

Technical Note

Projection of fMRI data onto the cortical surface using anatomically-informed convolution kernels

G. Operto,^{a,*} R. Bulot,^a J.-L. Anton,^b and O. Coulon^a

^aLaboratoire LSIS, UMR CNRS 6168, Marseille, France

^bCentre d'IRM fonctionnelle de Marseille, Marseille, France

Received 22 December 2006; revised 15 August 2007; accepted 20 August 2007

Available online 31 August 2007

As surface-based data analysis offer an attractive approach for intersubject matching and comparison, the projection of voxel-based 3D volumes onto the cortical surface is an essential problem. We present here a method that aims at producing representations of functional brain data on the cortical surface from functional MRI volumes. Such representations are for instance required for subsequent cortical-based functional analysis. We propose a projection technique based on the definition, around each node of the gray/white matter interface mesh, of convolution kernels whose shape and distribution rely on the geometry of the local anatomy. For one anatomy, a set of convolution kernels is computed that can be used to project any functional data registered with this anatomy. Therefore resulting in anatomically-informed projections of data onto the cortical surface, this kernel-based approach offers better sensitivity, specificity than other classical methods and robustness to misregistration errors. Influences of mesh and volumes spatial resolutions were also estimated for various projection techniques, using simulated functional maps.

© 2007 Elsevier Inc. All rights reserved.

Keywords: fMRI; Cortical surface; Surface-based analysis; Anatomical constraints

Introduction

Functional magnetic resonance images (fMRI) are generally presented and analysed in their original 3D space of acquisition with little regard to the structural features of the human brain. In this context, surface-based analyses of cortical data have gained interest as they advocate for anatomical considerations of the object under study. First, as the main part of the brain activity is generated by cortical pyramidal neurons, descriptions of the cortex should respect its sheet-like structure, despite the intrinsic voxel-based nature of magnetic resonance images and the highly convoluted nature of the cortex. Moreover, distances computed

in 3D voxel-based space do not take folds into account, and two points with low Euclidean distance can actually be located on the two opposite sides of a sulcus; with respect to this issue, geodesic distances, along the cortical surface, describe the anatomy better. Several papers in the literature (Clouchoux et al., 2005; Van Essen and Drury, 1997; Fischl et al., 1999; Toro and Burnod, 2003) advocate for this approach. As these works deal with visualisation, intersubjects registration or cortical localisation, little research has also been published on surface-based functional data analysis (Andrade et al., 2001; Flandin et al., 2002; Goebel and Singer, 1999; Kiebel et al., 2000). On this matter, one essential problem concerns the projection of functional data onto the anatomy, in other terms the mapping of functional voxel-based volumes onto cortical triangulated meshes. This is especially a crucial point to open the way to fMRI data analysis confined to the cortical ribbon.

Still, this problem has remained poorly studied and is rarely found in the literature. Some «geometric» simple methods do exist: some methods propose to average intensities along a normal direction or inside a sphere centered at each node of the mesh representing the cortical surface (e.g. when using free package BrainVISA), or to assign each node with the value of its containing voxel (Saad et al., 2004), or to compute trilinear interpolations (Andrade et al., 2001). Other methods try to embed some explicit anatomical information into the process: it can be done by defining each node's influence scope through voxels distances to and along the cortical mesh (Warnking et al., 2002), or computing a geodesic Voronoi diagram so that each node is integrated within an associated set of voxels defined by the local anatomical geometry (Grova et al., 2006).

From these methods stems a common interest in delineating 3D areas, sometimes overlapping, onto which the signal emanating from the nodes would possibly be dispersed, taking local anatomical features into account. Nevertheless, the problems remains very difficult because of numerous factors such as the volumes resolution, the partial volume effects, the sensitivity to errors from previous processing steps like segmentation, mesh extraction and registration, all of this being added to the highly folded nature of the surface. Therefore, a good interpolation method should show robustness to these various parameters.

* Corresponding author. ESIL, Case 925, 163 Avenue de Luminy, 13288 Marseille cedex 09, France. Fax: +33 4 91 82 85 51.

E-mail address: gregory.operto@univmed.fr (G. Operto).

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

In this study, we propose an original method to produce surface-based representations of the cortical activity directly on the cortical mesh, relying on an expected distribution of the BOLD signal in EPI volumes, taking into account as much anatomical information as possible. The method itself is described in the next section. The third section presents experiments we ran to validate our process and evaluate its robustness to segmentation and registration errors. Results are presented in the fourth section and the method is discussed and compared to others in the last section.

