

# Detecting Cortical Surface Regions in Structural MR Data

Biswajit Bose<sup>1,†</sup>, John Fisher<sup>1</sup>, Bruce Fischl<sup>1,2</sup>, Oliver Hinds<sup>3</sup> and Eric Grimson<sup>1</sup>

<sup>1</sup>MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA 02139, USA

<sup>2</sup>Massachusetts General Hospital, Harvard Medical School, Charlestown, MA 02129, USA

<sup>3</sup>MIT Department of Brain and Cognitive Sciences, Cambridge, MA 02139, USA

<sup>†</sup>cielbleu@csail.mit.edu

## Abstract

*We present a novel level-set method for evolving open surfaces embedded in three-dimensional volumes. We adapt the method for statistical detection and segmentation of cytoarchitectonic regions of the cortical ribbon (e.g., Brodmann areas). In addition, we incorporate an explicit interface appearance model which is oriented normal to the open surface, allowing one to model characteristics beyond voxel intensities and high gradients. We show that such models are well suited to detecting embedded cortical structures. Appearance models of the interface are used in two ways: firstly, to evolve an open surface in the normal direction for the purpose of detecting the location of the surface, and secondly, to evolve the boundary of the surface in a direction tangential to the surface in order to delineate the extent of a specific Brodmann area within the cortical ribbon. The utility of the method is demonstrated on a challenging ex-vivo structural MR dataset for detection of Brodmann area 17.*

## 1. Introduction

Detection and analysis of the structure of the human cerebral cortex is an important problem with diverse applications in medical diagnosis. For this purpose, the cortex is often represented as a 2D manifold in 3D space. While much work has focussed on modelling the shape of this manifold (by detecting sulci and gyri, for example), there has been some work on mapping the cortex using cytoarchitectonic features, *i.e.*, appearance features that correlate with functionally distinct regions of the cortex. This dates back to the seminal work by Brodmann [4], who manually identified and mapped out 52 Brodmann areas (BAs) from histological samples. Most work on cortical-region detection has been done on 2D slices of the cerebral cortex using histological techniques (*e.g.* [4, 12]), not directly in 3D from MR imaging data.

In this paper, we propose a 3D open-surface evolution method for detecting and segmenting a BA on the cortical surface from MR data. Since the Brodmann area lies on the 2D cortical surface, and has finite extent, we formulate the problem as the general one of evolving an open surface in both the normal and tangential directions in 3D. The open surface represents the Brodmann area; normal evolution is used to detect the location of the cortical surface (in 3D), while tangential evolution of the boundary of the open surface is used to segment the BA on the cortical surface (while keeping the location of the cortical surface in 3D fixed). Our method exploits the characteristic myeloarchitecture of the BA, as visible in a high-resolution MR volume. Figure 1 shows a slice through structural MR data of the primary visual cortex, where BA 17 is characterised by a dark band of myelinated fibres called the *stria of Gennari*.

While the problem of evolving closed surfaces has been studied in great detail [11, 13], there has been comparatively less work on evolving open surfaces [14]. Level-set methods are commonly used for detecting closed surfaces, because of their numerical stability and ability to handle changing topology. While closed surfaces can be represented implicitly as level-sets of higher dimensional functions, this is not the case for open surfaces whose boundary lies within the domain of interest (see Figure 2). Consequently, some care must be taken in formulating an extension of the level-set methodology to open surfaces.

We tackle the 3D open-surface detection problem as follows. We start with an initial open surface, and close it to get an initial closed surface. Following previous work [14], we then maintain two level-set functions: a primary function to evolve this closed surface, and an auxiliary function to (implicitly) keep track of the parts of the closed surface that are actually relevant, *i.e.*, that correspond to the open surface we wish to detect. Our key idea is to ensure that these two functions represent orthogonal surfaces, and to alternately perform two surface evolution tasks: normal evolution and tangential evolution. The first evolution task (normal evolution) involves moving each point on the




Figure 1. 2D slices through contrast-inverted structural MR data of the human cerebral cortex. The stria of Gennari, indicating Brodmann area 17, is marked in orange. A non-stria cortical region surrounding BA 17 is marked in green. Zoomed-in parts of the stria and non-stria are also shown: note the distinctive dark band of the stria.




