

Low-dimensional embedding of fMRI datasets

Xilin Shen and François G. Meyer*

Department of Electrical Engineering, University of Colorado at Boulder, USA

Received 13 September 2007; revised 7 January 2008; accepted 25 February 2008
Available online 7 March 2008

We propose a novel method to embed a functional magnetic resonance imaging (fMRI) dataset in a low-dimensional space. The embedding optimally preserves the local functional coupling between fMRI time series and provides a low-dimensional coordinate system for detecting activated voxels. To compute the embedding, we build a graph of functionally connected voxels. We use the commute time, instead of the geodesic distance, to measure functional distances on the graph. Because the commute time can be computed directly from the eigenvectors of (a symmetric version) the graph probability transition matrix, we use these eigenvectors to embed the dataset in low dimensions. After clustering the datasets in low dimensions, coherent structures emerge that can be easily interpreted. We performed an extensive evaluation of our method comparing it to linear and nonlinear techniques using synthetic datasets and in vivo datasets. We analyzed datasets from the EBC competition obtained with subjects interacting in an urban virtual reality environment. Our exploratory approach is able to detect independently visual areas (V1/V2, V5/MT), auditory areas, and language areas. Our method can be used to analyze fMRI collected during “natural stimuli”.

© 2008 Elsevier Inc. All rights reserved.

Keywords: fMRI; Laplacian eigenmaps; Embedding; Natural stimuli

Introduction

At a microscopic level, a large number of internal variables associated with various physical and physiological phenomena contribute to dynamic changes in functional magnetic resonance imaging (fMRI) datasets. fMRI provides a large scale (as compared to the scale of neurons) measurement of neuronal activity, and we expect that many of these variables will be coupled resulting in a low-dimensional set for all possible configurations of the activated fMRI signal. We assume therefore that activated fMRI time series can be parametrized by a small number of variables. This assumption is consistent with the usage of low-dimensional parametric models for detecting activated voxels (Pettersson et al., 1999). This assumption is

also consistent with the empirical findings obtained with principal components analysis (PCA) and independent components analysis (ICA) (Biswal and Ulmer, 1999; McKeown et al., 2003), where a small number of components are sufficient to describe the variations of most activated temporal patterns. Both PCA and ICA make very strong assumptions about the components: orthogonality and statistical independence, respectively. Such constraints are convenient mathematically but have no physiological justification, and complicate unnecessarily the interpretation of the components (Friston, 1998). A second limitation of PCA and ICA is that both methods only provide a linear decomposition of the data (Friston, 2005). There is no physiological reason why the fMRI signal should be a linear combination of eigen-images or eigen-time series. In practice, the first components identified by PCA are often related to physiological artifacts (e.g. breathing), or coherent spontaneous fluctuations (Raichle and Mintun, 2006). These artifacts can be responsible for most of the variability in the dataset. Stimulus triggered changes, which are much more subtle, rarely appear among the first components.

The contribution of this paper is a novel exploratory method to construct an optimal coordinate system that reduces the dimensionality of the dataset while preserving the functional connectivity between voxels (Sporns et al., 2000). First, we define a distance between time series that quantifies the functional coupling (Fox et al., 2005), or connectivity between the corresponding voxels. We then construct an embedding that preserves this functional connectivity across the entire brain. After embedding the dataset in a lower dimensional space, time series are clustered into coherent groups. This new parametrization results in a clear separation of the time series into: (1) response to a stimulus, (2) coherent physiological signals, (3) artifacts, and (4) background activity. We performed an extensive evaluation of our method comparing it to linear and nonlinear techniques using synthetic datasets and in vivo datasets.

Methods

Overview of our approach

Our goal is to find a new parametrization of an fMRI dataset, effectively replacing the time series by a small set of features, or coordinates, that facilitate the identification of task-related hemody-

* Corresponding author.

E-mail address: fmeyer@Colorado.Edu (F.G. Meyer).

URL: <http://ece.colorado.edu/~fmeyer> (F.G. Meyer).

Available online on ScienceDirect (www.sciencedirect.com).

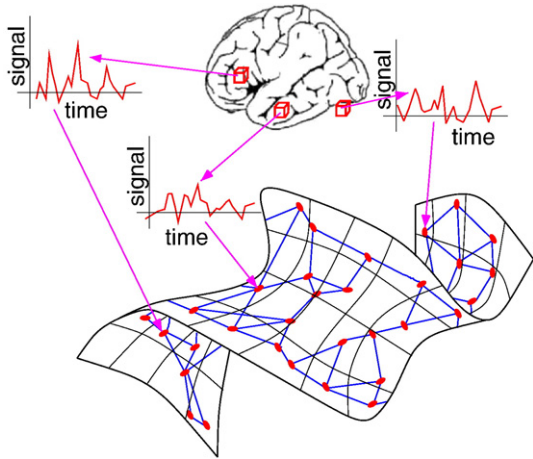


Fig. 1. The network of functionally correlated voxels, represented by a graph, encodes the functional connectivity between time series.

dynamic responses to the stimulus. The new coordinates will also be able to reveal the presence of physical or physiological processes that have an intrinsic low dimension. Such processes can be described with a small number of parameters (dimensions) in an appropriate representation (set of basis functions). These time series should be contrasted with noise time series that have a very diffuse representation in any basis. Example of such low-dimensional processes includes task-related hemodynamic responses, non-task-related physiological rhythms (breathing and heart beating motion). We expect that only a small fraction of all time series will have a low-dimensional representation. The remaining time series will be engaged in a spontaneous intrinsic activity (Raichle and Mintun, 2006). We call these time series *background time series*. As shown in our experiments, background time series are more complex, and cannot be well approximated with a small number of parameters.

We now introduce some notations that will be used throughout the paper. Let \mathbf{X} denote an fMRI dataset composed of T scans, each comprised of N voxels, which is represented as a $N \times T$ matrix,

$$\mathbf{X} = \begin{bmatrix} x_1(1) & \cdots & x_1(T) \\ \vdots & \vdots & \vdots \\ x_N(1) & \cdots & x_N(T) \end{bmatrix}. \quad (1)$$

Row i of \mathbf{X} is the time series $\mathbf{x}_i = [x_i(1), \dots, x_i(T)]$ generated from voxel i . Column j is the j th scan unrolled as a $1 \times N$ vector. In this work, we regard \mathbf{x}_i as a point in \mathbb{R}^T , with T coordinates.




Fig. 2. The eigenvector ϕ_k as a function defined on the nodes of the graph.

Algorithm 1: Construction of the embedding

Input:

- $\mathbf{x}_i(t), t = 0, \dots, T - 1, i = 1, \dots, N,$
- $\sigma; n_n; K$: dimension of the embedding,

Algorithm:

- (1) construct the graph defined by the n_n nearest neighbors of each \mathbf{x}_i
- (2) compute \mathbf{W} and \mathbf{D} . Find the first K eigenvectors, ϕ_k , of $\mathbf{D}^{-\frac{1}{2}}\mathbf{W}\mathbf{D}^{-\frac{1}{2}}$

Output: K coordinates of each $\mathbf{x}_i, \frac{1}{\sqrt{\lambda_k}} \{ \phi_k(i) / \sqrt{1 - \lambda_k} \}, k = 2, \dots, K + 1$

Fig. 3. Construction of the embedding.

We seek a new parametrization of the dataset that optimally preserves the local functional coupling between time series. Most methods of reduction of dimensionality used for fMRI are linear: each \mathbf{x}_i is projected onto a set of components ϕ_k . The resulting coefficients $\langle \mathbf{x}_i, \phi_k \rangle, k = 0, \dots, K - 1$ serve as the new coordinates in the low-dimensional representation. However, in the presence of nonlinearity in the organization of the \mathbf{x}_i in \mathbb{R}^T , a linear mapping may distort local functional correlations. This distortion will make the clustering of the dataset more difficult. Our experiments with in vivo data confirm that the subsets formed by the low-dimensional time series have a nonlinear geometry and cannot be mapped onto a linear subspace without significant distortion. We propose therefore to use a nonlinear map Ψ to represent the dataset \mathbf{X} in low dimensions. Because the map Ψ is able to preserve the local functional coupling between voxels, low-dimensional coherent structures can easily be detected with a clustering algorithm. Finally, the temporal and the spatial patterns associated with each cluster are examined and the cluster that corresponds to the task-related response is identified. In summary, our approach includes the following three steps:

- (i) Low-dimensional embedding of the dataset;
- (ii) Clustering of the dataset using the new parameterization;
- (iii) Identification of the set of activated time series.

The connectivity graph: A network of functionally correlated voxels

In order to construct the nonlinear map Ψ we need to estimate the functional correlation between voxels. We characterize the functional correlation between voxels with a network. Similar networks have been constructed to study functional connectivity in (Achard et al., 2006; Caclin and Fonlupt, 2006; Eguiluz et al., 2005; Fox et al., 2005; Sporns et al., 2000). We represent the network by a graph G that is constructed as follows. The time series \mathbf{x}_i originating from voxel i becomes the node (or vertex) \mathbf{x}_i of the graph¹. Edges between vertices quantify the functional connectivity. Each node \mathbf{x}_i is connected to its n_n nearest neighbors according to the Euclidean distance between the time series, $\|\mathbf{x}_i - \mathbf{x}_j\| = (\sum_{t=1}^T (\mathbf{x}_i(t) - \mathbf{x}_j(t))^2)^{1/2}$ (Fig. 1). The weight $W_{i,j}$ on the edge $\{i, j\}$ quantifies the functional proximity between voxels i and j and is defined by

$$W_{ij} = \begin{cases} e^{-\|\mathbf{x}_i - \mathbf{x}_j\|^2 / \sigma^2}, & \text{if } \mathbf{x}_i \text{ is connected to } \mathbf{x}_j, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

¹ We slightly abuse the notation here: \mathbf{x}_i is a time series, as well as a node on the graph.

