

Anatomical Surface Identification, Range-sensing and Registration for Characterizing Intrasurgical Brain Deformations

Michel Audette

Thesis submission document in partial fulfillment for the degree of Doctor
of Philosophy
Tuesday October 3rd, 2000

Committee: Terry Peters
Kaleem Siddiqi
Rob Kearney
Frank Ferrie
Greg Dudek
Jim Nemes

Abstract

We propose a method for estimating intrasurgical brain shift for image-guided surgery. This method consists of five stages: the identification of relevant anatomical surfaces within the MRI volume, range-sensing of the skin and cortex in the OR, computation of the range-MRI rigid transformation based on a calibration procedure and sensor-base tracking, non-rigid motion tracking over time of cortical range images, and lastly, interpolation of this surface displacement information over the whole brain volume via a realistically valued finite element model of the head. This research effort emphasizes the first four stages.

The first stage is a surface model which imbeds the 3D surface identification problem within the evaluation of a volumetric function, whose final zero-level isosurface coincides with the anatomical surface of interest. Contributions made by us include a novel approach to incorporate higher-level anatomical information in the model, as well as a post-processing stage consisting of a Euclidean Distance and Closest Point Map, which is later exploited to speed up the non-rigid surface registration. Quantitative validation of the cortical surface identification makes use of a digital anthropomorphic head phantom and a MRI simulator whose underlying model is derived from first principles. Further validation is provided with real healthy subject and patient scans.

The range-sensor is a commercial product which uses both laser-based triangulation and defocusing techniques to produce a 2D range profile, and is linearly swept across the skin or cortical surface by a commercial linear positioner to produce a 3D range image. The calibration and sensor-base tracking procedure suggests a means of relating the patient-MRI transformation, already computed in an IGNS setting, to a range-MRI transformation, through the use of the patient-space pointing device.

The surface registration technique is of the Iterative Closest Point type, where each iteration benefits from looking up, rather than searching for, explicit closest point pairs. These closest point pairs in turn are used to determine the non-rigid registration of the MRI-based cortical surface with its deformed range-based homolog. Quantitative validation is ongoing, and makes use of the digital head phantom, where the outer cortical surfaces subpatches taken from it can be explicitly identified and warped based on an analytical displacement function. Our validation also makes use of a novel deformable brain-shaped phantom, made of Polyvinyl Alcohol Cryogel.

1 Introduction

Image-guided neurosurgery (IGNS) is a technique whereby a model of a patient’s head, typically featuring skin, cortex and lesion surfaces, is constructed from a set of segmented MRI images and is then registered with the patient’s head in the OR, using a locating device or *probe*, whose position is tracked over time. The usefulness of this technology hinges on the accuracy of the transformation between the image and patient spaces. However, this transformation becomes less accurate as the brain shifts during surgery. Intrasurgical brain shift has been documented by Maurer [38] and by Skrinjar [50], who attribute it to the effect of gravity, to a gradual seepage of CSF, to cerebral blood volume manipulation, and to the effect of certain drugs. The amount of shift can extend to as much as 15mm, according to Maurer [38].

To alleviate this problem, we propose a method for estimating brain shift which is characterized by the following stages (see figure 1):

- semi-automatic identification (segmentation) of relevant anatomical surfaces within the MRI volume;
- range-sensing of the skin and cortex in the OR;
- computation of the range-MRI transformation, based on a calibration procedure and sensor base tracking;
- non-rigid motion tracking, over time, of the range image of the brain;
- and lastly, interpolation of the resulting surface displacement over the whole brain volume, assuming null displacement at the base of the skull, via a realistic finite element model (FEM).

Our research is an investigation of the first four stages, and aims to provide proof of concept of the fifth stage on the basis of a deformable phantom whose nonrigid motion mirrors that of the brain under surgery. We thus also propose a new phantom design, whose material properties and shape are good approximations of those of the human brain.

In contrast with existing FEM-based techniques [40, 50], we emphasize the estimation of a *dense* cortical surface displacement function, which will then be interpolated over the brain volume (assuming null displacement at the base of the head), whereas these techniques focus on the physical modelling which most realistically performs this interpolation, either assuming *sparse* displacement data (obtained via an intrasurgical pointing device) or altogether not providing any specifics about constraining the model with known displacements. Moreover, these head models are either generic or achieve patient specificity through a laborious manual segmentation stage, whereas our semi-automatic surface identification stage can be explicitly integrated with meshing software to make the finite element model patient-specific with little user interaction. This work can be seen as complimentary to ours, in constituting a patient-specific brain-shift estimation framework. Once the finite-element meshing is implemented in light of identified anatomical surfaces, such as skin, cortex, ventricle and lesion

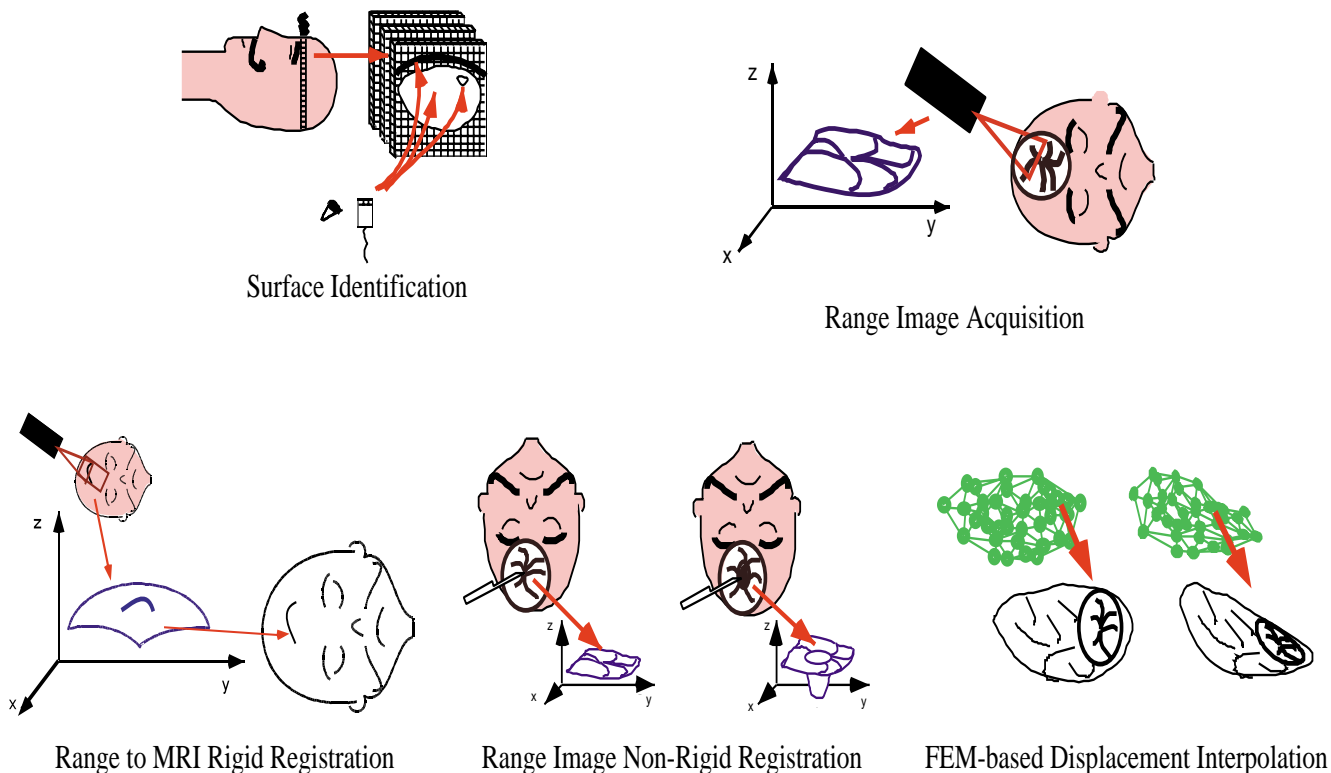


Figure 1: Brain shift estimation framework

boundaries, the use of published material properties of the brain and perhaps of pathological tissue, in conjunction with commercial displacement FEM software, will complete the framework proposed here.

This document is organized as follows. After the introduction, we discuss our work on each stage including results, namely: the semi-automatic identification of the brain surface (§ 2), the range-sensing of the exposed brain surface in the OR (§ 3), the calibration and sensor base tracking procedure for estimating the MRI-range transformation (§ 4), brain motion estimation (§ 5). The latter section also discusses our elastic brain phantom, in the context of validation. § 6 elaborates on the motivation and future work in the finite element modelling stage.

2 Model-Based Semi-Automatic Identification of the Outer Brain Surface in MRI

2.1 Introduction: Motivation for Surface Model

We begin by providing some motivation for using a semi-automatic model-based approach for anatomical surface identification (also known as *segmentation*). To this end, we contrast this

kind of approach with the more traditional and still prevalent manual technique. Furthermore, we make a distinction between a multi-slice (2D) and fully 3D technique, and justify employing the latter approach. We also distinguish between two kinds of surface models, namely *physically-based* and *surface evolution* models, and provide intuitive justification for employing the latter kind. Finally, the details of our model, emphasizing our contribution at the implementation and validation levels, are deferred until the following sections.

Currently, the traditional technique for identifying anatomical surfaces in IGNS involves identifying a point in a slice or volume corresponding to the tissue type of interest, and based on this “seed”, labelling software identifies the rest of the contiguous voxels which are similar intensity (based on user-defined intensity thresholds). Occasionally, a low contrast area at the tissue boundary will cause the labelling to bleed onto tissues of the wrong type, which requires special manual intervention in this area to alleviate the bleeding effect, and typically entails a labour-intensive slice-by-slice, rather than 3D, approach [24, 33]. Current tomographic volumes involve up to 256 of these slices, and each one may require this kind of interaction. Finally, the contiguous labelled voxels within these slices together comprise a volume of interest, coinciding with brain tissue for example, the outside of which (identified by the Marching Cubes algorithm [32]) constitutes the boundary of interest.

In contrast, we adopt a semi-automatic anatomical surface identification (segmentation) stage to replace this procedure. There are four main reasons for this approach.

- First, an immediate consequence of an automatic or semi-automatic technique is the *reproducibility* of the results, in contrast with the inter- and intra-subject variability documented for even expert anatomists[65].
- Moreover, another immediate consequence is the *reduced user intervention* (for a well-designed algorithm and interface).
- Third, a semi-automatic or automatic technique can incorporate *prior knowledge* into the model in order to avoid common pitfalls, such as low-contrast areas which tend to recur across scan subjects. However, because of the variability of brain anatomy across subjects and of the comparative cost of failure of an algorithm when measured against a small amount of user involvement, we consider a semi-automatic segmentation technique more attractive than a fully automatic one.
- Finally, a semi-automatic technique can easily process tomographic information in a fully *three-dimensional* manner, which offers advantages over a series of 2D slice operations. The justification for an inherently 3D algorithm is that either type of model, physical or surface evolution, features a term which favours a smooth contour (in 2D) or surface (in 3D), and a 3D framework produces a boundary which is smooth not only along an intra-slice contour, but across neighbouring slices. This characteristic is critical to an eventual shape-based registration with a range image of the same anatomical surface, whose sampling orientation has no relation with that of the MRI volume.




Figure 2: Imbedding function with 2D contour example, with inward-moving zero-level isocontour, based on synthetic image featuring 2 elliptical bodies: (a) zero-level isocontour (white), overlaid on image at $t = 0$, at $t = 15$ sec., and at $t = 60$ sec.; (b) imbedding function Ψ , with level set $\{\Psi = 0\}$ coinciding with isocontour in (a), at $t = 0$, at $t = 15$ sec., and at $t = 60$ sec.; (c) narrow band implementation of contour evolution model, with identical situation as in (b): the computation of Ψ is restricted within a narrow band near the evolving zero-level isocontour.

Our anatomical surface segmentation technique is an adaptation of a recently published surface evolution model[13, 61], which imbeds the identification of an implicit 3D surface \mathcal{S} into the estimation of an evolving volumetric function $\Psi(x, y, z, t)$, whose isosurface $\{\Psi = 0\}$ within the tomographic volume coincides with \mathcal{S} throughout. This concept of imbedding is illustrated in figure 2 for the analogous problem of lesser dimensionality, namely identifying a contour \mathcal{C} coinciding with an anatomical boundary in a 2D image.




Figure 3: Identification of elliptical boundaries as in figure 2, with a physically-based model without topological adaptation, courtesy of Johan Montagnat: initial configuration, 100th iteration, final iteration.

An initial contour $\mathcal{C}(t_0)$ is first defined by the user outside the boundary of interest, and the imbedding function $\Psi(x, y, t)$ is initialized as the *signed distance map* from $\mathcal{C}(t_0)$. The initial contour coincides with the isocontour $\{\Psi(x, y, t_0) = 0\}$. Interior points in the image domain are typically labelled negative, and the exterior points are labelled positive. The imbedding function $\Psi(x, y, t)$ evolves in such a way that its $\{\Psi = 0\}$ isocontour, or *zero-level set*, moves inwards until it coincides with the anatomical boundary of interest, at which point the zero-level set binds to the boundary, and the resulting contour is precisely this final zero-level set $\{\Psi(x, y, t_f) = 0\}$. To reduce the computational load of imbedding a problem over the whole image domain, we adopt the *Narrow Band* approach proposed by Adalsteinsson[1], whereby the computation of $\Psi(\mathbf{x}, t)$ is restricted to a band, a few pixels wide, in the neighbourhood of its evolving zero-level isocontour. This concept is illustrated in figure 2 (c) for the 2D contour identification pursued in this discussion.

A contour/surface evolution model possesses certain advantages over physically-based contour/surface models [55, 44], such as its capacity to capture a large variety of topologies from a given initialization and a relative insensitivity to initial conditions [35]. In contrast, physically-based models cannot easily cope with differences in topology with respect to their underlying model¹. A comparison of the examples in figures 2 and 3 illustrates this distinction. Moreover, the tensile properties of physical models may make them ill-suited for capturing highly involuted surfaces (such as that of the inner or outer cortex), as illustrated for the 2D invaginated circle in figure 4. From the outset, the justification for choosing a surface evolution model lay in its flexibility in capturing complex shapes, such as the cortex and arbitrarily diffuse pathologies, with minimum user intervention at the initialization stage.

Alternately, the 2D model could just as easily have been initialized with one or more

¹To be fair, significant progress has been made recently in making physical models more topologically adaptable[39, 21], although this comes at the cost of explicit collision detection, which can be expensive, and the extension from the identification of 2D contours to that of 3D surfaces is not as simple as the evolution framework adopted by us.




Figure 4: Demonstration of difference in tensile properties of evolution and physical models, on 2D invaginated circle: (a) contour evolution model at initialization ($t=0$), at 10 sec., and at 30 sec; (b) physical model at initialization, at 8th iteration, and at 30th iteration, courtesy of Johan Montagnat.

contours completely inside the boundary of interest, and have employed an outward moving isocontour as shown in figure 5. Finally, the method described here fully generalizes to identifying a 3D surface within a volume, requiring a inner or outer user-defined closed surface from which a signed distance map is computed with the volume as its domain, which serves as an initialization for the imbedding function $\Psi(x, y, z, t)$, as shown in figure 6.

The surface model proceeds from the signed distance map computed from the user-defined surface, and evolves according to the following expression [14, 61]:

$$\frac{\partial \Psi}{\partial t} = F(x, y, z) \|\nabla \Psi\| \left[\operatorname{div} \left(\frac{\nabla \Psi}{\|\nabla \Psi\|} \right) + \nu \right] + \nabla F \cdot \nabla \Psi, \quad (1)$$

This model features a *diffusive* term $\|\nabla \Psi\| \operatorname{div}(\nabla \Psi / \|\nabla \Psi\|)$ which tends to smooth out each isosurface, including the zero-level surface, a *hyperbolic* term $\|\nabla \Psi\| \nu$ which pushes the each isosurface forward in a discontinuity-preserving manner, and two image terms: a *speed* function F which slows down the zero-level surface near anatomical boundaries, and a so-called “*doublet*” term $\nabla F \cdot \nabla \Psi$ which prevents the zero-level surface from overshooting these boundaries. The sign of the ν constant in the hyperbolic term determines whether each isosurface is inward or outward moving.




Figure 5: Contour identification with outward-moving isocontour: zero-level isocontour (white), overlaid on invaginated circle image as in figure 4; at $t = 0$, at $t = 10$ sec. and at $t = 60$ sec..