Figure 2. Three types of curves: closed curves, open curves whose ends lie on the boundary of the domain, and open curves whose ends lie completely within the domain. Standard level-set methods are applicable to the first two cases. Though this is a 2D example, similar properties hold for surfaces in 3D.

closed surface in the local normal direction. For the second task, we show that evolving the auxiliary level-set function in its local normal direction leads to tangential evolution of the boundary of the open surface (on the closed surface). By alternating normal and tangential evolution, we are able to represent and detect open surfaces. The novelty of our method lies in the coupling of these two types of evolution. Previous authors have individually considered either only normal [14] or only tangential [7] evolution of open surfaces; the former presumes the boundary of the open surface is known, while the latter presumes the surface itself has been extracted.

In a surface evolution task, the motion of each point on the surface is determined by a force field. The force field typically consists of two terms: a data-fidelity term (or *likelihood*), which may be calculated from a statistical model of

surface appearance, and a regularisation term that smooths the surface. Two types of commonly used data terms are edge-based terms (*e.g.*, sum of magnitudes of image gradients calculated at the interface pixels) and region-based terms (*e.g.*, difference of mean intensities or statistical properties inside and outside the interface). Region-based terms are often preferred because they are more robust to noise, but they are defined only for closed surfaces, not for open surfaces. We propose a structured likelihood model defined on normal intensity profiles, that measures data fidelity not just at the interface, but also for a band surrounding the interface. We show that a model of profile vectors is useful for both normal and tangential evolution of the open surface. These normal profile vectors are ideally suited for capturing the laminar structure around the *stria* in BA 17, and distinguishing it from the surrounding cortical region.

Our main contributions are two-fold:

1. We present a novel level-set method for evolving open surfaces in both normal and tangential directions in 3D. This is in contrast to previously published methods for evolving open surfaces in only one of these two directions. We apply this method to detection of Brodmann area 17 in the cerebral cortex.
2. We propose a data-fidelity term for 3D surface evolution based on a statistical model of the appearance around the surface (and not just on the surface itself). This helps capture the layered structure of the cortex between the grey-white and pial surfaces.

## 2. Related Work

There are two broad classes of (closed) surface evolution methods: Lagrangian methods and Eulerian methods. Eulerian methods (such as level-set methods, *e.g.* [11, 13]), by implicitly representing the surface  $\mathcal{S}$  as the zero level-set of a higher dimensional function  $\Psi$ , are numerically stable, and can easily handle changing topology. For these reasons, they are often preferred over Lagrangian methods. However, Eulerian methods are not directly applicable when considering open surfaces, since (for example) the zero level-set of a 3D function is typically either a closed surface, or an open surface whose boundaries lie at the extremes of the domain (see Figure 2).

Open surfaces have been represented by considering the intersection of two closed surfaces (and hence two level-set functions) [3, 14]. Here, the open surface is a subset of the zero level-set of one level-set function, and the boundary of the open surface is indicated by the intersection of the two zero level-sets. We use a similar representation. In [14], normal evolution of an open surface is proposed for a stereo reconstruction problem. The tangential extent of the surface has been kept fixed, and it depends entirely

on the initialisation. Thus, once the 3D-shape of the open surface has been found, its boundary (and hence its extent) cannot be changed. Others have studied the evolution of curves on fixed manifolds [7]. In this case, the 3D-shape of the open surface—defined by the manifold—is fixed, but its boundary—defined by the evolving curve—can change. In [9], a method is suggested for evolving both surfaces and curves embedded on those surfaces; however, both the surface domain and appearance models have been simplified.

Most work on region detection in the brain has been done in 2D slices of the cerebral cortex, using histological techniques to characterise a region’s cytoarchitecture (e.g. [4, 12]). This is partly because the technology needed to get MR data at the required resolution for characterising the myeloarchitecture of the Brodmann areas was not available till very recently [1, 2]. Walters *et al.* [16] propose an automatic method for detecting the motion-sensitive region V5/MT+ by looking for the characteristic laminar structure in 2D slices through the cortex. We are not aware of any methods prior to ours that detect Brodmann areas directly in three dimensions.

