

# Structural Group Analysis of Functional Activation Maps

O. Coulon,<sup>\*†1</sup> J.-F. Mangin,<sup>†</sup> J.-B. Poline,<sup>†</sup> M. Zilbovicius,<sup>†</sup> D. Roumenov,<sup>†</sup>  
Y. Samson,<sup>†</sup> V. Frouin,<sup>†</sup> and I. Bloch<sup>\*</sup>

<sup>\*</sup>Departement TSI, Ecole Nationale Supérieure des Telecommunications, 46 Rue Barrault, 75631 Paris Cedex 13, France;  
and <sup>†</sup>Service Hospitalier Frederic Joliot, CEA, 91401 Orsay Cedex, France

Received September 24, 1999

**We present here a new method for cerebral activation detection over a group of subjects. This method is performed using individual activation maps of any sort. It aims at processing a group analysis while preserving individual information and at overcoming as far as possible limitations of the spatial normalization used to compare different subjects. We designed it such that it provides the individual occurrence of the activations detected at a group level. The localization can then be performed on the individual anatomy of each subject. The analysis starts with a hierarchical multiscale object-based description of each individual map. These descriptions are then compared, rather than comparing the images directly. The analysis is thus performed at an object level instead of voxel by voxel. It is made using a comparison graph, on which a labeling process is performed. The label field on the graph is modeled by a Markov random field, which allows us to introduce high-level rules of interrogation of the data. The process has been evaluated on simulated data and real data from a PET protocol.** © 2000

Academic Press

**Key Words:** functional neuroimaging; functional activation detection; scale space; Markov random fields.

## INTRODUCTION

Understanding the neural correlates of human brain function is a growing field of research. Due to the highly noisy nature of cerebral functional images, like positron emission tomography (PET) or functional magnetic resonance imaging (fMRI), brain activation detection has essentially been approached so far in terms of statistical analysis (Friston *et al.*, 1995; Worsley *et al.*, 1996a; McIntosh *et al.*, 1996; Ledberg *et al.*, 1998) using a common anatomical reference. Although they have been validated in a wide range of applications, these analyses lead to some problems in terms of

localization and/or detection with regard to anatomy, mainly because the spatial normalization performed to compare images from different subjects matches only gross (functional) features. As the comparison is made at a voxel level, and anatomical information being poorly considered, the process becomes very sensitive to intersubject variability and normalization limitations.

After a statistical analysis, it is also generally difficult to estimate from the group result the individual activated areas in terms of shape, extent, or position with respect to anatomical landmarks. This knowledge should help to study intersubject functional variability and would improve localization with regard to anatomy. Moreover, the contribution of each subject to the group result remains unknown and recent interest has been taken in the link between group and individual results (Friston *et al.*, 1999).

Among all the analysis methods proposed so far, several have considered looking at the data at different levels of scale (smoothing) or resolution (Poline and Mazoyer, 1994a,b; Worsley *et al.*, 1996b; Unser, 1995; Lindeberg *et al.*, 1999) to improve the detection sensitivity. Poline and Mazoyer (1994a) showed that a hierarchical description of the gray-level “landscape” provides a better structure extraction than the simple thresholding procedure usually performed. However, they analyze every level of scale separately. It has been proven in the field of computer vision that looking at all scales simultaneously provides a relevant and exhaustive description of the considered image structure (Koenderink, 1984; Koenderink and van Doorn, 1986; Pizer *et al.*, 1986; Lindeberg, 1993a; Vincken *et al.*, 1997). More precisely, making explicit the behavior of the image through the scales leads to this description. Such multiscale methods can provide an analysis of the image in terms of singular points, objects, or particular features such as cores or boundaries.

Recent works (Lindeberg *et al.*, 1999; Coulon *et al.*, 1997a) use that kind of multiscale approach to perform analyses of statistical maps. However, in the case of a group analysis, these works which use a multiscale

<sup>1</sup> Corresponding address: Dept. of Computer Sciences, University College London, Gower Street, London WC1E 6BT, UK. E-mail: O.Coulon@is.ucl.ac.uk.

method to analyze a single map have to face problems of intersubject variability and spatial normalization. To overcome these problems of variability, we advocate that the comparison of images should be done in terms of objects instead of voxels. Using objects permits the use of additional information such as individual anatomical location, shape, or extent.

A group analysis should thus be done using individual activation maps and comparing these maps at an object level, a higher level of description that can include semantic information.

We propose here a new method, based on a multi-scale hierarchical description of individual activation maps in terms of structures. This is followed by a graph-based comparison of these descriptions across subjects, rather than comparing directly the images in a stereotactic space at a voxel level. This comparison is the starting point of an activation detection process using Markov random fields. The method is designed to overcome, as far as possible, the problems induced by spatial normalization. After detection over a group of subjects, the method allows an easy way back to the individual structures and more generally permits high level, and in the future more informed, interrogation of functional data sets.

## METHOD

### Overview

The method presented here is applicable to any kind of "activation maps," for instance, PET/fMRI difference images or statistical  $t$  maps. It is divided into the three following steps, illustrated by Fig. 1: first, each individual map involved in the study is described by a hierarchical object-based multiscale representation, namely the scale-space primal sketch. Second, a graph is built, which contains all of the primal sketches and represents the potential repetitions of objects from one subject to another. Finally, a labeling process is performed on the graph, which aims to identify the objects representing functional activations and those representing noise. The label field on the graph is modeled as a Markov random field.

### The Scale-Space Primal Sketch

The scale-space primal sketch is a multiscale image description, proposed by Lindeberg (1993a), similar to the primal sketch suggested by Marr (1976, 1982). It is based on well-known properties of linear scale space to describe the first-order structure of images (Koenderink, 1984; Koenderink and van Doorn, 1986). We present here a brief description of linear scale space and the way the primal sketch is built. For more precise details, we invite the reader to refer to Lindeberg (1992) or to Coulon *et al.* (1997a,b) for the particular three-dimensional case applied to activation maps.

### Background on Linear Scale Space

The linear scale-space representation of a  $N$ -dimensional image  $f(\cdot)$  consists in building a one-parameter family of  $N$ -dimensional images  $L(\cdot, t)$  derived from the original under precise assumptions, namely homogeneity, isotropy, and causality (Koenderink, 1984; see also Alvarez *et al.*, 1993). The parameter  $t$  is continuous and is called *scale*. Under these assumptions, the scale-space representation is proved to verify a diffusion equation, the *heat equation*,

$$\frac{\partial L}{\partial t} = \frac{1}{2} \Delta L,$$

with the initial condition  $L(\cdot, 0) = f(\cdot)$ . In the continuous case, there is a strict equivalence between this representation and building  $L(\cdot, t)$  by convolution of the original image with Gaussian functions of increasing parameter  $t$ :