Materials and methods

Expected distribution of the signal

In order to represent the cortical activity from the whole 3D volume back to the surface where it was originated, we relied on an expected distribution of the BOLD signal around each node of the cortical mesh to address correlations between voxels and mesh nodes, with physiological and image-related motivations. Several distinct phenomena were considered: first, as a consequence of the columnar architecture of neurons inside the cortex (Mountcastle, 1978), we expect each cortical column to show homogeneous intrinsic activity, suggesting that this activity and the signal should be constant in the column orientation, i.e. normal to the surface within the cortical ribbon. Then, since neighborhood connections between columns lead to strong interactions along the surface, induced decaying activity around a functional activation focus can be measured (Johnson et al., 1993). Moreover, the image acquisition process inevitably leads the signal to sparsely spill outside the cortex, possibly mixing-up sources and bringing signal to anomalous areas. As for the BOLD signal measured with fMRI, it is shown that microvasculature organisation in the cortex is similar to the neuron columnar organisation (Duvernoy, 1999). Moreover, it has been suggested that the coupling between neural activity and the decrease in tissue oxygenation could be correlated at a columnar level (Thompson et al., 2003).

These observations resultingly inspired our method as it aims to figure out in which proportions neighboring voxels can be associated with the expected signal at a surface node and hence in which proportions their intensities should weigh on the interpolated node value. In respect with our assumptions, the distribution of activity on the cortical surface is thus depicted as a function of distance parameters along two main directions, normal and parallel (geodesic distance) to the surface, as shown in Fig. 1.

A scope of influence highly influenced by local anatomy hence gets defined for each node, within which voxels intensities will contribute in various proportions to the value on the surface. In this paper, these scopes are referred as convolution kernels, as the method computes a convolution product between functional




Fig. 1. Expected distribution of activity around a given point of the cortical surface presented with iso-influence curves.

volumes and some function related to anatomy, to recover from the volume the original signal from the surface.

Computing the convolution kernels

Preliminary steps consist in segmenting MR T1 anatomical images, extracting the white/gray matters interface triangulation and operating registration between functional volumes and anatomy. This is done using the BrainVISA package. To briefly summarize, a pipeline of processes is applied to the T1 MR image, including: bias correction, brain mask segmentation, hemisphere mask segmentation, and detection of the gray/white interface. A standard facet tracking algorithm is used to compute a first spherical mesh made of facets from the cortical mask, then the center of each facet is connected to the neighboring centers in order to build a spherical mesh of triangles (Mangin et al., 1995). A kernel is then computed for each node of the resulting mesh.

Each vertex gets actually associated to a mask describing in which proportions the intensities of surrounding voxels should weigh on the projected value. Such masks are always computed at the resolution of the functional images to be projected. Since these weights depend on distances, the method addresses the estimation of geodesic and normal distances of any voxel in relation to the surface. This is done by associating each voxel of the kernel to a node on the surface from which the distances are computed as explained hereafter. Computing a $15 \times 15 \times 15$ voxels mask at the resolution of $1 \times 1 \times 1 \text{ mm}^3$, or a $7 \times 7 \times 7$ mask at the resolution of $3 \times 3 \times 3 \text{ mm}^3$, the covered area is wide enough so any outer voxel has a null weight in respect to the current node.

The computation of normal distances relies on the propagation of a proximity information, starting from the voxel v_c containing the current node n_c . To perform this, we use a Fast Marching distance propagation algorithm adapted from the method described in Mauch and Breen (2000). The idea is to propagate a front that attributes to each surrounding voxel v a weight ω , depending on its closest node on the surface n_v . At each iteration, the voxel v_{current} of the front with the highest ω is considered processed and its neighbors v_i are added to the front, except those already processed. For each v_i , it is assumed that the associated n_{v_i} is in a close neighborhood of $n_{v_{\text{current}}}$. We define this neighborhood as the second order neighboring nodes of $n_{v_{\text{current}}}$ represented by the brown area in Fig. 3. It is possible to extend this second order neighborhood to a larger one if the resolution of the mesh is particularly fine, but a second order is enough for standard meshes generally used (20,000 to 30,000 nodes). Once n_{v_i} is found, its normal distance to the surface d_{v_i} is computed (with d_{v_i} being approximated by the Euclidean distance between v_i and n_{v_i}) and both are associated with v_i . The other neighbors of v_{current} are processed in a similar manner. Once every neighbor was assigned its proximity index to the surface, a specific weight ω_i is attributed to each of them, using two distinct weight functions, detailed below. Once all neighbors processed, the algorithm is iterated with the voxel having the highest ω_i . The algorithm is initialized with the nearest voxel v_c to the considered node n_c . This process is illustrated in Fig. 2. Fig. 3 details the first iterations of the algorithm.

Once every voxel of the kernel was assigned its normal node on the surface, each of them gets an estimated geodesic distance between n_c and the node n_v from which the normal distance was computed at first. A shortest path algorithm creates a geodesic distance map covering the whole area around n_c under the influence of its activity.



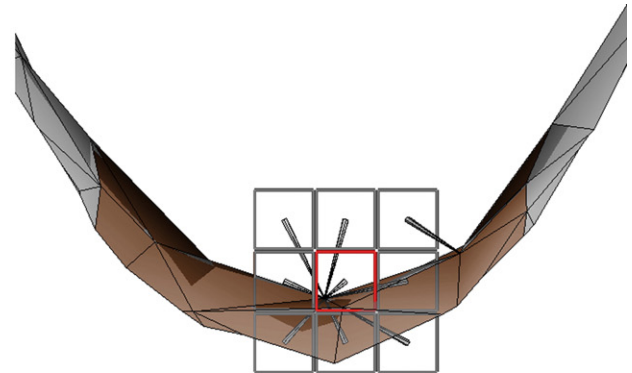



Fig. 2. (Left) Computation of normal (black lines) and geodesic (red line) distances using a Fast-Marching algorithm. (Right) Flowchart of the algorithm.