### 3. Framework for Interface Evolution

We assume that an initial open surface  $S_o$ , with boundary  $C_o$ , is provided to us as a triangle-mesh that does not self-intersect. To convert this explicit representation to an implicit one, we first close the open mesh by adding extra vertices and edges, as described in Appendix A<sup>1</sup>. We call the resulting closed surface the *primary* surface,  $S_p$ . Only a part  $S_o \subset S_p$  of the closed surface is *relevant*, i.e., represents the open surface.  $C_o$  is the curve that forms the boundary between the relevant part  $S_o$  and the remaining, non-relevant, part of  $S_p$ . To implicitly represent  $C_o$ , we use another closed surface, called the *auxiliary* surface,  $S_a$ . This surface has the property that  $C_o = S_p \cap S_a$ . Further,  $S_p$  and  $S_a$  are orthogonal where they intersect. Figure 3 illustrates the relationships among these manifolds<sup>2</sup>.

The core of the surface evolution process consists of alternating two tasks:

1. Normal evolution: the primary surface is evolved, while the auxiliary surface is fixed.
2. Tangential evolution: the auxiliary surface is reconstructed and then evolved, while the primary surface is fixed.

We now discuss these two tasks in detail. Figure 4 gives a graphical summary of our alternating surface evolution process.

<sup>1</sup>We provide the details of our automatic mesh-closing algorithm for the sake of completeness. In practice, this is a detail, and there are other ways of finding an initial closed surface (including interactive initialisation)

<sup>2</sup>Note that this figure uses a 2D example, since it is not easy to draw 3D surfaces.




Figure 3. Relationship between primary and auxiliary surfaces. Note that this is a 2D example, so open surfaces in 3D are replaced by open curves in 2D, and closed curves in 3D are replaced by pairs of points in 2D. The first figure shows the open surface we wish to detect in red. The second figure shows the arbitrary closure of this open surface in black. The closed surface (red and black together) form the primary surface  $S_p$ . The third figure shows the auxiliary surface  $S_a$  in cyan. This implicitly defines the curve  $C_o$ , shown using orange dots, which forms the boundary of the open surface on the primary surface.

#### 3.1. Normal Evolution of Open Surface

Given the closed surface  $S_p$ , it is evolved in the normal direction using standard level-set techniques. We represent  $S_p$  as the zero level-set of a 3D function  $\Psi_p$ , which we call the primary level-set function. Following [11], evolution of  $S_p$  involves solving the PDE:

$$(\Psi_p)_t = \mathcal{F}^p \cdot \nabla \Psi_p, \quad (1)$$

where  $\mathcal{F}^p(x)$  is a force field that specifies the velocity with which point  $x$  on the primary surface moves during curve evolution.

The force field is typically a sum of two terms: a data-fidelity term (or *likelihood*, in case of statistical models),  $\mathcal{F}_l^p$ , and a regularisation (or *smoothness*) term,  $\mathcal{F}_s^p$ . Note that only velocity in the direction normal to the surface matters, since  $\nabla \Psi_p$  is oriented along the normal to the zero level-set.

We consider the likelihood term in more detail in Section 4, where we discuss the relative advantages of different likelihood models. The smoothness term in our force field is the 3D analogue of 2D curve-shortening flow [11]:

$$\mathcal{F}_s^p = -b\kappa \mathcal{N}^p, \quad (2)$$

where  $b > 0$  is the weight factor,  $\kappa$  is the mean curvature of the surface at this point and  $\mathcal{N}^p$  is the outward normal. We use two different values of the weight  $b$  for different parts of the surface; we elaborate on this in Section 3.3, after outlining the tangential evolution process.

In practice, we set  $\Psi_p$  to be the signed distance function of the closed surface and follow the standard approach of velocity extension of the force field to a narrow-band around the zero-level set [11] to perform the level-set evolution.




Figure 4. An illustration of one round of the alternating steps of normal and tangential evolution. (a) Initial open surface (in red) and target ground-truth open surface (in blue). (b) The initial open surface is closed, forming a primary surface,  $\mathcal{S}_p$ , that consists of a relevant (red) and a non-relevant (black) region. Profile vectors at points on the primary surface are shown in gray. (c) Normal evolution will cause the primary surface to move in its normal direction. The auxiliary surface,  $\mathcal{S}_a$ , shown in light-blue, is orthogonal to the primary surface. (d) As a result of normal evolution, the open surface moves closer to the target location. However, its extent remains unchanged, since the auxiliary surface has not evolved. (e) The auxiliary surface is reinitialised. Next, tangential evolution will cause the auxiliary surface to move in *its* normal direction. (f) As a result of tangential evolution, the extent of the relevant region is better aligned with the ground truth.