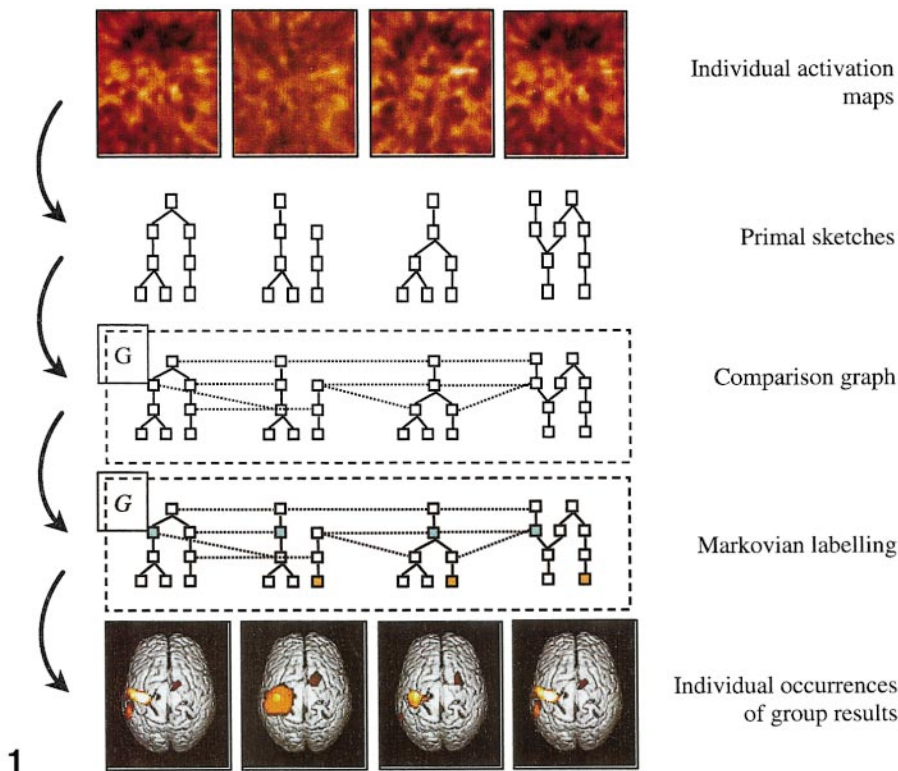
$$L(\cdot, t) = f(\cdot) \times G_t(\cdot),$$

$$\text{with } G_t(\vec{x}) = \frac{1}{(2\pi t)^{N/2}} \exp\left(-\frac{\vec{x} \cdot \vec{x}}{2t}\right).$$

The effect of these convolutions is a smoothing which progressively removes details, as illustrated in Fig. 2 for an activation map. Practically, one has various ways to build the scale space of an image, either by using the diffusion equation or by implementing the Gaussian convolution, although the accuracy of discrete scale spaces and the verification of causality depend strongly on the choice of the implementation. In our case, the scale space is built using approximation of Gaussian convolutions with recursive filtering (Deriche, 1993), shown to have an appropriate speed and precision. Theoretical properties of the linear scale space of an image have been widely investigated and are now well known (Koenderink, 1984; Florack *et al.*, 1992; Nielsen *et al.*, 1994). Particularly, for image description, a major topic concerns how to relate structures at different scales, namely the *deep structure* (Koenderink, 1984), and what is the behavior of singular points (local extrema and saddle points) or particular features (e.g., ridges) through the scales (Lindeberg, 1992; Pizer *et al.*, 1986; Lifshitz and Pizer, 1990; Fidrich, 1997; Morse *et al.*, 1994; Vincken *et al.*, 1997). The scale-space primal sketch is based on that type of description and is presented below.

### The Primal Sketch

The primal sketch is a multiscale hierarchical description making explicit the behavior through the scales of objects in the image. The only assumption



**FIG. 1.** An overview of the process.

**FIG. 2.** A slice of a 3D activation map and the corresponding scale levels  $t = 1, 2, 4, 8, 16, 32$ .

**FIG. 4.** A slice of 3D gray-level blobs extracted from an activation map at scale levels  $t = 1, 2, 4, 8, 16, 32$ .

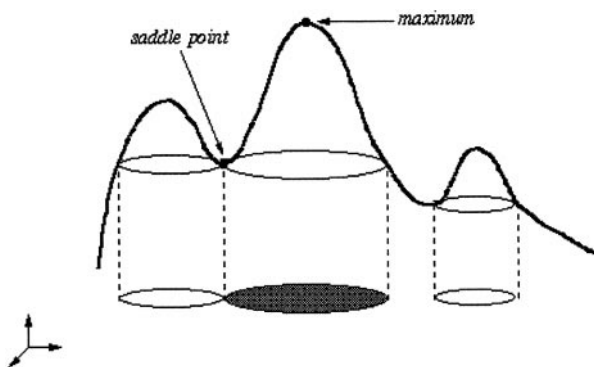
required to use it is that the objects of interest, functional activations, are blob-like.

The first step is to compute a linear scale space from the considered three-dimensional activation map. The scale parameter is continuous, but for obvious practical reasons we sample the scale space along the scale axis. This sampling is done in the classic logarithmic way, due to the fact that the number of events occurring in the scale space (e.g., changes in the number of singular points) increases logarithmically with the scale (Lindeberg, 1993b).

At each computed level of scales, objects—called *gray-level blobs*—are extracted in a fully automatic way, based on singular points in the image (local extrema and saddle points), as illustrated in Fig. 3 for the two-dimensional case.

Gray-level blobs have a mathematical definition (Lindeberg, 1993a), but their building is easy and intuitive to understand: from each local maximum, let a region gradually grow around this maximum, until it meets another region of the same kind or a point known as belonging to the “background.” The “growth”





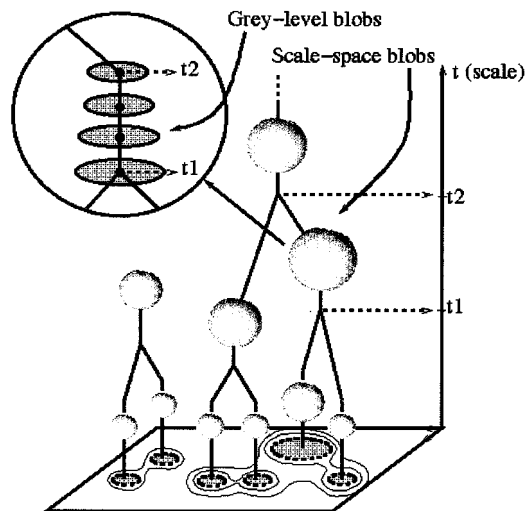
**FIG. 3.** Two-dimensional illustration of the gray-level blob definition reproduced, by permission of the publisher, from Lindeberg, 1993a).

of a blob is stopped at a saddle point, whose intensity defines the base level of the blob. Each gray-level blob is then defined by a local maximum and the associated saddle point. This definition is interesting in three ways: it provides an automatic extraction without any threshold, it guarantees the theoretical knowledge of gray-level blobs' behavior through the scales (since the behavior of singular points is known), and it associates one blob with each local maximum, which makes sense for our activation detection problem. We perform the gray-level blob extraction in a 3D fashion, following the same process as the one previously presented. Figure 4 illustrates the extraction for the different scale levels presented in Fig. 2.

In Fig. 4, one can see that the two blobs circled at scale level  $t = 4$  still exist at  $t = 8$  and  $t = 16$ . But at  $t = 32$  they have merged into a single blob. This roughly means that at  $t = 4$  the gray level landscape is locally composed of two peaks and that these two peaks lie on an underlying structure which is represented by the single blob at  $t = 32$ . The primal sketch aims to make this kind of information explicit. Therefore, gray-level blobs have to be linked from one scale to another, and events in the scale space have to be detected. These events can be theoretically itemized into four types (Lindeberg, 1992) and are called *bifurcations*. They refer to the relative behavior of extrema and saddle points (annihilation and creation of saddle point/extremum pairs) and can be annihilation of a blob, merge of two blobs, creation of a blob, or split of a blob, as illustrated by Figs. 5b–5e.



**FIG. 5.** The matching relations between gray-level blobs through the scales: (a) plain link, (b) annihilation, (c) merge, (d) creation, (e) split.



