ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage xx (2005) xxx – xxx

Automated 3-D extraction and evaluation of the inner and outer cortical surfaces using a Laplacian map and partial volume effect classification

June Sic Kim,[a] Vivek Singh,[b] Jun Ki Lee,[c] Jason Lerch,[a] Yasser Ad-Dab'bagh,[a] David MacDonald,[a] Jong Min Lee,[c] Sun I. Kim,[c] and Alan C. Evans[a,*]

[a]*McConnell Brain Imaging Centre, Webster 2B, Montreal Neurological Institute, McGill University, 3801 University Street, Montreal, Quebec, Canada H3A 2B4*
[b]*Deparment of Electrical Engineering, McGill University, Montreal, Quebec, Canada*
[c]*Deparment of Biomedical Engineering, Hanyang University, 17 Haengdang-dong Sungdong-gu, Seoul 133-791, South Korea*

Received 21 December 2004; revised 25 March 2005; accepted 29 March 2005

Accurate reconstruction of the inner and outer cortical surfaces of the human cerebrum is a critical objective for a wide variety of neuroimaging analysis purposes, including visualization, morphometry, and brain mapping. The Anatomic Segmentation using Proximity (ASP) algorithm, previously developed by our group, provides a topology-preserving cortical surface deformation method that has been extensively used for the aforementioned purposes. However, constraints in the algorithm to ensure topology preservation occasionally produce incorrect thickness measurements due to a restriction in the range of allowable distances between the gray and white matter surfaces. This problem is particularly prominent in pediatric brain images with tightly folded gyri. This paper presents a novel method for improving the conventional ASP algorithm by making use of partial volume information through probabilistic classification in order to allow for topology preservation across a less restricted range of cortical thickness values. The new algorithm also corrects the classification of the insular cortex by masking out subcortical tissues. For 70 pediatric brains, validation experiments for the modified algorithm, Constrained Laplacian ASP (CLASP), were performed by three methods: (i) volume matching between surface-masked gray matter (GM) and conventional tissue-classified GM, (ii) surface matching between simulated and CLASP-extracted surfaces, and (iii) repeatability of the surface reconstruction among 16 MRI scans of the same subject. In the volume-based evaluation, the volume enclosed by the CLASP WM and GM surfaces matched the classified GM volume 13% more accurately than using conventional ASP. In the surface-based evaluation, using synthesized thick cortex, the average difference between simulated and extracted surfaces was 4.6 ± 1.4 mm for conventional ASP and 0.5 ± 0.4 mm for CLASP. In a repeatability study, CLASP produced a 30% lower RMS error for the GM surface and a 8% lower RMS error for the WM surface compared with ASP.

Keywords: Cortical surfaces; Partial volume estimation; Laplacian map

Introduction

Accurate reconstruction of cerebral cortical surfaces from magnetic resonance images (MRI) is an important issue in quantitative brain analysis, visualization, cortical mapping, and multimodal registration. In particular, accurate reconstruction of the inner- and outer-surfaces of the cortex is essential to obtain reliable measurements of cortical thickness. Most MR images however, suffer from partial volume effects (PVE) due to the limited resolution of MRI scanners which can cause voxels on the boundary between tissues to be blurred, making accurate surface determination difficult. This problem is most pronounced in tightly folded sulci, where cerebrospinal fluid (CSF) is often hardly detected due to the fact that opposing sulcal banks can be closer than the MRI resolution, thus exhibiting a fused appearance. As a result, the boundary surface between gray matter (GM) and CSF, referred to hereafter as the pial surface, might not be correctly localized, and hence any derived morphometric information such as cortical thickness may be inaccurate. The worst case partial volume effect is having no CSF within identified sulci, and thus, no apparent sulci. Fig. 1 shows a surface deformed to fit the pial surface and gray-white boundary, hereafter referred to as the white matter (WM) surface. The WM surface contains well formed sulci, but the pial surface does not due to the fact that partial volume effects are resulting in little CSF being identified in buried sulci. This observation motivated the original ASP authors to use

* Corresponding author. Fax: +1 514 398 8925.
E-mail address: alan@bic.mni.mcgill.ca (A.C. Evans).
Available online on ScienceDirect (www.sciencedirect.com).

1053-8119/$ - see front matter © 2005 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2005.03.036

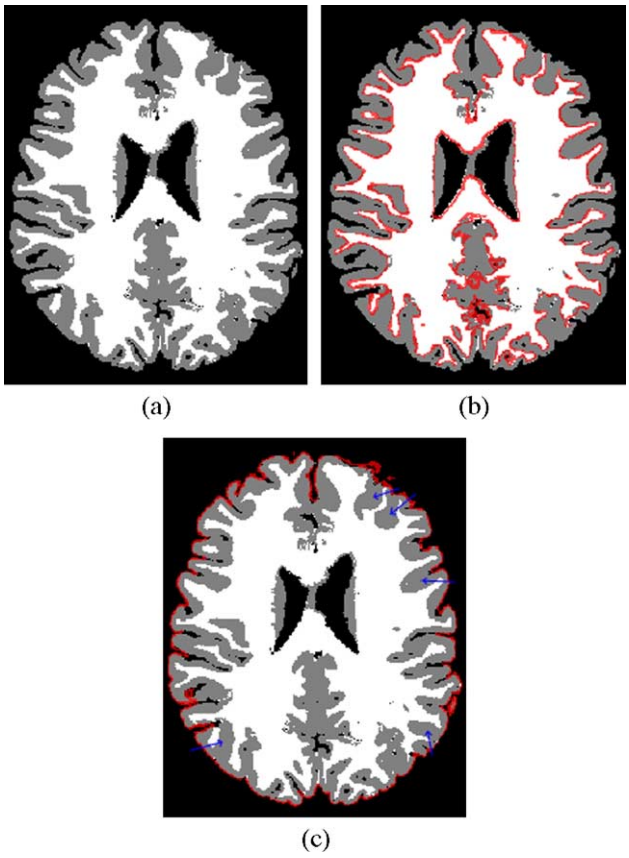

Fig. 1. Demonstration of partial volume effect within sulci. (a) Classified cross-section of an image obscured by partial volume effect, (b) the WM surface (red line) estimated on a discrete classification, (c) the GM surface (red line) estimated on a discrete classification. Blue arrows indicate examples of the obscured boundary between GM and CSF. The GM surface might not describe the folded sulci based on a discrete tissue classification if there is no constraint of sulcus.

information about the relative position of the WM surface in order to improve the pial surface extraction.

A number of methods currently exist for identifying the cortical surface. They can be classified into bottom-up approaches using edge detection and top-down approaches using model based deformation. One of the most common bottom-up approaches is the "Marching Cubes" algorithm which is of restricted utility for cortical surface extraction, since noise and the complex cortical folds make it difficult to obtain a smooth topologically correct surface reconstruction (Lorensen and Cline, 1987). Although there have been approaches to fix the topology of marching cubes surfaces (Han et al., 2002; Shattuck and Leahy, 2001), it is difficult to define the correct outer-cortical boundary because of tightly opposed sulcal banks.