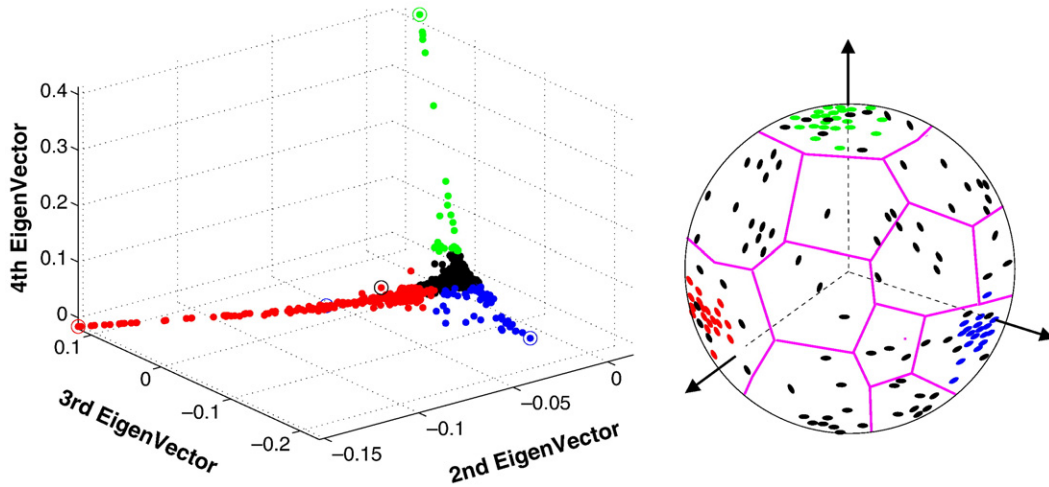


Fig. 4. Left: low-dimensional embedding of a block design dataset (see section Results). Each dot i is $\Psi(x_i)$, the image of the time series x_i through the mapping Ψ . Right: the $\Psi(x_i)$ are projected on the sphere, and the projections are clustered on the sphere.




Fig. 5. Left: synthetic brain: activation (white), non-activation (gray), outside the brain (black). Right: stimulus time series .

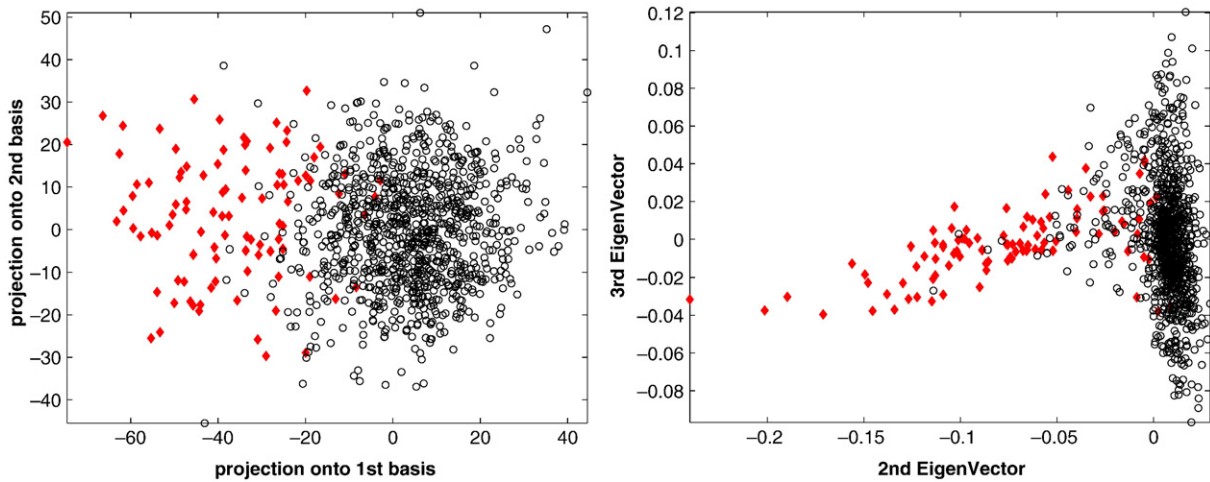


Fig. 6. Activated (red) and background (black) time series projected on the first two PCA components (left), and parametrized by Ψ (right).

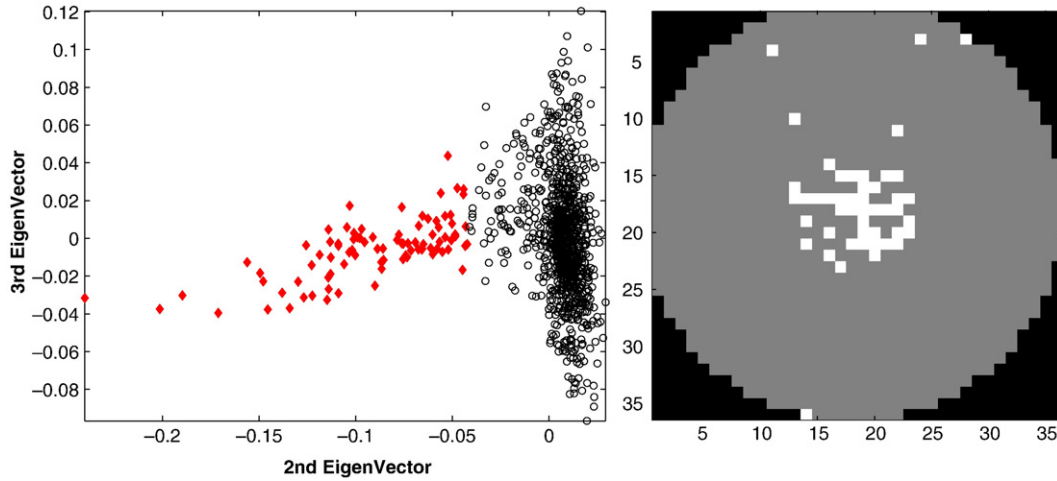


Fig. 7. Left: clustering of $\{\Psi(\mathbf{x}_i), i=1, \dots, N\}$ into 2 clusters: activated (red) and background (black). Right: activated (white) and background (gray) pixels.

The scaling factor σ modulates the definition of proximity measured by the weight (Eq. (2)). If σ is very large, then for all the nearest neighbors \mathbf{x}_j of \mathbf{x}_i , we have $W_{i,j} \approx 1$ (see Eq. (2)), and the transition probability is the same for all the neighbors, $P_{i,j} = 1/n_n$. This choice of σ promotes a very fast diffusion of the random walk through the dataset, and blurs the distinction between activated and background time series. On the other hand, if σ is extremely small, then $P_{i,j} = 0$ for all the neighbors such that $\|\mathbf{x}_i - \mathbf{x}_j\| > 0$. Only if the distance between \mathbf{x}_i and \mathbf{x}_j is zero (or very small) the transition probability is non zero. This choice of σ accentuates the difference between the time series, but is more sensitive to noise. In practice, we found that the universal choice

$$\sigma = n \times \min_{i < j} \|\mathbf{x}_i - \mathbf{x}_j\|$$

where $n \in (0, 5)$ (we used $n=2$ for all experiments) is usually optimal.

The weighted graph \mathbf{G} is fully characterized by the $N \times N$ weight matrix \mathbf{W} with entries $W_{i,j}$. Let \mathbf{D} be the diagonal degree matrix with entries $D_{i,i} = \sum_j W_{i,j}$. The spatial coordinates of the voxels i and j are currently not used in the computation of $W_{i,j}$. We know that spatial information can be useful: truly activated voxels tend to be spatially clustered. However, spatial proximity should be measured along the cortical ribbon, and not in the 3-D volume.

A new way to measure functional distances between voxels

Once the network of functionally connected voxels is created, we need to define a distance between any two vertices \mathbf{x}_i and \mathbf{x}_j in the network. This distance should reflect the topology of the graph, but should also be able to distinguish between strongly connected vertices (when the voxels i and j belong to the same functional region) and weakly connected vertices (when i and j have similar time series \mathbf{x}_i and \mathbf{x}_j , but belong to different functional regions). The Euclidean distance that is used to construct the graph is only useful locally: we can use it to compare voxels that are functionally similar, but we should use a different distance to compare voxels that may not be functionally similar. As shown in the experiments, the shortest distance $\delta(\mathbf{x}_i, \mathbf{x}_j)$ between two nodes \mathbf{x}_i and \mathbf{x}_j measured along the edges of the graph is

very sensitive to short-circuits created by the noise in the data. A standard alternative to the geodesic distance is the commute time, $\kappa(\mathbf{x}_i, \mathbf{x}_j)$, that quantifies the expected path length between \mathbf{x}_i and \mathbf{x}_j for a random walk started at \mathbf{x}_i (Bremaud, 1999).

We review here the concept of commute time in the context of a random walk on a graph. We show how the commute time can be computed easily from the eigenvectors of $\mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}$. Let us consider a random walk Z_n on the connectivity graph. The walk starts at \mathbf{x}_i , and evolves on the graph with the transition probability $\mathbf{P} = \mathbf{D}^{-1} \mathbf{W}$,

$$P_{i,j} = W_{i,j} / \left(\sum_j W_{i,j} \right). \quad (3)$$

If the walk is at \mathbf{x}_i , it jumps to one of the nearest neighbors, \mathbf{x}_j , with probability $P_{i,j}$. The walk first visits all voxels in the same functional area before jumping to a different functional area. Indeed, if voxels i and j are in the same functional area, and voxel k is in a different functional area, then we expect that $\|\mathbf{x}_i - \mathbf{x}_j\| \ll \|\mathbf{x}_i - \mathbf{x}_k\|$, and therefore $P_{i,j} \gg P_{i,k}$. This observation can be quantified by computing the average hitting time that measures the number of steps that it takes

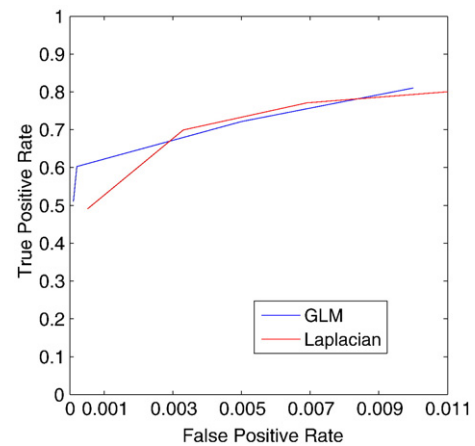


Fig. 8. ROC curves: comparison of our approach to the GLM.

