NeuroImage **17,** 1665–1683 (2002)
doi:10.1006/nimg.2002.1304

fMRI Retinotopic Mapping—Step by Step

J. Warnking,* M. Dojat,* A. Guérin-Dugué,† C. Delon-Martin,* S. Olympieff,‡
N. Richard,* A. Chéhikian,‡ and C. Segebarth*

**UMR INSERM/UJF, U438 RMN Bioclinique, LRC CEA, Centre Hospitalier Universitaire, Pavillon B, BP 217, 38043 Grenoble cedex 9, France; †Laboratoire CLIPS, Bâtiment B, 385 rue de la Bibliothèque, BP 53, 38041 Grenoble cedex 9, France; and ‡Laboratoire des Images et des Signaux (LIS), INPG, 46 av. Félix Viallet, 38031 Grenoble cedex, France*

Received February 28, 2002

fMRI retinotopic mapping provides detailed information about the correspondence between the visual field and its cortical representation in the individual subject. Besides providing for the possibility of unambiguously localizing functional imaging data with respect to the functional architecture of the visual system, it is a powerful tool for the investigation of retinotopic properties of visual areas in the healthy and impaired brain. fMRI retinotopic mapping differs conceptually from a more traditional volume-based, block-type, or event-related analysis, in terms of both the surface-based analysis of the data and the phase-encoded paradigm. Several methodological works related to fMRI retinotopic mapping have been published. However, a detailed description of all the methods involved, discussing the steps from stimulus design to the processing of phase data on the surface, is still missing. We describe here step by step our methodology for the complete processing chain. Besides reusing methods proposed by other researchers in the field, we introduce original ones: improved stimuli for the mapping of polar angle retinotopy, a method of assigning volume-based functional data to the surface, and a way of weighting phase information optimally to account for the SNR obtained locally. To assess the robustness of these methods we present a study performed on three subjects, demonstrating the reproducibility of the delineation of low order visual areas. © 2002 Elsevier Science (USA)

Key Words: **human cortex; vision; visual areas; surface maps; retinotopy; functional magnetic resonance imaging.**

INTRODUCTION

The human visual cortex is divided into several functional areas with distinct local neural properties (Zeki and Shipp, 1988). The positions of functionally specialized visual areas are only loosely linked to cortical anatomy and are subject to variability between individuals (Amunts *et al.,* 2000). Several of these areas are retinotopic, that is, their neurons respond to stimulation of limited receptive fields whose centers are organized to form a continuous mapping between the cortical surface and the visual field. The boundaries between most of the low order visual areas can be determined from their retinotopic properties: the local representation of the visual field on the cortical surface changes its orientation—the local visual field sign (VFS)—between adjacent visual areas (Sereno *et al.,* 1994). Thus, the knowledge of retinotopy, mapped by fMRI (Engel *et al.,* 1994), allows for a precise delineation of some low order retinotopic visual areas (Sereno *et al.,* 1995).

Precise delineation presents multiple interests, such as in establishing intersubject and interspecies comparisons of the visual system (Van Essen *et al.,* 2001), in improving our insight into its organization in humans (Tootell *et al.,* 1997; Hadjikhani *et al.,* 1998; Tootell and Hadjikhani, 2001; Wade *et al.,* 2002), in allowing for quantitative investigations of parameters such as the cortical magnification factor (Sereno *et al.,* 1995) or receptive field size (Smith *et al.,* 2001), and in constraining source localization in EEG/MEG imaging (Di Russo *et al.,* 2002). Furthermore, it greatly enhances interpretation of the visual responses in numerous cognitive experiments (Wandell, 1999; Tootell *et al.,* 1998a), and it opens clinical perspectives in permitting detailed investigation of the pathologic visual system (Baseler *et al.,* 1999; Morland *et al.,* 2001).

fMRI retinotopic mapping differs in at least two respects from a more "traditional" three-dimensional amplitude-based functional analysis: the analysis of retinotopy requires the interpretation of functional data in their local spatial context of the sheetlike, highly folded cortical gray matter. This context is not obvious in the three-dimensional Cartesian space in which the data are acquired. It is usually provided by an explicit model of the individual cortical surface used in a surface-based analysis of the functional data. Second, due to the Fourier-type paradigm commonly used for fMRI retinotopic mapping, the main parameter of interest for the functional analysis is the delay (phase) of the observed response, not its amplitude. The processing of this information differs conceptually from an analysis based on the response amplitude alone.

The basic principle of fMRI retinotopic mapping using phase encoding stimuli has been the subject of several publications (Engel *et al.,* 1994; De Yoe *et al.,* 1994; Sereno *et al.,*

1053-8119/02 $35.00
© 2002 Elsevier Science (USA)
All rights reserved.

1995). Some methodological aspects have been addressed in reports of applications of the method: optimization of the duty cycle and the pattern of stimuli (Tootell *et al.,* 1997; Hadjikhani *et al.,* 1998) and the equivalence between cross-correlation and Fourier transform in the delay analysis (Engel *et al.,* 1997). Other aspects have been dealt with in detail, but independently of their application to retinotopic mapping: the construction of the cortical surface model, in particular the segmentation of structural MRI data (Wells *et al.,* 1996; Teo *et al.,* 1997; Van Leemput *et al.,* 1999b; Dale *et al.,* 1999; Zeng *et al.,* 1999; MacDonald *et al.,* 2000; Zhang *et al.,* 2001; Shattuck *et al.,* 2001), the flattening or unfolding of the cortical surface model (Van Essen and Maunsell, 1980; Schwartz *et al.,* 1989; Carman *et al.,* 1995; Drury *et al.,* 1996; Teo *et al.,* 1997; Fischl *et al.,* 1999; Angenent *et al.,* 1999; Hermosillo *et al.,* 1999; Wandell *et al.,* 2000; Guérin-Dugué *et al.,* 2000) and the correction of topological errors in the surface (Shattuck and Leahy, 2001; Fischl *et al.,* 2001; Kriegeskorte and Goebel, 2001; Han *et al.,* 2002), the analysis of fMRI response delays (Saad *et al.,* 2001), and the smoothing of functional data along the cortical surface (Andrade *et al.,* 2001). Despite these efforts, implementing fMRI retinotopic mapping procedures remains a difficult endeavor, partly because a detailed description of the complete process, discussing all the steps from stimulus design to the processing of phase data on the surface, is still missing.

In this paper, we point out the challenges involved in retinotopic mapping and give a detailed description of our methodology for the complete processing chain, leading up to the delineation of low order visual areas. We combine methods proposed by other researchers in the domain and original ones: improved stimuli for the mapping of polar angle retinotopy and a method of assigning volume-based functional phase data to the surface, including an optimal weighting accounting for the uncertainty of the phase estimation.

BACKGROUND

Figure 1 gives an overview of the procedures involved in fMRI retinotopic mapping and the corresponding neural and physiological processes.

Neural and Physiological Processes

The processes that link stimulation of the retina to a physiologic response in terms of local variations of blood oxygenation in the visual cortex comprise a multitude of steps, some of which are not yet fully understood. In the context of retinotopic mapping, it is useful to distinguish three stages in this processing chain.

Receptive fields. The stimulus is processed through multiple neural stages to yield an input to a given population of cortical retinotopic neurons. During this processing, information is integrated over a certain region in the visual field. Using the concept of receptive fields, this integration can be summarized in a sensitivity profile in the visual field for each neuron of the population considered. The transformation of the stimulus by the spatial integration can then be described as a spatial widening or "blurring" of the stimulus (Fig. 1a). The resulting "blurred" stimulus is a fictive intermediate state between the "real" stimulus and its projection on the cortical surface according to retinotopy (in reality, integration and projection by retinotopy occur simultaneously, across several processing stages). Receptive field properties vary between different populations of neurons, are generally not precisely known, and may even depend on the stimulus (Sceniak *et al.,* 1999).

Retinotopy. Retinotopy links the position of each retinotopic neuron to the point in the visual field corresponding to the center of its receptive field. In this context, the position of neurons is best described in terms of two-dimensional coordinates on the cortical surface, an idealized, two-dimensional representation of the cortical sheet (rather than in three-dimensional Cartesian coordinates). There are mainly two reasons. First, for a given point on the cortical surface, receptive fields of neurons from different cortical layers are centered on the same point in the visual field. And second, adjacent points on the cortical surface represent adjacent points in the visual field. More precisely, the mapping between the cortical surface and the visual field is locally homeomorphic, that is, it is locally bijective and continuous and the inverse mapping is also continuous.

In terms of steps involved in the processing of the stimulus, retinotopy transforms the "blurred" stimulus (a spatiotemporal pattern of stimulation in the visual field) into a corresponding spatiotemporal pattern of neural activation on the cortical surface. The goal of retinotopic mapping experiments is to determine this correspondence.

Neurophysiology and anatomy. The neurophysiologic properties of the cortex, notably its metabolic and hemodynamic properties, link a given spatiotemporal pattern of neural activation to spatiotemporal variations of blood oxygenation. These variations are at the origin of the signal observed in BOLD fMRI. While we can assume the neural processing described above to be quasi-instantaneous in the context of fMRI, the variation of blood oxygenation exhibits a temporal lag and is smoothed, temporally as well as spatially, with respect to the neural activation. The resulting variation of blood oxygenation will be measured in three-dimensional Cartesian space. Anatomy defines the way the surface-based response is embedded in three dimensions. The relationship between the surface-based and volume-based representations of the functional response is not homeomorphic. Due to the strong folding of the cortex, gray matter points that are adjacent in three dimensions (e.g., on opposite banks of a sulcus) are not necessarily close to each other when their distance is measured along the cortical surface.

Retinotopic Mapping Procedures

The challenge of retinotopic mapping by fMRI is to accurately measure the properties of the second of these stages, while being insensitive to, or correcting for, the others. In the following, we describe the steps necessary to obtain retinotopic maps by fMRI and the methodological challenges involved. We distinguish five steps in the retinotopic mapping procedure.