In recent years, several excellent methods using a top-down approach or mixed approach with bottom-up methods have been developed for automatically extracting the cortical surface. Dale et al. introduced an automatic approach that is implemented in the freely-available FreeSurfer program (Dale et al., 1999; Fischl et al., 1999). This program initially finds the WM surface, and then fits smooth WM and pial surfaces using deformable models involving spring, curvature, and intensity-based terms. The authors developed morphological operations to improve the topology of a binary WM volume and then expanded a deformable surface toward the outer GM boundary. This approach was a significant advance, but the algorithm lacked a method for preventing self-intersecting topologies. To preserve topology, MacDonald et al. developed a cortical surface extraction procedure known as Anatomic Segmentation using Proximity (ASP) (MacDonald, 1998; MacDonald et al., 2000). This method is an improved version of Multiple Surface Deformation (MSD) for simultaneous deformation of multiple curves and surfaces to an MRI, with inter-surface constraints and self-intersection avoidance (MacDonald et al., 1994). ASP uses a topology-preserving deformation model with proximities, and automatically identifies the WM and pial surfaces of the cerebral cortex in a robust way with respect to partial volume effects by means of forcing the cortical thickness to lie within a defined anatomically-plausible range of values. Since the ASP algorithm deforms from a spherical polygonal model with proximities, it preserves the topology of a spherical sheet which, once deformed, serves as a model of the cortical sheet. But since this approach imposes a thickness constraint in order to insure topological correctness, it introduces a bias in the true calculation of cortical thickness in populations which may have much thicker or much thinner cortices than the constraints in the algorithm currently allow, such as children or specific diseased populations. Moreover, since the ASP method does not account for the poor differentiation between the putamen and the insula in the discretely classified image it uses as input, thickness values in the insular cortex are very inaccurate. Han et al. also introduced a topology preserving geometric deformable model (Han et al., 2001a,b), which reconstructs the pial surface by correcting the topology of a GM segmentation. This method estimates the position of sulci by assuming that they are equidistant, on either side, to the WM surface. While this approach was successful at preventing self-intersection, creating the representative topology based on the WM surface misplaces asymmetrically folded sulci. Kriegeskorte and Goebel developed a method which detects and removes topological errors as part of tissue classification (Kriegeskorte and Goebel, 2001). Their method uses a self-touching sensitive region growing process prioritized by distance-to-surface voxels considered for inclusion, and topologically corrects reconstructions of a cortical sheet. The use of region growing methods for detecting edges is an efficient and fast approach to surface extraction, however, the topology correction does not perfectly reconstruct the pial surface in the buried sulci where the intensity level in regions which contain CSF are very similar to that for GM due to partial volume effects.

In this paper, we introduce a fully automatic method to reconstruct the pial surface, called CLASP, which stands for Constrained Laplacian-based ASP, a modification of the original ASP algorithm (MacDonald et al., 2000). This method uses a more complex classification method with a statistical probabilistic anatomical map (SPAM) to estimate the boundary of the insula. Moreover, CLASP makes use of a novel geometric deformable surface model which estimates the pial surface through the use of partial volume effect information. We specifically focus on improving the accuracy of reconstructing the pial interface, as this surface presents the most serious morphological challenges to the conventional ASP method. To remove the necessity of a cortical thickness constraint, we incorporated a classification method for detecting the partial volume fraction of CSF in deep sulci. The pial interface is created by expanding from the WM surface, while the WM surface is extracted by the same technique used in the original ASP procedure. The expansion path for creating the pial surface is defined as a Laplacian map between the WM surface and a

skeletonized CSF map. The Laplacian approach is similar to a method presented by Jones (Jones et al., 2000) to estimate cortical thickness, but differs in that it is performed in voxel space and is used for surface expansion. This algorithm solves Laplace's equation with the cortical volume as the domain for solution of the differential equation, as well as separate boundary conditions at the gray-white junction and the gray-CSF junction. Distinct from other methods, we use measured CSF from partial volume classification to preserve the morphology of the boundary between GM and CSF, particularly in buried sulci, while the outer bound of the GM surface expansion is defined by the GM/CSF interface from a discrete classification. This framework ensures that the method will yield surfaces that are topologically equivalent to a sphere and that do not self-intersect, completely removing any requirement of prior information about a range of allowable thickness values as a constraint. Moreover, we use the Laplacian term so as to locally vary the expansion force of the pial surface by predicted cortical thickness at each vertex.

In order to compare the results, we made use of three evaluation procedures. These included a volume-based evaluation which compared GM voxels identified by discrete tissue classification with a GM voxel map created by each of ASP and CLASP. This voxel map is essentially cortex voxels bounded by the WM and pial surfaces. Secondly, a surface-based evaluation was used which compared a pial surface extracted by ASP or CLASP with a simulated, "ground truth" template surface. Lastly, we analyzed the consistency of the CLASP procedure using surfaces extracted from 16 scans of the same individual.

Method

The CLASP algorithm consists of several stages as follows: (1) acquired T1 MR images are preprocessed by intensity inhomogeneity correction and spatial normalization to stereotaxic space. (2) Preprocessed images are classified into GM, WM, and CSF tissues. (3) Processed volumes are divided into left and right hemispheres for reconstructing two hemispheric cortical surfaces. (4) The WM surface is reconstructed by deforming a spherical polygon model to the white matter boundary. (5) A Laplacian field is generated between the WM surface resampled to voxel space and a skeletonized CSF fraction image. (6) The GM surface is initiated from the WM surface and is expanded to the boundary between GM and CSF along the Laplacian field. These stages are described in greater detail in what follows.

Non-uniformity correction and registration

Intensity non-uniformity in raw MR images due to magnetic field inhomogeneity must be corrected so that volumes can be properly classified into GM, WM and CSF. The non-parametric N3 method, which iteratively maximizes the frequency content of tissue intensity distributions (Sled et al., 1998) was used here.

In order to spatially align the brains, automatic registration to stereotaxic space (Collins et al., 1994b; Evans et al., 1992, 1993; Talairach and Tournoux, 1988) was performed using the intensity-corrected image. This procedure computes the linear transformation parameters by gradient descent at multiple scales to maximize the correlation between the individual volume intensity and an average template volume (Collins et al., 1994a). Finally, a 3-D stereotaxic brain mask is used to remove extra-cerebral voxels.

Classification

Two different types of classified images are used during the surface extraction procedure. One is a discrete classification that is employed to find the boundary between WM and GM as well as GM and CSF. The other is a probabilistic partial volume classification used to detect buried CSF voxels in sulci and to differentiate subcortical and cortical gray matter. In discretely classified volumes, voxels classified as pure CSF are rare in folded sulci due to the fact that sulcal walls are often so close that voxels in these regions, even in high resolution images, have a higher content of GM than CSF, making the sulcal walls appear fused. Thus, the true pial surface is generally not found in the folded sulci. In order to compensate for this partial volume effect in these regions, we utilized partial volume classification, providing fractional estimates of the content of GM and CSF for such voxels. Improved localization of the pial surface is then possible since voxels in deep sulci which contain small amounts of CSF are given a mixed GM/CSF classification instead of a pure GM classification.

Discrete tissue classification

A k-nearest neighbor (k-NN) classification technique is used to automatically label each voxel by tissue type (background, CSF, GM, WM), using a non-uniformity corrected, high resolution MR image in stereotaxic space. This method makes use of stereotaxic spatial priors as well as intensity features to reduce the occurrence of anatomically implausible voxel classification (Cocosco et al., 2003). These 3-D prior volumes containing the probability of each voxel belonging to a given tissue class, were generated from a sample of several hundred MR volumes that were classified by a semi-automatic technique. The training set is defined for the k-NN classifier, by first selecting points from voxel positions with a high probability of containing a given tissue type and then pruning to remove outlying intensity values for each tissue class, thus customizing for a particular scan. Each voxel in the subject images is then independently classified into one of the defined tissue types resulting in a discrete three dimensional volume, where each voxel has one of the aforementioned labels.

Partial volume classification

The partial volume segmentation was performed using a statistical model of the partial volume effect known as the mixed model (Choi et al., 1991), in which the contents of a voxel are expressed as the weighted sum of random variables. The weights represent the fractional amount of a particular tissue, while the random variables represent the intensity distribution for each pure tissue class. The approach taken here is an extension of a method we have previously presented (Tohka et al., 2004) which has been modified to improve segmentation in deep sulci and to differentiate between cortical and subcortical (SC) gray matter. The segmentation procedure is divided into two steps, including the estimation of a context image (C^*) (i.e., voxels are given either a pure or mixed tissue label: GM, WM, CSF, SC, GMWM or GMCSF, SCWM, or SCGM), necessary to make the problem tractable in the single spectral case and a fraction estimation stage to estimate the amount of tissue in mixed class voxels. The iterative condition modes (ICM) algorithm is then used to solve for the context image, while maximum likelihood estimation is used to solve for the fraction values. If C represents the context image we are trying to estimate

and X represents the acquired image, an estimate of the context image can be made use a Maximum A Priori (MAP) criterion:

$$C^* = \arg\max_C \; p(C)\; p(C\,|\,X),$$

where the prior information, $p(C)$ is modeled as a Markov random field. One of the modifications we have made to our previous work is to use a spatially varying Markovian random field (MRF) as prior information, which acts in a smoothing fashion in most image areas, but adaptively modified its behavior in deep sulcal regions to encourage voxels with mixed tissue labels to occur adjacent to voxels that have pure tissue labels. Like our previous work, we employ a Potts model to represent the MRF:

$$p(C) \propto \mathrm{p}\left(-\beta \sum_{i=1}^{N} \sum_{k \in N_i} \frac{a_{ik}}{d(i,k)} \right).$$