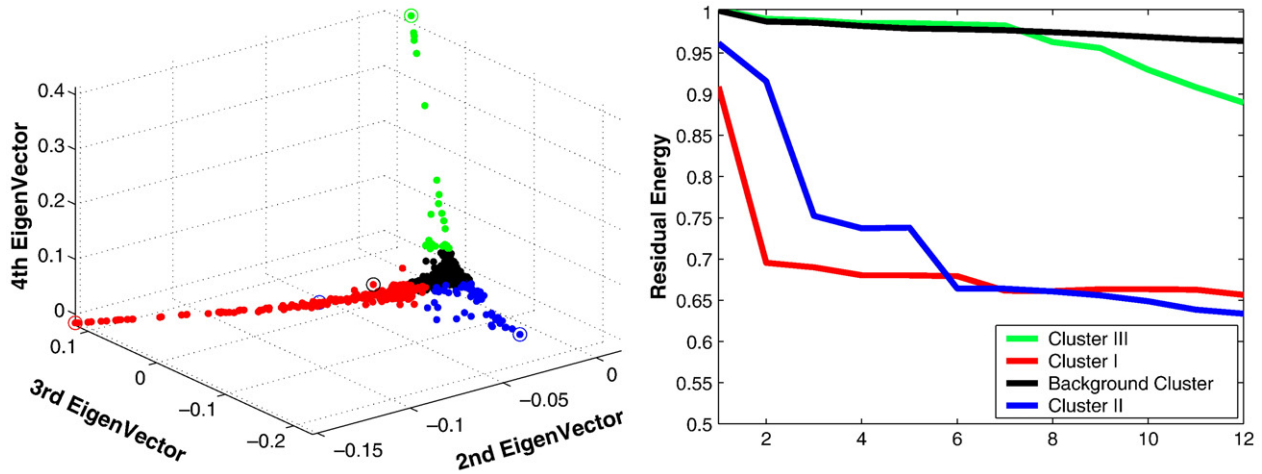


Fig. 9. Three-dimensional embedding: $\{\Psi(\mathbf{x}_i), i=1, \dots, 3084\}$: cluster I (red), cluster II (blue), cluster III (green), and background (black). Time series marked by a circle are shown in Fig. 11. Right: residual error $e_R(K)$ as a function of the number of coordinates K .

for the random walk started at x_i to hit x_j for the first time (Bremaud, 1999),

$$H(\mathbf{x}_i, \mathbf{x}_j) = E_i[T_j] \quad \text{with} \quad T_j = \min\{n \geq 0; Z_n = j\}.$$

The hitting time is not symmetric, and cannot be a distance on the graph. A proper distance is provided by a symmetric version of H , called the commute time, (Bremaud, 1999),

$$\kappa(\mathbf{x}_i, \mathbf{x}_j) = H(\mathbf{x}_i, \mathbf{x}_j) + H(\mathbf{x}_j, \mathbf{x}_i) = E_i[T_j] + E_j[T_i]. \quad (4)$$

As one would expect, $\kappa(\mathbf{x}_i, \mathbf{x}_j)$ increases with the geodesic distance $\delta(\mathbf{x}_i, \mathbf{x}_j)$. Unlike the geodesic distance, $\kappa(\mathbf{x}_i, \mathbf{x}_j)$ decreases when the number of paths between the nodes increases.

Commute time and clustering coefficient

The commute time is greatly influenced by the richness of the connections between any two nodes of the network. This concept can be quantified by the clustering coefficient. Let \mathbf{x}_i be a node with n_n neighbors. The neighbors of \mathbf{x}_i may, or may not, be neighbors of one another. To assess the transitivity of the connections, we can compute

the total number of edges, e_i , that exist between all the neighbors of \mathbf{x}_i . The clustering coefficient (Albert and Barabási, 2002) is $C_i = 2e_i / [n_n(n_n - 1)]$. The maximum value of C_i is 1 and is achieved when each neighbor of \mathbf{x}_i is connected to all the other neighbors of \mathbf{x}_i (they form a clique). If the average clustering coefficient, computed over all nodes of the network, is close to 1, then there will always be multiple routes between any two nodes \mathbf{x}_i and \mathbf{x}_j , and the commute time will remain small. Eguíluz et al. (2005) measured clustering coefficients in networks of functionally connected voxels in fMRI that were indicative of scale-free small-world networks. Achard et al. (2006) identified networks of richly connected hubs in the cortex, and have also shown that the functional network (as measured by fMRI) exhibits the “small world” property. The commute time provides a distance on the graph that takes into account the abundance of connections that may exist between two nodes of the graph.

A spectral decomposition of commute time

As explained in the Appendix, the commute time can be conveniently computed from the eigenvector ϕ_1, \dots, ϕ_N of the symmetric matrix $\mathbf{D}^{-1/2} \mathbf{P} \mathbf{D}^{-1/2}$. The corresponding eigenvalues are

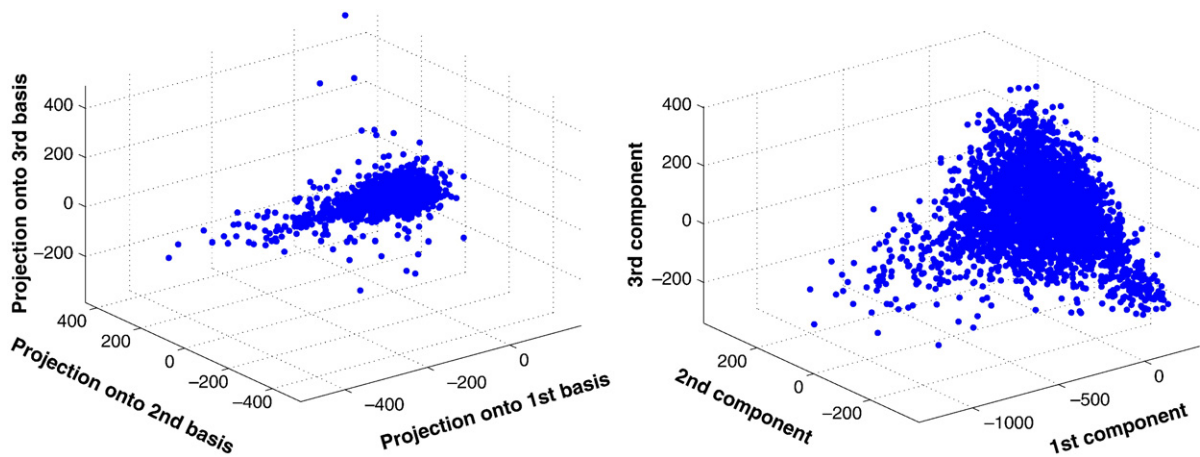


Fig. 10. Low-dimensional embedding obtained by PCA (left), and ISOMAP (right).

between -1 and 1 , and can be labeled such that $-1 \leq \lambda_N \dots \leq \lambda_2 < \lambda_1 = 1$. Each eigenvector ϕ_k is a vector with N coordinates, one for each node of the graph G . We therefore write

$$\phi_k = [\phi_k(1), \phi_k(2), \dots, \phi_k(N)]^T \quad (5)$$

to emphasize the fact that we consider ϕ_k to be a function defined on the nodes of the graph (Fig. 2). According to Eq. (16), the commute time can be expressed as

$$\kappa(\mathbf{x}_i, \mathbf{x}_j) = \sum_{k=2}^N \frac{1}{1 - \lambda_k} \left(\frac{\phi_k(i)}{\sqrt{\pi_i}} - \frac{\phi_k(j)}{\sqrt{\pi_j}} \right)^2, \quad (6)$$

where $\pi_i = d_i / \sum_{i,j} W_{i,j}$ is the stationary distribution associated with \mathbf{P} , $\pi^T \mathbf{P} = \pi^T$. The right hand side of Eq. (6) is the sum of $N-1$ squared contributions of the form $(\phi_k(i)/\sqrt{\pi_i} - \phi_k(j)/\sqrt{\pi_j})/\sqrt{1 - \lambda_k}$. Each contribution is the difference between two terms: $\phi_k(i)/\sqrt{\pi_i}$ and $\phi_k(j)/\sqrt{\pi_j}$, which are associated with nodes \mathbf{x}_i and \mathbf{x}_j , respectively.

Embedding of the dataset

We define a mapping from \mathbf{x}_i to a vector of size $N-1$,

$$\mathbf{x}_i \mapsto \frac{1}{\sqrt{\pi_i}} \left[\frac{\phi_2(i)}{\sqrt{1 - \lambda_2}}, \dots, \frac{\phi_N(i)}{\sqrt{1 - \lambda_N}} \right]^T. \quad (7)$$

The idea was first proposed in (Bérard et al., 1994) to embed manifolds. Recently, the same idea has been revisited in the machine learning literature (Belkin and Niyogi, 2003; Coifman and Lafon, 2006). According to Eq. (6), the commute time is simply the Euclidean distance measured between the new coordinates. In practice, we need not use all the eigenvectors in Eq. (7). Indeed, because $-1 \leq \lambda_N \dots \leq \lambda_2 < \lambda_1 = 1$ we have $1/\sqrt{1 - \lambda_2} > 1/\sqrt{1 - \lambda_3} > \dots > 1/\sqrt{1 - \lambda_N}$. We can therefore neglect $\phi_k(j)/\sqrt{1 - \lambda_k}$ in Eq. (7) for large k , and reduce the dimensionality of the embedding by using only the first K coordinates. Finally, we define the map Ψ from \mathbb{R}^T to \mathbb{R}^K ,

$$\mathbf{x}_i \mapsto \Psi(\mathbf{x}_i) = \frac{1}{\sqrt{\pi_i}} \left[\frac{\phi_2(i)}{\sqrt{1 - \lambda_2}}, \dots, \frac{\phi_{K+1}(i)}{\sqrt{1 - \lambda_{K+1}}} \right]^T. \quad (8)$$

The low-dimensional parametrization (Eq. (8)) provides a good approximation when the spectral gap $\lambda_1 - \lambda_2 = 1 - \lambda_2$, is large. The construction of the embedding is summarized in Fig. 3. Unlike PCA which gives a set of vectors on which to project the dataset, the embedding (Eq. (8)) directly yields the new coordinates for each time series \mathbf{x}_i . The new coordinates of \mathbf{x}_i are given by concatenating the i th coordinates of ϕ_k , $k=2, \dots, K+1$. The first eigenvector ϕ_0 is constant and is not used.