Iteration 1 : the Fast-Marching algorithm is initialized with the voxel containing the current node



Iteration 2 : the neighbors are added in the propagation front and assigned nodes on the surface, from a close neighborhood around the current node.



Iteration 3 : the voxel with the highest weight (i.e. the lowest distance) is removed from the front and its neighbors are processed



Iteration n : the range within which the nodes are selected varies in function to the currently processed voxel

Fig. 3. First iterations of the algorithm. The red voxel is the one whose neighbors (in dark-gray) are processed at the current step. The current front is illustrated with light-gray voxels. The normal nodes assigned to the voxels are selected from the brown area on the surface.




Fig. 4. Geodesic (left) and normal (right) weight functions.

The voxels characterized by two distances are then given weights by the combination of two functions described in Fig. 4. Geodesically, we estimate that a node located beyond 5 mm far from n_c is no longer under influence: aside from its physiological justification (Johnson et al., 1993), this user-specified distance allows to handle the smoothness of the final projection, similar to the variance of the basis functions defined by Kiebel et al. (2000). On the other hand, the normal distance weight depends directly on the anatomical structure the voxel is located in. The normal weight of a voxel is maximal if located in the cortical ribbon and decreases outside to reach zero at a user-specified normal distance of 3 mm (Fig. 4). As a result, each voxel v gets assigned a value ω corresponding to the product of its two weights. Finally, each mask is normalized so that the sum of all weights equals 1, for energy and intensity range preservation purposes. Fig. 5 displays a convolution kernel computed on a flat synthetic surface mesh, showing a plateau corresponding to the cortical ribbon. Fig. 6 shows on real cortical meshes three kernels with shapes strongly influenced by local anatomy.

It must be mentioned that locating a voxel in respect to the cortical ribbon, in order to determine its normal weight, is problematic in this context. Indeed, the use of a cortical mask would be suitable for this but is not a good algorithmic solution and generally fails in the depth of sulci where the quality of the cortex segmentation and the mesh normals can be questionable. On the other hand, an appropriate way is to provide each node an estimation of cortical thickness, beyond which voxels outside the cortical mesh should get lower weights. Even though our implementation can take such a cortical thickness texture as input, to date we chose to set this thickness parameter as a 3 mm constant value. This is the most relevant option considering the estimation of cortical thickness is a whole different problem by itself (Lerch, 2005), especially when working with $3 \times 3 \times 3$ mm³ functional voxels.




Fig. 5. Convolution kernel obtained for a node a flat surface with superimposed simulated cortical mask.

Projecting the functional data

Once the convolution kernels are computed for the whole mesh, we then integrate the voxel intensities onto the surface nodes: since any surrounding voxel has an estimated influence on the node, the value $s(n_s)$ assigned to the node n_s is the result of a linear combination expressed as follows:

$$s(n_s) = \sum_{v_i} \omega_s(v_i) \cdot I(v_i)$$

where v_i is the voxels area within which influences on n_s were evaluated, $\omega_s(v_i)$ is the computed weight for voxel v_i and $I(v_i)$ is v_i 's intensity on the functional image. The method is designed so that the costly step of computing the masks is done only once by subject, then any functional image with the same spatial resolution as the kernels and that is registered with anatomy can quickly be projected onto the cortical surface.

Experiments

Simulated data sets

We led a series of experiments in order to assess the benefits of projecting functional data using convolution kernels, in terms of detection power, robustness to various kinds of errors and sensitivity to resolution changes. Most of these experiments were previously suggested by Grova et al. (2006) for the Voronoi-based approach. In this validation frame, several types of meshes were generated. We first extracted cortical meshes from the T1-weighted volume at three different resolution levels: one «very-low-resolution» model with 10,000 nodes, one «low-resolution» model with 15,000 nodes and one «high-resolution» model with 25,000 nodes. The inter-node average distance for these meshes was respectively 3.03 mm, 2.56 mm and 2.07 mm.

The point was here to estimate the contribution of the kernel-based method and the influence of the mesh spatial resolution. Besides, different ranges of translations, from -5 mm to 5 mm, along x , y and z axes, were applied to the low and high resolution meshes, so as to simulate rigid registration errors in relation to the functional volume. Projections onto these meshes intended to address the robustness of our approach to this kind of errors, in comparison to other methods. Finally, in order to mimic geometric distortion and segmentation errors, we shifted each mesh node, along its normal direction (or its opposite), of a distance (up to 3 mm) given by a random scalars map previously computed and smoothed along the surface. The random normal distortion follows a uniform distribution between -3 and 3 mm, and the smoothing is performed using a diffusion process as presented in Chung and Taylor (2004) equivalent to a Gaussian smoothing of Full Width at




Fig. 6. Convolution kernels obtained for nodes on real cortical meshes. The kernel on the left shows a superimposed cortical mask.

