

Analysis of Brain Activation Patterns Using a 3-D Scale-Space Primal Sketch

Tony Lindeberg¹, Pär Lidberg^{1,2} and Per E. Roland²

¹Computational Vision and Active Perception Laboratory (CVAP)
Dept. of Numerical Analysis and Computing Science,
KTH, S-100 44 Stockholm, Sweden
<http://www.nada.kth.se/~tony>
Email: tony@nada.kth.se

²Division of Human Brain Research, Dept. of Neuroscience
Karolinska Institute, S-171 77 Solna, Sweden
Email: Per.Roland@neuro.ki.se

Technical report ISRN KTH/NA/P-98/XX-SE from KTH.

*To appear in Human Brain Mapping. Earlier version presented in
Proc. 3rd International Conference on Functional Mapping of the
Human Brain Mapping, Copenhagen, Denmark, May 19-23, 1997.
Neuroimage, Vol. 5, No. 4, p. 393, 1997.*

Abstract

A fundamental problem in brain imaging concerns how to define functional areas consisting of neurons which are activated together as populations. We propose that this issue can be ideally addressed by a computer vision tool referred to as the *scale-space primal sketch*. This concept has the attractive properties that it allows for automatic and simultaneous extraction of the spatial extent and the significance of regions with locally high activity. In addition, a hierarchical nested tree structure of activated regions and subregions is obtained.

The subject in this article is to show how the scale-space primal sketch can be used for automatic determination of the spatial extent and the significance of rCBF changes. Experiments show the result of applying this approach to functional PET data, including a preliminary comparison with two more traditional clustering techniques. Compared to previous approaches, the method overcomes the limitations of performing the analysis at a single scale or assuming specific models of the data.

Keywords: brain activation, human brain mapping, functional region, scale-space, primal sketch, scale selection, blob detection, multi-scale representation, computer vision

Contents

1	Introduction	1
2	Theory	2
2.1	Scale-space theory: Brief review	2
2.2	The scale-space primal sketch: A brief review	3
2.3	Proposed tool for analysing brain activation patterns	11
3	Methods	12
3.1	Choice of reference data for normalization	12
3.2	Brain activation data	14
4	Results	15
4.1	Experiments with synthetic image data	15
4.2	Experiments with 3-D PET images	18
5	Summary and discussion	33
A	Appendix	36
A.1	Definition of grey-level blob	36
A.2	Definition of scale-space blob	37
A.3	Scale-space blob events	38
A.4	Scale-space for discrete signals	39

1 Introduction

In functional brain imaging, one studies how activity patterns arise in the brain when subjects perform different types of tasks. A main purpose is to characterize functional contributions of local neuronal assemblies, based on the tight spatial coupling between regional synaptic activity and regional cerebral blood flow. The most common techniques today consist of measuring the cerebral blood flow by PET analysis and by measuring the blood oxygenation level detection (BOLD) by functional NMR. In this way, three-dimensional image data are obtained, in which the voxel values represent the degree of activity in corresponding regions.

A main assumption, which this overall approach relies upon, is that neurons in the cerebral cortex change their biochemical activity in large distinct populations. (This has been formalized as the cortical field activation hypothesis (Roland 1993).) To determine automatically from a given PET, SPECT or fNMR image what should be regarded as distinct populations of nerve cells (functional regions) has, however, turned out to be a non-trivial problem. The high noise level in the data results in complex hierarchical structures of blobs of different size and amplitude, in which a key problem concerns how to determine which blobs should be regarded as significant and which ones are due to noise.

Traditional approaches for analysing such data sets involve the computation of statistical measures, thresholding and the detection of peaks or significant clusters in the statistical parametric images so obtained (Friston et al. 1991, Friston et al. 1994, Worsley et al. 1992, Roland et al. 1993, Poline & Mazoyer 1993, Poline & Mazoyer 1994, Friston et al. 1996, Frackowiak et al. 1996, Holmes et al. 1996, Ledberg et al. 1998). Often these processing steps are combined with a spatial smoothing step to suppress high frequency noise and other rapid intensity variations or to compensate for anatomical misregistration (Poline & Mazoyer 1994, Worsley et al. 1996a, Worsley et al. 1996b).

The need for spatial smoothing that arises here is closely related to a more general multi-scale nature of real world data, which originates from the fact that objects in the world may appear in different ways depending upon the scale of observation. This inherent problem of analysing image data has been a main subject of research during the last decades. Specifically, the approach of *multi-scale analysis* has been developed, based on the idea that in the absence of any information about what scales are appropriate for describing a given data set, the only reasonable approach is to consider representations at *all scales* and to consider these scales *simultaneously* (Koenderink 1984, Lindeberg 1994).

Most current methods for analysing brain activation data on the other hand use a Gaussian filter of fixed size for smoothing. Such an approach favours blobs of size corresponding to the filter size, and may suppress other significant information in the image. (Poline & Mazoyer 1993, Poline & Mazoyer 1994) used a search over a range of Gaussian filter widths, but did not compute unified significance limits for the detected fields of blobs. A solution to this problem was proposed by (Worsley et al. 1996b), however, at the price of stronger assumptions for detecting significant peaks of activity.

The image structures we are interested in are regions of locally high activity cor-

responding to peaks, or bright blobs, in the intensity landscape. Interestingly, this problem of extracting of blob-like (*i.e.*, peak-like) structures from image data and selecting scales for these has been extensively studied in the field of computer vision, and a number of different approaches have been developed (Ehrich & Lai 1978, Blostein & Ahuja 1989, Lifshitz & Pizer 1990, Lindeberg 1993*a*, Lindeberg 1993*c*, Lindeberg 1994, Vincken et al. 1997). One such development, the scale-space primal sketch (Lindeberg 1993*a*, Lindeberg 1994), carries no assumptions about the data structure, but provides at the same time a unique detection of fields of increases and decreases at all scales, as well as a description of the extent and volume and hierarchical structure of the detected blobs and fields and finally a ranking of significant blobs and regions.

Concerning the usage of the “significance” in this article, it should, however, be emphasized that the purpose of this article is not to present any statistical theory for determining a significance level of the blobs in terms of a probability value for false positives, α . The word significance as used here should only be taken to mean significant in a scale-space sense (see section 2.2.2 for a precise definition). To develop a statistical model giving α of single blobs is a larger project beyond the scope of this article.

The subject of this article is to report the result of applying the scale-space primal sketch for hierarchical clustering of blob-like patterns in brain activation data. It will be shown that this computer vision tool has the attractive property that it allows for the extraction of significant blobs corresponding to functional fields, including the scales at which they occur, their spatial extent and their significance. Contrary to fixed-scale approaches, the scale-space primal sketch handles blob-like structures at all scales simultaneously and in a similar manner. Contrary to current statistical approaches, this concept does not make specific assumptions about the properties of the data.

The presentation will start with a brief review of two concepts the analysis will be based on — the notion of scale-space representation and a description of the scale-space primal sketch. Then, it will be shown how these tools are useful for automated analysis of brain activation patterns. An extensive expose of relations to previous work will be given later in section 5.

2 Theory

2.1 Scale-space theory: Brief review

To capture the abovementioned intrinsic multi-scale nature of image data, the notion of scale-space representation is essential. For any given D -dimensional signal $f: \mathbb{R}^D \rightarrow \mathbb{R}$, its scale-space representation $L: \mathbb{R}^D \times \mathbb{R}_+ \rightarrow \mathbb{R}$ is defined such that the representation at zero scale is equal to the original signal

$$L(\cdot; 0) = f(\cdot), \tag{1}$$

and the representations at coarser scales $t > 0$

$$L(\cdot; t) = g(\cdot; t) * f(\cdot) \tag{2}$$

are obtained by convolution with Gaussian kernels

$$g(x; t) = \frac{1}{(2\pi t)^{D/2}} e^{-x^T x/2t} \quad (3)$$

of increasing variance t . In terms of explicit integrals, the result of the convolution operation ' $*$ ' can with $x = (x_1, \dots, x_D)^T \in \mathbb{R}^D$ equivalently be written

$$L(x; t) = \int_{\xi \in \mathbb{R}^D} g(\xi; t) f(x - \xi) d\xi. \quad (4)$$

Within the class of possible multi-scale representations constructed by linear operations, it can be shown that this concept of scale-space representation has a very special role (Witkin 1983, Koenderink 1984, Babaud et al. 1986, Yuille & Poggio 1986, Koenderink & van Doorn 1987, Florack 1993, Lindeberg 1994). Based on natural symmetry requirements on the first stages of a vision system (referred to as scale-space axioms), several arguments can be presented showing that Gaussian kernels and their derivatives constitute a canonical choice when constructing a multi-scale representation.

A fundamental and characteristic property of this scale-space representation is that fine-scale structures will be successively suppressed when images are smoothed to gradually coarser scales in scale-space. An example of this effect is illustrated in the left column of figures 1(a)–(b), which show a detail of a brain activation pattern subject to various amounts of Gaussian blurring. In the right columns of these figures, a complementary illustration of this intensity is given by assigning an extremum region to each local maximum in the intensity pattern (the construction is based on the notion of grey-level blobs described next). As can be seen, the local maxima in the intensity landscape at fine scales mainly respond to noise, whereas the local maxima at coarser scales reflect image regions that we from an intuitive viewpoint may regard as natural clusters in the image data. A crucial problem in this context, however, concerns how to determine what structures should be regarded as significant and how to select appropriate scale levels for those. Specifically, there is a need to make such decisions automatically, without need for operator intervention. These are the main subjects that will be addressed next.

2.2 The scale-space primal sketch: A brief review

When computing a scale-space representation from an image, the information is only implicit in the intensity values. There is no explicit information about primitive image structures or the relations between image structures at different scales. The aim of the scale-space primal sketch is to build a representation which makes this information explicit, that is directly available, with emphasis on blob-like image structures.

2.2.1 Grey-level blob

The scale-space primal sketch is a tree-like representation of blobs (local extrema with extent) at *all* scales in scale-space. The most basic image primitive in this representation is a grey-level structure referred to as a grey-level blob. At each scale level, a (bright) grey-level blob is defined from each local maximum in the scale-space

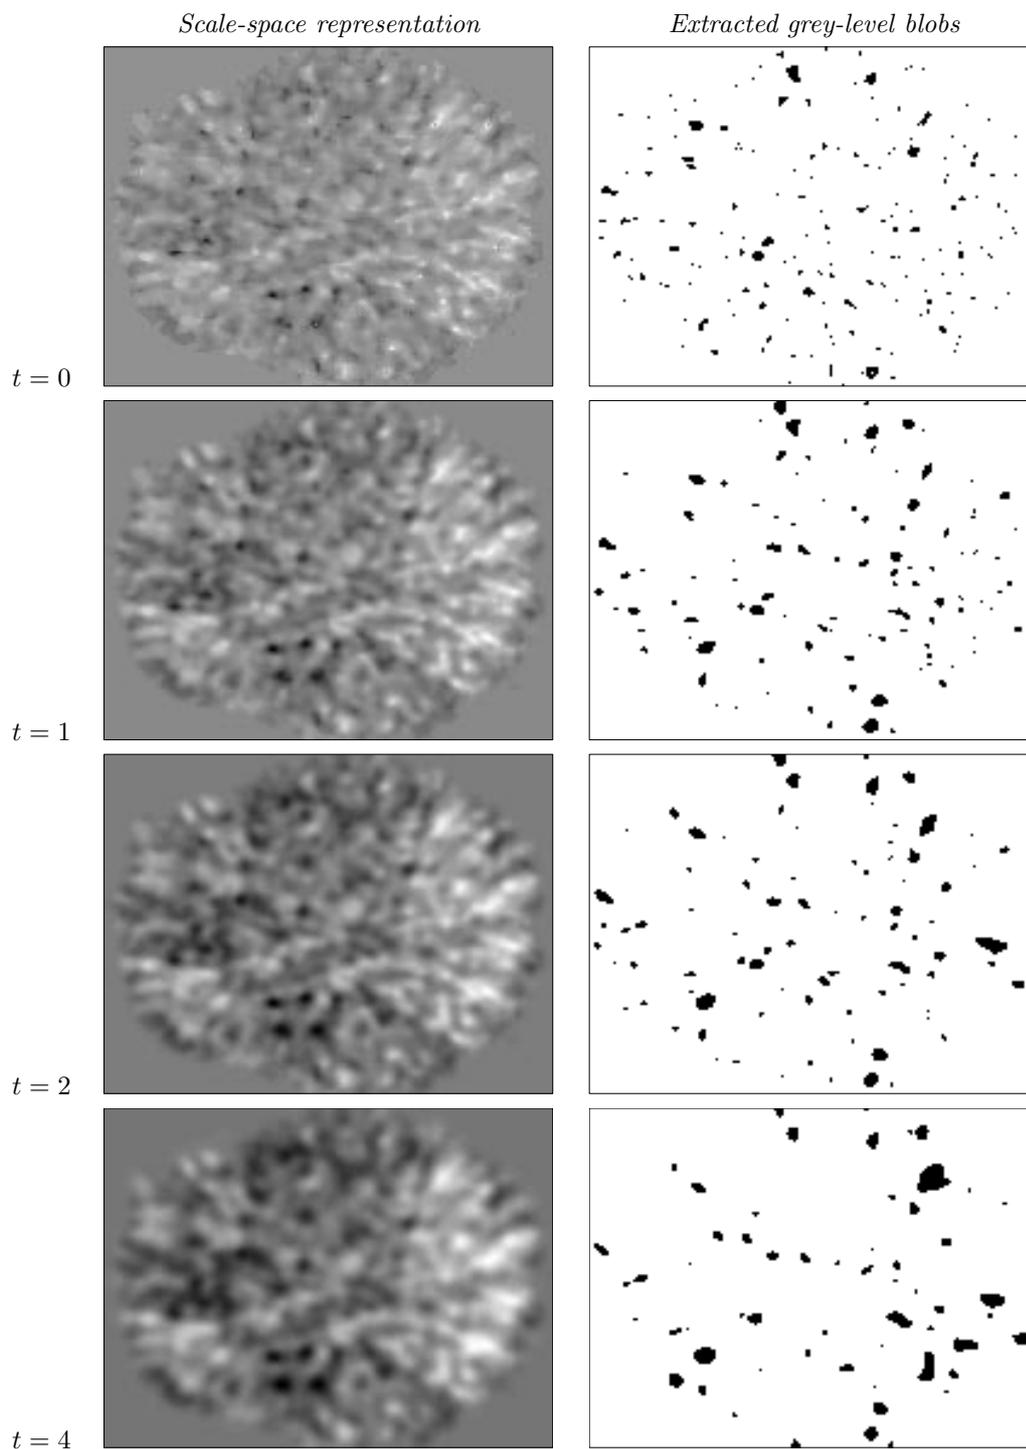


Figure 1: (a): (left column) Scale-space representation of a brain activation patterns at scale levels $t = 0, 1, 2$ and 4 . (right column) The result of extracting grey-level blobs from the image data in the left column. Essentially each dark region corresponds to a local maximum in the smoothed intensity landscape.

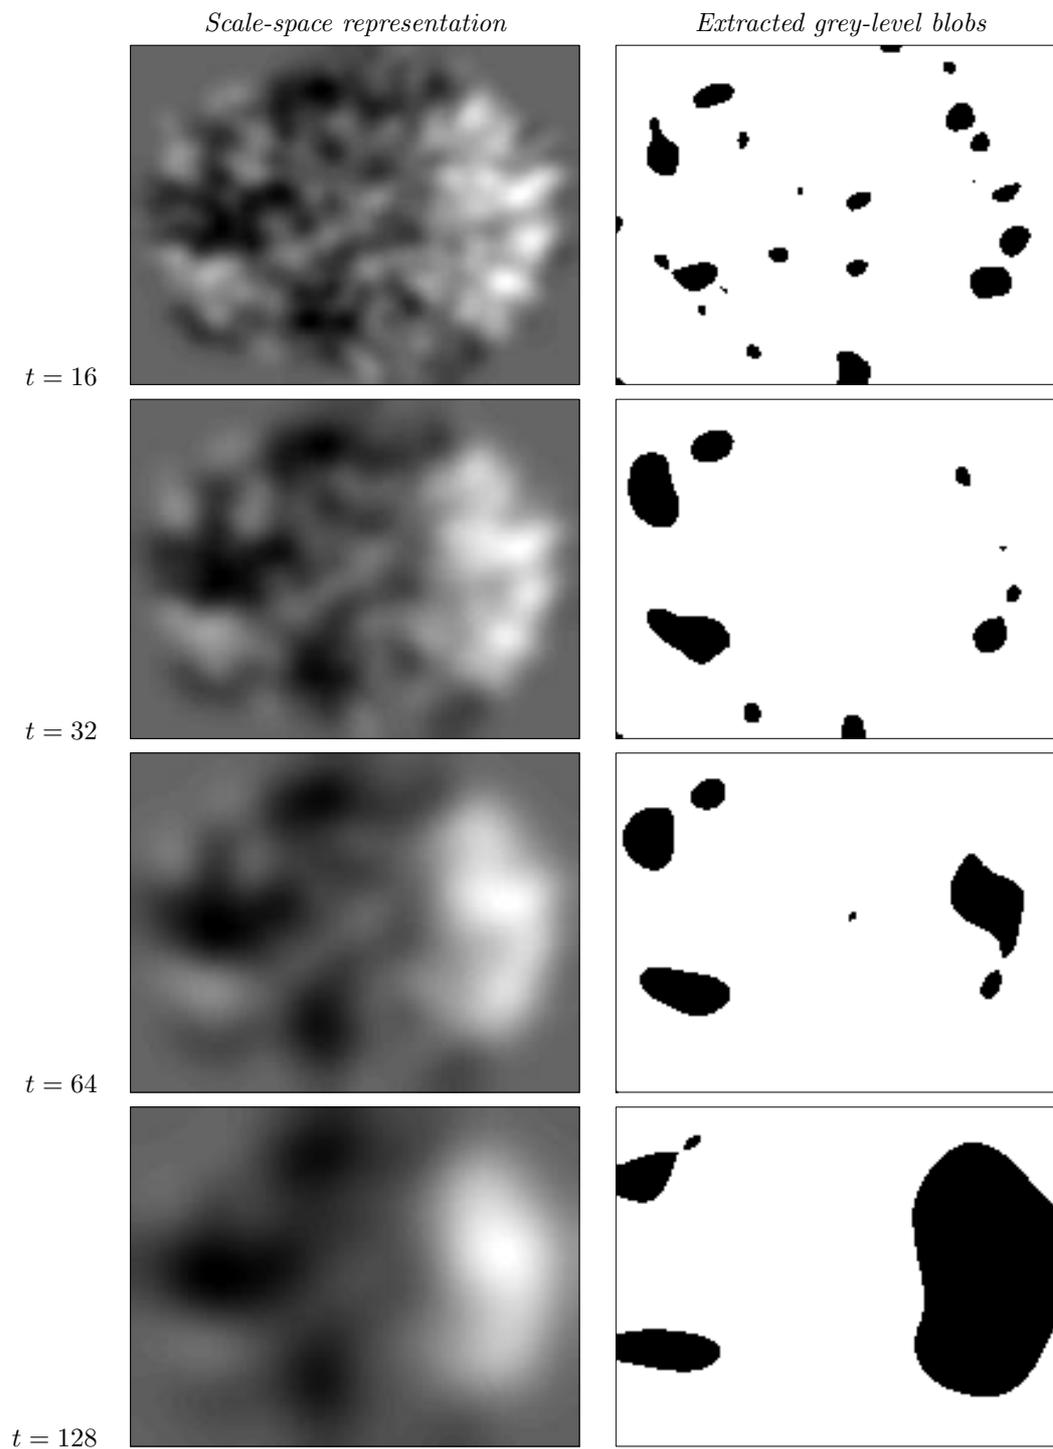


Figure 1:(b): (left column) Scale-space representation of a brain activation patterns at scale levels $t = 16, 32, 64$ and 128 . (right column) The result of extracting grey-level blobs from the image data in the left column. Essentially each dark region corresponds to a local maximum in the smoothed intensity landscape.

representation at that scale, and is associated with a spatial region with its extent delimited from the level curve through its *delimiting saddle point* as illustrated in figure 2 for the two-dimensional case.