**FIG. 6.** Symbolic representation of the scale-space primal sketch.

If one considers a gray-level blob at a given scale level, there is generally a corresponding blob at a slightly finer scale and at a slightly coarser scale. We want to link these gray-level blobs, and we use a simple overlap of the spatial support criterion for that (assuming that a gray-level blob hardly moves out of its support with a small increase of the scale parameter). Either this linking is possible or we detect a bifurcation. Performing such a linking is practically difficult, and the reader can find the exact algorithm in Coulon *et al.* (1997b). Between two bifurcations, a single gray-level blob evolves through the scales and is identified as a multiscale object called *scale-space blob*. The primal sketch is then composed of scale-space blobs linked by bifurcations, as described in Fig. 6.

This description is completed with measurements associated with each scale-space blob. These measurements are first defined for each gray-level blob and are related to the average and maximum intensity, contrast, extent, and geometrical volume of the blob. After a normalization step, to make them independent of the scale level, they are integrated along the scale, during the lifetime (Lindeberg, 1993b) of the corresponding scale-space blob. Precise details about the measurements and their normalization can be found in Coulon *et al.* (1997b). In Lindeberg *et al.* (1999) the authors also study the influence of various ways to perform the




FIG. 7. Direct links between two primal sketches  $P_1$  and  $P_2$ .

normalization. This results in a set of multiscale measurements for each scale-space blob. The primal sketch is then composed of scale-space blobs linked by bifurcations and the associated multiscale measurements.

It has been shown that scale-space blobs are a reliable and exhaustive way to describe objects of interest in activation maps (Coulon *et al.*, 1997a). However, measurements are not efficient enough to characterize which of these blobs represent actual activations, relevant to the experiment. We therefore need to look at several subjects to make the detection decision, and this is the aim of the process presented in the following sections.

### The Comparison Graph

From here, unless stated, “blob” will refer to “scale-space blob.” We aim at creating a *comparison graph* such that it contains the primal sketches of all the subjects involved in the analysis and such that this graph makes explicit the potential repetitions of a structure (i.e., a scale-space blob) over subjects. Indeed, repetition over subjects is the second criterion, along with intensity in the map, to decide whether an object is an activation or noise. The graph nodes are scale-space blobs of the primal sketches, and the graph links are built between blobs belonging to different primal sketches and which may represent the same structure (activation). At this point, we do not yet want to decide if these repetitions are real (i.e., associated with an activation repeated over subjects) or random. The detection decision is to be made at the next step. Therefore, the construction has to be exhaustive in the sense that no *real* repetition should be forgotten.

To compare different primal sketches we normalize them in a common referential frame, with usual normalization procedures (Ashburner and Friston, 1997). This is the only way we have, at the moment, to compare different subjects, but a longer term aim is to build the spatial referential using the subjects’ individual anatomy (Mangin *et al.*, 1995a) and a high-level description of this anatomy, in terms of landmarks and identified structures (Mangin *et al.*, 1995b). The com-

parison graph should be a convenient framework for this purpose.

The first criterion to link two blobs belonging to two different primal sketches is the overlap of their spatial support. If two blobs have this overlap, they may represent the same activation. We combine this spatial overlap criterion with a scale overlap (since scale-space blobs have a scale dimension), as far as the considered individual maps have the same inner scale. If these two criteria are fulfilled, we create a *direct* link  $[b_1-b_2]_{out1}$  between the two scale-space blobs  $b_1$  and  $b_2$  as illustrated by Fig. 7.

Since we want to introduce some flexibility in the position of activations, to overcome possible normalization problems, we have to allow blobs which are “close enough,” but without any spatial overlap, to be linked in the graph. Note for instance that very small focal activations may have no overlap. We therefore use the fact that information becomes more stable (but grosser) as scale increases: two blobs being linked by a direct link *out1* may not represent an activation but may roughly indicate an area in which there is a more focal activation. This situation suggests that at a finer scale there might be a blob that represents an activation, this blob being in the part of the primal sketch which is “under” the first blob. This experimental observation led us to define a second type of link, which we call *induced* links (or *out2* links for mathematical notations): if  $b_1$  and  $b_2$  are two blobs having no direct link between them, they have an induced link  $[b_1-b_2]_{out2}$  if they are “under” (in their primal sketches) two blobs  $c_1$  and  $c_2$  and if there is a direct link  $[c_1-c_2]_{out1}$ . Such links are created only when  $b_1$  and  $b_2$  overlap in the scale dimension, which reduces the combinational possibilities and keeps coherence in scale. Induced links are illustrated by Fig. 8.

We point out the fact that allowing blobs without spatial overlap to be linked is a key point of the process, since it is the solution for a greater flexibility in overcoming spatial normalization limitations. Notice also that the graph topology is not only determined by inter-primal-sketch links but also by intra-primal-sketch information. Particularly, we will see later that




FIG. 8. The links induced by the direct links of Fig. 7.

blobs belonging to the same primal sketch are considered neighbors in the graph.

### The Detection Model: Use of a Markov Random Field

Activation detection is performed using a labeling process that uses the intersubject comparison graph previously described. Our aim is to associate a positive label with each activation in the graph and a null label with structures of noninterest. An activation (i.e., a nonnull label) is associated with a localization and can therefore have only one occurrence in each of the individual primal sketches.

We present here the basic model used to perform the detection. This model is defined using the following set of rules:

- i. A blob representing an activation is likely to have high measurements.
- ii. Two blobs representing the same activation must be linked in the graph and have the same nonnull label.
- iii. Two blobs representing the same activation (same nonnull label) are likely to have spatial supports close to each other.
- iv. An activation is represented only once in each subject, i.e., one positive label must have only one occurrence per primal sketch.

One can see that these criteria define *local dependencies* only: the value of a blob label depends only on the characteristics of the blob and on the knowledge of the labels of its neighbors in the graph. In other words, the model specifies that *a blob represents an activation if it is endowed with high enough measurements in the map and if the corresponding structure is repeated often enough over the other subjects*. These contextual dependencies allow us to model the label field in the graph as a Markov random field (MRF). We briefly present in the following subsection how this Markovian hypothesis is used to perform the labeling in a Bayesian framework.

#### Maximum a Posteriori and Energy Minimization

We want to perform the labeling by maximizing the posterior probability  $P(X|Y)$ ,  $X$  being the *label field* on the graph and  $Y$  being the *data*, i.e., the characteristics of the blobs (measurements, scale) and the topology of the graph. In other words, we want to find the optimal labeling given the data. We follow a classical *maximum a posteriori* process. The Bayes rule,

$$P(X|Y) = \frac{P(Y|X) \cdot P(X)}{P(Y)},$$

tells us that this maximization is equivalent to maximizing the product  $P(Y|X) \cdot P(X)$ . The probability  $P(Y|X)$  expresses the measurements probability given

the labels, and the prior probability  $P(X)$  expresses the *a priori* knowledge (the model itself).

Considering the usual (and justified in our case) independence hypotheses,<sup>2</sup> one has  $P(Y|X) = \prod_{s=1}^N P(y_s|x_s)$ , where  $s = 1 \dots N$  is the  $N$  sites (blobs) in the graph and  $y_s$  and  $x_s$  are the realizations of  $Y$  and  $X$  at site  $s$ . Finally, if  $P(Y|X)$  is not null, which we will verify, one can write

$$P(Y|X) \propto \exp\left(-\sum_{s=1}^N V(y_s|x_s)\right) = \exp(-U(Y|X)),$$

where  $V(y_s|x_s)$  is a data-driven *potential* function.

