Supervised Labeling of Brain Sulci using Multidimensional Scaling

Meena Mani
Visages Lab
IRISA INRIA Rennes
memani@irisa.fr

Abstract

In the last 15 years, two strategies have emerged to address the variability of sulci in labeling problems. The first is the use of a spatial distribution model of sulci to match candidate brains; the second is the explicit or implicit use of graphs to match the linkages between sulci. We base this paper on the second idea, using multidimensional scaling (MDS) to recreate the structural relationship between sulci in an embedded space. We implement this in a supervised learning framework where unlabeled sulci, projected into the atlas space using an out-of-sample procedure, are given class designations based on nearest neighbor search. We take a heuristic approach in selecting the distance measure for the MDS distance matrix. Preliminary evaluation showed that a pure shape feature distance cannot discriminate sulci adequately but that distances based on spatial features give encouraging results. We evaluate 3 spatial distances designed to accommodate the full spatial variation (position/orientation/scale). In experiments using a leave-one-out strategy, 90% of the 180 sulci drawn from 10 sulcal classes are successfully classified as are data from 6 patients with cortical tumors.

1 Introduction

The cortical surface is characterized by alternating ridges and furrows, the gyri and sulci (Figure 2). Internal changes in the brain, either due to aging or pathology, alter the cortical surface. Quantifying and studying these changes is a first step in the differential diagnosis of disease. Image-guided neurosurgery, where the sulci function as landmarks or provide pathways to access internal tissue, is another application that would benefit from a systematic study of gyri and sulci.

In this paper, we describe an automatic method to label sulci across subjects, an important first step in this study. This is a difficult problem since cortical sulci are highly variable not just across individuals but even between the hemispheres of a brain [1]. They vary in shape, scale, placement and branching morphology. They may be continuous (present as one uninterrupted segment) in some individuals, fragmented (exist as multiple segments) in others and altogether absent in yet others. This poses a problem for feature selection and classification. Figure 1 illustrates how the variability can make feature selection difficult.

19th century illustrations such as those from Horsley [2], trace the wide variations along a sulcal fold. A whole nomenclature has developed since then to account for the branch variations possible along a single sulcus. While accurate sulcal identification can be a challenge even for expert neuroanatomists, there are sulci that are to some degree more consistent, and for which anatomical correspondence can be established across subjects. These are the larger primary sulci (Figure 2(a)). The localization of these sulci allows us to generate a probabilistic map which can be used to label candidate sulci. A graph can also be constructed and unlabeled sulci (or the more variable secondary

and tertiary sulci) can be identified against this reference. These two ideas have been incorporated into automated and semi-automated labeling methods in various ways.

A straightforward implementation of the probabilistic atlas paradigm can be seen in Le Goualher et al. [3, 4]. Statistical probabilistic anatomical maps (SPAM) [5] give the probability for each sulcal class so that at any given location, unlabeled sulci are assigned the most probable label for that location. A different spatial distribution model is used by Lohmann et al. [6]. A *point distribution model* computes the shape of sulcal basins across a training set. Any unlabeled sulcus can be expressed as a linear combination of the eigenvalues generated from the PCA of this shape covariance matrix; an optimization over the linear function would give the best label.

Spatial distribution models give spatial bounds but this is not adequate to discriminate between the sulci in a local region. They are usually combined with graphs which model connections between sulci thus giving local context. In the combined strategy, the spatial information is used to supply spatial priors [7, 8], localization constraints or to narrow the search space in an optimization or matching process [9].

An example of a graph approach, where a model graph is matched to a random graph of sulci from a test subject, is presented in Rivière et al. [10]. Sulcal features, represented as the node attributes and neighborhood relations modeled as the edges of a graph are matched by training a set of neural networks to optimize the potential function at each graph node. Yang and Kruggel [9] also apply graph matching to labeling. A similarity function that uses statistical information about spatial extent, orientation, shape and neighborhood structure is optimized using a genetic algorithm.

Mechouche et al. [11] use graph-matching in a very different way. An atlas represented via a semantic ontology is matched to pre-segmented candidate brains using descriptive logic. The main engine for the matching is a constraint solver (CSP), software that uses advanced numerical techniques that are computationally expensive.

Distinct from these approaches, deformable atlas matching [12], a method widely used in brain imaging to deform a template to a candidate brain, has also been been applied to the sulcal labeling problem. Elastic or fluid deformation [13], in general, ensures a close match to local regions, but directly applying these algorithms to the cortical surface with its complex folding patterns which vary from subject to subject, may give many local optima [10, 3, 6, 14]. Some authors have sought to adapt this method to the cortical surface by, for instance, first aligning sulcal landmarks which act as surface constraints [15] or using a multi-resolution strategy which matches the lengths of sulci at different scales [16].

The learning, graph and warping approaches typically use large feature sets and computationally expensive algorithms. In [17], 8000 features at different scales are used to learn likelihoods to generate a probability map. Under a Bayesian formulation, the best curve is computed by dynamic programming using the probability map and shape priors. In [18], a large sample space of candidate curves is generated to learn shape priors which are used to define the potential functions of a graphical model. In [10], once again, a large number of training features are generated by oversampling the sulcal segments to be identified. Yang and Kruggel [9] set up graph matching as a computationally expensive combinatorial optimization problem.