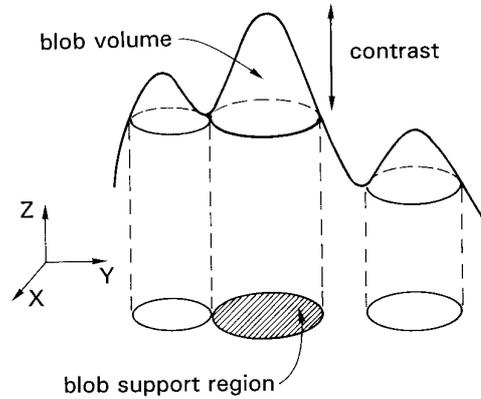


Figure 2: Illustration of the definition of a grey-level blob in the two-dimensional case.

Intuitively, this construction can be visualized out as follows: Imagine the image function as a flooded intensity landscape. If the water level sinks gradually, peaks will appear. At some instances two different peaks become connected. The corresponding elevation levels or intensity levels are called the *base-levels* of the blobs, and are used for delimiting the spatial extent of the blobs. The *support region* of the blob is defined to consist of those points that have an intensity exceeding the base-level and can be reached from the local maximum point without descending below the base-level of the blob. Finally, the grey-level blob consists of the (here, three-dimensional) volume bounded by the intensity surface and the level surface through the base level.

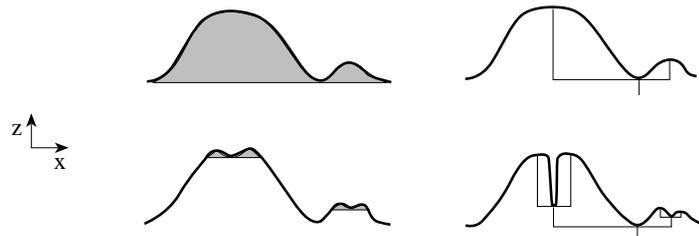


Figure 3: (left) A high-contrast large peak with two superimposed low-contrast fine scale peaks will not be detected as a grey-level blob if the signal is considered at one scale only. (right) A single noise spike can also substantially affect the relational tree.

By proceeding this water level analogy further, the successive merges of grey-level blobs leads to a *grey-level blob tree* with hierarchical relations defined between blobs of different size and amplitude (Ehrlich & Lai 1978, Lindeberg 1994); see figure 3. b Corresponding constructions can be carried out for signals of arbitrary dimension, and for a D -dimensional signal, the grey-level blob will be a $(D + 1)$ -dimensional entity

with extent in the D spatial dimensions as well the intensity dimension. A precise mathematical definition of this concept, including the case with discrete data, is given in appendix A.1. An algorithmic description of how to extract grey-level blobs from image data can be found in section 9.1 in (Lindeberg 1994).

2.2.2 Scale-space blob

The concept of a grey-level blob at a single scale is extremely noise sensitive, as is the detection of local extrema. For example, after a distortion of a large peak by a few superimposed local extrema, the large scale peak will not be detected as one unit; only the fine scale blobs will be found. Also the grey-level blob tree will be noise sensitive when considered at a single level of scale, since the hierarchical relations between different blobs are determined directly by the intensity levels in the valleys of the original signal (see figure 3).

A powerful approach to obtain more stable image features is by examining the behaviour of grey-level blobs in scale-space. In general, for every grey-level blob existing at a certain scale in scale-space, there will be corresponding grey-level blobs at slightly coarser and finer scales. By *linking* such grey-level blobs *across scales*, one obtains *scale-space blobs*. In our case with three-dimensional image data, these scale-space blobs will be *five-dimensional* entities with extent in the three spatial dimensions, the intensity dimension as well as the scale dimension. An implicit assumption behind this approach is that significance should be reflected by stability in scale-space.

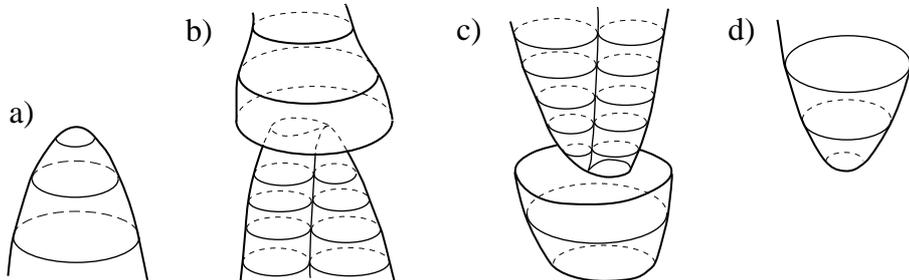


Figure 4: Generic blob events in scale-space: (a) annihilation, (b) merge, (c) split, (d) creation.

Formally, the linking is performed by considering trajectories of local maxima across scales, which are referred to as *extremum paths* (whose existence can be established by using the implicit function theorem; see section 8.2 in (Lindeberg 1994)). The extent of the scale-space blob in the scale direction is delimited by so-called catastrophes, where either the local maximum or the delimiting saddle point participate in a *bifurcation event*. By the use of singularity theory (Poston & Stewart 1978), it can be shown that there are only four possible types of generic bifurcation events (Lindeberg 1994, section 8.4)

- annihilation: a blob disappears,
- merge: two blobs merge into one,

- split: one blob splits into two,
- creation: a new blob appears.

In the scale-space primal sketch, these bifurcation events define hierarchical relations between scale-space blobs at different scales, as illustrated in figure 4.

A formal definition of a scale-space blob as well as a more detailed description of the bifurcation events can be found in the appendix.

2.2.3 Measuring the significance of scale-space blobs

The basic idea for extracting significant image structures of a D -dimensional signal using the scale-space primal sketch is to *rank* scale-space blobs on significance based on the *scale-space blob volumes* they occupy in $(D + 1)$ -dimensional (here, five-dimensional) scale-space. The motivations for this approach are that the following components are natural to include in a significance measure (Lindeberg 1993a):

- **spatial extent** x : in the absence of further information, a grey-level blob having large spatial extent may be treated as more significant than a similar grey-level blob with smaller extent,
- **contrast** z : in the absence of further information, a grey-level blob having a high contrast may be treated as more significant than a similar grey-level blob with lower contrast,
- **lifetime** t : in the absence of further information, a grey-level blob having a long lifetime in scale-space may be treated as more significant than a similar grey-level blob having a shorter lifetime.

Then, for each scale-space blob, the scale is selected at which the (normalized) grey-level blob volume assumes its maximum over scales.

By construction, the relative ranking of blobs obtained in this way, in decreasing order of significance, will be invariant to translations, rotations and uniform rescalings in the spatial domain, as well as affine transformations of the intensity values.

In a similar way, the selected scale levels are invariant under translations and rotations of the spatial domain, as well as affine transformations of the intensity domain, whereas under rescalings of the spatial domain, the selected scale levels are transformed with a similar scaling factor as the spatial rescaling.

2.2.4 Normalization of the scale-space blob volume

When measuring the scale-space blob volume, it is of crucial importance that this volume is measured in appropriate units, such that the significance of image structures at different scales can be compared without bias to structures at fine or coarse scales.

Measuring the scale parameter. Concerning the scale dimension, a natural requirement is that the expected lifetime of a scale-space blob should not vary over scales. This leads to a transformation function of the form (Lindeberg 1993b)

$$\tau_{\text{eff}}(t) = C_1 + C_2 \log(p(t)) \quad (5)$$

where

$$p(t) = \{\text{expected density of extrema at scale } t\}, \quad (6)$$

and C_1 and C_2 are constants, and C_1 without loss of generality can be set to zero. For a large class of continuous signals, it holds that the number of local extrema decreases with scale approximately as $t^{-\alpha}$ ($\alpha > 0$). This gives

$$\tau_{\text{eff}}(t) = -C_2 \alpha \log(t), \quad (7)$$

where we can further without loss of generality set $C_2 = -1/\alpha$.

In practice all brain images are discrete and of finite extent. At fine scales, this leads to interference with the *inner scale* of the signal given by its sampling density, and at coarse scales there will be interference with the *outer scale* of the signal given by its finite size. Then, an appropriate transformation function for discrete signals can be obtained by

$$\tau_{\text{eff}}(t) = \log \left(\frac{p_{\text{ref}}(0)}{p_{\text{ref}}(t)} \right), \quad (8)$$

where $p_{\text{ref}}(t)$ denotes the average density of local extrema as function of scale t (Lindeberg 1993b). Figure 5 shows an example of a transformation function defined in this way, based on simulation results from three-dimensional white Gaussian noise. Alternative choices of reference data will be in section 3.1.

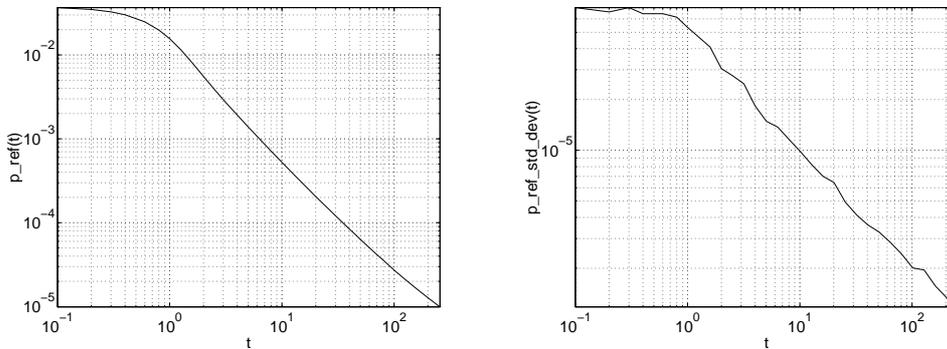


Figure 5: The density of local extrema as function of the scale parameter t accumulated for 220 white noise pictures of size $128 \times 128 \times 128$ pixels: (left) mean value $p_{\text{ref}}(t)$, (right) standard deviation $p_{\sigma}(t)$.

Measuring the grey-level blob volume. Similarly, the grey-level blob volumes need to be transformed. A straightforward approach to normalization is to subtract the expected mean value $V_m(t)$ from the grey-level blob volume G_{vol} and to divide by the expected standard deviation $V_{\sigma}(t)$:

$$V_{\text{prel}} = \frac{G_{\text{vol}} - V_m(t)}{V_{\sigma}(t)}. \quad (9)$$

Here, the reference functions $V_m(t)$ and $V_\sigma(t)$ should again be measured from reference data, which should have similar properties as the image data that is to be analyzed. Figure 6 shows experimental results based on white noise. Other choices of reference data are described in section 3.1.

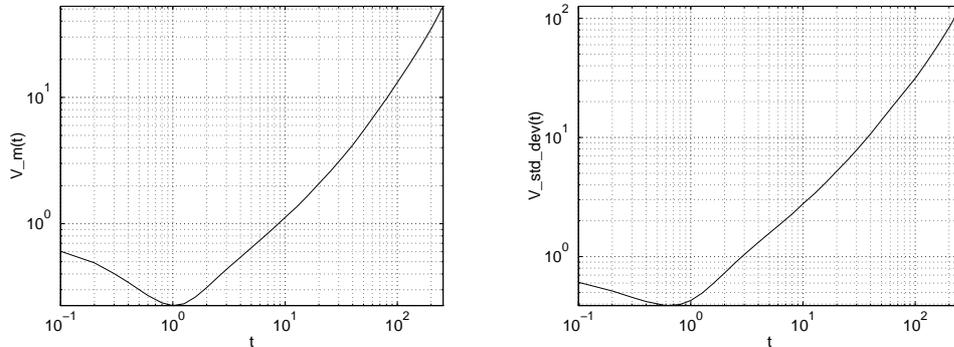


Figure 6: Expected grey-level blob volume as function of the scale parameter t as accumulated over 220 white noise pictures of size $128 \times 128 \times 128$: (left) mean value $V_m(t)$, (right) standard deviation $V_\sigma(t)$.

When computing the scale-space blob volume, however, it should be noted that the entity $V_{\text{prel}}(t)$ in (9) is not suitable for integration, since it can take negative values. A normalization that empirically has given reasonable results is to let

$$V_{\text{eff}} = \begin{cases} 1 + V_{\text{prel}} & \text{if } V_{\text{prel}} \geq 0, \\ e^{V_{\text{prel}}} & \text{otherwise.} \end{cases} \quad (10)$$

In order to adapt the amplitude of the signal to the reference data, V_m and V_σ are rescaled linearly from a least-squares fit between the actual and the expected behavior of these entities.

Computing the scale-space blob volume. Given the transformed entities τ_{eff} and V_{eff} , the scale-space blob volume is then defined by

$$S_{\text{vol}} = \int_{t_{\text{appear}}}^{t_{\text{disappear}}} V_{\text{eff}}(t) d\tau_{\text{eff}}(t) \quad (11)$$

where t_{appear} and $t_{\text{disappear}}$ are the appearance and disappearance scales of the scale-space blob, respectively.

2.2.5 Computing the scale-space primal sketch

When computing the scale-space primal sketch for a given image, the following algorithmic steps are involved:

- Sample a given scale interval uniformly in effective scale.
- Extract grey-level blobs at each scale.

- Match grey-level blobs between scales.
 - This matching is based on overlap and inclusion of extrema in the blob support regions. For generic image data, it can be shown that there are five types of matching relations (see figure 7). If this is not the case, then the scale sampling is refined and a new matching step is initiated.
- Build a hierarchical data structure based on the matching relations between blobs at different scales.
- Rank structures on significance based on the scale-space blob volumes, and select the scales at which the blob responses are maximal with respect to scale.

For a detailed algorithmic description, the reader is referred to (Lindeberg 1994, ch. 9).

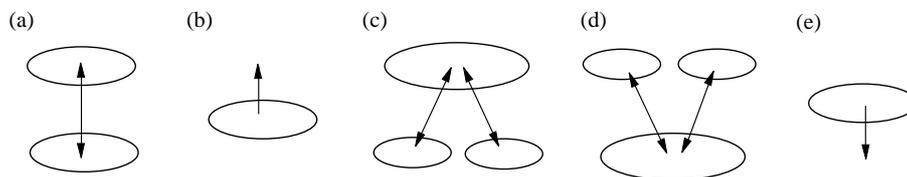


Figure 7: Elementary matching situations between blobs at adjacent scales; (a) plain link, (b) annihilation, (c) merge, (d) split, and (e) creation.

2.3 Proposed tool for analysing brain activation patterns

As an image representation, the scale-space primal sketch is computed under extremely weak conditions about the signal, and with no strong assumptions about the specific shape of the image features. There are essentially no parameters to tune, and the only parameter to adjust is how many regions should be returned or alternatively a threshold on the significance. The main assumption is that the image structures should be blob-like.

In earlier work (Lindeberg 1993a, Lindeberg 1994), it has been shown that the scale-space primal sketch allows coarse segmentations to be generated in the sense that the localization of boundaries can be poor due to the natural shape distortion occurring in scale-space. Such segmentations will, however, be safe in the sense that the regions, which are given by scale-space blobs with high significance, really are useful as landmarks for significant structures. Specifically, the scale-space blobs obtained from the scale-space primal sketch give information about:

- the approximate location and extent of relevant regions in the images,
- appropriate scales for treating these regions in further processing.