where β is the MRF weighting parameter, i is the index of a voxel from the image X, N is the total number of voxels, N_i is the 26 neighborhood of i, k is the index of a voxel from N_i, and $d(i,k)$ is the Euclidean distance between voxel i and k. The value of a_{ik} is determined by the relationship between the voxel label at k, C_k, and the label at i, C_i. In our previous work, a_{ik} solely worked to encourage spatial smoothing. In our present approach, its behavior is modulated to encourage GM/CSF and CSF labeled voxels to occur between GM labeled voxels in deep sulcal areas:

$$a_{ik} = \begin{cases} -2 & c_i = c_k \\ -1 & c_i, c_k \text{ share a component and} \\ & c_i \notin \{GMCSF, CSF\} \text{ or } c(X,i) \geq 0 \\ f(M(X,i)) & c_i, c_k \text{ share a component and} \\ & c_i \notin \{GMCSF, CSF\} \text{ and } c(X,i) < 0 \\ 1 & \text{otherwise} \end{cases},$$

where M is an MRF-influence image used to probabilistically identify deep sulcal areas, and f is a sigmoid function that modulates the influence of the parameter. The influence image is based on an average of two pieces of information. The first of these is image curvature calculated from a T1 scan blurred with a Gaussian kernel (FWHM = 4 mm) (Thirion and Gourdon, 1993), where negative curvature areas represent sulcal regions. The second is the medial surface of the union of GM and CSF from a discretely classified image. This latter image is determined by calculating a distance transform of the inverse of the GM + CSF union (Borgefors, 1984). A gradient magnitude image is then calculated from the distance transformed image and the resultant image is geometrically smoothed (Kimia and Siddiqi, 1996) to reduce noise. Finally, its intensity values are transformed so that medial regions roughly correspond in value to the negative areas in the curvature image. Examples illustrating these two types of information are shown in Fig. 2.

A significant increase in voxels identified as containing CSF using the partial volume method as opposed to discrete classification can be seen in Fig. 3, particularly in buried sulci.

The other modification made was to include a subcortical gray matter tissue (SC) class in the partial volume segmentation, along with classes for mixed WM/SC and SC/GM. These classes were not used in the entire image, however, but only in regions where subcortical tissue (thalamus, globus pallidus, putamen, and claustrum) and the insular cortex were expected based on SPAMs generated from 40 images in the ICBM data set. Specifically any voxel that had any probability > 0 of coming from one of these structures was classified using the extra classes. In such regions, the WM/GM class was also eliminated. Identification of all structures prior to creating the SPAMs was done using the ANIMAL non-linear registration package (Collins et al., 1994a,b). The mean and variance values for the SC class were calculated by using the intensity of points with >50% likelihood of coming from a subcortical structure and applying a robust parameter estimation technique (Tohka et al., 2004). An example of a classification using the extra classes is shown in Fig. 4.

Skeletonization of the CSF fraction image

In CLASP, the GM surface is determined by expanding outward from the WM surface until the GM/CSF boundary is found. The process is guided (in a manner that is outlined in the next section) by the use of partial volume CSF in order to insure correct sulcal morphology. To facilitate this, the partial volume CSF calculated is binarized, setting any voxel containing CSF to 1 and all others to 0, and then skeletonized. The skeleton is created using a 2-subfield connectivity-preserving medial surface skeletonization algorithm. (Ma and Wan, 2001; Ma et al., 2002). An example of a CSF skeleton is shown in Fig. 3.

Reconstruction of the WM surface

The classified MR image is separated into volumes representing each cortical hemisphere, by identifying the midsagittal plane

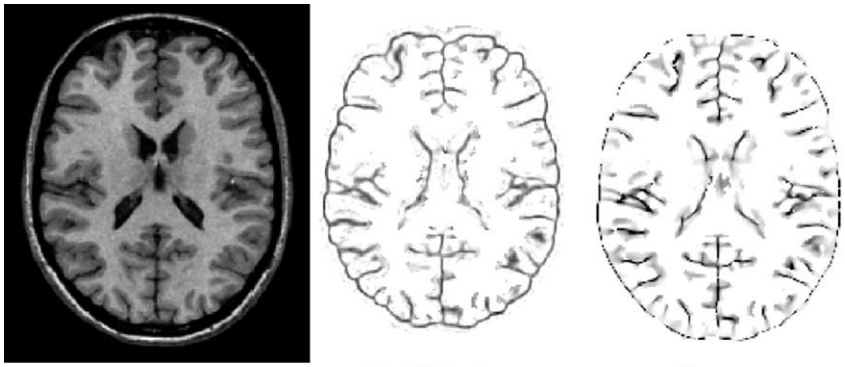

Fig. 2. A T1 weighted image, a medial surface image and negative curvature image, respectively. Note how sulci are generally localized in the darker areas.

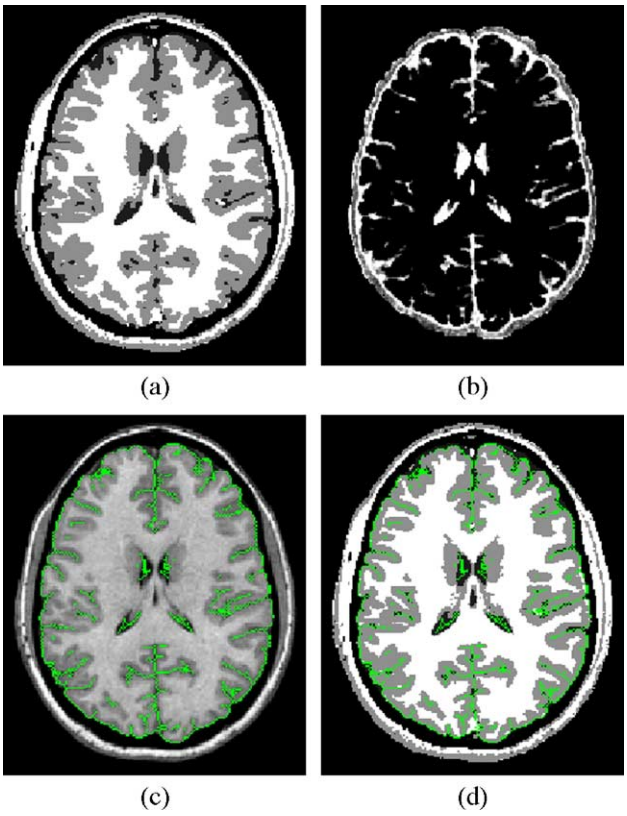

Fig. 3. Classification and skeletonized CSF: (a) Discrete classification, (b) PVE classified CSF, (c) the skeletonized CSF (green) overlapped with its T1 image, and (d) the skeletonized CSF (green) overlapped with INSECT classification. Note that skeletonized CSF voxels represent buried sulci very well while INSECT classification does not.

which passes through anterior and posterior commissures in spatially normalized stereotaxic space.

The WM surface in each hemisphere is reconstructed in the same manner as the ASP algorithm. As described in Classification, since segmentation of the subcortical structures is not perfect, the WM surface is corrected using gradient map values around the boundary of the insula calculated from a T1 image. Due to the presence of high gradient boundaries in surrounding structures, it is possible that the WM surface at the insula could be localized to an improper boundary. We reduce the possibility as follows: (1) since the segmentation for the subcortical structures based on SPAMs gives a good initial localization of the WM surface, the WM surface needs only to be slightly deformed. (2) Voxels in the gradient map close to the boundary between GM and CSF are masked out using a binary mask image, created by calculating the intersection of the gradient of a CSF fraction image and the gradient of a classified GM tissue volume. (3) The gradient image is also blurred by a 1-mm FWHM kernel to reduce noise. We add a gradient term, $T_{Gradient}$, to the original set of ASP constraints, to fit the WM surface:

$$T_{Gradient} = \sum_{v=1}^{n_v} w_G \cdot (G_M - G(x_v))^2$$

where $G(x_v)$ is the gradient map as described above, G_M is a local maximum value in the gradient image, and w_G is a weight value for the gradient term.

Cross-sections of the white matter surface with various levels of correction for the insula region are show in Fig. 5. Without subcortical segmentation, the WM surface is not properly fitted in subcortical area (Fig. 5a). With subcortical segmentation, the WM surface reconstructs the boundary of insula but there are mismatched vertices due to misclassification in the subcortical segmentation (Fig. 5b). The WM surface is corrected to a qualitatively more likely boundary by means of adjustment with a gradient image term (Fig. 5c).