Half Maximum equal to 21 mm. The smoothing is performed only to preserve a minimum regularity to the mesh after distortion, as it would even in the case of segmentation errors or misregistration. Again this aimed at evaluating the method's sensitivity to such noise, typically existing in real images.

We also simulated an fMRI activity map, in a similar fashion as in Grova et al. (2006) by drawing first an activated area on the cortical mesh, previously extracted from the anatomical T1 volume. From this area on the mesh, defined around a manually selected node, a set of voxels of 1.5 cm^3 were picked out within the cortical ribbon, considered as activated voxels and taken as 3D ground truth. This binary volume at anatomical resolution ($1 \times 1 \times 1 \text{ mm}^3$) was then smoothed using a 3D Gaussian kernel of Full Width at Half Maximum (FWHM) 5 mm, and down-sampled to the resolution of EPI volumes ($61 \times 73 \times 61$ matrix, isotropic voxel size $3 \times 3 \times 3 \text{ mm}^3$) in order to reproduce the spatial distribution of signal observed in real fMRI results. Fig. 7 illustrates a simulated activity map superimposed on a T1-contrast MRI volume.

Finally, the 3D voxel-based activation maps were projected onto these various meshes to give surface-based cortical activation maps by using four different projection techniques: sphere-based interpolation with 3 mm radius, with 5 mm radius (with iteration step 1 mm), interpolation along the outside normal line (on 3 mm, with iteration step 1 mm) and kernels-based projection.

Besides, the methods were compared when projecting high-resolution functional volumes. This was performed by computing kernels at anatomical resolution, then projecting the previously created binary volume on the surfaces at low, very-low and high resolutions.

All the values are normalized between 0 and 1 before being compared from one method to another.

Comparison metric

Receiving-Operator Characteristic (ROC) analysis is a relevant approach to quantify the sensitivity and specificity of a binary classifier as its discrimination threshold varies. In this context, According to ROC methodology, each detection – i.e. each thresholded projection, in this context – presents, for a given threshold value, a certain amount of true positive and false negative nodes regarding the relative ground truth. As a result, the ROC curve plots sensitivity and specificity of the projection for each threshold between 0 and 1. The Area Under the ROC Curve is typically worthy of consideration in ROC analysis as it renders the various sensitivities and specificities computed for one projection in one single value, which can hence be used as a comparison index between the methods. In particular, it is demonstrated that this value is related to the probability for an active node to show higher intensity in the projected map than an inactive node. The results that will be presented in this paper, which aimed at comparing the different methods for projection of functional voxel-based data onto the surface, mainly rely on these Area Under Curve values. However, as active and inactive nodes exist in very different numbers on the mesh, we use the $\text{AUC}_{\text{close}}$ iterative estimator (Grova et al., 2006) which randomly selects inactive nodes in a close neighborhood around the active area for the estimation of sensitivity and specificity. The resulting value gets less biased as it is computed from equal numbers of active and inactive nodes.




Fig. 7. (Left) Simulated functional map used along the different experiments. (Right) Projection of this functional map presented on a slightly inflated cortical mesh.

Results

We projected simulated functional data onto the cortical surfaces at three different resolutions. For each projection, we computed an AUC_{close} value. The results are presented in Table 1. They tend to illustrate kernel-based projection's best accuracy in recovering, out from a simulated volume, the surface-based activation map initially taken as ground truth, in comparison to other methods, and for each resolution. With very-low resolution meshes, all the methods give results very close or equal to unity. This may result only from the low number of nodes for reproducing the variations of intensity across voxels. As the nodes distribution is sparser, the nodes are more inclined to get the same value as the most homogeneous areas in the volume (i.e. either 0 or 1), therefore being much less influenced neither by the volume smoothing. It may also be due to overlapping between interpolation kernels, which varies from one resolution to

another; still, at functional resolution, this parameter may not be of greater influence.

We projected the same simulation activation volumes on cortical meshes that were applied various shifts along the three axes. For each of the four compared methods, we computed an AUC_{close} value for every shift range (Fig. 8). One can appreciate the robustness of kernel-based projection which shows good sensitivity and specificity even using misregistered cortical meshes. Moreover, translation axes do not affect the results the same way: noticeably, shifts along z axis, in negative values (from bottom to top), make the detection rapidly fail, followed by the normal-based interpolation. This particular case is actually highlighting the anatomically-inspired nature of these two methods compared to sphere-based interpolation. Indeed, according to the activation location along the precentral gyrus, such translations from bottom to top bring the activation cluster underneath the cortical mesh. Hence, only methods which integrate voxels independently from their anatomical location could reach this cluster. Since kernel-based interpolation gives preference to outside voxels (belonging to the cortical ribbon), the activation cluster in this case is not described well by the projection.

On the other hand, the mesh spatial resolution has no particular influence on these results. From a general point of view, the AUC_{close} values appear higher using the low-resolution mesh than the high-resolution mesh, for the various reasons previously mentioned. However, the same observations regarding the relative better robustness of kernel-based interpolation to small translations are rigorously reproducible from one resolution to another.