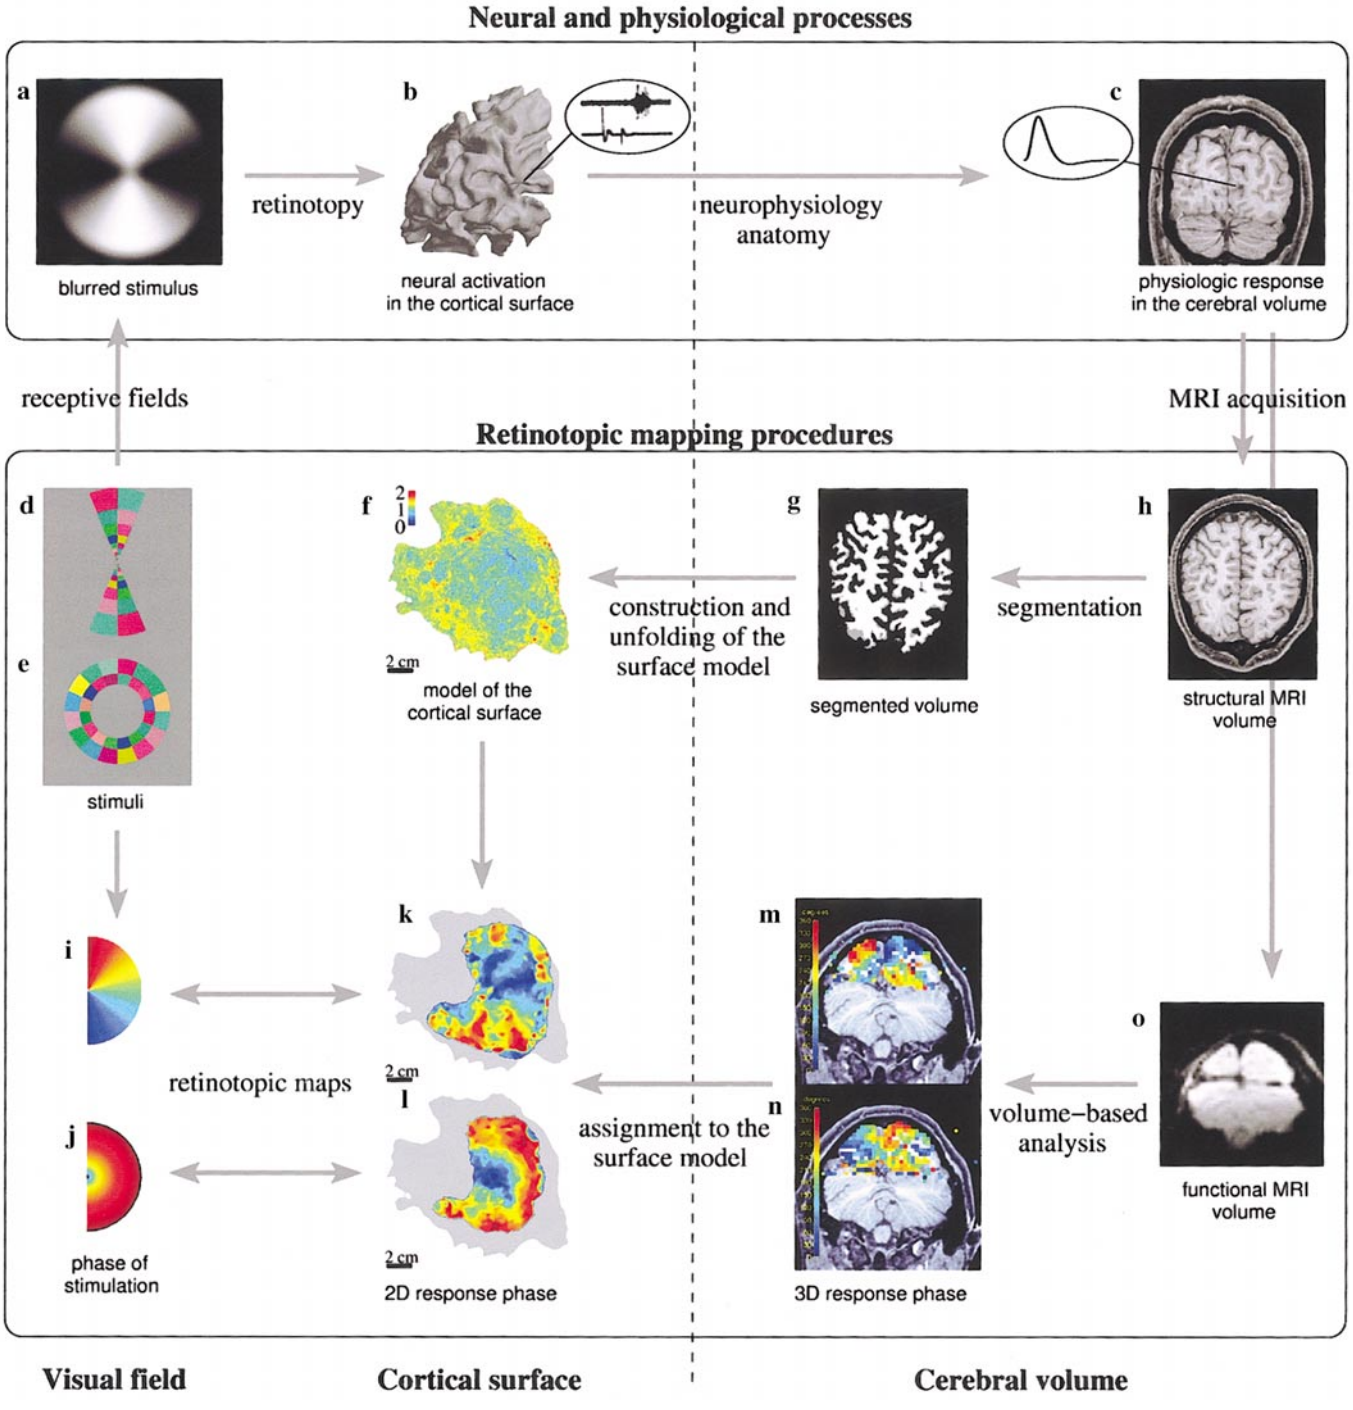

FIG. 1. Overview of fMRI retinotopic mapping procedures and the neural and physiological processes involved. Functional retinotopic mapping data are obtained in response to stimuli encoding the polar angle (d) and eccentricity (e) coordinates of the visual field. Concerning the neural processing of visual information, we separately consider the spatial integration of visual information, represented by the receptive field size of cortical neurons, and the point-to-point relationship between the centers of the receptive fields and the cortical surface, represented by retinotopy. The effect of receptive fields of cortical neurons can be described in this context as a spatial smoothing of the stimulus in the visual field (a). Retinotopy projects the "blurred" stimulus onto the cortical surface, where it gives rise to corresponding neural activation (b). Neurophysiological processes transform the neural activation into variations of blood oxygenation. Anatomy defines the way these variations are embedded in the three-dimensional Cartesian space of the cerebral volume (c) (note the difficulty of illustrating the "true" neural and physiological responses, and the "true" cortical surface or volume, as opposed to measured responses and reconstructed representations). A high-resolution structural MR volume is acquired (h) and segmented and the gray matter/white matter interface is dilated to approach the center of the cortex (g). A model of the cortical surface is reconstructed from the segmentation and is unfolded. The color coding overlaid here represents the distortion of local area induced in the flattening process, expressed as the ratio of triangle surfaces in 2D with respect to 3D. A ratio of 1 means no distortion (f). BOLD-sensitive functional MR volumes are acquired (o) and analyzed in 3D to obtain for each voxel an estimation of the response phase, encoding the locally represented visual field polar angle (m) and eccentricity (n). These data are assigned to the surface model obtained previously and smoothed slightly ($\sigma = 1.5$ mm). The response phase to polar angle and eccentricity stimuli is displayed on the flattened representation of the cortical surface (k and l, respectively). Together with the information about the corresponding phase of stimulation (i and j), they form the desired retinotopic maps.

Acquisition and segmentation of structural data. In order to obtain a two-dimensional map of retinotopy, an explicit model of the cortical surface is required. The first step in the construction of this model is the acquisition of a high-resolution, high contrast-to-noise structural MR image of the brain. This volume is subsequently segmented to obtain representations of the different brain tissues. Accurate and automatic segmentation of the structural volume to obtain a topologically correct representation of the cortical surface is a major challenge. In the context of construction of a model of the cortical surface, several descriptions of brain segmentation algorithms exist (Teo *et al.,* 1997; Dale *et al.,* 1999; Joshi *et al.,* 1999; Zeng *et al.,* 1999; Germond *et al.,* 2000; MacDonald *et al.,* 2000; Shattuck *et al.,* 2001). Most commonly, voxels are labeled as one of three tissue types: white matter (WM), gray matter (GM), or cerebrospinal fluid (CSF).

Construction and unfolding of the surface model. A model of the cortical surface is extracted from the segmented volume, based on one or both of the borders between cortical GM and the adjacent tissues. Ideally, it represents the center of the cortical GM, where functional activation is expected. The ensuing surface-based analysis of functional data requires that the surface model be anatomically and topologically correct. Topological defects usually need to be corrected manually, although recently efforts have been made to automatically correct some of those errors (Shattuck and Leahy, 2001; Fischl *et al.,* 2001; Kriegeskorte and Goebel, 2001; Han *et al.,* 2002).

A major advantage of a surface-based analysis—convenient display—can only be reaped by unfolding the surface model, a procedure presenting a challenge in terms of computational complexity. The unfolding needs to be homeomorphic to be useful. This means in particular that partially folding the surface onto itself in the flattened representation must be avoided. Homeomorphic flattening of the surface is only possible if the folded model is free from topological errors.

Stimuli for retinotopic mapping. Stimuli for retinotopic mapping are designed to encode the position in the visual field by a unique pattern of temporal activation. This is achieved by means of slowly moving periodic stimuli consisting of concentric expanding or contracting rings and clockwise or counterclockwise rotating wedges, presented while the subject is fixating their center or apex (Engel *et al.,* 1994; DeYoe *et al.,* 1994). These stimuli link each position along a visual field coordinate (eccentricity/polar angle) to a unique delay of the periodic stimulation. This delay is usually quantified as a phase in the frequency domain. This encoding is robust with respect to the (unknown) spatial and temporal smoothing applied by the visual fields and the hemodynamic response, provided the temporal and spatial frequency of the stimuli is low enough. The hemodynamic delay creates a phase shift of the response that needs to be corrected for. This is done by comparing the responses to two stimuli moving in opposite directions for each visual field coordinate (four stimuli in total) (Sereno *et al.,* 1995).

Acquisition and volume-based analysis of functional data. BOLD-sensitive fMRI data are acquired throughout the occipital lobe. The MR images usually present distortions with respect to the structural data acquired, due to susceptibility artifacts and gradient nonlinearities. As discussed below, misalignment between functional and structural images can severely degrade the two-dimensional representation of the functional data. One of the challenges in the acquisition of fMRI data is to minimize susceptibility artifacts, while maintaining a high level of sensitivity to BOLD contrast and short acquisition times. Distortions due to gradient nonlinearities are independent of the acquisition sequence. Depending on the gradient hardware, appropriate correction of the data may be necessary (Wald *et al.,* 2001).

The volume-based analysis in retinotopic mapping experiments involves estimating the response phase for all voxels. Importantly, the individual uncertainties of these phase estimations can be quantified and can be taken into account in the subsequent analysis. The responses of the stimuli of opposite directions of movement are combined to correct for the phase shift induced by the hemodynamic delay.

Assignment of functional data to the surface model. The task is to obtain a surface-based representation of the cortical response from the functional data acquired in three-dimensional Cartesian space. Two issues need to be addressed. The first concerns the problems inherent in the reduction of dimensionality, assigning volume-based data to the surface model. This step is necessarily nonhomeomorphic. In the presence of misalignment (local and/or global) between functional and structural data, the assignment may induce large errors in the two-dimensional representation of the data. These errors are best exemplified by the case of assigning functional data to the wrong bank of the calcarine sulcus. The issue of alignment needs to be addressed at the moment of data acquisition or by appropriate correction of distorsions prior to assignment of data to the surface model. The second issue is the potential mismatch between the data at the individual voxels and the original cortical response, due to noise and the distance between voxel centers and surface elements. Sources of this mismatch need to be identified and their respective contributions estimated and taken into account in the context of the assignment of phase information.

Processing of the Retinotopic Maps

Retinotopic mapping is usually not a goal in itself. Consequently, the obtained maps are to be processed further to extract information such as the position of the borders between retinotopic functional areas. We will use the example of delineation of visual areas to illustrate and evaluate our retinotopic mapping procedures.

Between adjacent retinotopic visual areas the visual field sign (VFS, Sereno *et al.,* 1994) changes, which allows for a reconstruction of their borders (Sereno *et al.,* 1995). The visual field sign designates the orientation of the representation of the visual field on the cortical surface. To determine the visual field sign, it is convenient to calculate the ratio of an oriented area measured using the local representation of the visual field coordinates with respect to the same area measured using a locally isometric parametrization of the

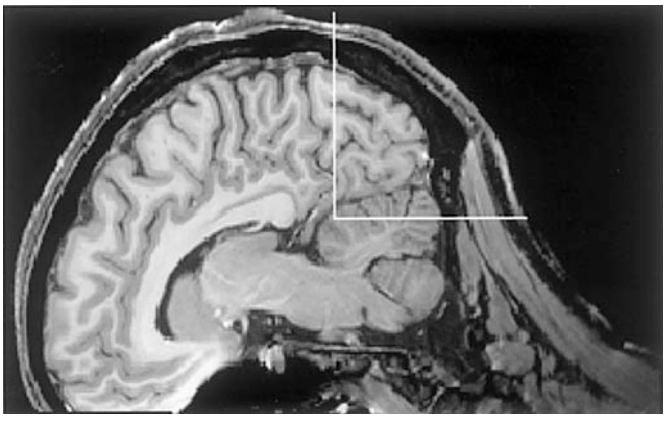

FIG. 2. Schematic display of the part of the cortical surface selected for unfolding.

surface.[1] We refer to this quantity as the visual field ratio (VFR). The visual field sign is then the sign of the VFR, and the visual area borders correspond to contour lines of zero VFR. The size of the zone of small absolute VFR around visual area borders gives an immediate visual impression of the uncertainty of the position of the delineated borders, information that is absent from the VFS. Most of the spatial features of the VFR are present in the representation of the polar angle coordinate whose gradient reverses direction at the borders between visual areas. In contrast, the eccentricity gradient is smooth across visual area borders.

Two-dimensional processing of the retinotopic maps presents a computational challenge linked to the representation of the data on an irregular two-dimensional grid embedded in three dimensions. Standard image processing approaches are therefore not always easily implemented.

METHODS

Acquisition and Segmentation of Structural Data

Acquisition and data preprocessing. All MRI data were acquired on a Philips Intera 1.5 T system equipped with a Powertrack 6000 gradient system (23 mT/m with a slew rate of 105 T/m · s). Structural data were acquired by means of a spoiled 3D GRE Flash sequence, TR of 23.7 ms, TE of 6.9 ms, flip angle of 28°, and an isotropic resolution of 1 mm. The body coil was used for RF excitation and a volume head coil for signal detection. Slices were oriented approximately parallel to the calcarine sulcus, inclined by about 45° with respect to AC–PC. To optimize the contrast-to-noise ratio, three volumes of $256 \times 256 \times 160$ voxels were acquired in the same scanning session for a total acquisition time of 35 min. Head motion was constrained by means of small sandbags to the right and left of the subject's head. Residual motion was corrected for by realigning the three volumes using the SPM software (Ashburner and Friston, 1997). In the ensuing analysis, only the mean of the realigned structural images was used.