Laplacian map

The human cerebral cortex is a thin layer of GM at the outer surface of the brain and is topologically equivalent to a sheet. Surface extraction methods that expand from the WM surface to the pial surface have used an image term such as the local gradient of the image so as to push the surface toward the correct edge (Dale et al., 1999; Han et al., 2001a,b; MacDonald et al., 1994). However, if the deforming surface is far from an edge, or the direction to the edge is misleading, there may be difficulty locating the edge correctly. In our study, the direction towards CSF voxels is determined by a Laplacian map similar to a map of boundary distance. If we give a minimum value to WM voxels and a maximum value to the skeletonized CSF and background, the computation of the Laplacian provides smoothly increasing potential surfaces between WM and CSF. Its gradient produces an expanding route for the pial surface similar to a distance map, which labels voxels by the distance from the WM surface, and guarantees topology preservation. The Laplacian map has two major advantages over a distance map. First, the Laplacian map always generates the same distance ratio from the starting surface (a preset minimum value) to the ending face (a preset maximum value) even though cortical thickness varies across the cortex. Thus gradients of the Laplacian map in the thicker parts of the cortex are smaller than those in the thinner parts. We introduce a novel term to control the degree of expansion of the pial surface by using the gradient of the Laplacian map. In contrast to conventional approaches using Laplacian maps, the pial surface is expanded slowly in regions of thinner cortex and expanded faster in the thicker parts of the cortex. This ensures that every vertex on the surface moves to the outer boundary of GM at a uniform rate, proportional to the local thickness. This is important in performance terms because

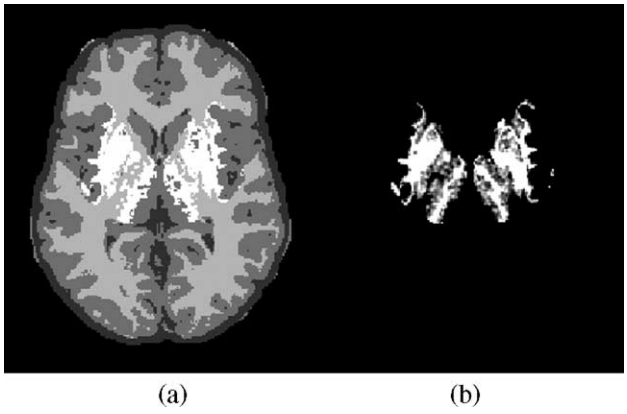

Fig. 4. A partial volume classification (a), illustrating areas (b) that have been classified as subcortex (intensity indicates probability of subcortex). Light gray and white areas of (a) are classified as WM. This helps in correctly localizing the boundary of the insula.

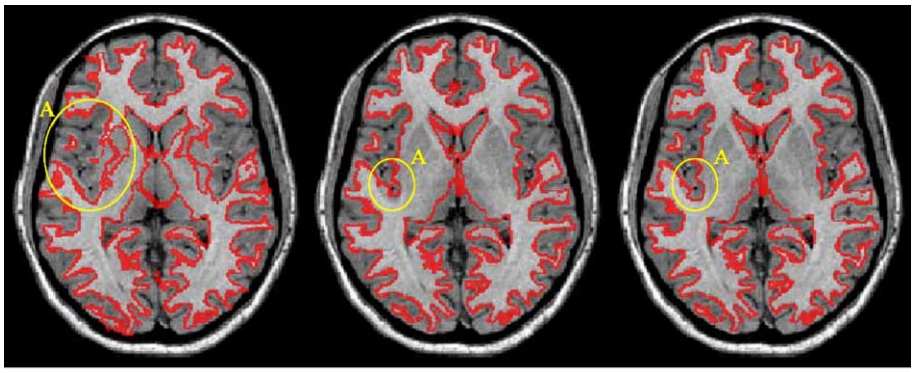

Fig. 5. Subcortical segmentation and the WM surface of (a) ASP, (b) CLASP before correction, and (c) CLASP after correction. Note the fitted surface in region A.

vertices globally converge faster with the thickness constraint than without the constraint. Moreover, if some of the vertices were able to reach the final pial surface earlier than the others, those vertices would be moved by other model constraint terms such as the stretch and the self-intersect terms rather than the Laplacian term. This may not only misplace the vertices from the Laplacian map but also may terminate the iterations though some vertices are not expanded enough. The Laplacian map also guarantees a one-to-one mapping between the two surfaces without intersection of paths as shown in Fig. 6a.

Another important fact is that the stream lines of the Laplacian map have characteristics similar to those of cortical striations in that they are unique as well as non-intersecting and thus seem better suited to capture cortical features than simply

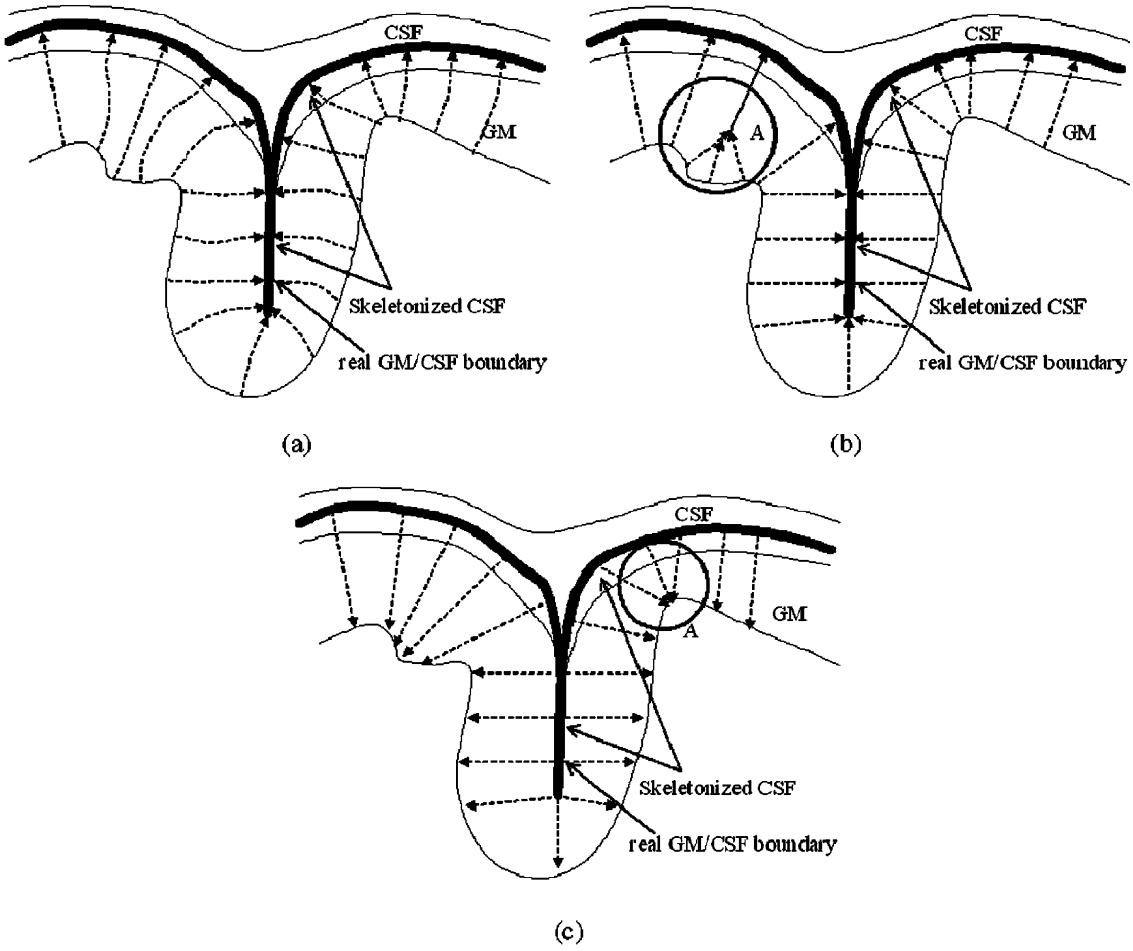

Fig. 6. Demonstration of the deformation route using (a) the Laplacian map, (b) the distance map from the WM surface, and (c) the distance map from the pial surface. While the distance maps may lead neighboring vertices to collide as shown in region A of (b) and (c), the Laplacian map smoothly expands vertices and reduces the collision.

using a distance map. The distance map between surfaces is also not unique in that calculating the distance from the WM surface outward (Fig. 6b) or from the skeletonized CSF inward (Fig. 6c) can produce different values. Fig. 6b shows that vertices in the region denoted "A" may not expand outward to the skeletonized CSF and can map to the same location as they should but across the bulge shown at A. In another situation, many vertices on the pial surface may map to the same location on the WM surface as shown in Fig. 6c (region A).