We projected simulated functional volumes onto the high-resolution cortical mesh, on which simulated geometric distortion was previously applied, each mesh node being translated along its normal direction of a random distance. Applied distance shifts were generally ranged under 3 mm which is reasonable according to the functional resolution. We computed AUC_{close} values using low, very-low and high-resolution meshes. Results are detailed in Table 1. Kernel-based projection is presented as sensitive to these errors as other methods. Once again, this may directly depend on whether the surface was locally shifted towards the outside or the inside. As mentioned previously, the method is designed so as to integrate intensities from the cortical ribbon, i.e. outside the cortical mesh. White matter voxels have very little influence on the projected values. Therefore, whereas translations towards the inside may not affect the integration of relevant voxels, those towards the outside may be responsible for disturbing the detection. To illustrate this point, the random distortion map was derived into one with positive scalars only (the nodes are all drawn of 1 mm outside the pial interface) and one with negatives only (the nodes are all drawn of 1 mm inside the white matter) and the experiment was reiterated using these two consequent distortion maps. Resultingly, our approach shows a better robustness when nodes are pulled to the inside than to the outside, whereas it has no influence using sphere-based methods.

We finally projected simulated functional volumes with $1 \times 1 \times 1 \text{ mm}^3$ voxels onto the cortical surface, using sequentially the low, very-low and high-resolution meshes. Instead of the same volumes as previously, we used the simulated volumes before the down-sampling step. The volumes hence showed the smoothness of a typical functional map but at higher resolution. We computed AUC_{close} values which are presented in Table 1. From a general point of view, all methods tend to show higher detection accuracy when projecting high resolution volumes. Of course, this effect

Table 1

AUC_{close} values from various experiments: estimating influences of mesh and volume spatial resolutions, robustness to random errors and to uniform shifting of the mesh nodes

	Sphere 5 mm	Sphere 3 mm	Normal outward	Kernels
<i>Detection accuracy for each method and various mesh spatial resolutions</i>				
Very low resolution (10, 000 nodes)	1.000	0.997	1.000	1.000
Low resolution (15, 000 nodes)	0.999	0.978	0.997	1.000
High resolution (25, 000 nodes)	0.999	0.966	0.965	0.989
<i>Detection accuracy projecting $1 \times 1 \times 1 \text{ mm}^3$ functional maps</i>				
Very low resolution (10, 000 nodes)	0.997	0.976	0.997	1.000
Low resolution (15, 000 nodes)	1.000	0.992	0.998	1.000
High resolution (25, 000 nodes)	0.997	0.987	0.959	0.988
<i>Detection robustness to simulated random distortion errors</i>				
Very low resolution (10, 000 nodes)	1.000	0.991	0.975	0.986
Low resolution (15, 000 nodes)	0.983	0.991	0.975	0.967
High resolution (25, 000 nodes)	0.975	0.966	0.947	0.979
<i>Detection accuracy after shifting all nodes of 1 mm outwards</i>				
Very low resolution (10, 000 nodes)	1.000	0.975	0.964	0.975
Low resolution (15, 000 nodes)	0.994	0.988	0.976	0.970
High resolution (25, 000 nodes)	0.963	0.944	0.942	0.951
<i>Detection accuracy after shifting all nodes of 1 mm inwards</i>				
Very low resolution (10, 000 nodes)	0.999	1.000	0.996	0.998
Low resolution (15, 000 nodes)	0.992	0.987	0.991	0.997
High resolution (25, 000 nodes)	0.992	0.986	0.994	0.984