What is the connection to PCA?

Because each ϕ_k is also an eigenvector of the graph Laplacian, $\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1}\mathbf{P}\mathbf{D}$, it minimizes the “distortion” induced by ϕ and measured by the Rayleigh ratio (Belkin and Niyogi, 2003),

$$\min_{\phi, \|\phi\|=1} \frac{\sum_{i,j} W_{i,j} (\phi(i) - \phi(j))^2}{\sum_i d_i \phi^2(i)}, \quad (9)$$

where ϕ_k is orthogonal to the previous eigenvector $\{\phi_0, \phi_1, \dots, \phi_{k-1}\}$. The gradient of ϕ at the vertex \mathbf{x}_i on the graph can be

computed as a linear combination of terms of the form $(\phi(i) - \phi(j))$, where j and i are connected. Therefore the numerator of the Rayleigh ratio (Eq. (9)) is a weighted sum of the gradients of ϕ at all nodes of the graph, and quantifies the average local distortion created by the map ϕ . A function that minimizes (Eq. (9)), will still introduce some global distortions. Only an isometry will preserve all distances between the time series, and the isometry which is optimal for dimension reduction is given by PCA. However, as shown in the experiments, the first PCA components are usually not able to capture the nonlinear structure formed by the set of time




Fig. 11. Time series from the cluster I (A), cluster II (B), and cluster III (C).




Fig. 12. Eigenvectors ϕ_2 and ϕ_3 visualized as images.




Fig. 13. Top row: voxels are color-coded according to the cluster color (except for the background cluster). We interpret cluster I (red) as voxels in the visual cortex recruited by the stimulus. Bottom row: activation maps obtained using the linear regression model ($p=0.001$).




Fig. 14. Three-dimensional embedding. Top: one-trial (blue) and two-trial (red) conditions are well separated. Bottom: cluster I (blue), cluster II (red) and background (green). Time series marked by a circle are shown in Fig. 16.

series. As a result, PCA fails to reveal the organization of the dataset in terms of low-dimensional activated time series and background time series. Thirion and Faugeras (2003) have used kernel PCA to analyze the distributions of the coefficients of a model fitted to fMRI time series. A block design fMRI dataset from a macaque monkey is studied. By visual inspection, the authors show that their method can organize the coefficients according to the relative strength of the activation. The eigenvectors of the Laplacian have also been used to construct maps of spectral coherence of fMRI data in (Thirion et al., 2006).

How many new coordinates do we need?

We can estimate K , the number of coordinates in the embedding (Eq. (8)), by calculating the number of ϕ_k needed to reconstruct the low-dimensional structures present in the dataset. As opposed to PCA, the embedding defined by Eq. (8) is not designed to minimize the reconstruction error, it only minimizes the average local distortion (Eq. (9)). Nevertheless, we can take advantage of the fact that the eigenvectors $\{\phi_k\}$ constitute an orthonormal basis for the set of functions defined on the vertices of the graph (Chung, 1997). In particular, we make the trivial observation that the scan at

time t , $\mathbf{x}(t)=[x_1(t), \dots, x_N(t)]^T$, is a function defined on the nodes of the graph: $x_i(t)$ at node \mathbf{x}_i is the value of the fMRI signal at voxel i , $x_i(t)$. We can therefore expand $\mathbf{x}(t)$ using the ϕ_k ,

$$\mathbf{x}(t) = \sum_{k=1}^K \langle \mathbf{x}(t), \phi_k \rangle \phi_k + \sum_{k=K+1}^N \langle \mathbf{x}(t), \phi_k \rangle \phi_k = \hat{\mathbf{x}}^K(t) + \mathbf{r}(t),$$

where $\hat{\mathbf{x}}^K(t) = \sum_{k=1}^K \langle \mathbf{x}(t), \phi_k \rangle \phi_k$ is the approximation to the t th scan using the first K eigenvectors, and $\mathbf{r}(t)$ is the residual error. We can compute a similar approximation for all the scans ($t=1, \dots, T$), and compute the temporal average of the relative approximation error at a given voxel i ,

$$\varepsilon_i(K) = \frac{\sum_{t=1}^T (\mathbf{x}_i(t) - \hat{\mathbf{x}}_i^K(t))^2}{\sum_{t=1}^T \mathbf{x}_i^2(t)}. \tag{10}$$

Finally, one can compute the average of $\varepsilon_i(K)$ over all a group of voxels in the same functional area \mathcal{R} , $\varepsilon_{\mathcal{R}}(K) = (\sum_{i \in \mathcal{R}} \varepsilon_i(K)) / |\mathcal{R}|$. We expect $\varepsilon_{\mathcal{R}}(K)$ to decay fast with K if the time series within \mathcal{R} are

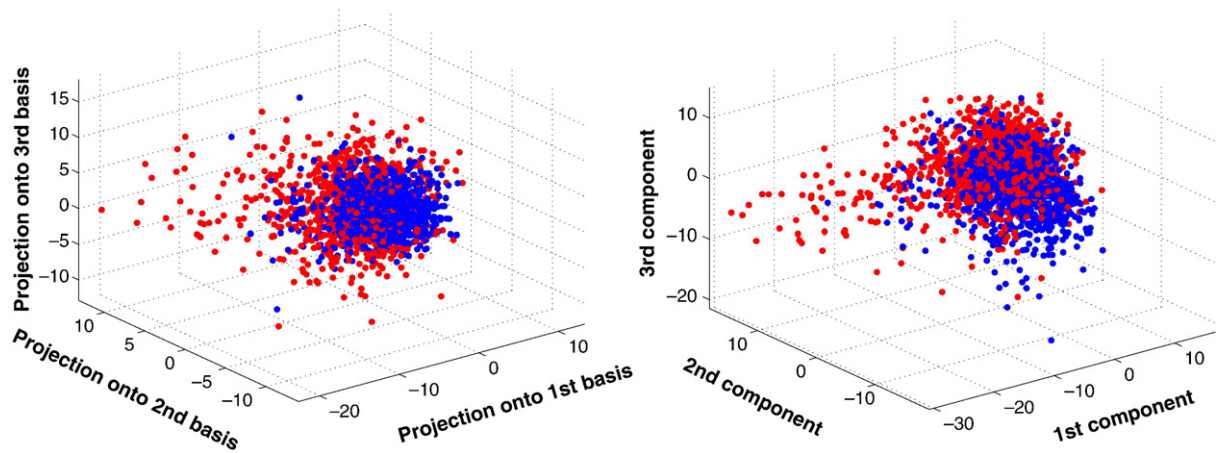


Fig. 15. Low-dimensional embedding obtained by PCA (left), and ISOMAP (right). One-trial condition (blue) and two-trial condition (red) are mingled.

well approximated by ϕ_1, \dots, ϕ_K . In practice, for each region \mathcal{R} we find $K_{\mathcal{R}}$ after which $\varepsilon_{\mathcal{R}}$ stops having a fast decay. We then choose K the number of coordinates to be the largest $K_{\mathcal{R}}$ among all regions \mathcal{R} . Examples are shown in Fig. 9-right, where K is chosen at the “knee” of the curves $\varepsilon_{\mathcal{R}}(K)$.

k-Means clustering

The map Ψ defined by (8) provides a set of new coordinates, $\Psi(\mathbf{x}_i)$, for each \mathbf{x}_i . We cluster the set $\{\Psi(\mathbf{x}_i), i=1, \dots, N\}$ into a small number of coherent structures and a large background component. We use a variation of the *k*-means clustering algorithm for this task. The low-dimensional parameterization of the dataset usually has a “star” shape (Fig. 4-left), where the out-stretching “arms” of the star are related to activated time series or strong physiological artifacts, and the center blob corresponds to background activity. Our goal is to segment each of the “arms” from the center blob. We project all the $\Psi(\mathbf{x}_i)$ on a unit sphere centered around the origin (Fig. 4-right). We then cluster the projections on the sphere: the distance between two points on the sphere is measured by their angle. The center component (shown in black in Fig. 4-left) is usually spread all over the sphere, and is mixed with the branches (Fig. 4-right). The time series from the background component can be separated from the other time series by measuring the distance of $\Psi(\mathbf{x}_i)$ to the origin (Fig. 4-left). The number of clusters can be chosen to be equal to $K+1$. This choice is based on experience that each eigenvector (each dimension) will contribute to an independent arm, and the background time series will contribute to the last cluster. This choice may over-segment the dataset. This is usually obvious from the visual inspection scatter plot, and the corresponding spatial maps. We iteratively refine the estimate of the number of clusters, by merging small clusters at each iteration.

Results

In this section we describe the results of experiments conducted on synthetic and in vivo datasets. We construct the embedding according to the algorithm described in Fig. 3, and the clustering algorithm, described in *k*-Means clustering, divides the embedded dataset into coherent groups. We interpret the coherent structures in terms of a task-related hemodynamic response, and physiological artifacts. Voxels that correspond to task-related activation are identified and activation maps are generated accordingly. We evaluate our approach using two different criteria. First, we compare the parameterization created by our approach with the parameterization produced by PCA. The comparison is based on our ability to identify and extract well defined structures from the new parameterization. Our second criterion is the comparison of the activation maps obtained with our approach with the ones generated by the General Linear Model (GLM). Five datasets are selected for the analysis: a synthetic dataset, a block design dataset, an event-related dataset, and two datasets from the EBC competition University of Pittsburgh (2007).

