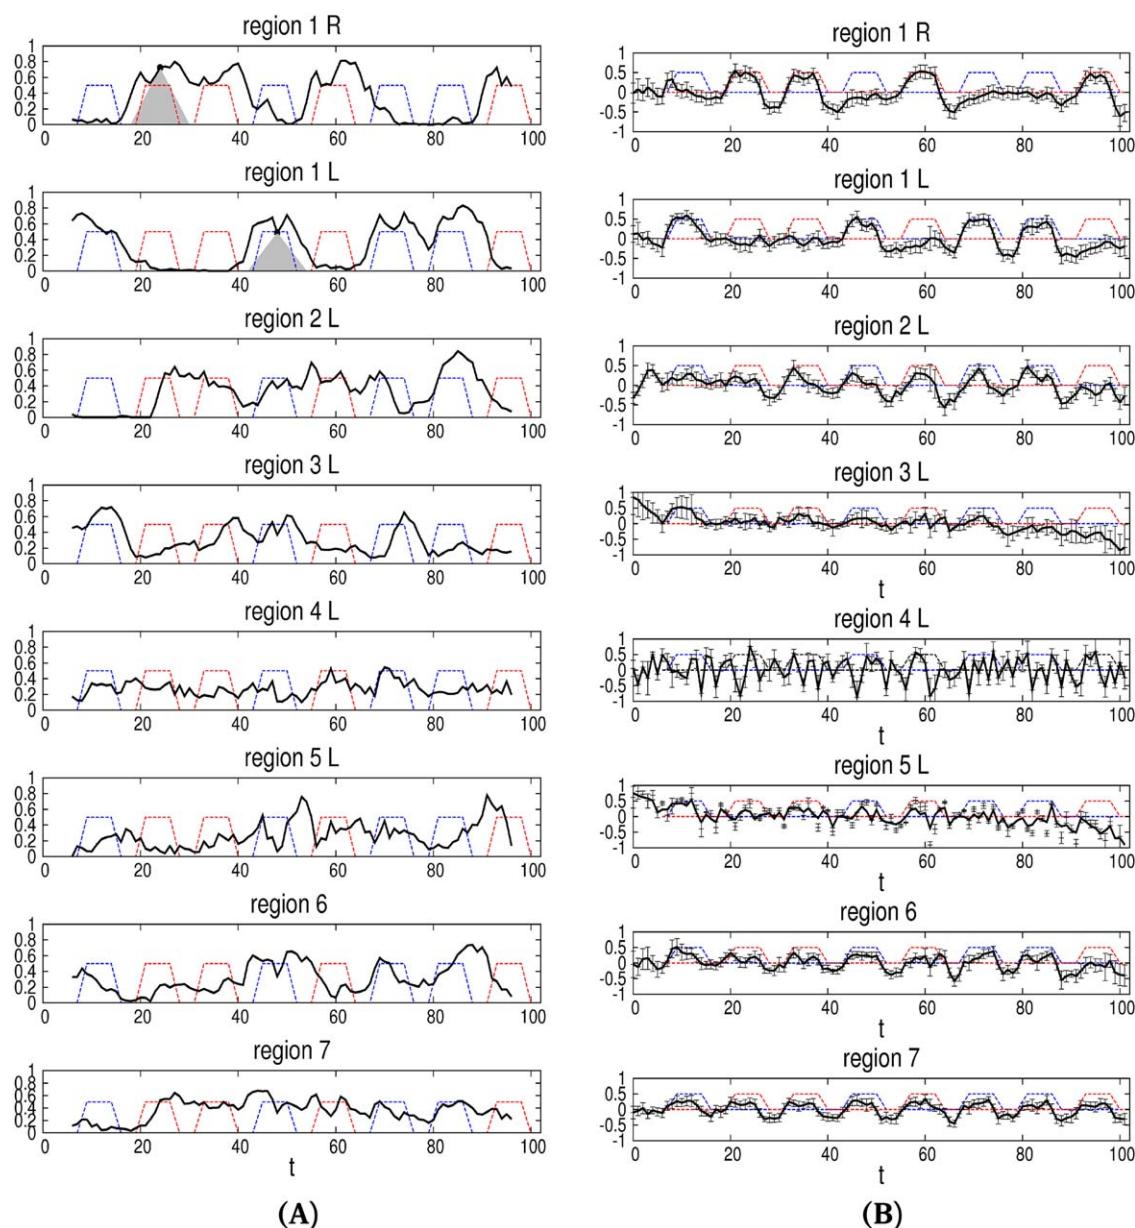


Figure 6.