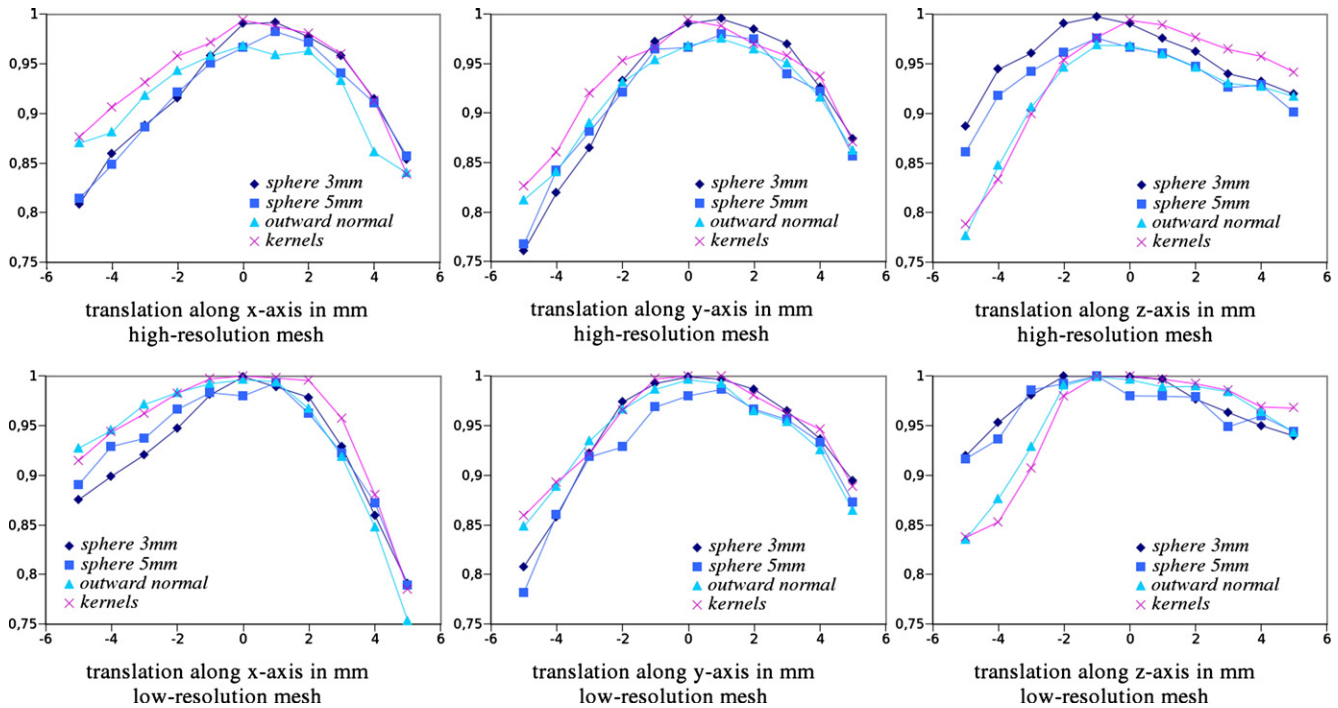


Fig. 8. Detection robustness to small registration errors: AUC values plotted in function to shifts applied to low and high resolution meshes.

could be related to the observations on mesh resolution influence: when using a high-resolution volume, the relative spatial resolution of the mesh logically decreases, in comparison to when projecting low-resolution volumes. Still, high-resolution kernels embed more information about the local anatomy around each node, which may lead to this observed stability in AUC_{close} values. Further validation could include a measure of significance of these results. This would involve experiments on a large number of meshes and activation clusters at various locations.

As an illustration of the process, Fig. 9 presents projections of functional maps issued from a functional experiment. A volume-based SPMt map has been computed for a simple (finger movement–rest) contrast with a spatial resolution of $3 \times 3 \times 3 \text{ mm}^3$. Coregistration between functional and anatomical data has been performed using SPM2 (see <http://www.fil.ion.ucl.ac.uk/spm/software/spm2/> and Ashburner and Friston, 2000). This map has been projected onto the corresponding cortical surface to illustrate the behavior of our method. The figure presents the same projection performed with a standard sphere-based interpolation (of radius 5 mm), for visual comparison. It can be seen that the nodes in the depth of the sulcus show higher values on the kernel-based projection than on the sphere-based projection: indeed, in that area, the sphere-based method averages the intensities of many voxels from the white matter as pictured in Fig. 9, therefore decreasing the projected value. This shows once again the benefits of integrating anatomical information into the projection process.

Discussion

Choice of the orthogonality criterion

In order to compute geodesic and normal distances for any voxel in a node’s neighborhood, each voxel is assigned what was

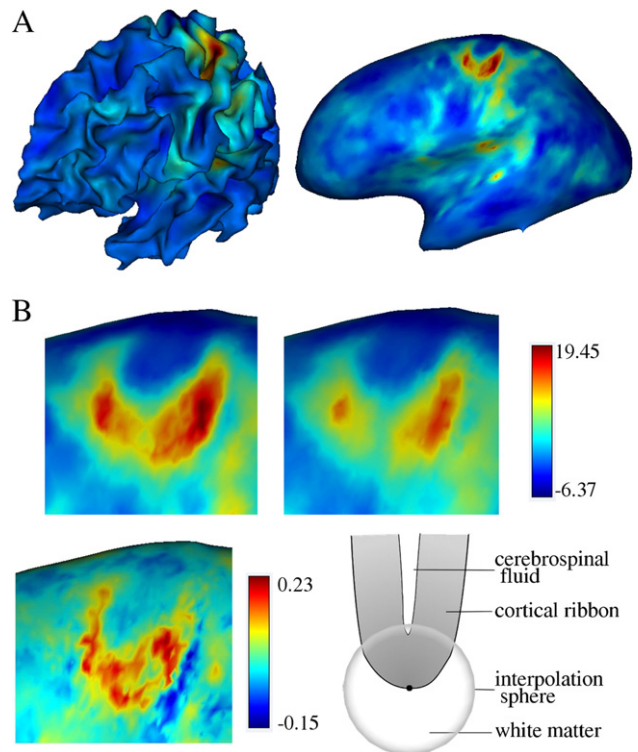


Fig. 9. A—Projection of a SPMt map onto its matching cortical mesh, presented on original and inflated surfaces. B—(top) Magnified views of projected activations: (left) using kernel-based method (right) using 5 mm sphere-based method. (Bottom left) difference between the two textures (bottom right) sphere-based interpolation in the depth of a sulcus.