The main subject of this article is to show how such descriptors and the principles behind their construction are useful for analysing functional brain images.

It will be shown that scale-space blobs extracted from the scale-space primal sketch allows for simultaneous determination of the significance and the extent of

activated region. In addition, a hierarchical description is obtained of activated regions at multiple scales.

Support for the proposed approach will be presented by an experimental study on real PET experiments, as well as a comparison with two statistical techniques, showing that the highest ranked blobs in two examples correspond to intuitively reasonable and biologically meaningful regions. In contrast to other methods, this approach does not depend upon strong (unrealistic) statistical assumptions about the data.

3 Methods

3.1 Choice of reference data for normalization

When ranking image structures on significance based on the scale-space blob volume, it is of crucial issue that all scale levels are treated in a uniform manner. In the original work on the scale-space primal sketch (Lindeberg 1993a, Lindeberg 1994), the normalization was based on the conservative approach of using white noise reference data i.e., image data without correlation between disjunct image regions. The motivation for using this approach was to compensate for the rate to which accidental groupings may occur in scale-space.

For the image data we obtain from PET analysis (details of image acquisition will be given later), the situation is different in the sense that we can expect the input data to be highly correlated in the spatial domain. In particular, the common use of low-pass filtering as a pre-processing step to the analysis introduces a bias to structures at a certain range of scale. For this reason, it is of interest to consider other choices of image data for normalization, and in this work we will study the following approaches:

- uncorrelated white Gaussian noise as described previously,
- real difference images obtained from the same camera as the data to be analysed,
- real Student-t images obtained by summing up difference values and dividing by the standard deviation voxel by voxel,
- simulated Student-t images obtained by the following procedure:
 - the autocorrelation function is estimated from difference images from real PET experiments,
 - a linear filter is computed, such that convolution of white noise with this filter gives an image with a similar correlation function,
 - a Student-t image is obtained by summing up a fixed number of such noise images and dividing by the number of images as well as the standard deviation.

Figure 8 and figure 9 show calibration curves for the density of local extrema $p_{\text{ref}}(t)$ as well as the average grey-level blob volume $V_m(t)$ and its standard deviation $V_\sigma(t)$ computed at different scales.

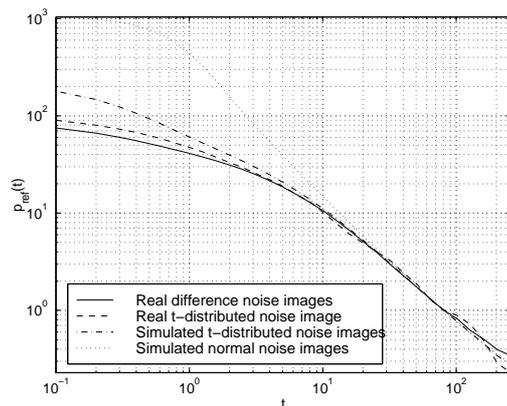


Figure 8: Experimental results showing the density of local extrema as function of (effective) scale, τ , for four types of image data.

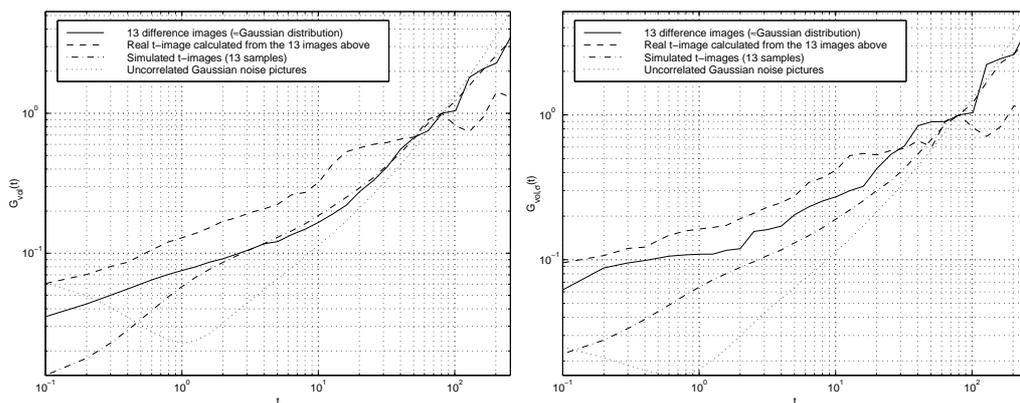


Figure 9: Experimental results showing the grey-level blob volume as function of (effective) scale, τ , for four different types of image data: (a) mean value $V_m(t)$, (b) standard deviation $V_\sigma(t)$.

Note the large difference between normalization based on white Gaussian noise and correlated PET data. It is of particular interest to observe the significant difference between the results from the real and the simulated t-images (see the dashed and the dot-dashed curves). Ideally, if the PET data would have been generated from a simultaneous normal distribution (implying that not only the individual voxels have a 1-D normal distribution, but any set of K voxels constitutes a K -dimensional normal distribution), then the shape of the auto-correlation function would have been sufficient to specify the statistical properties of the data. The difference in scale-space behaviour between the real and synthetic images on the other hand suggests that are limitations in using the autocorrelation function as the only statistical descriptor for simulation analysis as is the case for certain approaches.

3.2 Brain activation data

The image data we have used for normalizing and testing the scale-space primal sketch algorithm have been obtained from two data sets.

The first data set was obtained from an experiment of *tactile-visual matching* of shape in eight normal male volunteers (Hadjikhani & Roland 1998). The subjects matched ellipsoids drawn from two identical sets of 15 ellipsoids all having the same volume, surface properties and weight (Roland & Mortensen 1987). One ellipsoid was presented to the right hand and while this was tactually explored, the subject was looking at an ellipsoid which was either identical or slightly different in shape from the one felt by the hand. The task being to decide by a two-alternative false choice whether the ellipsoids were identical or different. A control condition was obtained also, during which the subject looked at the little shelf on which the ellipsoid was visually presented, but without any ellipsoid to see. During the control, the subjects moved their right hand as if they were exploring an ellipsoid, however no ellipsoid was present in the right hand. From these eight subjects, a Student' t statistical image was made showing the differences between the control and the tactile-visual matching.

In this first experiment, the regional cerebral blood flow (rCBF) was measured in 3-D acquisition mode with an ECAT-EXACT-HR PET scanner (Wienhard et al. 1994) with 15 mCi of ^{15}O -labelled butanol. The sinograms were reconstructed with a ramp filter having a cut-off frequency of 0.5 cycles, the reconstructed image was subsequently filtered with a 4.2 mm full width half maximum 3-D isotropic Gaussian filter. The rCBF was calculated by an autoradiographic procedure, based on the frames between 0 and 60 seconds (Meyer 1989). Individual magnetic resonance images of each of the subjects and the corresponding rCBF images were standardized anatomically by the Human Brain Atlas (Roland et al. 1994). To reduce the variance, the global blood flow was normalized to 50 ml/100 g/min. Each condition (of test and control) was repeated three times in each subject. By a pairwise subtraction of control and tactile visual matching, images for the first, second and third repetition and mean subtraction image of tactile-visual (TV) minus control rCBF were calculated. Subsequently, these intra-subject mean subtraction images were used to calculate a descriptive Student-t image (Roland 1993).

The second data set was from a *reaction time task* with nine participating male subjects. A 2 mm in diameter stylus suddenly protruded through a hole and indented the tip of the right index finger. This occurred at random intervals such that the subjects could not predict when exactly this was going to happen. Their task was to, as soon as possible, press a response key with help in the right hand. During the experiment, the subjects had their eyes open and fixated a small circular light source. The control was a rest during which the subjects had their eyes closed.

In the second experiment, the rCBF was measured with an eight ring (15-slice) PET camera (PC 2048-15B Scanditronics) having an in-plane spatial resolution of 4.5 mm and an interslice distance of 6.75 mm. 70 mCi of ^{15}O -butanol was injected, and the rCBF was calculated as before. The sinograms were reconstructed with a 4 mm Hanning filter, but no further filtering was done. The rCBF of rest condition was subtracted from that of the reaction time condition. The resulting subtraction images were anatomically standardised as described above. Subsequently, these intersubject

subtraction's images were used to calculate a descriptive Student t-image (Roland 1993).

In addition, a set of noise data was computed from the visual-tactile crossmodality experiment by permuting the matching repetitions images and subtracting the rCBF of the matching from the repetition, i.e. repetition 1 from repetition 2, etc. In this way, based on these subtraction images, 20 Student's t-images were computed, which were used as reference data for the abovementioned normalization of the scale-space data for the tactile-visual matching. Similarly, a set of noise data was computed from the reaction time experiment, by permuting the order of the reaction time images and the control images for subtraction. In this way, 20 Student's t-images were computed, which were used as reference for the normalization of the reaction time scale-space data.

4 Results

4.1 Experiments with synthetic image data

As a first illustration and as a validation of the approach, we applied the scale-space primal sketch to synthetic two-dimensional image data generated from the superposition of Gaussian blobs with randomly selected position, size and amplitude.

Figure 10(a) shows such a grey-level image with 10 Gaussian blobs to which white Gaussian noise has been added. In figure 10(b) the 10 strongest blob responses are displayed after suppression of multiple responses to the same Gaussian blob. Each blob is represented by the boundary of the blob support region at the representative scale of the scale-space blob, and the boundaries of the 10 strongest blobs are superimposed on a bright copy of the grey-level image. Figure 10(c) shows the intensity data before the noise was been added, and in figure 10(d) the blob boundaries have been superimposed on this data set.

Table 1 lists basic data about the 10 synthetic Gaussian blobs, including their position, size (given as the scale parameter of the Gaussian blob), and volume (the mass of the Gaussian blob). As can be seen from this table, the six Gaussian blobs with indices 1–6 have grey-level volumes (mass) significantly larger than the grey-level volumes of the four other ones. These blobs are also the blobs that stand out visually, after the noise has been added. Table 2 lists corresponding data for the 10 most significant scale-space blobs extracted by the scale-space primal sketch, after the suppression of multiple scale-space blob responses to the same Gaussian blob.

As can be seen, the six most dominant Gaussian blobs in the grey-level image give rise to corresponding scale-space blobs, which are all among the six strongest scale-space blob responses. The ranking of the four other blobs, of smaller size, volume and subjective visibility, is on the other hand strongly influenced by spurious structures in the noise pattern. For the six dominant blobs, we see that the spatial extents of the support regions of the scale-space blobs conform quite well to the visual extents of the superimposed Gaussian blobs, and that the extremum points of the grey-level blobs correspond well to the center points of the Gaussian blobs. In these respects, the qualitative behaviour of the method is intuitively reasonable.

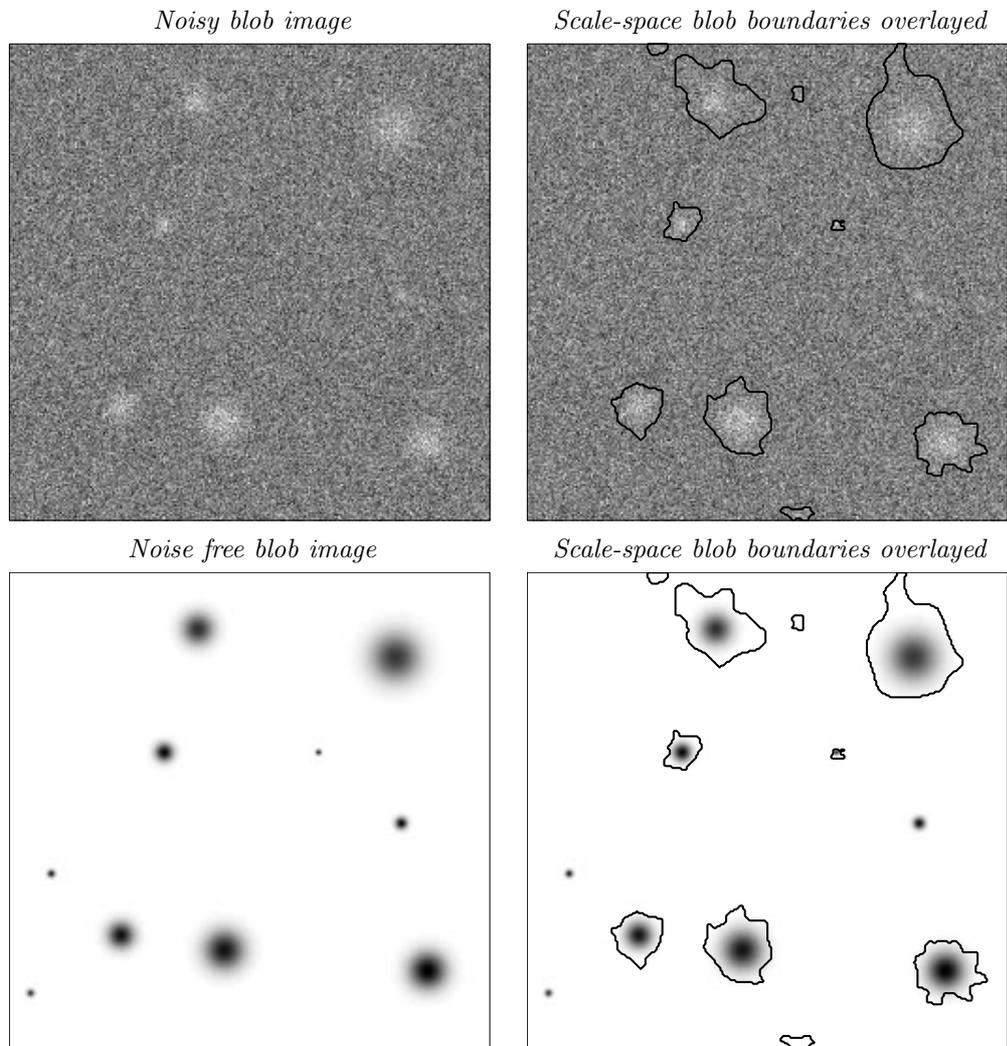


Figure 10: The result of applying the scale-space primal sketch to a synthetic two-dimensional image with 10 Gaussian blobs (with randomly selected position, size and amplitude) to which white Gaussian noise has been added. (a) grey-level image, (b) boundaries of the 10 most significant blobs overlaid on a bright copy of the grey-level image, (c) grey-level image without noise, (d) boundaries of the 10 most significant blobs overlaid on a bright copy of the grey-level image before the noise was added.

Synthetic Gaussian blobs				
Index	Center	Size	Volume	Ranking
1	205, 210	79.2	767.1	5
2	114, 53	51.7	594.1	2
3	222, 42	45.6	568.5	1
4	100, 225	33.4	334.7	6
5	59, 61	23.2	267.8	3
6	82, 159	11.0	133.0	4
7	208, 121	4.8	57.5	>10
8	22, 94	1.9	21.1	>10
9	11, 30	1.8	18.1	>10
10	164, 159	1.1	11.6	9

Table 1: Positions, scale values and grey-level volumes of the 10 synthetic Gaussian blobs in figure 10, sorted in decreasing order of their grey-level volumes. The last columns gives the significance ranking of the blob, as induced by the scale-space blob volume of a corresponding scale-space blob detected by the scale-space primal sketch.

Scale-space blobs from synthetic Gaussian blob image					
Ranking	Significance	Scale	Area	Extremum	Index
1	23.0	5.9	871	222, 41	3
2	20.3	8.9	988	115, 54	2
3	14.8	8.9	521	59, 61	5
4	10.0	4.2	264	82, 159	6
5	9.9	29.3	2097	205, 210	1
6	9.9	18.2	1264	99, 225	4
7	5.8	2.0	54	69, 255	-
8	4.9	1.4	41	145, 228	-
9	4.8	0.8	27	164, 159	10
10	4.8	3.2	85	143, 3	-

Table 2: The 10 most significant scale-space blobs extracted from the synthetic image in figure 10, as ranked according to their scale-space blob volumes. The columns show from left to right; the blob ranking, the scale-space blob volume, the selected scale, the area of the support region of the grey-level blob at the selected scale, the spatial position of the maximum at the selected scale, and the index of the Gaussian blob in table 1.

4.2 Experiments with 3-D PET images

The scale-space primal sketch algorithm was then tested on two real data sets. The first was obtained from a study in which subjects matched the shapes of tactically felt ellipsoids with visually presented ellipsoids (Hadjikhani & Roland 1998). From that experiment a descriptive Student's-t image of tactile-visual "matching – control" (see section 3) was subjected to scale-space analysis. The normalization of the scale-space blob volume was based on the set of difference images of "matching – matching".

Figure 11 shows the 100 most significant blobs isolated in the descriptive Student's-t image of tactile-visual "matching – control". Figure 12 shows a "corner plot" obtained by exposing the three orthogonal planes through the center of the most significant scale-space blob. In this picture also the 49 lower ranked blobs are apparent. Figure 13 and figure 14 show corresponding results for the second and third most significant scale-space blobs, respectively. It is obvious that the scale-space primal sketch algorithm captures the visual and somatosensory areas expected to be engaged in the somatosensory—visual cross-modality matching of shape (Roland & Larsen 1976, Seitz et al. 1991, Gulyás & Roland 1994, Logothetis et al. 1995). The third most significant scale-space blob is located in the insula claustrum, also known to be engaged in cross modal matching (Shindy et al. 1994, Hadjikhani & Roland 1998). Table 3 gives a summary of the significance values, the selected scale levels as well as the physiological brain regions corresponding to the 12 strongest blob responses.