### 3.2. Tangential Evolution of Open Surface

The intersection of  $\mathcal{S}_p$  and  $\mathcal{S}_a$  is a closed curve that marks the boundary of the relevant region on  $\mathcal{S}_p$ . We do tangential evolution by keeping  $\mathcal{S}_p$  fixed and evolving  $\mathcal{S}_a$ . This is also done using a level-set method.

We represent  $\mathcal{S}_a$  as the zero level-set of a 3D auxiliary function  $\Psi_a$ .  $\Psi_a$  is negative in the relevant (striated) region of  $\mathcal{S}_p$ , and positive in the non-relevant (non-striated) region. We construct  $\mathcal{S}_a$  so that it is orthogonal to  $\mathcal{S}_p$  where they intersect. Please see Appendix B for details of how this construction is done. The orthogonality property ensures that normal evolution of the primary function  $\Psi_p$  will not change the boundary of the relevant region (to a first-order approximation). At the same time, *normal* evolution of the auxiliary function  $\Psi_a$  will induce *tangential* evolution of the boundary of  $\mathcal{S}_o$  on  $\mathcal{S}_p$ . As a consequence, we hold one surface fixed while performing level-set evolution of the other. This is in contrast to previous methods [3, 14], where both level-sets are evolved together to achieve normal evolution, and no tangential evolution is performed.

Similar to normal evolution, the force field  $\mathcal{F}^a$  for tangential evolution also consists of a sum of terms. There is a likelihood term,  $\mathcal{F}_l^a$ , and a smoothness term,  $\mathcal{F}_s^a$ . The likelihood term is based on a region-model, and is discussed in further detail in the next section. The smoothness term is the same as Equation 2, except that the normals are obtained from the auxiliary function,  $\Psi_a$ , instead of the primary function,  $\Psi_p$ . In addition, a third term—a separating force,  $\mathcal{F}_{sep}^a$ —is used to discourage the evolving auxiliary surface  $\mathcal{S}_a$  from causing extra intersections with the primary surface  $\mathcal{S}_p$ . At each point on  $\mathcal{S}_a$ , the magnitude of the separating force is inversely proportional to the distance from the nearest point on  $\mathcal{S}_p$ , and its direction points away from this point.<sup>3</sup>

<sup>3</sup>The separating force term has no effect where the primary and auxiliary surfaces intersect, since it is in a direction tangential to the evolving auxiliary surface. This is important, since we do not want the separating

### 3.3. Interaction between Normal and Tangential Evolution

So far from our description, it might seem that normal evolution and tangential evolution are independent of each other. However, there is a subtlety that makes them dependent. In general, we want a strong smoothness term on the primary surface, to prevent it from becoming too jagged. The standard smoothness term is a curve length penalty, which tends to shrink the primary surface. Near the boundary of the open surface, though, we wish to expand the primary surface instead of shrinking it, so that the relevant region can grow if necessary. To solve this problem, we use a much weaker smoothness term in the vicinity of the boundary  $C_o$ . In other words, the weight  $b$  in Equation 2 is small near  $C_o$ , and larger elsewhere. As  $C_o$  changes due to tangential evolution, the force field for normal evolution changes accordingly.

## 4. Likelihood Model for Surface Appearance

We now discuss our statistical model for appearance of the open surface. This gives us a likelihood, that we use in the data-fidelity terms for normal and tangential evolution.

Two commonly used data terms for curve/surface-evolution are edge-based [5] and region-based [6] terms. Region-based terms are more robust to noise and incorrect initialisation. However, while edge-based terms can be defined for both open and closed surfaces, region-based terms are only defined for closed surfaces.

