Mindboggling Shape Analysis and Identification

Specific Aims

2.A. Long-term objectives

This proposal is motivated by two long-term research objectives involving brain morphology characterization and parcellation. By parcellation, we mean anatomical labeling of macroscopic anatomical features such as gyrii; parcels are also referred to as *regions of interest* (ROIs).

1. Characterize the shapes and sizes of brain structures and ROIs and their co/variations in healthy subjects and patients. This knowledge will be instrumental in understanding the relationship of macroscopic structure to microstructure, connectivity, physiology, and function for genetic, behavioral, developmental, and clinical research.
2. Establish a consistent, automatic, and fast method for parcellating brains based on probabilistic feature identification, with an accuracy and precision comparable to that of manual labeling, a method that is inconsistent, tedious, and very time-consuming.

2.B. Specific goals

The specific aim of this proposal is to automatically identify/match brain features based on a geometric and parametric analysis of their shapes, by means of a Bayesian framework derived from the face recognition literature. Mindboggle [1-3], a freely available software package for performing automated brain parcellation, will serve as the software infrastructure for implementing the Bayesian framework. The secondary aim is to further develop Mindboggle to automatically label an entire brain based on these probabilistic matches.

Features will be extracted from human brain MR image data by morphological image processing algorithms. Image processing will entail skeletonization of brain cortical folds using an exact medial axis method [4-6], combined with revised Mindboggle algorithms for fragmenting these skeletons. The shapes of these features will be analyzed using geometric and parametric approaches. Geometric analysis of shapes will include gross shape descriptors such as volume, extent, orientation, etc. Parametric analysis will employ quantitative shape discrepancy metrics derived by an active surfaces model [7,8]. These measures of similarity between shapes will be applied to a large dataset of manually labeled brain data and incorporated in a Bayesian framework for the purpose of estimating the probability of a given shape corresponding to a particular brain structure.

In section 3 (Background and Significance), we will give examples to convey the significance of identifying brain structures and parcellating brains. We will then mention parametric shape analysis, common parcellation methods and their evaluations, the use of multiple atlases, and recent related developments. These topics relate indirectly to the specific goals of this proposal by providing the background for, and significance of, brain shape analysis and identification.

In section 4 (Preliminary Studies), we outline the processing steps and evaluation of the original version of Mindboggle. This relates to the specific goals of shape analysis and identification by describing the context (software environment) within which this work will take place.

In section 5 (Research Design and Methods), we will describe our manually labeled data, the automated methods to analyze the shapes of structures extracted from those data, and how we will evaluate the identification of structures and ROIs based on these analyses. The specific goals of shape analysis and identification will be directly addressed by establishing shape descriptors, a Bayesian framework relating shapes to structures, and its incorporation into Mindboggle to identify these structures anatomically.

2.C. Ancillary contributions

In addition to software and algorithmic contributions, we will make other contributions to the biomedical community, including:

1. Manually edited labels for the OASIS [9] brain set for research and educational purposes (5.A.3)
2. Morphology database of co/variations within and across brains (3.A.1, 5.A.4)

Background and Significance

3.A. Significance of brain structure identification and brain parcellation

Identifying structures and labeling ROIs extracted from magnetic resonance images (MRI) is essential for a broad range of research applications, including morphometric analysis of brain structures, ROI analyses of functional imaging data, connectivity-based analyses of structures, determination of receptor distribution and number in positron emission tomography (PET) studies, and partial volume correction in PET research. Because a full parcellation of the brain would mean a complete identification of all structures within a brain, we will focus on parcellation methods in this section, with shape analysis addressed in 3.E.

3.A.1. Morphometric analysis of MRIs

Brain morphometry, the study of sizes and shapes of brain regions, has become an exciting area for computational anatomists, neuroscientists, and neurologists alike. With regard to clinical applications, various neuropsychiatric disorders are accompanied by volumetric differences in brain structures. These range from enlarged ventricles in schizophrenia [10], decreased cerebral and hippocampal volumes in seizure disorders [11], decreased hippocampal volumes in post-traumatic stress disorder [12 --59], cerebrum and cerebellum volume changes in Fragile X [13] and Downs Syndrome [14], ventricular volume changes in normal aging [15], and altered temporal lobe volumes in Alzheimer's disease [16]. However, these findings have been difficult to reproduce [10], particularly in mood disorders [17,18]. For example, there is a great disparity in the literature regarding volumetric changes in unipolar and bipolar depression [18]. Undoubtedly one source of the wide variability is in the method used to identify structures and determine the ROIs amongst different groups. An automated structure identification and parcellation method such as the proposed revision of Mindboggle will allow for a user-independent automated determination of the volumes (as well as shapes) of all significant regions within the brain for comparative and longitudinal studies.