In Figures 15(a)–(e) a comparison is done between the results of the 50 most significant blobs from the scale-space primal sketch and the statistical cluster analysis method of (Roland et al. 1993). The left column shows the boundaries of the scale-space blobs overlaid on the intensity image, whereas in the right column these boundaries have been overlaid upon a cluster image, in which each connected region of Student-t values above the statistical significance threshold has been assigned a unique colour.

The significance level was set such that there was an overall probability of less than 0.05 of finding one or more false positive clusters within the space of the brain (Roland et al. 1993, Roland & Gulyás 1996). As can be seen from Figures 15(a)–(e), the scale-space primal sketch algorithm captures the clusters in the somatosensory, visual and frontal cortex. In addition, the scale-space primal sketch also generates responses at coarser scales and shows the hierarchical structure of the activations, such that the broad support regions for the focal activations are also seen.

13 clusters were found with the method of (Roland et al. 1993), based on 2000 simulations of noise images generated from the image data from the eight subjects doing cross-modal matching. The scale-space primal sketch algorithm captured 11 of these clusters within the 12 highest ranked blobs (Table 3). Clusters #12 and #13 were not significant according to normal conventions, since $p > 0.05$. However, even clusters of doubtful statistical significance were included to illustrate the correspondence between the ranking of blobs and clusters. The blob ranked as number three did not correspond to any significant cluster ($p > 0.8$). However, the underlying Student-t image contained two clusters of 230 mm^3 and 75 mm^3 with average t-values of 3.2 each and situated closely together corresponding to the location of blob #3. Clusters ranked #5, #6, #7 and #10 by the (Roland et al. 1993) method were all captured

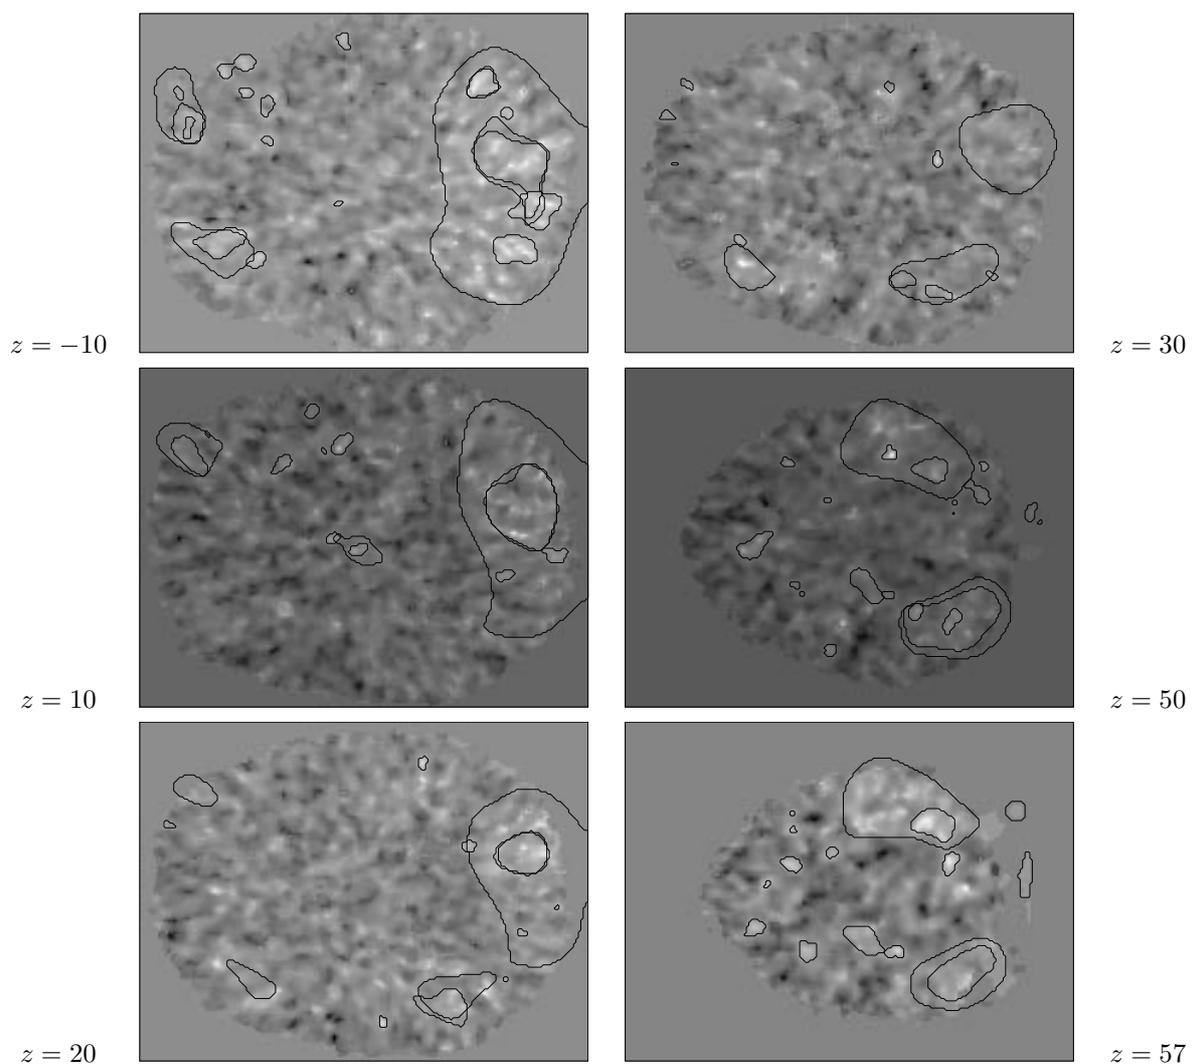


Figure 11: The result of applying the scale-space primal sketch to a three-dimensional Students-t image image from the visual-tactile crossmodality experiment. The results are displayed as six horizontal slices through the data volume with the boundaries of the 100 most significant scale-space blobs overlaid.

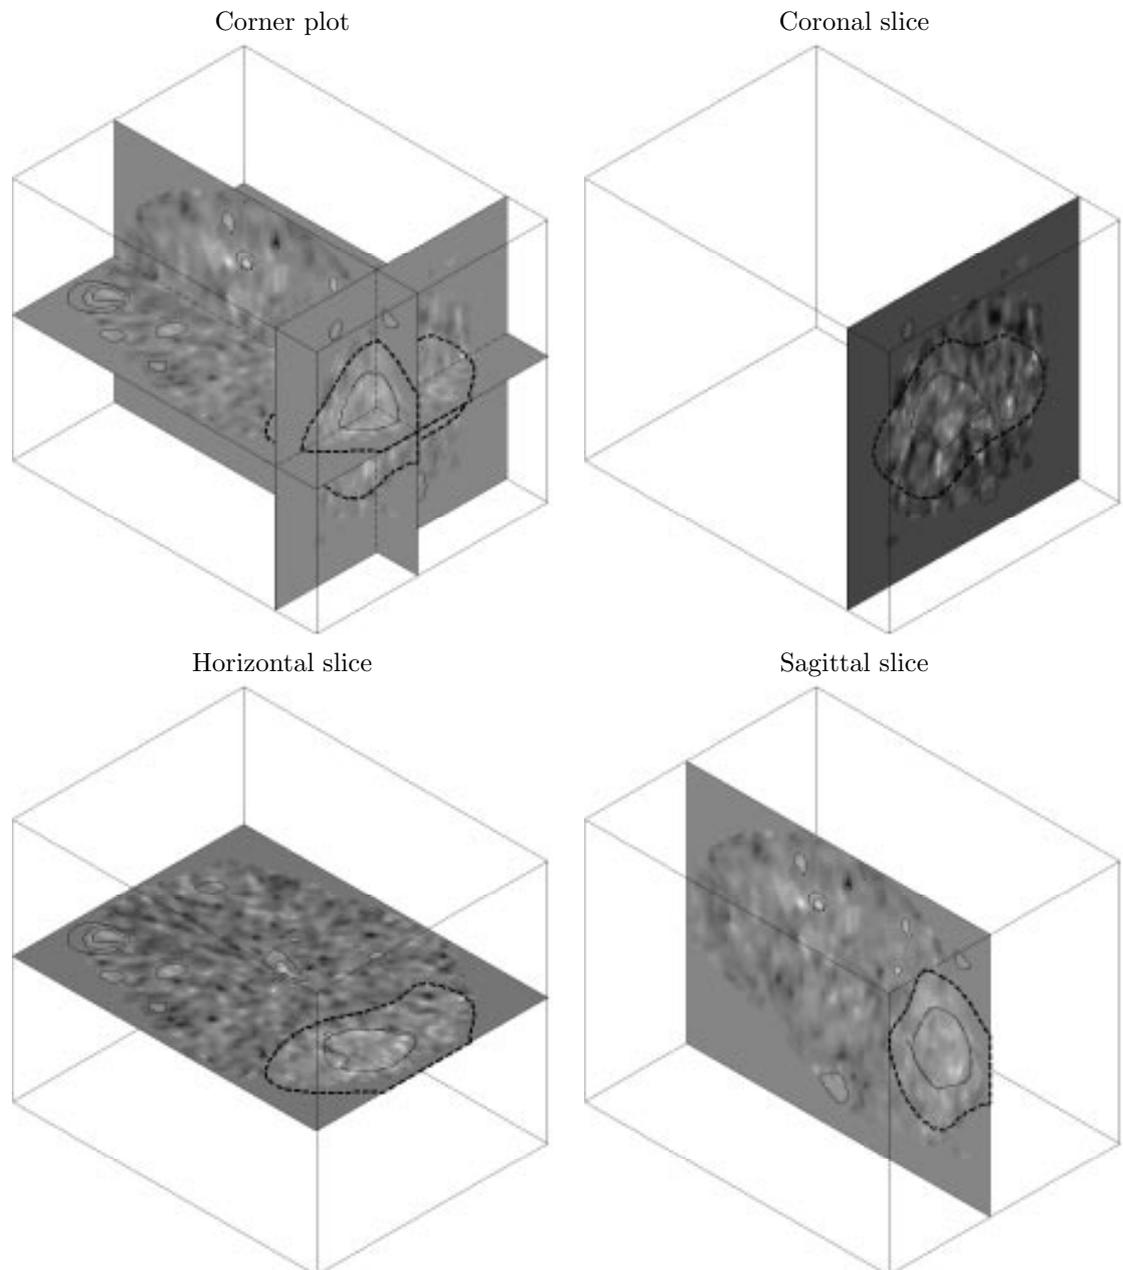


Figure 12: Three-dimensional illustration of scale-space blob number 1 from the visual-tactile cross-modality matching task. The figure shows the intersection of the center of the scale-space blob with three orthogonal planes.

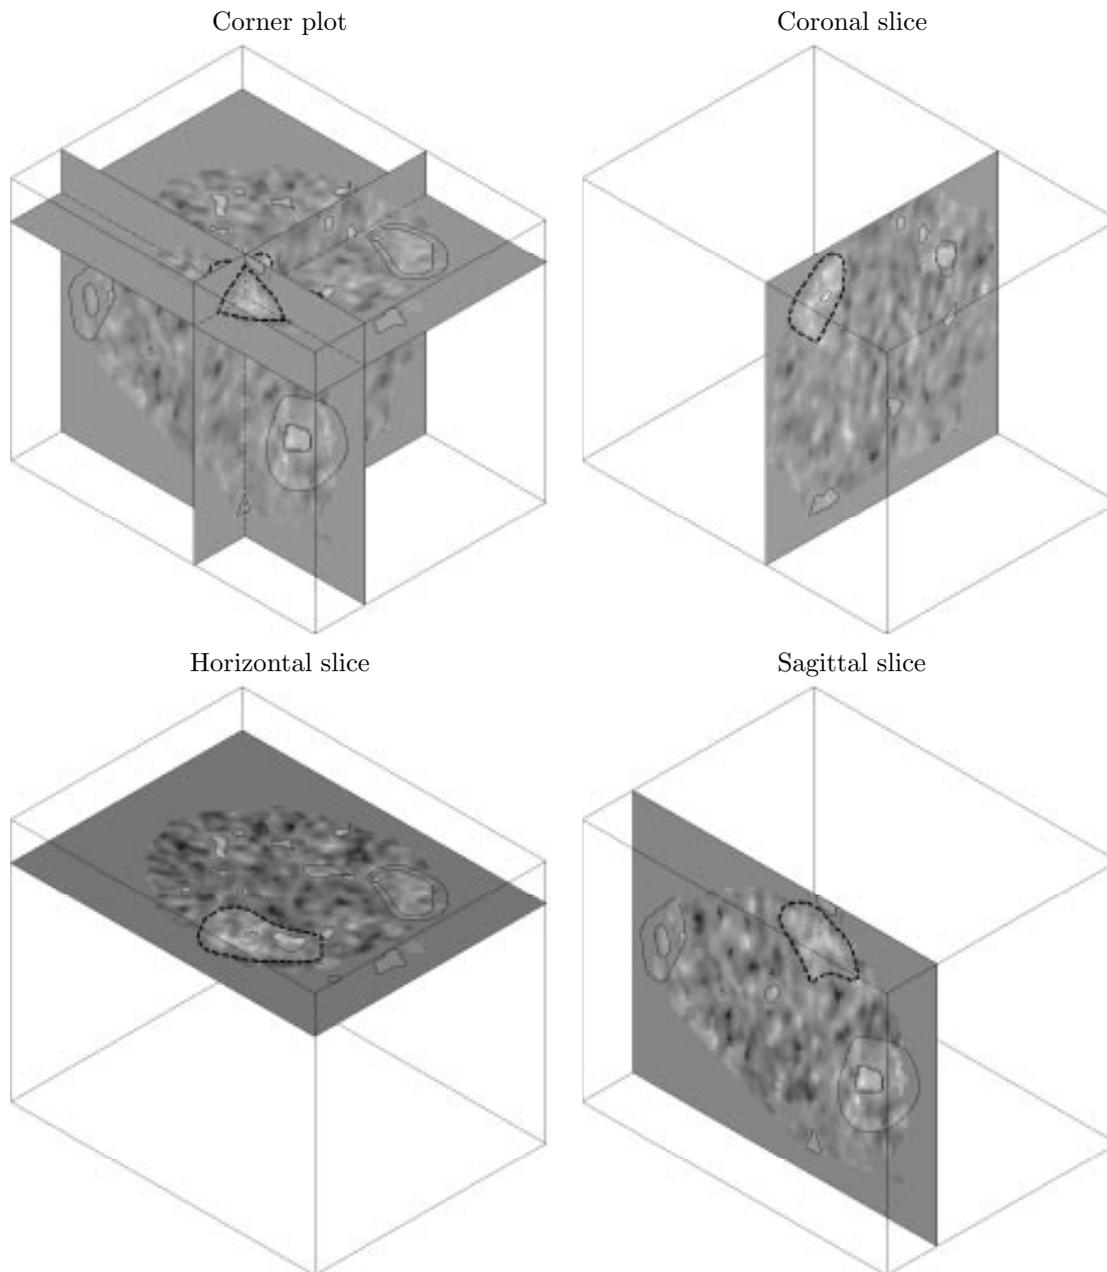


Figure 13: Three-dimensional illustration of scale-space blob number 2 from the visual-tactile cross-modality matching task. The figure shows the intersection of the center of the scale-space blob with three orthogonal planes.

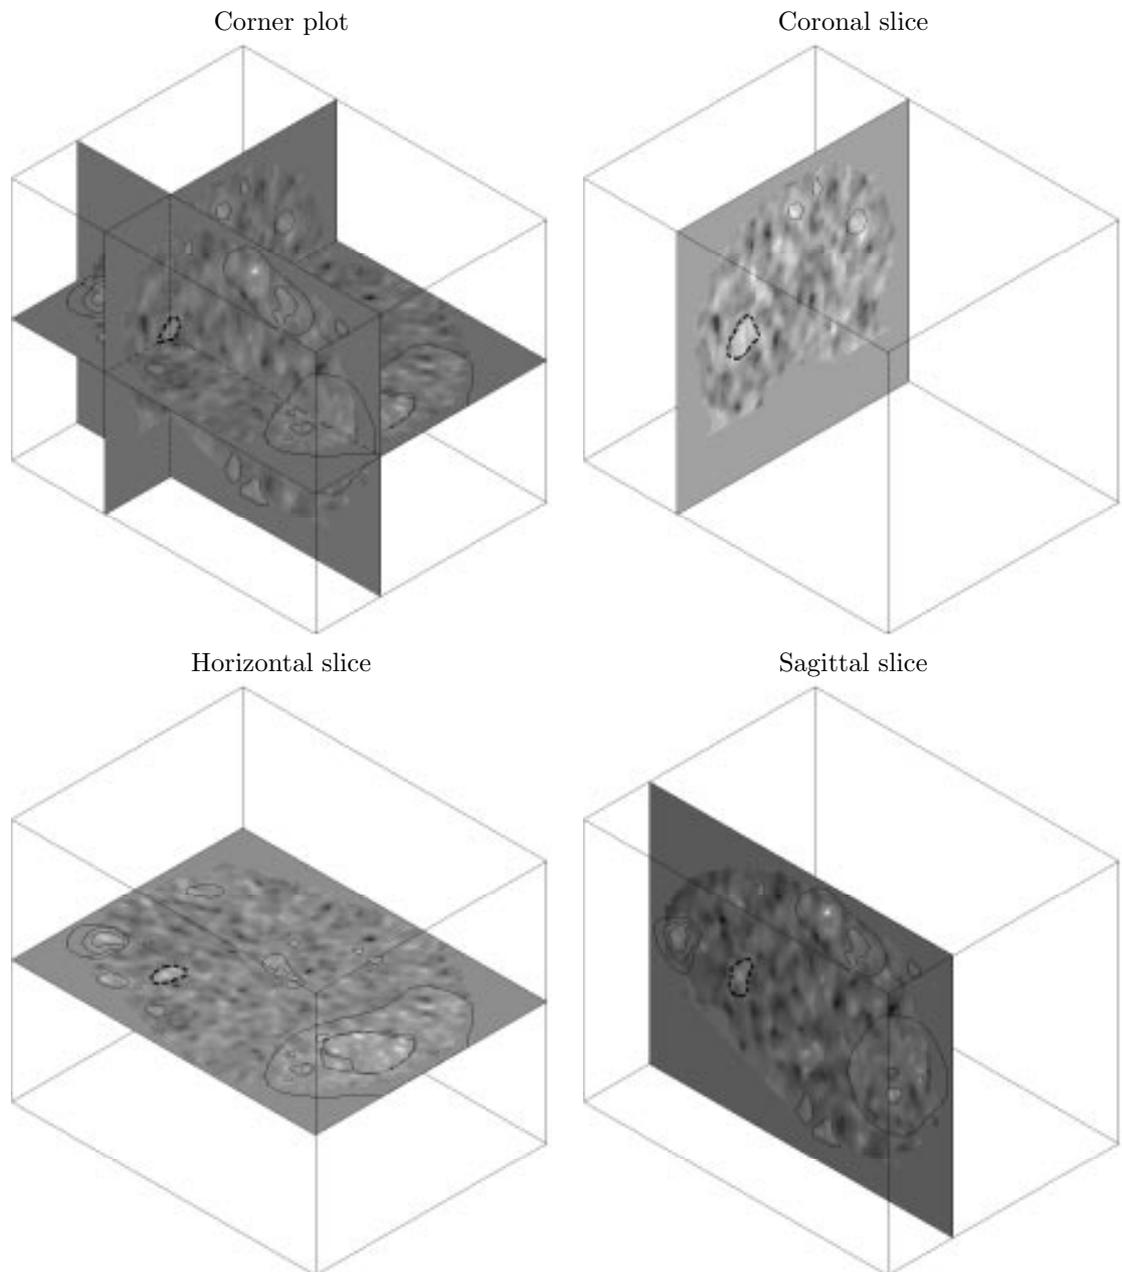


Figure 14: Three-dimensional illustration of scale-space blob number 3 from the visual-tactile cross-modality matching task. The figure shows the intersection of the center of the scale-space blob with three orthogonal planes.