Segmentation. The details of the algorithm employed are beyond the scope of this paper. Briefly, voxels are labeled sequentially, starting with CSF and proceeding to GM and WM. The image intensity distributions of the tissues are modeled as normal distributions (Wells *et al.*, 1996). Their parameters are estimated separately for each slice to take into account inhomogeneities in the *z* direction. More sophisticated techniques incorporating Expectation-Maximization and Markov random field models could be introduced to better account for the full three-dimensional bias field distortion (Van Leemput *et al.*, 1999a; Zhang *et al.*, 2001; Shattuck *et al.*, 2001). The sulci are initially detected as dark and narrow regions using a morphological operator (Guérin-Dugué *et al.*, 2000). This allows for an estimation of their intensity distribution and subsequent refinement of their labeling. In a similar fashion, voxels close to sulci are considered belonging to GM, providing an initial estimation of their intensity distribution. This labeling is in turn refined in a region growing process seeded around the sulci, aggregating voxels based on their intensity and on the topology of their neighborhood. Using the same procedure, WM is segmented starting from voxels close to the GM. After this first segmentation, voxel labeling is refined over several cycles, iteratively updating the estimation of the intensity distributions for all tissues and the voxel labeling. The whole process is controlled by two parameters, one indicating the maximum intensity to take into account and the other tuning the attribution of voxels that have borderline intensities between GM and WM. For a given acquisition sequence, these parameters usually need not be readjusted manually.

After the segmentation, the interface between the volumes labeled GM and WM is extended to represent approximately the center of the GM by a series of constrained region growing steps applied to CSF and WM. At each of these steps, only voxels initially labeled as GM can be reaffected, and a layer of at least one voxel of GM is imposed between WM and CSF. Many of the topological defects initially present in the volume labeled as WM disappear during this postprocessing step. However, to obtain a topologically correct model of the occipital lobe, manual editing is still required.

Construction and Unfolding of the Surface Model

The first step in the construction of a model of the cortical surface is the selection of the brain region to be represented. This depends on the region studied, but also on the way the surface should be visualized: by flattening or inflation. Inflation allows for the representation of an entire hemisphere without cutting, but not all of that surface is visible at once. Flattening displays all of the model at once, but requires surface cuts if the intrinsic curvature of the surface is too strong, thus loosing in that case some of the connectivity information. In the context of studies pertaining to the retinotopic visual areas, we found it useful to model only part of the cortical surface situated in the occipital lobe. This allows us to completely flatten this surface without the need for further cuts. The entire region under investigation can then

[1] This is the Jacobian of the visual field representation on the surface.

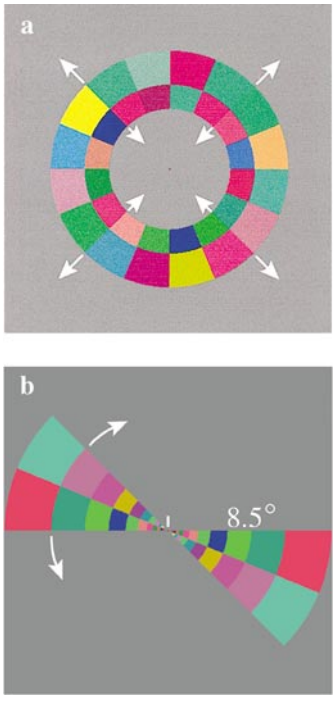

FIG. 3. Stimuli used for mapping retinotopy with respect to eccentricity (a) and polar angle (b). Four stimuli are presented: expanding and contracting rings and clockwise and counterclockwise rotating wedges. The subject is asked to fixate the central dot while maintaining attention to the stimulus and reporting the occurrence of a pair of bright yellow/dark blue checks that appear at random moments and positions.

be visualized simultaneously, while faithfully representing connectivity information throughout the surface.

The portion of the surface to be unfolded is defined manually for each of the two hemispheres from the segmented volume. It is delimited by two perpendicular planar cuts. One is made approximately parallel to, and just posteroventral of, the parietooccipital sulcus, and the other is approximately parallel to, and about 3 cm anteroventral of, the calcarine sulcus (Fig. 2). Within the delimited region, a triangulated model of the interface between voxels labeled WM and GM is created using the marching cubes algorithm (Lorensen and Cline, 1987).[2] Due to the postprocessing steps applied to the segmented volume, this model approximately represents the center of the GM. Its nodes are initially positioned on a regular grid. The model is subsequently smoothed slightly by iteratively displacing each node a fraction of the distance to the mean position of its nearest neighbors. Ten iterations of smoothing are applied, displacing nodes at each iteration a tenth of the distance to the center of its neighbors.

The flattening algorithm employed is described elsewhere (Guérin-Dugué *et al.*, 2000). Further details will be provided in a forthcoming paper. This algorithm is a modified version of a multidimensional scaling like algorithm, called Curvilinear Components Analysis (Demartines and Herault, 1997).

[2] Slight modifications from the original algorithm were made to avoid holes in the surfaces generated.

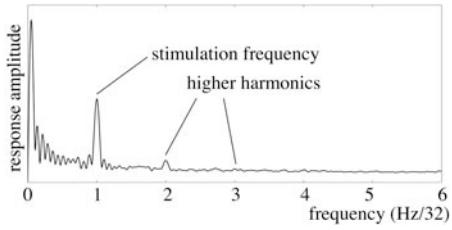

FIG. 4. Average spectral response observed at voxels activated above a threshold of SNR > 4.5 during one functional run of a retinotopic mapping experiment.

Briefly, the approximate geodesic distances from each node to all its neighbors within a tenth order neighborhood are calculated. These distances provide information about the local structure of the surface. Information about the global structure is provided by 10 "representative" nodes that are selected automatically using the K-means technique (MacQueen, 1967). These nodes serve as "anchors." For each of the anchors, the distances to all the other nodes of the surface are determined using the Dijkstra algorithm (Dijkstra, 1959). This leads to a sparse distance matrix, which contains only about 2% of all the mutual distances between nodes. The unfolding is initialized by a projection of all nodes from their 3D positions to the plane formed by the two first principal components of their spatial distribution. The 2D node positions are then iteratively updated during 5000 iterations. Nodes are selected one at a time and all the neighbors of the node currently chosen are repositioned, according to the mismatch observed between the distances in the plane and the geodesic distances. At each iteration, each of the anchors is selected once, updating the global structure of the unfolded

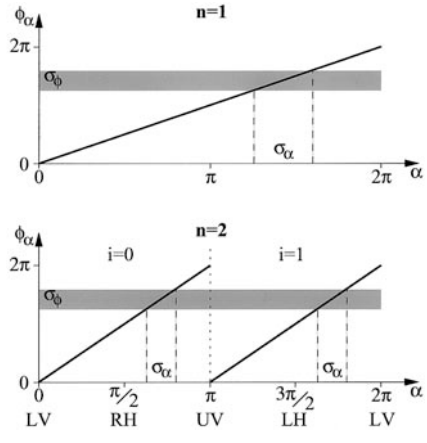

FIG. 5. Illustration of the response phase ϕ_α as a function of visual field position (polar angle α) for two different stimuli: one single rotating wedge ($n = 1$) and two opposing rotating wedges ($n = 2$). Positions of the visual field meridians (lower vertical, LV, right horizontal, RH, upper vertical, UV, and left horizontal, LH) are indicated. The decoding is not unique in the case $n = 2$. *A priori* information is needed about which hemifield is locally represented ($i = 0$, right hemifield; $i = 1$, left hemifield). The uncertainty of the estimation of visual field position (σ_α) as a function of the uncertainty of response phase estimation (σ_ϕ) is indicated by the dashed lines. Uncertainty is reduced by a factor of 2 in the case $n = 2$.

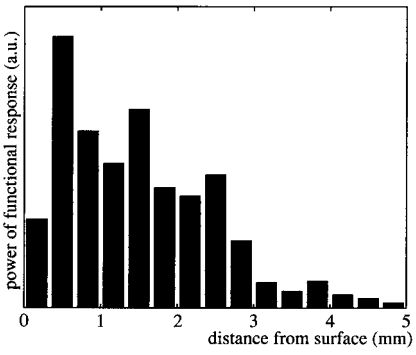

FIG. 6. Histogram of the functional response power (subject M.D., reference data set, left hemisphere, polar angle scan) as a function of distance from the surface model. Only voxels with an SNR exceeding 2 were included in the analysis.

model, and a fixed number of "ordinary" nodes are selected randomly, refining the local structure. The global structure of the unfolded model settles rapidly. After 500 iterations, the anchors are dropped from the calculations, and node positions are only refined locally. The method proposed is computationally efficient mainly due to two aspects: first, the distance matrix is sparse and therefore only a fraction of the mutual distances needs to be calculated. And second, for each node selected during the unfolding process, the positions of all neighbors are updated. This contributes to a higher mobility of the nodes compared to an approach displacing only the selected node, like classical stochastic gradient descent techniques. As a result, the model converges rapidly toward a low-distortion representation of the unfolded surface, at the same time avoiding local minima that would lead to the surface being folded onto itself. The flattening produced is homeomorphic over the vast majority of the surface area.

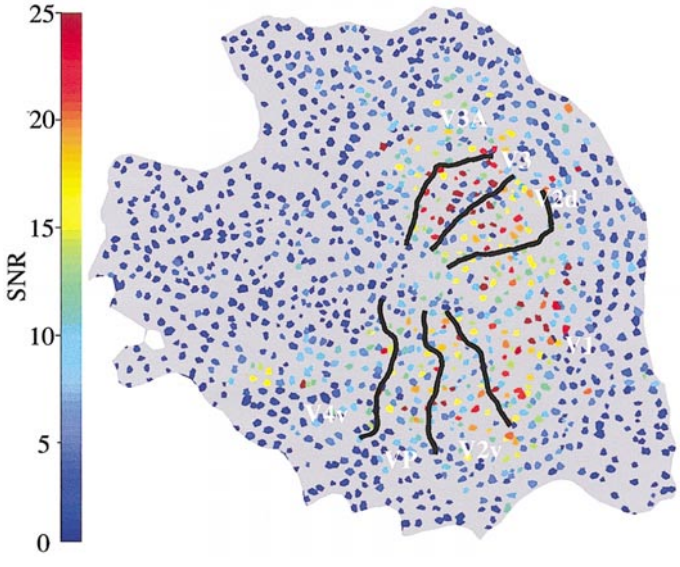

FIG. 7. SNR of functional information (subject M.D., reference data set, left hemisphere, polar angle scans) projected to the surface. Only voxels exceeding a threshold of SNR > 2 have been projected. Each triangle is colored according to the maximum SNR at any of its nodes. The maximum SNR observed was above 40, but a ceiling of 25 was imposed for display purposes. The automatically delineated borders between retinotopic visual areas are superposed for orientation.

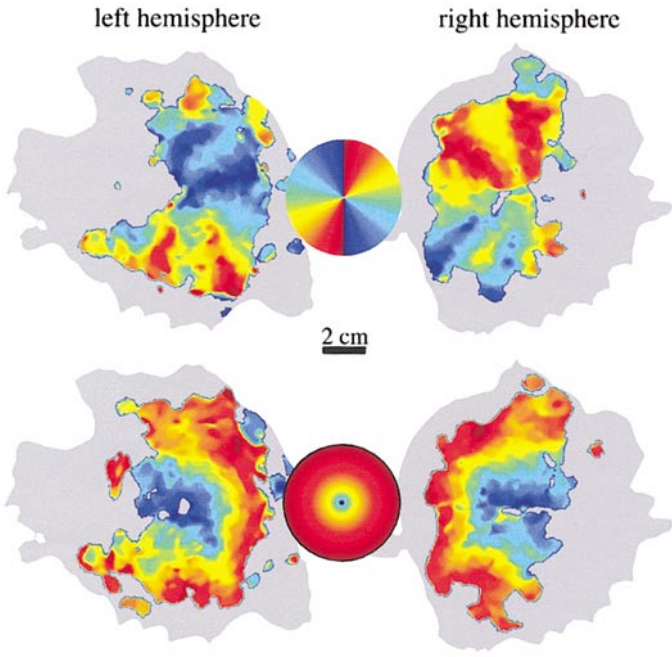

FIG. 8. Final results of the assignment of phase data to the surface (subject M.D., reference data set, both hemispheres, polar angle and eccentricity scans). The color legends represent a circular area of the visual field with an eccentricity of 8.5°. The processing details were those described in the text. The phase data displayed are smoothed slightly (σ = 1.5 mm, corrsponding to a FWHM of 3.53 mm) at assignment. The mask applied is derived from the corresponding amplitude data (assigned using σ = 2.5 mm; FWHM, 5.89 mm) applying a threshold of SNR > 12. Note that the amplitude data represent the effective SNR after smoothing and not the average SNR (see Appendix).

The step of unfolding necessarily induces some amount of distortion of distances and/or angles, because intrinsic curvature is present in the cortical surface. As far as display purposes are concerned, this distortion is a small price to pay to be able to view the entire surface under investigation at once. The functional analysis can be performed entirely based on the folded model and is therefore not hampered by this distortion.

Stimuli for Retinotopic Mapping

The stimuli we used are similar in design to those described by other groups (Engel *et al.*, 1994; Sereno *et al.*, 1995; Tootell *et al.*, 1997). Eccentricity is mapped by a slowly expanding or contracting ring, and polar angle is mapped by two rotating wedges (Fig. 3). For the stimuli mapping eccentricity, the speed of expansion or contraction varies linearly with eccentricity (exponentially with time). Due to the approximately exponential cortical magnification factor (Engel *et al.*, 1994; Tootell *et al.*, 1998b) this stimulus produces a wave of activation on the cortical surface traveling at approx-

imately constant speed. When the ring reaches maximum eccentricity, it wraps around to be replaced by a new one at minimum eccentricity, and vice versa.