The Markov hypothesis tells us that  $P(x_s|x_{S \setminus s}) = P(x_s|x_{\nu_s})$ , where  $S$  stands for the site set: the probability of realization of  $X$  on a site depends only on the field realization in a neighborhood  $\nu_s$  of this site. If the *positivity condition* is respected, i.e., no realization has a null probability ( $\forall X, P(X) > 0$ ), the Hammersley–Clifford theorem (Besag, 1974; Geman and Geman, 1984) tells us that the Markov field is equivalent to a Gibbs field, i.e., one can write


$$P(X) = \frac{\exp(-U(X))}{Z},$$

where  $Z$  is a normalization constant (the *partition function*).  $U(X)$  is an energy function and is defined as a sum of contextual potential functions  $V_c$  on the *cliques* in the graph. Cliques are sets of sites that are all neighbors in the graph. The *order* of the clique is the number of elements of this set. Thus, one has  $U(X) = \sum_{c \in C} V_c(X|_c)$ , where  $c \in C$  are the cliques in the graph. The maximization of the posterior probability  $P(X|Y)$  can then be written as the minimization of the following energy:

$$U(X|Y) = \sum_{s=1}^N V(y_s|x_s) + \sum_{c \in C} V_c(X|_c).$$

Such a representation is of great interest because we do not need to know the exact form of the label field distribution: the model is directly defined through the potential functions. If two sites are neighbors in the MRF graph, there is an interaction between their la-

<sup>2</sup> These hypotheses are: (a) Independence of the observations.  $P(Y|X) = \prod_s P(y_s|x_s)$ , i.e., the measurements on two different blobs are independent: true in every case except when the two blobs belong to the same primal sketch and share the same bifurcation. In this case, the dependency is complex and is a function of the surrounding noise and of the third blob involved in the common bifurcation. In practice, this case is shown not to disrupt the model. (b) Observations depend only on the label of the site at which they are made.  $P(y_s|X) = P(y_s|x_s)$ , i.e., the measurements on one blob depend only on this blob, which is true.



**FIG. 9.** The second-order inter-primal-sketches clique potential function.

bels, and the aim is to define properly these interactions with the  $V_c$  functions. Let us point out the fact that such a representation is not a strict probabilistic interpretation of the problem and should be seen as a convenient way to define the above-mentioned model.

More precisely, rule (i) is modeled using the data-driven potential  $V(y_s|x_s)$ , rules (ii) and (iii) are modeled using potential functions on second-order cliques, and rule (iv) appears in potential functions on cliques defined over individual primal sketches. The minimization of the energy defined by these functions is then a competition process between configurations in the graph decreasing the energy and those increasing the energy. The potential functions must then express which are the suitable configurations we encourage (associated with negative potentials) and which is the one we want to avoid (associated with positive potentials). We present in the following subsections our choice for these potential functions.

### Second-Order Clique Potential Functions

Rules (ii) and (iii) concern blobs which do not belong to the same primal sketch and are linked in the graph. A set of two such blobs defines a second-order clique in the graph. We then express the two rules through a potential defined on second-order interprimal sketch clique,  $C_2 \subset C$ . This potential function is particularly important since it aims at giving the same label to blobs being neighbors in the graph. The idea here is to measure, for two neighboring blobs in the graph having the same positive label, how similar these two blobs are. The more similar they are, the lowest the potential function is, and then the more it encourages them to have the same positive label. We therefore define the potential  $V_{C_2}(b_1, b_2)$  between two blobs  $b_1$  and  $b_2$ , with labels  $I_{b_1}$  and  $I_{b_2}$ , belonging to different primal sketches and being linked in the graph by an *out1* or *out2* link:

- if  $I_{b_1} > 0, I_{b_2} > 0, I_{b_1} = I_{b_2}$ ,
- $[b_1-b_2]_{out1} \Rightarrow V_{C_2} = -K_{out1} \cdot ((e^{-f(b_1, b_2)} - 1)/(e^{-1} - 1)) - K_{out2}$ ,
- $[b_1-b_2]_{out2} \Rightarrow V_{C_2} = -K_{out2} \cdot e^{-f(b_1, b_2)}$ ,
- else  $V_{C_2} = 0$ ,

where  $I_b$  stands for the label of blob  $b$ ,  $K_{out1}$  and  $K_{out2}$  for constant weights of the potential, and  $f(b_1, b_2)$  for a *similarity function*. The potential function  $V_{C_2}$  is shown in Fig. 9 for the case  $I_{b_1} > 0, I_{b_2} > 0, I_{b_1} = I_{b_2}$ .

The similarity function  $f(b_1, b_2)$  is the second way, in the whole process, to overcome the limitation of spatial normalization. If one could identify any location in the brain according to individual anatomy, we would be able to tell how far  $b_1$  and  $b_2$  are from each other with respect to anatomical individual landmarks. Measuring this “anatomical distance” is the aim of the similarity function  $f$ , which would somehow define a new kind of spatial normalization. At the moment we use a more conventional type of measurement: if  $b_1$  and  $b_2$  are linked with an *out1* link,  $f(b_1, b_2)$  is defined by an overlap rate between 0 and 1, as follows,

$$f(b_1, b_2) = \frac{2 \cdot \#(b_1 \cap b_2)}{\#b_1 + \#b_2},$$

where  $\#b$  stands for the area of blob  $b$ . If  $b_1$  and  $b_2$  are linked with an *out2* link,  $f(b_1, b_2)$  is defined by a Euclidean distance function:


$$f(b_1, b_2) = \min[\|x_2 - x_1\|, x_1 \in b_1, x_2 \in b_2].$$

A longer term aim is to define a real anatomy-related similarity function.

### Intraprimal Sketch Maximal Clique Potential Functions

Rule (iv) is a strong condition on the label occurrences inside the graph. It is modeled through a potential function  $V_{C_{ps}}$ , defined over each primal sketch. This definition means that we consider primal sketches inside the graph as special kind of cliques,  $C_{ps} \subset C$ , which, as previously noticed, implies that the topology in the graph is such that two blobs belonging to the same primal sketch are neighbors. Figure 10 shows three different intra-primal-sketch configurations in the graph.

Figure 10a illustrates the ideal configuration with respect to constraint (iv): no activation is represented more than one time for a subject, i.e., no positive label



**FIG. 10.** Three different intra-primal-sketch configurations. Light gray, dark gray, and black represent three different nonnull labels.



has more than one occurrence in the primal sketch. We associate the potential  $V_{C_{ps}} = 0$  with this configuration. Figure 10b shows a limit case that we tolerate in the graph: an activation having one occurrence within a subject's map can be represented by two focal components within another map. This is tolerated if these two components are close enough in the graph, which for us means that the two concerned blobs are "under" the same bifurcation. Therefore, we associate the potential  $V_{C_{ps}} = 0$  with this configuration as well. Figure 10c represents the nonsuitable configuration: a positive label has more than one occurrence in the graph (and we are not in the situation described by Fig. 10b). The potential associated with this situation is then positive and defined by

$$V_{C_{ps}} = N_{ps} \cdot K_{ps} \cdot \sum_{l \neq 0, l \in L} n_l$$

where  $N_{ps}$  is the number of primal sketches involved in the process,  $K_{ps}$  the weight of this potential,  $n_l$  the number of occurrences in the primal sketch of any nonnull label  $l$ , and  $L$  the set of labels concerned by the configuration described in Fig. 10b. The farther away we are from the ideal configuration, the higher  $V_{C_{ps}}$  is. This potential function is made proportional to the number of primal sketches involved in the process to keep the whole analysis independent of this parameter. The influence of function  $V_{C_2}$  is implicitly proportional to the average number of neighbors of a blob in the graph which have the same label as this blob, and this average number increases with  $N_{ps}$  in a linear way. We therefore balance  $V_{C_{ps}}$  and  $V_{C_2}$ .

### Data-Driven Potential Functions

We then have to model constraint (i), that is to say the consistency between the label given to a blob and the measurements associated with this blob. This is done with the data-driven potential function  $V(y_s|x_s)$ , where  $x_s$  stands for the label at site  $s$  and  $y_s$  the measurements associated with the blob at this site. In our case, we simply define it with  $y_s$  being the measurement relative to the maximum intensity (normalized and integrated as mentioned under "The Primal Sketch.") Since we have no *a priori* about measurements of blobs representing noise, we define  $V(y_s|0) = 0$ .