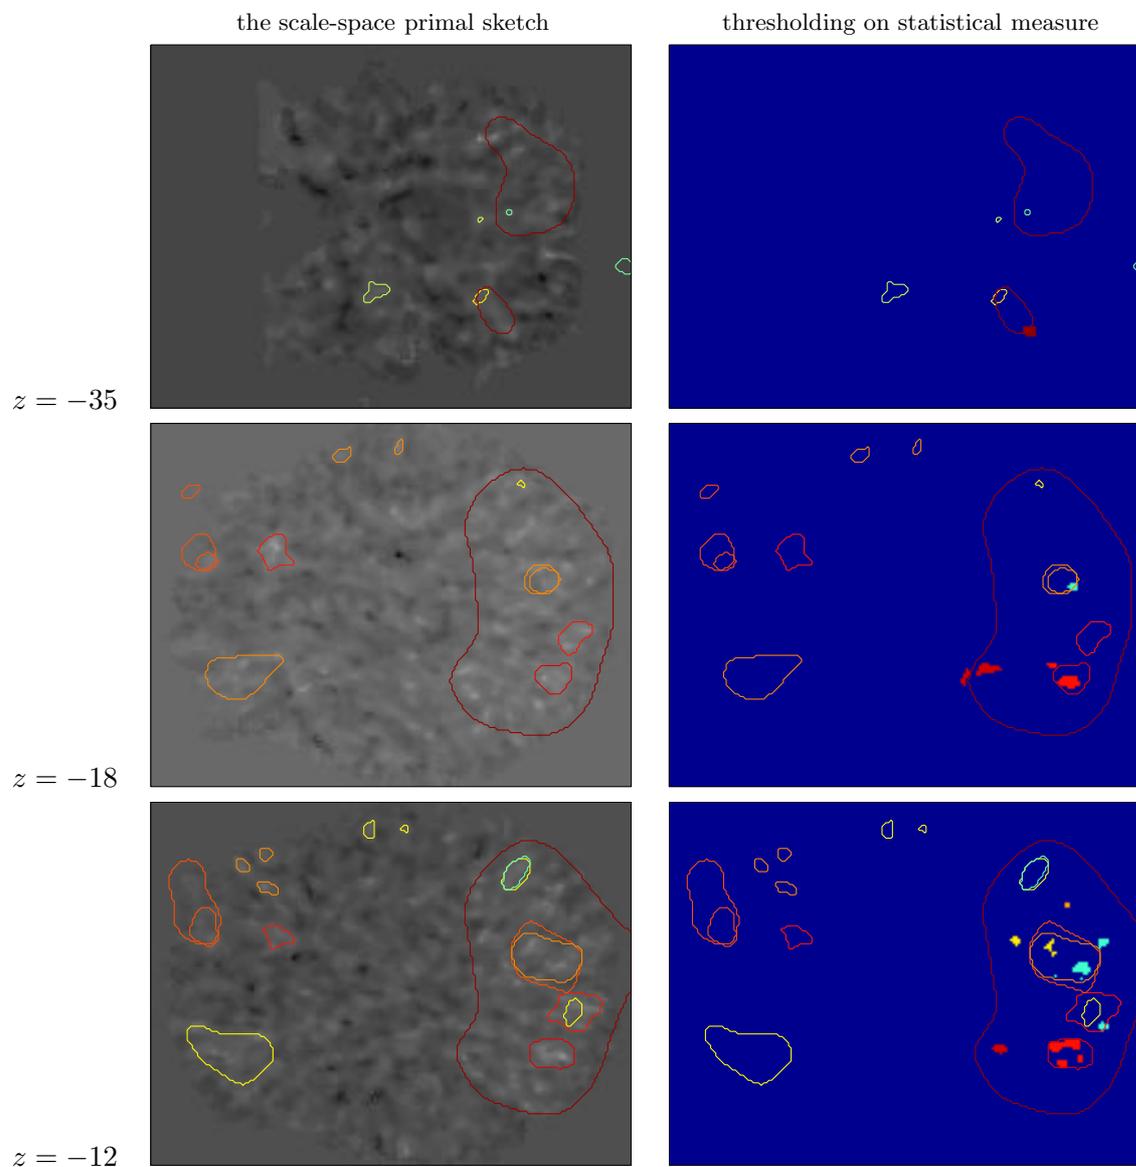


Figure 15: (a): Comparison between the results obtained from scale-space primal sketch and a standard technique for analysing PET activation patterns. The left column shows three horizontal slices through the data volume and the boundaries of the 50 strongest blob responses. In the right column, these boundaries have been overlaid upon a cluster image, in which each connected region of Student-t values above the significance threshold has been assigned a unique colour.

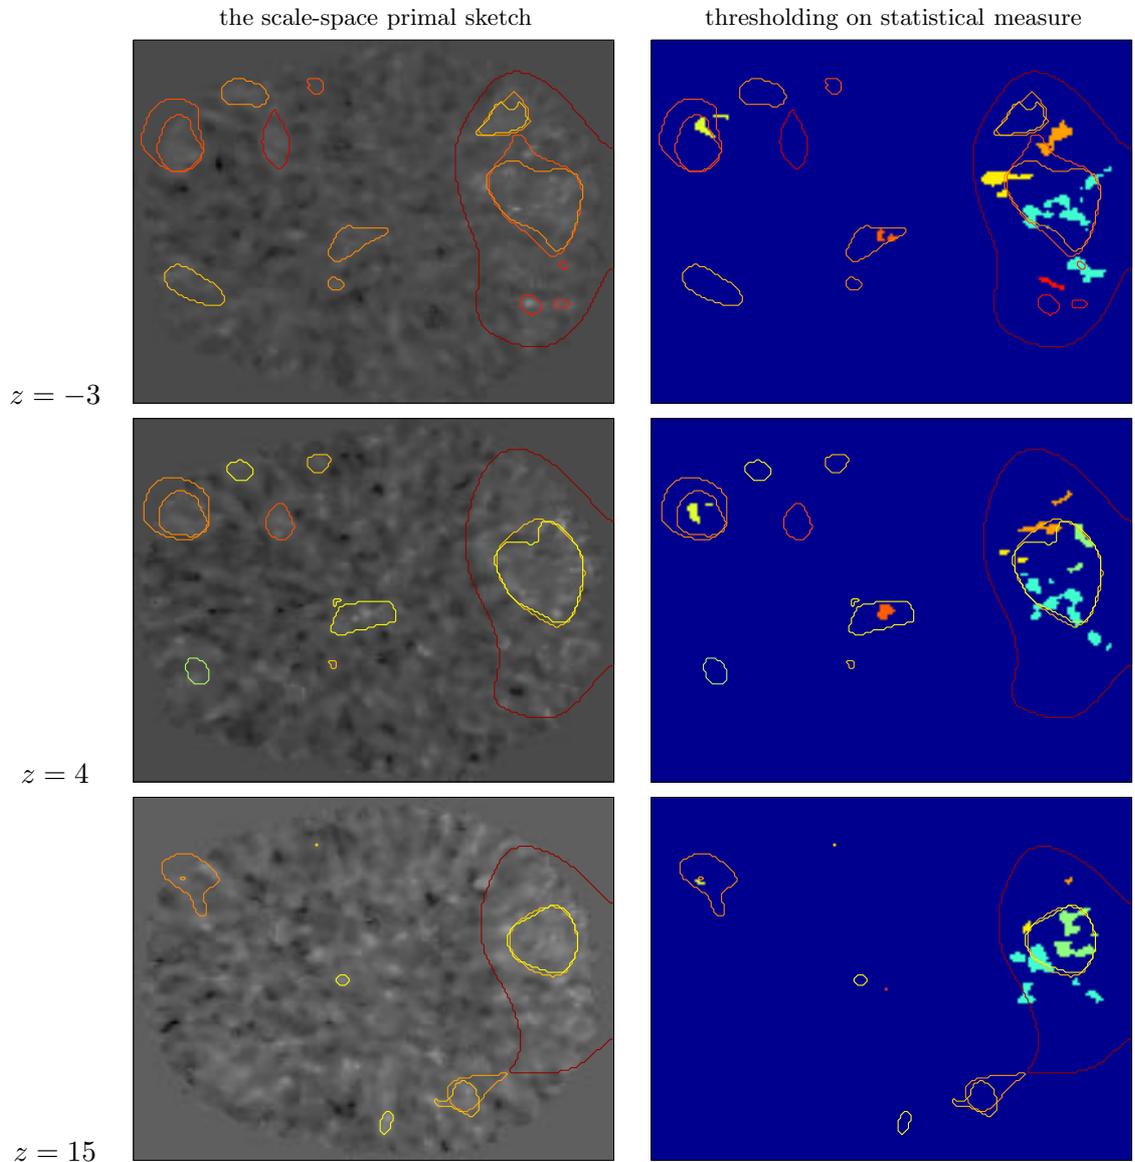


Figure 15: (b): Comparison between the results obtained from scale-space primal sketch and a standard technique for analysing PET activation patterns. The left column shows three horizontal slices through the data volume and the boundaries of the 50 strongest blob responses. In the right column, these boundaries have been overlaid upon a cluster image, in which each connected region of Student-t values above the significance threshold has been assigned a unique colour.

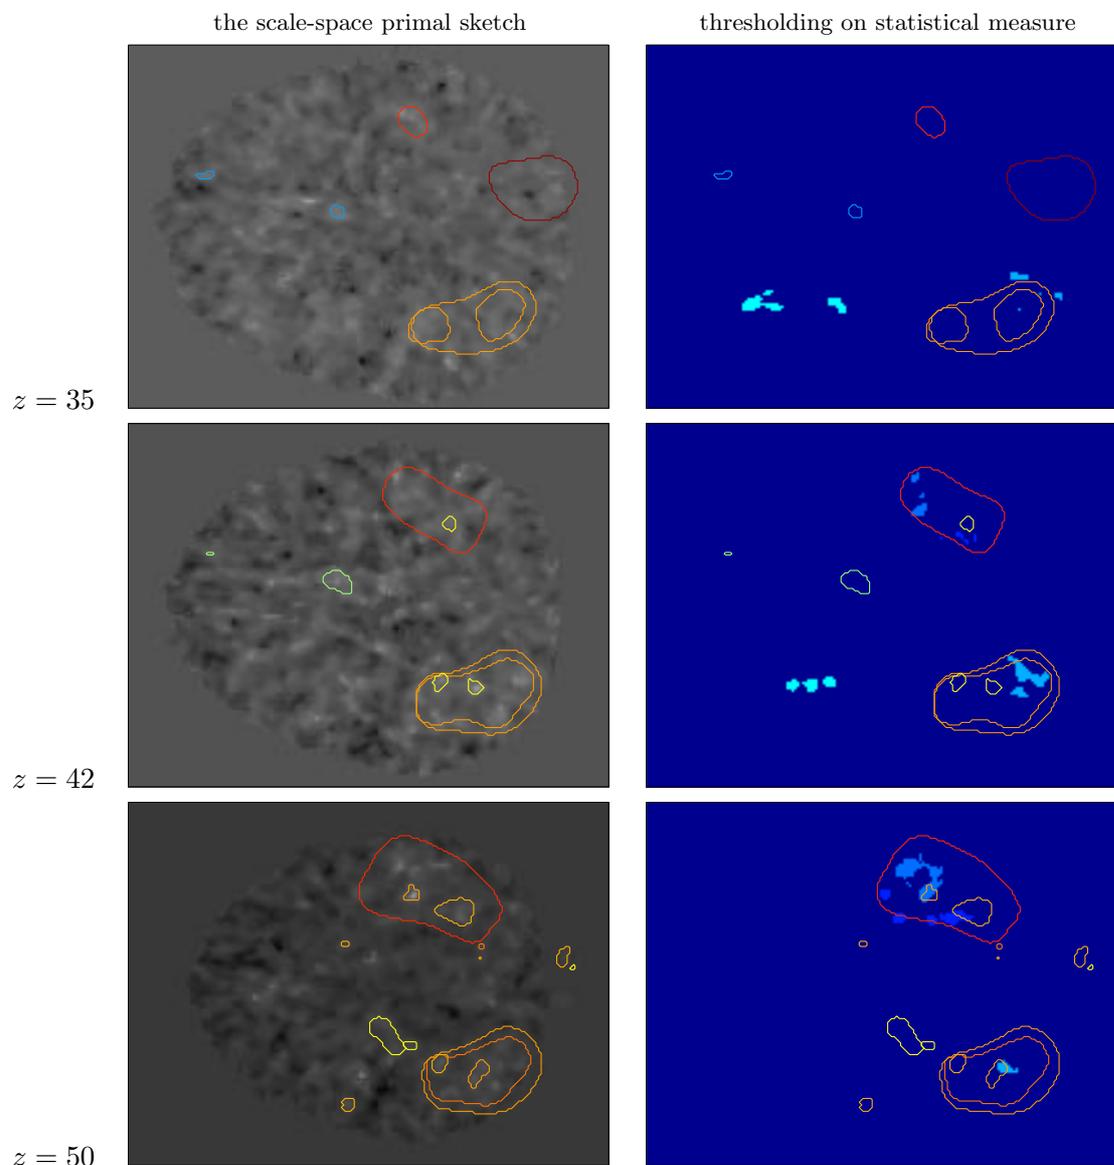


Figure 15: (c): Comparison between the results obtained from scale-space primal sketch and a standard technique for analysing PET activation patterns. The left column shows three horizontal slices through the data volume and the boundaries of the 50 strongest blob responses. In the right column, these boundaries have been overlaid upon a cluster image, in which each connected region of Student-t values above the significance threshold has been assigned a unique colour.

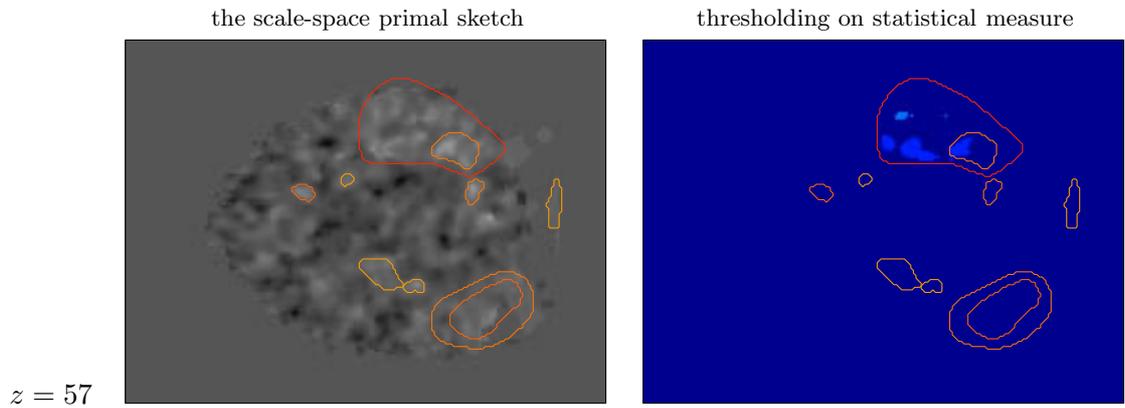


Figure 15: (d): Comparison between the results obtained from scale-space primal sketch and a standard technique for analysing PET activation patterns. The left column shows three horizontal slices through the data volume and the boundaries of the 50 strongest blob responses. In the right column, these boundaries have been overlaid upon a cluster image, in which each connected region of Student-t values above the significance threshold has been assigned a unique colour.