Constraints for expanding the GM surface

The CLASP algorithm uses a series of terms to achieve a satisfactory match between the deforming surface mesh and the GM/CSF interface in the volume image while adhering to the shape constraints of the model.

Objective function

As with the conventional ASP, the CLASP algorithm uses an objective function that is minimized in order to fit the model constraints and image data. The domain of the function is the set of vertex coordinates, and the range is a scalar value representing the goodness of fit of the polyhedron to the target data. The objective function, $O(S)$, similar to the function used in ASP, is defined generally as a weighted sum of $(N = 4)$ terms:

$$O(S) = \sum_{k=1}^{4} T_k(S),$$

where each term $T_k(S)$ measures some aspect of the polyhedral pial surface, S. The four terms of $T_k(S)$ are $T_{\text{Laplacian}}$ defined by the Laplacian map, $T_{\text{intensity}}$ determining the boundary between GM and CSF, T_{stretch} which equalizes the distance among polygon vertices and $T_{\text{self-proximity}}$ that prevents polygon intersection. Each term includes a weighting coefficient.

CLASP substitutes the Laplacian term and an intensity term for the boundary distance term and the vertex–vertex proximity term of ASP. Other terms for stretch and self-proximity are left as they were. The following sections provide more detail about these terms.

Laplacian term

The Laplacian term presented here is based on the Laplacian map. The Laplacian map is generated by solving Laplace's equation, $\nabla^2 G(\bar{x}_v) = 0$, in a grid volume $G(\bar{x}_v)$, where $\bar{x}_v$ is a point in space. The grid volume is defined as follows:

$$G(\bar{x}_v)\begin{cases} 0 & \text{if } x_v \text{ is inside of the WM surface} \\ 10 & \text{if } x_v \text{ is on the skeletonized CSF or background} \\ 5 & \text{otherwise} \end{cases}$$

Each vertex moves into the location of higher value in the Laplacian map which is expressed mathematically as:

$$T_{\text{Laplacian}} = \sum_{v=1}^{n_v} w_L \cdot \left(L(\bar{x}_v) - L_M \cdot d_L(\bar{x}_v)\right)$$

where n_v is the number of vertices in polyhedral mesh, $L(\bar{x}_v)$is a value of the Laplacian map at the location $\bar{x}_v$, and L_{M} is the maximum value of the Laplacian map. $d_L(\bar{x}_v)$ is the Euclidean distance between vertex v on the expanding surface and its linked vertex on the initial surface, that is approximately 1 and w_L is a weight value for the Laplacian term. In our study, the discrete Laplacian map with a minimum value of 0 and a maximum value of 10 is estimated by iterative Jacobi relaxation (Press and Flannery, 1988), so values of the Laplacian map are always positive and increase at a rate proportional to the distance between the WM surface and the skeletonized CSF inside of the brain mask. The movement direction, defined by the gradient of the Laplacian term at each vertex, is calculated as:

$$T'_{\text{Laplacian}} = \sum_{v=1}^{n_v} w_L \cdot \left(-SIGN\left(L'(\bar{x}_v)\right) \cdot \left(\left|L'(\vec{\bar{x}}_v)\right| + L_M\right)\right).$$

The relation between the Laplacian term and the distance from the WM surface to the skeletonized CSF is shown in Fig. 7. Vertices move slower in thinner cortex than in thicker cortex.

Intensity term

The intensity term determines the boundary between GM and CSF while the Laplacian term determines expansion to that boundary. As described in Reconstruction of the WM surface, the Laplacian map is defined in the region between the WM surface and skeletonized CSF voxels. The skeletonized CSF voxels help define the morphology of sulci, however, they do not indicate the boundary between GM and CSF (Fig. 3b). In order to find the pial surface, the voxel tissue labels around each vertex are examined. Vertices which lie within GM voxels continue to expand, but vertices within CSF are clamped back to the nearest

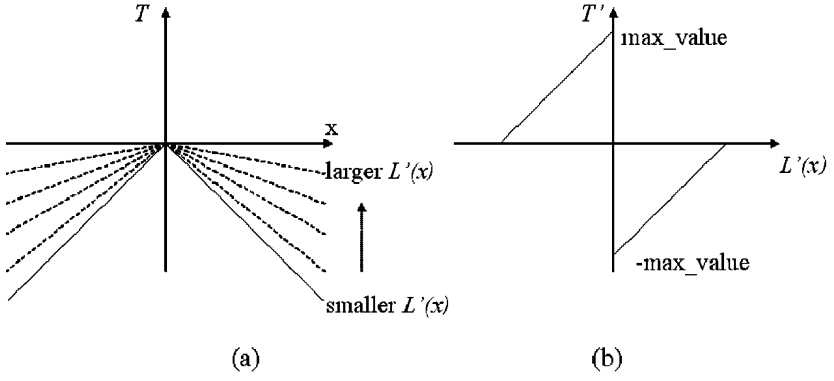

Fig. 7. Shape of the Laplacian term: (a) the Laplacian term, (b) derivative of the Laplacian term. The Laplacian constraint varies according to the derivative of a Laplacian field.

GM voxel along the expansion direction. The function, which detects classified voxels and restricts the expansion, is a modification of the vertex–vertex proximity term of ASP. In ASP, the vertex–vertex proximity term constraints points on two surfaces to be a certain distance apart. However, in CLASP, the intensity term is only activated to shrink the vertices if vertices are found within CSF. Otherwise, the term is inactivated and vertices are expanded along the Laplacian map. The intensity term is defined by:

$$T_{\text{Intensity}} = \sum_{v=1}^{n_v} w_i \cdot d_i(\bar{x}_v)$$

where n_v is the number of vertices in polyhedral mesh, d_i is the distance between vertex v on the expanding surface and its linked vertex on the WM surface. w_i is a weight value. This weight is zero to inactivate the term if the vertex v is within CSF, otherwise, it takes a constant value.

Stretch term

The stretch term measures lengths between vertices, and increases if vertices are stretched or compressed relative to an average length of the current pial surface.

$$T_{stretch} = \sum_{v=1}^{n_v} \sum_{j=1}^{m_v} \left(\frac{d_s\left(\bar{x}_v, \bar{x}_{n_{v,j}}\right) - L}{L} \right)^2$$

where n_v is the number of vertices in polyhedral mesh, m_v is the number of neighbors of vertex v, L is an average length of the current pial surface, and d_s is the distance between vertex v, $\bar{x}_v$, and its neighboring vertices, $\bar{x}_{n_{v,j}}$.

Self-proximity term

The self-proximity term measures the proximity of pairs of nonadjacent polygons in a surface,

$$T_{self-proximity} = \sum_{i=1}^{n_p-1} \sum_{j=1+1}^{n_p} \begin{cases} \left(d_{\min}(P_i,P_j) - d_{i,j}\right)^2, & \text{if } d_{\min}(P_i,P_j) < d_{i,j} \\ 0, & \text{otherwise,} \end{cases}$$

where n_p is the number of polygons, $d_{\min}(P_i,P_j)$ is the smallest Euclidean distance between the ith polygon, P_i, and the jth polygon, P_j, and $d_{i,j}$ is a distance threshold. Pairs of adjacent polygons are not included in the above equation, as the distance $d_{\min}(P_i,P_j)$ is a constant zero value for any deformation of the polyhedra in this case.