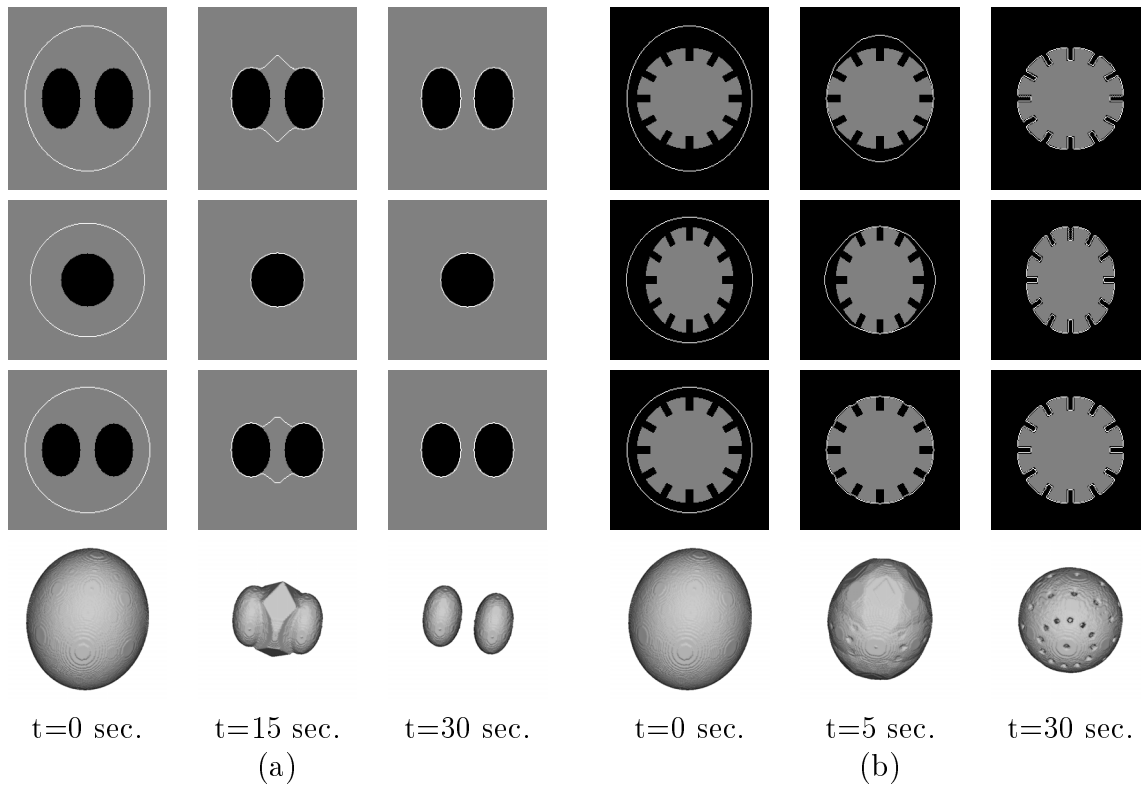


Figure 6: 3D surface identification with synthetic volume featuring (a) a pair of ellipsoids and (b) an invaginated sphere: multiplanar and rendered views of $\Psi(x, y, z) = 0$ at initial, intermediate and final instants.

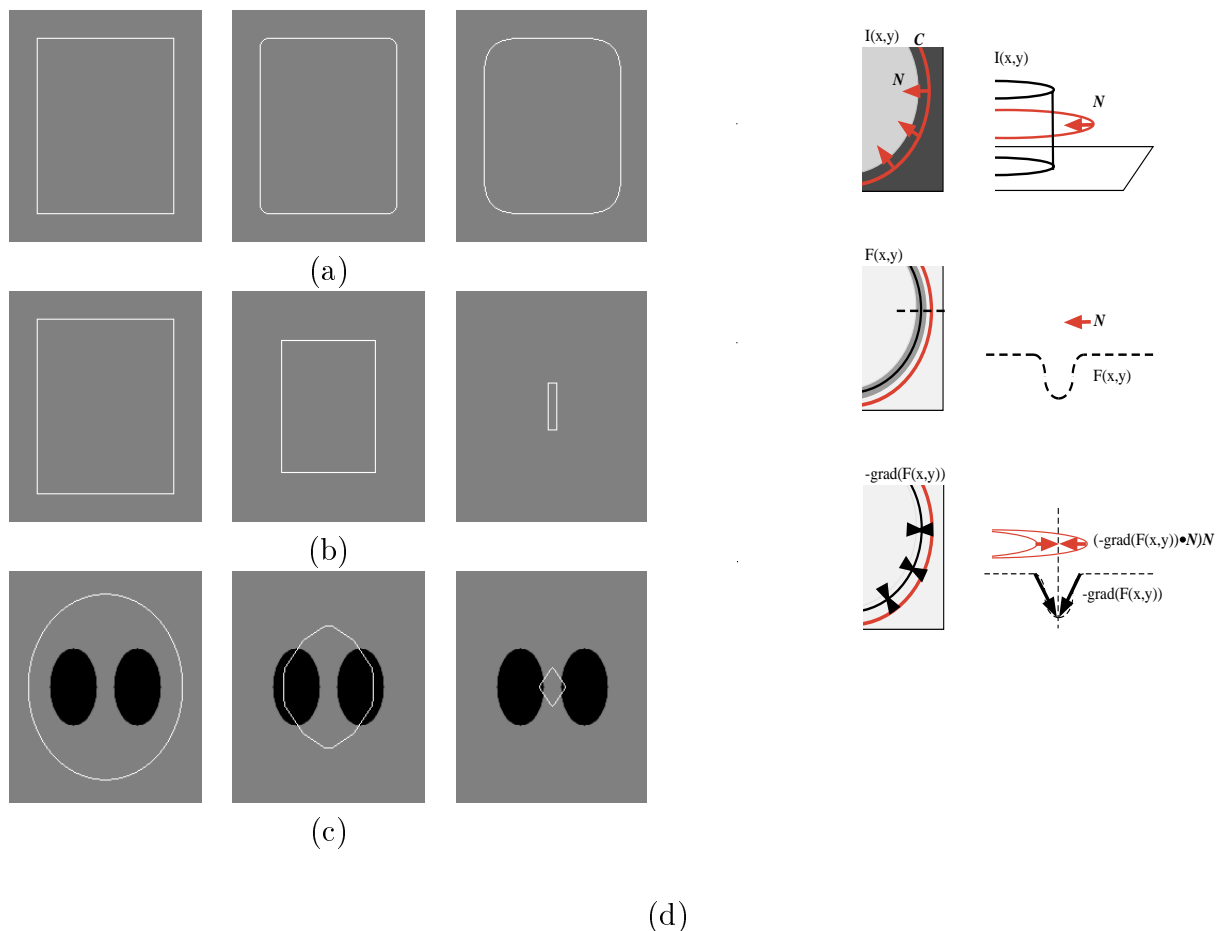


Figure 7: 2D illustration of each term at work: (a) diffusive: at 0 sec., at 10 sec., and at 120 sec.; (b) hyperbolic: at 0 sec., at 10 sec., and at 30 sec.; (c) speed and doublet: consider the model behaviour in the absence of these two terms (compare with fig. 2): at 0 sec., at 15 sec., and at 30 sec.. (d) Illustration in 2D of doublet term: raw 2D signal $I(x, y)$ with contour \mathcal{C} approaching boundary, speed function $F(x, y)$ derived from filtered signal, and doublet term $(\nabla F \cdot \vec{N}) \vec{N}$.

An intuitive illustration of each term in expression (1) appears in figure 7, featuring once again an analogy of lesser dimension, an evolving 2D contour. These illustrations fully extend to each term of the 3D surface evolution model. In figure 7, the diffusive term smooths the corners of the rectangular zero-level contour, nudging it towards an elliptical shape (and ultimately towards a circle, i.e.: tending towards constant curvature κ). The hyperbolic term shrinks the rectangle while retaining its discontinuities. These effects can be seen as being applied to each isocontour in $\Psi(x, y, t)$, including the zero-level contour.

The speed term slows the zero-level contour down at the tissue boundary in the image, and the doublet term constitutes a local trough, centered at this boundary, which prevents the zero-level contour from overshooting it: in their absence, the zero-level contour never halts at the underlying boundary.

The doublet term is further considered in figure 7 (d), and is best explained in 3D surface or 2D contour evolution representation. The level set representation of expression (1) corresponds to the evolution of a surface \mathcal{S} , along the inward normal direction:

$$\frac{\partial \mathcal{S}}{\partial t} = \left(FH - \nabla F \cdot \vec{N} \right) \vec{N}, \quad (2)$$

where H is the mean curvature of \mathcal{S} , and \vec{N} is its inward normal. If we consider the 2D analogical situation in figure 7 (d), we have an inward moving zero-level contour \mathcal{C} whose evolution is given by $\partial \mathcal{C} / \partial t = \left(F\kappa - \nabla F \cdot \vec{N} \right) \vec{N}$, where this time κ is the curvature and \vec{N} the inward normal of \mathcal{C} . In both the surface and contour expressions, the term $-\nabla F \cdot \vec{N}$ corresponds to the doublet term in level set representation. As shown here, the speed term F has a local trough at the boundary of interest, and the vector $-\nabla F$ points toward the minimum of this trough. For an inward moving contour, the dot product $-\nabla F \cdot \vec{N}$ outside the boundary is positive, and the doublet term is in the direction of the inward normal and pushes the zero-level contour towards the minimum. Inside the boundary, if the inward moving contour overshoots it, $-\nabla F \cdot \vec{N}$ becomes negative, and the doublet term points outwards and draws the contour back towards the minimum. This feature may be needed for a continuously-valued image-based speed function, such as one inversely proportional to a finite gradient magnitude, if this function in practice is never zero-valued.

2.2 Implementation Details of our Surface Model

To recapitulate the overview presented in the previous section, our anatomical surface identification technique is a model-based technique, known as a surface evolution model, which is characterized by imbedding the identification of a 3D surface \mathcal{S} into the evaluation of a function Ψ whose domain spans the tomographic volume:

$$\{X \in \mathfrak{R}^3 \text{ such that } \mathcal{S}(t) \equiv X : \Psi(X, t) = 0\}. \quad (3)$$

$$\frac{\partial \Psi}{\partial t} = F(x, y, z) \|\nabla \Psi\| \left[\text{div} \left(\frac{\nabla \Psi}{\|\nabla \Psi\|} \right) + \nu \right] + \nabla F \cdot \nabla \Psi, \quad (4)$$

where $\text{div}(\nabla \Psi / \|\nabla \Psi\|)$ represents the mean curvature H of an isosurface of the function Ψ . This model features a *diffusive* smoothing term $\|\nabla \Psi\| \text{div}(\nabla \Psi / \|\nabla \Psi\|)$, a *hyperbolic* motion term $\|\nabla \Psi\| \nu$, and two image terms: a *speed* function F and a *doublet* term $\nabla F \cdot \nabla \Psi$ which serve to bind the evolving surface to the anatomical boundary of interest. Moreover, the model is initialized by one or more user-defined 3D surfaces which fully contain or are fully

contained by the anatomical boundary, and the model evolves in a manner which nudges the zero-level isosurface inwards or outwards until the image terms bind it to the boundary. The imbedding function Ψ is initialized as a *signed distance map* from the user-defined surface(s).

In most existing implementations of surface evolution models, the speed term is strictly a function of the image gradient, i.e.:

$$F = \frac{1}{1 + \|\nabla \tilde{I}(x, y, z)\|^n}, \quad (5)$$

where typically $n=1, 2$ or 3 and $\tilde{I}(x, y, z)$ usually is a Gaussian-filtered image (MR) volume. However, as we will see in the next section, this approach has some limitations when it comes to identifying the brain surface, such as a lack of T1 contrast between grey matter and the sagittal sinus and between grey matter and muscle tissue, particularly near the eyes, possibly entailing a bleeding effect outside the brain surface. To alleviate this problem, we propose to endow our model with higher-level anatomic information, namely by first running a *voxel classification* algorithm on the tomographic volume(s), and to compute a speed function which restricts the surface model to expand only within white and grey matter, according to the following discrete relation:

$$\begin{aligned} F(x, y, z) &= 1.0 \text{ if } \{(x, y, z) \in WM\} \text{ or } \{(x, y, z) \in GM \text{ and near } WM\}, \\ &= 0.0 \text{ otherwise.} \end{aligned} \quad (6)$$

In practice, our implementation features the following stages, as illustrated in figure 8 and 9:

- Preprocessing of the MR volume(s) for the subsequent computation of the image-based speed term. Given that we want to eventually *register* the identified cortical surface, we prefer to preprocess the MR volume with an anisotropic diffusion operator rather than an isotropic filter, to limit unwanted blurring across relevant anatomical surfaces. The preprocessing we use is Kimia & Siddiqi's *curvature flow*[25], which is implemented within the level-set framework, as described in §2.5.1.
- A image-based speed term computation stage designed to allow the expansion of the model through white matter, to capture its boundary: $F(x, y, z) = 1.0$ if $\{(x, y, z) \in WM\}$ and $F(x, y, z) = 0.0$ otherwise.
- The identification of the WM boundary, which starts as a set of user-defined small spheres imbedded in white matter as the initial zero-level isosurface at $t = 0$, and computes the surface evolution model over (simulated) time by alternating between the following two methods:
 - The initialization of $\Psi(t_{l+1})$ as a signed distance map from the $\{\Psi(t_l) = 0\}$ isosurface at the previous iteration, which is computed with the Fast Marching method. This method is briefly described and validated in §2.5.3.

- The numerical implementation of expression (4) as a finite difference equation at the voxel $\mathbf{i} = (i, j, k)$ and at time $t_{l+1} \equiv (l + 1)\Delta t$, namely:

$$\begin{aligned} \Psi(\mathbf{i}, (l + 1)\Delta t) = & \Psi(\mathbf{i}, l\Delta t) + \\ & \Delta t \left\{ F(\check{\mathbf{i}}) \frac{\|(\Psi_x(\mathbf{i}), \Psi_y(\mathbf{i}), \Psi_z(\mathbf{i}))\|}{[H(\Psi(\mathbf{i})) + \nu]} + \right. \\ & \left. (F_x(\check{\mathbf{i}}), F_y(\check{\mathbf{i}}), F_z(\check{\mathbf{i}})) \cdot (\Psi_x(\mathbf{i}), \Psi_y(\mathbf{i}), \Psi_z(\mathbf{i})) \right\} . \end{aligned} \quad (7)$$

The numerical implementation of the second method follows standard classical level sets procedure. The label $\check{\mathbf{i}}$ denotes the voxel on the $\{\Psi(t_l) = 0\}$ isosurface which is closest to the voxel of interest \mathbf{i} . Since the image information only coincides with the zero-level, it must be “extended” from zero-level isosurface voxels to other voxels where expression (7) is evaluated.

Note that the narrow band restriction applies to both methods, whereby each is computed only within a thin volumetric shell close to the evolving $\{\Psi(t) = 0\}$ isosurface.

- A recomputation of the image-based speed term, according to expression (6), to allow expansion of the surface model through grey matter “close” to the WM boundary, to capture the outer brain surface.
- The expansion of the $\{\Psi(t) = 0\}$ isosurface to coincide with the outer cerebral and cerebellar surface, again based on alternating between signed distance map computation and the numerical evaluation of the surface evolution model.
- Post-processing of the outer brain surface to produce a Distance/Closest Point Map, for the subsequent non-rigid registration with the intra-operative range-based surface. This is simply the Fast Marching method alluded to in stage 1 of the surface model, applied not to computing a distance map from the evolving zero-level isosurface within a narrow band, but from the final surface everywhere in the volume. As will be seen, the identity of the final surface point closest to each voxel is relevant to the subsequent range-MRI surface registration in the OR.

2.3 Integrating Anatomical Information into the Model

Surface evolution models have recently been demonstrated on tomographic scans of the brain[36, 62]. These techniques are founded on an outward moving zero-level isosurface initialized completely within the brain (typically a few small spheres imbedded in white matter). An inward-moving surface initialized at the skin surface is conceivable, but not likely to produce good results at the base of the brain, which features muscle and bone tissue boundaries distal to the brain which can easily trap the evolving surface model away from the brain. Hence the predominance of outward-moving approaches initialized from within.