(A) Average density of regions 1–7 as a function of the time window. For the L/R symmetric regions 2–5 we only show the L curve, as the R curve is nearly identical. The time t represents the center of a window of 12 volumes, as shown by the shaded gray areas that correspond to the two windows 19–30 and 43–

54 studied in detail above. Dashed blue/red lines represent the experimental paradigm with L/R hand clenching. (B) Average BOLD signal for voxels in the seven regions. Error bars represent standard deviation. [Color figure can be viewed at wileyonlinelibrary.com]

over all 91 time windows, resulting in a time-averaged density map. To identify voxels that are often clustered during the experiment, we cut this map at a threshold of 0.25, which yields approximately 300 voxels above threshold. Voxels passing the threshold are shown in Figure 5A.

We have spatially grouped these voxels into twelve small regions around each time-averaged density peak. By putting pairs of regions that are approximately left/right symmetric under the same label, we obtain seven regions, which are depicted with different colors in Figure 5.

ICA performance on many subjects

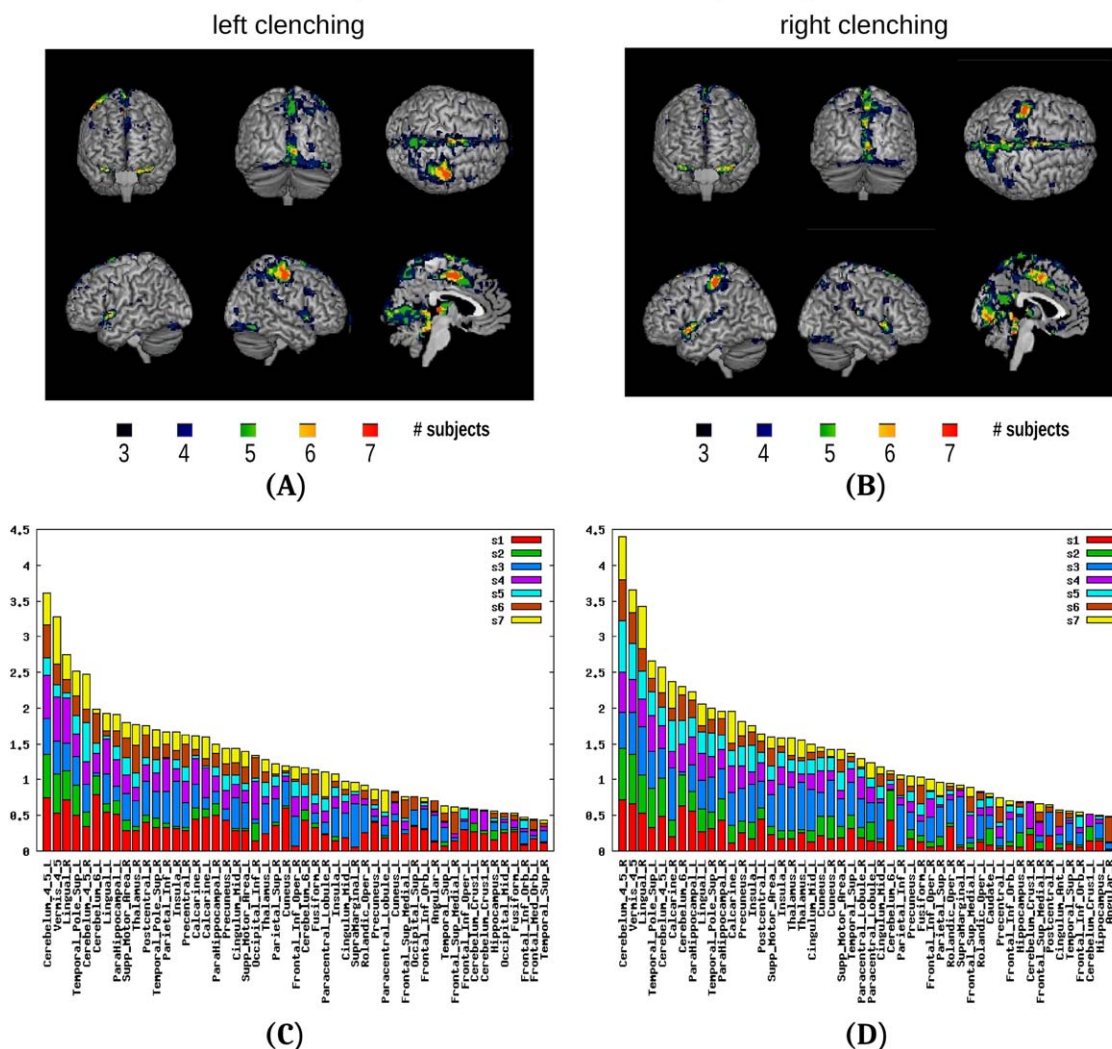


Figure 7.

(A,B) ICA map for the seven subjects for window 19–30 (left-hand clenching, panel A) and 43–54 (right-hand clenching, panel B). In different colors, we show the voxels that simultaneously pass a value of $z = 0.7$ for a different number of subjects: (blue) 3 subjects, (dark green) 4 subjects, (light green) 5 subjects, (orange) 6 subjects, (red) 7 subjects. For each subject, we have taken the quadratic sum of the z -scores of all independent components to produce a global z . The z map of each subject has

Tracking the Clusters Temporal Evolution

Since the method is especially suitable to discover and track short-lived patterns of activity, we investigated the temporal evolution of the clusters by computing the average density of the voxels for different regions in each time window. In Figure 6A we plot the average density for the regions 1–7 overlapped with the task time-line (red = right-

hand movement, blue = left-hand movement). Note that each point in the curves actually represents a 12-volume time window, which is exemplified by showing the time windows covered by two points in Figure 6A (shaded areas). We find that the density in the frontal (3, 5 in Fig. 6), visual (2) and parietal (4) regions is bilateral, and density of the two sides exhibits strong correlation. By contrast, the density of the primary motor region (1) is strongly