When optimizing these stimuli, essentially four parameters can be adjusted independently: temporal frequency, the number of rings and wedges, the duty cycle of the stimulation, and the pattern (color, contrast) of the stimuli.

Temporal frequency. The temporal frequency is limited by two principal constraints: the presence of strong low-frequency noise (baseline drift) and the low-pass filtering due to the hemodynamic response. We chose a temporal period of 32 s for both polar angle and eccentricity stimulation. This period allows for a full return to baseline between activations and thus leads to maximal response amplitude (Bandettini *et al.,* 1993; Friston *et al.,* 1994). At the same time, the frequency is high enough to avoid the strong low-frequency noise caused, for example, by subject movement and scanner instability (Fig. 4). Physiologic noise from cardiac pulsation and respiratory events also needs to be considered. Both processes occur at frequencies sufficiently high not to interfere directly with the stimulation frequency. But due to the discrete temporal sampling of functional data, these signals may be aliased back into the spectrum. In our case of a repetition time of 1.28 s for functional volumes, typical (fundamental) frequencies for respiration are not aliased. The aliased cardiac noise does not coincide with the stimulation frequency for cardiac frequencies in a range from 50 to 90 pulsations per minute.

Number of rings/wedges. The number of rings or wedges the stimulus contains determines the correspondence between the observed response phase and the position (eccentricity or polar angle, respectively) in the visual field. This correspondence is one of the factors determining the accuracy of the final retinotopic map, the other one being the accuracy of the response phase measurement. We will describe this correspondence for the case of the polar angle stimulus, the case of the stimulus mapping eccentricity being similar. A polar angle stimulus consisting of n equally spaced rotating wedges can be described by the phase of the periodic stimulation

$$\tilde{\phi}_\alpha = (n\tilde{\alpha}) \bmod 2\pi$$

corresponding to each polar angle $\tilde{\alpha} \in [0, 2\pi)$ of the visual field. This stimulus creates a unique phase encoding of polar angles within each of n regions of the visual field, subtending $2\pi/n$ of polar angle each. If $n > 1$, the encoding of the entire visual field is not unique since any two visual field positions that are an integer multiple of $2\pi/n$ apart have identical phases (they are stimulated simultaneously). Let $\phi_\alpha \in [0, 2\pi)$ be the estimated phase of the response. The corresponding estimated position of stimulation in the visual field can then be calculated as

$$\alpha = \frac{\phi_\alpha}{n} + i\,\frac{2\pi}{n},$$

where $i \in (0 \ldots n - 1)$ indicates which of the n uniquely stimulated regions of the visual field is represented locally (Fig. 5). The estimation of i needs to be based on *a priori* information about retinotopy. Assuming that i can be determined accurately, the uncertainty σ_α of the estimated visual field position due to the uncertainty σ_ϕ of the phase measurement is then given by

$$\sigma_\alpha = \frac{\sigma_\phi}{n}.$$

Thus, the accuracy of the polar angle retinotopic map is proportional to the number of wedges, provided the accuracy of phase estimation is constant and provided sufficient *a priori* knowledge is available to unwrap the response phase. More generally, the accuracy is proportional to the local derivative of the function translating the response phase observed into a visual field position.

Note that the choice of the number of wedges (rings) links the temporal frequency of stimulation to an angular (radial) speed of the rotation (expansion or contraction). An increase of the number of elements leads to a slower displacement of the stimulus in the visual field at constant temporal frequency.

Retinotopic stimuli with more than one element have been used before. Engel *et al.* (1994) used two rings in their original report on the Fourier method for retinotopic mapping. Later, the same group used three wedges to map polar angle (and a single ring to map eccentricity) (Engel *et al.,* 1997). In a recent report, the group uses the standard stimuli comprising a single ring and a single wedge (Press *et al.,* 2001). To our knowledge, the relationship between the number of rings/wedges and the accuracy of the retinotopic maps obtained has not been mentioned.

We chose to use two wedges for the polar angle stimulus, because distinguishing between the two visual hemifields using *a priori* information is particularly simple. Only the contralateral hemifield is represented in the low order visual areas of each hemisphere, so that positions that are sufficiently far from the vertical meridians can be uniquely identified. Positions close to the vertical meridians can be correctly attributed, since retinotopy varies smoothly over the cortical surface (Sereno *et al.,* 1994), and the representations of the upper and lower vertical meridians are surrounded by representations of the upper and lower quadrants, respectively.

Retinotopy with respect to eccentricity is less important for the delineation. It contains little information about the position of the visual area limits. Since distinguishing between the responses to several rings using *a priori* information seems to introduce more difficulties than a higher accuracy would resolve, we chose to use only a single ring stimulus.

Duty cycle. A case has been made for stimuli with a very low duty cycle, to improve the responses in visual areas where neurons have large receptive fields (Tootell *et al.,* 1997). We found that the amplitude of the response at the fundamental stimulation frequency decreased for very thin stimuli and obtained better results with somewhat higher duty cycles of 25% for the polar angle stimuli (two wedges subtending 45° each) and 17% for the eccentricity stimuli

(one ring, two checks wide). This corresponds to a width of two checks of the radial checkerboard, creating a line of high local contrast at the center of the stimulus.

Pattern. Rings and wedges consist of a radial checkerboard. The aspect ratio of the checks is kept constant by scaling their height linearly with eccentricity. In order to maximize local luminance and color contrast, neighboring checks are of complementary color. The stimulus had the same mean luminance as the gray background. Equiluminance values for three color combinations (red/green, blue/red, and green/blue) are measured for each subject using a minimum of motion test (Anstis and Cavanagh, 1983) and are used for individual luminance correction. Colors change at a frequency of 4 Hz. During all scans, the subject's task is to fixate a central dot on the screen, while focusing attention to the stimulus. To control and maintain attention, the subject is asked to press a button at each occurrence of a given pair of yellow and blue checks, which appears at random positions and times in the stimulus at a mean frequency of one target every 6 to 8 s.

Stimuli are back-projected on a translucent screen situated outside the magnet. The subject views this screen at a distance of about 150 cm via an angled mirror. Our stimuli cover eccentricities from 0.2° to 8.5°. The movement of the stimuli occurs in small steps of four images per second and appears almost smooth. Four retinotopic functional scans are acquired, one for each of the two directions of motion for each of the two stimulus types. The stimuli start 10 s before the actual acquisition begins, to be able to detect responses from the beginning of the acquisition period. The start of the stimuli is triggered by a signal from the scanner.

Acquisition and Volume-Based Analysis of Functional Data

Acquisition. The goal in the choice of an acquisition sequence for functional data based on BOLD contrast is to provide fast T_2^*-sensitive imaging, while reducing the distortion often present in single shot EPI sequences due to the narrow "bandwidth" in the phase encoding direction. A means of reducing loss of phase coherence during the echo train is to shorten it, using segmented EPI. However, at fixed echo time, required for T_2^* sensitivity, this leads to an unused delay between the RF pulse and the start of the echo train. Thus, measurement time is to a certain extent wasted, at the expense of the SNR achievable during a fixed measurement time. The 3D PRESTO sequence shifts the echoes, acquiring during this lapse the echo train corresponding to the preceding RF pulse (Liu *et al.,* 1993; van Gelderen *et al.,* 1995). The expected relative distortions between functional and structural data due to susceptibility artifacts are reduced by a factor of 4.5 with respect to single shot EPI, while the minimal scan time remains essentially unchanged. This reduction of distortion comes at the cost of a slightly lower sensitivity due to the decrease in repetition time and a corresponding decrease in longitudinal magnetization. At the same time, the decreased repetition time has the benefit of reducing the signal contributions from large blood vessels presenting a long T_1 relaxation time, leading to a lower sensitivity to macrovascular artifacts.

Functional data were acquired in the same scanning session as the structural scans; 12 slices oriented approximately perpendicular to the calcarine sulcus were scanned by means of a 3D PRESTO sequence, acquiring 21 echoes for each RF pulse, with a repetition time of 28 ms, an echo time of 40 ms, a flip angle of 14°, and a resolution of $3 \times 3 \times 4$ mm^3. A single loop surface coil, positioned inside the volume coil used for the structural scans, was used for signal reception. The body RF coil was used for excitation. The acquisition time was 7 min 16 s per functional run containing 341 volumes of $64 \times 64 \times 12$ voxels. Thus, 25 functional volumes are acquired during each period of stimulation.

Data import and preprocessing. Functional volumes are converted from the proprietary scanner image format to SPM/Analyze, taking into account relative position at acquisition of the functional volumes with respect to the structural data. To faithfully represent subvoxel offsets and rotations between functional and structural images, a linear spatial transformation matrix is stored for each acquired volume. This transformation matrix is taken into account in all subsequent processing steps. We consistently observed an overall shift of the functional images with respect to the structural volumes by about one voxel in phase encoding direction (right–left). This shift is corrected for by simple translation of the functional volumes.

If gradient nonlinearity is a problem, appropriate correction might be needed at this stage to allow for accurate assignment of the functional data to the surface model (Wald *et al.,* 2001). In our case, gradient nonlinearity did not noticeably affect images. According to the figures provided by the manufacturer of our gradient system, nonlinearity does not exceed 1% over a volume of 25 cm diameter (1.4% over 53 cm).

Finally, some of the effects of head movement during functional scans are removed by realigning the functional volumes with respect to the one acquired closest in time to the structural scan using the SPM package (Friston *et al.,* 1995). Data are resampled using a windowed sinus cardinal kernel. Since the acquisition is done with a 3D technique, rather than with a multiple slice method (as is usually done with EPI acquisitions), there is no difference in timing of the acquisition of the slices to be taken into account.

3D analysis of retinotopic mapping data. The estimation of the response phase and amplitude is done in three dimensions on a voxel-per-voxel basis. With the timing used here, the response to the periodic stimulus is close to sinusoidal, with only a small amount of energy present in higher harmonics (Fig. 4). However, the amplitude of the signal present at those harmonics varies greatly between individual voxels and is generally very low. We chose to base our analysis on the fundamental frequency only.

Both, signal amplitude and phase can be calculated from the complex valued Fourier transform at the stimulation frequency (Engel *et al.,* 1997):

$$F_{\nu_0}(\vec{x}_j) = \sum_{k=1}^{N} f(\vec{x}_j, t_k) \exp(\pm i 2\pi \nu_0 (t_k - t_H)) , \qquad (1)$$

where F_{ν_0} is the volume of complex Fourier components at the frequency ν_0 of the stimulus, $\vec{x}_j$ are the voxel positions,