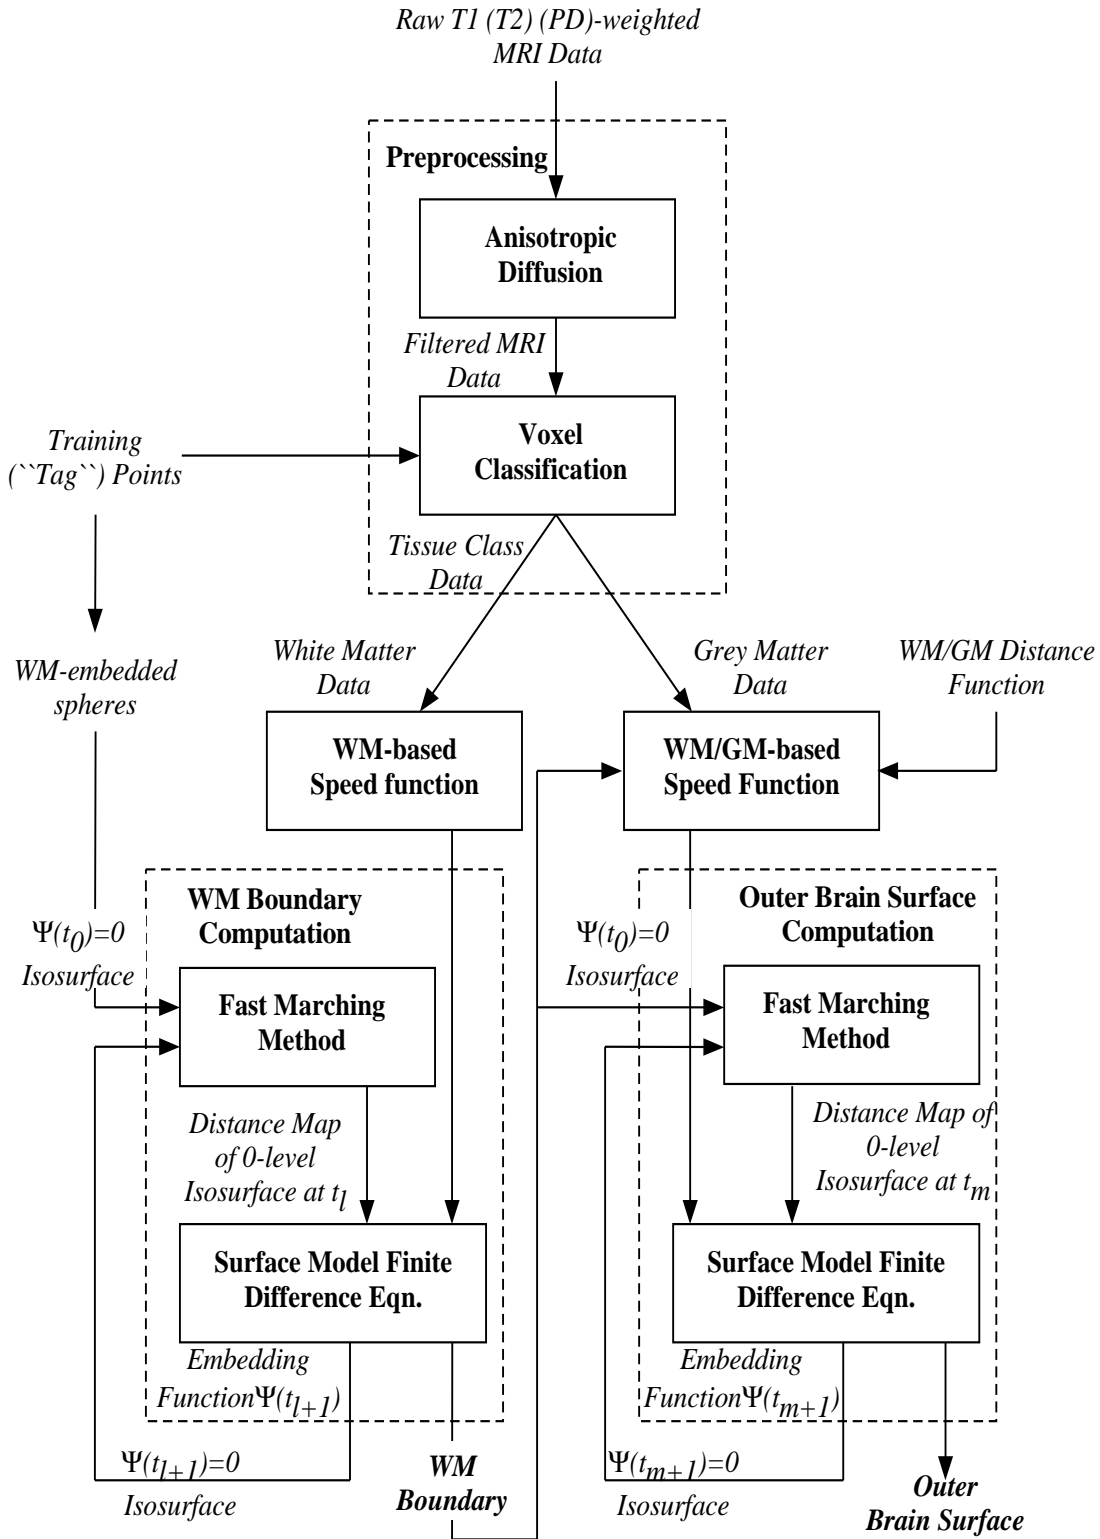


Figure 8: Algorithmic description of our surface evolution model framework.

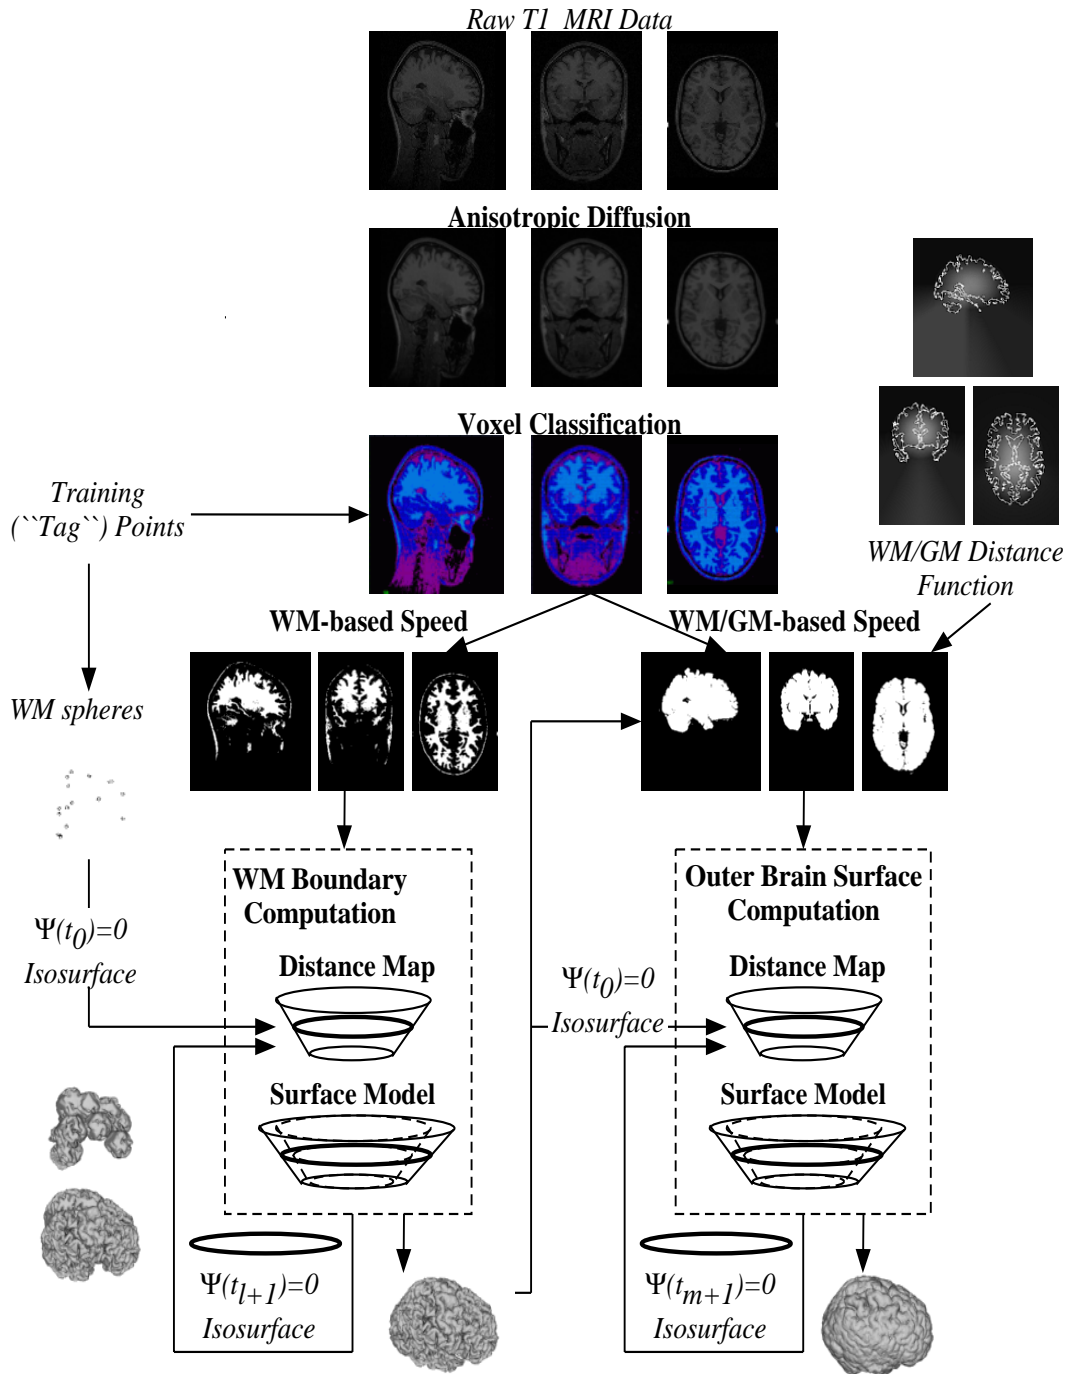


Figure 9: Pictorial description of our surface evolution model framework.


The problem faced by this approach however is how to prevent the evolving surface from getting snagged in the white-grey matter boundary, also characterized by a strong gradient, albeit a weaker one than the outer cortical surface (grey matter-CSF boundary), as illustrated figure 10(a). A simple approach would be to threshold gradients and only slow down the evolving surface at very strong gradients, i.e.: outer cortex and lesion boundaries. But automatically finding such a threshold is elusive, as is guaranteeing that there won't be gaps in the halting force on the outer cortex, where the local intensity gradient falls below this threshold. A bleeding effect which ensues, through such gaps, is illustrated in figure 10(b).

Zeng[62] and MacDonald[34] each propose a solution featuring two evolving surfaces which are bound to each other by a distance constraint and which expand to segment the inner and outer surfaces of a healthy cortex. This elegant idea confers on their respective models higher-level anatomical information which stabilizes their convergence to inner and outer cortical surfaces. However, the assumption of the white matter boundary always being contained by the grey matter boundary may break down in the presence of pathologies. Moreover, the tight distance threshold which is appropriate for the outer cortex, useful to prevent the surface model from bleeding into the eyeballs or sagittal sinus, is inappropriate for the cerebellum and for the area near the ventricles, which are characterized by comparatively large expanses of grey matter.


For these reasons, we adopt a different approach to providing our surface model with higher-level anatomical knowledge, namely we integrate the results of a *voxel classification algorithm*[16, 28] into the image-based speed term. Briefly, a classifier maps an image voxel, represented for D modalities as a D -dimensional feature vector in \mathfrak{R}^D , $\mathbf{x} = (x_1, x_2, \dots, x_D)$, onto a tissue class C_j , where $j = 1, \dots, c$, and c is the total number of tissue classes in the scan. This concept is illustrated on figure 11. We are interested in a technique whereby the user can provide anatomical knowledge by simply clicking on a few (5-20) *tag points* for each class. The classifier training points for white matter can in turn also serve as the centers of the outward moving spheres used in the subsequent level sets surface identification. We could also use spheres centered at grey matter points, but as will be seen, we operate in the following stages: move outward from imbedded WM points to the WM boundary, determine GM points "close" to this WM boundary, move to the outer boundary of these GM points.

Our motivation here is to select a classifier which can exploit multiple modalities, which does not degrade too badly if only one (typically T1) modality is available (which is often the case in a clinical setting, although the scanned tissues may feature contrast agents), which is able to deal with the presence of pathologies, and which can function with limited supervision. We investigate a variation of Fuzzy C-Means (FCM) clustering, called the *semi-supervised FCM* (ssFCM) algorithm, which requires only limited supervision and which in fact has been used to classify voxels of brain scans featuring pathologies[7]. We also investigate the *Minimum Distance* (MD) classifier[28, 22], which can function with a relatively small training set and constitutes the first iteration of the ssFCM classifier. Consideration of the two allows us to easily establish the benefit of clustering after the initial partition of parameter space.

Our method has the following advantages:



(a)



(b)

Figure 10: Tri-planar view of evolving cortical surface model with its image term solely based on gradient magnitude. (a) With a low gradient threshold (allowing points of low gradient magnitude to halt the zero-level surface), the surface may converge to the WM boundary rather than the outer brain boundary. (b) With a higher threshold, the surface can bleed through the outer brain boundary, and never actually converge to a solution.

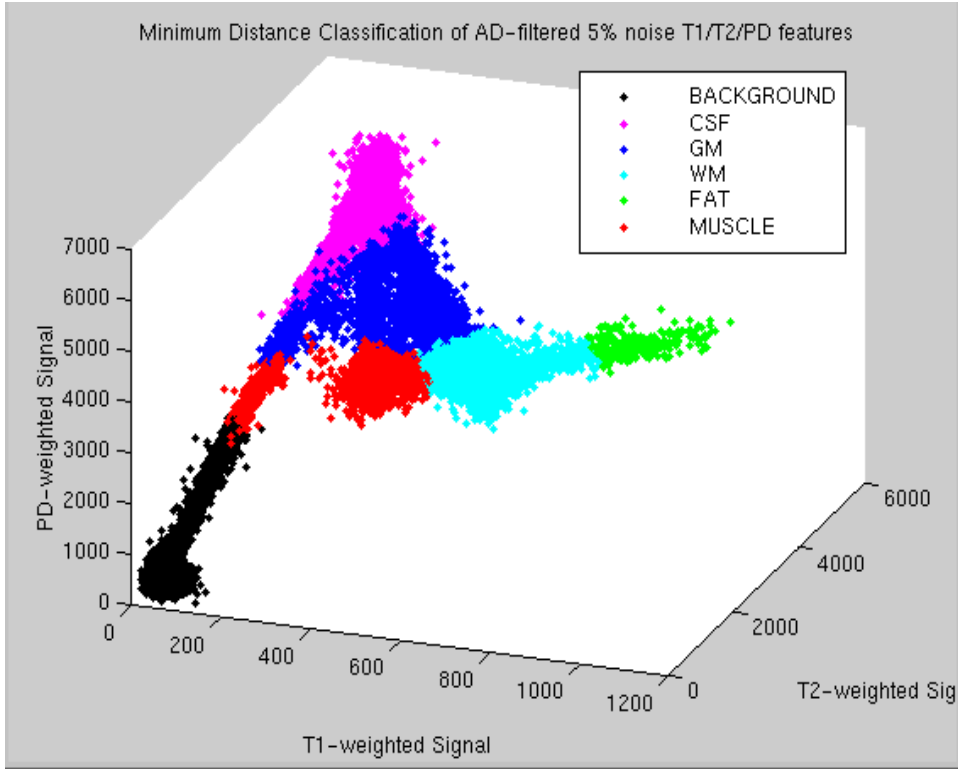


Figure 11: Illustration of voxel classification, for T1-, T2- and PD-Weighted feature space, based on the middle coronal slice of simulated data with 5% noise, preprocessed with anisotropic diffusion ($t = 0.5sec$).

- Integrating tissue class information into the speed term allows us to constrain our model in a manner comparable to Zeng and MacDonald: we can restrict the zero-level isosurface to evolve only in voxels classified as white matter or as grey matter and deemed close to white matter, thereby *preventing undesired bleeding effects*. Moreover, there is no reason for this distance tolerance, which enforces grey/white matter boundary proximity, to be a constant over the whole volume. As illustrated in figure 12, this **WM/GM Distance Function** is set to some maximum value at the centroid, and tends to a minimum parameter at the outer edges of the cortex. Moreover, the minimum value at the edge of the cortex also varies spatially.
- Secondly, our technique, which is rooted in a speed function rich in anatomical information, is highly *flexible*. Pathologies are easily dealt with, provided they can be discriminated from white and grey matter[9].
- Next, this approach to constituting a speed function also provides us with an efficient way of taking into account *multiple scan modalities*, if available: the classifier simply maps the parameter space encompassing T1-, T2- and/or proton density-weighted volumes to a tissue class volume spatially coinciding with these tomographic volumes.




Figure 12: Illustration of WM/GM Distance Function: (a) Distance Function mapped to a grey-level intensity; (b) speed function based on membership of voxels classified as WM or GM, using T1 alone without a contrast agent (note that there little discrimination from T1 alone between GM and muscle, skin or sagittal sinus voxels); (c) speed function based on WM/GM voxels, as in (b), but excluding GM voxels which exceed the Distance Function shown in (a).