In this paper we present a simpler approach that retains the idea of a graph, but as a global representation of sulcal class data. Structural patterns of a distance graph act as a reference against which unlabeled sulci can be matched. An example of a distance graph is shown in Figure 2(b).

Multidimensional scaling (MDS) is a natural choice for this implementation. It gives a geometric interpretation which can be used to reproduce the structural relationships between sulci in a low dimensional space. We can thus build a reference atlas against which unlabeled sulci, using the same relational constraints, can be identified using nearest neighbor criterion (NN).

An important consideration with graph-based methods such as MDS, is the distance measure. In preliminary work [19] we used a geometric shape representation for open curves. The results indicate that shape distance, i.e. the geodesic distance between sulcal curves, is not a good feature for our classification scheme. A pure shape feature is obtained after removing rotation, translation and scaling. As part of our heuristic approach, we looked at each of these separately to determine their effects. We found that the position-based feature, which complements the spatial structure inherent

in our pattern matching approach, gave very promising results. Here, we extend this study to three position based distance measures.

This paper is organized as follows: In Section 2 we explain our classification method which uses MDS and NN for labeling the major sulci. In Section 3, we describe the distance measures used. Section 4 has the experimental details and results. In the discussion that follows, Section 5, we consider the performance on both healthy and pathological data and offer insights on the success of our classification scheme.

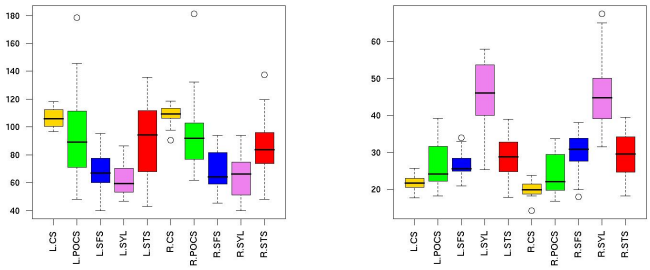

(a) length boxplot (b) depth boxplot

Figure 1: **The variability in sulcal features** is seen here in the length and depth boxplots for 18 subjects. Of the 10 types of sulci shown, only the sylvian fissure (in violet, see (b)) can be identified by a single feature, its depth measurement. This is not surprising since the deep sylvian fissure is one of the most easily identifiable parts of the cortical surface. The other classes of sulci cannot be separated out solely on a length or depth measurement.

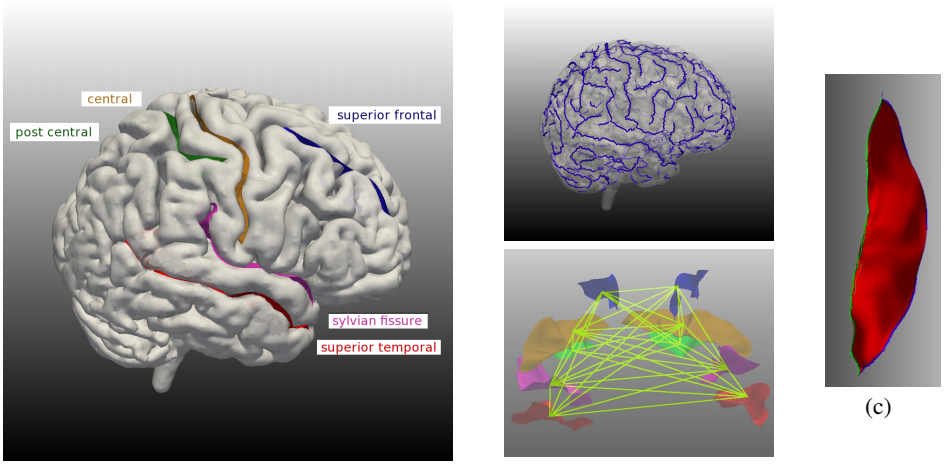

(a) (b) (c)

Figure 2: (a) The five major sulci used in this study. (b) Top: Sulcal curves, in blue, form a complex and widely varying pattern on the cortex. Bottom: Graph representation where the distances are the edges and the sulci are the nodes. (c) An individual sulcus. The top curve, which points towards the outer cortex, is in blue. The bottom is green.

2 Methods

Multidimensional scaling (MDS) gives a geometric interpretation which can be used to reproduce the structural relationships between sulci in a low dimensional space. We can thus build a reference atlas against which unlabeled sulci, using the same relational constraints, can be identified using nearest neighbor criterion (NN). The approach we take to classify unlabeled sulci, falls under the ambit of supervised learning. The steps are as follows:

1. First, using multidimensional scaling, we embed a map using dissimilarity data computed from a database of labeled sulci.
2. Next, once again using MDS, unlabeled sulci are introduced into the embedded space. Since this is an *out-of-sample* procedure, the original embedding remains unchanged.
3. Finally, the unlabeled points are assigned using nearest neighbor classification.

In describing this combination of steps as supervised learning, we would like to emphasize two points: First, each of the labeled nodes in the embedded space represents a class average rather than the individual data points of a class. MDS is not used as a precursor to clustering, but only to provide a common space to place labeled and unlabeled points. Second, the reference MDS map is generated solely from labeled sulci and not from a combination of labeled and unlabeled sulci as is done in semi-supervised learning. Each of the components of the labeling scheme is discussed below.