There is a direct connection between a likelihood model of observed appearance in a statistical framework, and the corresponding data-fidelity energy in a variational optimisation framework (*e.g.* [10]). This allows the development of statistical likelihood models for surface evolution. Following this approach, we propose likelihood models for *intensity profiles* along the normal direction at points cen-

force to affect the location of the evolving boundary.



tred on the open surface (as illustrated by the gray lines in Figure 4b). Such profile-based likelihood models can be thought of as a generalisation of edge-based models to include regions near the edge. Each observed normal profile vector,  $\mathbf{x}_o$ , is normalised in intensity<sup>4</sup> by first subtracting the component-wise mean from each component and then scaling the component intensities of the profile, to obtain a unit vector,  $\mathbf{x}$ . This helps reduce the effects of the varying bias field in the MR volume. We model these normalised profiles as being generated from a low-dimensional multivariate Gaussian vector  $\alpha$  plus additive Gaussian noise  $\mathbf{n}$ :

$$\mathbf{x} = \Phi\alpha + \mathbf{n}, \quad (3)$$

where  $\Phi$  is an orthogonal matrix of basis vectors. The parameters of this model are learnt from a training set of profile vectors via probabilistic principal component analysis (PPCA) [15].

Separate PPCA models,  $M_1$  and  $M_2$ , are learnt for profiles from the relevant (stria) and non-relevant (non-stria) regions respectively. An example set of normalised profile samples from stria and non-stria regions, and their corresponding PPCA likelihood models—mean and basis vectors—are shown in Figure 5. Note that the typical profile obtained from either region of the (contrast-inverted) images varies from high intensity near the pial surface to low intensity near the grey-white interface. This characteristic shape of the profile is used for localisation of the surface during normal evolution. The stria model shows an additional dip in intensity in the middle, due to the dark band formed along the cortical surface in BA 17 by the myelinated stria fibres. It is this difference between the models that allows us to discriminate between the two regions during tangential evolution.

We thus have two probability distributions,  $p(\mathbf{x}|M_1)$  and  $p(\mathbf{x}|M_2)$ , over profile vectors, corresponding to the stria and non-stria regions. For normal evolution, we calculate a mixture distribution assuming equal prior probability of these two models. We use the gradient of the log-likelihood under this mixture distribution as the data-term  $\mathcal{F}_l^p$  for evolving the primary surface.

For tangential evolution, we use the two likelihood models  $M_1$  and  $M_2$  in a region-based evolution framework. We assume that the voxels in the true underlying regions are i.i.d. samples from the corresponding relevant-region or non-relevant-region model. This leads to a 2-class classification problem on the primary surface, in which the evolving curve  $C_o$  defines the boundary between the classes. To minimise classification error, each voxel should be assigned to the model  $k$  which has higher posterior probability:  $p(\mathbf{x}|M_k)$ . This posterior can further be written as a product of the likelihood and a class prior. Hence the energy

<sup>4</sup>Note that the dimensionality (*i.e.*, number of components) of the profile vector is not changed during normalisation.

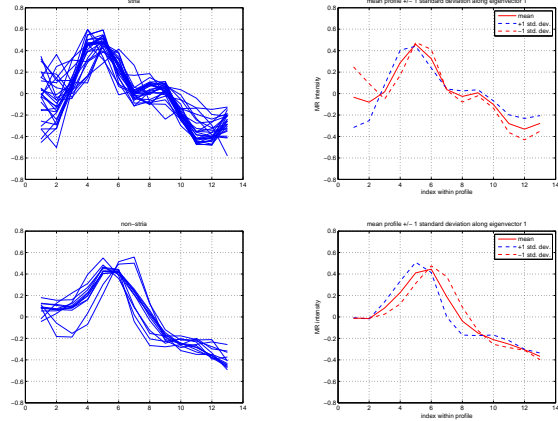


Figure 5. Sample normal profile vectors, and PPCA model (mean vector  $\pm$  first basis vector) for stria (top 2 images) and non-stria (bottom 2 images) regions. Each profile vector is of length 13, *i.e.*, comprising 6 voxels on either side of the surface.

function  $E_l^a$  to be minimised with respect to the curve  $C_o$  is the sum of log-likelihood and log-prior differences over the two regions,  $\mathcal{S}_o$  and  $\mathcal{S}_p - \mathcal{S}_o$ :

$$E_l^a = \sum_{\mathcal{S}_o} \log \frac{p(\mathbf{x}|M_1)p(M_1)}{p(\mathbf{x}|M_2)p(M_2)} + \sum_{\mathcal{S}_p - \mathcal{S}_o} \log \frac{p(\mathbf{x}|M_2)p(M_2)}{p(\mathbf{x}|M_1)p(M_1)}. \quad (4)$$