- Moreover, this approach is easily adapted to *other anatomical applications*, at the cost of a different voxel classification training, and alternate use of voxel class information, possibly involving some sort of distance constraint between two classes, possibly not.

The ability to cope with conflicting requirements in white to grey matter boundary proximity is handled by the nature of the WM/GM Distance Function:

- It is tightest at the outer cortex (above the centroid along the proximal-distal axis), e.g.: near the sagittal sinus.
- It is somewhat looser at the base of the brain where the grey matter tends to be thicker than in the cortex.
- It is more permissive still at the cerebellum, where the grey matter is thickest.
- It is loosest at the centroid of white matter, where bleeding of the surface model is less a concern, and which coincides with grey matter in the area of basal ganglia, neighbouring the ventricles[12].

One point worth emphasizing here is that the WM boundary, to which proximity of GM voxels is established, is the one identified by first running the surface model only on exclusively on WM voxels. An alternative would be to simply compute a distance map from all voxels putatively identified as WM by the classifier, but this leaves the surface

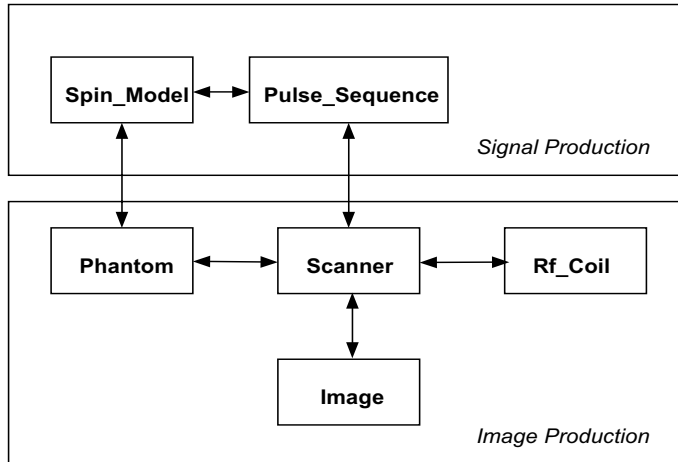


Figure 13: An overview of the simulation process (adapted from [29], courtesy of G. Bruce Pike).

model susceptible to isolated points misclassified as WM, which would drastically reduce the effectiveness of our coupling. Our way of proceeding ensures that only voxels classified as WM which are contiguous with user-selected points (known to be true WM points) are retained to anchor the identification of the outer cerebral and cerebellar surfaces.

2.4 Validation Tools: MRI Simulator and Realistic Digital Brain Phantom

We propose to validate the various stages of our brain surface identification technique with a quantitative approach based on realistic synthetic data, where ground truth is known, and qualitative tests with *in vivo* subject data. The latter data comes in the form of healthy subject scans obtained from the ICBM² database, as well as MNI patient scans featuring pathologies. While patient data can be exploited quantitatively by comparing image processing results against manual analyses performed by experts, the latter process is known to be labour-intensive and suffers from both intra- and inter-expert variabilities[30, 65]. In contrast, our quantitative validation stage makes use of two complementary software tools which were developed at the MNI and which eliminate these variabilities: a MRI simulator specifically designed by Kwan et al. for the purpose of image analysis evaluation[30], and a high-resolution, volumetric, anthropomorphic digital brain phantom elaborated by Collins et al.[17].

The MRI simulator features *Signal* and *Image Production* modules, the former of which

²International Brain Mapping Consortium[57].

is comprised of *Spin Model* and *Pulse Sequence* objects, while the latter is comprised of *Phantom*, *Scanner*, *Rf Coil* and *Image* objects[30]. The architecture of the system, emphasizing the interaction between objects, appears in figure 13, while an illustration of the 3D brain phantom and typical T1-, T2- and PD-weighted images resulting from it appear in figure 14. The Signal model is dedicated to computing tissue intensities on the basis of the magnetization properties, made explicit in the Spin Model, and of the signal resulting from Pulse Sequence object. The Image Production module is controlled by the Scanner object and generates tomographic images from computed tissue intensities, and featuring imaging effects such as noise, RF inhomogeneity, and partial volume. The Phantom object synthesizes the patient anatomy in the scanner as described above, the RF Coil object models signal reception, and finally the resulting tomographic scan is stored in the Image object.

The design of the digital brain phantom is motivated by the limitations of *in vivo* studies, namely the implausibility of establishing with certainty the exact boundaries of tissues in a tomographic scan (without slicing and staining the underlying anatomy), and those of previously existing phantoms, both physical and digital, whose geometrical simplicity has limited their predictive ability and anatomical relevance[17]. This phantom is based on single-subject average dataset (compiled from 27 scans) which was semi-automatically classified and subsequently manually corrected by a trained neuroanatomist. Each voxel is labelled as one of 10 classes, namely³: *Background*, *Grey Matter*, *White Matter*, *Cortico-Spinal Fluid*, *Fat*, *Muscle*, *Skin*, *Skull*, *Glial Matter*, and *Other*. The phantom comes in two forms, respectively conveying *discrete* and *fuzzy* tissue membership. While the latter is judged to represent partial volume effects most realistically, we have chosen to work with the discrete version because of our requirement of precise boundaries in characterizing the accuracy of the various stages of our anatomical surface identification technique.

2.5 Validation of the Surface Model

2.5.1 Anisotropic Diffusion Validation

The validation of the anisotropic diffusion preprocessing comprises a comparison of the Kimia level set-based technique[25] adopted by us,

$$\begin{aligned} \psi_t &= H \|\nabla\psi\| \text{ where} \\ H &= [\psi_{xx}(\psi_y + \psi_z) + \psi_{yy}(\psi_x + \psi_z) + \psi_{zz}(\psi_x + \psi_y) \\ &\quad - 2\psi_x\psi_y\psi_{xy} - 2\psi_x\psi_z\psi_{xz} - 2\psi_y\psi_z\psi_{yz}] / [\psi_x^2 + \psi_y^2 + \psi_z^2]^{3/2}, \end{aligned} \quad (8)$$

with the classical Gaussian filtering, as well as a series of tests to determine an optimal scale factor under typical imaging conditions. In contrast with our surface evolution model

³However, the T1, T2 and PD values of the latter 4 classes mirror those of 4 of the first 6, which is equivalent to a labelling according to 6 tissue classes. The realistic assignment of magnetization properties of the latter 4 classes was not judged to be essential to the applicability of the phantom[18]. For this reason, all illustrations of the phantom in this thesis feature 6 classes, distinguished on the basis of magnetization properties.




Figure 14: (a) Illustration of 3D digital brain phantom, with discrete class membership mapped to grey-level; (b) typical T1-weighted volume resulting from simulation, with discrete phantom in (a) used as input; (c) T2-weighted volume with same phantom input; (d) PD-weighted volume with same phantom input.

application, in this application the function $\psi(x, y, z)$ is not initialized with a signed distance map, but with the raw tomographic volume. The smoothed output is simply the value of ψ at the simulated time duration, based on a finite differences discretization of expression (8). Our comparison exploits Koenderink's Diffusion Equation formulation for Gaussian filtering[27],

$$\psi_t = \nabla^2 \psi, \quad (9)$$

which allows us to easily look at the two formulations at identical scales, *characterized by the time duration of each diffusion*. The relationship between the spatial extent of the convolution implementation and the time duration of the diffusion implementation of the Gaussian kernel is given by[27]:

$$2\sigma^2 = 4t. \quad (10)$$




Figure 15: Illustration of effect of scale parameter, for (a) simulated T1 volume (with null noise and inhomogeneity values) on (b) Gaussian and (c) anisotropic diffusion.

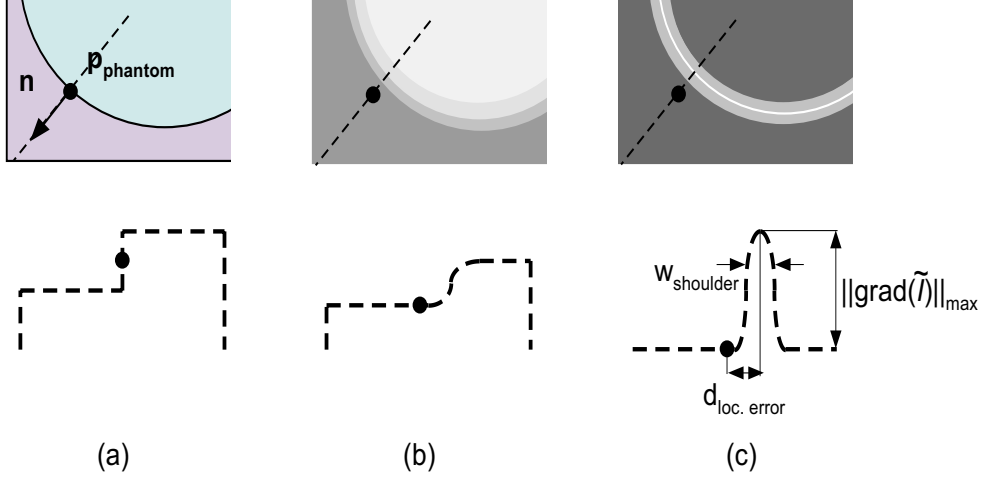


Figure 16: 2D Illustration of measurements for validation and choice of scale of anisotropic diffusion: (a) phantom class data, with 1D profile along normal direction, through known boundary point; (b) filtered simulated T1-weighted MR data, with 1D profile as in (a); (c) magnitude of the gradient of filtered T1-weighted data, with 1D profile as (a) and (b), with measurements shown.

Figure 15 provides a qualitative illustration of the difference between anisotropic and Gaussian diffusion, over a range of scales. The quantitative comparison between the two types of diffusion comprises the evaluation of the strength and localization of the edges (an indication of the position of the outer brain boundary ultimately determined by our surface model) detected in simulated T1 volumes (under various noise and RF inhomogeneity conditions), given the known position of the brain boundary in the underlying digital phantom, as well as the evaluation of the voxel classification technique from T1 volumes (again under various noise conditions), given the known class membership of the voxel of the underlying digital phantom. The two tests are done at different scales for both formulations. The second test is deferred to the validation of the voxel classifier, and is discussed in §2.5.2.

The edge test uses as a starting point the position of all tissue boundary of the digital phantom, assumed to be ground truth, and computes three measurements from the diffused T1 volume, under different noise, inhomogeneity and diffusion scale characteristics, in the immediate neighbourhood of this boundary point, as illustrated in figure 16: average value of the *maximal gradient amplitude*, in this immediate neighbourhood of known boundary points, average *localization error*, along the boundary normal direction, between this maximum gradient point and the boundary point and the average *width of the shoulders* (coinciding with 80% of the maximum) of the gradient magnitude profile, along the normal direction.

This validation study, whose results appear in figure 17, shows that over the range of scales of interest, anisotropic diffusion better retains gradient magnitude information which

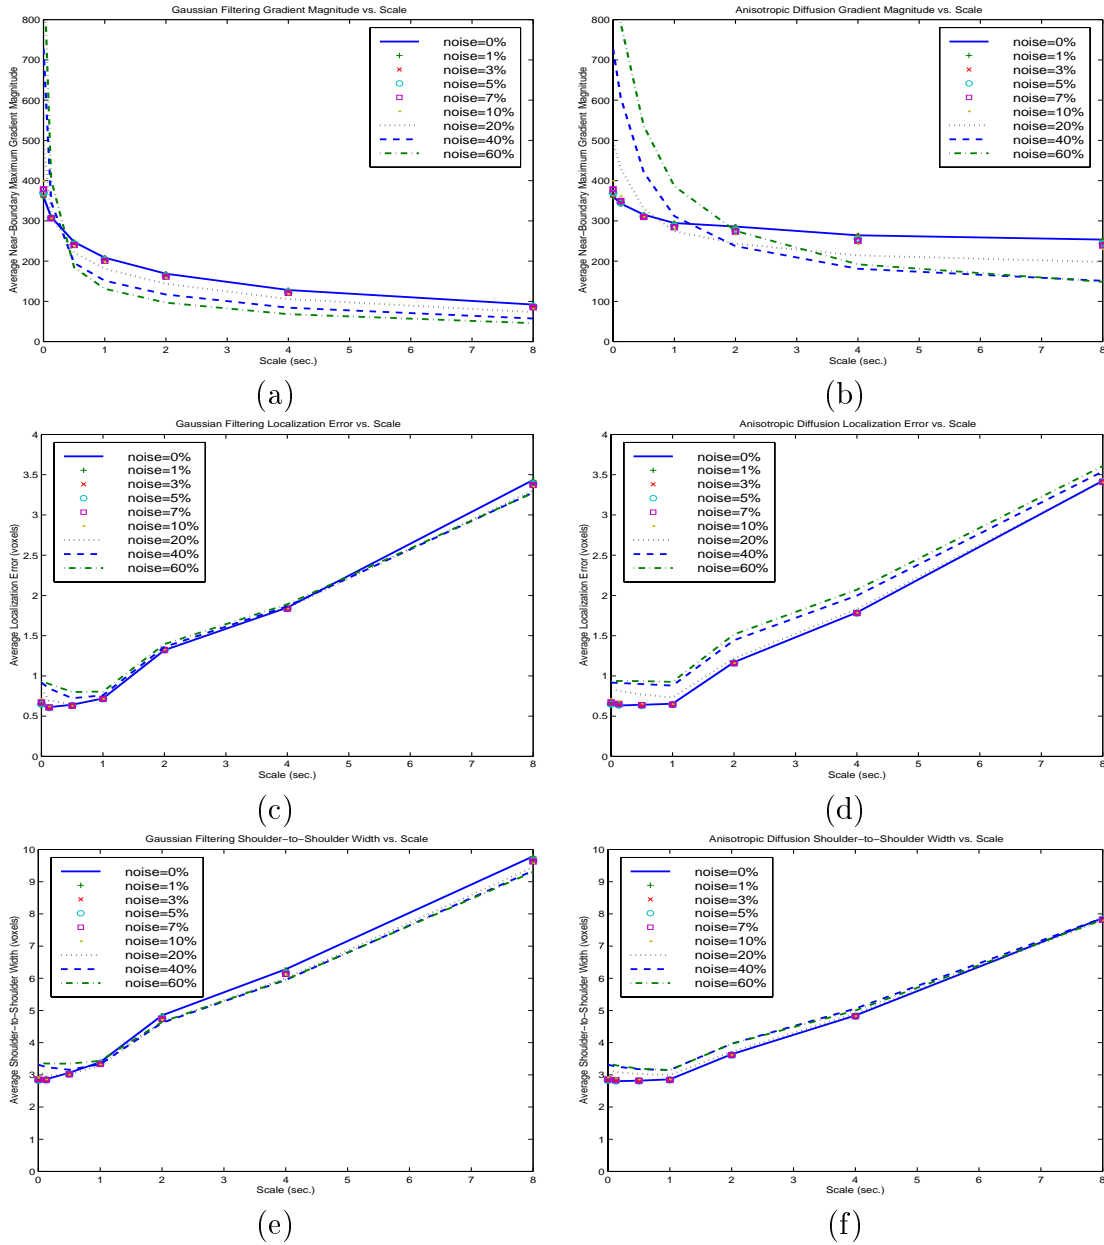


Figure 17: Measurements statistics for validation and choice of scale of anisotropic diffusion, under different noise and 0% RF inhomogeneity conditions: average local maximum gradient magnitude statistics for (a) Gaussian and (b) anisotropic diffusion; average localization error of local gradient maximum from known boundary position, for (c) Gaussian and (d) anisotropic diffusion; average shoulder (80% of gradient magnitude maximum) width value along the boundary, for (e) Gaussian and (f) anisotropic diffusion.

can be used to find tissue boundaries. Moreover, this information also is comparably or better localized for anisotropic diffusion, depending on the choice of statistic used to quantify position accuracy. A comparison of the two diffusion techniques as they relate to voxel classification is pursued in the next section.

2.5.2 Validation of Semi-Automatic Voxel Classification

The objectives of this validation study are to determine the feasibility of using either the semi-supervised Fuzzy C-Means (ssFCM) or Minimum Distance (MD) algorithm to map a T1 value or a feature vector (T1-T2 or T1-T2-PD) to the correct tissue class C_i (which is known a priori for this study, by virtue of the digital phantom), and in particular to choose one method over the other, based on which modalities are available, and based on assumptions about noise (we assume once again that RF inhomogeneity can be corrected as in [51]). The effect of noise and of the choice of filtering method (including a choice of scale), on classifier performance is also featured here.