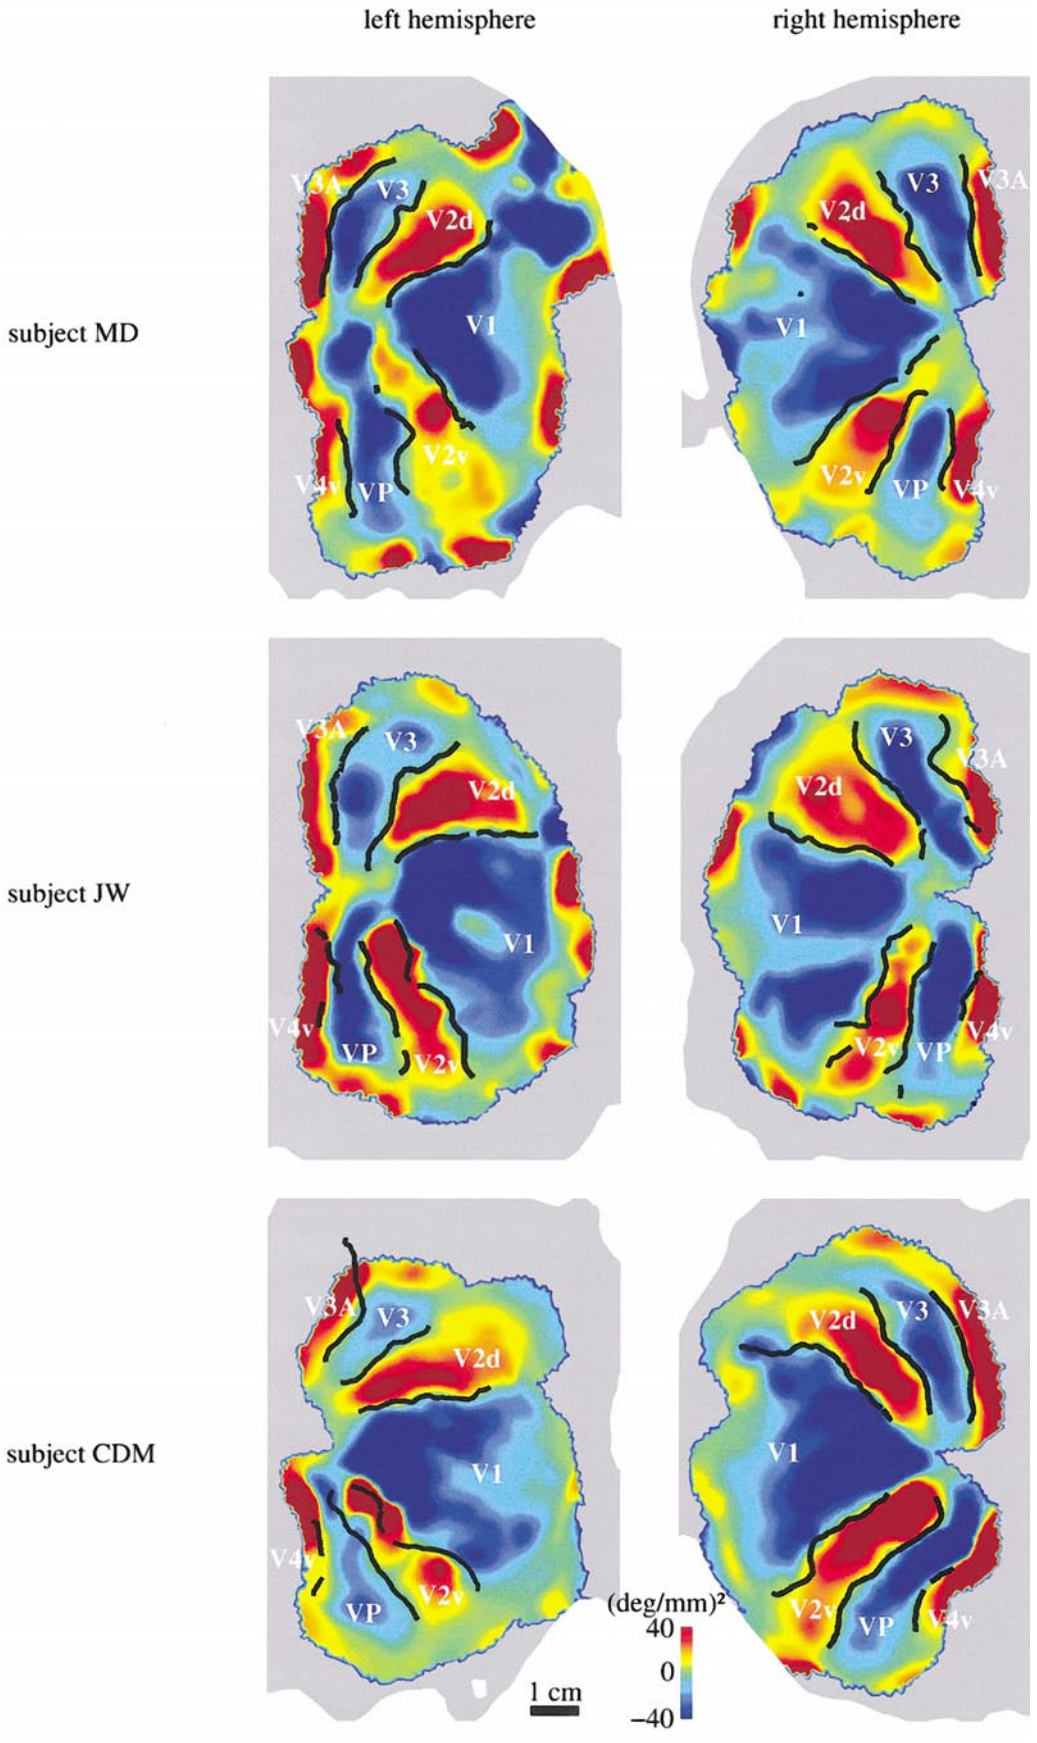
left hemisphere
right hemisphere
subject MD
subject JW
subject CDM
V3A
V3
V2d
V1
V2v
VP
V4v
(deg/mm)²
40
0
−40
1 cm

N is the number of volumes acquired during one functional run, f is the raw functional data (three spatial and one temporal dimensions), t_k are the instances of acquisition of the functional volumes counted from onset of the acquisition, and t_H is an estimation of the mean hemodynamic delay. Since the phase of F_{ν_0} is to be a measure of the position of the stimulus in the visual field, the sign in the exponent depends on the direction of motion of the stimulus. We chose the positive sign for the expanding and counterclockwise rotating stimuli. Furthermore, we assume that the stimulus position at the beginning of the scan is to be associated with a response phase of zero; otherwise an additional phase offset is required in the exponential function.

For the ensuing analysis the absolute amplitude of the response is of little interest *per se*. Rather, we are interested in the response amplitude because it reflects the uncertainty of the phase measurement. The standard deviation of the phase error (when expressed in radians) is the inverse of the SNR of the response amplitude at the stimulation frequency (see Appendix). The noise at the stimulation frequency is not accessible directly. Instead, we base its measurement on the assumptions that noise above the stimulation frequency is approximately white and that the response observed at frequencies other than the stimulation frequency contains only noise. Consequently, the noise is estimated for each voxel independently as the standard deviation of the real and imaginary spectra over a range of frequencies. Frequencies below the stimulation frequency (containing baseline drifts) and all harmonics of the stimulation frequency are excluded from the noise calculation.

The relation between the observed response phase and the position of stimulation in the visual field depends on the hemodynamic delay. Since the hemodynamic delay may vary as a function of the position on the cortex (Kastrup *et al.*, 1999), it needs to be measured and corrected for locally. This can be achieved by comparing the responses to two stimuli moving in opposite directions and being otherwise identical. Once response amplitude and phase are known for each of the two scans mapping one visual field coordinate, the two phases are combined on a voxel-per-voxel basis. For any given voxel, the local difference between estimated (t_H) and actual hemodynamic delay offsets the phase observed in the two scans by equal amounts. Thus, the receptive field positions estimated from the delayed responses are biased by the same amount, but in opposite directions with respect to the true value, due to the opposite directions of motion. The two phases obtained at a voxel may exhibit a phase jump with respect to each other. Phase unwrapping is straightforward, if the bias introduced by the mismatch between actual and estimated hemodynamic response is small with respect to the temporal period of the stimulus. Therefore, the estimation of the hemodynamic delay t_H should be close to the mean hemodynamic delay observed. We used $t_H = 5$ s. The arithmetic mean of the unwrapped phases yields an unbiased estimation of the true corresponding spatial position. The uncertainty of the combined phase can be calculated by error propagation from the known individual uncertainties.

The results of the analysis are stored as two pairs of parametric data volumes containing phase and SNR information for each of the two visual field coordinates. These data are subsequently assigned to the model of the cortical surface for further two-dimensional processing.

Assignment of Functional Data to the Surface Model

The task of assigning functional data to the model of the cortical surface is to estimate the original cortical response to the stimuli from the available three dimensional data. This estimate will be represented as functional data for each node of the surface model. We describe the general estimation of surface data at each node as a linear combination of the three dimensional data at all the voxels.[3] The contributions of each voxel to each node need to be chosen based on the expected match between the data observed at the voxels and the underlying cortical response. We consider three mechanisms contributing to a mismatch between the cortical response at a given point of the surface and the data acquired at a (more or less distant) voxel: 1. Noise at the acquisition: even if the signal present in a voxel exactly matches the cortical response, the addition of noise introduces errors to the estimation of the response characteristics. 2. Distance perpendicular to the cortical surface: the signal in a voxel that is not centered on the surface may contain contributions other than the response at the nearest point of the surface. 3. Distance along the cortical surface: the cortical response varies along the surface, and the correlation between responses at two points of the surface depends on their distance and the spatial properties of the response.

In the following we discuss each of these mechanisms and derive a scheme for assigning functional data to the cortical surface, notably integrating an optimal weighting with respect to local SNR.

Noise at acquisition introduces an error in the estimation of the response phase. Partial volume effects, local physiologic noise, and the limited amplitude of the cortical response all contribute to this error.[4] Importantly, the uncertainty of the phase estimation is known for each voxel. For voxels with

[3] Note that this formulation is general, as we allow the contributions of the individual voxels to be a function of the data.

[4] In this context, only those partial volume effects are considered whose only effect is to reduce the SNR of the voxel. This concerns voxels containing local gray matter and neighboring tissues not exhibiting a functional response (WM, CSF). Voxels including partial volumes of gray matter from opposite banks of a sulcus or gyrus remain problematic.

FIG. 9. Reproducibility of the retinotopic mapping experiments. For each of the three subjects, two complete functional and structural data sets were acquired and processed independently by two operators (J.W. and M.D.). The automatically delineated borders between the retinotopic visual areas from the data sets processed by M.D. are shown superposed on the VFR maps obtained in the exams processed by J.W. For display purposes, only the parts of the surfaces containing the delineated visual areas are shown.

moderate to high SNR, phase noise is close to Gaussian, and the standard deviation of the phase error is equal to the inverse of the SNR of the response amplitude (see Appendix). Only for very low SNR does the response amplitude tend to be overestimated since Gaussian noise is no longer an appropriate model (Gudbjartsson and Patz, 1995). Voxels with an SNR < 2 were therefore excluded from the analysis. For the remaining voxels, the phase uncertainty can be taken into account by weighting their respective contribution in the subsequent analysis. In the Appendix, we show the optimal relative weighting of phase measurements to be the square of the respective SNRs. Note that this result relies on statistical independence of the data.

The second source of error is the distance of voxels perpendicular to the surface. Due to the columnar organization of retinotopy in the cortex, a small offset of the voxel from the surface does not introduce additional errors in the phase estimation. The influence of noise due to partial volumes of WM or CSF in the voxel has already been taken into account above. Given the cortical thickness of about 2 mm and voxel size of 3 mm in some directions, voxels more than 2.5 mm away from their respective closest node of the surface model, cannot be expected to contain useful information, and are discarded. Concerning the relative weighting of the remaining voxels, the need to avoid signal from large draining veins at the pial surface (macrovascular artifact) might be an issue, depending on the acquisition sequence used. The PRESTO sequence used here is intrinsically minimally sensitive to these artifacts, due to fast repetition times and a corresponding saturation of blood signal. We therefore do not take additional measures to avoid macrovascular artifacts. A second—and more problematic—issue is the assignment of data from voxels containing partial volumes of gray matter from either bank of a sulcus or gyrus, notably the calcarine sulcus. These voxels potentially contain two independent signals. While the phase estimated from the sum of the two signals is mostly dominated by the stronger one (which presumably originates from the closest bank), phase errors of up to $\pm 90°$ with respect to this stronger signal may occur. Weighting data as a function of distance from the surface can reduce the impact of these voxels. However, a lot of information from voxels that are not problematic is also suppressed, increasing noise-induced errors in the final maps. We obtained optimal results including all of the data within the distance threshold mentioned above (data not shown). We therefore directly attribute data to the node closest to the center of the respective voxel.[5]

Finally, the distance along the surface needs to be taken into account when combining the projected phase data into a smooth representation of retinotopy. Ideally, the way this is done depends on the spatial properties of the functional response. We do not model those properties explicitly. Instead, we assume that the phase image contains essentially low spatial frequencies and that smoothing with a two-dimensional Gaussian filter along the surface is appropriate. Gaussian smoothing, a simple operation on a regular grid in Cartesian space, is less straightforward to implement for data represented on the irregularly sampled, folded model of the cortical surface.[6] An approach based on the analogy between Gaussian smoothing and heat diffusion has been proposed (Andrade *et al.*, 2001). However, its iterative nature makes it difficult to incorporate the optimal weighting presented above, since data are no longer statistically independent after the first iteration. We therefore use an approach that is simpler, both conceptually and computationally, by calculating all mutual geodesic distances involved and constructing the filter explicitly. Note that this is computationally tractable, because only a fraction of the surface nodes have been assigned functional information in the preceding step. More specifically, the weights w_{ij} of the Gaussian filter are calculated individually for each node i and each projected voxel j as a function of the geodesic distance d_{ij} from node i to the node closest to voxel j:

$$w_{ij} = \begin{cases} \exp\left(-\dfrac{1 d_{ij}^2}{2\sigma^2}\right), & d_{ij} \le r_{\max} \\ 0, & d_{ij} > r_{\max} \end{cases}, \qquad (2)$$

where σ is the standard deviation of the filter and $r_{\max}$ is a cutoff radius introduced to limit distance calculations. There is a trade-off in the choice of the cutoff radius between the time needed to construct this filter and the artifacts introduced due to truncation. We obtained good results with a cutoff radius of 2.5 times the standard deviation of the filter applied. Thus, on the rim the filter has dropped to about 4% of its center height. Distances along the surface are calculated using the Dijkstra algorithm (Dijkstra, 1959). This algorithm overestimates geodesic distances, which means that the parameters of the smoothing filters reported here overestimate the width of the filter actually applied. The choice of the standard deviation of the filter depends on the purpose of the data assignment. For visualization, both eccentricity and polar angle data are only slightly smoothed using a filter of $\sigma = 1.5$ mm.

In summary, the phase $\bar{\phi}_i$ assigned to node i is given by

$$\bar{\phi}_i = \frac{\sum_{j \in V'} w_{ij}\, \mathrm{SNR}_j^2\, \phi_j}{\sum_{j \in V'} w_{ij}\, \mathrm{SNR}_j^2}, \qquad (3)$$

where SNR_j and ϕ_j are, respectively, the SNR and phase observed at voxel j and V' is the set of voxels whose center is at most 2.5 mm from the closest node of the surface and whose response exceeds an amplitude threshold of SNR > 2.

So far we have disregarded the fact that the phases determined from the voxel responses may present phase wrap-

[5] Accuracy could be increased by projecting to the closest point of the surface (not necessarily a node). However, in our case of a dense surface mesh with typical internode distances of 0.9 mm the expected gain is small.