2.1 Classical Multidimensional Scaling (CMDS)

Multidimensional scaling is a method that allows us to compute an optimal configuration, X, for a set of n observations using only their interpoint distances. If $\Delta = [\delta_{ij}]$ is a non-negative, symmetric matrix of dissimilarities computed from a set of (hidden) coordinates, X^d of rank d, with MDS, we can represent X^d as X^p, $p \leq d$, such that the distance matrix D computed from it is close to Δ. An intermediate step would be to derive the scalar product matrix, $B = XX^T$, from the squared dissimilarities by double centering. An eigen decomposition, $B = \Gamma\Lambda\Gamma^T$, gives us Λ, a diagonal matrix of eigenvalues, and Γ, a matrix of orthonormal eigenvectors. X, the square root of XX^T, is then $\Gamma\Lambda^{\frac{1}{2}}$. Since B is of rank p, X is also of rank p. Thus, we see that in CMDS, the coordinates, X, of the embedded configuration are implicit and need to be extracted from the Euclidean inner product matrix B. The eigenvalue problem solves the optimization:

$$\min_{x_i \in \mathbb{R}^p} \sum_{i=2}^{n} \sum_{j=1}^{i-1} (\|x_i - x_j\|^2 - \delta_{ij}^2). \tag{1}$$

2.2 Out-of-sample Embedding

For classification, our objective is to introduce k unlabeled sulci, $y_1,\ldots,y_k$, without disturbing an existing configuration. Using CMDS to embed $n + k$ points, however, gives an entirely new map computed from inner products which are centered with respect to the centroid of the $n + k$ points. The optimization we seek instead is:

$$\min_{y \in \mathbb{R}^d} 2 \sum_{i=1}^{n} \sum_{j=1}^{k} (\|x_i - y_j\| - a_{i(n+j)})^2 + \sum_{i=1}^{k} \sum_{j=1}^{k} (\|y_i - y_j\| - a_{(n+i)(n+j)})^2, \tag{2}$$

where $A = [a_{ij}]$ is the augmented $(n + k) \times (n + k)$ dissimilarity matrix. We can get y_i, the coordinates for the k points with this while the original labeled points, x_i, which were previously computed, remain fixed. This is an *exact solution* for the out-of-sample problem for CMDS but it is a nonlinear optimization problem. If, though, the second term of the objective function above is dropped, the resulting function is convex. Solutions to this give an approximate out-of-sample embedding. The treatment is detailed in Trosset and Priebe [20].
We use Equation 2 later, in the context of a leave-one-out strategy to assess the utility of different data features in terms of generalisation error (as assessed with cross validation). In this instance the x correspond to the training set and the y the test set.

2.3 Nearest Neighbor Classification

For unclassified points, embedded in the out-of-sample step described above, class membership is determined based on their proximity to the class nodes. A minimum euclidean distance criterion is used.

In the nearest neighbor framework, we can also compute, as a confidence measure, the relative probability of an unlabeled sulcus being assigned to a class. Let us consider the entire embedded space as a Voronoi tesselation with a training node at the center of each cell. A *softmin* function over distances in this space will give a continuous probability that decreases the further a point is from a node. If d is the euclidean distance between an unlabeled point i and a node j, for n classes, the relative probability with which an unlabeled point selects a labeled node as its nearest neighbor is:

$$p_{ij} = \frac{e^{\frac{-d_{ij}^2}{\sigma^2}}}{\sum_{k=1}^{n} e^{\frac{-d_{ik}^2}{\sigma^2}}}. \tag{3}$$

The variance, σ, can be heuristically determined from training set data. Such computations cannot be made through a spectral embedding alone.

3 Feature Set

An important consideration with graph-based methods such as MDS, is the distance measure. In preliminary work [19] we evaluated position, shape, orientation, length and mean depth. We found that the position-based feature, which complements the spatial structure inherent in our pattern matching approach, gave very promising results. Here, we extend this study to three distance measures based on the position distance.

3.1 Distance Measures Evaluated

We evaluate three pairwise spatial distance measures between individual sulci. Since there are many different configurations to consider, with curves at various angles and of varying length, a measure that is robust to this variability will give the best results.

1. **Distance from mid-point**: The two ends of a sulcus were used to determine a mid-point (not usually on the sulcus). Pairwise distances were then computed from the mid-points of two sulcal curves.
2. **Mean closest point (MCP)**:

$$\begin{aligned} d_{MCP} &= \text{mean}(d_m(C_i, C_j), d_m(C_j, C_i)) \\ \text{where } d_m(C_i, C_j) &= \textbf{mean}_{i \in C_i} \min_{j \in C_j} d(i, j) \end{aligned}$$

   This distance (d_{MCP}) is described in [21, 22]. For each of n equally spaced points along curve C_i, the minimum distance to a point in curve C_j is measured. d_m is the mean of the set of these minimum distances. Since this is a directed distance, $d_m(C_i, C_j) \neq d_m(C_j, C_i)$ in general. The symmetric distance is found by computing the average of $d_m(C_i, C_j)$ and $d_m(C_j, C_i)$. See Figure 3.
3. **Median closest point (NCP)**: This is similar to the mean closest point, but we compute a median for the set of minimum distances instead of a mean.