Constraint (i) tells that if the measurements of a blob are low this blob has few chances to represent an activation, and if the measurements are high the blob is likely to represent an activation. We then define two values  $y_{low}$  and  $y_{high}$  such that

- $y_s < y_{low} \Rightarrow V(y_s|x_s \neq 0) = N_{ps} \cdot K_d$ ,
- $y_s > y_{high} \Rightarrow V(y_s|x_s \neq 0) = 0$ ,
- $y_{low} \leq y_s \leq y_{high} \Rightarrow V(y_s|x_s \neq 0) = (N_{ps} \cdot K_d) \cdot (y_{low} - y_{high}) \cdot (y_s - y_{high})$ ,




FIG. 11. The data-driven potential function.

where  $K_d$  is a constant weight. This function is illustrated by Fig. 11. It is proportional to  $N_{ps}$  for reasons similar to those already explained for  $V_{C_{ps}}$ .

This potential function, in particular  $y_{low}$  and  $y_{high}$ , is empirically deduced from measurement observation, but will be directly derived from the distribution of the considered measurements over populations of scale-space blobs representing activations and noise. Furthermore, we use only the maximum intensity measurements as a first definition, but we aim at taking into account several types of measurements to define the data-driven potentials. A brief study of what measurements represent the data the best is presented in Coulon *et al.* (1997a).

### Minimization

Once all potential functions are defined, the total energy function is defined. To minimize it, we use a stochastic algorithm, the *Gibbs sampler with annealing* (see Geman and Geman, 1984, Section XI, for details on the algorithm), whose convergence has been proven in the case of Gibbs random fields (Geman and Geman, 1984).

After minimization, the process gives a set of positive labels, each one representing an activation and having an occurrence in a number of primal sketches. We therefore know the occurrence, or the absence of occurrence, of each activation for any subject. This occurrence can then be mapped on the individual anatomy of the subject for localization considerations.

## EVALUATION AND RESULTS

In order to assess the performance of the system we simulated individual activation maps. These activation maps were of two sorts: pure noise to evaluate a rate of false positive and noise plus activations to study the performance of the system in limit cases. The detection process has also been tested on experimental data from a PET experiment. All experiments have been done with the same set of parameters, given in Table 1. These parameters have been chosen empirically, first on the basis of the relative influence we wanted be-



**TABLE 1**  
Values of the Parameters

$y_{low}$	$y_{high}$	$K_d$	$K_{out1}$	$K_{out2}$	$K_{ps}$
2	8	0.3	1.8	0.5	1

tween all the different terms of the model, then with adjustment to have results as good as possible on the simulations. A single set of parameters is needed for both simulations and real data, and the model is very stable with variations of the parameters around the chosen values.

### Simulations

#### Pure Noise

Activation maps were simulated using white noise convoluted with a Gaussian kernel of variance  $\sigma = 0.72$ , which is equivalent to a full width at half-maximum (FWHM) of 2 pixels. Five hundred activation maps of size  $64 \times 64 \times 48$  voxels were simulated and divided into 50 groups of 10 maps. For each group we then computed the primal sketch of each map and performed the group analysis using the 10 primal sketches.

The number of detected false positive was then counted, with the number of occurrences (between 1 and 10) of these false positive. Ideally, since we are using pure noise images, we should not detect any activation. Over the 50 groups, the estimated average number of false positives per group was 2.03, with an average number of occurrences of 5.63 per false positive. The estimated average number of false positives per map was 1.138.

We define the local energy  $U_l$  associated with a positive label  $l$  as

$$U_l = \sum_{s \in S_l} V(y_s | x_s) + \sum_{c \in C_l} V_c(X|c),$$

where  $S_l$  is the set of sites having the label  $l$  and  $C_l$  the set of cliques including these sites.

It is interesting to look at the local energy associated with these false positives since it gives a relative measure of consistency between the activation and the model we propose: the lower the energy is, the more consistent it is. First, the average energy is  $-0.54$ , which is much higher than the energy associated with real activations in real data (see below). Second, we notice that some of the false positives have a positive energy. This shows that the energy minimization does not reach a global minimum. This comes from the definition of the energy which leads to a non-convex-energy landscape and from the fact that the theoretical

conditions of convergence of our stochastic minimization algorithm cannot be used practically (in particular the decrease rate of the temperature (Geman and Geman, 1984)). Nevertheless, an activation with positive energy can be removed and the resulting global configuration is a “better” minimum. By doing so, the average number of false positive per group falls to 1.32, with an average number of occurrences equal to 5.74. The average number of false positives per map is then 0.76.

What we measure here is not an error rate but an average number of misclassified objects (false positive) per map. (For information, a single map contains several hundreds of objects.) These average numbers of false positive are acceptable since the detected false positives are associated with high energy and relatively low number of occurrences. This information together with the knowledge of neuroscientists interpreting the results should lead to a discrimination of errors. Moreover, the presence of a real activation in the group will reduce the chances of having false positives in the neighborhood of this activation, due to interactions between the activation and other structures, both in the primal sketches and in the comparison graph. Finally, a better evaluation of the data-driven potential functions should improve these results.

#### Noise and Activations

Using 10 maps similar to the previous ones, we added two synthetic “activations” in each map. These activations are defined by a Gaussian function of standard deviation equal to 5. Their positions vary in each image around a fixed position according to

$$(x, y, z) = (x_0 + \epsilon \cdot V_m, y_0 + \epsilon \cdot V_m, z_0 + \epsilon \cdot V_m),$$

with  $(x_0, y_0, z_0)$  the reference position of the object,  $\epsilon$  a uniform law of amplitude  $[-1; 1]$ , and  $V_m$  the maximum variation around  $(x_0, y_0, z_0)$ .

The amplitude of each object can vary as well according to

$$I_{\max}(\text{object}) = k \cdot I_{\max}(\text{noise}),$$

with  $I_{\max}(\text{object})$  the maximum intensity of the object,  $I_{\max}(\text{noise})$  the maximum intensity of noise (measured in the images), and  $k$  the ratio between both.

The rate of false negatives has been studied with  $V_m$  and  $k$  varying, to test the resistance of the process to intersubject variations of amplitude and to low-intensity activations. The results are presented in Table 2. One can see that the process is resistant to low-intensity signal and important variations of localization, which was one of the aims of the method.

**TABLE 2**

Number of Occurrences ( $N_{occ}$ ) and False Negatives ( $N_{FN}$ ) of Simulated Activations Depending on Intensity and Position Variations of These Activations

$V_{max}$	$k$	$N_{occ}(\text{object1})$	$N_{occ}(\text{object2})$	$N_{FN}$
3	1.25	10	10	0
5	1.25	10	10	0
10	1.25	10	10	0
15	1.25	—	8	1
3	1	10	10	0
3	0.8	10	10	0
3	0.6	10	10	0

Furthermore, these results show the role that induced *out2* links play in the process: for high values of  $V_m$  such as 10 or 15, a lot of occurrences of the objects have no overlap from one map to another. Their matching in the graph is therefore performed only with *out2* links.

### Real Data

The process presented here has been tested on a PET motor protocol, including 10 subjects and 12 images per subject. For each subject, an individual statistical  $t$  map was first computed using the *SPM* software (Functional Imaging Laboratory, London; Friston *et al.*, 1995). A primal sketch was then built from each of the maps, and the 10 primal sketches were compared using the labeling process.