hand movement, blue = left-hand movement). Note that each point in the curves actually represents a 12-volume time window, which is exemplified by showing the time windows covered by two points in Figure 6A (shaded areas). We find that the density in the frontal (3, 5 in Fig. 6), visual (2) and parietal (4) regions is bilateral, and density of the two sides exhibits strong correlation. By contrast, the density of the primary motor region (1) is strongly

lateralized and clearly shows a correlation with the left-hand/right-hand movement. The remaining regions show a density that is fluctuating but never vanishing during the experiment. This means that they almost constantly belong to some cluster during the whole experiment. In most time windows, regions 1 (L or R depending on the clenching hand), 2 (both L and R), 6, 7 are assigned to the same cluster (usually the first), implying that our approach is able to recognize when anatomically distant regions simultaneously participate to a single coherent process. In Figure 6B we plot the average BOLD signals of voxels belonging to the same seven regions together with their standard deviation. One can clearly see that regions 1, 2, 6, 7 have a time series that is similar and strongly correlated with the paradigm. Regions 4, 5 have a different time series, and in most time windows they are assigned to a different cluster (as we saw, e.g., for window 19–30). Region 3 shows a correlation with the task only in some windows and correspondingly it is clustered together with 1, 2, 6, 7 only in these windows.

Comparison with ICA

We now compare the results of CDPC with those of ICA. Similar to what was observed in the analysis of artificial data, in the motor task experiment ICA detects as meaningful not only regions that are clearly involved with the task, but also several other areas that are likely to be false positives. For both windows 19–30 and 43–54, we found that the z scores of ICA were very low ($z < 1$), which implies the failure of ICA of finding a significant signal. Therefore, in order to produce maps of significant voxels we could not use a conventional threshold (e.g., $z > 2.3$), and we used instead a threshold of $z = 0.7$ allowing to gather a sizable number of significant voxels.