3.2 Preliminary Features Evaluated

In our preliminary evaluation [19], in addition to the position feature, we used four other descriptors, pure shape, mean depth, length and orientation. These contain physical information not captured by the relative position of a 3D sulcal curve alone. The features, their class averages and the distances between the class averages are summarized below. It is these distances that we used as the distance feature in our characterization.

1. **Elastic Shape of a Curve** We use the square-root velocity description for shapes of 3D sulcal curves: $q(t) = \frac{\dot{\beta}(t)}{\sqrt{\|\dot{\beta}(t)\|}}$ where $\beta : [0, 1] \rightarrow \mathbb{R}^3$ is a unit-length, parameterized curve such that its speed $\|\dot{\beta}(t)\|$ is non-zero everywhere.
2. **Orientation** Using singular value decomposition (SVD), the relative orientations, O, between sulcal curves were computed for each subject.

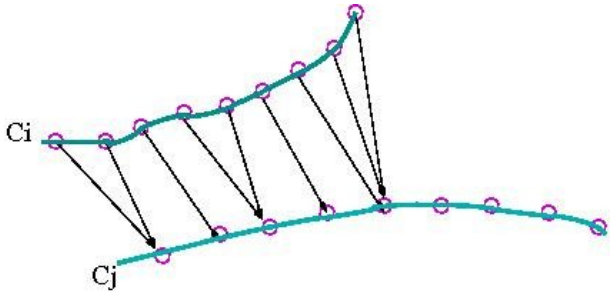

Figure 3: The set of *n* minimum distances from *Ci* to *Cj* used in the mean and median closest point computation. This is a directed distance; the symmetric distance is used in the distance matrix computations.

3. **Length** The length of a sulcus was computed using $\int_0^1 \|\dot{\beta}(t)\|dt$, the length of the parameterized open curve β.
4. **Mean Depth** The depth profile for each sulcus was first obtained by resampling 100 points along the length of the sulcus and computing the distance between points on the top sulcal curve with corresponding points on the bottom curve. The mean depth was then calculated by averaging these. The top and bottom curve are shown in Figure 2(c).

Class Averages
The class average for each of these descriptors was calculated by averaging across the subjects. For the relative orientation, the average was obtained by computing a SVD of $\sum_i O_i$. The Karcher mean gave us the mean shape.

Distances
The distance matrix was computed using the feature distances between two sulci. For position we computed a euclidean distance. For depth and length, we used an absolute distance. For orientation, we used $d_o = \frac{\|\log(O_1 O_2^T)\|}{\sqrt{2}}$, where $\|.\|$ is the Frobenius norm. For shapes, the geodesic distance was used.
The composite distance between any two sulci was also computed by using a combination of the scaled distances, $d = w_1 d_s + w_2 d_p + w_3 d_d + w_4 d_l + w_5 d_o$, where the w_i could be chosen by trial and error.

Table 1: Feature set used in preliminary analysis

	Descriptors	Class Averages	Distance features
1	Spatial Position		euclidean distance
2	Shape	Karcher mean	geodesic distance
3	Orientation	SVD($\sum_i O_i$)	$\frac{\|\log(O_1 O_2^T)\|}{\sqrt{(2)}}$
4	Length		$\|l_1 - l_2\|$
5	Mean Depth		$\|\bar{d}_1 - \bar{d}_2\|$

(1) Training set represented by class averages
(2) Individual unlabeled sulci constitute test sulci
(3) Distance matrix computed using feature distances between two sulci

4 Data Analysis

To assess the capability, i.e. the performance and limitations of the method we propose, we evaluated the following:
1) Preliminary analysis of spatial position, shape, orientation, length and depth
2) Evaluation of spatial position distances

3) Evaluation of spatial position distances with a larger data set
4) Partial sulci, i.e., sulci that are split into segments
5) Tumor data
Clustering, an alternative method, has also been used to group the data. The results of the two approaches are discussed in Section 5.

4.1 Data and Initial Processing

Subjects: We used a database of 18 T1-MR 3D SPGR brain images of healthy subjects, matched for sex (male), handedness (right), and age (35 ± 10 years).
Sulci used: Five major sulci, the central, postcentral, superior frontal, sylvian fissure and superior temporal, (shown in Figure 2(a)), were extracted from each hemisphere. This gave us $18 \times 10 = 180$ sulci for classification.
Segmentation: The sulci were extracted using the *active ribbon* method [4, 23]. A 1D curve, skeletonized on the superior surface of the sulcal groove, evolves through the depth of the groove tracing out a 2D median surface in the process. This algorithm has been adapted from the active contour model [24].
Final form of the data: Performance was evaluated on two separate classifiers, one constructed with the top sulcal line and the other with the bottom line of a sulcus (see Figure 2(c)).