Three different conditions (four images per condition) were performed by the subjects: a rest condition ( $R$ ), a periodic auditory-cued right-hand movement ( $M_1$ ), and a self-paced right-hand movement ( $M_2$ ). Individual maps were made using the 12 images of each subject and looking at the contrast ( $M_1 - R$ ). A group analysis was performed using the *SPM* software and used as a reference to validate our results. Numerous activations were found at a very significant level in the group map in expected brain regions.

After the labeling process, 11 positive labels were found. Table 3 shows the identification of these detected objects, sorted by increasing associated energy. The more negative this energy is, the better the structure fits with the model. Individual occurrences of these activations are presented in the Appendix, together with the *SPM* group analysis. Expected activations were detected together with two false positives, both outside the brain and caused by border effects during the blob extraction. (They are therefore easy to eliminate and do not prevent the correct interpretation of the maps. An addition to the process can be to use a mask of the brain and to take into account during the analysis only the blobs which are located inside this mask.) Some of the activations (labels 17, 10, 5, 18) had

no occurrence in one or two of the individual maps, and we could check that these activations had very low intensity in the corresponding individual maps. A classical threshold on the individual maps yielded poor results, either too selective or too noisy, and this shows a crucial advantage of our process: the detection is processed for each subject taking into account not only the intensity in the map but also the knowledge of the other subjects' maps. On the contrary, the number of occurrences of labels 7, 14, and 9 is greater than the number of subjects. It means that these activations were detected as one connected component for most of the subjects and two connected components for some of them. This illustrates the situation described in Fig. 10b and the definition of potential function  $V_{c_{ps}}$ .

This study also showed that the local energy function associated with each label is a good measure of the consistency with the model. Particularly, when a positive label corresponds to no real activation, its local energy might be high enough to discriminate it from labels associated with real activations. In our case, the best activation had an energy of  $-44.34$  while the worst had an energy of  $-2.82$ . These values confirm the observation made about false positive in synthetic data (see "Simulations") having an average energy of  $-0.54$ , the minimum energy found over the 50 groups being  $-5.46$ . However, let us notice that the local energy is not an absolute measurement but should be considered a clue when one looks at the result of a group analysis.

Figure 12 shows the occurrence of the right-hand motor activation on four subjects, on a rendering of their individual anatomy (Mangin *et al.*, 1998). Their location is very coherent since they are all in the right-hand sensorimotor area in front of the higher knob of the left central sulcus (Yousry *et al.*, 1998). One can see

**TABLE 3**

Result of the Structural Analysis of PET Data

Label	Energy	No. of occurrences	Brain region
7	$-44.34$	14	R cerebellum
14	$-38.99$	15	L sensorimotor area (hand)
9	$-22.93$	11	L sensorimotor area (hand and SMA)
3	$-20.83$	10	L and R thalamus
20	$-10.95$	10	FP, outside brain
6	$-9.77$	10	R auditory cortex
17	$-6.53$	8	L posterior insula
10	$-6.30$	8	FP, outside brain
5	$-5.38$	9	R parietal cortex
13	$-3.31$	10	L cerebellum
18	$-2.82$	9	L SMA

*Note.* Activations detected are sorted by increasing associated energy, with their localization (FP means false positive).



**FIG. 12.** Four three-dimensional renderings of individual anatomies with the corresponding occurrences of the primary motor activation.

that the shape and extent vary across subjects, although the four corresponding blobs have been detected as representing the same activation (same positive label in the graph). This can be caused by the intersubject variability, but also by the choice of the scale we used to represent the activation. Indeed, the objects labeled during the process are multiscale objects (scale-space blobs) and the choice of the representation scale changes their size and shape. The detection process automatically chooses a range of scale (the interval during which the blob exist) but not a single scale of observation. Automatic scale selection is an important issue in multiscale methods, and it will be addressed in the near future. At the moment, for the display of the results, we use the criterion proposed by Lindeberg (1993a): for a scale-space blob the scale of display is selected at which the (normalized) gray-level blob measurement assumes its maximum over scale.

### Comments

In simulations the process was shown to be sensitive enough to detect structures with an important inter-subject variation of position. It was also shown to be resistant to poor signal-to-noise ratio. Faced with the intersubject variability, this good flexibility preserves

a sensible false-positive rate: although the matching constraints between subjects are loosened, the increasing combinatory possibilities do not lead to a high false-positive rate. Moreover, the difficulty of not having a theoretical control of the risk error is compensated by the explicit information available at the end of the detection: each detected activation is described by its occurrence on the individual anatomy of subjects, which makes the interpretation easier. The number of occurrences of these activations within the group, and above all their local energy, provides useful information about the relevance of the results.

The results on real data are more difficult to validate since we have no “ground truth.” Apart from obvious activations, should a difference with a standard group analysis be interpreted as an error or as an effect of intersubject variability? Further experiments should provide a better validation, especially with fMRI data, for which individual observations are more reliable.

### CONCLUSION

We have presented here a new method to analyze brain functional images that considers the detection at a structural level and permits a way back to individual results after detection over a group of subjects. The user then has an interpretation of the data which allow him to know:

- What are the group results?
- What is the behavior of each subject relative to these group results? Are there “subgroups,” e.g., six subjects who show an activation and four other subjects who do not?
- What is the exact localization of the group results at an individual level? Knowing the individual occurrence of an activation, the user can check how and where it fits in the anatomy of the subject. This would improve the localization and the interpretation of the results and possibly provide a better understanding of the correlation between anatomy and function.

The method uses the power and comprehensiveness of multiscale methods to describe image structure, by looking at their whole scale space without any *a priori* hypothesis about scales of interest and without any “coarse-to-fine” strategy. A major difference from classical methods is the comparison of several subjects at an object level instead of at a voxel level, which permits one to introduce higher level criteria for the analysis and which is a way to overcome the intersubject variability effects.

Spurious spatial normalization effects can be reduced by means of the comparison graph combined with an appropriate definition of similarity between blobs from different subjects. With this similarity function one can introduce more flexible rules to decide whether two blobs represent the same activation. For



instance, one could use individual anatomical or functional localization. We illustrated that by allowing the matching of structures that have no spatial overlap.

The process has been proved to be able to efficiently detect expected activations on a PET data set. It is promising for functional MRI studies, since fMRI provides much more reliable individual maps than PET. Further research still has to be undertaken to solve outstanding questions, particularly concerning the choice of the optimal scale used to *represent* (as opposed to detect) a scale-space blob. Detection and interpretation of the results are in fact not completely independent of this representation scale since the extent of the reported activations depends on it. Second, a precise evaluation of the data-driven potential functions still has to be further investigated. For the latter, it is important to provide data-driven potential functions directly and properly evaluated from the distribution of the measurements associated to the blobs. In that way, the detection will be more accurate, especially with fMRI for which we have a better confidence in each subject map.

Although it is difficult to relate the proposed analysis to standard statistical analyses, it is worth noting that there is some kind of analogy with analyses that use random effects linear models (Friston *et al.*, 1999), in the sense that activation detection is performed using a subject-by-subject variability rather than on a scan-by-scan variability. Note that our approach is not opposed to classical statistical methods, since it uses individual maps made with these methods. We advocate that both approaches are complementary: statistical methods present the great advantage of summarizing the information related to a particular question in a single map. This quality makes these statistical maps particularly interesting to analyze.

Finally, we would like to point out the fact that using a Markovian model for the detection allows the user of such a system to interrogate the data in ways that can be designed according to what is looked for. It is very easy to define new potential functions in which one can introduce, for instance, *a priori* information about a precise expected location or about the search for a network of activations instead of isolated activation. That makes the system able to investigate functional data sets at a much higher level than what has been done so far.

## APPENDIX

We present here (Figs. 13–16) the group results from the SPM group analysis and the individual occurrences of the Markovian group analysis. All results are presented on a template 3D rendering from the SPM software. For examples of results on individual anatomy, see Fig. 12.