By calculus of variations, one finds that the force field that minimises this energy is given by a log-posterior-ratio:

$$\mathcal{F}_l^a = \log \frac{p(\mathbf{x}|M_1)p(M_1)}{p(\mathbf{x}|M_2)p(M_2)} \mathcal{N}^a. \quad (5)$$

This force is only defined at points along the curve,  $C_o$ , on the auxiliary surface  $\mathcal{S}_a$ . To evolve the entire auxiliary surface, we use velocity extension ([11]), whereby we set the force at each point in the remainder of  $\mathcal{S}_a$  as equal to the force at the corresponding closest point on  $C_o$ .

## 5. Results

We ran our experiments on 250 $\mu$ m isotropic ex-vivo MR brain data obtained using a multiecho FLASH pulse sequence at 7T [1]. At this resolution, the stria of Gennari shows up as a dark band 1-2 voxels thick in the contrast-inverted image.

For our experiments, an initial open surface mesh  $\mathcal{S}_o$  was marked manually about half-way between the grey-white interface and the pial surface, with parts of it lying in BA 17 and parts of it outside. Note that such an initialisation could also be obtained automatically, using techniques that localise the grey-white and pial surfaces (*e.g.*, [8]). A ground-truth cortical surface was also obtained from a neuro-anatomist. This surface was provided as a triangle mesh consisting of over 3000 manually labelled vertices.

Each vertex was marked as stria or non-stria. Part of the ground truth mesh was used for training the stria and non-stria likelihood models, while a separate part was used for testing and evaluation.

For the training part of the mesh, normal profile vectors were obtained at the vertices, and PPCA models for stria and non-stria learnt from them. The localisation accuracy of the mesh vertices, as provided by the expert, was  $\pm 1$  voxel ( $250\mu\text{m}$ ). In comparison, the main discriminating feature between stria and non-stria profile vectors is a dark band only 1-2 voxels thick. Hence, the error in the ground truth would offset the profiles and seriously reduce our accuracy in differentiating stria from non-stria. Note that such misalignment is unavoidable, because of the volume of data needed for precisely marking the cortical region, and the tedious nature of the manual labelling task.

To reduce the effect of this error, the training profiles were realigned with the learnt PPCA models, by repeatedly shifting (offsetting) each profile vector in the normal direction by a few voxels and calculating the likelihood under the PPCA model for each shift. Each profile vector was then replaced by the shifted version at which maximum likelihood occurred, and the PPCA models were reestimated. The resulting models were tighter fits, with significantly lower variances along the low-dimensional basis vectors.

In our experiments, normal and tangential evolution were performed alternately in rounds. One round of normal evolution consisted of 20 level-set iterations, while one round of tangential evolution consisted of 500 level-set iterations<sup>5</sup>. A round might have fewer iterations than the above if the corresponding evolution converges.

In our result, the RMS distance between the detected BA 17 boundary and ground-truth boundary<sup>6</sup> was about 1.75mm. This corresponded to a average classification rate of 80% for the stria and non-stria regions (*i.e.* 4 out of 5 voxels in these regions were correctly classified). On average, the total distance moved by a point on the BA boundary, between the initialisation and the final result, was 15mm (60 voxels).

Figure 6 shows zoomed-in views of sample 2D slices through our 3D results to illustrate one round each of normal evolution and tangential evolution. The images correspond to MR data of size 10mm by 7mm. Only the primary surface,  $S_p$ , is shown in each slice (in blue and yellow). To avoid cluttering, the auxiliary surface is not shown, except that it implicitly passes through the boundaries of the visible relevant region.

We note a few salient points. As shown in row (d), if

<sup>5</sup>The separating force term caused tangential evolution to require more iterations for a similar amount of curve motion: a number of iterations were needed just to re-orient the auxiliary surface after normal evolution, before the relevant-region boundary started moving.