Synthetic data

The synthetic datasets were designed by blending activation into background, non-activated, time series that were extracted from a real in vivo dataset (described in In vivo data I: Block design dataset). We discarded time series exhibiting large variance. Activated time series were constructed by adding an activation pattern $f(t)$ to the background time series. The activation pattern $f(t)$

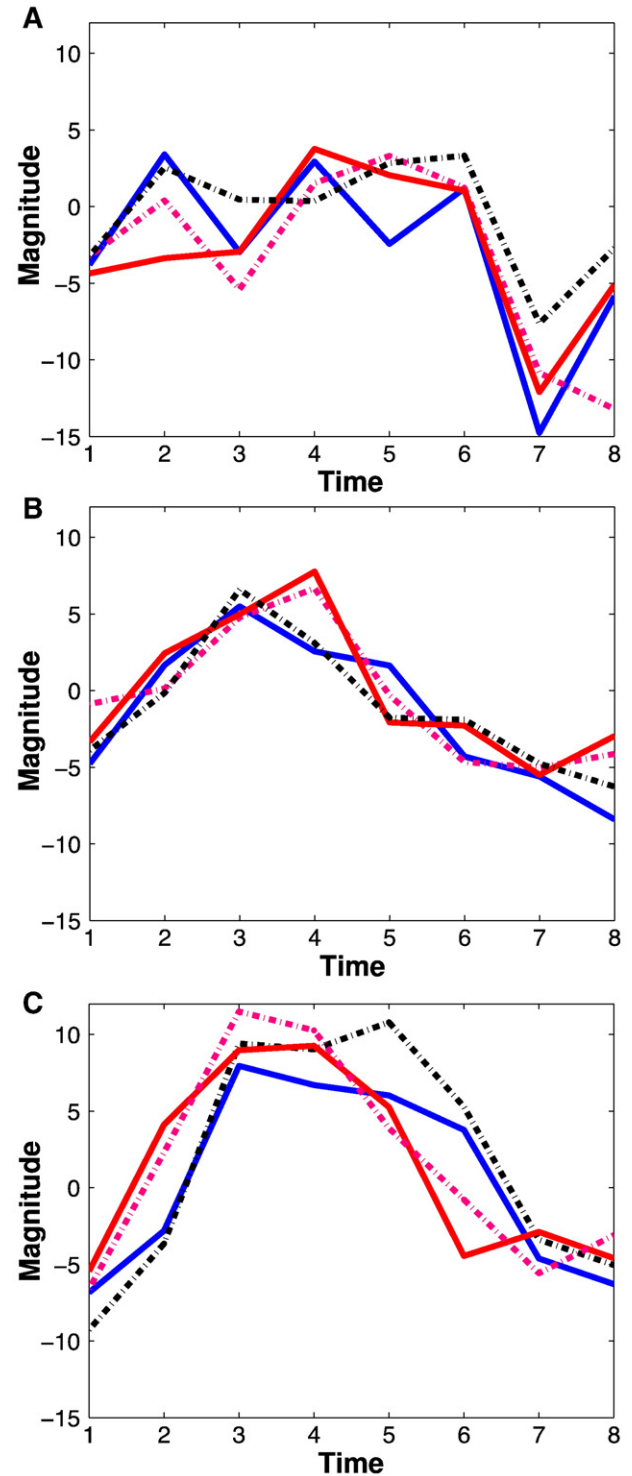


Fig. 16. Four representative time series from cluster I (A), cluster II with one-trial condition (B), and from cluster II with two-trial condition (C).

was obtained by convolving the canonical hemodynamic filter $h(t)$ used in SPM (Friston, 2005)

$$h(t) = \alpha \left(\frac{t}{d_1} \right)^{\alpha_1} e^{-(t-d_1)/b_1} - c \left(\frac{t}{d_2} \right)^{\alpha_2} e^{-(t-d_2)/b_2}, \quad (11)$$

with the stimulus time-series $g(t)$ (Fig. 5-left). The two parameters α and b_1 were randomly distributed according to two uniform distributions, on $[0.8, 1.2]$ and on $[5, 10]$ respectively. The other parameters were fixed and chosen as follows, $a_1=6$, $a_2=12$, $b_2=0.9$, and $c=0.35$. By varying b_1 and α independently, we generated a family of hemodynamic responses with different peak dispersions and scales. We generated 20 independent realizations according to this design. Each dataset consisted of a white disk of activated voxels inside a circular gray brain of background voxels (Fig. 5-left). Black voxels were in the air. There were altogether 1067 voxels inside the circular brain, with 97 activated voxels, (9% activation). Fig. 6-left shows the projections on the first two principal axes of the 1067 time series of one realization. The projections are color-coded according to their status: activated (red diamond) and background (black circle). The parameterization given by our approach is shown in 6-right. We used only $K=2$ coordinates in (8) for Ψ . Activated time series are distributed along a thin strip that extends outward from the main cluster. This low-dimensional structure is compact and easy to identify. In comparison, the two dimensional representation given by PCA (Fig. 6-left) is less conspicuous: activated time series (red) overlap with background time series (black). After embedding the dataset into two dimensions, the dataset is partitioned into two clusters. Fig. 7-left shows the result of clustering: the labels (red for activated, black for background) are based on the clustering only. The corresponding activation map is shown in Fig. 7-right. We compared our algorithm with a linear model equipped with the perfect knowledge of the hemodynamic response $h(t)$ (with $b_1=1$ and $\alpha=1$). A Student t-test was applied to the regression coefficient to test its significance, and voxels with a p-value smaller than a threshold were considered activated. Fig. 8 provides a quantitative comparison of our approach with the linear model using a receiver operator characteristic (ROC) curve. The true activation rate (one minus the type II error) is plotted against the type I error (false alarm rate). The ROC curve was computed over 20 trials. Each trial included different activation strengths α . In this experiment, the linear model has access to an oracle, in the form of the perfect knowledge of the hemodynamic response $h(t)$, and should therefore perform very well. In fact, if the

noise added to $f(t)$ were to be white, we know from the matched filter theorem that the linear model would provide the optimal detector. Here, the noise is extracted from the data, and is probably not white (Bullmore et al., 2001). As shown in Fig. 8, our approach performs as well as the GLM for a type I error in the range $[0.003, 0.009]$. At low type I error, our approach misses activations.

In vivo data I: Block design dataset

We apply our technique to a block design dataset that demonstrates activation of the visual cortex (Tanabe et al., 2002). A flashing checkerboard image was presented to a subject for 30 s, and a blank image was presented for the next 30 s. This alternating cycle was repeated four times. Images were acquired with a 1.5T Siemens MAGNETOM Vision equipped with a standard quadrature head coil and an echoplanar subsystem (TR=3 s, xy dimension: 128×128 , voxel size= $1.88 \times 1.88 \times 3$ mm, 12 contiguous slices). Eighty images were obtained. We analyze a volume that contains the calcarin cortex (Brodmann areas 17) (Fig. 12). There are altogether 3084 intracranial time series in the volume. The linear trend of each time series is removed. Fig. 9-left displays the image of the time series by the embedding, $\Psi(\mathbf{x}_i)$, $i=1, \dots, 3084$. The time series are color-coded according to the result of the clustering. The dataset is partitioned into four clusters. For comparison purposes, we show the embedding generated by ISOMAP (Tenenbaum et al., 2000) (Fig. 10-left), and the embedding obtained by projecting the time series on the first three PCA axes (Fig. 10-right). ISOMAP makes it possible to reconstruct a low-dimensional nonlinear structure embedded in high dimension. The algorithm computes the pairwise geodesic distance between any two points in the dataset, and uses multidimensional scaling to embed the dataset in low dimension. As explained in A new way to measure functional distance between voxels, the geodesic distance is not appropriate for fMRI data that are very noisy. For this reason an embedding based on commute time is more robust. The representation given by PCA and ISOMAP, shown in Fig. 10 left and right, are less conspicuous than the representation obtained by our approach. No low-dimensional structures are apparent in these two plots. The plot of




Fig. 17. Top row: activation maps: cluster I (blue), cluster II (red). Bottom row: activation maps obtained using the linear regression model ($p=0.005$).




Fig. 18. Subject 14: views of the embedded dataset using ϕ_3, ϕ_4, ϕ_5 (left), ϕ_4, ϕ_5, ϕ_6 (center), and ϕ_5, ϕ_6, ϕ_3 (right). The labels are obtained after clustering the datasets into six clusters. The background activity (maroon cluster) is centered around 0.

the residual error (Fig. 9-right) indicates that eigenvectors ϕ_2 and ϕ_3 create the largest drop in the residual energy, and should be the most useful. We use $K=3$ coordinates: ϕ_2 , ϕ_2 and ϕ_4 to embed the dataset. In order to interpret the role of the clusters, we select several time series from each cluster (identified by circle in the scatter plot in Fig. 9) and plot them in Fig. 11. The time series from the red cluster (Fig. 11-A) are typical hemodynamic responses to a periodic stimulus. The time series at the tip of the red cluster, marked by a red circle, is the red curve in Fig. 11-A. It is the strongest activation pattern. Time series in the middle of the red cluster (blue and black circles) exhibit weaker activation (blue and black curves). We interpret cluster I as voxels activated by the stimulus. The embedding has organized the time series according to the strength of the activation: strong activation at the tip and weak activation at the base of the branch (close to the background activity). The time series from the blue cluster have a high frequency (Fig. 11-B), and are grouped together inside the brain (Fig. 12). These time series could be related to non-task-related physiological responses, such as a pulsating vein. Finally, the time series from the green cluster are less structured (Fig. 11), and are scattered across the region of analysis (Fig. 12). We were not able to interpret the physiological role of these time series.

So what do the eigenvectors look like?