[6] Note that conventional isotropic smoothing in three-dimensional Cartesian space is not an option as averaging across sulci must be avoided.

ping. The stimuli are designed such that minimal and maximal eccentricities, or the lower and upper vertical meridians, are stimulated simultaneously, creating an ambiguity between the lowest and highest response phases. Note that the ambiguity between right and left visual hemifields is resolved by the fact that only the contralateral hemifield is represented in the low order visual areas of each hemisphere. This reduces the (complex) task of unwrapping phase over large surface areas to the simpler one of unwrapping noise induced jumps at relatively isolated points.[7] Since retinotopy varies smoothly along the cortical surface, phase discontinuities can be attributed to phase jumps. Specifically, we compare the phase of each voxel k with the phase assigned from all other voxels $j \neq k$ to its closest surface node i:

$$\Delta\phi_k = \phi_k - \frac{\sum_{j\in V, j\neq k} w_{ij}\,\mathrm{SNR}_j^2\,\phi_j}{\sum_{j\in V, j\neq k} w_{ij}\,\mathrm{SNR}_j^2}. \quad (4)$$

This assignment is performed using a width of $\sigma = 4.5$ mm for the Gaussian filter. A mismatch $\Delta\phi_k$ of over 170° is considered indicative of phase wrapping. This method proved robust and efficiently corrects most isolated phase wraps. Only after unwrapping of the voxel phases are their data assigned to the surface using Eq. 3.

Processing of the Retinotopic Maps—Delineation of the Retinotopic Visual Areas

The delineation of the retinotopic visual areas is based on the fact that the orientation of the visual field representation on the cortical surface changes between adjacent areas. The term orientation here refers to a mirror image versus non mirror image representation. The polar visual field coordinates (r, α), expressed as the observed response phases (ϕ_r, ϕ_α), establish at each point of the cortical surface a two-dimensional coordinate system. The orientation of this coordinate system with respect to a local parametrization (u, v) of the surface is most conveniently determined from the Jacobian of the mapping $\phi_r = \phi_r(u, v)$, $\phi_\alpha = \phi_\alpha(u, v)$. The Jacobian can be interpreted as the ratio of an oriented area measured in coordinates (ϕ_r, ϕ_α) with respect to the same area measured in coordinates (u, v). We refer to this ratio as the VFR:

$$\mathrm{VFR} = \frac{\partial(\phi_r, \phi_\alpha)}{\partial(u, v)} \quad (5)$$

$$= \nabla_u\phi_r \cdot \nabla_v\phi_\alpha - \nabla_v\phi_r \cdot \nabla_u\phi_\alpha, \quad (6)$$

where $\nabla_u\phi_r = \partial\phi_r / \partial_u \ldots$ are the phase gradients with respect to the local parametrization.

The central step in the calculation of the VFR is the detection of the phase gradients. Gradient detection tends to be very sensitive to high-frequency noise in the data. We therefore assigned data to the surface using large Gaussian filters with standard deviations of 3.5 and 7 mm to the polar angle and eccentricity phases, respectively. To effectively reduce noise, the largest filter compatible with the spatial frequency spectrum of the signal should be applied. Retinotopy with respect to eccentricity is much smoother than its polar angle counterpart, meaning that the signal in the eccentricity phase maps contains essentially very low spatial frequency components. Therefore, stronger smoothing can be applied without significantly deteriorating the retinotopic eccentricity map.

The calculation of the phase gradients requires a local two-dimensional coordinate system (a parametrization of the surface) with respect to which the gradients can be calculated. The parametrization used for the gradient calculations should correctly reflect local distances and angles in the surface. A global parametrization of the surface can be provided by the two-dimensional Cartesian coordinates of the nodes in a flattened representation of the cortical surface. However, depending on the size of the flattened surface and the accuracy desired, the distortions induced in the flattening process may make this parametrization unsuitable for gradient calculation. Rather, we chose to locally map each node together with its first-order neighbors to a plane. The central node is placed at the origin. Neighbors are placed preserving their distance to the central node as well as the proportions of the angles at the central node (Welch and Witkin, 1994; Andrade *et al.*, 2001). If the Jacobian of the mapping is constrained to be positive definite, this mapping is unique except for rotations in the plane. Due to this rotational degree of freedom, gradients calculated with respect to this mapping are not comparable across nodes. In the case of calculation of the visual field sign, this is not a limitation, since the VFR is calculated based on local gradients only and is invariant under rotation. Therefore, we did not constrain the angular position of nodes in the plane. Given this local mapping, the response phase can be fitted locally with a first-order function of the node coordinates. The coefficients of the first-order terms in the fit are the phase gradients.

Given these gradients, the calculation of the VFR according to Eq. 6 is straightforward.

From the VFR map, candidates for the visual areas are detected as contiguous regions exceeding a certain threshold on the absolute VFR ($|\mathrm{VFR}| > 8$ (deg/mm)^2) and on the SNR of the smoothed eccentricity and polar angle phase maps (SNR > 15). These candidates are ordered by size, based on the cumulative response power observed during all functional scans within each of the delineated regions. The low order visual areas are then detected among the candidates, starting with V1. V1 is selected as the largest candidate with negative VFR. The two largest candidates with positive VFR adjacent to V1 are then labeled V2d and V2v. Proceeding in this manner, areas up to V3A and V4 are labeled automatically. In a last step, the limits between those areas are drawn automatically as contour lines of zero VFR that border two adjacent visual areas. The delineation of the visual areas

[7] This would not be the case if a higher number of wedges or rings had been used, complicating the data analysis in those cases.

from the VFR map is entirely performed in the two-dimensional Cartesian space of the flattened surface representation. Distortions induced by the unfolding process have little or no bearing on these processing steps. The procedure is automatic and for most data sets it produces correct results with the standard set of parameters given here.

RESULTS

In order to assess the reproducibility of the methods described, we independently acquired two complete data sets, each comprising the high-resolution structural and the functional retinotopic mapping data, for each of three healthy adult human volunteers (M.D., J.W., C.D.M.). Subjects gave prior informed consent to participate in the study. Data analysis was performed according to the procedures detailed above. In the development of those methods, automation has been an important goal. Some manual interaction still remains, essentially concerning the correction of the segmentation and the adjustment of some processing parameters. To include the effect of operator dependence in the assessment of the reproducibility, data were analyzed independently by two researchers (J.W. and M.D.). To distinguish between the two data sets for each subject, we call the three exams processed by J.W. "reference" exams and the remaining three processed by M.D. "control" exams.

Segmentation of Structural Data

The segmentation and postprocessing of the high-resolution structural images took about 6 min on a SPARC Ultra10 workstation and produced accurate results for most of the volume. However, some errors requiring manual correction remained. Editing systematically concerned the WM medially to the posterior horns of the lateral ventricles. Remaining errors were mostly localized near the calcarine sulcus, where the contrast-to-noise ratio between GM and WM is low. Topological defects were rapidly detected by tentatively constructing and unfolding a model of the GM/WM interface. Anatomical errors were localized by comparing original and segmented images side by side. The entire process of localizing and correcting all errors required about 3 h per hemisphere. Typically, from 1 to 3% of the voxels labeled GM or WM in the region to be unfolded were edited.

The portions of the 12 hemispheres that were selected for unfolding had sizes ranging from 89 to 126 cm^2, with a median size of 109 cm^2 and an average node density of 157 nodes/cm^2. Construction and flattening of the surface models took an average of 130 s per hemisphere on a SPARC Ultra10 workstation. In order to test if flattening was homeomorphic, we calculated the Jacobian of the projection to the plane. The fraction of the surface exhibiting a negative definite Jacobian ranged from 0.4 to 5.4‰, with a median of 1.1‰.

The two cortical surface models reconstructed for each hemisphere were very similar, albeit not identical. The two flat maps of visual area boundaries were therefore not directly comparable. Instead, the results for both exams were represented on the flat maps obtained from the reference exams.[8]

Stimuli, Acquisition, and 3D Analysis of Functional Data

The results of the three-dimensional analysis for subject M.D. (reference data set) are depicted in Figs. 1m and 1n for polar angle and eccentricity data, respectively. The stimuli evoked robust activation throughout the low order visual areas. The superposition of functional and structural data was checked visually and found to be good. The sensitivity of the acquisition sequence proved to be sufficient, the SNR of the spectra obtained allowing for a phase estimation with a mean standard deviation of about 12°. The sensitivity was higher dorsally than ventrally, presumably due to the sensitivity profile of the single loop surface coil used.

Assignment of Functional Data to the Surface Model

Figure 6 shows an example of the distribution of response power in one polar angle scan (subject M.D., reference data set) as a function of the distance from the left hemisphere surface model. Only voxels exceeding an SNR threshold of 2 were considered. At the distance threshold of 2.5 mm used here, the voxels that were not projected represented about 25% of the total response power.

The spatial distribution of the surface nodes receiving information can be appreciated from the flat SNR map immediately after projection of the functional data, prior to smoothing. A representative example of this distribution is shown in Fig. 7, depicting the SNR of the combined functional information from both polar angle scans of subject M.D. (reference data set, left hemisphere) after projection to the surface. For all of the surfaces processed, the information projected was sufficiently dense to provide for a good sampling of retinotopy.

We did not find evidence of erroneous assignment of data between the opposite banks of the calcarine sulcus. The final results of the assignment of these data are presented in Fig. 8.

Delineation of the Retinotopic Visual Areas

The default set of parameters allowed for a correct automatic identification of six retinotopic visual areas (V1, V2d/v, V3, VP, V3A, V4v) from the VFR map in 7 of the 12 hemispheres processed. "Correctness" was assessed visually, comparing the delineated areas to the VFR map. In the remaining 5 hemispheres V4v or V3A was not identified correctly with the standard set of parameters, requiring manual adaptation of the VFR or SNR thresholds. In the 2 left hemispheres of subject J.W., V3A could not be delineated by our algorithm. The border between V3 and V3A was therefore drawn manually, based on the VFR map. It was projected in the same fashion as the other visual area borders.

[8] This choice of which one of the surfaces served as reference had little effect on the results, since the two surfaces were close to each other (the distance was smaller than 1 mm for over 90% of the total surface area).

Only the borders between the delineated visual areas are shown. SNR was not sufficient to delineate the anterior limits of V3A and V4v reproducibly. The decreased sensitivity of the single loop surface coil and the lower retinotopic specialization of neurons in these regions presumably account for the low SNR observed.

Reproducibility

In order to assess the reproducibility of the entire processing chain, the visual area borders obtained in the control exams were projected to the reference surfaces. First, the data from the two exams were brought into a common reference frame by coregistering the two anatomical volumes using SPM (Ashburner and Friston, 1997). Next, the points defining the piecewise linear area borders on the control surfaces were projected in three dimensions from their original positions to the closest point on the reference surface. Finally, local linear transformations between the three-dimensional and flattened representations of the reference surfaces were calculated for each triangle and applied to the projected points. The projected visual area borders were displayed on the VFR map obtained from the reference exam (Fig. 9). The limits delineated for the reference surfaces are not shown for clarity of the display, but their position can be easily inferred from the VFR map. The projected area borders closely follow the VFR pattern of the reference surface.

At some places, "jumps" appear in the projected visual area limits, although displayed on their original surface they are continuous (see, for example, the limits between V1 and V2v and between V2v and VP for subject J.W., right hemisphere). These jumps are artifacts generated by the projection from the control surface to the reference surface. A more sophisticated procedure that continuously maps the two surfaces to each other would allow for an uninterrupted (and more accurate) visualization of the projected visual area border.

DISCUSSION AND CONCLUSION

The methods described allow for a delineation of visual areas with an excellent reproducibility. Note that systematic errors that might be inherent in the method are not addressed here. To assess them, an independent "gold standard" method, to precisely identify the borders between visual areas, would be necessary. However, the delineation of visual areas is usually not a goal in itself, but serves to help interpret other cognitive or sensory functional mapping results. As long as the positioning of those results relative to the retinotopic mapping data is accurate, their interpretation will be accurate too. We therefore hold that the variability of the maps observed in this study is apt to represent the size of potential errors in the interpretation of functional data, provided the latter is acquired using a MRI sequence with little distortion with respect to the retinotopic mapping.

The manual correction of segmentation errors is by far the most operator-dependent step in the analysis of retinotopic mapping data presented here. As such, it potentially adds variability to the visual area maps beyond the mere variability of the raw data. However, given the good agreement between the models obtained independently by different operators, the influence on the results should be small. We are currently working to improve our segmentation procedure (Richard *et al.*, 2002) and hope to further increase the reproducibility of the surface models obtained in the near future.

The PRESTO acquisition sequence reduces the distortions present in the functional data with respect to an EPI sequence, provides for a good alignment of the data with respect to the surface models and is expected not to be sensitive to macrovascular artifacts. These advantages come at the cost of a slight loss of SNR in the acquired images, with respect to a single shot EPI acquisition. The acquisition time of 1.28 s for each volume is small with respect to the period of stimulation and the resulting variation of the image contrast during this time does not introduce significant artifacts or phase uncertainties.