Selection of each weight

Weights for the model terms were experimentally selected. To do this, 10 T1-weighted 3-D MRI datasets were randomly tested with different weights for each of the constraints. The weights for the Laplacian term and the intensity term were changed individually while the weights for the stretch term and the self-proximity term remained fixed. A voxel-based comparison between cortical GM obtained by the classification procedure in Classification and cortical GM defined between the WM/GM and GM/CSF interfaces produced by CLASP with each of the various parameter sets being evaluated using a Kappa statistic. Since MacDonald already established the optimal weights for the stretch and self-proximity constraints, we use the values he specified (MacDonald et al., 2000). The weights for the Laplacian and intensity constraints added in this paper were also evaluated. After many empirical tests, the ranges of the weights were limited to 10^{-4}–10^{-2} and 10^{-8}–10^{-4} for the Laplacian and intensity constraints, respectively. First, we fixed the weight of the intensity constraint at 10^{-7} and changed the weight of the Laplacian constraint. Next, the weight of the intensity constraint was changed with the best weight for the Laplacian constraint fixed. Finally, 10^{-3} and 10^{-6} were selected for the Laplacian and intensity constraints, respectively.

Evaluation

To evaluate the CLASP algorithm, we used T1-weighted images of 70 pediatric brains with tightly folded gyri, possessing an age range of 16 ± 2.8, each of 1.0 mm × 1.0 mm × 1.0 mm resolution and 181 × 217 × 181 voxel dimension. In each brain, the pial surface was extracted by both the CLASP and the ASP methods. The accuracy of each method was evaluated by both volume- and surface-based comparisons as detailed below. In addition, a repeatability test was performed with 16 MRI scans of the same subject.

Volume-based comparison

In the volume-based evaluation, we compared classified cortical GM with voxels in between the GM and WM surfaces. This comparison evaluated fidelity based on input information to CLASP and ASP because they both extract cortical surfaces based on classified volumes. It is possible that the reconstructed surface does not exactly represent the GM/CSF boundary not only in sulci but also in gyri because of the distance between adjacent vertices and local minima in the optimization. Thus, in this section, we evaluated how well the GM surface was expanded to the gyri. The surface in buried regions is evaluated in the next sections.

After extracting cortical surfaces a GM map was created by resampling the surfaces to a 1 × 1 × 1 mm volume and filling in voxels between the WM and GM surfaces for both ASP and CLASP. The classified GM maps were masked such that only cortical gray matter voxels were present. The two GM maps from the classification and the surface extraction were then overlapped and compared.

In the overlapped volume, three statistics were used: (1) percentage of matched GM voxels to total voxels between the classified image and the surface-masked image (True-Positive, TP), (2) percentage of background voxels not identified by surface masking but classified as GM voxels (False-Negative, FN), and (3) percentage of voxels identified as GM by the surface-based method but classified as background (False-Positive, FP). Equations for the three statistics are:

$$TP = \sum_{v=1}^{N_{Gv}} \left(\frac{G_c(v) \cdot G_s(v)}{N_{Gv}} \times 100 \right),$$

$$FN = \sum_{v=1}^{N_v} \left(\frac{G_c(v) - G_c(v) \cdot G_s(v)}{N_v} \times 100 \right),$$

$$FP = \sum_{v=1}^{N_v} \left(\frac{G_s(v) - G_s(v) \cdot G_c(v)}{N_v} \times 100 \right),$$

where v is the voxel index, N_v is the total voxel count in a volume, N_{Gv} is the number of INSECT GM voxels, and $G_C(v)$ and $G_S(v)$ are GM voxel maps (values of GM voxels are 1 and values of other voxels are 0) of the classified volume and surface-masked volumes, respectively.

Cross-sectional GM voxel maps produced for a single individual by each method are shown in Figs. 8a–c, while maps showing differences between classified and surface GM are also shown in Figs. 8d–f. Fig. 8e shows less difference voxels than Fig. 8d. Fig. 8f displays GM voxels which are extracted by CLASP but not by ASP. To create the surface masks a scan line filling algorithm was used to resample the surfaces to voxel space. This causes voxels on the surface to receive a GM label regardless of how much of the surface actually encapsulates a boundary voxel, which may result in a higher FN rate. CLASP showed a 13% improvement in accuracy in terms of GM matching compared with the conventional ASP as shown in Table 1. Note that the FP errors are largely reduced in CLASP (0.8%) compared with conventional ASP (18.7%). Since CLASP and ASP algorithms use classified volumes as input, they do not reconstruct the subcortical region, cerebellum, and brain stem correctly. Thus, those non-cerebral voxels reduce TP ratios.

Although the volume-based comparison evaluates how the surfaces reconstruct cerebral volume, it does not specifically show an improvement in the extraction of deep sulcal regions. We thus introduce additional comparisons to evaluate accuracy and precision in the next sections.

Table 1
Volume-based evaluation for pediatric brains

	TP	FN	FP
(1) Conventional ASP (matched voxels to whole voxels)	58.1%	4.8%	18.7%
(2) Laplacian-based ASP (matched voxels to whole voxels)	71.1%	3.3%	0.8%
Difference rate of (2)−(1)	13% (increased)	−1.5% (decreased)	−17.8% (decreased)
t test probability	0.00071	0.00035	0.000006

Surface-based comparison

It is difficult to define a gold standard for MRI surface reconstruction, as it is virtually impossible to produce surface meshes from volumes that could be defined as “ground truth”. In order to circumvent this problem, we made use of the surface extraction output of another algorithm to generate volumes where the ground truth boundaries for the WM and pial surfaces are known. While the phantom surface might not be a true gold standard, this approach provides surfaces that will have many of the important features of a real cortex and not be biased by one of the algorithms under study.

For the surface-based evaluation, it is required to make MR images from the phantom surfaces since the CLASP and ASP

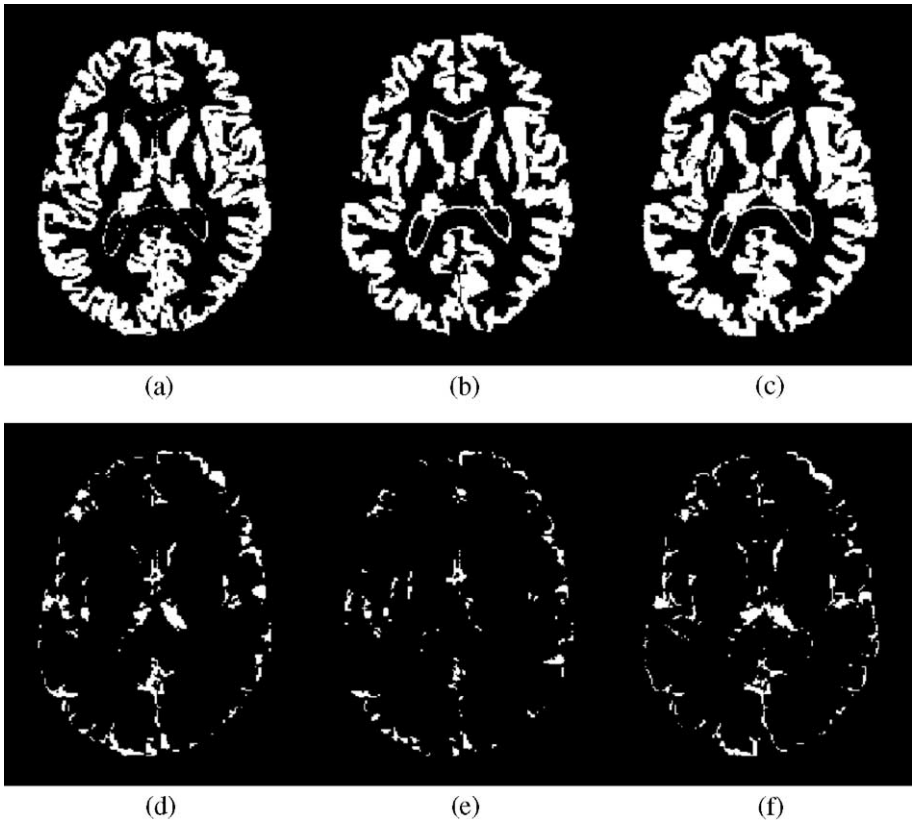

Fig. 8. Extracted GM voxel maps and difference voxel maps: (a) GM by INSECT, (b) GM by ASP, (c) GM by CLASP, (d) a difference map between INSECT and ASP [(a)–(b)], (e) a difference map between INSECT and CLASP [(a)–(c)], and (f) a difference map between CLASP and ASP [(c)–(b)]. Note that voxels of (e) are less than voxels of (d).

procedures start from an MRI. We used an MRI simulator which generates a realistic MR image incorporating partial volume effects and noise from a phantom tissue volume (Collins et al., 1998; Kwan et al., 1996). The surface-based evaluation was divided into several steps as follows:

(1) Pial and WM surfaces were first estimated using another surface reconstruction algorithm (FreeSurfer) (Dale et al., 1999). Surfaces from the software were used as a standard to be compared to the CLASP surfaces extracted from simulated MRI although the FreeSurfer surface did not mean "ground truth".
(2) The surfaces were linearly expanded with an experimentally chosen scaling factor of 1.2. The conventional ASP showed a tendency to incorrectly estimate the pial surface in cortex thicker than a given thickness range, so the surfaces were scaled up and evaluated to determine if this problem was solved using CLASP.
(3) A phantom including 4 labeled tissues (GM, WM, CSF, and background) was created from the surfaces. WM voxels were defined inside of the WM surface, and GM voxels were set between the pial and WM surfaces. To create partial volume effects, voxels on the pial surface were given partial volume fractions of 75% for GM and 25% for CSF. Fractions of 75%/25% were selected because these were middle numbers between no partial volume effect (100%/0%) and an ambiguous evenly mixed status (50%/50%). CSF voxels were also defined in the voxels between the inside of the exterior brain mask and outside of pial surface to simulate extra-pial CSF. All other voxels were set to a background label.
(4) The MRI simulator was given parameters of TR = 18 ms, TE = 10 ms, slice thickness = 1 mm, the same resolution and volume size as the original classified image, in the simulation of a T1 MR image from the phantom.
(5) Using the simulated MR image, the pial surfaces were reconstructed by ASP and CLASP.
(6) Finally, we measured the distance between the two surfaces resulting from (2) and (5), based on the average Euclidean distance from the vertices of the phantom surface to the nearest vertices of the recreated surface.

The ASP algorithm had a much higher average error (4.3 ± 1.6 mm) than did CLASP (0.6 ± 0.6 mm). This evaluation was performed on a scaled phantom volume to synthesize thick cortex, so the difference indicates that ASP would likely have more difficulty in finding the correct boundary in thick cortical areas when compared to CLASP. Even though we introduced partial volume effects for voxels on the pial surface, the voxels were classified as GM because of the higher GM fraction. It was thus difficult for ASP to find the real boundary in the folded sulci. Also, ASP approximates the boundary in these regions by means of the a priori vertex–vertex constraint, while CLASP makes use of a PV classification, resulting in better correspondence with the "ground truth" boundary. Fig. 9 demonstrates the difference between surfaces before and after MRI simulation. The extraction by conventional ASP resulted in a poorly expanded surface compared with the phantom mesh because of the vertex–vertex constraint, as shown in Fig. 9a. The CLASP surface, by contrast, is expanded better, as shown in Fig. 9b.

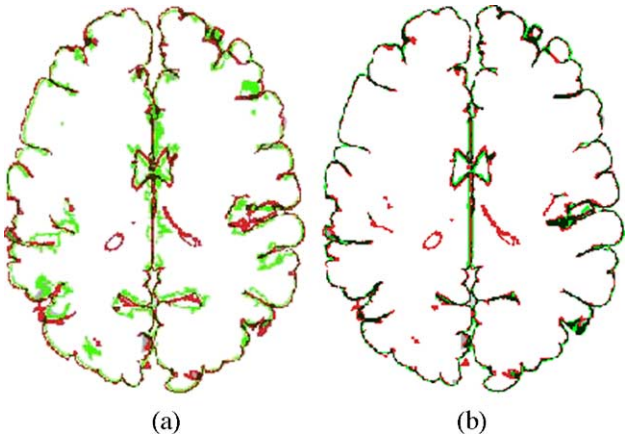

Fig. 9. Surface-based evaluation: (a) the conventional ASP, (b) the CLASP. The red line is the outer cortex of the phantom, while the green line is the surface re-extracted from the simulated MRI of the phantom. Voxels of which the GM surface by ASP or CLASP fits the phantom are shown as black color. Note the improved match between the red line and the green line.

Evaluation of repeatability

This section examines the reproducibility of the whole procedure including classification and surface extraction. In order to quantify the normal variance expected for reconstructing cortical surfaces, 16 different T1 MRIs with 1 mm isotropic sampling were acquired from the same subject over a short period of time (Holmes et al., 1998) and were averaged to create a template image. The pipeline for reconstructing cortical surfaces as described in this paper was run on each acquisition.

The root mean square (RMS) distance between the cortical surface of the averaged template and each surface for the 16 MRIs were measured (Fig. 10). Averages RMS errors from the 16 scans were 0.59 mm, 0.85 mm, 0.61 mm, and 0.66 mm for the GM surface of CLASP, the GM surface of ASP, the WM surface of CLASP, and the WM surface of ASP, respectively. Compared with the ASP algorithm, the CLASP algorithm reduced the RMS error in the repeatability test by 31% and 8% for GM and WM surfaces, respectively.

Qualitative demonstration

One of the major applications for cortical reconstruction is to measure cortical thickness. We proposed the method to find buried sulcal CSF voxels using partial volume estimation. This approach might result in regionally different measurement of cortical thickness. This can be seen by the distance between adjacent regions on the WM and pial surfaces shown in Fig. 11. The GM and WM surfaces of CLASP are properly fitted on the GM boundary while the GM surface of ASP misses the boundary between CSF and GM and even intersects the WM surface.

Discussion and summary

We have described a novel Laplacian-based modification of the ASP surface extraction algorithm, which we have shown to be more accurate at finding the boundaries of the cerebral cortex compared with the original ASP method. The main difference is

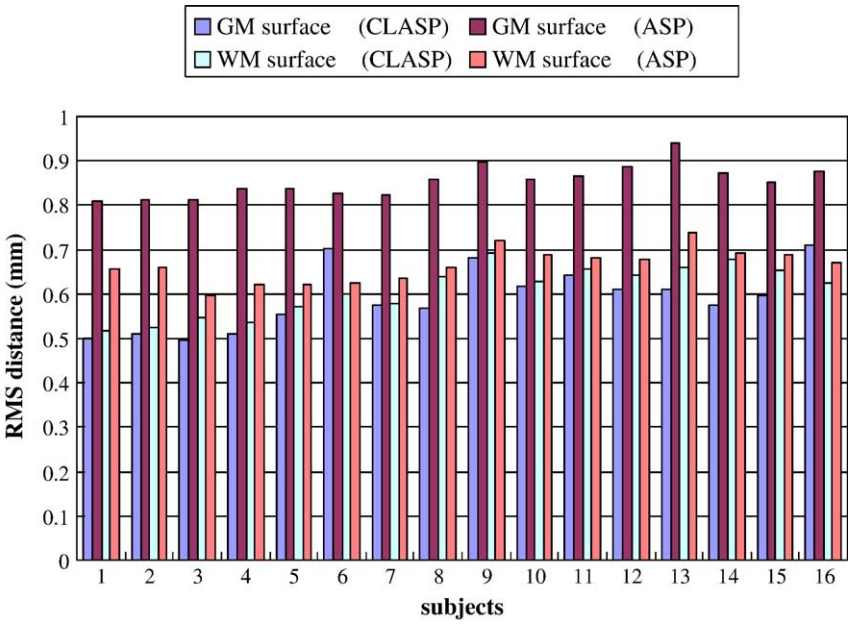

Fig. 10. RMS of distance between the surface of an averaged MRI and each surface of 16 scans. Note the reduce RMS errors for the CLASP surfaces.

that the algorithm makes use of a partial volume classification to insure topology preservation. Another difference is that it classifies subcortical structures with SPAMs to reconstruct the insular region properly. The accurate estimation of the outer surface of the cerebral cortex is important in measuring cortical thickness, but is made difficult because of partial volume effects. The conventional ASP was based on a hard classified volume, and boundary voxels in the deep sulci were often classified as GM though the voxels have some fraction of CSF. ASP estimated the boundaries of the sulci by incorporating a boundary constraint and a vertex–vertex constraint which limited the allowable cortical thickness. However, cortical thickness varies considerably by age and disease and thus these thickness limitations needed to be relaxed. The CLASP algorithm introduced in this paper overcomes these problems by incorporating a PVE classification as input to guide the deformation model. The topology of the cerebral cortex is preserved because the pial surface is expanded using a Laplacian field where the upper boundary of the field represents the partial volume CSF surface. Using volume- and surface-based evaluations showed

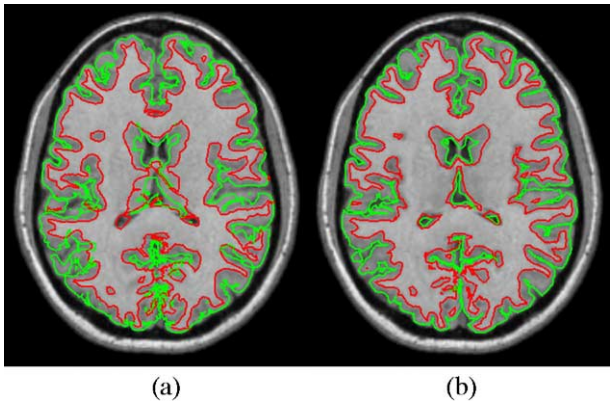

Fig. 11. Cortical surfaces of (a) ASP and (b) CLASP. There are intersections between WM and GM surfaces in ASP. Note that the surfaces of CLASP represent the cortical boundary better than ASP.

improved reconstruction of the outer cortical surface. Another improvement of this study was to use a Laplacian map with a thickness constraint. Gradient values of the Laplacian map express how thick GM is on the surface vertices. The vertices were expanded with different ratio of movement so that all vertices were able to reach the final pial surface at the same time.