<sup>6</sup>taken to be the mid-point of the region where the ground-truth labeling was ambiguous.

tangential evolution is performed first, without any normal evolution, it converges to a local minimum. In contrast, the result of performing both normal and tangential evolution (row (f)) is much more accurate. Note that even though the initial relevant region (row (c)) overlaps significantly with the ground truth non-stria region, the final relevant region after both normal and tangential evolution (row (f)) does not have any holes in it (*i.e.*, the topology is correct). This is mainly because of strong smoothing term used while evolving the auxiliary surface. The apparent discontinuities in the relevant region shown are due to visualisation of a 3D surface in 2D slices. Further, note that after normal evolution (row (e)), parts of the primary surface (in the relevant region) actually move away from the ground truth surface. This is not necessarily an error, since the evolved surface in this region actually matches the typical stria profile better.

Figure 7 illustrates the advantage of alternating normal and tangential evolution. The original surface is shown in (a), followed by the result of one round of tangential evolution in (b). Tangential evolution converges early in this case, because the original primary surfaces deviates from the cortical surface (due to its arbitrary closure). The next round of normal evolution fits the cortex better, as seen in (c). Finally, a subsequent round of tangential evolution further extends the relevant region, as shown in (d).

## 6. Conclusion

We have proposed a novel open-surface evolution algorithm that evolves a surface in both the normal and tangential directions in 3D. The evolution is done in a level-set framework, by alternately updating two implicit functions.

We have also proposed a principled statistical model for modelling structured appearance of surfaces along the normal direction. This model fits naturally within the level-set framework, and generalises the traditional edge-based appearance model. For normal evolution, a mixture of relevant and non-relevant profiles is used for learning an appearance model, while for tangential evolution, a region-based likelihood term is used to solve the classification problem of relevant versus non-relevant region.

The method has been applied to automatic detection of the primary visual cortex (Brodmann area 17), and looks promising for detection of both the location and extent of this cortical region. This opens up possibilities for future automatic analysis of cortical structure directly from high resolution MR data—possibly even acquired *in vivo*—without requiring histological samples. Our formulation can also be applied for locating other cortical regions, as well as to detecting the grey-white and pial surfaces.

Possible future directions of research include incorporating a global shape prior on the open surface, and modelling variation in scale (*i.e.*, length) of profile vectors due to variable thickness of the cortex.



Figure 6. Multiple 2D slices through our (3D) results. The 4 columns are 4 different slices through the MR data. The first row simply shows the MR intensities. In row (b), ground truth labelling is shown in orange (striated Brodmann area 17) and green (non-striated region outside BA 17). Results of surface evolution are shown in subsequent rows, in blue (BA 17: relevant region) and yellow (non-relevant). The third row shows the initial surface before normal and tangential evolution. Row (d) shows the (rather poor) result of performing tangential evolution alone, without any normal evolution. Row (e) shows the results of normal evolution alone. Row (f) shows results after one round each of both normal and tangential evolution are performed. Gap between orange and green curves is due to a region where ground truth labelling was ambiguous, as determined by the expert labeller.

## Acknowledgements

BB would like to thank Thomas Yeo for pointing him towards this medical imaging problem, Ayres Fan and Kush Varshney for tips on level-set methods, and Wanmei Ou, Serdar Balci and the anonymous reviewers for their comments. Funding support was provided in part by the National Center for Research Resources (P41-RR14075, R01 RR16594-01A1 and the NCCR BIRN Morphometric Project BIRN002, U24 RR021382), the National In-

stitute for Biomedical Imaging and Bioengineering (R01 EB001550), the National Institute for Neurological Disorders and Stroke (R01 NS052585-01), and by the National Institutes of Health through the NAMIC Grant U54 EB005149.

## References

- [1] J. Augustinack, A. van der Kouwe, M. Blackwell, D. Salat, C. Wiggins, M. Frosch, G. Wiggins, A. Potthast, L. Wald,




Figure 7. A 2D slice through our (3D) results, illustrating the need for alternating normal and tangential evolution. Shown here are (a) an initial configuration, (b) the results of one round of tangential evolution, (c) the next round of normal evolution and (d) the following round of tangential evolution. Colour scheme used is the same as in Figure 6.

and B. Fischl. Detection of entorhinal layer II using 7 Tesla magnetic resonance imaging. *Annals of Neurology*, 57(4):489–494, 2005.