referred to as a normal node on the surface, from which the two distances, then the weights, were computed. These associations result from a Fast-Marching propagation algorithm. Relating voxels to surface nodes using geodesic distances is a relevant way to deal with the folded nature of cortex, rather than euclidean distances. However, the original orthogonality criterion we use presents various advantages in comparison to others. For instance, the closest node criterion generally fails in configurations where several nodes are almost equidistant to the considered voxel, like inside sulci, and a slight shift of the mesh can hence lead to totally opposite node associations. This could also lead to irrelevant cases, e.g. when two neighboring voxels are assigned nodes belonging to opposite sides of a sulcus. Relying instead on the mesh normals directly faces similar drawbacks. According to this, this fashion of propagating node associations and distances among the voxels implies a geometrical coherence between these nodes, by giving implicit preference to nodes on the same side of sulcus as the node at the center of the current kernel. This consequently increases the method's robustness to errors in registering the functional volumes to anatomy.

Comparisons with other methods

Averaging voxels intensities within a sphere centered at each node can mix up signals from different structures like opposite sides of a sulcus, or gray and white matters. In our approach, the shape of the interpolation area around each node is strongly influenced by local anatomy, which prevents from integrating intensities from other anatomical structures by penalizing, through the geodesic and normal weights, voxels not belonging to the area of interest. Moreover, by sphere-based interpolation, voxels inside the sphere contribute equally to the value attributed to the node whereas one main assumption behind kernel-based projection is that voxels should weigh in various proportions on the projected value regarding the relevance of their location towards the considered node.

Averaging voxels intensities along the normal line outside the cortical surface precisely prevents from blurring information issued from separate sources. First, only outside voxels are considered, giving preference to signals from the cortical ribbon. Then, the normal information selects the voxels to average, which is a first step in embedding anatomical information such as local geometry. However, the method suffers from several limitations: first, the number of considered voxels remains low, regarding the spatial resolutions of the mesh and the functional volume. Much information is hence ignored by the process since only one single line is taken into account. Then, the method gets the projection sensitive to normals estimation quality, which is questionable especially inside sulci. In this case, kernel-based projection overcomes this issue by propagating the normal proximity information among voxels.

Voronoi-based projection (Grova et al., 2006) embeds more anatomical features. Indeed, the average of voxels intensities around each node is performed within a Voronoi cell whose shape is constrained by the local geometry of the cortical ribbon. This implies more voxels contribute to each projected value, in comparison to interpolation along normal line. Also, using a cortical mask to filter out voxels outside of the cortex and parcelling the cortical ribbon in relation to nodes scopes, the method prevents from mixing up signals from different anatomical structures. Nevertheless, associating a voxel to one surface node

exclusively is a questionable point. In our approach, as each voxel is bound to receive signal from various nodes, one voxel's intensity can be distributed back to several nodes, their influences depending on their distances, geodesic and normal, to a given node. As a consequence, contrary to Voronoi cells, neighboring nodes have overlapping scopes, still without integrating signals from other anatomical regions, penalized by low weights. One could argue that a Voronoi-based projection could be followed by a surface-based smoothing to «distribute» the activity. But the smoothing would distribute only in a close neighborhood on the surface whereas our method allows the intensity of a voxel inside a sulcus to be projected on distant nodes (in the geodesic sense) belonging to opposite sides of the sulcus. This issue is particularly important considering the relatively low resolution of functional data.

One advantage of the Voronoi approach resides in its ability to integrate a maximum amount of information, including voxels possibly distant from the surface which belong though to some Voronoi cells. In particular, this works out some cases of failed segmentation, for instance when a blurry gray/white matter interface leads to lower gyral ridges. Moreover, the use of a functional mask filters out the voxels containing no information. In our method, the non-significant voxels are subsequently filtered-out by being located outside of any vertex's influence scope, whereas the user-specified decay parameter allows to take into account the possibly distant, relevant voxels.

Kernel-based projection therefore gathers advantages from overlapping interpolation spheres and anatomically-informed Voronoi cells. From a general point of view, the implementation allows to emphasize any of these features independently: by setting a null geodesic weight function, only normal influence is considered and the interpolation gets analogous to outside normal line averaging. By choosing symmetric gate functions for normal and geodesic influences, we get some sorts of spheres distorted by the local geometry. The user can hence give preference to any of these influence factors.

Conclusion

In this paper, we propose a method allowing the projection of functional images onto the cortical surface. Relying on physiological and image-related hypotheses, it determines, for each mesh node, a specific interpolation area, strongly influenced by local anatomy, within which voxels intensities are averaged. Applying it to activation maps can serve visualisation purposes, as well as cortical localisation of activation foci through the use of a surface-based coordinate system (Clouchoux et al., 2005; Fischl et al., 1999), in comparison to localisation in 3D normalized spaces. Finally, this algorithm is freely available and distributed with the BrainVISA package (<http://brainvisa.info>).