Figure 18 provides a comparison of the two filtering options as they relate to the Minimum Distance voxel classification. This study uses Minimum Distance rather than semi-supervised Fuzzy C-Means because of the nonlinear effect of clustering on the partitioning of parameter space, given that we want to isolate the effect of noise distribution, diffusion type or scale factor. Under all but severe noise conditions (20% and on), anisotropic diffusion is comparable at very small scales and is otherwise superior to Gaussian filtering. Moreover, we would not expect to use a classifier with any success under those conditions severe noise conditions.

An observation worth making about figure 18 is that for higher scales, $t \geq 2.0$, at low noise values ranging from 0 to 5%, the classification doesn't appear to gain much from the added consideration of T2- and PD-weighted scans. This anomaly can be cleared up by considering erroneous voxels summed in accumulators at the pixels in one plane, say the sagittal plane, as shown in figure 19. In this case, the intensity of the pixel is proportional to the number of misclassified voxels at the same (y, z) coordinate, over all x values. It's clear that misclassified voxels in the multimodality case are grouped in bands, coinciding with the borders of tissue classes, and that much of the misclassification in this case is related to border effects. It stands to reason that these border effects would worsen with increasing scale, which is seen on all combinations of modalities. Moreover, the consideration of more than one modality improves the classification of brain tissue voxels, particularly at finer scales. This is borne out by considering separate classification statistics for white and grey matter voxels, which will appear in the thesis.

Finally, figure 20 illustrate the fragility of the convergence properties of the semi-supervised Fuzzy C-Means algorithm, which in turn justifies our decision to ultimately opt for the Minimum Distance classifier. Figures (a), (b) and (c) illustrate the classification results at various noise levels and at various scales of anisotropic diffusion, respectively based on a T1-weighted scan alone, on a T1/T2 feature and on a T1/T2/PD feature. Figure (d) illustrates the migration of class centroids due to clustering, from points marked "O" to points marked "X", for T1- and T2-weighted data with 5% noise, processed with anisotropic diffusion at scale

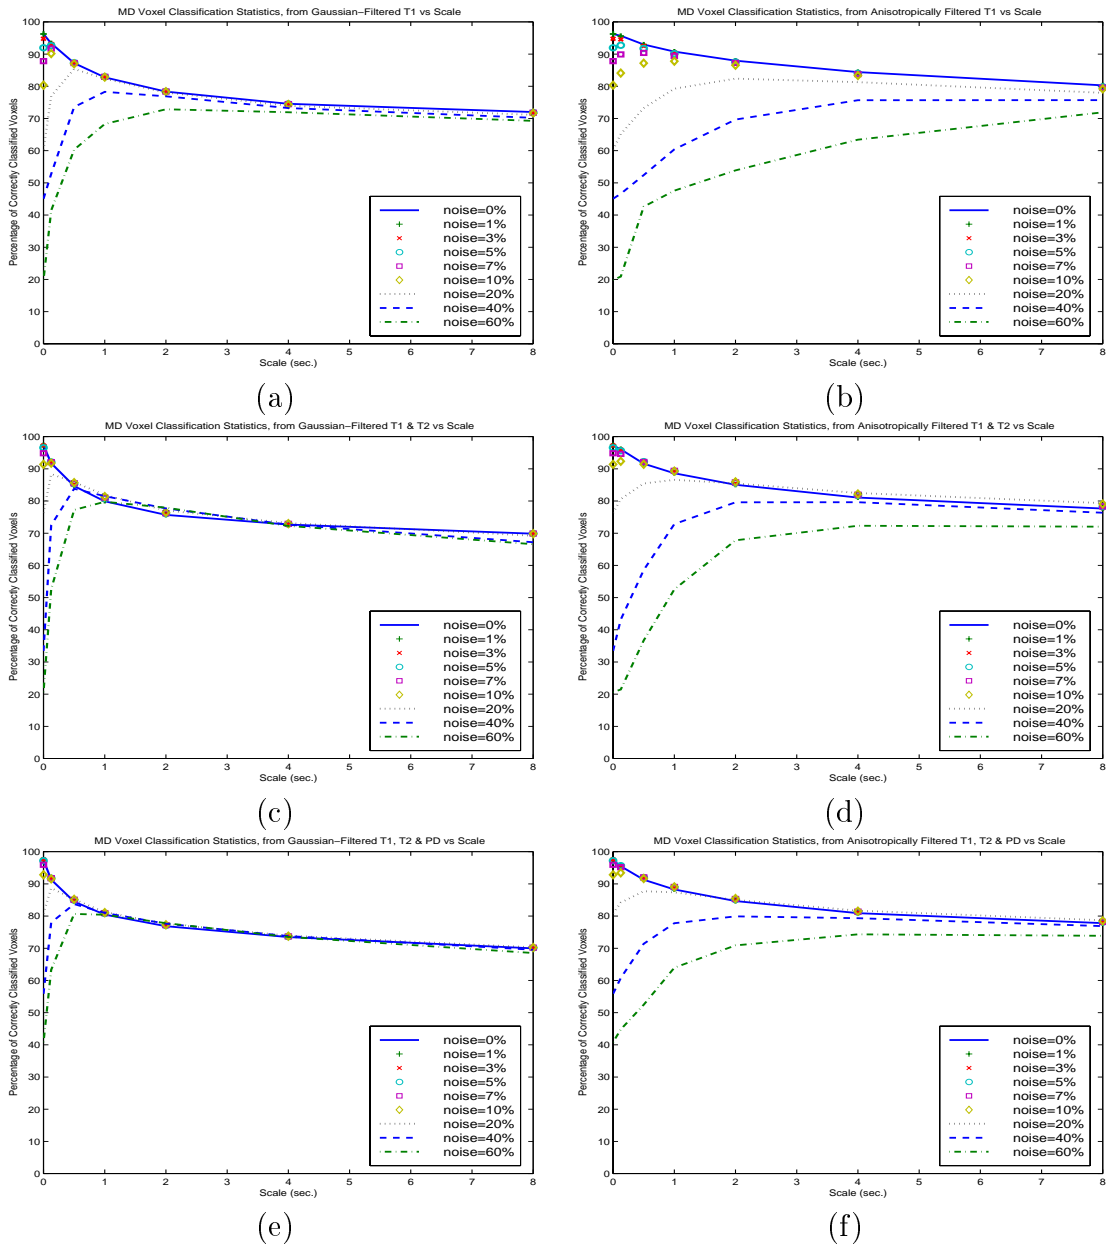


Figure 18: Illustration of sensitivity of MD classification to noise levels, to preprocessing method and scale, and to number of modalities: based on T1-weighted scan alone, at various noise levels, under (a) Gaussian filtering (GF) and (b) anisotropic diffusion (AD); based on T1- and T2-weighted scans, under (c) GF and (d) AD; based on T1-, T2- and proton density-weighted scans, under (e) GF and (f) AD.




Figure 19: Spatial distribution of misclassified voxels, based on accumulators in the $y - z$ (sagittal) plane. The brighter the pixel, the larger the number of misclassified voxels having that x position. Results for (a) T1-weighted scan alone and (b) T1- and T2-weighted scans.

$t = 2.0$ (this coincides with the dip in graph (b) at this scale value). Border effects between CSF and grey matter, which worsen with increasing scale, can draw the GM class centroid towards the CSF area of parameter space, causing in turn a significant portion of GM voxels to be misclassified as WM or MUSCLE.

In short, this validation study confirms the choice of anisotropic diffusion as optimal preprocessing, based on subsequent classifier performance. This study suggests a choice of scale between $t = 0.125$ and $t = 0.5$, which corresponds to the $\sigma = 0.5$ to $\sigma = 1.0$ range. Finally, the fragility of semi-supervised Fuzzy C-Means technique, particularly in response to border effects, justifies our decision adopt the Minimum Distance classifier.

2.5.3 Signed Distance Computation Validation

The validation of the signed distance computation is a comparison between our method of choice for computing the distance to the surface, the Fast Marching (FM) method, and a gold standard obtained by a “brute force” (BF) exhaustive search for the brain boundary point closest to each voxel. In other words, the BF technique explicitly finds the closest surface point for a given voxel by determining the Euclidean distance between this voxel and each brain boundary point, and associates with this voxel the distance to and the label of the closest surface point⁴, and this is compared with the distance and closest point label obtained via the Fast Marching technique.

Briefly, the idea behind the Fast Marching technique is that the computation of distance from a surface or contour can be idealized, as shown for a simple 2D contour illustration in figure 21, as equivalent to determining the “arrival time” of an outward moving surface (or contour), assuming it has monotonic unit speed. In other words, if we describe our initial

⁴On a 200 MHz Pentium, the BF approach took more than 2 weeks to produce the equivalent map computed in less than 20 minutes by the FM method, for the $181 \times 217 \times 181$ scan volume comprising the digital phantom, although improvements to the exhaustive search have proposed, such as the k-D tree[63].

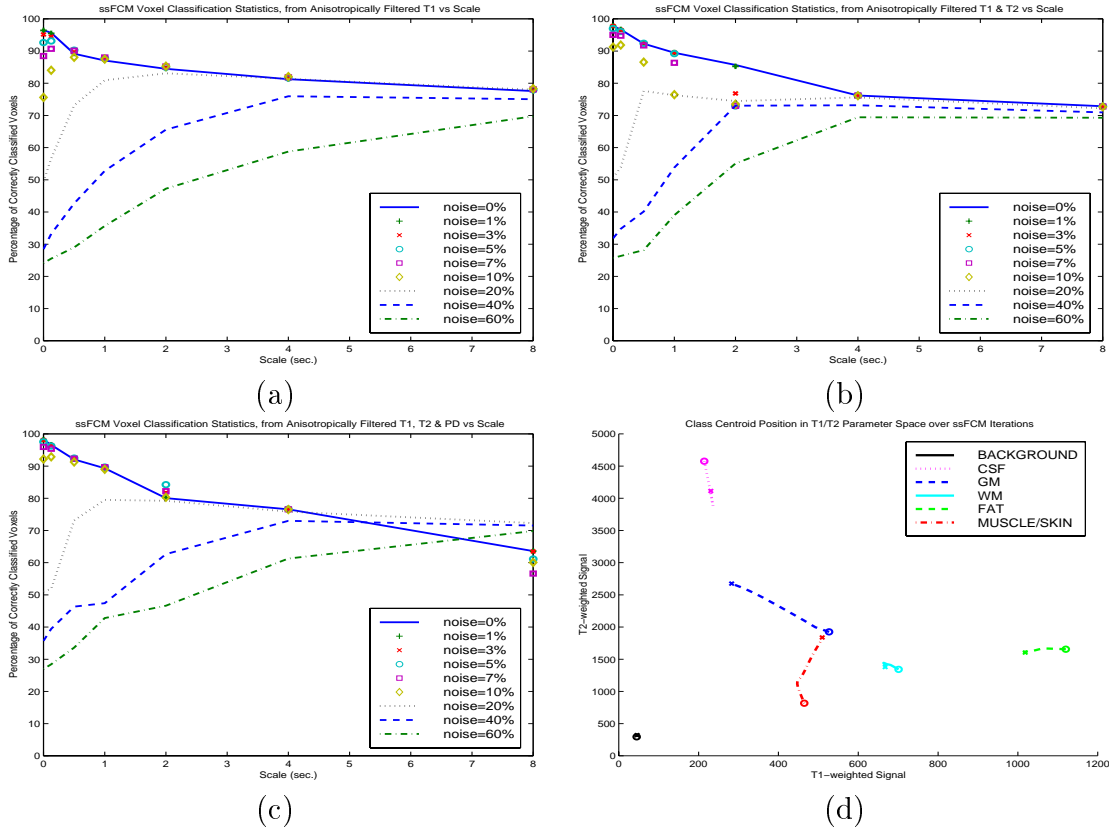


Figure 20: Illustration of sensitivity of FCM classification to noise levels, to preprocessing method and scale, and to number of modalities: based on (a) T1-weighted scan alone, (b) on T1- and T2-weighted scans, and (c) on T1-, T2- and proton density-weighted scans, all after anisotropic diffusion; (d) illustration of migration of centroids over successive iterations caused by border effects at noise=5% and scale $t=2.0$ sec., where a “o” illustrates the position of the class centroid at the first iteration (coinciding with MD classification) and an “x” indicates the final (25th) iteration. Note the final position of MUSCLE/SKIN class centroid nearly coinciding with the initial GM centroid position.

surface (contour) as a front with unit speed, and picture individuals with a stopwatch at each voxel (pixel) in the volume (image), those who are at a point traversed by the front at time $t = 1$ are in effect 1 pixel away from the original surface (contour). Those which see the front at $t = 2$ are 2 pixels away, and so on. Similarly, we can describe the distance to inside points in terms of the arrival time of an inward moving front originating at the initial surface and assign a negative sign to interior points, to produce the signed distance map required to initialize Ψ . The motion of the surface or contour can be expressed as the




Figure 21: Computation of Distance and Closest Point Map for 2D contour, featuring the “arrival time” of the contour: \mathbf{p}_0 is the closest contour point for \mathbf{p}_1 and \mathbf{p}_2 .

*Eikonal equation*⁵:

$$\|\nabla T\|F = 1. \quad (11)$$

Here, T is the arrival time of the surface, this expression states in effect that the gradient of this arrival time is inversely proportional to the speed of the surface.

This expression leads to a highly efficient and numerically stable method for estimating the monotonic motion of a surface, and assuming uniform unit speed, for numerically estimating the distance map to a surface. The method uses *upwind differencing*[47] to ensure numerical stability, which can be likened to a front being swept outward, whose motion estimation is based on points where the front has been and whose motion has already been computed, rather than on points where the front is going. It exploits a *minheap*[48] structure to efficiently implement this type of difference equation. Details appear in [49], pp. 86-100.. Finally, this algorithm is easily modified to not only store the *distance to the closest surface point* at a given voxel, but the *label of the closest surface point* as well, which allows us to compute a Closest Point Map from the identified outer brain surface.

The brain boundary in question is obtained by first binarizing the digital phantom volume, by setting to 1 the grey, white and glial matter voxels and setting to 0 all other voxels, and then computing the gradient everywhere. Because the volume is binary, the gradient should be null everywhere but on the boundary. Moreover, we restrict our attention to the

⁵In 1D, this equation is derived in [49]: considering $T(x)$ the arrival time and using *distance = rate × time*, we have $1 = F \frac{dT}{dx}$. In multiple dimensions, ∇T is orthogonal to the level sets of T , and its magnitude is inversely proportional to speed.

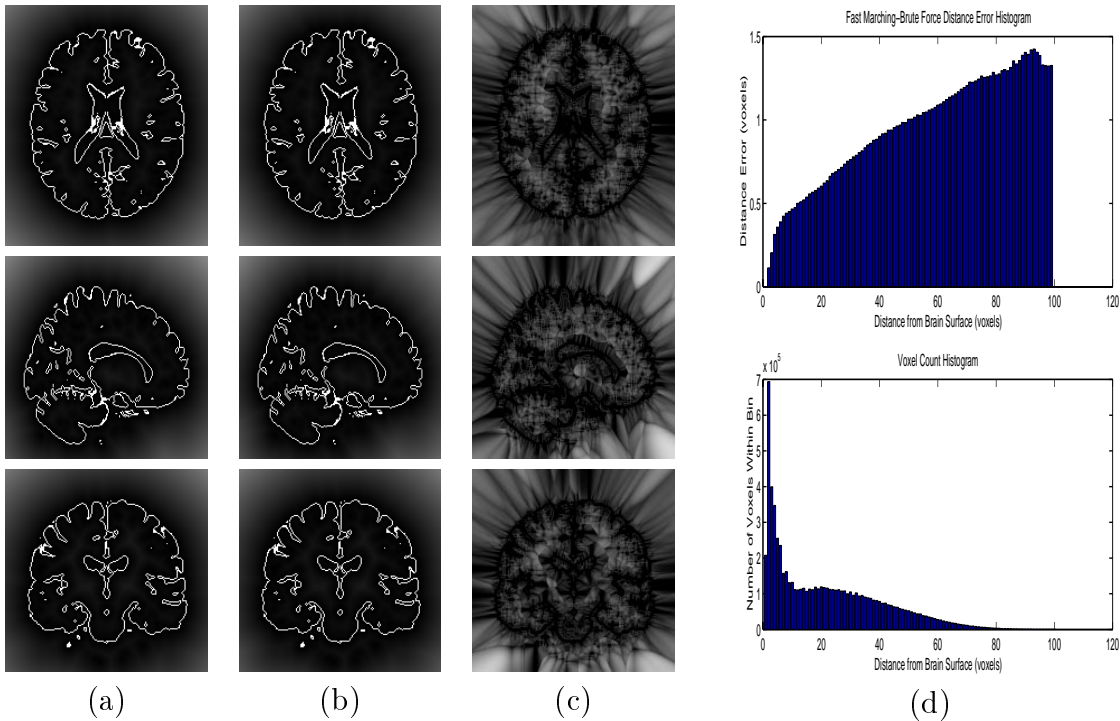


Figure 22: (a) Distance Map to brain surface produced by Fast Marching (FM) technique, with distance values ranging in the volume from 0 to 100 pixels are mapped to grey-levels 0 to 192, and brain boundary appearing in white; (b) Distance Map to brain surface produced by Brute Force (BF) technique, with same distance to grey-level mapping as in a; (c) difference between respective Distance Maps, shown as grey-level: a voxel error ranging from 0.0 to 2.0 voxels maps to 0-255; (d) the difference between FM and BF, averaged over voxels of similar distance (0-1.0 voxel, 1.0-2.0 voxels, ..., 99.0-100.0 voxels) to the brain surface: the average FM-BF difference value for each bin appears in the top histogram, whereas the number of points summed in each bin is shown in the lower histogram.

boundary voxels which are taken within the voxels set to 1, which correspond to inside edge points.