- [2] E. Barbier, S. Marrett, A. Danek, A. Vortmeyer, P. van Gelderen, J. Duyn, P. Bandettini, J. Grafman, and A. Koretsky. Imaging cortical anatomy by high-resolution MR at 3.0T: Detection of the stripe of Gennari in visual area 17. *Magnetic Resonance in Medicine*, 48(4):735–738, 2002.
- [3] M. Bertalmio, G. Sapiro, and G. Randall. Region tracking on level-sets methods. *IEEE Trans. Medical Imaging*, 18(5):448–451, May 1999.
- [4] K. Brodmann. Vergleichende Lokalisationslehre der Großhirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues, 1909.
- [5] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *Intl. J. Comp. Vis.*, 22(1):61–79, 1997.
- [6] T. Chan and L. Vese. Active contours without edges. *IEEE Trans. Image Processing*, 10(2):266–277, 2001.
- [7] L.-T. Cheng, P. Burchard, B. Merriman, and S. Osher. Motion of curves constrained on surface using a level-set approach. *Journal of Comp. Physics*, 175(2):604–644, 2002.
- [8] B. Fischl and A. Dale. Measuring the thickness of the human cerebral cortex from magnetic resonance images. *Proceedings of the NAS*, 97(20):11050–11055, 2000.
- [9] H. Jin, A. Yezzi, and S. Soatto. Region-based segmentation on evolving surfaces. In *Proceedings ECCV*, 2004.
- [10] J. Kim, J. Fisher, A. Yezzi, M. Cetin, and A. Willsky. A non-parametric statistical method for image segmentation using information theory and curve evolution. *IEEE Trans. Image Processing*, 14(10):1486–1502, 2005.
- [11] S. Osher and R. Fedkiw. *Level Set Methods and Dynamic Implicit Surfaces*. Springer-Verlag, 2003.
- [12] A. Schleicher, K. Amunts, S. Geyer, P. Morosan, and K. Zilles. Observer-independent method for microstructural parcellation of cerebral cortex: A quantitative approach to cytoarchitectonics. *Neuroimage*, 9(1):167–177, 1999.
- [13] J. Sethian. *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science*. Cambridge University Press, 1999.
- [14] J. E. Solem and A. Heyden. Reconstructing open surfaces from image data. *Intl. J. Comp. Vis.*, 69(3):267–275, 2006.
- [15] M. Tipping and C. Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B*, 61(3):611–622, 1999.
- [16] N. Walters, S. Eickhoff, A. Schleicher, K. Zilles, K. Amunts, G. Egan, and J. Watson. Observer-independent analysis of high-resolution MR images of the human cerebral cortex: in vivo delineation of cortical areas. *Human Brain Mapping*, 28(1):1–8, 2007.

## A. Closing an Open Surface

Closing an arbitrary open mesh is non-trivial. The closed mesh should have no boundary edges (and thus be topologically equivalent to the surface of a sphere). We calculate oriented<sup>7</sup> normals at each vertex of the open mesh, and move by an amount  $\epsilon$  along each normal to create a new vertex. New edges are added to connect the new vertices among themselves in a manner paralleling the connection of the old vertices. Vertices on the original boundary are then connected to vertices on the new boundary. For small enough  $\epsilon$ , this algorithm will close an open triangle mesh without producing self-intersections.

## B. Label Extension for Constructing Auxiliary Surface

The auxiliary surface is reinitialised after every round of normal or tangential evolution, to ensure that it is orthogonal to the primary surface, by a process that we call *label extension* (due to its similarity to velocity extension). First, points on  $\mathcal{S}_p$  are labelled as relevant or non-relevant using the sign of the current auxiliary function,  $\Psi_a$ . Then, a fast-marching method is used to fill the volume with a binary label field that indicates whether the closest point on  $\mathcal{S}_p$  is labelled relevant or non-relevant. Solving the eikonal PDE on this label field gives the new auxiliary function  $\Psi_a$ , whose zero level-set,  $\mathcal{S}_a$ , has the desired orthogonality property.

Label extension is also used to initialise the auxiliary surface,  $\mathcal{S}_a$  at the very beginning of the evolution process.

<sup>7</sup>By oriented, we mean that all the normals must point outwards on one side of the surface.