The stimuli presented provided robust retinotopic activation and led to high-quality retinotopic maps. There might still be room for improvement, however. As reasoned under Methods, the uncertainty of the retinotopic maps decreases linearly with increasing number of rings or wedges in the stimulus. This is true until the spatial low-pass characteristics of the finite size receptive fields degrade the observed response. Polar angle stimuli with more than two wedges might therefore be desirable. In that case, *a priori* information about retinotopy will likely not suffice to resolve the ambiguity between the wedges. Rather, stimuli without a simple periodicity, and consequently analysis at multiple frequencies, might prove to be a solution.

The assignment of functional data to the surface model recovered most of the information from the three-dimensional analysis while remaining robust with respect to small voluntarily induced misalignments of structural and functional data (data not shown). Using a more realistic model of the cortex, taking its thickness into account, might allow to identify voxels that are prone to partial volume effects involving signal from opposite banks of sulci or gyri, notably the calcarine sulcus. These voxels could then be weighted appropriately, which might further improve the assignment. Note, however, that discarding data does not address the issue fundamentally. Rather, the functional response of those voxels would need to be described as a mix of two reference functions in a framework modeling the signal on both sides of the sulcus or gyrus.

In summary, we have given a step-by-step description and discussion of the methods involved in fMRI retinotopic mapping, leading to a map of the borders between retinotopic visual areas of the examined subject. Several issues that are specific to the processing of functional retinotopic mapping data have been discussed in detail. The advantage of using multiple elements in stimuli for retinotopic mapping has been presented. The polar angle stimuli described reduce the noise-induced uncertainty of the retinotopic maps by a factor of 2 with respect to the stimuli generally used. We proposed a method of data assignment that optimally takes into account the influence of limited SNR on the response phase measured. The methods described lead to a highly reproduc-

ible delineation of the retinotopic visual areas as assessed by a reproducibility study performed on three subjects.

APPENDIX

Uncertainty of the Phase Measurement

At each voxel j, a real discrete-time signal $f(\vec{x}_j, t_k)$ is observed at N instances t_k. We model the observed signal as the sum of zero-mean white Gaussian noise and a harmonic response at the known stimulation frequency ν_0. The phase and amplitude of the response are calculated from the real time courses' complex Fourier component F_{ν_0}:

$$F_{\nu_0}(\vec{x}_j) = \sum_{k=1}^{N} f(\vec{x}_j, t_k) \exp(i2\pi\nu_0 (t_k - t_H)) , \tag{7}$$

where t_H is the expected hemodynamic delay. Due to the linearity of the Fourier transform, $F_{\nu_0}(\vec{x}_j)$ is the sum of a complex signal with amplitude $\tilde{A}_j$, phase $\tilde{\phi}_j$, and a zero-mean Gaussian noise, whose real and imaginary parts are independent random variables entirely characterized by their standard deviations of $\sigma_{\text{imag},j} = \sigma_{\text{real},j}$, noted hereafter σ_j. Noise in the raw data is assumed to be white at frequencies above ν_0. The noise can then be measured at those frequencies assumed to contain no signal. We define the local signal-to-noise ratio SNR_j as

$$\text{SNR}_j = \frac{\tilde{A}_j}{\sigma_j} . \tag{8}$$

The actually measured phase ϕ and amplitude A of the response at the stimulation frequency ν_0 are

$$\phi_j = \arg(F_{\nu_0}(\vec{x}_j)) \tag{9}$$

$$A_j = |F_{\nu_0}(\vec{x}_j)| . \tag{10}$$

The noise of the measured phase and amplitude is Gaussian only in the limit of high SNRs, because Eqs. (9) and (10) are nonlinear. However, at moderate signal to noise ratios of $\text{SNR} > 2$ the mean values and variances of measured amplitude and phase can be expressed in reasonably good approximation as (Gudbjartsson and Patz, 1995)[9]

$$\langle \phi \rangle = \tilde{\phi} \tag{11}$$

$$\sigma_\phi^2 \approx \frac{\sigma^2}{\tilde{A}^2} = \text{SNR}^{-2} \tag{12}$$

$$\langle A \rangle \approx \sqrt{\tilde{A}^2 + \sigma^2} \tag{13}$$

[9] The explicit dependence on $\vec{x}_j$ has been dropped in these and subsequent equations to increase readability, although all variables vary over the imaged volume.

$$\sigma_A^2 \approx \sigma^2 , \tag{14}$$

where $\langle\ \rangle$ denotes the first moment of the distribution of the measured values. The measured amplitude is biased by the noise, and a corrected amplitude $A' = \sqrt{|A^2 - \sigma^2|}$ could be calculated, which is essentially unbiased in terms of its mean value (Gudbjartsson and Patz, 1995). However for a $\text{SNR} > 2$, its mean square error $e_{A'}^2 = \langle (A' - \tilde{A})^2 \rangle$ is higher than the one of the uncorrected amplitude. We therefore chose to base our SNR measurement on the uncorrected amplitude. The measured phase is unbiased, assuming that phase wrapping due to noise can be corrected perfectly.

The response observed is in the following characterized by the measured phase ϕ and the measured SNR of the response, $SNR = A/\sigma = \sigma_\phi^{-1}$, representing the inverse of the phase measurement error.

Averaging Signals with Nonstationary Noise

For N independent measurements (SNR_j, ϕ_j) of the response phase at voxels j pertaining to the same point of the cortical surface, the average phase can be determined as the weighted sum of the individual phases. The optimal weights depend on the individual uncertainties of the phase measurements. Consider the general weighted average phase, with weights β_j:

$$\phi_{\text{avg}} = \frac{\sum_{j=1}^{N} \beta_j \phi_j}{\sum_{j=1}^{N} \beta_j} . \tag{15}$$

Since for moderate and high SNRs, the phase noise of the individual ϕ_j is approximately Gaussian, and since the individual measures are assumed to be statistically independent, the uncertainty σ_{avg} of this average can be calculated according to Gauss's equation for propagation of random error from the individual uncertainties σ_j (Gauss, 1863):

$$\sigma_{\text{avg}}^2 = \sum_{j=1}^{N} \left(\frac{\partial \phi_{\text{avg}}}{\partial \phi_j} \sigma_j \right)^2 \tag{16}$$

$$= \frac{\sum_{j=1}^{N} (\beta_j \sigma_{\phi_j})^2}{(\sum_{j=1}^{N} \beta_j)^2} . \tag{17}$$

This uncertainty depends on the weights chosen. There is an optimal choice of weights, which is found to be

$$\beta_j = \sigma_j^{-2} = \text{SNR}_j^2 , \tag{18}$$

leading to an optimal combined phase

$$\phi_{\text{avg}} = \frac{\sum_{j=1}^{N} \text{SNR}_j^2 \phi_j}{\sum_{j=1}^{N} \text{SNR}_j^2} \tag{19}$$

with minimal uncertainty of

$$\sigma_{\text{avg}} = \left(\sum_{j=1}^{N} \text{SNR}_j^2\right)^{-1/2}. \quad (20)$$

Note that adding any information with positive SNR improves the quality of the phase estimation. If SNRs are estimated accurately, and phase noise is Gaussian, no data should be excluded from the analysis. Conversely, data not fulfilling these hypotheses might degrade the overall phase estimation.

The average phase ϕ_{avg} can be regarded as a new phase measurement with a signal to noise ratio SNR_{avg} of

$$\text{SNR}_{\text{avg}}^2 = \sum_{j=1}^{N} \text{SNR}_{\text{j}}^2. \quad (21)$$

It is useful to represent the response phases and SNRs in terms of the noise normalized signal power $P = \text{SNR}^2$ and weighted phase $\Phi = P\phi$. Averaging independent phase measurements with optimal weights then reduces to summing up the individual signal powers and weighted phases:

$$\Phi_{\text{avg}} = \sum_{j=1}^{N} \Phi_j \quad (22)$$

$$P_{\text{avg}} = \sum_{j=1}^{N} P_j. \quad (23)$$

Two caveats are in place here: First, individual measurements with very low SNR (SNR < 2) should be excluded from the data processing. Due to noise, the amplitude, and therefore SNR, is systematically overestimated, especially at low SNR. Noisy measurements therefore contribute more strongly to the weighted phase than they should, bearing the risk of degrading the SNR of the averaged value, rather than adding information. Second, in this calculation we assume the (P_j, Φ_j) to be measurements of the same observable. In the context of retinotopic mapping this is true if they pertain to points on the cortical surface whose mutual distance is small with respect to the local smoothness of retinotopy.

Smoothing Phase Data

The reasoning exposed above for averaging phase information can be extended to the convolution of phase information along the surface with a smoothing kernel. We determine a set of voxels V' to be included in the analysis depending on their SNR and distance from the surface. A general smoothing kernel can then be defined by a set of mutual weights w_{ij} as a function of the geodesic distance from each node i to the node closest to the voxels $j \in V'$. Smoothing the response phase ϕ, taking into account the signal power at each point, can then be written as

$$\bar{\phi}_i = \frac{\sum_{j\in V'} w_{ij} P_j \phi_j}{\sum_{j\in V'} w_{ij} P_j}. \quad (24)$$

$\bar{\phi}_i$ represents the smoothed phase of voxel i. Contrary to the previous deduction, we do not claim here that this filter is optimal (we believe it is close to optimal if the w_{ij} match the spatial properties of the phase signal). But its choice can be motivated by a number of observations: in the case of the constant filter ($\mathbf{w} \equiv 1$) we find the previous result of the optimal average of the phases; in the case of constant response power ($\mathbf{P} \equiv p$) we apply the classical smoothing filter defined solely by the weights w_{ij}; it can be shown that the noise propagated by the filter (24) is never stronger than that propagated by the classical filter alone; points containing no data ($P_j = 0$) do not influence the smoothed phase. The result of the filtering procedure is undefined if no data are present at all ($\mathbf{P} \equiv 0$).

Again, the uncertainty of the smoothed phase $\bar{\sigma}$ can be calculated using Gauss's equation for propagation of random error. Care must be taken, however, in the interpretation of the errors assigned to the individual terms contributing to the overall error. Smoothing combines data from mutually distant points of the surface. The individual phase measures involved do not pertain to the same observable. The error attributed to the phase measurement ϕ_j must then be interpreted as the error σ_{ij} made when representing the phase at point i by the measurement at voxel j:

$$\bar{\sigma}_i^2 = \sum_{j\in V'} \left(\frac{\partial \bar{\phi}_i}{\partial \phi_j} \sigma_{ij}\right)^2 \quad (25)$$

$$= \frac{\sum_{j\in V'} (w_{ij} P_j \sigma_{ij})^2}{\left(\sum_{j\in V'} w_{ij} P_j\right)^2}. \quad (26)$$

Rigorous analytic calculation of this error would require an explicit model of the data in a Bayesian framework, since the σ_{ij} it depends on the (classically unknown) spatial properties of the actual signal present. Instead we assume that the w_{ij} reflect the spatial properties of the signal such that the error we make representing the phase at point i by the measurement at voxel j is $\sigma_{ij} = (w_{ij} P_j)^{-1/2}$.[10] In that case we can write this filter in analogy to (22) and (23) as

$$\bar{\Phi}_i = \sum_{j\in V'} w_{ij} \Phi_j \quad (27)$$

$$\bar{P}_i = \sum_{j\in V'} w_{ij} P_j \quad (28)$$

$$\bar{\phi}_i = \frac{\bar{\Phi}_i}{\bar{P}_i}. \quad (29)$$

[10] Note that this implies a normalization of the smoothing weights such that $w_{ii} = 1$.

The kernel used in our implementation is normalized to unity at its center. The value of the filtered phase (29) is independent of this scaling, but the filtered power of the response (28) is not. The normalization of the smoothing kernel dictates the interpretation of the obtained power maps. The power map smoothed with the present kernel is the inverse of the actual local phase variance after smoothing (provided the kernel is adapted to the spatial properties of the phase signal). Smoothing with a kernel whose integral is normalized to unity would produce a map interpolating the phase errors present locally in the raw data.

ACKNOWLEDGMENTS

J.W. was supported by Ph.D. grants from the "Deutscher Akademischer Austausch Dienst (DAAD)" and the French Ministry of Research and Technology. We acknowledge financial support from the Région Rhône-Alpes (Programme de Recherches Thématiques Prioritaires Neurosciences et Cognisciences).

REFERENCES

Amunts, K., Malikovic, A., Mohlberg, H., Schormann, T., and Zilles, K. 2000. Brodmann's areas 17 and 18 brought into stereotaxic space—Where and how variable? *NeuroImage* **11**(1), 66–84.

Andrade, A., Kherif, F., Mangin, J. F., Worsley, K. J., Paradis, A. L., Simon, O., Dehaene, S., Le Bihan, D., and Poline, J. B. 2001. Detection of fMRI activation using cortical surface mapping. *Hum. Brain Mapp.* **1**2(2), 79–93.

Angenent, S., Haker, S., Tannenbaum, A., and Kikinis, R. 1999. On the Laplace-Beltrami operator and brain surface flattening. *IEEE Trans. Med. Imaging* **1**8(8), 700–711.

Anstis, S. M., and Cavanagh, P. 1983. A minimum motion technique for judging equiluminance. In: Mollon, J. D., Sharpe, L. T. (Eds.), Colour Vision: Physiology and Psychophysics. Academic Press, London, pp. 155–166.