In application to functional sequences of raw data registered with the anatomy, the produced surface-based representations of the BOLD images are essential in the framework of cortical-restricted statistical analysis methodologies, such as Cortical Surface Mapping (Andrade et al., 2001). The experiments ran on various simulated data showed better sensitivity, and robustness to registration and geometric distortion errors, comparing our method to sphere-based interpolation and average along outside normal line. Still, inferring a surface-based representation of the signal from the cortical surface, on which it was originated, back from voxel-based volumes, is a difficult problem mainly due to

low resolution of functional images. Its context will get more compliant as images resolutions will increase and preliminary issues get solved (e.g. better segmentation of corticals interfaces, estimation of cortical thickness, lower partial volume effect, etc.). According to this, future research will address surface-based functional data analysis, in which this work takes part as a preliminary step.

References

- Andrade, A., Kherif, F., Mangin, J.-F., Worsley, K., Paradis, A.-L., Simon, O., Dehaene, S., Poline, J.-B., 2001. Detection of fMRI activation using cortical surface mapping. *Human Brain Mapping* 12, 79–93.
- Ashburner, J., Friston, K.J., 2000. Voxel-based morphometry—The methods. *NeuroImage* 11, 805–821.
- Chung, M.K., Taylor, J., 2004. Diffusion smoothing on brain surface via finite element method. *Proceedings of the 2004 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, Arlington, VA, USA, 15–18 April 2004. IEEE, pp. 432–435.
- Clouchoux, C., Coulon, O., Riviere, D., Cachia, A., Mangin, J.-F., Regis, J., 2005. Anatomically constrained surface parameterization for cortical localization. In: Duncan, J., Gerig, G. (Eds.), *LNCS—Medical Image Computing and Computer-Assisted Intervention—MICCAI 2005: 8th International Conference*, vol. 3750. Springer-Verlag, Berlin, pp. 344–351.
- Duvernoy, H.M., 1999. *The Human Brain—Surface, Blood Supply and Three-Dimensional Sectional Anatomy*. Springer.
- Fischl, B., Sereno, M.I., Tootell, R., Dale, A.M., 1999. Cortical surface-based analysis, ii: inflation, flattening, and a surface-based coordinate system. *NeuroImage* 9, 195–207.
- Flandin, G., Kherif, F., Pennec, X., Malandain, G., Ayache, N., Poline, J.-B., 2002. Improved Detection Sensitivity in Functional MRI Data using a Brain Parcelling Technique. *MICCAI 2002: 5th International Conference*. Springer-Verlag, Berlin.
- Goebel, R., Singer, W., 1999. Cortical surface-based statistical analysis of functional magnetic resonance imaging data. *NeuroImage* 9, S64.
- Grova, C., Makni, S., Flandin, G., Ciuciu, P., Gotman, J., Poline, J.-B., 2006. Anatomically informed interpolation of fMRI data on the cortical surface. *NeuroImage* 31, 1475–1486.
- Johnson, P.B., Ferraina, S., Caminiti, R., 1993. Cortical networks for visual reaching. *Experimental Brain Research*, vol. 97, pp. 361–365.
- Kiebel, S.J., Goebel, R., Friston, K.J., 2000. Anatomically informed basis functions. *NeuroImage* 11 (6.1), 656–667.
- Lerch, J., 2005. *In-vivo Analysis of Cortical Thickness using Magnetic Resonance Images*, PhD thesis, Department of Neurology and Neurosurgery, McGill University, Montreal, Canada.
- Mangin, J.-F., Frouin, V., Bloch, I., Régis, J., Lopez-Krahe, J., 1995. From 3D magnetic resonance images to structural representations of the cortex topography using topology preserving deformations. *J. Math. Imaging Vis.* 5, 297–318.
- Mauch, S., Breen, D., 2000. A fast algorithm for computing the closest point and distance function. Technical report, unpublished, CalTech.
- Mountcastle, V.B., 1978. An organizing principle for cerebral function: the unit module and the distributed system. *The mindful brain: cortical organization and the group selective theory of higher brain function*, page 7.
- Saad, Z., Reynolds, C., Argall, B., Japee, S., Cox, R.W., 2004. Suma: an interface for surface-based intra- and inter-subject analysis with afni. *Proc. IEEE Intl. Symp. Biomedical Imaging*, p. 1510.
- Thompson, J.K., Peterson, M.R., Freeman, R.D., 2003. Single-neuron activity and tissue oxygenation in the cerebral cortex. *Science* 299 (5609), 1070–1072.
- Toro, R., Burnod, Y., 2003. Geometric atlas: modeling the cortex as an organized surface. *NeuroImage* 20 (3), 1468–1484.
- Van Essen, D.C., Drury, H.A., 1997. Structural and functional analyses of human cerebral cortex using a surface-based atlas. *J. Neurosci.* 17 (18), 7079–7102.
- Warnking, J., Dojat, M., Gurin-Dugu, A., Delon-Martin, C., Olympie, S., Richard, N., Chhikian, A., Segebarth, C., 2002. fmri retinotopic mapping—Step by step. *NeuroImage* 17 (4), 1665–1683.