To our knowledge, there is no publicly available repository of data regarding morphological covariability of structures in individual human brains. To address one of our long-term research objectives, that of characterizing this variability in healthy subjects and patients for basic and clinical research, we will conduct statistical analyses (5.A.4) on manually edited versions of Freesurfer-labeled, publicly available (OASIS) brains. The manually edited label sets and statistical results will not only help to construct the Bayesian priors for the Bayesian framework in this proposal, but will serve as valuable resources for the scientific community.

3.A.2. ROI analysis of fMRI data

Group statistical analysis of multi-subject fMRI blood-oxygen-level dependent (BOLD) activity data is hindered by noise in the acquired data [19] and the striking amount of regional, neuroanatomical variability in the location of sulci and gyri across subjects (3.C.1), even after transformation into a standard space (3.C.2). Standard voxel-based based fMRI analyses address these problems by spatially smoothing data [20] in three-dimensional (volume element) space. This results in the smearing of functional data across topologically and functionally distant regions. For example, data from auditory cortex in the superior temporal plane are blended with data from motor and/or somatosensory cortices in the opercula located immediately above the temporal plane. While these areas lie in close proximity in 3-dimensional space, they are far apart on the cortical sheet, separated by the insular cortex. Furthermore, a nonlinear normalization procedure (see 3.C.3), meant to align anatomical regions across subjects, leaves a profound amount of residual inter-subject anatomical variability [21,22] while further blurring individual anatomical distinctions. Thus, standard voxel-based procedures compromise the anatomical localization of BOLD responses. An ROI-based procedure using parcellated imaging data would overcome this shortcoming, by pooling functional data within ROIs defined by individual anatomical landmarks prior to performing group statistical analyses. The ROI-based analysis eliminates the need for voxel-based spatial smoothing and non-linear normalization and has been shown to significantly increase statistical power [22].

3.A.3. Connectivity-based analysis of MRIs

Diffusion-weighted imaging is a growing area in neuroscience, and very recently, researchers are beginning to learn about connectivity between cortical regions, and classifying tracts based on the cortical regions to which they connect [23,24]. Mindboggle could aid in this classification by automatically classifying the ROIs on either end of a fiber tract.

3.A.4. ROIs for PET studies

Parcellation of brain MRIs into ROIs is also useful in PET and SPECT studies. PET has a lower intrinsic resolution than MRI and so determination of small regions or boundaries between large regions can be difficult on typical PET scans. Currently, many centers acquire MRI scans along with the PET scans. The ROIs are drawn on MRIs, the MRI and PET images are coregistered, and the ROIs are transferred from the MRI to the PET images [25-27]. The mean activity within the ROIs is then used to determine the activity in the identified region. Automated determination of ROIs by Mindboggle could be transferred onto coregistered PET images for regional quantification of radioligand binding.

3.A.5. Partial volume correction in PET imaging

PET imaging has a resolution of between 4-8 mm at full width half maximum depending on the camera. As a result of this limited resolution, quantifying the activity in a given voxel in a PET image is hampered in two ways [28]. First, the signal measured within an individual PET voxel may arise from radioactivity from several tissue types; gray matter typically gives rise to higher activity, whereas white matter and cerebrospinal fluid typically give rise to lower activity. This is of concern for gray matter structures which abut white matter or cerebrospinal fluid and also for regions that have a mixture of gray and white matter, such as the thalamus. The correction for this effect requires additional information such as a high-resolution MRI to determine the relative contribution of the tissue components. The second error is the "spillover" of activity between regions where activity from the ROI spills over to adjoining regions and activity within the ROI is contaminated by activity from surrounding regions. Both types of errors, called partial volume errors, are well-characterized problems in PET imaging that result in the inaccurate determination of activity within an ROI. Practically, it has been difficult to correct for this error because every voxel in the structural imaging dataset has to be assigned a label. When implemented, partial volume correction can result in a 40 to 60% increase in measured regional PET activity [29,30].