Voxels that pass this z threshold for at least three subjects are shown in Figure 7A,B for window 19–30 and 43–54, respectively. The regions that are found for more than three subjects are many. In both windows, we again see the left/right M1, the SMA, and the cerebellum. In addition, ICA finds strong signal in many other areas: the parahippocampal gyrus, the thalamus, the calcarine, the central fissure, and the temporal lobe. We have normalized the z maps and superimposed the AAL atlas, averaging the z values over each AAL region for each subject. In Figure 7C,D, we show the corresponding AAL histograms. ICA detects signal in several regions, including many regions that are not plausibly related to the task: either because they are outside gray matter, or because they are in regions not associated to sensorimotor tasks (e.g., temporal pole, parahippocampal gyrus, insula). Furthermore, the BOLD time series of many of these regions could not be easily related with the main task events. Overall, this shows that ICA, while being able to spot the signal in areas that are strongly related to the task, is also much more sensitive than CDPC to spurious signal.

DISCUSSION AND CONCLUSIONS

The knowledge of brain function could benefit from a deeper understanding of how transient couplings between brain regions underlie the execution of different activities [Calhoun et al., 2014; Hutchison et al., 2013; Kopell et al., 2014]. However, investigations in this subject are hindered by the current limitations of fMRI data analysis methods. In this work we have proposed CDPC, a novel fMRI data analysis approach specifically tailored to study short-lived brain activity patterns. Our method builds on the DPC algorithm, which allows finding and grouping together voxels that—irrespective of their brain location—have a similar time development even during a short-time interval.

We first tested CDPC on simulated data, then we applied it on a real data set. The test of CDPC on simulated fMRI data provided clear-cut results about the effectiveness of our method. We probed our method against different demanding conditions. We created simulated data having a large data space (the size of a brain image was of the order of 100k voxels), a low signal-to-noise ratio (down to $S/N = 2$), a significant spatial and temporal correlation of the noise, and a small number of images (12 volumes). Despite these quite challenging conditions, CDPC was always able to detect the signal even at the lower values of S/N , while it proved very robust against false positives. By contrast, an established method for unsupervised fMRI analysis (ICA) provided less reliable results on the same data, as it detected a significantly higher number of false positives, even overcoming the number of true positives for low S/N ratios.

As the analysis of simulated data yielded evidence of the effectiveness of the method, we then turned to test CDPC on real fMRI data. We studied a motor task in which participants were required to perform clenches with either the left or the right hand, so as to be able to compare the results of CDPC analysis with the clear expectations one can have about the brain areas involved in the task [Ehrsson et al., 2003; Lotze et al., 1999, 2000; Mayka et al., 2006; Stippich et al., 2002; Witt et al., 2008]. We analyzed time windows of 12 volumes, matching the conditions of the simulated data. Focusing on a single subject and on only two time windows centered on a left-hand and right-hand clench, we found a dominant cluster involving two expected regions: the primary motor cortex (contralateral with respect to the left/right clenching), and the primary visual cortex. A direct analysis of the BOLD signal confirmed that voxels in these regions had very similar time series in the time windows under analysis, directly showing the presence of a coherent activity involving these distant areas. Analysis of all time windows in the same subject revealed that these regions are involved in coherent patterns in the different trials of the experiment (Fig. 6). The BOLD signal of most of the regions detected in multiple trials shows a clear relationship with the different task phases (right/left clenching,

rest). Interestingly, CDPC also detected two regions (left lateral parietal lobe and left dorsolateral frontal lobe) that showed a pattern of BOLD response not clearly related to the task phases. The fact that they exhibited a highly coherent signal throughout the task is a hint for their task relevance, but we do not currently have enough evidence to assign a functional role to these regions. Nevertheless, this shows how CDPC would allow detecting potentially task-relevant regions even if their BOLD signal does not have an obvious correlation with the task at hand. Finally, when all subjects are considered, CDPC detected the motor area in all subjects and the visual cortex in most of the subjects, again only by relying on a single time window of 12 volumes. Consistently with simulation results, also ICA detected both regions but in a context of much noisier results.