The distance map for each technique is shown in figures 22 (a) and (b), in tri-planar grey level image format, where a maximum distance of 100 voxels in the volume maps to a grey-level of 192 (out of a possible 255), and where the boundary itself is overlaid in white. The difference between the two maps is shown qualitatively in tri-planar format in figure 22 (c), where 255 represents a maximum difference of 2.0 voxels. Moreover, the top histogram appearing in figure 22 (d) provides a quantitative illustration of the difference between the two techniques: the difference value at each voxel was binned according to the reference (Brute Force) distance to the brain surface (0.0-1.0 voxel, 1.0-2.0 voxels..., 99.0-100.0 voxels) and the average computed for each bin. The size of the summation is shown in the bottom histogram in figure 22 (d). As one would expect, particularly for a distance




Figure 23: Closest Point Maps computed by 3D (a) FM and (b) BF techniques, for the brain boundary, labelled white, with surface labels 1-207459 mapped to grey levels 64-192, and with ambiguously labelled voxels shown as black (0); in (c), a bar graph illustrates the percentage of labels in agreement, as a function of distance to surface, based on the same voxel populations presented in figure 22 (d); (d) FM Closest Point Map with ambiguities due to digitization artefacts eliminated with a distance threshold of 2 voxels.

computation which is based on a propagative front emanating from the original surface, the error is cumulative over distance, and monotonically increasing (except for the very tail which factors in very few data points).

However, perhaps more relevant than the discrepancy in computed distances between the two methods is the comparison of the closest point labels found by each, which is illustrated in figure 23. The greater relevance of labels is attributable to the distance value itself only being used in managing the narrow band (which, for a band of a half-width of 5 pixels, corresponds to a mean discrepancy of less than .35 voxels), whereas the label of the closest

surface point is used in the segmentation for computing the sign of the signed distance function and for extending the image information away from the zero-level isosurface, and in the registration with the range-based cortical surface later on, in the form of a Closest Point Map.

Moreover, it is geometrically feasible for a voxel to have 2 or more closest surface points, and our Fast Marching technique has an explicit check for this case, whereupon the segmentation uses the first label (one label is required everywhere in the narrow band for the image extension and for the sign computation) but the registration can dispense with ambiguous voxel-surface point matches. The voxels which possess two or more closest points are shown as black. In theory (with an analytical surface, or with voxels of infinite resolution), they should tend to the *medial surface*[11] of the volume circumscribed by the brain boundary (white), but other streaks appear, due to the digitization of the boundary. Those which are artefacts can be detected with a simple check on the distance between the positions of the 2 or more surface points matched with an ambiguous voxel.

2.5.4 Outer Brain Surface Identification Validation

The object of the final validation of the model is to quantitatively estimate the accuracy of the outer brain surface identification with synthetic data and to qualitatively demonstrate its performance on real data, involving both healthy ICBM subjects and MNI patient data involving pathologies.

The quantitative estimation of the model accuracy is currently being conducted. We use as a reference the distance map computed from the known brain boundary by the brute force method discussed in the preceding section. We then run the surface model on simulated MRI data featuring 5% noise, and from the final boundary, we can look up the distance to the underlying boundary and compute overall statistics about the proximity of the model to this boundary. Furthermore, the sensitivity of this performance to each of the 4 parameters which determine the WM/GM Distance Threshold can be ascertained, by varying each one about its “optimal” value. Each optimal value was established by trial and error with real subject data.

Further validation with real data uses the same WM/GM Distance Threshold parameters, throughout. We currently report results on healthy subject data from the ICBM database, which appear in figure 24 and are presently conducting tests on patient data from the MNI.

3 Range-sensing

The three-dimensional coordinates of the visible surfaces of the skin and evolving cortex are computed by a commercial range-sensor made by Vitana Corp. (Ottawa, Canada), which uses both laser-based *triangulation* and proprietary *defocusing* techniques to estimate range [60]. Laser-based range-sensing is the industrial standard for quickly and accurately acquiring dense, quantitative Our rangefinder is mounted on a commercial linear positioner,

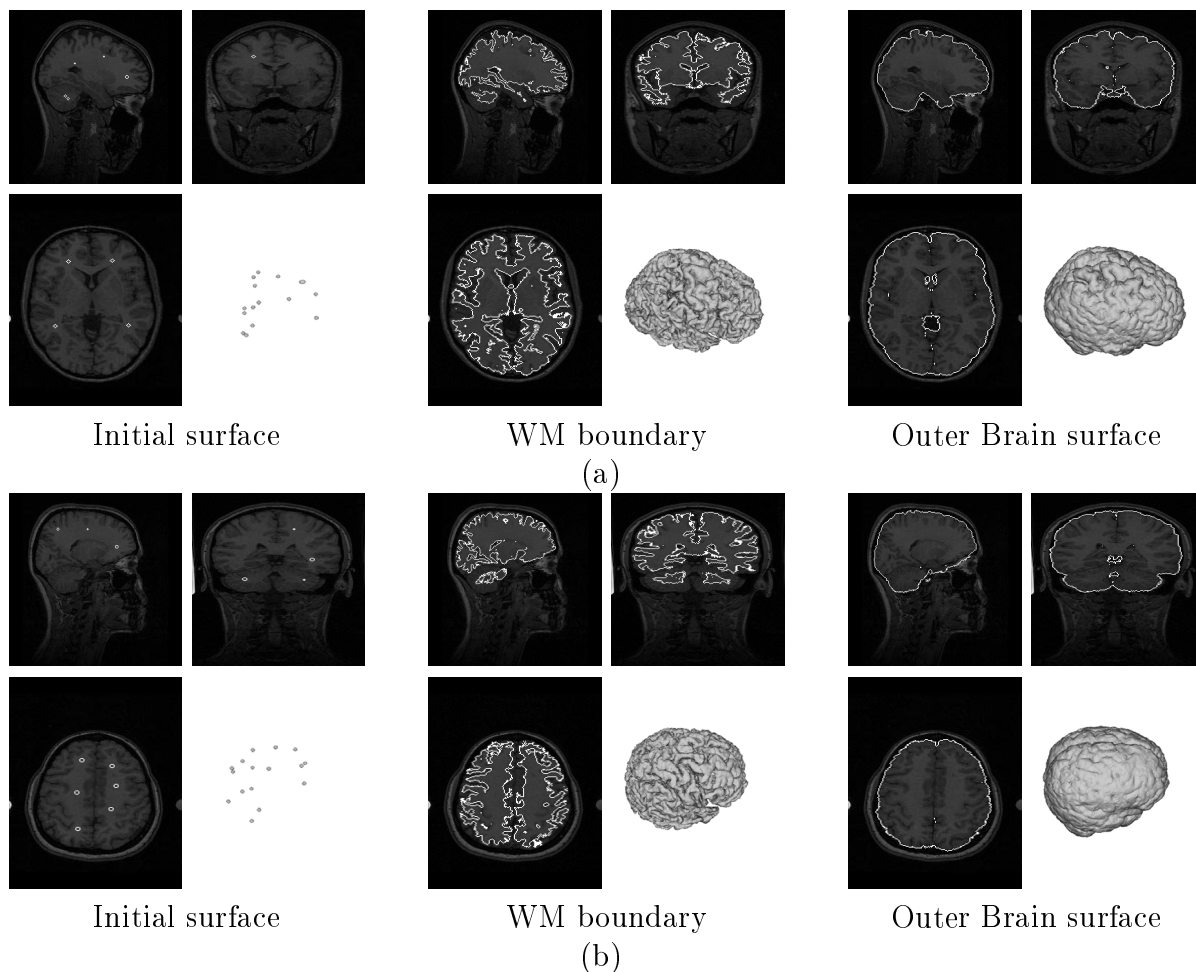


Figure 24: Results on real T1-weighted data: (a) ICBM subject 00100; (b) ICBM subject 00101.

made by Intelligent Actuator. Typical results produced by our sensor are shown in figure 25.

Laser-based triangulation involves projecting a thin ray of laser light onto an object, capturing the reflected light by a charge-coupled device (CCD), and inferring depth from the pixel position of the reflected point and from the geometric relationship between the position of the laser source and the CCD. The relationship which summarizes the triangulation is illustrated on figure 26(a) and expressed as [60]:

$$z = \frac{d l}{p + l \tan \theta}. \quad (12)$$

In equation (12), z is the range value of the scene point, d is the width separating the laser source from the CCD optical center, l is the distance from the lens to the image plane, θ is a fixed angle expressing the attitude of the laser source, and p is the pixel position of




Figure 25: Range image of the left hemisphere of a human brain.

the reflected point. The x -component of position is obtained by analogous triangulation. These relationships are characterized by a factory *calibration* procedure, after which each laser peak detected at pixel (i, j) on the CCD is interpreted as the (x, z) position of the point of the scene illuminated by the incident laser ray. Ordinarily, in order to get range information at regular samples of a surface, this ray would have to be swept along two axes normal to each other and spanning a plane which is aligned with a local plane tangent to the anatomical surface of interest. However, our sensor dispenses with one sweep axis by optically spreading the ray into a *laser plane*. The intersection of this plane with the imaged object instantaneously produces a 2D range-map of a “slice” of the surface, which appears as a curve of 256 points, or *laser profile*, on the CCD⁶. A 3D surface is then obtained by linearly sweeping the sensor, ideally in a direction normal to the laser plane, across the surface of interest. The coordinate y is given by the positioner, and the three coordinate axes form an orthogonal basis.

Defocusing is implemented by allowing the laser light to pass through two holes, at a predetermined distance d apart, rather than one. The CCD sees two laser profiles instead of one, and range is determined by measuring the space b (in each pixel column) between the two images b_1 and b_2 of a laser point, as illustrated in figure 26(b), according to the expression:

$$z = \frac{1}{o^{-1} + b/(dl)}, \quad (13)$$

where l and d are the lens-image plane and laser diode-optical center distances, as before. The

⁶Laser profiles are acquired at a 60 Hz rate: collecting a 256×256 range image, consisting of 256 profiles spanning the surface of interest, requires just over four seconds.

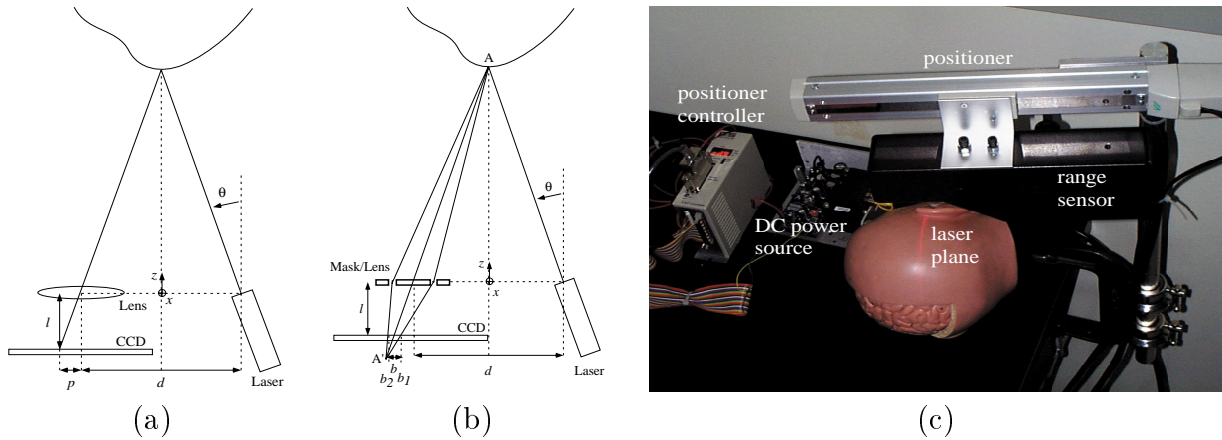


Figure 26: Range-sensing: (a) laser triangulation; (b) defocusing; (c) equipment

range measurements obtained by triangulation and by defocusing are combined by assigning weights to the two expressions for z given in expressions (12) and (13) [60].

Finally, the sensor has a depth of field of 135 mm (from 108 mm to 243 mm), which for a 256-pixel CCD corresponds to roughly .5 mm pixel accuracy. The field of view at a typical range is around 10-12 cm (depending on the proximity of the sensor to the surface), which gives it a .4-.5 mm pixel accuracy width-wise. The positioner which linearly sweeps the sensor across the skin or cortical surface possesses .02 mm repeatability and is capable of 150 mm travel at speeds up to 330 mm/sec. During this linear travel, our program acquires 256 laser profiles at a constant rate, when the positioner is known to move at a constant speed. The *equal spacing* of our samples, resulting from the CCD being regular (in the i - k pixel coordinates, not in x - z spatial coordinates) and from constant-rate acquisition along the y -axis, is exploited later on in the registration process. Both the sensor and the positioner are supplied with Windows-based software interfaces and dedicated control hardware. Both software interfaces support a standalone program calling vendor-supplied libraries. A picture of our range-sensing system appears in figure 26(c).

4 Probe-based Calibration and Sensor Base Tracking Procedure

The objective of the calibration and sensor base tracking procedure is to compute the rigid transformation which relates range image (inner) coordinates to MRI coordinates. This transformation can then be applied to the outer brain surface identified in the first stage, to relate it to range images of the exposed brain surface in the OR. An overview of the calibration geometry is illustrated in figure 27 (a). The goal of the calibration stage, which takes place outside of the OR, is to relate the inner range-sensor reference to its outer side-plates. The goal of the sensor base tracking stage is to establish the position of the outer side-plates at time t_n , $n \geq 0$ in the OR, with respect to the probe reference, and use this

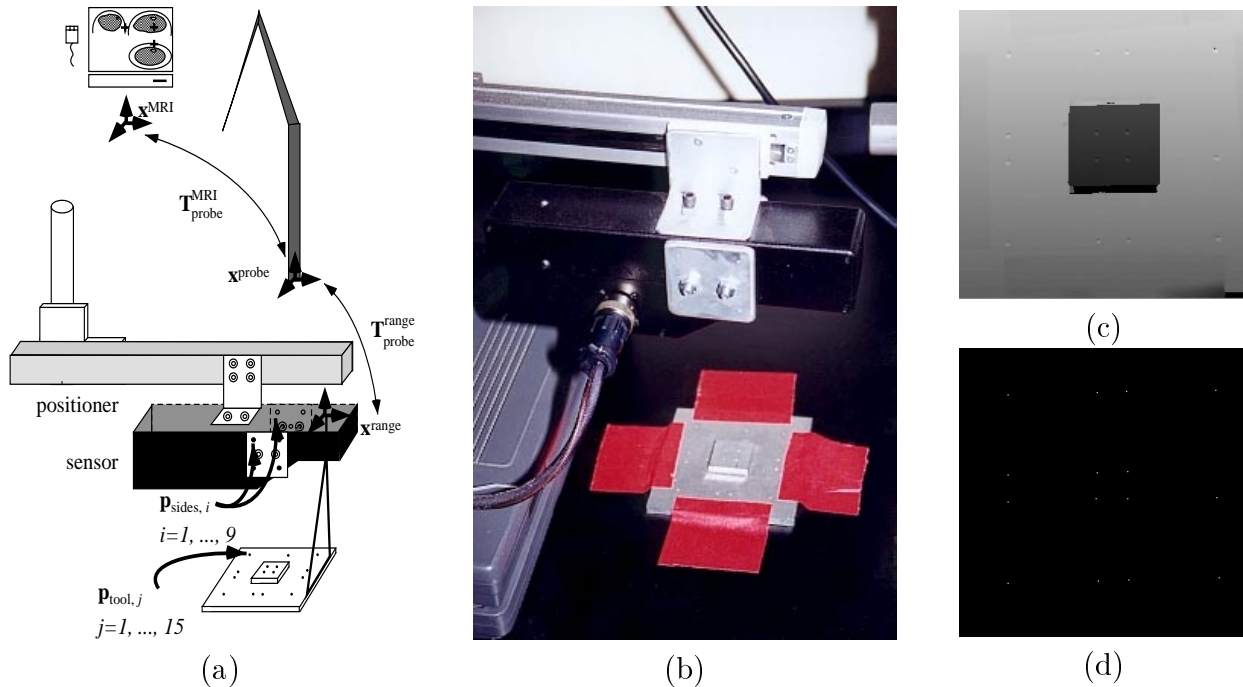


Figure 27: Calibration procedure: (a) overview of geometry; (b) picture of calibration tool and side-plate bolted to sensor; (c) range image of calibration tool displayed as grey level image (the greater the range value, the whiter the pixel); (d) detected range pits, after component analysis and pruning by distance-to-neighbour constraints.

information along with the calibration results to express the inner range-sensor reference with respect to the probe reference.