3.B. Parametric shape analysis with deformable models

Shape analysis for the description and comparison of structures can be based on points, curves, surfaces, or volumes. The geometric (Euclidean distance-based) shape analysis methods used by the current version of Mindboggle treat shapes as discrete voxels or point distributions, and result in simple comparisons such as average distance and overlap measures. A parametric shape analysis that treats shapes as continuous surfaces to better characterize their form and texture would result in better structure identification and ROI delineations than relying solely on geometric approaches.

Physically based deformable models, referred to as snakes, are based on elasticity theory. Lagrangian equations of non-rigid motion are expressed in terms of position functions in Euclidean 3-D space. The deformable model representation is parametric and deformation of the model is applied on its surface via minimization of individual equations of motion for each set of Cartesian coordinates. This energy expresses the equilibrium of internal and external forces on the surface. Such deformable models were introduced by Kass *et al.* [31] as 2-D explicit deformable contours and generalized to the 3-D case by Terzopoulos *et al.* [32].

A second family of deformable models was introduced by Osher and Sethian [33]. These models are based on an implicit formulation of the surface to deform, embedded in a level set function. In an original approach, Caselles *et al.* [34] proposed a reformulation of the snake deformable model as the definition of a geodesic curve whose length is constrained by image-based data information.

The above deformable models were applied to the task of image segmentation. A separate family of segmentation methods was proposed by Mumford and Shah [35] to segment an image into smooth areas and a "finite" set of contours. This variational approach was reformulated recently by Chan and Vese [36] into a level-set like deformable model approach.

3.C. Brain parcellation methods

Approaches to the identification of brain structures determination of ROIs range from fully manual to fully automated. Fully manual to semi-automated approaches involve varying degrees of feature identification and hand drawing on a computer screen. Fully automated approaches predominantly take two forms: registration of an image to a standard template in a parcellated space, and nonlinear coregisteration of an individual atlas brain image and an individual target brain image.

3.C.1. Manual and semi-automated parcellation

Manually assigning anatomical labels requires expertise and is tedious and time-consuming. Labeling is sensitive to the significant morphological variation across brains [37-46] and labeling inconsistencies within and across labelers [47-50]. The great variability across brains demands tailored labeling for each individual brain. To indicate the level of investment required to manually label brain anatomy, the Center for Morphometric Analysis (CMA) at the Massachusetts General Hospital expects at least a month of training to train new technicians to the point of acceptable inter-rater reliability using their CardViews [47] labeling protocol and software; once trained, it takes hours to days to manually label a single brain. The assistant who labeled half of the brain images of section 4.B trained at the CMA for two months and took two weeks to label each brain. It took an entire five-month period to simply label 10 of the 22 brains we used to evaluate Mindboggle. At this rate, performing a modest imaging study with 20 subjects and 20 controls would require 20 months devoted strictly to labeling. For these reasons, fast, consistent, and accurate automation of anatomical labeling is needed.

3.C.2. Template-guided parcellation

Automated brain labeling is universally attempted by first registering labeled (atlas) and unlabeled (target) brain images in the same standardized space, then transferring atlas labels to the target brain. The standardized space consists of a three-dimensional coordinate system usually containing a template brain, the registration target. Registration consists of a spatial transform to the template in that coordinate system. The transformation vectors provide data on shape and volume differences. Many analyses are performed in standardized space, such as voxel-level statistical maps.

One of the most popular, but *inaccurate,* methods for automatically labeling a brain is to linearly register a target brain to an atlas space where the labels reside. The widely used Talairach coordinate system is an attempt to coregister brains in a piece-wise linear manner about two medial structures [51,52]. Linear and piece-wise linear coregistration is not sufficient to align the local features of two brains and so one or the other method is routinely implemented as a prelude to nonlinear coregistration, most often by warping (nonlinearly deforming) one volume to another [53-56].