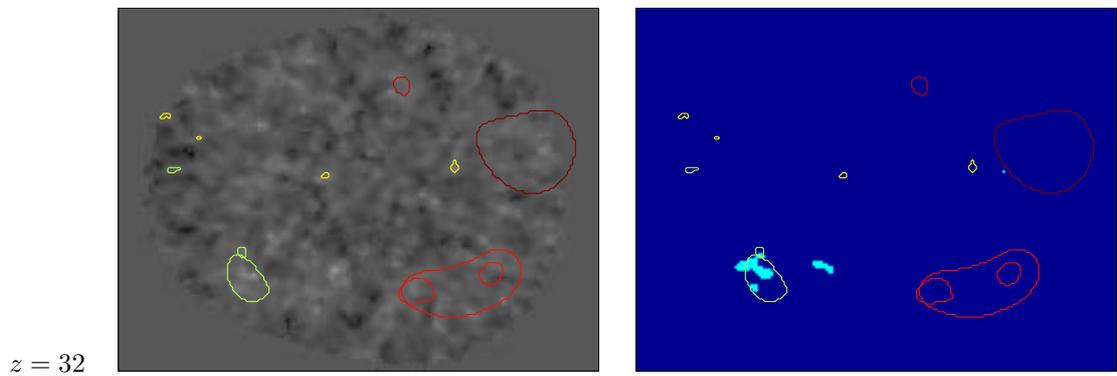


Figure 15: (e): Alternative illustration of the cross-section at $z = 32$ with the boundaries of the 100 strongest scale-space blobs drawn instead of the 50 strongest ones.

Visual-tactile cross-modality matching										
Blob ranking	Cluster ranking	Significance value	Selected scale	Talairach x, y, z	Blob size (mm ³)	Cluster size (mm ³)	Cluster p	Cluster	Brain region	
1	1	172.6	89.3	-11, -83, 2	249 443	12 540	0.0005	R & L visual areas		
2	4	48.4	50.2	-41, -35, 2	36 954	3 131	0.0005	L sensory motor areas		
3	ns	19.6	8.4	33, 20, 0	2 208			L anterior insula & claustrum		
4	3	17.4	7.3	-26, -79, -13	2 725	1 669	0.005	R fusiforme gyrus		
5	16	16.8	1.8	-21, -52, 41	952	470	0.45	R post. lobe cerebellum		
6	11	16.1	3.8	25, 28, 19	1 602	545	0.02	L orbitofrontal		
7	4	14.7	5.5	-7, -90, -14	2 479	1 197	0.0005	R occipital pole		
8	n.s.	14.0	2.5	31, 57, -2	17 772			L sup. frontal gyrus		
9	8	13.1	7.2	-33, -63, 47	39 978	955	0.0005	R intraparietal cortex		
10	n.s.	13.0	0.2	13, 8, 57	392			L sup. frontal gyrus		
11	(2)	12.9	83.8	40, -36, 53	44 542	3 131	0.0005	multiple response to blob #2		
12	(8)	12.8	32.6	-33, -61, 48	17 642	955	0.0005	multiple response to blob #8		

Table 3: Scale-space analysis and cluster analysis of the PET data from the visual-tactile crossmodality matching task. The columns show from left to right; ranking of the scale-space blobs, ranking of the clusters, scale-space blob volume, selected scale for the scale-space blobs, Talairach coordinates of the maximum associated with the scale-space blob, volume of the support region of the scale-space blob at the selected scale, volume from cluster analysis, the probability that the cluster is a false positive according to the model in (Ledberg et al. 1998), and the corresponding brain regions for the 12 strongest blob responses (“n.s.” means non-significant by cluster analysis).

within larger size blobs, i.e., blob #1 and #2. These clusters were also captured as smaller blobs at finer scales, but of lower rank, albeit within the 100 highest ranked blobs. The cluster ranked as #9 by the cluster method was not captured within the 100 highest ranked blobs. The cluster ranked as #12 appeared as blob ranked as #52.

From these data it is apparent that the scale-space algorithm picks out relevant regions which are found statistically significant, one exception being the left anterior insula and claustrum. The insula-claustrum has, however, been found significant in another comparison (Hadjikhani & Roland 1998). Sometimes the cluster and the scale-space blob do not exactly cover the same volume. At large scales, several clusters, for example located within the visual cortex, are encompassed in a large scale-space blob, and the large scale-space blob has a hierarchical substructure (see for example, blobs ranked #1, #4 and #7).

A second test of the scale-space primal sketch was made on the PET data from the tactile reaction time experiment. Figures 16–18 show corner plots through the centers of the three most significant scale-space blobs after the suppression of a multiple response (#3) to a region (#1) already detected. The strongest blob responses occurred in the left sensory motor hand region, the right supra-marginal gyrus and the supplementary and cingulate motor areas, all known to be engaged in sensory motor activities (Roland & Zilles 1996). For the detection of significant clusters, we did 7500 simulations to generate noise images from the nine subjects doing the somatosensory reaction time task. Four clusters appeared significant at $p < 0.06$, and were captured by the scale-space primal sketch (Table 4). The cluster ranked as #3 was included in blob #2. However, even clusters of lower rank and of less statistical significance are listed in Table 4. From these results it can be seen that all the 10 highest ranked clusters were captured by the 12 highest ranked blobs (cluster #6 was included in the blob #5). Again, however, blobs were detected at different scales, subdividing the image structure into hierarchical regions which are not apparent from the cluster analysis. One of the blobs ranked just below the 12 highest ones was the cluster of the mid-brain reticular formation, previously shown to be important for this type of task (Kinomura et al. 1996).

In the tables, the p -values of the clusters signify the the probability that the cluster of a particular volume is a false positive. Generally, the blob volumes are larger than the volumes of the corresponding clusters. This is not surprising, since most the scale-space blobs are usually detected at much coarser scales than the inner scale¹ of image at which the cluster analysis is performed.

In summary, activations were found in the following areas for the PET data from the reaction time experiment: the left somatosensory and motor region for the right hand, the right and left supplementary motor area and adjacent parts of the mesial prefrontal cortex, the left middle frontal gyrus, the right middle frontal gyrus and the left premotor cortex and the left angular gyrus. These domains of the cortex are either engaged in the motor control of the right hand or in the perception of the somatosensory signals or the regulation of the attention.

¹As a common pre-processing stage to both the scale-space analysis and the cluster analysis, all PET images were pre-filtered with filters of 4.2 and 5 mm full width half maximum, respectively. This pre-smoothing stage defines the inner scale of the data.

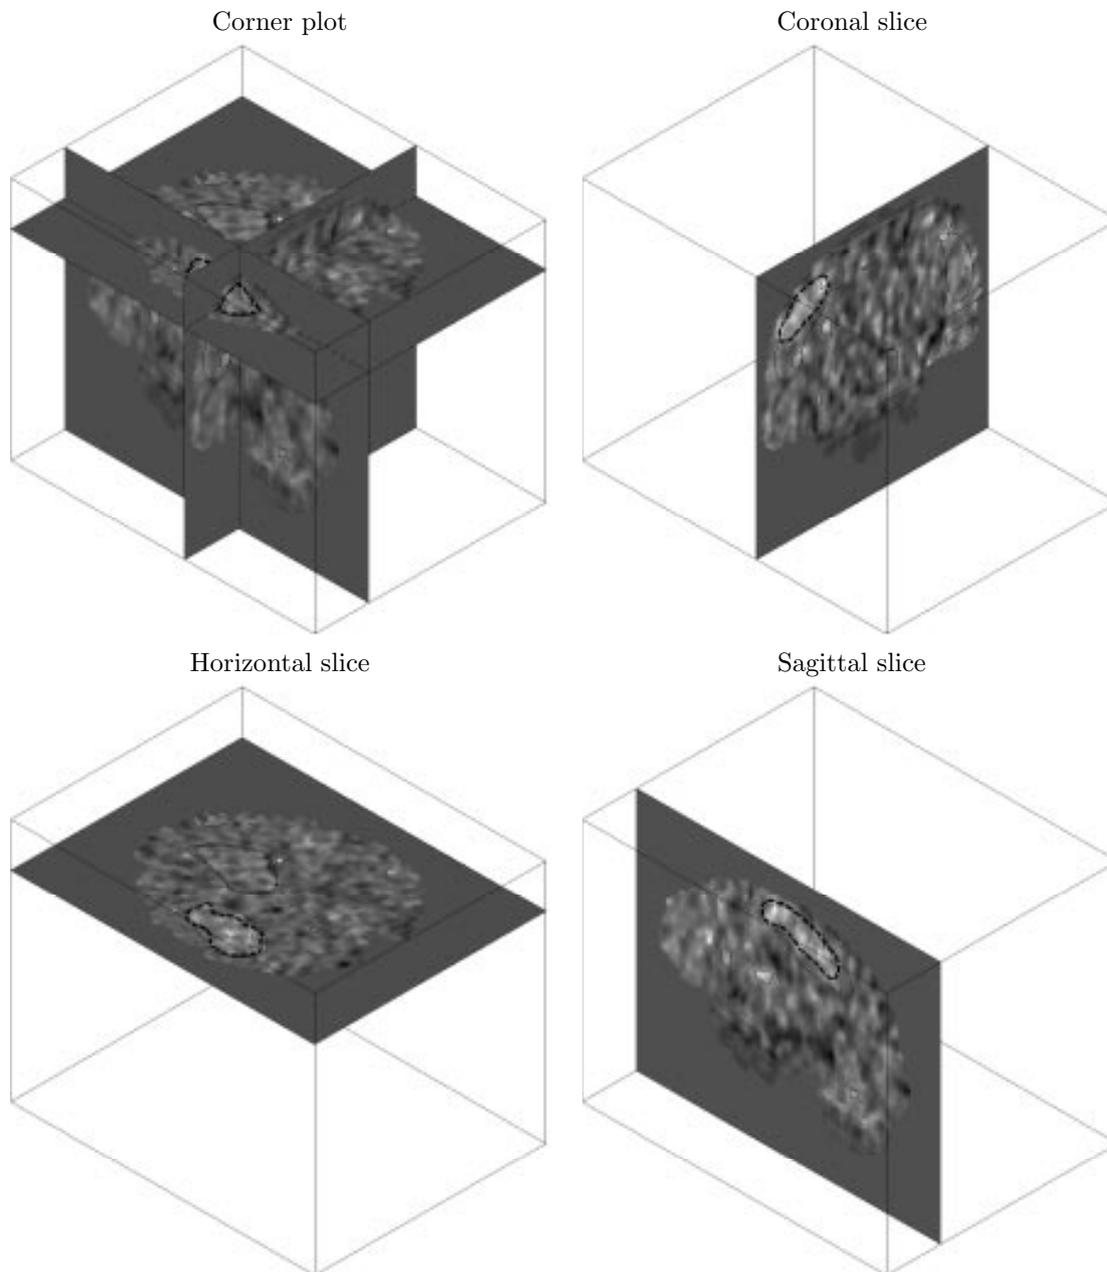


Figure 16: Three-dimensional illustration of scale-space blob number 1 from the tactile reaction time task. The figure shows the intersection of the center of the scale-space blob with three orthogonal planes.

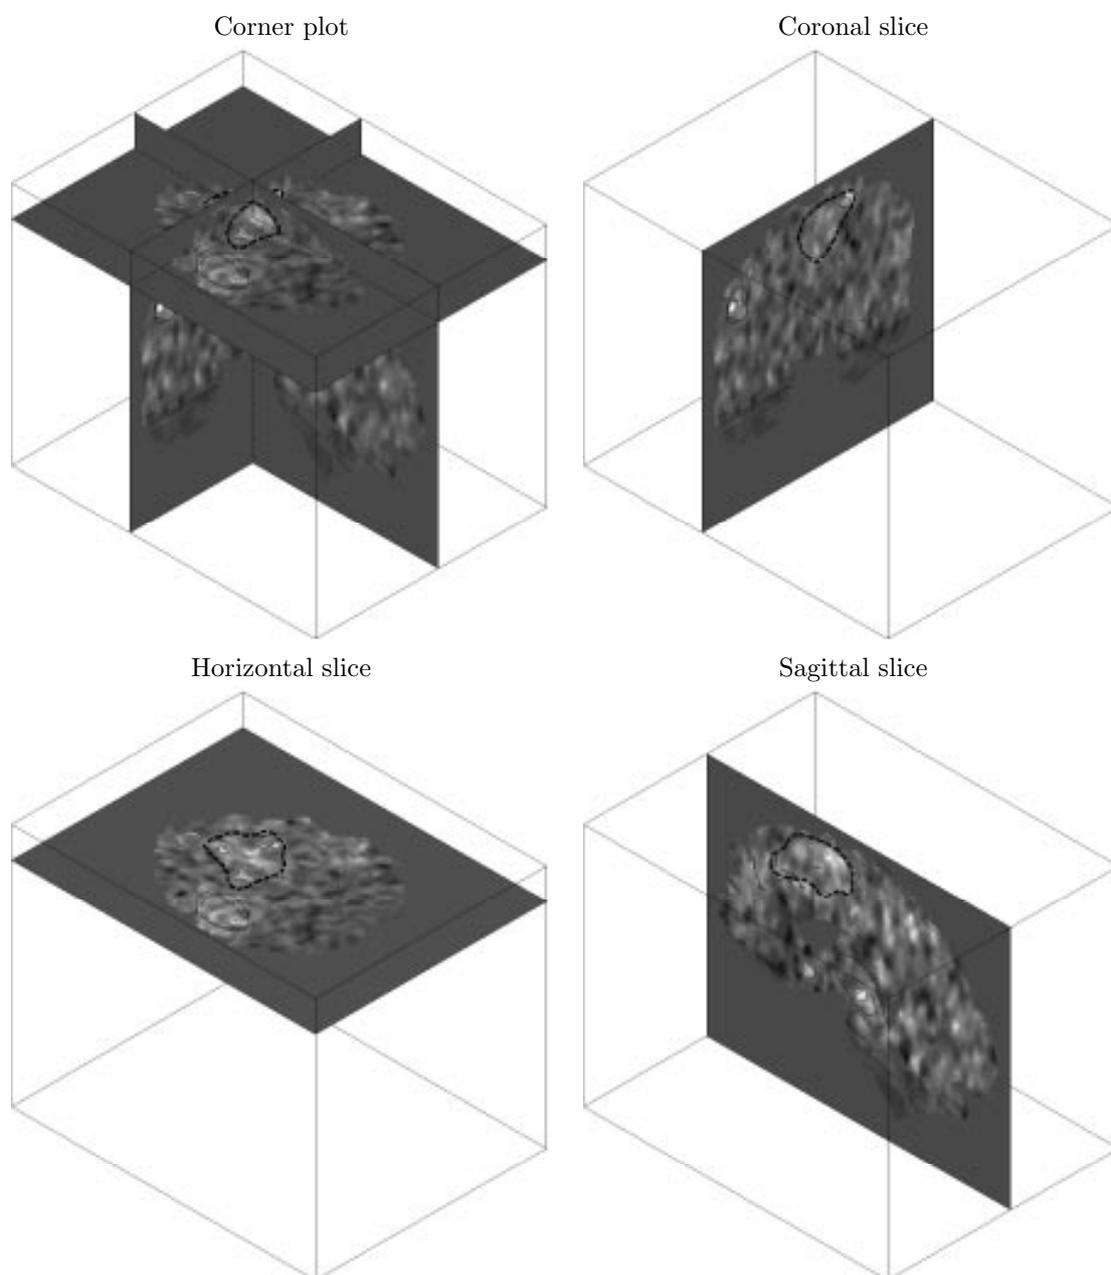


Figure 17: Three-dimensional illustration of scale-space blob number 2 from the tactile reaction time task. The figure shows the intersection of the center of the scale-space blob with three orthogonal planes.