In short, the study of both simulated and real data showed that CDPC can accurately identify coherent signals in short-time windows, with little noise. This demonstrates that well-defined spatio-temporal patterns can be reliably identified by CDPC even in very short fMRI data series. Two ingredients are critical for the successful performance of CDPC. One is the intrinsic power of the DPC clustering algorithm [Rodriguez and Laio, 2014]. The other is our choice of suitable noise filters that allow for an optimized detection of genuine signal. In particular, the noise filter embedded in the definition of the density is essential. In its absence, one might face several undesired consequences: (1) appearance of spurious clusters of noisy origin, (2) bad source discrimination, as the true clusters would be contaminated with spurious signal, (3) appearance of very large clusters, including many voxels that do not really participate to a coherent process.

While the experiment and the simulation described in this work provide a first validation of the method, only future applications of CDPC to more complex fMRI experiments will reveal its potential and limits. We can here point out some possible limitations. First, our method assigns voxels to a single cluster. This turned out to be appropriate for the test cases discussed in this article, where clusters are spatially well-separated. For more complex, overlapping situations, we would need to extend the method in the direction of a fuzzy clustering allowing a probabilistic assignation of a voxel to more than a single cluster or a hierarchical clustering able to reliably split clusters into subclusters. Second, the type of noise filter we introduced cannot be guaranteed to remove all sources of noise. In particular, head-movement related artifacts might survive the filtering, as they can be a source of spurious coherence between spatially near voxels [Power et al., 2012] and they can be present in the high-frequency domain we investigate. A refinement of our filtering procedure might be required in order to optimally cope with such artifacts. More in general, we based the discrimination of signal and noise on heuristic principles, rather than rigorous statistical arguments. Unfortunately, the current lack of advanced theoretical models of noise in fMRI data

implies that we cannot use a statistical test to discriminate signal from noise, since a realistic null model would be required for this aim [Eklund et al., 2015]. Only further development in this area will enable a refinement of our method in the direction of stricter statistical rigor.

Besides its suitability to investigate short-time windows, our approach offers other advantages. First, it is completely unsupervised, as it does not rest upon specific assumptions on signal shape, timing, strength, or the number of possible clusters to be found. Second, in contrast to univariate methods for single-trial unsupervised analysis like paradigm-free mapping [Caballero et al., 2011; 2013] CDPC does not focus on “activation events,” but rather on coherent patterns of activation across multiple brain networks. Third, the procedure requires limited computational power: a single-subject, single time-window analysis runs in less than 5 min on a standard desktop PC. As a result, we can implement a “sliding window” analysis applying the procedure to all time-windows of the same length centered at successive volumes, capturing the growth and decay of short-lived coherent activation patterns. Finally, the coherence-based character of CDPC allows integrating this approach within the framework of ongoing research efforts trying to bridge macroscopic functional connectivity in fMRI with microscopic effective connectivity induced by neural dynamics such as neural oscillations [Deco and Kringelbach, 2016; Kopell et al., 2014].

We can envisage several applications of our method that we have not explored in this manuscript. Our method can be used to construct a dynamical network description of brain dynamics, identifying a set of brain regions that play the role of network nodes, and specifying for each pair of nodes a (time-dependent) connectivity strength. The average density allows identifying the regions that are recurrently recruited in a process and thus are apt to form the nodes of a network, while the participation of these nodes to the same clusters as the time window is moved can provide the basis for computing a time-dependent connectivity structure. Furthermore, the ability of CDPC of robustly identifying the relevant connectivity structure even in short-time windows would allow tracking long range network reconfigurations that may be associated to strategy changes in learning or problem solving, even when these changes happen only rarely or once per subject [Badre et al., 2010; Reverberi et al., 2005; Schuck et al., 2015; Seyed-Allaei et al., 2010]. Finally, the metric may be extended to include other factors of functional (or pathological) relevance beyond the coherence of BOLD signals—for example, a voxel-specific sensitivity to a given drug. Thus, the robustness of within subject network identification would allow to effectively explore between subject differences of functional network and how these may be predictive of cognitive variation or clinical status [Baldassarre et al., 2012; Barch et al., 2013; Martin et al., 2012; Prat and Just, 2011; Reverberi et al., 2012].