3.C.3. Warping-based parcellation

Atlas-based approaches that parcellate an individual's brain in a nonlinear manner consist predominantly of warping algorithms. The primary problem with relying solely on warping to nonlinearly register one brain to another is that without sufficient constraints, there are many ways to reshape a brain to look like another without regard for anatomical borders. Indeed, image correspondence is often mistaken for anatomical correspondence [57,58]. Because point correspondence from one cortex to another is ill-defined, some degree of manual intervention is often used to initially assign corresponding points or curves about which deformations are to be performed [59-64]. It is also questionable to assume that intervening points between two well-defined landmarks correspond to intervening points between matching landmarks. The point correspondence problem is simply revisited at a smaller scale; if different anatomical structures happen to exist between the corresponding pairs of landmarks in two brains, there may be no point-to-point correspondence. Coregistration by warping assumes that topography is preserved from brain to brain, where topography refers to the existence of component structures (such as sulci) and their relative spatial relationship along the cortical surface. This assumption may not be justifiable, as there may be missing or interrupted sulci in some brains and the sequence of the sulci may be different between brains [39]. Arguably, the most popular warping software packages that have been demonstrated to completely label a brain using an atlas are SPM [65,66], ANIMAL [67,68], AIR [69,70], and Rueckert's Image Registration Toolkit [71,72].

3.C.4. Alternative parcellation methods

There are two chief alternatives to warping: Bayesian approaches and feature identification with label filling (Bayesian-warping and feature-warping hybrids also exist). Each alternative has one leading software package for fully automated, complete parcellation of the human cortex: Freesurfer [73,74] and BrainVisa [75-77]. Freesurfer uses an anisotropic, nonstationary Markov random field to model the spatial relationships between neighboring labeled structures, and does not presently rely on warping [78-80]. BrainVisa identifies sulcal structures and assigns labels to and between these structures.

3.D. Comparative evaluation of Mindboggle and alternative methods

Evaluations that have been made of warping/parcellation strategies are in general difficult to assess. They are usually demonstrated with restricted label sets or sparse landmarks, and evaluated under artificial conditions or most commonly by visual inspection where image correspondence can be mistaken for anatomic correspondence. We are aware of only a few studies that have compared different nonlinear registration algorithms [74-80].

When compared with manual labels, Mindboggle was shown to give results that compared favorably [2] against SPM2, ANIMAL, and AIR, as well as linear registration with FLIRT (see section 4). We have just completed and submitted for publication a far larger comparative study that includes FLIRT and 14 nonlinear brain image registration methods: AIR, ANIMAL, ART, Diffeomorphic Demons, JRD-fluid, FNIRT, IRTK, ROMEO, SICLE, SyN, and four SPM5 algorithms: SPM2-type as well as template-based Normalization, Unified Segmentation, and DARTEL Toolbox. The participants/coauthors have provided guidance in the use of their software through extensive communications, and several have even upgraded or written additional software or parameter files for the study. The participants include: Jesper Andersson, Babak Ardekani, John Ashburner, Brian Avants, Ming-Chang Chiang, Gary Christensen, Louis Collins, Pierre Hellier, Joo Hyun Song, Mark Jenskinson, Claude Lepage, Daniel Rueckert, Paul Thompson, Tom Vercauteren, and Roger Woods, and their lab members.

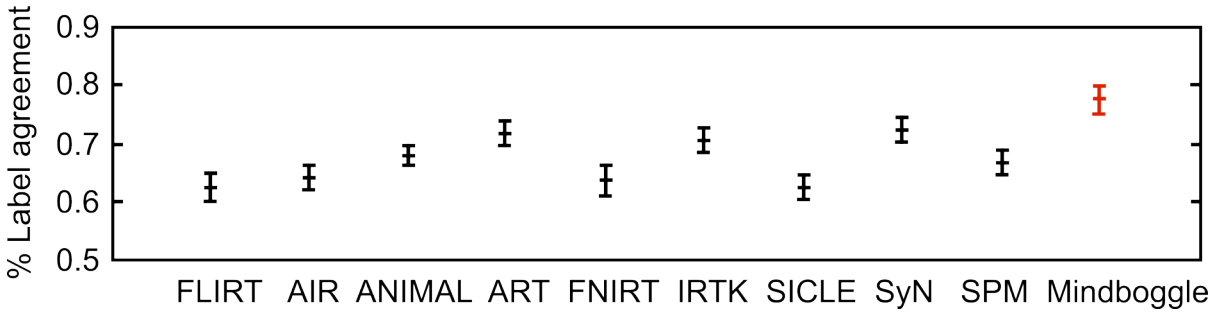

Fig. 1: Accuracy of nonlinear brain registration methods for automated anatomical labeling
Mindboggle accuracy is at least one standard deviation better than every other method that we have evaluated. The vertical axis is the percent of voxels in the manually labeled brain volume whose automatically assigned labels agree with the manual labels (error bars extend two standard deviations). Nine other methods were run 90 times each on different brain-to-brain pairs from a set of 10 brains used to evaluate Mindboggle in 2005 [1]. Mindboggle spreads labels to completely fill a target brain mask, whereas the other methods deform labels, which usually leads to incomplete overlap with target manual labels.