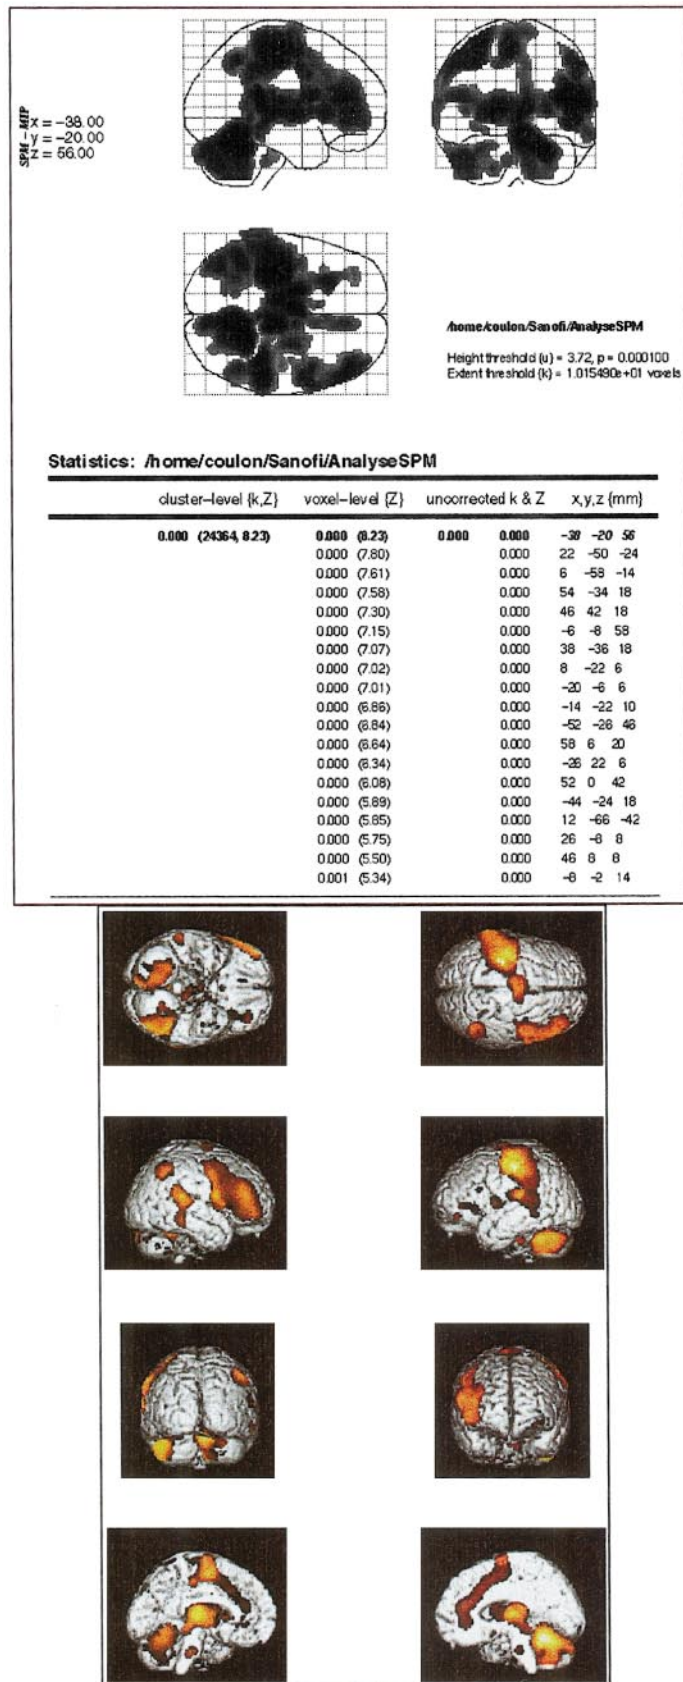
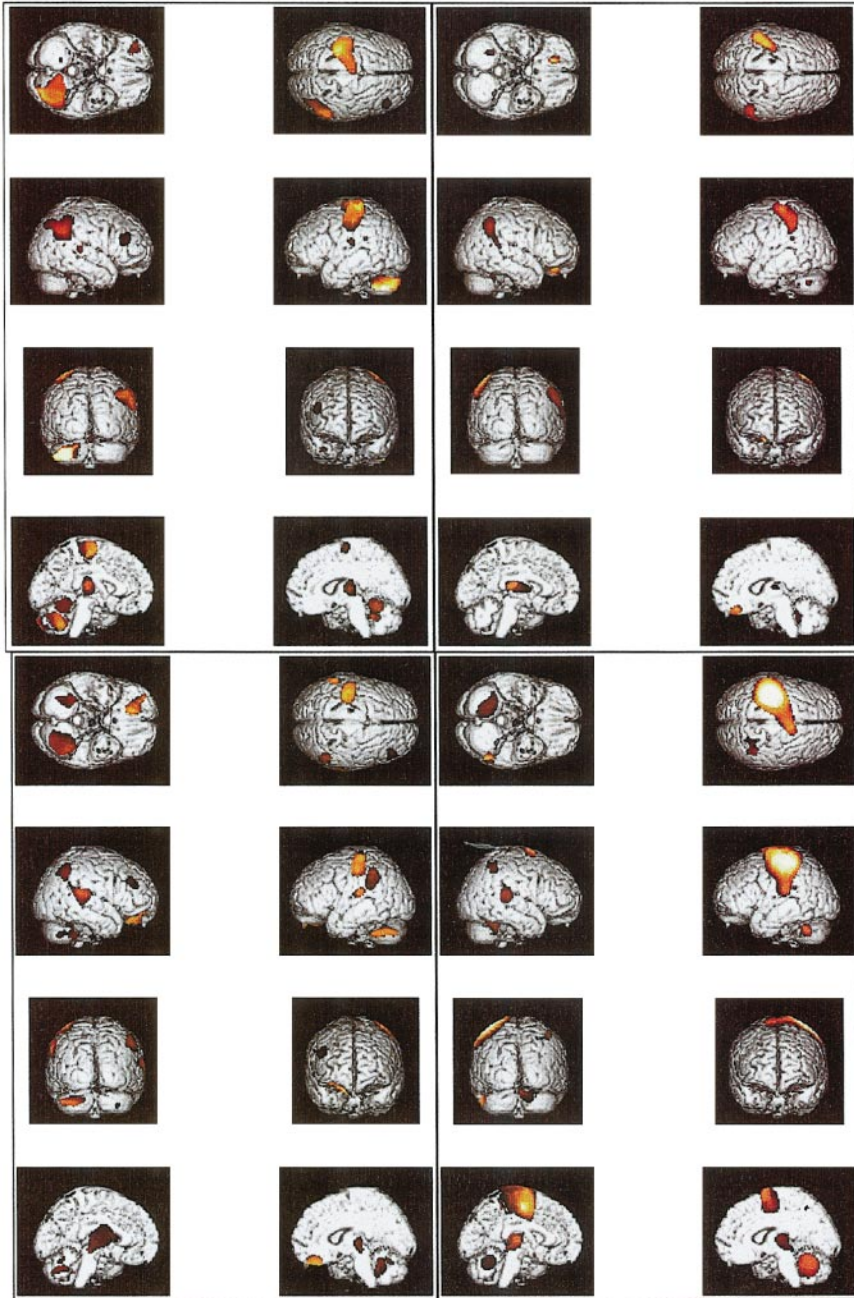



FIG. 13. Group results from the SPM group map.





**FIG. 14.** Individual occurrences for subjects 1, 2, 3, and 4.



**FIG. 15.** Individual occurrences for subjects 5, 6, 7, and 8.




FIG. 16. Individual occurrences for subjects 9 and 10.