The calibration procedure (see figure 27(a) and (b)) consists of the following steps (assuming that the sensor is in stable configuration through the calibration procedure, and that the sensor is brought to *home position*, $y = 0$, after each range image acquisition, as it is currently programmed to do), which have been implemented but not thoroughly validated:

1. *Range-sensing of a wedding-cake (i.e.: non-coplanar) calibration tool* (see figure 27(c)), featuring an array of hemispherical milled holes, or *pits* of 2mm radius (corresponding to the size of the probe tip), producing an range image such as in figure 27 (b) (note one missing pit, to break the symmetry and allow us to ascertain the orientation of the tool within the range image).
2. *Automatic identification of hemispherical pits $p_{tool, j}^{range, t_{cal}}$* (see figure 27) *within range image*, assuming a hemispherical template of 2mm radius, by *template matching* [46]. We assume a z -reference which is estimated by taking the median of 8-neighbouring range points 8 pixels removed from the current coordinate. We then *threshold* the correlation values in order to retain only those which correlate highly with the template (using a histogram of these values to select the threshold), then perform a *connected*

component analysis [46] in order to establish the contiguous regions of high “pitness”. Within each such region, we find the point of maximal correlation value. We also eliminate false positives (near discontinuities within the tool) on the basis of known distances to neighbouring pits (using projections on 2D image plane: i.e.: neglecting the z -component of the range data).

3. *Probe-space identification of pits in the tool*, $\mathbf{p}_{\text{tool},j}^{\text{probe},t_{\text{cal}}}$, by insertion of the probe tip.
4. *Registration of range pits* $\mathbf{p}_{\text{tool},j}^{\text{range},t_{\text{cal}}}$ *with probe-space coordinates of the pits* $\mathbf{p}_{\text{tool},j}^{\text{probe},t_{\text{cal}}}$: the transformation from probe reference to internal range reference (for the calibration procedure) $\mathbf{T}_{\text{probe}}^{\text{range},t_{\text{cal}}}$ is determined from the point pairs by Arun’s SVD technique [2].
5. *Probe-space identification of pits in sensor side-plates*: $\mathbf{p}_{\text{sides},i}^{\text{probe},t_{\text{cal}}}$.
6. *Apply probe-range transformation to side-plate pits*:

$$\mathbf{p}_{\text{sides},i}^{\text{range}} = \mathbf{T}_{\text{probe}}^{\text{range},t_{\text{cal}}} \mathbf{p}_{\text{sides},i}^{\text{probe},t_{\text{cal}}} . \quad (14)$$

Note that the points $\mathbf{p}_{\text{sides},i}^{\text{range}}$ shouldn’t depend on time, since the side-plates are bolted to the sensor and the positioner carriage is always in home ($y = 0$) position when we track the side-plates with the probe.

The sensor base tracking procedure in the OR at t_n , $n \geq 0$ (initially and otherwise when sensor base tracking is required) is as follows:

1. *Probe-space identification of pits in sensor side-plates*, to characterize the new sensor base position: $\mathbf{p}_{\text{sides},i}^{\text{probe},t_n}$.
2. *Match probe-space pits with the points* $\mathbf{p}_{\text{sides},i}^{\text{range}}$ *to determine the probe reference-range transformation* $\mathbf{T}_{\text{range}}^{\text{probe},t_n}$ *for the configuration at time* t_n .
3. *The range data comprising surface* S_B *are then related to MRI-space anatomical surfaces* by considering the probe-MRI transformation computed in typical IGNS fashion, along with the probe reference-range transformation:

$$\mathbf{x}_{S_B,t_n}^{\text{MRI}} = \mathbf{T}_{\text{probe}}^{\text{MRI},t_n} \mathbf{T}_{\text{range}}^{\text{probe},t_n} \mathbf{x}_{S_B,t_n}^{\text{range}} . \quad (15)$$

In general, $\mathbf{T}_{\text{range}}^{\text{probe},t_n}$ can change over time (particularly if the sensor base is retracted out of the way of the surgeon in between range scans), whereas $\mathbf{T}_{\text{probe}}^{\text{MRI},t_n}$ will tend to be stable due to the use of a Mayfield clamp which generally keeps the patient’s head in place, although some slippage may occur and can be ascertained with a *tracker* [56] affixed to the patient’s head.

Validating both the sensor base tracking involves the skin phantom shown in figure 28. This is currently underway. An absolute measure of the error can be computed by automatically detecting the position of the glue-on landmarks in the range image, by template

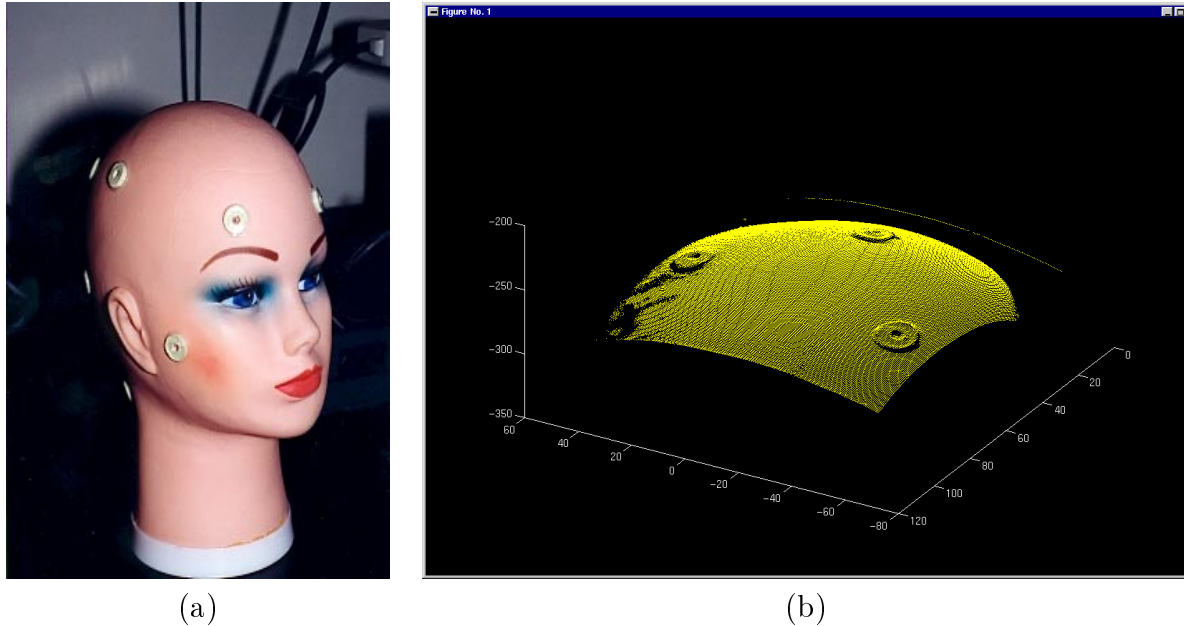


Figure 28: (a) “Poupoune”, the exotic skin phantom; (b) range image of the right side of the head of Poupoune.

matching, and comparing this position to that which is predicted by the position of the landmarks in MR space, as determined by the user, and transformed according to the rotation and translation computed by this procedure. The validation which is under way will superimpose on this range image the probe-based positions of the glue-on landmarks, as well as their image-space homologs, based on the probe-range and MRI-range transformations computed by the calibration and sensor base tracking procedure. A quantitative estimate of the error can then be obtained by extracting the centers of the landmarks on the range image, by template matching, and comparing them with the transformed probe- and MRI-based landmark centers.

5 Surface Displacement Estimation

5.1 Non-rigid Brain Motion Estimation

Our approach to estimating non-rigid brain deformation is by first establishing a baseline for this movement (i.e.: the preoperative outer brain surface computed in §2 rigidly transformed to range space by the sensor base tracking procedure), and measuring incremental non-rigid motion with an Iterative Closest Point (ICP) [8, 31] surface registration procedure. Such a registration procedure works as follows: starting from an initial (possibly rough) alignment between two surfaces S_A and S_B , iteratively uses the set of *closest point pairs* (as opposed to features of similar shape, for example) between the two surfaces to determine a new incremental transformation. The incremental transformation \mathbf{T}_k corresponding to the k th

iteration is then applied to one of the surfaces (which may already be transformed), namely $\mathbf{T}_{k-1}S_B$, gradually improving its alignment with S_A . The iterative process halts when the registration error between S_A and $\mathbf{T}_k S_B$ is less than some preset tolerance.

In general, ICP registration procedures make use of formal search algorithms to find the set of closest point pairs which determine the transformation \mathbf{T}_k , usually with some kind of preprocessing of the data which improves the efficiency of the search (for a survey see [4]). As emphasized in § 2, we make use of the Fast Marching technique to produce a *Closest Point Map* of the final outer brain surface in the MRI scan. This map allows us to do away with a search for closest points by simply *looking them up* and to produce explicit point pairs which can take advantage of closed-form transformation algorithms, thereby making each iteration of our ICP technique essentially *instantaneous*.

If the motion we are modelling were rigid, the iterative transformation \mathbf{T}_k would simply consist of a rotation \mathbf{R}_k and a translation \mathbf{t}_k which minimize the least-squares k th iteration registration error:

$$\min_{\mathbf{R}_k, \mathbf{t}_k} \sum_{i=1}^N \|\mathbf{x}_{A_i} - (\mathbf{R}_k \mathbf{x}_{B_i} + \mathbf{t}_k)\|^2 . \quad (16)$$

However, the motion we want to model is inherently non-rigid, so that the displacement must be expressed as a function with piecewise-local support, such as a spline. Furthermore, this motion can be expressed as a 2D surface spline, since it takes place within the domain of the range image, which is an open surface. Moreover, because the range image consists of equidistant data, the surface spline defined over its domain involves no re-sampling of irregular data. Lastly, we exploit the regularity of the range domain by using extremely efficient *recursive smoothing splines* [58, 59] to model the displacement function.

In practice, our ICP algorithm features a rigid motion stage which can compensate significant global tissue shift, followed by non-rigid motion characterization. We use Arun’s SVD technique[2] to compute the rigid transformation which estimates global shift. Our non-rigid characterization draws on the work of Unser et al [58, 59], which expresses the fitting of interpolating, smoothing or approximating splines as a sequence of filtering stages, all of which can be implemented as recursive filters, provided that the data are sampled at regular positions. We adopt a smoothing spline approach, whose first stage is a convolution with a smoothing kernel ⁷

$$S_\lambda^n(Z) = 1 / \left(B_1^n(Z) + \lambda(-Z + 2 - Z^{-1})^{\frac{n+1}{2}} \right) , \quad (17)$$

where $B_1^n(Z)$ is the *indirect B-spline transform* given by Unser [58], n is the order of the spline fitting and λ is the regularization parameter of the smoothing. This stage produces B-spline coefficients and is followed by a convolution with the indirect transform $B_1^n(Z)$ to yield the smoothed output. The filters can simply be cascaded to implement smoothing in 2D.

⁷Note: Z here relates to the Z -transform, not to be confused with the depth axis of the range-sensor.

We adopt a first-order fit to limit overshoots due to a higher-order continuity constraint, near the boundary of the craneotomy, outside of which the displacement should be null after rigid registration. These operations in Z -space translate to local operations in 1D involving *immediately neighbouring* range image points, independently of the value of the smoothing parameter λ . Just prior to each spline fitting operation, in order to define a displacement vector at image coordinates where range information is lacking, due to possible occlusion or low signal value, an average of valid neighbouring displacement vectors, weighted according to the proximity of the corresponding range point, is propagated and iteratively refined. Lastly, in order to make the non-rigid motion estimation well-behaved, the smoothing parameter λ is initially set very high, and is gradually lowered as the process iterates, progressively resolving finer-level motion detail. Just as for the rigid ICP stage, each iteration of the non-rigid surface motion estimation is essentially instantaneous.

The justification for emphasizing computational efficiency here is two-pronged: *clinical acceptability* and the inherent *temporal underdetermination* in estimating non-rigid motion. Clinical acceptability hinges on allowing the surgeon to easily deploy the system, get the surface displacement quickly, use this surface information with efficient FEM schemes [3, 20] to estimate volumetric displacement, and proceed with the surgery. The issue of underdeterminacy implies that it makes sense to enable the acquisition of a large number of displacement maps over time if needed, just as it does to use dense point displacement information to characterize non-rigid motion between any two successive instants.

5.2 Validation of Surface Motion Estimation

Quantitative assessment of the accuracy of our non-rigid surface registration is based on synthesizing range images from subpatches of the known cortical boundary of the digital phantom, and applying an analytical displacement function to each subpatch. The results of our motion estimation based on the ICP technique can then be compared with the analytical displacement function. This study is currently under way. The analytic expression for surface motion is expressed in spherical coordinates. We are investigating

- a simple global flattening of the subpatch where the amplitude of the motion is linearly varying with ϕ and θ from the centre of the patch to its extremities, and
- a simple global flattening, as above, with a spatial sinusoid added to it.

This is applied to subpatches coinciding with the top, back, front, and left and right sides of the head, namely subpatches centered at $\phi = 0$, and at $\phi = \pi/4$ and $\theta = 0$, $\theta = \pi/2$, $\theta = \pi$, and $\theta = -\pi/2$.

A more physically realistic assessment is possible through a brain-shaped elastic phantom, which is described in appendix A. Non-rigid surface motion is the result of turning the set screws under the support plate of our elastic phantom, triggering a deformation of up to 15mm at the top. We can image the cortical surface with our range-sensor, while maintaining the support plate fixed with respect to the range-sensor/positioner reference (see figure




Figure 29: Typical non-rigid registration results: (a) MRI level-set surface rigidly registered to the range image of the deformed phantom, in range space (both decimated by a factor of 2); (b) MRI level-set surface non-rigidly registered to the range image of the deformed phantom (both decimated by a factor of 2); (c) slice through (a), at $y = 79.592$; (d) slice through (b), at $y = 79.592$; (e) evolution of RMS point-pair distance with ICP iteration.

30(b)). We then compare the results of a strictly rigid transformation computation with a sequence of rigid and non-rigid ICP registrations. Each option uses a rough manual initial alignment based on the corners of the range domain. This comparison is currently based on visual inspection of how the original MRI level-set surface, subject to the inverse of the range-to-MRI transformation, aligns with the range data of the deformed cortical surface, and on the two plots of the RMS distance between point-pairs against iteration number. The abrupt improvement of RMS distance on the second curve indicates the initiation of the non-rigid stage, with an earlier rigid stage termination than the rigid-body stage illustrated by the first curve. If the displacement vectors between the original (many-to-one) closest point pairs are used to determine the non-rigid spline-based transformation, we find that they distort the displacement map by inflating the components (essentially along x and y) normal to the shortest path at each point between the two surfaces. We obtain better results by

considering only the displacements of *mutually* closest (i.e.: one-to-one) pairs, which can be determined fairly easily from the original set of point pairs, and iteratively propagating this information everywhere else, and this approach is what is reflected in figure 29. However, we should again emphasize that the RMS error was computed for all closest point pairs, which are many-to-1, and that a measure based on mutually closest pairs would give a better measure of the quality of the registration.

6 Finite Element Modelling

The four preceding stages provide a dense nonrigid displacement map of the part of the outer cortex which is visible in the OR. Obviously, a surgeon is interested in what goes on below this surface as well. To this end, the last stage uses known displacements of the outer cortical surface to constrain a finite element model with realistic material properties. The displacement finite element method numerically solves for unknown displacements, deformations, stresses, forces and possibly other variables of a solid body. An exact solution would require *force* and *momentum equilibrium* at all times everywhere in the body, but the finite element method replaces this requirement with the weaker one that equilibrium must be maintained in an average sense over a finite number of divisions of the volume of the body [6, 23]. The actual division of complex geometries into simple elements, such as tetrahedra and hexahedra, corresponds to the *meshing* problem[43], and is an still active research area.