As explained in A new way to measure functional distances between voxels, we consider the eigenvectors ϕ_k to be functions of the nodes i of the graph. We can therefore represent ϕ_k as an image: each voxel x_i is color-coded according to the value of $\phi_k(i)$. The majority of the values of ϕ_2 are positive (deep red, in Fig. 13-left). A few voxels take negative values (yellow and cyan in Fig. 13-left). The nodal lines (where ϕ_2 changes sign) are localized around the area of activation in the visual cortex. We can check in Fig. 9 that the activated time series (red cluster) have a negative ϕ_2 coordinate. In fact, ϕ_2 is known as the Fiedler vector (Chung, 1997) and is used to optimally split a dataset into two parts. (Fig. 13). We computed the activation map obtained using a GLM. The regressor is computed by convolving the haemodynamic response defined by Eq. (11) with the experimental paradigm. The activation map (thresholded at $p=0.001$) is shown in Fig. 12, and is consistent with the activation maps obtained by our approach.

In vivo data II: Event-related dataset

We apply our method to an event-related dataset. Buckner et al. (2000) used fMRI to study age-related changes in functional anatomy. The subjects were instructed to press a key with their right index finger upon the visual stimulus onset. The stimulus lasted for 1.5 s. Functional images were collected using a Siemens 1.5-T Vision System with an asymmetric spin-echo sequence sensitive to BOLD contrast (volume TR=2.68 s, xy dimension: 64×64 , voxel size: 3.75×3.75 mm, 16 contiguous axial slices). Each run consists of 128 TRs. For every 8 images, the subjects were presented with one of the two conditions: (i) the one-trial condition where a single stimulus was presented to the subject, and (ii) the two-trial condition where two consecutive stimuli were presented. The inter-stimulus interval of 5.36 s. was sufficiently large to guarantee that the overall response would be about twice as large as the response to the one-trial condition. We analyzed one run. After discarding the first and last four scans, the run included 15 trials (8 one-trial and 7 two-trial conditions) of 8 temporal samples. Time series from the one-trial and two-trial conditions were averaged separately. Therefore, each voxel gave rise to two

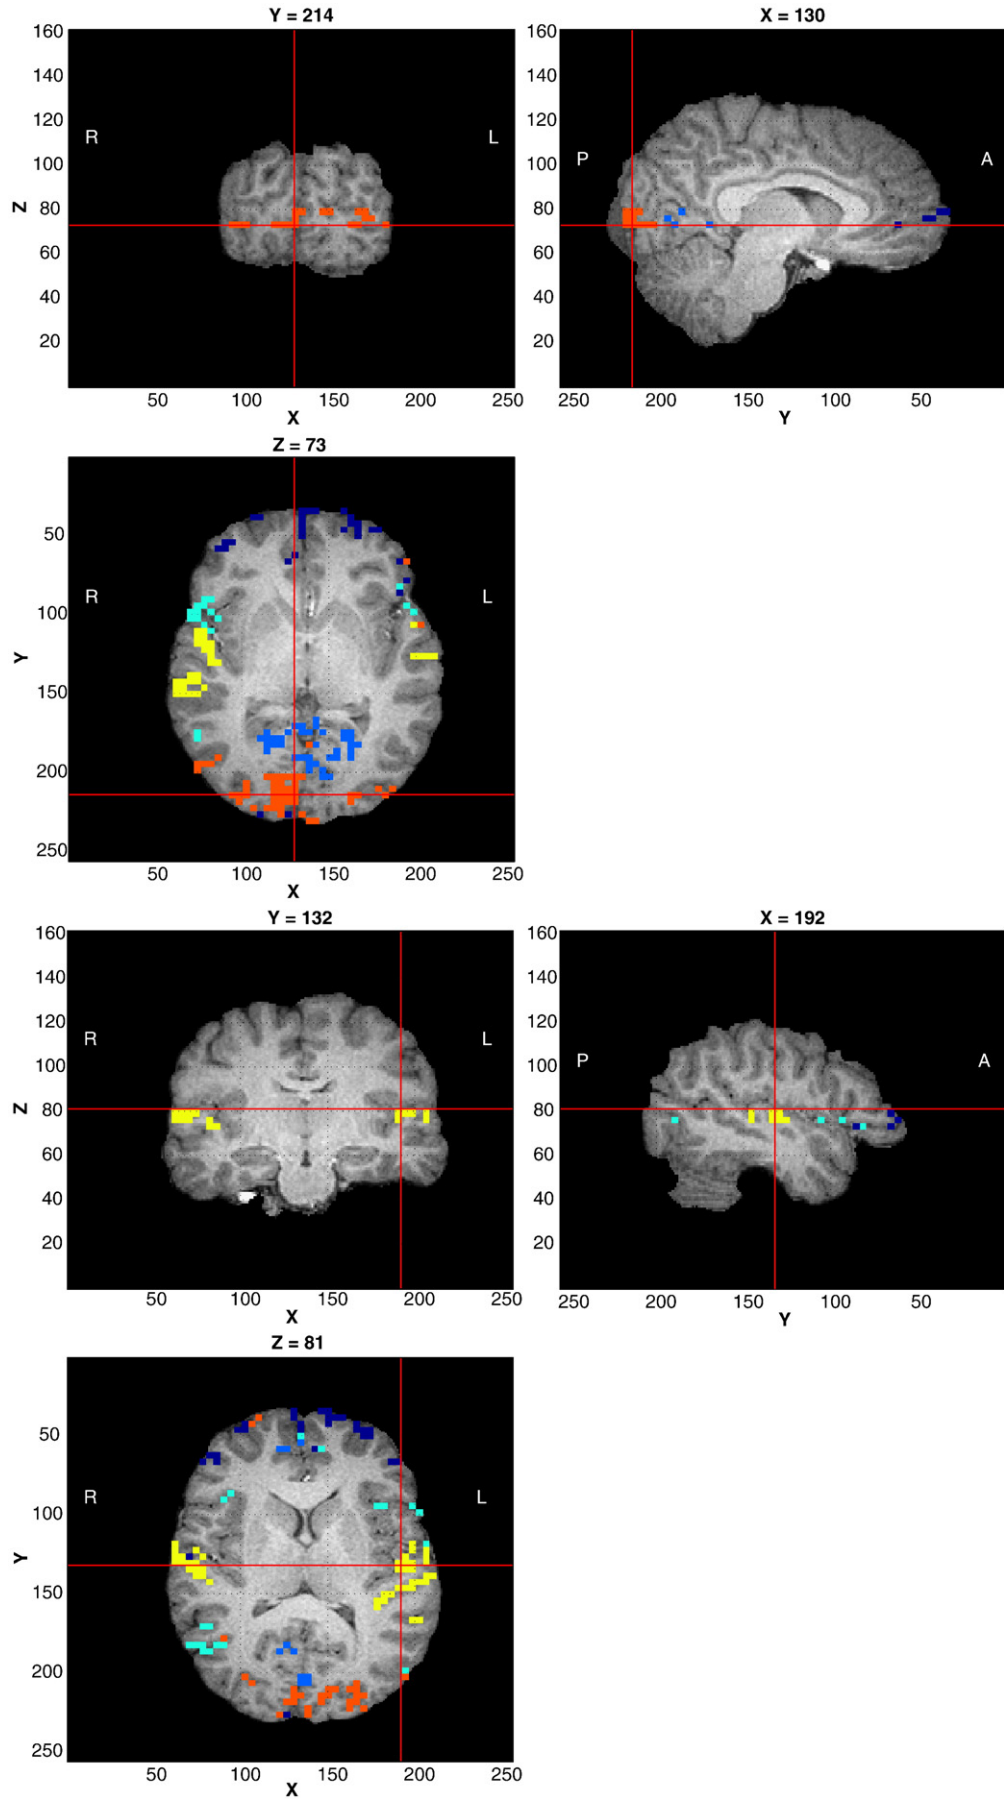
average time series of 8 samples. The linear trend was removed from all average time series. The results published in (Buckner et al., 2000) show activation in the visual cortex, motor cortex, and cerebellum. We focus our analysis on four contiguous axial slices (7, 8, 9 and 10) that extend from the superior caudate nucleus to the midlevel diencephalon. We extract a region (1025 intra-cranial voxels) that extends from the occipital posterior horn of the lateral ventricles until the end of the occipital lobe. The total number of time series included in our analysis is 2050: two time series (one-trial and two-trial conditions) of $T=8$ samples for each voxel.

Embedding the dataset in 3 dimensions

We display in Fig. 14-left the 2050 time series after embedding the dataset in three dimension. The one-trial condition time series (blue) form a blob at the center. The two-trial (red) conditions time series form a “V” (Fig. 14-right). Both branches of the “V” are nearly one dimensional and are aligned with ϕ_2 and ϕ_4 . The branch aligned with ϕ_2 also includes one-trial condition time series at the base of the branch. The branch aligned with ϕ_4 contains only two-trial condition time series. The dataset was partitioned into three clusters (Fig. 14-right). We compare our embedding to the embedding generated by ISOMAP (Fig. 15-left) and PCA (Fig. 15-right). No low-dimensional structures emerge from the representations given by PCA or ISOMAP, and two-trial and one-trial time series are mixed together. The eigenvectors ϕ_2 and ϕ_4 create the largest drop in the residual energy, and are the most useful coordinates. We use $K=3$ coordinates: ϕ_2 , ϕ_3 and ϕ_4 to embed the dataset. To determine the role of the clusters, we select three groups of four time series (identified by circles in Fig. 14-right) and we plot them in Fig. 16. Time series from cluster I all have an abrupt dip at $t=7$. The corresponding voxels are located along the border of the brain (Fig. 17). The original time series (before averaging) suffer from a sudden drop at time 95, which could be caused by a motion artifact, that affects the average time series. There are two groups of time series selected from cluster II: two two-trial condition time series located at the tip of the branch, and two one-trial condition time series located at the border with the background cluster. Time series from cluster II have a shape similar to an hemodynamic response (Fig. 16-B and C), and the corresponding voxels are located in the visual cortex (Fig. 17). Therefore, we hypothesize that cluster II contains time series recruited by the stimulus. Interestingly, the embedding has organized the time series along the branch of cluster II according to the strength of the activation: from two-trial condition (strong response) at the tip, to one-trial condition (weak response) at the stem (close to the background time series). This is a remarkable result since no information about the stimulus, or the type of trial was provided to the algorithm. Fig. 17 shows the location of the voxels corresponding to the time series of cluster I (blue) and II (red). For comparison purposes we computed the activation map obtained using the GLM. The averaged time series from the two-trial condition are used for the regression analysis. We use the hemodynamic response function defined by (Dale and Buckner, 1997), $h(t) = ((t-\delta)/\tau)^2 e^{-(t-\delta)/\tau}$, where $\delta=2.5$, $\tau=1.5$. The regressor was given by $g(t) = h(t) * s(t)$, where $s(t)$ is the stimulus time series. We thresholded the p -value at $p=0.005$, and the activation maps are shown in Fig. 17. The activation maps constructed by our approach are consistent with the maps obtained with a GLM.