4.2 Evaluating Dimensionality, Assessing Fit

MDS reproduces a high dimensional input in a lower dimensional space. In Figure 4, an *elbow* in the scree plot at $d = 3$, represents the optimal dimension beyond which only minor reductions in stress are acheived. The Shepard diagram, which plots the scaled MDS distances against the original dissimilarity data, was then used to assess fit for 2 and 3D data. A good fit is indicated by a minimal spread in the scatter plot and evidence of a monotonically ascending relation. A comparison between the two Shepard plots, shows that the output distances are highly correlated to the input dissimilarities for the 3D data. As a final test, the MDS map we obtain, reflects the actual arrangement of sulci in the brain.

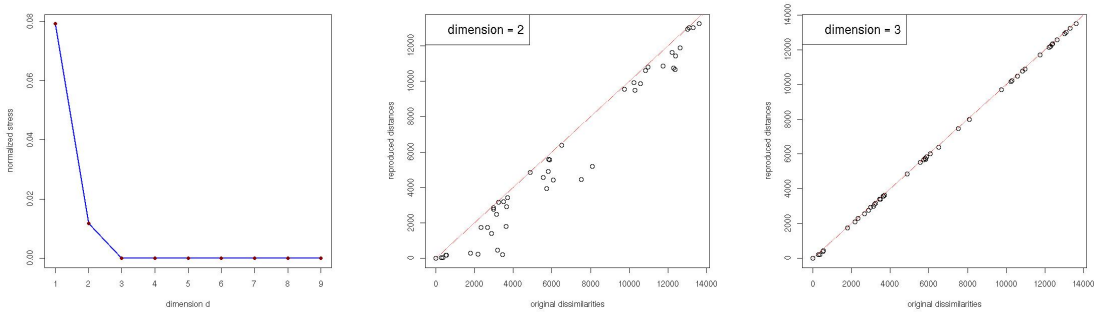

Figure 4: Evaluating the number of dimensions required to represent the data. The scree plot (left), shows that 3 dimensions is adequate for the MDS map. The Shepard plot for 2D data (middle) shows more spread than the corresponding 3D plot (right).

4.3 Classification

Our goal was to label pre-segmented sulci without making assumptions about predefined relationships or using semantic information. For this we designed a 10-class MDS classifier. We evaluated the classification using leave-one-out (LOO) cross-validation. The 18 sulci for each of the 10 sulcal classes were split 17:1 into a training set and a testing set.

4.3.1 Preliminary Characterization

The best results for the 18 LOO iterations were obtained when spatial distances were used for the dissimilarities. Classification was also done for shape, using geodesic paths for distances, and for scaled combinations of shape, position, length, depth and orientation. The shape classifier labeled all the central sulci, which are very regular, but was less successful with the other classes. Combining the features did not improve the success rate. Rather, the results achieved by the spatial distance classifier were compromised (Figure 5). The depth and length classifiers could not discriminate right and left brain sulci, and while the orientation information improved results in cases where the orientations of the sulci differed, these were offset by the errors introduced using an *average* orientation as a feature.

4.3.2 Spatial Distances

For the distance measured from the mid-point of the sulcal curve, we got an overall true positive rate of 90% and 90.6% for the top and bottom sulcal classifiers respectively (Table 2), and were able to correctly label all the superior frontal sulci (100%). The central and postcentral sulci run parallel to each other as do the sylvian fissure and superior temporal sulcus. These were sometimes misclassified, one for the other.

The results for the mean(d_{MCP}) and median(d_{NCP}) closest point distance classifiers were 86% and 84% respectively (Figure 6). The MCP and NCP calculations underestimate distances when two sulci are perpendicular. It is perhaps because of this uneven application of distances that we got poorer results.

4.3.3 Partial Sulci

While the central sulcus and sylvian fissure are highly continuous [1], some of the other large sulci are interrupted by buried gyri or split into 2 or 3 pieces as an artifact of the segmentation. To simulate this, we split the test sulci into two at the mid-point. Of the $180 \times 2 = 360$ pieces tested using LOO cross-validation with the mid-point distance classifier, 77% were correctly labeled. This is in line with published results [10]. The MCP and NCP classifiers performed poorly here.

4.3.4 Tumor Data

The algorithm was tested on data from 6 patients with medium to large tumors (see Figure 9). The lesions had displaced sulci or had otherwise altered the arrangement of the surrounding cortical tissue. The sulci in the unaffected regions were successfully labeled, i.e. they had the same error rate as for the healthy subjects. The value in our method lies in the fact that we were also able to classify sulci in the region surrounding the tumor.

5 Discussion

Our supervised MDS spatial distance classifier offers several advantages.

First, it is a fast and lightweight implementation of CMDS and NN search. CMDS is an $O(N^3)$ time algorithm. For NN, we need to search through the entire training set to classify each new point. Both of these operations are computationally intensive. By computing class averages in the feature space, we get a compact representation of the training data, thereby significantly reducing the computational load.
Also, because we use an out-of-sample strategy, we further limit the number of computations. For the training distance matrix, we need only make $\frac{n(n-1)}{2}$ one-time measurements, where n is the number of classes. These, along with the MDS reference map are computed off-line. The only runtime computations are those that involve unlabeled sulci.