ACKNOWLEDGMENTS

We thank Dante R. Chialvo for his careful reading of the manuscript and his important suggestions. MA thanks Arianna Montorsi and Riccardo Zecchina for their kind support.

REFERENCES

- Badre D, Kayser AS, D'Esposito M (2010): Frontal cortex and the discovery of abstract action rules. *Neuron* 66:315–326.
- Baert AL, Moonen CTW, Bandettini PA (2000): *Functional MRI*. New York: Springer.
- Baldassarre A, Lewis CM, Committeri G, Snyder AZ, Romani GL, Corbetta M (2012): Individual variability in functional connectivity predicts performance of a perceptual task. *Proc Natl Acad Sci* 109:3516–3521.
- Barch DM, Burgess GC, Harms MP, Petersen SE, Schlaggar BL, Corbetta M, Glasser MF, Curtiss S, Dixit S, Feldt C, Nolan D, Bryant E, Hatley T, Footer O, Bjork JM, Poldrack R, Smith S, Johansen-Berg H, Snyder AZ, Van Essen DC (2013): Function in the human connectome: Task-fMRI and individual differences in behavior. *NeuroImage* 80:169–189.
- Bassett DS, Wymbs NF, Porter MA, Mucha PJ, Carlson JM, Grafton ST (2011): Dynamic reconfiguration of human brain networks during learning. *Proc Natl Acad Sci* 108:7641–7646.
- Baumgartner R, Ryner L, Richter W, Summers R, Jarmasz M, Somorjai R (2000): Comparison of two exploratory data analysis methods for fMRI: Fuzzy clustering vs. principal component analysis. *Magn Reson Imaging* 18:89–94.
- Beckmann CF, Smith SM (2004): Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Trans Med Imaging* 23:137–152.
- Braun U, Schäfer A, Walter H, Erk S, Romanczuk-Seiferth N, Haddad L, Schweiger JI, Grimm O, Heinz A, Tost H, Meyer-Lindenberg A (2015): Dynamic reconfiguration of frontal brain networks during executive cognition in humans. *Proc Natl Acad Sci* 112:11678–11683.
- Bullmore E, Brammer M, Williams SCR, Rabe-Hesketh S, Janot N, David A, Mellers J, Howard R, Sham P (1996): Statistical methods of estimation and inference for functional MR image analysis. *Magn Reson Med* 35:261–277.
- Caballero Gaudes C, Petridou N, Dryden IL, Bai L, Francis ST, Gowland PA (2011): Detection and characterization of single-trial fMRI bold responses: Paradigm free mapping. *Hum Brain Mapp* 32:1400–1418.
- Caballero Gaudes C, Petridou N, Francis ST, Dryden IL, Gowland PA (2013): Paradigm free mapping with sparse regression automatically detects single-trial functional magnetic resonance imaging blood oxygenation level dependent responses. *Hum Brain Mapp* 34:501–518.
- Calhoun VD, Adali T, Pearlson GD, Pekar JJ (2001): A method for making group inferences from functional MRI data using independent component analysis. *Hum Brain Mapp* 14:140–151.
- Calhoun VD, Miller R, Pearlson G, Adali T (2014): The chronnectome: time-varying connectivity networks as the next frontier in fmri data discovery. *Neuron* 84:262–274.
- Coulon O, Pizzagalli F, Operto G, Auzias G, Delon-Martin C, Dojat M (2011) Two new stable anatomical landmarks on the central sulcus: definition, automatic detection, and their relationship with primary motor functions of the hand. In: *IEEE EMBC 2011, 33rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. pp 7795–7798.