Complex natural stimuli

We demonstrate here that our method can be used to analyze fMRI datasets collected in natural environments where the subjects



are bombarded with a multitude of uncontrolled stimuli that cannot be quantified (Golland et al., 2007; Hason et al., 2004; Haynes and Rees, 2006; Malinen et al., 2007; Meyer and Stephens, 2008). During such experiments, the subjects are submitted to an abundance of concurrent sensory stimuli, which makes the analysis with inferential methods impossible. Exploratory techniques can help unravel the different neural processes involved during the experiments (Malinen et al., 2007). Unlike ICA, our approach does not posit the existence of a mixture model. Our approach merely explores the connectivity in the dataset, and proposes a new parameterization that preserves connectivity.

The Experience Based Cognition competition (EBC) (University of Pittsburgh, 2007) offers an opportunity to study complex responses to natural environments. The EBC datasets comprise three 20-min runs (704 TRs in each run) of subjects interacting in an urban virtual reality environment. Subjects were audibly instructed to complete three search tasks in the environment: looking for weapons (but not tools) taking pictures of people with piercing (but not others), or picking up fruits (but not vegetables). The data was collected with a 3T EPI scanner (TR=1.75 s, xy dimension: 64×64 , voxel size= 3.28×3.28 mm, 34 slices with a thickness of 3.5 mm). We analyze the second runs of subjects 14 and 13. For each subject, the matrix X is composed of $N=4843$ intra-cranial voxels at $T=704$ TRs. We first remove the non-regionally specific variance captured by the first eigenmodes of a singular value decomposition of the dataset. We then compute ϕ_k , $k=2, \dots, 10$ using $n_n=100$ and $\sigma=2 \times \min_{i < j} \|\mathbf{x}_i - \mathbf{x}_j\|$. After embedding the dataset into four dimensions, we cluster the voxels. Figs. 18 and 20 display the datasets after embedding. Because we cannot display four dimensions, we show the projections of the dataset on three consecutive coordinates. All the coordinates contribute to the spread the dataset along elongated arms, which facilitates the clustering. Voxels that do not correspond to the background activity (the maroon cluster in Figs. 18 and 20) are superimposed on anatomically registered structural images and colored according to their cluster label (see Figs. 19 and 21). For both subjects, the clusters are connected regions (see Figs. 19 and 21), compactly organized around functional areas related to the processing of visual, and auditory stimuli (music, cellphone ringing, dog roaring) in the virtual environment. It is important to emphasize that our method never enforces any form of spatial proximity, and is purely based on functional connectivity.

For subject 14 (Fig. 19), the orange cluster corresponds to activation in the calcarine cortex associated with V1/V2 representations of the lower visual fields, while the light blue cluster corresponds to representations of the upper visual fields. Activation in lateral areas (visual motion areas, MT/V5) is also present, as well as activity in the posterior convexial cortex (area VP). The activation is predominantly in the right hemisphere. Interestingly, the two clusters located in the visual cortex (light blue and orange) have very similar ϕ_3 and ϕ_4 coordinates (see Fig. 18-left). The cyan cluster corresponds to activation in the right frontal gyrus (Broca area) associated with language comprehension. The yellow clusters are located in the right and left superior temporal gyri and medial temporal gyri (Wernicke area). These regions correlate with activation in the auditory cortex and language areas. Finally, the dark blue cluster corresponds to activation in the prefrontal cortex. A very similar pattern of activity (Fig. 21) was obtained for subject 13. The

blue and orange clusters, located in the calcarine cortex, correspond to V1 and V2 areas. Again, these two clusters, both located in the visual cortex, have similar ϕ_2 and ϕ_6 coordinates. The green cluster is located in the medial temporal gyrus (Wernicke area) and is associated with language processing. We replaced the embedding constructed by our method with the parametrization produced by PCA, using the same pre-processing steps. PCA was unable to produce any meaningful activation maps (results not shown). In fact, the clustering would usually not converge. Interestingly, the activation maps obtained with these natural stimuli are very similar to the extrinsic network, found by Golland et al. (2007), that is composed of areas dedicated to the processing of sensory information: auditory, visual, somatosensory and language.

Discussion

We proposed a new method to compute a low-dimensional embedding of an fMRI dataset. The embedding preserves the local functional connectivity between voxels, and can be used to cluster the time series into coherent groups. Our approach, based on a spectral decomposition of commute time, appears to be more robust than a method based on the computation of the geodesic distance between time series (Tenenbaum et al., 2000). Our approach is able to detect independently visual areas (V1/V2, V5/MT), auditory and language areas that are recruited when subjects interacted in an urban virtual reality environment. We believe that this approach offers a new approach for the analysis of natural stimuli. The method still requires the visual inspection of the spatial patterns of activations. This is a standard limitation of exploratory methods. We are currently exploring methods to combine our approach with standard functional atlases.

Choice of the parameters

There are only two parameters that determine the embedding: n_n the number of nearest neighbors and the scaling factor σ . We described in The connectivity graph: A network of functionally correlated voxels a universal choice for σ . The number of nearest neighbors can also be chosen according to a universal strategy, n_n can be assigned to the largest power of ten smaller than the number of time samples (TRs) in the time series. For instance, we used $n_n=7$ and $n_n=9$ for the event-related and block design datasets, respectively. The analysis of the EBC datasets that contain 704 TRs each required a much larger number of nearest neighbors: $n_n=100$.

Computational complexity

The complexity of our method is determined by the combined complexity of the nearest neighbor search and the eigenvalue problem. We use a brute force approach to search for nearest neighbors. We use the restarted Arnoldi method for sparse matrices implemented by the Matlab function `eigs` to solve the eigenvalue problem. The algorithm requires $\mathcal{O}(K) + \mathcal{O}(K^2)$ storage, and $\mathcal{O}(NK^2) + \mathcal{O}(K^3)$ computations. Our MATLAB code will be made available shortly from our webpage.

Fig. 19. Subject 14. Top: V1/V2 (orange); representations of the upper visual fields (light blue); Broca area (cyan). Bottom: Wernicke area (yellow); prefrontal cortex (dark blue).