Second, our results suggest that we can very quickly assign a given sulcus to a small region of the left or right hemisphere. 179 out of 180 sulci were correctly identified when we relaxed the classification criteria to include the second nearest sulcus. For most of the misclassifications, the first and second nearest sulci had almost the same probability of being assigned to the correct class (see Section 2.3).

As an alternative to our method, we consider clustering, a method which is also often combined with MDS. When we tried to group the sulci using spectral clustering, 123/180 (68%) sulci were correctly classified (see Figure 8(a)). Clustering may not be the best choice for sulcal data since distance-based clustering algorithms do not use the structural pattern information from the graphs. The post- and precentral sulci, for instance, both equidistant and parallel to the central sulcus but on opposite sides, group together. Also, relying solely on distances when the data is inherently noisy could also lead to incorrect grouping. Spectral clustering has been successfully applied to DTI data [22], for which, the large number of fibers warrant the use of an unsupervised or semi-supervised approach. Sulci, by contrast, are orders of magnitude fewer in number.

Third, the classifier is robust to normal population variation. The results we obtained, were for data that had not been spatially normalized and thus representative of the full sample variation. Also, the results for both the variable top and stable bottom sulci, each independently classified, were comparable (Table 2).

The relational paradigm allows us to identify anatomical regions for which, because of inherent variability, a feature set is difficult to select. We were even able to correctly identify displaced and distended sulci, taken from tumor patients. The preliminary success with labeling tumor datasets strongly suggests that we can design a general-purpose tool, i.e. one that can be used to identify both normal and pathological sulcal data. In general, diagnostic tools seek to differentiate healthy tissue from disease–tumor detection being an example. Here, we have demonstrated a method which can *unify* class data for applications where it is important to do so.

The classifier works best when only a small number of well separated sulcal classes are used in the training set. We saw the performance degrade from 90.6% to 84.3% when two additional classes, the left and right precentral sulcus were added (Figure 7). This is a limitation of the current design. There are ways to extend the classification however. We can, for instance, train with different sets of small, well-separated sulcal classes. There are many ways in which these individual results can be combined to influence the final solution.

6 Concluding Remarks

In this paper, we describe a novel scheme to label sulci which uses MDS in a supervised learning framework. Our method is conceptually simple and easy to implement, using for the most part standard, off-the-shelf algorithms. It is computationally fast because we use a compact representation of data and an out-of-sample procedure. It also requires no registration allowing us to use standalone sulci from a database. We considered several distance measures in the preliminary evaluation; some, like the scale distance, have not been widely used in medical imaging applications. Our method has the flexibility to be used both for healthy and diseased data. While there is a need for such an application, the conventional approach has been to design tools to differentiate rather than unify anatomy on the basis of group differences. Finally, our method gives impressive results, accurately labeling 91% of the 180 sulci we tested.

By training with different sets of sulci, we can extend the classification and this will be the direction for future work. Another direction for future work is to apply this classification tool in different contexts using it to label white matter fibers bundles and other anatomical structures.

Table 2: Confusion Matrix for Spatial Distance LOOCV (for mid-point distances).
The true positive rate (blue), gives the % of the sulci correctly identified in 18 LOO tests. The off-diagonal terms (red) give the % of false negatives.

Top Sulcal Curve

True Label		Test Label 1	2	3	4	5	6	7	8	9	10
Left central	1	94.5	16								
L. postcentral	2	5.5	84								
L. sup. frontal	3			100							
L. sylvian fissure	4				78	11					
L. sup. temporal	5				22	89					
Right central	6						94.5	22			
R. postcentral	7						5.5	78			
R. sup. frontal	8								100		
R. sylvian fissure	9									84	16
R. sup. temporal	10									16	84

Bottom Sulcal Curve

True Label		Test Label 1	2	3	4	5	6	7	8	9	10
Left central	1	89	16								
L. postcentral	2	11	84								
L. sup. frontal	3			100							
L. sylvian fissure	4				84	11					
L. sup. temporal	5				16	89					
Right central	6						94.5	22			
R. postcentral	7						5.5	78			
R. sup. frontal	8								100		
R. sylvian fissure	9									94.5	5.5
R. sup. temporal	10									5.5	94.5

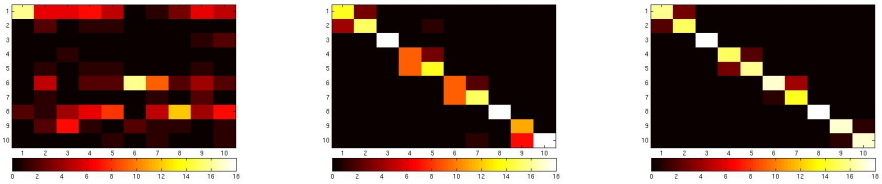

Figure 5: Confusion map for classification: the shape distance results (left) improve considerably when the shape and spatial distances (middle) are combined; the spatial distance alone (right) gives the best perfomance.

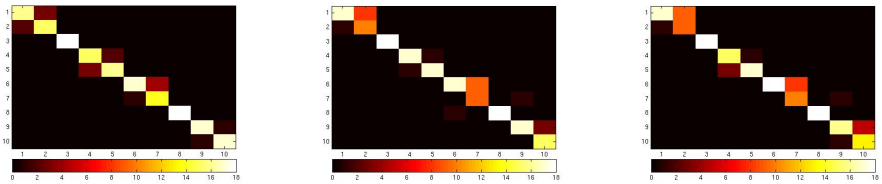

Figure 6: Confusion map for spatial distance classification: the mid-point distance classifier (left) gave the best results. The MCP (middle) and NCP (right) classifiers gave roughly similar results.