- Craddock RC, James GA, Holtzheimer PE, 3rd, Hu XP, Mayberg HS (2012): A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Hum Brain Mapp* 33:1914–1928.
- Dale A (1999): Optimal experimental design for event-related fMRI. *Hum Brain Mapp* 8:109–114.
- Daubechies I, Roussos E, Takerkart S, Benharrosh M, Golden C, D'ardenne K, Richter W, Cohen JD, Haxby J (2009): Independent component analysis for brain fMRI does not select for independence. *Proc Natl Acad Sci* 106:10415–10422.
- Deco G, Kringelbach ML (2016): Metastability and coherence: Extending the communication through coherence hypothesis using a whole-brain computational perspective. *Trends Neurosci* 39:125–135.
- Dimitridou E, Barth M, Windischberger C, Hornik K, Moser E (2004): A quantitative comparison of functional MRI cluster analysis. *Artif Intell Med* 31:57–71.
- Ehrsson HH, Geyer S, Naito E (2003): Imagery of voluntary movement of fingers, toes, and tongue activates corresponding body-part-specific motor representations. *J Neurophysiol* 90:3304–3316.
- Eklund A, Nichols T, Knutsson H (2015): Can parametric statistical methods be trusted for fMRI based group studies? *arXiv* 1511.01863.
- Friston KJ, Holmes AP, Worsley KJ, Poline J-P, Frith CD, Frackowiak RSJ (1995): Statistical parametric maps in functional imaging: A general linear approach. *Hum Brain Mapp* 2:189–210.
- Genovese CR, Noll DC, Eddy WF (1997): Estimating test-retest reliability in functional MR imaging. I: Statistical methodology. *Magn Reson Med* 38:497–507.
- Hansen LK, Larsen J, Nielsen FA, Nielsen A, Strother SC, Rostrup E, Savoy R, Svarer C, Paulson OB (1999): Generalizable patterns in neuroimaging: How many principal components?. *Neuroimage* 9:534–544.
- Henson, RNA (2007) Efficient experimental design for fMRI. In: Penny WD, Friston KJ, Ashburner JT, Kiebel SJ, Nichols TE, editors. *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. New York: Academic Press.
- Himberg J, Hyvärinen A (2003) Icasto: Software for investigating the reliability of ICA estimates by clustering and visualization. In: *IEEE 13th Workshop on Neural Networks for Signal Processing*. pp 259–268.
- Hutchison RM, Womelsdorf T, Allen EA, Bandettini PA, Calhoun VD, Corbetta M, DallaPenna S, Duyn JH, Glover GH, Gonzales-Castillo J, Handwerker DA, Keilholz S, Kiviniemi V, Leopold DA, de Pasquale F, Sporns O, Walter M, Chang C (2013): Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage* 80:360–378.
- Jezzard P, Matthews PM, Smith SM (2001) *Functional MRI an Introduction to Methods*. Oxford: Oxford University Press.
- Kalcher K, Boubela RN, Huf W, Biswal BB, Baldinger P, Sailer U, Filzmoser P, Kasper S, Lamm C, Lanzenberger R, Moser E, Windischberger C (2013): RESCALE: Voxel-specific task-fMRI scaling using resting state fluctuation amplitude. *Neuroimage* 70:80–88.
- Kopell NJ, Gritton HJ, Whittington MA, Kramer MA (2014): Beyond the connectome: The dynamome. *Neuron* 83:1319–1328.
- Lai SH, Fang M (1999): A novel local PCA-based method for detecting activation signals in fMRI. *Magn Reson Imaging* 17:827–836.
-
-