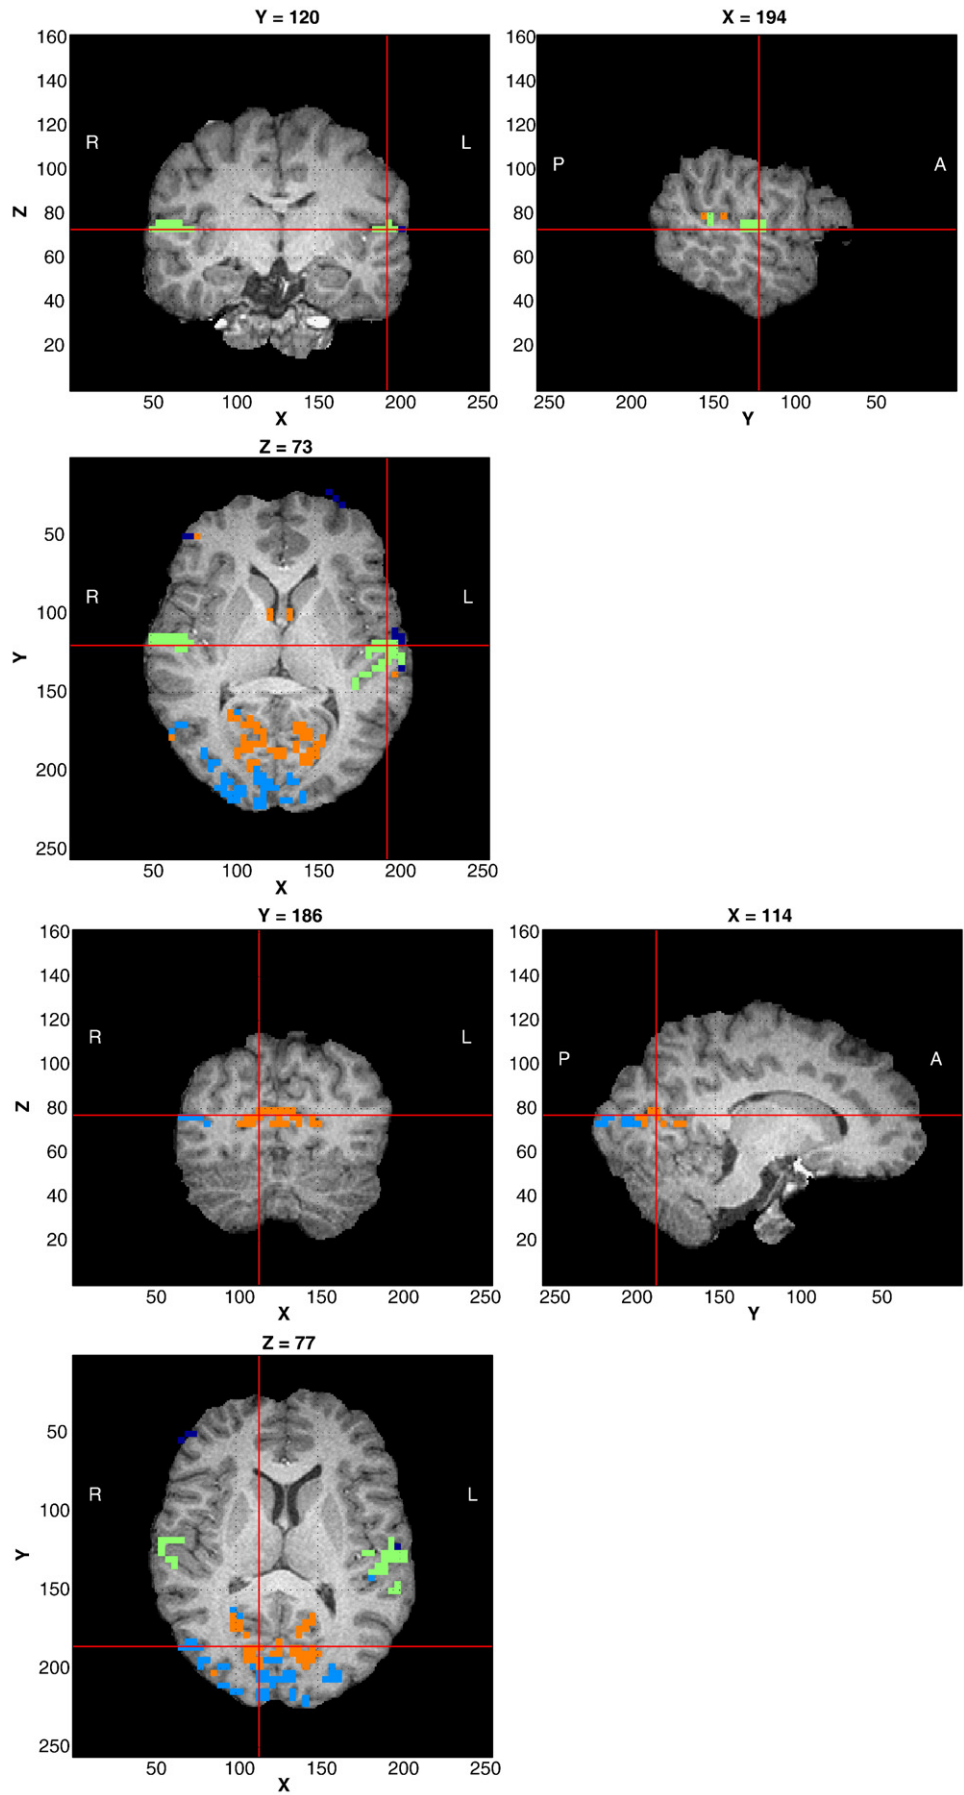


Fig. 20. Subject 13: views of the embedded dataset using ϕ_2, ϕ_3, ϕ_6 (left), ϕ_3, ϕ_6, ϕ_8 (center), and ϕ_6, ϕ_8, ϕ_2 (right). The labels are obtained after clustering the datasets into five clusters. The background activity (maroon cluster) is centered around 0.



Acknowledgments

The authors are very grateful to the reviewers for their comments and suggestions. This work was initiated at the Institute for Pure and Applied Mathematics, UCLA which provided partial support for the authors. The authors are very grateful to Raphy Coifman, Stéphane Lafon, and Mauro Maggioni for introducing them to diffusion maps, and to Ivo Dinov for making the event-related fMRI dataset available for this study. FGM was on sabbatical at the Center for the Study of Brain, Mind and Behavior, Princeton University, when he performed the analysis of the EBC datasets, and he is very grateful to all the members of the center for their support and insightful discussions.

Appendix

The hitting time satisfies the following one-step equation,

$$E_i [T_j] = 1 + \sum_{k:k \neq j} P_{i,k} E_i [T_k]. \quad (12)$$

Iterating this equation yields an expression of $E_i [T_j]$ in terms of powers of P . Let us define the fundamental matrix (Bremaud, 1999),

$$\mathbf{Z} = \mathbf{I} + \sum_{k \geq 1} \mathbf{P}^k - \mathbf{\Pi}. \quad (13)$$

then we have,

$$E_i [T_j] = (Z_{j,j} - Z_{i,j}) / \pi_j. \quad (14)$$

The proof is a straightforward consequence of (12), and can be found in (Bremaud, 1999). Note that $\mathbf{Z} = (\mathbf{I} - (\mathbf{P} - \mathbf{\Pi}))^{-1}$ is also the Green function of the Laplacian, which explains the connection with the graph Laplacian. We can consider the eigenvectors ϕ_1, \dots, ϕ_T of the symmetric matrix

$$\mathbf{D}^{\frac{1}{2}} \mathbf{P} \mathbf{D}^{-\frac{1}{2}}, \quad (15)$$

and write \mathbf{P} in (14) in terms of the eigenvectors. We obtain

$$\kappa(i,j) = \sum_{k=2}^T \frac{1}{1 - \lambda_k} \left(\frac{\phi_k(i)}{\sqrt{\pi_i}} - \frac{\phi_k(j)}{\sqrt{\pi_j}} \right)^2. \quad (16)$$

Finally, we note that $\mathbf{D}^{1/2} \mathbf{P} \mathbf{D}^{-1/2} = \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}$.

References

- Achard, S., Salvador, R., Whitcher, B., Suckling, J., Bullmore, E., 2006. A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. *J. Neurosci.* 26 (1), 63–72.
- Albert, R., Barabási, A., 2002. Statistical mechanics of complex networks. *Rev. mod. phys.* 74, 47–97.
- Belkin, M., Niyogi, P., 2003. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computations* 15, 1373–1396.
- Bérard, P., Besson, G., Gallot, S., 1994. Embeddings Riemannian manifolds by their heat kernel. *Geom. funct. anal.* 4 (4), 373–398.
- Biswal, B.B., Ulmer, J., 1999. Blind source separation of multiple signal sources of fMRI data sets using independent component analysis. *J. comput. assist. tomogr.* 23 (3), 265–271.
- Bremaud, P., 1999. *Markov Chains*. Springer Verlag.
- Buckner, R., Snyder, A., Sanders, A., Raichle, M., Morris, J., 2000. Functional brain imaging of young, nondemented, and demented older adults. *J. cogn. neurosci.* 12, 24–34.
- Bullmore, E., Long, C., Suckling, J., Fadili, J., Calvert, G., Zelaya, F., et al., 2001. Colored noise and computational inference in neurophysiological (fMRI) time series analysis: resampling methods in time and wavelet domain. *Hum. brain mapp.* 78, 61–78.
- Caclin, A., Fonlupt, P., 2006. Effect of initial fMRI data modeling on the connectivity reported between brain areas. *NeuroImage* 33, 515–521.
- Chung, F., 1997. *Spectral Graph Theory*. CBNS-AMS.
- Coifman, R., Lafon, S., 2006. Diffusion maps. *Appl. comput. harmon. anal.* 21, 5–30.
- Dale, A.M., Buckner, R.L., 1997. Selective averaging of rapidly presented individual trials using fMRI. *Hum. Brain Mapp.* 5, 329–340.
- Eguiluz, V., Chialvo, D., Cecchi, G., Baliki, M., Apkarian, A., 2005. Scale-free brain functional networks. *Phys. rev. lett.* 94 (018102).
- Fox, M., Snyder, A., Vincent, J., Corbetta, M., Van Essen, D., Raichle, M., 2005. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc. Natl. Acad. Sci.* 102 (27), 9673–9678.
- Friston, K.J., 1998. Modes or models: a critique on independent component analysis for fMRI. *Trends cogn. sci.* 2 (10), 373–375.
- Friston, K., 2005. Models of brain function in neuroimaging. *Ann. Rev. Psychol.* 56, 57–87.
- Golland, Y., Bentin, S., Gelbard, H., Benjamini, Y., Heller, R., Nir, Y., Hasson, U., Malach, R., 2007. Extrinsic and intrinsic systems in the posterior cortex of the human brain revealed during natural sensory stimulation. *Cereb. Cortex* 17, 766–777.
- Hason, U., Nir, Y., Levy, I., Fuhrmann, G., Malach, R., 2004. Intersubject synchronization of cortical activity during natural vision. *Science* 303, 1634–1640.
- Haynes, J.-D., Rees, G., 2006. Decoding mental states from brain activity in humans. *Nat. Neurosci.* 7.
- Malinen, S., Hlushchuk, Y., Hari, R., 2007. Towards natural stimulation in fMRI—issues of data analysis. *NeuroImage* 35, 131–139.
- McKeown, M., Hansen, L., Sejnowski, T., 2003. Independent component analysis of functional MRI: what is signal and what is noise? *Curr. opin. neurobiol.* 13, 620–629.
- Meyer, F., Stephens, G., 2008. Locality and low-dimensions in the prediction of natural experience from fmri. In: Platt, J., Koller, D., Singer, Y., Roweis, S. (Eds.), *Advances in Neural Information Processing Systems 20*. MIT Press, Cambridge, MA.
- Petersson, K.M., Nichols, T., Poline, J.-B., Holmes, A., 1999. Statistical limitations in functional neuroimaging II. Signal detection and statistical inference. *Phil. Trans. R. Soc. Lond. B* (354), 1261–1281.
- Raichle, M., Mintun, M., 2006. Brain work and brain imaging. *Ann. rev. neurosci.* 29, 449–476.
- Sporns, O., Taroni, G., Edelman, G., 2000. Theoretical neuroanatomy: relating anatomical and functional connectivity in graphs and cortical connection matrices. *Cereb. Cortex* 10, 127–141.
- Tanabe, J., Miller, D., Tregellas, J., Freedman, R., Meyer, F., 2002. Comparison of detrending methods for optimal fMRI pre-processing. *NeuroImage* 15, 902–907.
- Tenenbaum, J., de Silva, V., Langford, J., 2000. A global geometric framework for nonlinear dimensionality reduction. *Science* 290, 2319–2322.
- Thirion, B., Fugeras, O., 2003. Dynamical components analysis of fMRI data through kernel PCA. *Neuroimage* 20, 34–49.
- Thirion, B., Dodel, S., Poline, J.-B., 2006. Detection of signal synchronizations in resting-state fMRI datasets. *Neuroimage* 2, 321–327.
- University of Pittsburgh, 2007. The experience based cognition project, <http://www.etc.pitt.edu>.

Fig. 21. Subject 13. Top: Wernicke area (green). Bottom: V1/V2 (orange); representations of the upper visual fields (light blue); lateral areas (light blue): visual motion areas (MT/V5).