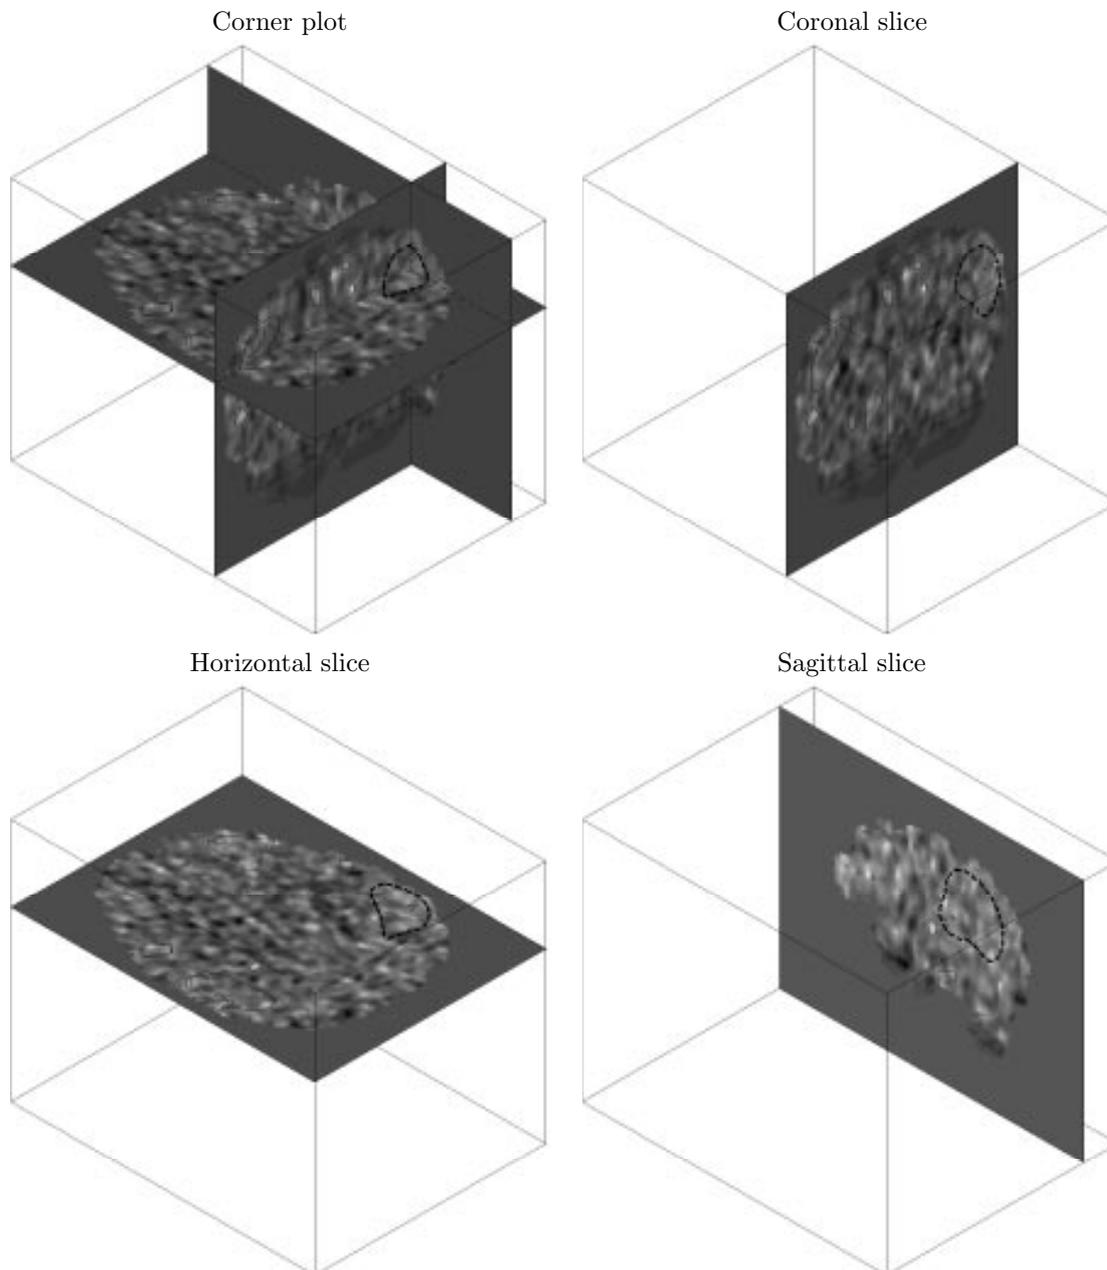


Figure 18: Three-dimensional illustration of scale-space blob number 3 from the tactile reaction time task. The figure shows the intersection of the center of the scale-space blob with three orthogonal planes.

Tactile reaction time									
Blob ranking	Cluster ranking	Significance value	Selected scale	Talairach x, y, z	Blob size (mm ³)	Cluster size (mm ³)	Cluster p	Brain region	
1	1	28.6	4.0	36, -30, 52	20 720	6 492	0.0002	L sensory motor areas	
2	4	8.8	18.1	-51, -42, 34	53 928	480	0.06	R supramarginal	
3	(1)	7.9	2.9	36, -30, 52	14 792			multiple response to blob #1	
4	8	7.2	0.0	-31, -5, 54	2 032	328	0.25	R premotor	
5	2	6.9	9.6	-1, 1, 56	16 560	1 504	0.0002	SMA cingulate motor areas	
6	7	5.6	2.9	-33, -5, 54	3 400	464	0.07	R post. prefrontal	
7	12	5.6	2.9	44, 49, 2	1 016	248	0.5	L sup. frontal gyrus	
8	n.s.	4.8	0.0	-17, -1, 17	424			R amygdala	
9	5	4.4	1.9	-39, -48, 26	1 144	472	0.06	R angular gyrus	
10	n.s.	4.3	0.0	-29, 29, -5	568			R orbitofrontal	
11	10	3.9	1.9	46, -60, 10	1 720	272	0.4	L angular gyrus	
12	9	3.9	9.6	32, 41, 40	1 288	280	0.4	L sup. frontal sulcal cortex	

Table 4: Scale-space analysis and cluster analysis of the PET data from the tactile reaction time task. The columns show from left to right; ranking of the scale-space blobs, ranking of the clusters, scale-space blob volume, selected scale for the scale-space blobs, Talairach coordinates of the maximum associated with the scale-space blob, volume of the support region of the scale-space blob at the selected scale, volume from cluster analysis, the probability that the cluster is a false positive according to the model in (Lidberg et al. 1998), and the corresponding brain regions for the 12 strongest blob responses (“n.s.” means non-significant by cluster analysis).

5 Summary and discussion

The scale-space primal sketch provides an unambiguous mathematical tool, for the analysis of signals in 3-D images in the presence of noise. The scale-space primal sketch allows the detection of the localization, volume and hierarchical structure of any blob-like entity in the image. In addition, the scale-space primal sketch algorithm automatically selects optimal scales for detecting individual blobs and it provides a ranking of the blobs detected. The algorithm works without problems on synthetic data in the presence of added noise, and it has been tested on real PET data from which it was able to detect and rank all the clusters considered statistically significant.

The scale-space primal sketch is computed from image data almost entirely without specific assumptions about the structures to be analysed. This is in contrast to all existing methods of quantitative brain activation analysis except (Holmes et al. 1996). The only prerequisites are that the structures to be detected are blob-like, and that normalization can be performed on reference data.

Concerning limitations of this approach, one might envisage structures in the brain which are not blob-like and which could be activated, for example thin unfolded gray matter structures. One such structure is the claustrum. To be able to capture such elongated structures, explicit focus on more specialized methods for detecting ridge-like image features (Pizer et al. 1994, Eberly et al. 1994, Koenderink & van Doorn 1994, Lindeberg 1996) would probably be beneficial.

The ranking of blobs is dependent on the images used for normalization. The density of local extrema as a function of the scale-space parameter t can be sensitive to the choice of normalization data. Using real difference images seemed to give the best solution. This also has the advantage that one is free from assumptions about the statistical structure of the image data. Provided that the assumption of a simultaneous normal distribution is satisfied, simulated Student's t -images could be used as reference. Our experiments, however, indicate that this assumption may not be fully satisfied. Images obtained from simulation experiments seem to deviate from the real data, especially at finer scales (cf, Figures 8 and 9).

The scale-space primal sketch was compared to two statistical methods of cluster analysis in 3-D images (Roland et al. 1993, Ledberg & Roland 1996). The comparison shows that the results of the scale-space primal sketch are to a large extent compatible with the results of the cluster based method, although the results also suggest that the scale-space primal sketch may be more sensitive. It may seem natural to use a method sensitive to the extent of functional activations as a comparison with this extent and localization of the blobs. It is, however, beyond the scope of the present article to compare the scale-space primal sketch to all currently used methods for the detection of statistically significant functional activations. Such comparisons can only be roughly indicative of the performance of the primal sketch. A major difference is that the current methods all rely on statistical assumptions of stationarity, and some on many additional assumptions. Also whereas the current methods are relatively efficient in controlling for the probability of false positives, α , the sensitivity of the current methods can only be estimated under strong assumptions. We plan to address these issues in future work, including the statistical foundations of the scale-space primal sketch.

Clearly, an approach such as cluster analysis is sub-optimal compared to the scale-space primal sketch, since it only considers data at a certain level of scale. Consequently, one cannot expect to find 100 % concordance between the two methods. The 3-D cluster analysis cannot, at the level of filtering used, detect the large blobs of the visual or somatosensory cortex. These might have been detected if the input data had been filtered with a large Gaussian filter, but at the expense of detecting the smaller focal activations. This is exactly the dilemma of current statistical methods (Friston et al. 1991, Friston et al. 1994, Friston et al. 1996, Worsley et al. 1992). Moreover, all current statistical methods with the exception of (Worsley et al. 1996b) employ thresholding of statistical images. Consequently, the extent of the detected clusters depend on the level of thresholding. The scale-space primal sketch on the contrary depends on the mathematical description and delimitation of functional activated fields. Finally, some currently methods are dependent on a certain filter size for the theoretical assumptions underlying the methods be valid. This excludes the detection of smaller blobs, for example confined to sub-cortical structures. Specifically, the approach of detecting connected components in thresholded statistical data introduces a certain degree of arbitrariness, since the global connectivity between clusters can be influenced by a small number of local pixel values. The scale-space primal sketch in contrast detects blobs no matter of the scale in which these blobs occur.

A natural question concerns which blobs represent true increases of rCBF and whether all blobs having a significance value above a certain threshold are biologically meaningful. This issue can be separated into a statistical question of which blobs are due to noise and the question of whether blobs at all scales are biologically meaningful. The large scale blobs in the real PET data from the first experiment comprised the visual areas and the somatosensory areas. With respect to the particular task of matching somatosensory and visual shape information, the engagement of (many) somatosensory and visual areas is expected. Furthermore, such large activations are sometimes seen in data which have been heavily smoothed, for example (Dupont et al. 1993, Grasby et al. 1994, Shallice et al. 1994, Kapur et al. 1995). The large blobs provide support regions for smaller and more focal activations located at places expected to be engaged in the representation and perception of objects (Hadjikhani & Roland 1998).

The scale-space primal sketch algorithm is automatic and there are no parameters to tune. The only free parameter is the number of blobs to be extracted from the image. This number has here been set arbitrarily and without statistical testing of the number and sizes of blobs that may be noise blobs. There are four current statistical approaches, which bear some resemblance to the scale-space primal sketch. (Poline & Mazoyer 1993) used a search over a range of Gaussian filter widths for signals in Student's t-images. For each filter size, one or more blobs could be extracted from an image slice. Blobs detected at a small scale were removed from the sample to prevent them being detected again at coarser levels of scale-space, and a hierarchical structure of nested regions was defined at each scale (Poline & Mazoyer 1994). The significance level was then subsequently determined from Monte-Carlo simulations. (Worsley et al. 1996a) used a model for PET activations, in which it was assumed that the signal consisted of just one peak and that its shape was Gaussian and that the PET data can be described as a Gaussian field. An extension of the latter method

(Worsley et al. 1996b) also permits a calculation of a unified p -value for 4-D local maxima.

There are close relationships between these approaches and the scale-space primal sketch. The hierarchical decomposition used by (Poline & Mazoyer 1994) is similar to the grey-level blob tree defined at each scale; the main difference appears to be how grey-level blob volumes are defined at each level. The scale-space primal sketch does not suffer from the problems of analysing data at a single scale, which gives a preference to regions of a specific size and limits fixed-scale methods such as those proposed by (Worsley et al. 1996a) and (Roland et al. 1993). Our method analyses all scales simultaneously, while the method proposed by (Poline & Mazoyer 1994) only combines the results obtained from applying a detection method at each single scale. The idea in (Worsley et al. 1996b) to detect 4-D maxima with respect to scale and space bears very close relationship to approaches for automatic scale selection in the computer vision literature (Lindeberg 1994). In the scale-space primal sketch, scale levels are selected from the scales at which the normalized grey-level blob volume assumes maxima with respect to scale (with complementary selection of spatial points along the extremum paths). A more general methodology for scale selection for feature detection developed in (Lindeberg 1993c)(Lindeberg 1994, chapter 13) is based on the joint maximization of normalized responses to differential operators with respect to scale and space. In these respects, the scale-space primal sketch unifies components of the approaches by (Poline & Mazoyer 1994, Worsley et al. 1996b) into a common framework. Three main differences are that the scale-space primal sketch, in addition, (i) involves a linking of topologically similar image structures over scales, (ii) makes the hierarchical relations explicit between image structures at different scales, and (iii) includes the stability of image structures across scales into the significance measure. Both of the approaches by (Poline & Mazoyer 1994) and (Worsley et al. 1996b) carry further assumptions. In contrast, the scale-space primal sketch carries no assumptions on the structure of the image data or the shape of the signals. The signals are detected and ranked automatically on the basis of both peak amplitude and extent simultaneously at all scales, and the scale-space primal sketch gives rules for subdividing regional activations into a hierarchy of unambiguously defined parts. Its drawback is that, at the present stage, no p -value is attached to each blob. Such a value could, in principle, be defined and we will address this issue in our future work.

Acknowledgements

This work has been supported by a grant from the European Commission under the BioTech programme BIO-CT 96-0177 (P. E. R.), the Swedish Research Council for Engineering Sciences, TFR, (T. L.), and the Swedish Foundation for Strategic Research (T. L. and P. E. R.).

A preliminary version of this work was presented in (Lidberg et al. 1996). Independently of our work, a parallel application of the scale-space primal sketch to the analysis of brain activation data has been presented by (Coulon et al. 1996, Coulon et al. 1997).

A Appendix

A.1 Definition of grey-level blob

A precise mathematical definition of the grey-level blob concept can be stated as follows (Lindeberg 1994, section 7.1.2): Consider the case with bright blobs on dark background, and assume a continuous generic (non-degenerate)² grey-level function $f: \mathbb{R}^D \rightarrow \mathbb{R}$ at a fixed level of scale. Consider a local maximum $A \in \mathbb{R}^D$. For any grey-level $z < f(A)$ let

$$X_z^{(A)} = \{ \text{the connected component of } \{(x; \zeta) \in \mathbb{R}^D \times \mathbb{R} : z \leq \zeta \leq f(x)\} \text{ that contains } (A, f(A)) \}, \quad (12)$$

and define the sets $G_z^{(A)}$ and $H_z^{(A)}$ as follows: A point $(B, \zeta_0) \in X_z^{(A)}$ belongs to $G_z^{(A)}$ ($H_z^{(A)}$) if and only if there exists a path $p_{(A, f(A)), (B, \zeta_0)}$ from $(A, f(A))$ to (B, ζ_0) such that (i) every point on the path belongs to $X_z^{(A)}$, and (ii) the derivative of ζ along this path $\zeta'|_{p_{(A, f(A)), (B, \zeta_0)}} < 0$ ($\zeta'|_{p_{(A, f(A)), (B, \zeta_0)}} \leq 0$). The base-level of the blob $z_{base}(A)$ is then defined as the maximum value of z such that

$$z_{base}(A) = \max_{z < f(A)} z : \overline{G_z^{(A)}} \neq H_z^{(A)}, \quad (13)$$

where the notation \overline{C} stands for the closure of a set C . $z_{base}(A)$ is the grey-level value of the *delimiting saddle point* $S = S_{delimit}(A)$ associated with A . The grey-level blob associated with the local maximum A is the set of points

$$G_{blob}(A) = \overline{G_{z_{base}(A)}^{(A)}}, \quad (14)$$

with the (three-dimensional) *grey-level blob volume*

$$G_{vol}(A) = \int_{(x; z) \in G_{blob}(A)} dx dz. \quad (15)$$

The projection of this region onto the spatial plane is called the *support region*,

$$D_{supp}(A) = \{x \in \mathbb{R}^D : (x; \zeta) \in G_{blob}(A) \text{ for some } \zeta\}, \quad (16)$$

and the difference in grey-level between the extremum point and the base-level gives the *blob contrast*

$$C_{blob}(A) = f(A) - z_{base}(A). \quad (17)$$

It is worth stressing that the grey-level blob is treated as an object with extent both in space and grey-level. Whereas the definition is expressed in terms of signals defined on a continuous domain, it can be extended to discrete signals by replacement of \mathbb{R}^D by \mathbb{Z}^D , and by letting the paths be given by a suitable connectivity concept (e.g., eight-connectivity for a two-dimensional square grid and 26-connectivity on a three-dimensional rectangular grid). In the discrete case, the derivative condition $f'|_{p_{A, B}} < 0$ is replaced by a difference condition $f(x^{(k+1)}) - f(x^{(k)}) < 0$ along the path $\{x^{(k)}\}$.

²Unless otherwise stated, the signals are throughout assumed to be *Morse*, i.e., all critical points are assumed to be non-degenerate, and all critical values are assumed to be distinct.

A.2 Definition of scale-space blob

A.2.1 Extremum path and saddle path

Consider a critical point (a local maximum, a local minimum or a saddle point) $x_0 \in \mathbb{R}^D$ at some scale $t_0 \in \mathbb{R}_+$ in the scale-space representation of a D -dimensional signal. The critical point is given by

$$\nabla L|_{(x_0; t_0)} = \{\partial_{x_i} L\}|_{(x_0; t_0)} = 0. \quad (18)$$

The implicit function theorem ensures that if the Hessian matrix

$$\mathcal{H}L|_{(x_0; t_0)} = \{\partial_{x_i x_j} L\}|_{(x_0; t_0)} \quad (19)$$

is non-singular at this point, then there exists some smooth function $r_0 : I_{t_0} \rightarrow \mathbb{R}^D$

$$x = r_0(t) \quad (20)$$

such that $x_0 = r_0(t_0)$, and for every t in some neighbourhood I_{t_0} of t_0 the point $(r_0(t); t)$ is a critical point for the mapping $x \mapsto L(x; t)$.

By continuation, such local paths can be extended to curves as long as the Hessian matrix remains non-singular. It can be easily shown that the type of critical point remains the same as long the Hessian matrix is non-singular.

In other words, if $(x_0; t_0)$ is a local maximum (minimum/saddle), then there exists a curve through this point, such that every point on the curve is a local maximum (minimum/saddle) at that scale. The curve is delimited by two scale levels t_{min} and t_{max} , at which the Hessian matrix degenerates (except for the boundary cases $t_{min} = 0$ or $t_{max} = \infty$). At all interior points the extremum point is non-degenerate. Such a curve $r_0 : [t_{min}, t_{max}] \rightarrow \mathbb{R}^D$ is called an *extremum path (saddle path)* (Lindeberg 1994, section 8.1).