- Logothetis NK (2008): What we can do and what we cannot do with fMRI. *Nature* 453:869–878.
- Lotze M, Erb M, Flor H, Huelsmann E, Godde B, Grodd W (2000): fMRI evaluation of somatotopic representation in human primary motor cortex. *NeuroImage* 11:473–481.
- Lotze M, Montoya P, Erb M, Hülsmann E, Flor H, Klose U, Birbaumer N, Grodd W (1999): Activation of cortical and cerebellar motor areas during executed and imagined hand movements: An fMRI study. *J Cogn Neurosci* 11:491–501.
- Martin A, Barnes KA, Stevens WD (2012): Spontaneous neural activity predicts individual differences in performance. *Proc Natl Acad Sci* 109:3201–3202.
- Mattay VS, Frank JA, Santha AK, Pekar JJ, Duyn JH, McLaughlin AC, Weinberger DR (1996): Whole-brain functional mapping with isotropic MR imaging. *Radiology* 201:399–404.
- Mayka MA, Corcos DM, Leurgans SE, Vaillancourt DE (2006): Three-dimensional locations and boundaries of motor and premotor cortices as defined by functional brain imaging: A meta-analysis. *NeuroImage* 31:1453–1474.
- McKeown MJ, Makeig S, Brown GG, Jung TP, Kindermann SS, Bell AJ, Sejnowski TJ (1998): Analysis of fMRI data by blind separation into independent spatial components. *Hum Brain Mapp* 6:160–188.
- Pizzagalli F, Auzias G, Delon-Martin C, Dojat M (2013): Local landmark alignment for high-resolution fMRI group studies: Toward a fine cortical investigation of hand movements in human. *J Neurosci Meth* 218:83–95.
- Power JD, Barnes KA, Snyder AZ, Schlaggar BL, Petersen SE (2012): Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage* 59:2142–2154.
- Prat CS, Just MA (2011): Exploring the neural dynamics underpinning individual differences in sentence comprehension. *Cereb Cortex* 21:1747–1760.
- Reverberi C, Toraldo A, D’Agostini S, Skrap M (2005): Better without (lateral) frontal cortex? Insight problems solved by frontal patients. *Brain* 128:2882–2890.
- Reverberi C, Gørgen K, Haynes J-D (2012): Distributed representations of rule identity and rule order in human frontal cortex and striatum. *J Neurosci* 32:17420–17430.
- Rodriguez A, Laio A (2014): Clustering by fast search and find of density peaks. *Science* 322:1492–1496.
- Schuck NW, Gaschler R, Wenke D, Heinzle J, Frensch PA, Haynes J-D, Reverberi C (2015): Medial prefrontal cortex predicts internally driven strategy shifts. *Neuron* 86:331–340.
- Seyed-Allaei S, Amati D, Shallice T (2010): Internally driven strategy change. *Think Reason* 16:308–331.
- Smith AM, Lewis BK, Ruttimann UE, Frank QY, Sinnwell TM, Yang Y, Duyn JH, Frank JA (1999): Investigation of low frequency drift in fMRI signal. *Neuroimage* 9:526–533.
- Stippich C, Ochmann H, Sartor K (2002): Somatotopic mapping of the human primary sensorimotor cortex during motor imagery and motor execution by functional magnetic resonance imaging. *Neurosci Lett* 331:50–54.
- Thirion B, Flandin G, Pinel P, Roche A, Ciuciu P, Poline JB (2006): Dealing with the shortcomings of spatial normalization: Multi-subject parcellation of fMRI datasets. *Hum Brain Mapp* 27:678–693.
- Turner R, Howseman A, Rees GE, Josephs O, Friston K (1998): Functional magnetic resonance imaging of the human brain: Data acquisition and analysis. *Exp Brain Res* 123:5–12.
- Tzourio-Mazoyer N, Landeau B, Papathanassiou D, Crivello F, Etard O, Delcroix N, Mazoyer B, Joliot M (2002): Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *Neuroimage* 15:273–289.
- Welvaert M, Durnez J, Moerkerke B, Verdoolaege G, Rosseel Y (2011): NeuRosim: An R package for generating fMRI data. *J Stat Softw* 44:1–18.
- Witt ST, Laird AR, Meyerand ME (2008): Functional neuroimaging correlates of finger-tapping task variations: An ALE meta-analysis. *NeuroImage* 42:343–356.
- Yee S-H, Gao J-H (2002): Improved detection of time windows of brain responses in fMRI using modified temporal clustering analysis. *Magn Reson Imaging* 20:17–26.