Ashburner, J., and Friston, K. 1997. Multimodal image coregistration and partitioning—A unified framework. *NeuroImage* **6**(3), 209–217.

Bandettini, P. A., Jesmanowicz, A., Wong, E. C., and Hyde, J. S. 1993. Processing strategies for time-course data sets in functional MRI of the human brain. *Magn. Reson. Med.* **30**(2), 161–173.

Baseler, H. A., Morland, A. B., and Wandell, B. A. 1999. Topographic organization of human visual areas in the absence of input from primary cortex. *J. Neurosci.* **19**(7), 2619–2627.

Carman, G. J., Drury, H. A., and Van Essen, D. C. 1995. Computational methods for reconstructing and unfolding the cerebral cortex. *Cereb. Cortex* 5(6), 506–517.

Dale, A. M., Fischl, B., and Sereno, M. I. 1999. Cortical surface-based analysis. I: Segmentation and surface reconstruction. *NeuroImage* **9**(2), 179–194.

Demartines, P., and Herault, J. 1997. Curvilinear Component Analysis: A self-organising neural network for non-linear mapping of data sets. *IEEE Trans. Neural Networks* 8(1), 148–154.

DeYoe, E. A., Bandettini, P., Neitz, J., Miller, D., and Winans, P. 1994. Functional magnetic resonance imaging (FMRI) of the human brain. *J. Neurosci. Methods* 54(2), 171–187.

Di Russo, F., Martinez, A., Sereno, M. I., Pitzalis, S., and Hillyard, S. A. 2002. Cortical sources of the early components of the visual evoked potential. *Hum. Brain Mapp.* **1**5(2), 95–111.

Dijkstra, E. W. 1959. A note on two problems in connexion with graphs. *Num. Math.* **1**, 269–271.

Drury, H. A., Van Essen, D. C., Anderson, C. H., Lee, C. W., Coogan, T. A., and Lewis, J. W. 1996. Computerized mappings of the cerebral cortex: A multiresolution flattening method and a surface-based coordinate system. *J. Cogn. Neurosci.* 8(1), 1–28.

Engel, S. A., Glover, G. H., and Wandell, B. A. 1997. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb. Cortex* 7(2), 181–192.

Engel, S. A., Rumelhart, D. E., Wandell, B. A., Lee, A. T., Glover, G. H., Chichilnisky, E. J., and Shadlen, M. N. 1994. fMRI of human visual cortex. *Nature* **369**(6481), 525.

Fischl, B., Liu, A., and Dale, A. M. 2001. Automated manifold surgery: Constructing geometrically accurate and topologically correct models of the human cerebral cortex. *IEEE Trans. Med. Imaging* 20(1), 70–80.

Fischl, B., Sereno, M. I., and Dale, A. M. 1999. Cortical surface-based analysis. II: Inflation, flattening, and a surface-based coordinate system. *NeuroImage* **9**(2), 195–207.

Friston, K. J., Ashburner, J., Frith, C. D., Poline, J.-B., Heather, J. D., and Frackowiack, R. S. J. 1995. Spatial registration and normalization of images. *Hum. Brain Mapp.* **3**(3), 165–189.

Friston, K. J., Jezzard, P., and Turner, R. 1994. Analysis of functional MRI time-series. *Hum. Brain Mapp.* **1**, 153–171.

Gauss, C. F. 1863. Theoria combinationis observationum erroribus minimis obnoxiae. Pars prior. In *Werke*, Vol. 4, pp. 1–26. Dietrich, Göttingen.

Germond, L., Dojat, M., Taylor, C., and Garbay, C. 2000. A cooperative framework for segmentation of MRI brain scans. *Artif. Intell. Med.* 20(1), 77–93.

Gudbjartsson, H., and Patz, S. 1995. The Rician distribution of noisy MRI data. *Magn. Reson. Med.* **34**(6), 910–914.

Guérin-Dugué, A., Olympieff, S., Gisert-Lopez, J., Chéhikian, A., Warnking, J., Rubin, C., and Segebarth, C. 2000. *Représentation Plane du Cortex Visuel en Imagerie Fonctionnelle à Résonance Magnétique.* In 12ème congrès Reconnaissance Formes et Intelligence Artificielle (RFIA'2000), Vol. 2, pp. 29–38.

Hadjikhani, N., Liu, A. K., Dale, A. M., Cavanagh, P., and Tootell, R. B. 1998. Retinotopy and color sensitivity in human visual cortical area V8. *Nature Neurosci.* **1**(3), 235–241.

Han, X., Xu, C., Braga-Neto, U., and Prince, J. L. 2002. Topology correction in brain cortex segmentation using a multiscale, graph-based algorithm. *IEEE Trans. Med. Imaging* 21(2), 109–121.

Hermosillo, G., Faugeras, O., and Gomes, J. 1999. Unfolding the cerebral cortex using level set methods. In, *Scale-Space Theories in Computer Vision* (M. Nielsen, P. Johansen, O. F. Olsen, J. Weickert, Eds.), Vol. 1682 of *Lecture Notes in Computer Science*, pp. 58–69. Springer-Verlag, Berlin.

Joshi, M., Cui, J., Doolittle, K., Joshi, S., Van Essen, D., Wang, L., and Miller, M. I. 1999. Brain segmentation and the generation of cortical surfaces. *NeuroImage* **9**(5), 461–476.

Kastrup, A., Kruger, G., Glover, G. H., Neumann-Haefelin, T., and Moseley, M. E. 1999. Regional variability of cerebral blood oxygenation response to hypercapnia. *NeuroImage* **10**(6), 675–681.

Kriegeskorte, N., and Goebel, R. 2001. An efficient algorithm for topologically correct segmentation of the cortical sheet in anatomical mr volumes. *NeuroImage* **14**(2), 329–346.

Liu, G., Sobering, G., Duyn, J., and Moonen, C. T. 1993. A functional MRI technique combining principles of echo-shifting with a train of observations (PRESTO). *Magn. Reson. Med.* **30**(6), 764–768.

Lorensen, W. E., and Cline, H. E. 1987. Marching cubes: A high resolution 3D surface construction algorithm. *Comput. Graphics* (Proc. of SIGGRAPH '87) 21(4), 163–169.

MacDonald, D., Kabani, N., Avis, D., and Evans, A. C. 2000. Automated 3-D extraction of inner and outer surfaces of cerebral cortex from MRI. *NeuroImage* **1**2(3), 340–356.

MacQueen, J. 1967. Some methods for classification and analysis of multivariate observations. In *Proc. of the Fifth Berkely Symposium on Mathematical Statistics and Probability* (L. M. LeCam, and J. Neyman, Eds.), pp. 281–297. *Vol. 1. Statistics.* Univ. of California Press, Berkely, CA.

Morland, A. B., Baseler, H. A., Hoffmann, M. B., Sharpe, L. T., and Wandell, B. A. 2001. Abnormal retinotopic representations in human visual cortex revealed by fMRI. *Acta Psychol.* **107**(1–3), 229–247.

Press, W. A., Brewer, A. A., Dougherty, R. F., Wade, A. R., and Wandell, B. A. 2001. Visual areas and spatial summation in human visual cortex. *Vision Res.* **41**(10, 11), 1321–1332.

Richard, N., Dojat, M., and Garbay, C. 2002. Situated cooperative agents: A powerful paradigm for MRI brain scans segmentation. In *Proc. ECAI* (F. Van Harmelen, Ed.), pp. 33–37. IOS Press, Amsterdam.

Saad, Z. S., Ropella, K. M., Cox, R. W., and De Yoe, E. A. 2001. Analysis and use of fMRI response delays. *Hum. Brain Mapp.* **13**(2), 74–93.

Sceniak, M. P., Ringach, D. L., Hawken, M. J., and Shapley, R. 1999. Contrast's effect on spatial summation by macaque V1 neurons. *Nature Neurosci.* 2(8), 733–739.

Schwartz, E., Shaw, A., and Wolfson, E. 1989. A numerical solution to the generalized mapmaker's problem: Flattening nonconvex polyhedral surfaces. *IEEE Trans. Pattern Anal. Machine Intell.* **11**(9), 1005–1008.

Sereno, M. I., Dale, A. M., Reppas, J. B., Kwong, K. K., Belliveau, J. W., Brady, T. J., Rosen, B. R., and Tootell, R. B. 1995. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* **268**(5212), 889–893.

Sereno, M. I., McDonald, C. T., and Allman, J. M. 1994. Analysis of retinotopic maps in extrastriate cortex. *Cereb. Cortex* **4**(6), 601–620.

Shattuck, D. W., and Leahy, R. M. 2001. Automated graph-based analysis and correction of cortical volume topology. *IEEE Trans. Med. Imaging* **20**(11), 1167–1177.

Shattuck, D. W., Sandor-Leahy, S. R., Schaper, K. A., Rottenberg, D. A., and Leahy, R. M. 2001. Magnetic resonance image tissue classification using a partial volume model. *NeuroImage* **13**(5), 856–876.

Smith, A. T., Singh, K. D., Williams, A. L., and Greenlee, M. W. 2001. Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex. *Cereb. Cortex* **11** (12), 1182–1190.

Teo, P. C., Sapiro, G., and Wandell, B. A. 1997. Creating connected representations of cortical gray matter for functional MRI visualization. *IEEE Trans. Med. Imaging* **16**(6), 852–863.

Tootell, R. B., and Hadjikhani, N. 2001. Where is 'dorsal V4' in human visual cortex? Retinotopic, topographic and functional evidence. *Cereb. Cortex* **11**(4), 298–311.

Tootell, R. B., Hadjikhani, N. K., Mendola, J. D., Marrett, S., and Dale, A. M. 1998a. From retinotopy to recognition: fMRI in human visual cortex. *Trends Cogn. Sci.* 2(5), 174–183.

Tootell, R. B., Hadjikhani, N. K., Vanduffel, W., Liu, A. K., Mendola, J. D., Sereno, M. I., and Dale, A. M. 1998b. Functional analysis of primary visual cortex (V1) in humans. *Proc. Natl. Acad. Sci. USA* **95**(3), 811–817.

Tootell, R. B., Mendola, J. D., Hadjikhani, N. K., Ledden, P. J., Liu, A. K., Reppas, J. B., Sereno, M. I., and Dale, A. M. 1997. Functional analysis of V3A and related areas in human visual cortex. *J. Neurosci.* **17**(18), 7060–7078.

Van Essen, D. C., Lewis, J. W., Drury, H. A., Hadjikhani, N., Tootell, R. B., Bakircioglu, M., and Miller, M. I. 2001. Mapping visual cortex in monkeys and humans using surface-based atlases. *Vision Res.* **41**(10–11), 1359–1378.

Van Essen, D. C., and Maunsell, J. H. 1980. Two-dimensional maps of the cerebral cortex. *J. Comp. Neurol.* **191**(2), 255–281.

van Gelderen, P., Ramsey, N. F., Liu, G., Duyn, J. H., Frank, J. A., Weinberger, D. R., and Moonen, C. T. 1995. Three-dimensional functional magnetic resonance imaging of human brain on a clinical 1.5-t scanner. *Proc. Natl. Acad. Sci. USA* **92**(15), 6906–6910.

Van Leemput, K., Maes, F., Vandermeulen, D., and Suetens, P. 1999a. Automated model-based bias field correction in MR images of the brain. *IEEE Trans. Med. Imaging* **18**(10), 885–896.

Van Leemput, K., Maes, F., Vandermeulen, D., and Suetens, P. 1999b. Automated model-based tissue classification of MR images of the brain. *IEEE Trans. Med. Imaging* **18**(10), 897–908.

Wade, A. R., Brewer, A. A., Rieger, J. W., and Wandell, B. A. 2002. Functional measurements of human ventral occipital cortex: Retinotopy and colour. *Philos. Trans. R. Soc. London B* **357**, 963–973.

Wald, L., Schmitt, F., and Dale, A. M. 2001. Systematic spatial distortion in MRI due to gradient non-linearities. *NeuroImage* (Proc of HBM 2001) **13**(6), S50.

Wandell, B. A. 1999. Computational neuroimaging of human visual cortex. *Annu. Rev. Neurosci.* 22, 145–73.

Wandell, B. A., Chial, S., and Backus, B. T. 2000. Visualization and measurement of the cortical surface. *J. Cogn. Neurosci.* **12**(5), 739–752.

Welch, W., and Witkin, A. 1994. Free-form shape design using triangulated surfaces. *Comput. Graphics* (Proc. of SIGGRAPH '94) **28**, 247–256.

Wells, W. M. III, Grimson, W. E. L., Kikinis, R., and Jolesz, F. A. 1996. Adaptive segmentation of MRI data. *IEEE Trans. Med. Imaging* **15**, 429–442.

Zeki, S., and Shipp, S. 1988. The functional logic of cortical connections. *Nature* **335**(6188), 311–317.

Zeng, X., Staib, L. H., Schultz, R. T., and Duncan, J. S. 1999. Segmentation and measurement of the cortex from 3D MR images using coupled-surfaces propagation. *IEEE Trans. Med. Imaging* **18**(10), 927–937.

Zhang, Y., Brady, M., and Smith, S. 2001. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans. Med. Imaging* **20**(1), 45–57.