A brief description of the comparative evaluation study follows (see Fig.1):

Subjects: Brain images of 80 normal subjects (male & female, mixed handedness) were acquired from different laboratories, including LONI's LPBA40 40 individual, manually labeled brains. For each subject, there is one T1-weighted M.R. scan and a corresponding manually labeled volume.
Processing: Non-brain matter was removed from each MRI by applying a mask constructed from the manual labels for that brain. All of these brain images were linearly registered to MNI space using FLIRT (9 DOF with trilinear interpolation), then each of these images was linearly registered (6 DOF) to each of the remaining images in its group (four groups of 10, 12, 18, and 40 subjects), resulting in 2,168 linearly registered brain images.
Registration: Each of the nonlinear registration methods was then applied to each of the 2,168 brain images above to register it to each of the 80 brain images in its group, resulting in 2,168 nonlinearly registered brains for each method.
Evaluation: Volume overlap (see 4.A), surface overlap, volume similarity, and distance error measures were used to evaluate how well individual anatomical regions register to one another, assuming the manual label sets as silver standards.

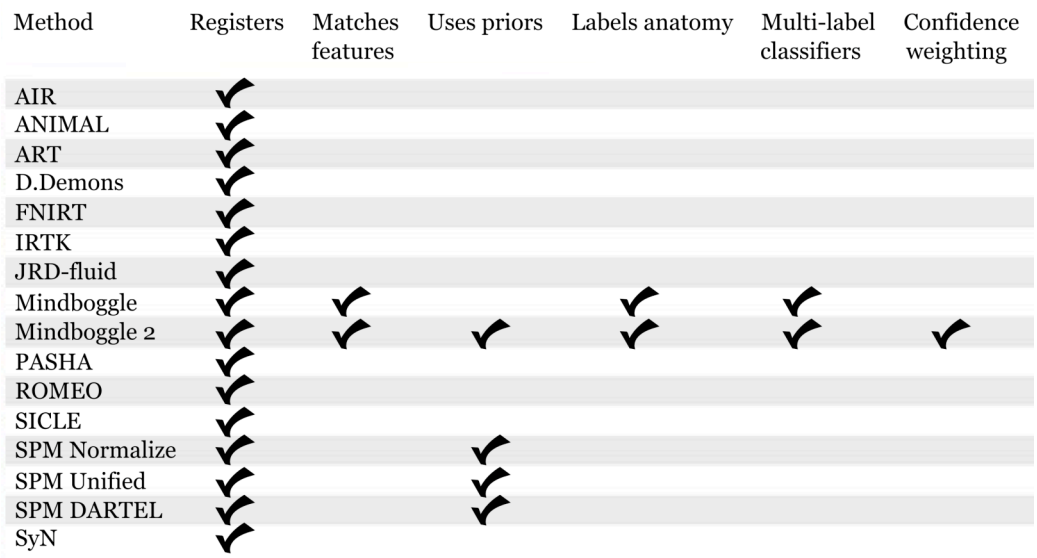

Fig. 2: Feature comparison chart of brain registration methods
This figure compares the features supported by the brain registration methods in our evaluation study (see text). Mindboggle2 refers to our proposed Mindboggle that incorporates the Bayesian framework of this proposal.

Freesurfer's cortical parcellation algorithm has been evaluated primarily with a volume similarity measure and mean distance maps [79,80], both of which underestimate anatomical misregistration. As an example of how a volume similarity measure can fail, one could label a region, offset a copy of the region so that it is disjoint from its original, and the corresponding regions have a similarity of one (equal volume) yet have zero overlap. The mean distance map is a point-for-point estimate of mismatch, but the true indicator of mismatch between two corresponding regions is the offset between their boundaries, since their internal point-to-point mapping is ill-defined, as stated above. The gyral parcellation of the cortical surface in BrainVisa's 2005 pipeline has only been evaluated by visual inspection; its validity rests on how many of the sulci are correctly identified, which has been estimated at only around 75% [77]. To our knowledge, neither Freesurfer's, nor BrainVisa's, cortical parcellation methods have been tested against any other parcellation methods other than linear coregistration with an atlas. Both became available at the same time as Mindboggle, and they have not been adequately compared against one another, primarily because Freesurfer and BrainVisa produce 2-D surface-based representations of the labeled cortex whereas Mindboggle labels in the native 3-D brain space. In addition, the parcellation schemes and atlas label definitions between the three packages differ significantly.