Our intent here is simply to provide proof of concept with our elastic brain phantom. The meshing issue should be resolved by public domain software, such as QMG[45], provided that a triangulated surface of the phantom can be input in QMG format, and that the rigid part of the meshing, namely the imbedded disk and rod and support plate, can be meshed interactively. The finite element modeling per se will use a simple viscoelastic ABAQUS model. We plan on using known properties of PVA-C as input to ABAQUS. Validation will be a comparison between the displacement predicted by ABAQUS, and the measured movement of imbedded glass beads, tracked over successive MRI scans (the rigid transformation between successive scans can be ascertained by glue-on landmarks on the support plate).

7 Summary

We've presented a method for estimating intrasurgical brain shift, featuring five stages: semi-automatic outer brain surface identification, range-sensing of the exposed brain surface in the OR, a procedure for the rigid registration of the range and MRI spaces, non-rigid cortical motion tracking, and realistic finite-element modelling of the patient's head.

The surface identification scheme is based on the level-set framework, and features two of our contributions. First, we integrate voxel classification results into the image-based speed function, providing it with higher-level anatomical information in flexible manner. We also stabilize the estimation of the outer brain surface, against possible bleeding effects

due to lack of T1 contrast between grey matter and nearby muscle, skin and sagittal sinus voxels. This is done by sweeping the model outward in two stages, one which finds the white matter boundary, and another which pushes the model through grey matter voxels deemed close to the previous boundary, to capture the outer brain surface. Second, we feature a postprocessing stage where a Closest Point Map is computed for the outer brain surfaces, which subsequently makes each iteration of the Iterative Closest Point registration essentially instantaneous. Validation makes use of a digital anthropomorphic phantom and of a MRI simulator, both developed at the MNI.

The range-sensing stage also represents a contribution of ours, in that it is the first application of laser-based range-sensing to brain shift estimation. The rigid transformation between range and MRI coordinates is handled by a calibration and sensor based tracking procedure. Validation of the procedure is under way, and makes use of a skin phantom with glue-on landmarks.

The non-rigid surface registration is an Iterative Closest Point technique. Our contribution here is in using a Closest Point Map and recursive splines to efficiently find closest point pairs and to compute the transformation determined by these pairs. A qualitative validation, using an elastic brain phantom, is featured here. A quantitative study, based on synthesizing a range image from the known cortical boundary of the digital phantom and warping it analytically, is currently underway.


A simple proof of concept of the finite-element model stage is planned, with the elastic brain phantom, and will involve existing public domain meshing software, as well as commercial FEM software. These results will be compared against displacements determined from a series of MRI scans of the phantom, to track the movement of imbedded glass beads.

A Non-Clinical Validation: Elastic Brain Phantom


For the purpose of reproducing non-rigid cortical movement outside the OR, we have implemented a brain-shaped phantom with elastic material properties. We use a jello mold in the shape of the brain (obtained from the Brainstorms website [10]), into which we pour *PolyVinyl Alcohol Cryogel* (PVA-C), as illustrated in figure 30(a). The latter is a relatively *viscous liquid*, which upon freezing and thawing, becomes an *elastic solid* [15] as shown in figure 30(b). Furthermore, PVA-C can sustain several freeze-thaw cycles to acquire more rigidity, and the PVA concentration can also be manipulated to that effect, producing a Young’s modulus E in the .1 to .7 MPa spectrum [52]. This is comparable to values published in the literature for gray and white matter material properties.⁸

This phantom rests on a support plate of plexiglass (usable within an MR scanner), and features a moving assembly consisting of a small disk and rod, also of plexiglass, imbedded

⁸Zhou [64], for example, lists bulk moduli K of .219 GPa and .349 GPa, for gray and white matter, as well as a Poisson’s ratio ν of .4996 for each. These values translate to values of E of .525 MPa and .837 MPa respectively, based on the formula $K = E/[3(1 - 2\nu)]$ for isotropically elastic material [37]. The values found elsewhere in the literature are closer to the lower end of the PVA-C spectrum [40].



(a)



(b)

Figure 30: Elastic PVA-C brain phantom: (a) liquid stage (b) final result after freezing and thawing.

within the elastic material, as well as some glass beads used to validate the surface displacement/FEM approach to volumetric motion estimation. The other end of the rod (outside the elastic material) is attached to an external plate which rests on the opposite side of the support plate. The position of the moving assembly can be modified by means of some set screws pushing the external plate away from the support plate, thus drawing the embedded disk toward the support plate and triggering a compression of the elastic material.

References

- [1] D. Adelsteinsson & J.A. Sethian, A Fast Level Set Method for Propagating Interfaces, *J. Computational Physics*, No. 118, pp. 269-277, 1995.
- [2] K.S. Arun, T.S. Huang & S.D. Blostein, Least-squares Fitting of Two 3-D Point Sets, *IEEE Trans. Pattern Analysis & Machine Intelligence*, Vol. 9, No. 5, pp. 698-700, May 1987.
- [3] O. Astley & V. Hayward, Real-Time Finite-Elements Simulation of General Visco-elastic Materials for Haptic Presentation, *IROS '97, IEEE/RJS Int. Conference on Intelligent Robots and Systems*, Sept. 1997.
- [4] M.A. Audette, F.P. Ferrie & T.M. Peters, An Algorithmic Overview of Surface Registration Techniques for Medical Imaging, *Medical Image Analysis*, Vol. 4, No. 3, pp. 201-217, Sept. 2000.
- [5] M.A. Audette, K. Siddiqi, & T.M. Peters, Level-set Surface Segmentation and Fast Cortical Range Image Tracking for Computing Intrasurgical Deformations, *Med. Image Comput. & Comp.-Ass. Interv. (MICCAI99)*, Cambridge, England, pp. 788-797, 1999.

- [6] K.-J. Bathe, *Finite Element Procedures in Engineering Analysis*, Prentice-Hall, 1982.
- [7] A.M. Bensaid et al., Partially Supervised Clustering for Image Segmentation, *Pattern Recognition*, Vol. 29, No. 5, pp. 859-871, 1996.
- [8] P.J. Besl & N.D. McKay, A Method for Registration of 3-D Shapes, *IEEE Trans. Patt. Anal. & Mach. Intell.*, Vol.14, No.2, pp. 239-256, 1992.
- [9] P.A. Bottomley et al., A Review of ^1H Nuclear Magnetic Resonance Relaxation in Pathology: Are T_1 and T_2 Diagnostic?, *Med. Phys.*, Vol. 14, No. 1, pp. 1-37, Jan/Feb. 1987.
- [10] Brainstorms, www.brainstorms.com/store/showitem.asp?k=pr148&f=0, 1999.
- [11] S. Bouix & K. Siddiqi, Divergence-Based Medial Surfaces, *Proc. Euro. Conf. Comp. Vision*, 2000.
- [12] M.B. Carpenter, *Core Text of Neuroanatomy*, 4th ed., Williams & Wilkins, Hong Kong, 1996.
- [13] V. Caselles, R. Kimmel & G. Sapiro, Geodesic Active Contours, *Proc. IEEE Int. Conf. Computer Vision*, pp. 694-699, 1995.
- [14] V. Caselles et al., Minimal Surfaces: a Three Dimensional Segmentation Approach, *Technion EE Pub. No. 973*, 1995.
- [15] K.C. Chu & B.K. Rutt, Polyvinyl Alcohol Cryogel: An Ideal Phantom Material for MR Studies of Arterial Flow and Elasticity, *Magnetic Resonance in Medicine*, No. 37, pp. 314-319, 1997.
- [16] L.P. Clarke et al., MRI Segmentation: Methods and Applications, *Magnetic Resonance Imaging*, Vol. 13, No. 3, pp. 343-368, 1995.
- [17] D.L. Collins et al., Design and Construction of a Realistic Digital Brain Phantom, *IEEE Trans. Medical Imaging*, Vol. 17, No. 3, pp. 463-468, June 1998.
- [18] D.L. Collins, personal communication.
- [19] R.M. Comeau et al., Intraoperative Ultrasound for Guidance and Tissue Shift Correction in Image-guided Neurosurgery, *Medical Physics*, Vol. 27, No. 4, April 2000.
- [20] S. Cotin, H. Delingette & N. Ayache, Real-Time Elastic Deformations of Soft Tissues for Surgery Simulation, *IEEE Trans. Visualization & Computer Graphics*, No. 5, pp. 62-73, Jan.-Mar. 1999.
- [21] H. Delingette & J. Montagnat, New Algorithms for Controlling Active Contours Shape and Topology, European Conference on Computer Vision (ECCV'00), to appear, June 2000.

- [22] K. Fukunaga, *Introduction to Statistical Pattern Recognition*, 2nd ed., Academic Press, 1990.
- [23] Hibbitt, Karlsson & Sorensen, Inc., *ABAQUS Theory Manual*, 1995.
- [24] ISG Technologies Inc., *Viewing Wand Image-Guided Surgical System- Operator's Guide*, 1997.
- [25] B.B. Kimia & K. Siddiqi, Geometric Heat Equation and Nonlinear Diffusion of Shapes and Images, *Computer Vision and Image Understanding*, Vol. 64, No. 3, pp. 305-322, Nov. 1996.
- [26] R. Kimmel, N. Kiryati, and A.M. Bruckstein. Sub-pixel Distance Maps and Weighted Distance Transforms, *J. Mathematical Imaging and Vision*, No. 6, pp. 223-233, 1996.
- [27] J.J. Koenderink, The Structure of Images, *Biological Cybernetics*, pp. 363-370, 1984.
- [28] V. Kollokian, *Performance Analysis of Automatic Techniques for Tissue Classification in Magnetic Resonance Images of the Human Brain*, Master's Thesis, Dept. Computer Science, Concordia University, Nov. 1996.
- [29] R.K.-S. Kwan, A.C. Evans & G.B. Pike, An Extensible MRI Simulator for Post-Processing Evaluation, *Proc. Visualization in Biomedical Computing, 4th Int. Conf. VBC'96*, pp. 135-140, Springer-Verlag, 1996.
- [30] R.K.-S. Kwan, A.C. Evans & G.B. Pike, MRI Simulation-Based Evaluation of Image-Processing and Classification Methods, *IEEE Trans. Medical Image Analysis*, Vol. 18, No. 11, Nov. 1999.
- [31] S. Lavallée & R. Szeliski, Recovering the Position and Orientation of Free-form Objects from Image Contours Using 3D Distance Maps, *IEEE Trans. Patt. Anal. & Mach. Intell.*, Vol. 17, no. 4, pp. 378-390, 1995.
- [32] W.E. Lorensen & H.E. Cline, Marching Cubes: A High Resolution 3D Surface Construction Algorithm, *Computer Graphics*, Vol. 21, No. 3, pp. 163-169, July 1987.
- [33] D. MacDonald, *Display - Program for Display and Segmentation of Surfaces and Volumes*, McConnell Brain Imaging Center document, 1995.
- [34] D. MacDonald, D. Avis & A.C. Evans, Proximity Constraints in Deformable Models for Cortical Surface Identification, *Med. Image Comput. & Comp.-Ass. Interv.- MICCAI'98*, pp. 650-659, 1998.
- [35] R. Malladi, J.A. Sethian & B.C Vemuri, Shape Modeling with Front Propagation: A Level Set Approach, *IEEE Trans. Patt. Anal. & Mach. Intell.*, Vol. 17, No. 2, Feb. 1995.

- [36] R. Malladi & J.A. Sethian, A Real-Time Algorithm for Medical Shape Recovery, *Proc. IEEE Int'l. Conf. Computer Vision*, pp. 304-310, 1998.
- [37] L.E. Malvern, *Introduction to the Mechanics of a Continuous Medium*, Prentice-Hall, 1969.
- [38] C.R. Maurer et al., Measurement of Intraoperative Brain Surface Deformation Under a Craniotomy, *Med. Image Comput. & Comp.-Ass. Interv. - MICCAI '98*, 1998.
- [39] T. McInerney & D. Terzopoulos, Topology Adaptive Deformable Surfaces for Medical Image Volume Segmentation, *IEEE Trans. Medical Image Analysis*, Vol. 18, No. 10, pp. 840-850, Oct. 1999.
- [40] M. I. Miga et al., A 3D Brain Deformation Model Experiencing Comparable Surgical Loads, *Proc. 19th Int. Conf. IEEE EMBS*, pp. 773-776, 1997.
- [41] K. Miller, Constitutive Model of Brain Tissue Suitable for Finite Element Analysis of Surgical Procedures, *J. Biomechanics*, No. 32, pp. 531-537.
- [42] V.S. Nalwa & T.O. Binford, On Detecting Edges, *IEEE Trans. Patt. Anal. & Mach. Intell.*, Vol. PAMI-8, No. 6, pp. 699-711, Nov. 1986.
- [43] S. Owen, A Survey of Unstructured Mesh Generation Technology, available online at www.andrew.cmu.edu/user/sowen/survey/index.html, 1999.
- [44] A. Pentland & S. Sclaroff, Closed-Form Solutions for Physically Based Shape Modelling and Recognition, *IEEE Trans. Patt. Anal. & Mach. Intell.*, Vol. 13, No. 7, pp. 715-729, July 1991.
- [45] QMG Mesh Generation, <http://www.cs.cornell.edu/home/vavasis/qmg-home.html>, 2000.
- [46] A. Rosenfeld & A.C. Kak, *Digital Picture Processing*, 2nd ed., Academic Press, 1982.
- [47] E. Rouy & A. Tourin, A Viscosity Solutions Approach to Shape-From-Shading, *SIAM J. Numerical Analysis*, Vol. 29, No. 3, pp. 867-884, June 1992.
- [48] R. Sedgewick, *Algorithms in C*, 3rd ed., Addison-Wesley, 1998.
- [49] J.A. Sethian, *Level Set Methods and Fast Marching Methods: Evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*, 2nd ed., Cambridge University Press, 1999.
- [50] O. Škrinjar, D. Spencer & J. Duncan, Brain Shift Modeling for Use in Neurosurgery, *Med. Image Comput. & Comp.-Ass. Interv.- MICCAI'98*, pp. 641-649, 1998.

- [51] J.G. Sled, A.P. Zijdenbos & A.C. Evans, A Nonparametric Method for Automatic Correction of Intensity Nonuniformity in MRI Data, *IEEE Trans. Medical Imaging*, Vol. 17, No. 1, pp. 87-97, Feb. 1998.
- [52] K. Surry, www.irus.rrri.on.ca/~kath, 1999.
- [53] R. Szeliski & S. Lavallee, Matching 3-D Anatomical Surfaces with Non-Rigid Deformations using Octree-Splines, *Int. J. Computer Vision*, Vol. 18, No. 2, pp. 171-186, 1996.
- [54] J. Talairach, G. Szikla & P. Tournoux et al., *Atlas d'anatomie stéréotaxique du télencéphale*, Masson, Paris, 1967.
- [55] D. Terzopoulos, & D. Metaxas, Dynamic 3D Models with Local and Global Deformations: Deformable Superquadrics, *IEEE Trans. Patt. Anal. & Mach. Intell.*, Vol. 13, No. 7, 703-714, 1991.
- [56] Traxtal Technologies, <http://www.traxtal.com>.
- [57] ICBM, <http://nessus.loni.ucla.edu/icbm/index0.html>.
- [58] M. Unser, A. Aldroubi, & M. Eden, B-Spline Signal Processing: Part I- Theory, *IEEE Trans. Signal Proc.*, Vol. 41, No. 2, pp 821-833, Feb. 1993.
- [59] M. Unser, A. Aldroubi, & M. Eden, B-Spline Signal Processing: Part II- Efficient Design and Applications, *IEEE Trans. Signal Proc.*, Vol. 41, No. 2, pp. 834-848, Feb. 1993.
- [60] Vitana Corporation, *ShapeGrabber Reference Manual*, 1997.
- [61] A. Yezzi Jr. et al., A Geometric Snake Model for Segmentation of Medical Imagery, *IEEE Trans. Medical Imaging*, Vol. 16, No. 2, pp. 199-209, Apr. 1997.
- [62] X. Zeng et al., Segmentation and Measurements of the Cortex from 3D MR Images, *Med. Image Comput. & Comp.-Ass. Interv.- MICCAI'98*, pp. 519-530, 1998.
- [63] Z. Zhang, Iterative Point Matching for Registration of Free-Form Curves and Surfaces, *Int. J. Comp. Vision*, Vol. 13, No. 2, pp. 119-152, 1994.
- [64] C. Zhou, T.B. Khalil, & A.I. King, A New Model Comparing Impact Responses of the Homogeneous and Inhomogeneous Human Brain, *Society of Automotive Engineers, Inc. report #952714*, 1995.
- [65] A. Zijdenbos, R. Forghani & A. Evans, Automatic Quantification of MS Lesions in 3D MRI Brain Data Sets: Validation of INSECT, *Med. Image Comput. & Comp.-Ass. Interv.- MICCAI'98*, pp. 439-447, 1998.