A.2.2 Scale-space blob

Concerning grey-level blobs, this result means that a unique linking of grey-level blobs across scales can be performed as long as both the extremum point and the saddle point determining the extent of the grey-level blob remain non-singular. In summary, a scale-space blob is defined as the union of all grey-level blobs associated with the extremum points along a segment of an extremum path where such a unique linking can be performed.

To express this statement precisely, let $[t'_{min}, t'_{max}] \in [t_{min}, t_{max}]$ be a (maximal) subset of an extremum path, along which the delimiting saddle point $S_{delimit}(r_0(t))$ associated with the extremum point $r_0(t)$ is always non-degenerate. At some distinct scales it may happen that the delimiting saddle point jumps from one saddle path to another. In such non-Morse³ situations, when two saddle points have the same grey-level, both saddle points are required to be non-degenerate. At the end points, either of $r(t'_{min})$ and $S_{delimit}r(t'_{min})$ and also either of $r(t'_{max})$ and $S_{delimit}r(t'_{max})$ are degenerate critical points (unless $t_{min} = 0$ or $t_{max} = \infty$).

³Generically, these events occur at isolated scales, and then only two different critical points have the same critical values.

Then, the scale-space blob associated with this segment $r'_0 : [t'_{min}, t'_{max}] \rightarrow \mathbb{R}^D$ is the set

$$S_{blob}(r'_0) = \overline{\{(x; z; t) \in \mathbb{R}^D \times \mathbb{R} \times \mathbb{R}_+ : (t'_{min} < t < t'_{max}) \wedge ((x; z) \in G_{blob}(r'_0(t)))\}}, \quad (21)$$

where $G_{blob}(r'_0(t))$ is the grey-level blob associated with the extremum point $r'_0(t)$ in the scale-space representation L at scale t (Lindeberg 1994, section 8.1.2).

A.3 Scale-space blob events

The implicit function theorem used in previous appendix section guarantees that linking of non-degenerate critical points is a well-defined operation. When the Hessian matrix becomes singular, *bifurcations* may occur. In summary, the following result holds at such singularities (Koenderink & van Doorn 1986, Lifshitz & Pizer 1990, Lindeberg 1992, Johansen 1994, Damon 1996): In $D \geq 2$ dimensions, the only generic (structurally stable) bifurcations are annihilations and creations of pairs consisting of one extremum point and one saddle point. A canonical model of this so-called *fold singularity* is for a local maximum in D dimensions given by the polynomial

$$x_1^3 + 3x_1(t - t_0) - \sum_{i=2}^D (x_i^2 + t - t_0). \quad (22)$$

The positions of the critical points are given by

$$(x_1(t), x_2(t), \dots, x_D(t)) = \pm(\sqrt{t_0 - t}, 0, \dots, 0) \quad (t \leq t_0) \quad (23)$$

i.e. the critical points merge along a parabola. Concerning scale-space blobs, this classification means that two distinct types of cases can be distinguished, depending on whether the saddle point involved in the bifurcation is part of one or two grey-level blobs. A saddle point delimiting the extent of only one grey-level blob is said to be *non-shared*, while a saddle point belonging to two grey-level blobs is said to be *shared*. Hence, in the generic case, there are four following cases are possible at a structurally stable bifurcation (Lindeberg 1994, section 8.4.2) (see Figure 4 for an illustration):

- *blob annihilation* — annihilation of an extremum-saddle pair where the saddle path is *non-shared before* the bifurcation,
- *blob merge* — annihilation of an extremum-saddle pair where the saddle path is *shared* with another scale-space blob *before* the bifurcation,
- *blob split* — creation of an extremum-saddle pair where the saddle path is *shared* with another scale-space blob *after* the bifurcation,
- *blob creation* — creation of an extremum-saddle pair where the saddle path is *non-shared after* the bifurcation.

A.4 Scale-space for discrete signals

Given a discrete signal $f : \mathbb{Z}^D \rightarrow \mathbb{R}$, its separable discrete scale-space representation $L : \mathbb{Z}^D \times \mathbb{R}_+ \rightarrow \mathbb{R}$ is defined as the result of convolving f by the D -dimensional discrete Gaussian kernel (Lindeberg 1990, Lindeberg 1994)

$$T_D(\xi; t) = \prod_{i=1}^D T_1(\xi_i; t), \quad (24)$$

where $\xi = (\xi_1, \dots, \xi_D)$, $T_1 : \mathbb{Z} \times \mathbb{R}_+ \rightarrow \mathbb{R}$ is the one-dimensional discrete analogue of the Gaussian kernel, $T_1(n; t) = e^{-t} I_n(t)$ and I_n is the modified Bessel function of integer order (Abramowitz & Stegun 1964).

Equivalently, this scale-space family can be obtained as the solution to the semi-discretized diffusion equation

$$\partial_t L = \frac{1}{2} \nabla_{2N+1}^2 L, \quad (25)$$

where

$$(\nabla_{2D+1}^2 L)(x; t) = \sum_{i=1}^D L(x + e_i; t) - 2L(x; t) + L(x - e_i; t), \quad (26)$$

and e_i denotes the unit vector in the i th coordinate direction. This is the scale-space concept that underlies all the implementations described in this presentation.

References

- Abramowitz, M. & Stegun, I. A., eds (1964), *Handbook of Mathematical Functions*, Applied Mathematics Series, 55 edn, National Bureau of Standards.
- Babaud, J., Witkin, A. P., Baudin, M. & Duda, R. O. (1986), ‘Uniqueness of the Gaussian kernel for scale-space filtering’, *IEEE Trans. Pattern Analysis and Machine Intell.* **8**(1), 26–33.
- Blostein, D. & Ahuja, N. (1989), ‘Shape from texture: integrating texture element extraction and surface estimation’, *IEEE Trans. Pattern Analysis and Machine Intell.* **11**(12), 1233–1251.
- Coulon, O., Bloch, I., Frouin, V., & Mangin, J.-F. (1996), The structure of PET activation maps, in ‘Proc. 3rd International Conference on Functional Mapping of the Human Brain’, Copenhagen, Denmark. Neuroimage vol. 5 no. 4 p 391, 1997.
- Coulon, O., Bloch, I., Frouin, V. & Mangin, J.-F. (1997), Multi-scale measures in linear scale-space for characterizing cerebral functional activations in 3D PET difference images, in B. M. ter Haar Romeny, L. M. J. Florack, J. J. Koenderink & M. A. Viergever, eds, ‘Scale-Space Theory in Computer Vision: Proc. First Int. Conf. Scale-Space’97’, Springer Verlag, New York, Utrecht, The Netherlands, pp. 188–199.
- Damon, J. (1996), Local Morse theory for Gaussian blurred functions, in J. Sparring, M. Nielsen, L. Florack & P. Johansen, eds, ‘Gaussian Scale-Space Theory: Proc. PhD School on Scale-Space Theory’, Kluwer Academic Publishers, Copenhagen, Denmark.

- Dupont, P., Orban, G., Vogels, R., Bormans, G., J., N., Schiepers, C., De Roo, M. & Mortelmans, L. (1993), 'Different perceptual tasks performed with the same visual stimulus attribute active different regions of the human brain – A positron emission tomography study', *Proc. Nat. Acad. Sci US* **90**, 10927–10931.
- Eberly, D., Gardner, R., Morse, B., Pizer, S. & Scharlach, C. (1994), 'Ridges for image analysis', *J. of Mathematical Imaging and Vision* **4**(4), 353–373.
- Ehrich, R. W. & Lai, P. F. (1978), Elements of a structural model of texture, in 'Proc. PRIP', IEEE CS Press, pp. 319–326.
- Florack, L. M. J. (1993), The Syntactical Structure of Scalar Images, PhD thesis, Dept. Med. Phys. Physics, Univ. Utrecht, NL-3508 Utrecht, Netherlands.
- Frackowiak, R. S. J., Zeki, S., Poline, J.-B. & Friston, K. J. (1996), 'A critique of a new analysis proposed for functional neuroimaging', *European Journal of Neuroscience* **8**, 2229–2231.
- Friston, K. J., Frith, C. D., Liddle & Frackowiak, R. S. J. (1991), 'Comparing functional PET images: The assessment of significant change', *J Cereb Blood Flow Metab* **11**, 81–95.
- Friston, K. J., Holmes, A., Poline, J.-B., Price, C. J. & Frith, C. D. (1996), 'Detecting activations in PET and fMRI: Levels of inference and power', *Neuroimage* **4**, 223–235.
- Friston, K. J., Worsley, K. J., Frackowiak, R. S. J., Mazziotta & Evans, A. C. (1994), 'Assessing the significance of focal activations using their spatial extent', *Human Brain Mapping* **1**, 210–220.
- Grasby, P. M., Frith, C. D., Friston, K. J., Simpson, J., Fletcher, P. C., Frackowiak, R. S. J. & Dolan, R. J. (1994), 'A graded task approach to the functional mapping of brain areas implicated in auditory-verbal memory', *Brain* **117**, 1271–1282.
- Gulyás, B. & Roland, P. (1994), 'Processing and analysis of form, colour, and binocular disparity in the human brain: Functional anatomy by positron emission tomography', *European J. of Neuroscience* **6**, 1811–1828.
- Hadjikhani, N. & Roland, P. (1998), 'Cross-modal transfer of information between the tactile and the visual representations in the human brain — A PET study'. Submitted.
- Holmes, A. P., Blair, R. C., Watson, J. D. G. & Ford, I. (1996), 'Non-parametric analysis of statistical images from functional mapping experiments', *J. of Cerebral Blood Flow* **16**, 7–22.
- Johansen, P. (1994), 'On the classification of toppoints in scale space', *J. of Mathematical Imaging and Vision* **4**, 57–67.
- Kapur, S., Craik, F., Jones, C., Brown, G., Houle, S. & Tulving, E. (1995), 'Functional role of the prefrontal cortex in retrieval of memories: A PET study', *NeuroReport* **6**, 1880–1884.
- Kinomura, S., Larsson, J., Gulyás, B. & Roland, P. (1996), 'Activation by attention of human reticular formation and thalamic intralaminar nuclei', *Science* **271**, 512–515.
- Koenderink, J. J. (1984), 'The structure of images', *Biological Cybernetics* **50**, 363–370.
- Koenderink, J. J. & van Doorn, A. J. (1986), 'Dynamic shape', *Biological Cybernetics* **53**, 383–396.
- Koenderink, J. J. & van Doorn, A. J. (1987), 'Representation of local geometry in the visual system', *Biological Cybernetics* **55**, 367–375.
- Koenderink, J. J. & van Doorn, A. J. (1994), 'Two-plus-one-dimensional differential geometry', *Pattern Recognition Letters* **15**(5), 439–444.

- Ledberg, A., Åkerman, S. & Roland, P. E. (1998), 'Estimation of probabilities of 3-D clusters in functional brain images', *Neuroimage* **6**. (in press).
- Ledberg & Roland (1996), 'Intersubject variation of cerebral activations in a somatosensory detection task', *NeuroImage* **3**(3), 330.
- Lidberg, P., Lindeberg, T. & Roland, P. (1996), Analysis of brain activation patterns using a 3-D scale-space primal sketch, in 'Proc. 3rd International Conference on Functional Mapping of the Human Brain', Copenhagen, Denmark, p. 393. *Neuroimage* vol. 5 no. 4 p 393, 1997.
- Lifshitz, L. & Pizer, S. (1990), 'A multiresolution hierarchical approach to image segmentation based on intensity extrema', *IEEE Trans. Pattern Analysis and Machine Intell.* **12**(6), 529–541.
- Lindeberg, T. (1990), 'Scale-space for discrete signals', *IEEE Trans. Pattern Analysis and Machine Intell.* **12**(3), 234–254.
- Lindeberg, T. (1992), 'Scale-space behaviour of local extrema and blobs', *J. of Mathematical Imaging and Vision* **1**(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. of Computer Vision* **11**(3), 283–318.
- Lindeberg, T. (1993b), 'Effective scale: A natural unit for measuring scale-space lifetime', *IEEE Trans. Pattern Analysis and Machine Intell.* **15**(10), 1068–1074.
- Lindeberg, T. (1993c), On scale selection for differential operators, in K. H. K. A. Høgdra, B. Braathen, ed., 'Proc. 8th Scandinavian Conf. on Image Analysis', Norwegian Society for Image Processing and Pattern Recognition, Tromsø, Norway, pp. 857–866.
- Lindeberg, T. (1994), *Scale-Space Theory in Computer Vision*, The Kluwer International Series in Engineering and Computer Science, Kluwer Academic Publishers, Dordrecht, Netherlands.
- Lindeberg, T. (1996), Edge detection and ridge detection with automatic scale selection, in 'Proc. IEEE Comp. Soc. Conf. on Computer Vision and Pattern Recognition, 1996', IEEE Computer Society Press, San Francisco, California, pp. 465–470.
- Logothetis, N., Pauls, J. & Poggio, T. (1995), 'Shape representation in the inferior temporal cortex of monkeys', *Curr. Biol.* **5**, 552–563.
- Meyer, E. (1989), 'Simultaneous correction for tracer arrival delay and dispersion in CBF measurements by the $H_2^{15}O$ autoradiographic method and dynamic PET', *J. Nucl Med* **30**, 1069–1078.
- Pizer, S. M., Burbeck, C. A., Coggins, J. M., Fritsch, D. S. & Morse, B. S. (1994), 'Object shape before boundary shape: Scale-space medial axis', *J. of Mathematical Imaging and Vision* **4**, 303–313.
- Poline, J.-B. & Mazoyer, B. (1993), 'Cluster analysis in individual functional brain images: Some new techniques to enhance the sensitivity of activation detection methods', *J Cereb Blood Flow Metab* **13**, 425–437.
- Poline, J.-B. & Mazoyer, B. (1994), 'Analysis of individual positron emission tomography activation maps by detection of high signal-to-noise-ratio pixel clusters.', *Human Brain Mapping* **2**, 103–111.
- Poston, T. & Stewart, I. (1978), *Catastrophe Theory and its Applications*, Pitman, London.

- Roland, P. E. (1993), *Brain Activation*, John Wiley and Sons, New York.
- Roland, P. E., Graufelds, J., Wåhlin, J., Ingelman, L., Andersson, M., Ledberg, A., Pedersen, J., Åkerman, S., Dabringhaus, A. & Zilles, K. (1994), 'Human brain atlas: For high-resolution functional and anatomical mapping', *Human Brain Mapping* **1**, 173–184.
- Roland, P. E. & Gulyás, B. (1996), 'Assumptions and validations of statistical tests for functional neuroimaging', *European. J. of Neuroscience* **8**, 2232–2235.
- Roland, P. E., Levin, B., Kawashima, R. & Åkerman, S. (1993), 'Three-dimensional analysis of clustered voxels in ^{15}O -butanol brain activation images', *Human Brain Mapping* **1**, 3–19.
- Roland, P. E. & Mortensen, E. (1987), 'Somatosensory detection of microgeometry, macrogeometry and kinaesthesia in man', *Brain Res. Rev.* **12**, 1–42.
- Roland, P. E. & Zilles, K. (1996), 'Functions and structures of the motor cortices in humans', *Curr. Opin. Neurobiol.* **6**, 773–781.
- Roland, P. & Larsen, B. (1976), 'Focal increase of cerebral blood flow during stereognostic testing in man', *Arch. Neurol.* **33**, 551–558.
- Seitz, R., Roland, P., Bohm, C., Greitz, T. & Stone-Elander, S. (1991), 'Somatosensory discrimination of shape-tactile exploration and cerebral activation', *Eur. J. Neurosci.* **3**, 481–492.
- Shallice, P., Fletcher, P., Frith, C. D., Grasby, D., Frackowiak, R. S. J. & Dolan, R. J. (1994), 'Brain regions associated with acquisition and retrieval of verbal episodic memory', *Nature* **368**, 633–635.
- Shindy, W., Posley, K. & Fuster, J. (1994), 'Reversible deficit in haptic delay tasks from cooling prefrontal cortex', *Cerebr Cortex* **4**, 443–450.
- Vincken, K., Koster, A. & Viergever, M. (1997), 'Probabilistic multiscale image segmentation', *IEEE Trans. Pattern Analysis and Machine Intell.* **19**(2), 109–120.
- Wienhard, K., Dahlbom, M., Eriksson, L., Michel, C., Bruckbauer, T., Pietrzyk, U. & Heiss, W. (1994), 'The ECAT-EXACT HR: Performance of a new high resolution positron scanner', *J. Comp. Ass. Tomogr* pp. 110–118.
- Witkin, A. P. (1983), Scale-space filtering, in 'Proc. 8th Int. Joint Conf. Art. Intell.', Karlsruhe, West Germany, pp. 1019–1022.
- Worsley, K. J., Evans, A. C., Marrett, S. & Neelin (1992), 'A three-dimensional analysis for CBF activation studies in human brain', *J. Cereb. Blood Flow Metab.* **12**, 900–918.
- Worsley, K. J., Marret, S., Neelin, P. & Evans, A. C. (1996b), 'Searching scale space for activation in PET images', *Human Brain Mapping* **4**, 74–90.
- Worsley, K. J., Marret, S., Neelin, P., Vandal, A. C., Friston, K. J. & Evans, A. C. (1996a), 'A unified statistical approach for determining significant signal in images of cerebral activation', *Human Brain Mapping* **4**, 58–73.
- Yuille, A. L. & Poggio, T. A. (1986), 'Scaling theorems for zero-crossings', *IEEE Trans. Pattern Analysis and Machine Intell.* **8**, 15–25.