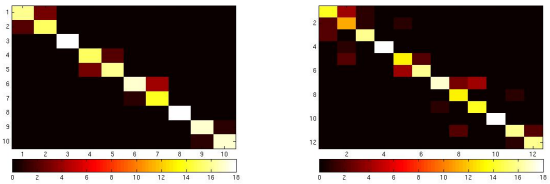

(a) 10 sulcal classes (b) 12 sulcal classes

Figure 7: Labeling 10 vs 12 classes: The results for the spatial distance deteriorated when the two precentral sulci were added, increasing the number of classes to 12. The precentral sulcus was sometimes miscategorized as the adjacent sylvian fissure, the central or superior frontal sulcus.

References

[1] Ono, M., Kubik, S., Abernathy, C.D.: Atlas of the Cerebral Sulci. Thieme Medical Publisher (1990)

[2] Horsley, V.: On the topographical relations of the cranium and surface of the cerebrum. In Cunningham, D.J., ed.: Contribution to the surface anatomy of the cerebral hemispheres. Royal Irish Academy (1892) 306–355

[3] Le Goualher, G., Collins, D.L., Barillot, C., Evans, A.C.: Automatic identification of cortical sulci using a 3D probabilistic atlas. In: MICCAI. (1998) 509–518

[4] Le Goualher, G., Procyk, E., Collins, D., Venugopal, R., Barillot, C.: Automated extraction and variability analysis of sulcal neuroanatomy. IEEE Trans. Med. Imag. **18**(3) (1999) 206–217

[5] Evans, A.C., Collins, D.L., Neelin, P., Kamber, M., Marrett, S.: Three-dimensional correlative imaging: Applications in human brain mapping. In Thatcher, R., Hallett, M., Zeffiro, T., John, E., Huerta, M., eds.: Advances in Functional NeuroImaging: Technical Foundations. Academic Press (1994) 145–162

[6] Lohmann, G., von Cramon, Y.: Automatic labeling of the human cortical surface using sulcal basins. IEEE Trans. Med. Imag. **4** (2000) 179–188

[7] Perrot, M., Rivière, D., Mangin, J.F.: Identifying cortical sulci from localizations, shape and local organization. In: 5th Proc. IEEE ISBI, Paris, France (May 2008) 420–423

[8] Perrot, M., Rivière, D., Tucholka, A., Mangin, J.F.: Joint Bayesian Cortical Sulci Recognition and Spatial Normalization. In: IPMI. (2009) 176–187

[9] Yang, F., Kruggel, F.: A graph matching approach for labeling brain sulci using location, orientation, and shape. Neurocomputing (2009) 179–190

[10] Rivière, D., Mangin, J.F., Papadopoulos-Orfanos, D., Martinez, J.M., Frouin, V., Régis, J.: Automatic recognition of cortical sulci of the human brain using a congregation of neural networks. Medical Image Analysis **6**(2) (2002) 77–92

[11] Mechouche, A., Morandi, X., Golbreich, C., Gibaud, B.: A Hybrid System Using Symbolic and Numeric Knowledge for the Semantic Annotation of Sulco-Gyral Anatomy in Brain MRI Images. IEEE Transactions on Medical Imaging **28**(8) (August 2009) 1165–1178

[12] Bajcsy, R., Broit, C.: Matching of deformed images. IEEE Conf. on Pattern Recognition (1982) 351–353

[13] Christensen, G.E., Rabbitt, R.D., Miller, M.I.: Deformable templates using large deformation kinematics. IEEE Trans. Imag. Proc. **5**(10) (1996) 1435–1447

[14] Hellier, P., Barillot, C.: Cooperation between local and global approaches to register brain images. In: IPMI. Volume 2082., Springer-Verlag (2001) 315–328

[15] Vaillant, M., Davatzikos, C.: Hierarchical matching of cortical features for deformable brain image registration. In: IPMI, Springer-Verlag (1999) 182–195

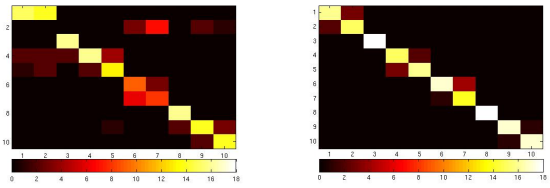

(a) Clustering (b) Labeling with MDS

Figure 8: (Spectral) clustering vs the MDS labeling method: With clustering (a), 121/180 (68%) sulci were correctly grouped (cluster label assignments were based on majority vote). With our MDS method (b), 163/180 (91%) sulci were correctly labeled. Clustering uses only distance information whereas our labeling algorithm uses both distance and structural pattern information.