CLASP does not require any assumptions about or constraints on cortical thickness as opposed to conventional ASP, so the algorithm is robust to variations of cortical thickness and topology. Moreover, CLASP still retains the following features from ASP: (1) the use of proximity constraints excluded self-intersecting surface configurations; (2) the stretch constraint had the polygon vertices evenly distributed; and (3) A 1:1 correspondence between polygon vertices on a sphere to the boundary surfaces, allowing efficient transformation and mapping of the vertices reciprocally.

It is hard to define a gold standard for surface reconstruction. Even comparisons with manually chosen landmarks in an MR image are affected by subjective decisions about which structures to choose and ambiguity in about the extent of buried sulci. We thus have introduced automated and objective methods for the evaluation. Although the volume-based and surface-based evaluation methods are affected by image resolution and the properties of the MRI simulator, they do provide a good basis to evaluate relative performance. In addition, the repeatability evaluation using 16 scans of one subject provides a relatively objective and unbiased evaluation of precision.

In summary, the intrinsic advantage of the CLASP algorithm compared with conventional ASP arises from its use of geometric information based on a probabilistic classification, allowing partial volume effects to be accounted for. This algorithm expands the pial surface from the WM surface and robustly reconstructs the surface while preserving the topology of the cerebral cortex. The reconstructed surfaces created by CLASP guarantee a more precise extraction and thus allows for a more detailed analysis of subtle cortical morphological features. This more general applicability of CLASP is important for studies of pediatric brain development and cortical degeneration in brain disease. The advantage of the

CLASP algorithm is to uses simple geometric constraints and is based on neuroanatomical knowledge, resulting in surfaces that provide a valuable model for morphological experiments.

In the future, the comparison of CLASP to other cortical surface reconstruction methods to assess its strengths and weaknesses in terms of measuring cortical thickness, functional brain mapping, visualization, surface labeling, and other applications would be of great utility.

Acknowledgments

This work was supported by the Post-doctoral Fellowship Program of Korea Science and Engineering Foundation (KOSEF).

References

Borgefors, G., 1984. Distance transformations in arbitrary dimensions. Comput. Vis. Graph. Image Process. 27, 321–345.

Choi, H.S., Haynor, D.R., Kim, Y., 1991. Partial volume tissue classification of multichannel magnetic resonance images—A mixed model. IEEE Transact. Image Imaging 10 (3), 395–407.

Cocosco, C.A., Zijdenbos, A.P., Evans, A.C., 2003. A fully automatic and robust brain MRI tissue classification method. Med. Image Anal. 7 (4), 513–528.

Collins, D.L., Peters, T.M., Evans, A.C., 1994a. An automated 3D non-linear deformation procedure for determination of gross morphometric variability in human brain. SPIE, 2359 (Rochester).

Collins, D.L., Neelin, P., Peters, T.M., Evans, A.C., 1994b. Automatic 3d intersubject registration of MR volumetric data in standardized Talairach space. J. Comput. Assist. Tomogr. 18 (2), 192–205.

Collins, D.L., Zijdenbos, A.P., Kollokian, V., Sled, J.G., Kabani, N.J., Holmes, C.J., Evans, A.C., 1998. Design and construction of a realistic digital brain phantom. IIEEE Trans. Med. Imag. 17 (3), 463–468.

Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis 1: segmentation and surface reconstruction. NeuroImage 9, 179–194.

Evans, A.C., Collins, D.L., Milner, B., 1992. An MRI-based stereotactic atlas from 250 young normal subjects. J. Soc. Neurosci. Abstr. 18, 408.

Evans, A.C., Collins, D.L., Mills, S.R., Brown, E.D., Kelly, R.L., Peters, T.M., 1993. 3D statistical neuroanatomical models from 305 MRI volumes. Proc. IEEE-Nuclear Science Symposium and Medical Imaging Conference, pp. 1813–1817.

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

Han, X., Xu, C., Prince, J.L., 2001. A topology preserving deformable model using level sets. IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR2001), pp. 765–770.

Han, X., Xu, C., Tosun, D., Prince, J.L., 2001. Cortical surface reconstruction using a topology preserving geometric deformable model. 5th IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA2001). , pp. 213–220.

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

Holmes, C.J., Hoge, R., Collins, L., Woods, R., Toga, A.W., Evans, A.C., 1998. Enhancement of MR images using registration for signal averaging. J. Comput. Assist. Tomogr. 22 (2), 324–333.

Jones, S.E., Buchbinder, B.R., Aharon, I., 2000. Three-dimensional mapping of the cortical thickness using Laplace's equation. Hum. Brain Mapp. 11, 12–32.

Kimia, B.B., Siddiqi, K., 1996. Geometric heat equation and nonlinear diffusion of shapes and images. Comput. Vis. Image Underst. 64 (3), 305–322.

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

Kwan, R.K.S., Evans, A.C., Pike, G.B., 1996. An extensible MRI simulator for post-processing evaluation. Lect. Notes Comput. Sci. 1131, 135–140.

Lorensen, W.E., Cline, H.E., 1987. Marching cubes: a high resolution 3-d surface construction algorithm. Comput. Graph. Quart. Rep. SIGGRAPH-ACM 21 (4), 163–169.

Ma, C.-M., Wan, S.-Y., 2001. A medial-surface oriented 3-d two-subfield thinning algorithm. Pattern Recogn. Lett. 22, 1439–1446.

Ma, C.-M., Wan, S.-Y., Lee, J.-D., 2002. Three-dimensional topology preserving reduction on the 4-subfields. IEEE Trans. Pattern Anal. Mach. Intell. 24 (12), 1594–1605.

MacDonald, D., 1998. A method for identifying geometrically simple surfaces from three-dimensional images, PhD thesis, School of Computer Science, McGill University, Montreal, Canada.

MacDonald, D., Avis, D., Evans, A.C., 1994. Multiple surface identification and matching in magnetic resonance images. In Proc. SPIE 2359, 160–169.

MacDonald, D., Kabani, N., Avis, D., Evans, A.C., 2000. Automated 3-D extraction of inner and outer surfaces of cerebral cortex from MRI. NeuroImage 12, 340–356.

Press, W.H., Flannery, B.P., 1988. Numerical recipes in C. Cambridge University Press, Cambridge.

Shattuck, D.W., Leahy, R.M., 2001. Graph-based analysis and correction of cortical volume topology. IEEE Trans. Adv. Packaging 20 (11), 1167–1177.

Sled, J.G., Zijdenbos, A.P., Evans, A.C., 1998. A non-parametric method for automatic correction of intensity non-uniformity in MRI data. IEEE Trans. Med. Imag. 17, 87–97.

Talairach, J., Tournoux, P., 1988. Co-planar stereotaxic atlas of the human brain. Thieme Medical Publishers, New York.

Thirion, J.P. and Gourdon, A., 1993, The marching lines algorithm: new results and proofs. Technical Report 1881-1, Intitute National de recherche en informatique et en automatique.

Tohka, J., Zijdenbos, A.P., Evans, A.C., 2004. Fast and robust parameter estimation for statistical partial volume models in brain MRI. Neuroimage 23 (1), 84–97.