## REFERENCES

- Alvarez, L., Guichard, F., Lions, P. L., and Morel, J.-M. 1993. Axioms and fundamental equations of image processing. *Arch. Rat. Mech.* **123**:199–257.
- Ashburner, J., and Friston, K. J. 1997. In *Human Brain Function* (R. S. J. Frackowiak, K. J. Friston, C. D. Frith, R. J. Dolan, and J. C. Mazziotta, Eds.), Chap. 3, pp. 43–58. Academic Press, San Diego.
- Besag, J. 1974. Spatial interaction and the statistical analysis of lattice systems. *J. R. Stat. Soc. B* **36**:192–236.
- Coulon, O., Bloch, I., Frouin, V., and Mangin, J.-F. 1997a. Multiscale measures in linear scale-space for characterizing cerebral functional activations in 3D PET difference images. In *First International Conference on Scale-Space Theory in Computer Vision* (B. ter Haar Romeny, L. Florack, J. Koenderink, and M. Viergever, Eds.), No. 1252 in LNCS, Utrecht, pp. 188–199. Springer-Verlag, Berlin.
- Coulon, O., Bloch, I., Frouin, V., and Mangin, J.-F. 1997b. Analysis of PET images using 3D scale-space primal sketch. Technical Rep. 97 D 001. Images Department, Ecole Nationale Supérieure des Télécommunications, Paris.
- Deriche, R. 1993. Recursively implementing the Gaussian and its derivative. Technical Rep. RR-1893. INRIA, Sophia-Antipolis.
- Fidrich, M. 1997. Following feature lines across scales. In *First International Conference on Scale-Space Theory in Computer Vision* (B. ter Haar Romeny, L. Florack, J. Koenderink, and M. Viergever, Eds.), No. 1252 in LNCS, Utrecht, pp. 140–151. Springer-Verlag, Berlin.
- Florack, L., ter Haar Romeny, B. M., Koenderink, J. J., and Viergever, M. 1992. Scale-space and the differential structure of images. *Image Vision Comput.* **10**:376–388.
- Friston, K. J., Holmes, A. P., Poline, J.-B., Worsley, K. J., Frith, C. D., and Frackowiak, R. S. J. 1995. Statistical parametric maps in functional imaging: A general linear approach. *Hum. Brain Mapp.* **2**:189–210.
- Friston, K. J., Holmes, A. P., Price, C. P., Büchel, C., and Worsley, K. J. 1999. Multi-subject fMRI studies and conjunction analyses. *NeuroImage* **10**:385–396.
- Geman, S., and Geman, D. 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* **6**:721–741.
- Koenderink, J. J. 1984. The structure of images. *Biol. Cybernet.* **50**:363–370.
- Koenderink, J. J., and van Doorn, A. J. 1986. Dynamic shape. *Biol. Cybernet.* **53**:383–396.
- Ledberg, A., Åkerman, S., and Roland, P. E. 1998. Estimation of probabilities of 3-D clusters in functional brain images. *NeuroImage* **8**:113–128.
- Lifshitz, L., and Pizer, S. 1990. A multiresolution hierarchical approach to image segmentation based on intensity extrema. *IEEE Trans. Pattern Anal. Mach. Intell.* **12**:529–540.
- Lindeberg, T. 1992. Scale-space behaviour of local extrema and blobs. *J. Math. Imag. Vision* **1**:65–99.
- Lindeberg, T. 1993a. Detecting salient blob-like image structures and their scales with a scale-space primal sketch: A method for focus of attention. *Int. J. Comput. Vision* **11**:283–318.
- Lindeberg, T. 1993b. Effective scale: A natural unit for measuring scale-space lifetime. *IEEE Trans. Pattern Anal. Mach. Intell.* **15**:1068–1074.
- Lindeberg, T., and Eklundh, J.-O. 1992. Scale-space primal sketch: Construction and experiments. *Image Vision Comput.* **10**:3–18.
- Lindeberg, T., Lidberg, P., and Roland, P. E. 1999. Analysis of brain activation patterns using a 3-D scale-space primal sketch. *Hum. Brain Mapp.* **7**:166–194.
- Marr, D. 1976. Early processing of visual information. *Philos. Trans. R. Soc. B* **275**:483–524.
- Marr, D. 1982. *Vision*. Freeman, San Francisco.
- Mangin, J.-F., Régis, J., Bloch, I., Frouin, V., Samson, Y., and Lopez-Krahe, J. 1995a. A MRF based random graph modelling the human cortical topography. In *First International Conference on Computer Vision, Virtual Reality and Robotics in Medicine, CVRMed'95*, No. 905 in LNCS, Nice, pp. 177–183. Springer-Verlag, Berlin.



- Mangin, J.-F., Frouin, V., Bloch, I., Régis, J., and Lopez-Krahe, J. 1995b. From 3D magnetic resonance images to structural representations of the cortex topography using topology preserving deformations. *J. Math. Imag. Vision* **5**:297.
- Mangin, J.-F., Coulon, O., and Frouin, V. 1998. Robust brain segmentation using histogram scale-space analysis and mathematical morphology. In *First International Conference on Medical Image Computing and Computer Assisted Intervention, MICCAI'98* (W. M. Wells, A. Colchester, and S. Delp, Eds.), No. 1496 in LNCS, Cambridge, MA, pp. 1230–1241. Springer-Verlag, Berlin.
- McIntosh, A., Bookstein, F., Haxby, J., and Grady, C. 1996. Spatial pattern analysis of functional brain images using partial least squares. *NeuroImage* **3**:143–157.
- Morse, B., Pizer, S., and Fritsch, D. 1994. Robust object representation through object-relevant use of scale. In *SPIE'94 Image Processing*, Vol. 2167, pp. 104–115, Newport Beach, CA, USA, SPIE Publisher.
- Nielsen, M., Florack, L., and Deriche, R. 1994. Regularisation and scale-space. Technical Rep. RR-2352. INRIA, Rocquencourt.
- Pizer, S. M., Koenderink, J., Lifshitz, L., Helmink, L., and Kaasjager, A. 1986. An image description for objects definition, based on extremal regions in the stack. In *Information Processing in Medical Imaging 1986*, Boston.
- Poline, J.-B., and Mazoyer, B. 1994a. Enhanced detection in brain activation maps using a multi-filtering approach. *J. Cereb. Blood Flow Metab.* **14**:639–642.
- Poline, J.-B., and Mazoyer, B. 1994b. Analysis of individual brain activation maps using hierarchical description and multiscale detection. *IEEE Trans. Med. Imag.* **13**:702–710.
- Unser, M., Thévenaz, P., Lee, C., and Ruttimann, U. E. 1995. Registration and statistical analyses of PET images using the wavelet transform. *IEEE Eng. Med. Biol.* **14**:603–611.
- Vincken, K., Koster, A., and Viergever, M. 1997. Probabilistic multiscale image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **19**:109–120.
- Worsley, K. J., Marrett, S., Neelin, Neelin, P., Vandal, A. C., Friston, K. J., and Evans, A. C. 1996a. A unified statistical approach for determining significant signals in images of cerebral activation. *Hum. Brain Mapp.* **4**:58–73.
- Worsley, K. J., Marrett, S., Neelin, P., and Evans, A. C. 1996b. Searching scale-space for activations in PET images. *Hum. Brain Mapp.* **4**:74–90.
- Yousry, T. A., Schmid, U. D., Alkadhi, H., Schmidt, D., Peraud, A., Buettner, A., and Winkler, P. 1997. Localisation of the motor hand area to a knob on the precentral gyrus, a new landmark. *Brain* **120**:141–157.