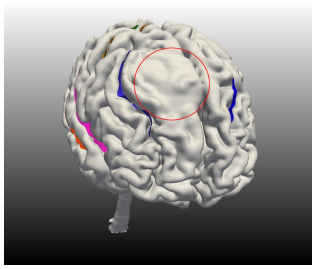

(a)right superior frontal sulcus

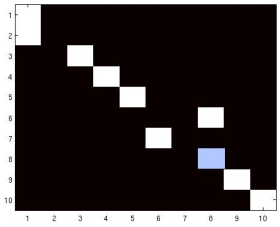

(b) displaced sulcus identified

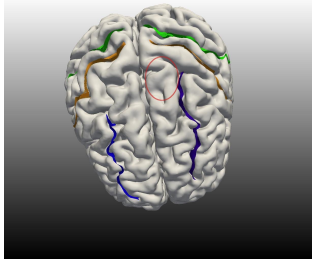

(c)left superior frontal & central

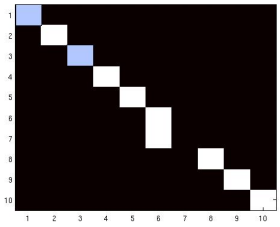

(d) displaced sulci identified

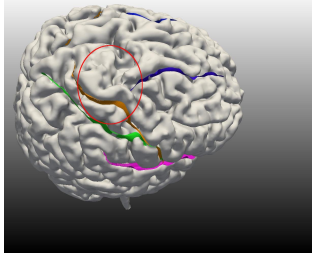

(e)right central & superior frontal

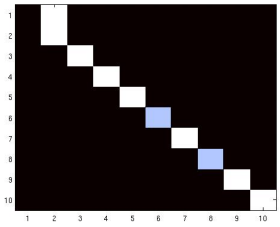

(f) displaced sulcus identified

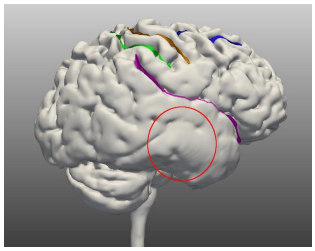

(g)right superior temporal

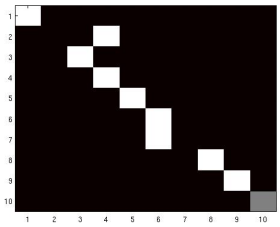

(h) no sulcus due to tumor

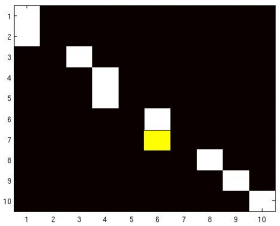

(i)displaced sulcus misidentified

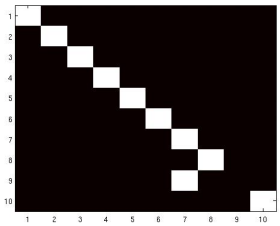

(j) deep tumor, sulcus unaffected

Figure 9: **Labeling sulci displaced by tumors for 6 subjects.** The affected sulci are shown on the left; the labeling results, on the right and bottom. The labeling error rate in the unaffected cortical regions is the same as for the healthy subjects. In the affected regions, the sulci are correctly labeled (shaded mauve) or do not exist (shaded grey) in 5/6 cases. It is misidentified in only one case, (i) (shaded yellow).

[16] Jaume, S., Macq, B.M., Warfield, S.K.: Labeling the Brain Surface using a Deformable Resolution Mesh. In: MICCAI. (2002) 451–458

[17] Tu, Z., Zheng, S., Yuille, A.L., Reiss, A.L., Dutton, R.A., Lee, A.D., Galaburda, A.M., Dinov, I., Thompson, P.M., Toga, A.W.: Automated Extraction of the Cortical Sulci Based on a Supervised Learning Approach. IEEE Trans. Med. Imag. **26**(4) (2007) 541–552

[18] Shi, Y., Tu, Z., Reiss, A.L., Dutton, R.A., Lee, A.D., Galaburda, A.M., Dinov, I., Thompson, P.M., Toga, A.W.: Joint Sulcal Detection on Cortical Surfaces with Graphical Models and Boosted Priors. IEEE Trans. Med. Imag. **28**(3) (2009) 361–373

[19] Mani, M., Srivastava, A., Barillot, C.: The Labeling of Cortical Sulci using Multidimensional Scaling. The Midas Journal (2008) http://hdl.handle.net/10380/1502.

[20] Trosset, M.W., Priebe, C.E.: The Out-of-Sample Problem for Classical Multidimensional Scaling. Computational Statistics & Data Analysis **52**(10) (2008) 4635–4642

[21] Corouge, I., Gerig, G., Gouttard, S.: Towards a Shape Model of White Matter Fiber Bundles Using Diffusion Tensor MRI. In: ISBI. (2004) 344–347

[22] O'Donnell, L.J., Westin, C.F.: Automatic tractography segmentation using a high-dimensional white matter atlas. IEEE Trans. Med. Imag. **26**(11) (2007) 1562–1575

[23] Le Goualher, G., Barillot, C., Bizais, Y.: Modeling cortical sulci with active ribbons. IJPRAI **11**(8) (1997) 1295–1315

[24] Kass, M., Witkin, A.P., Terzopoulos, D.: Snakes: Active contour models. International Journal of Computer Vision **1**(4) (1988) 321–331