3.E. Multiple atlases

To mitigate the problem of morphological variability across brains, a given brain may be compared with not one, but many other brains. There has been a rising trend in the neuroimaging community to make a single comparison between a given brain registered to a "common space" and a label set in that space; sometimes this label set is acquired from a single individual [52,88], and sometimes it is a probabilistic label set derived from multiple coregistered individual brains. There has been a recent upsurge in the availability and use of probabilistic atlases [89-92]. These atlases provide valuable information about the average shape and extent of regions. However, publications are appearing that demonstrate the superiority of labeling with multiple individual labeled brain image datasets over labeling with probabilistic, averaged, or indirect (propagated) label data [93,94]. Mindboggle's accuracy was shown to improve significantly by using multiple, individual atlases rather than a single, individual atlas [3] using a simple majority voting rule.

3.F. Historical context

Data are becoming available and communities are getting organized to address difficulties in finding and using software and data. Thousands of brain images, with multiple acquisitions per subject, such as OASIS, IXI, BIRN, and IBSR have recently become publicly available. They will provide extensive training and test data for our work. Parcellated brains have become available (e.g., [95]) that could provide an early development label set while we create our own label sets. NITRC [96], the Neuroimaging Informatics Tools and Resources Clearinghouse, released in October 2007, provides a center for finding and comparing software and sharing ideas for the neuroimaging community. Mindboggle is now registered on NITRC's website and will be actively maintained on the site with its own forum, mailing list, Subversion repository, etc.

Preliminary Studies

4.A. Mindboggle

Mindboggle [2,3] has been downloaded hundreds of times in over 21 different countries from its own website [1] (prior to registering it with nitrc.org), its second publication [3] has been flagged as a "highly accessed" article on BioMed Central's website (7,188 times since Oct. 2005; they state that "overall statistics indicate that your article will have been accessed on PubMed Central a roughly equivalent number of times…").

Mindboggle was developed to automate anatomical labeling in the guise of a feature-based matching problem, without making any assumptions of topographical or even topological preservation or one-to-one mapping (3.C.3). Mindboggle employs morphological image processing steps to reduce brain images to sets of sulcal pieces. Combinations of pieces derived from an atlas brain are matched against pieces from a target brain. The atlas label boundaries are then transformed in the target brain space to account for the differences in shape and size between the matching pieces, and these boundaries are filled with labels.

4.A.1. Preprocessing steps

Atlas and unlabeled target brains are presently preprocessed by third-party software prior to Mindboggle labeling, using Oxford's FMRIB Software Library (FSL) algorithms [100]. Preprocessing consists of stripping the skull, registering to a common space, segmenting tissues, and cropping the cerebellum, brainstem and subcortex. Any reasonable software that performs these steps will suffice; we have also used SPM in past evaluations. Mindboggle is quite robust to imaging artifacts and differences in preprocessing methods.

4.A.2. Create sets of pieces from brain images

The current version of Mindboggle performs the following steps to skeletonize a brain image, then fragment the skeleton to make reasonable, cohesive shapes (Fig. 1). The resulting set of pieces is used to find structural correspondences across brains:

1. Extract a non-white matter volume from the tissue-segmented input image volume.
2. Erode the outer cortical surface until only sulcal folds and subcortical regions remain.
3. For each axial slice, apply a majority-surround 2-D filter to remove isolated pixels.
4. For each axial slice, successively thin sulci (non-white matter) to 1-pixel-wide curves.
5. Divide the brain into hemispheres by a plane that is warped along its normals to medial skeletal points, using a modified Kohonen neural network, or Self-Organizing Map [101].
6. For each axial slice, find contiguous curves by segmenting each slice into 2-D sets of 8-connected pixels.
7. Assemble neighboring 2-D curves successively across slices to construct 3-D sulcal pieces.
8. Apply k-means clustering to fragment each piece (seed points set about bounding box for each piece).
9. Recombine those fragments that would form new pieces with low surface-to-volume ratios and remove very small pieces.

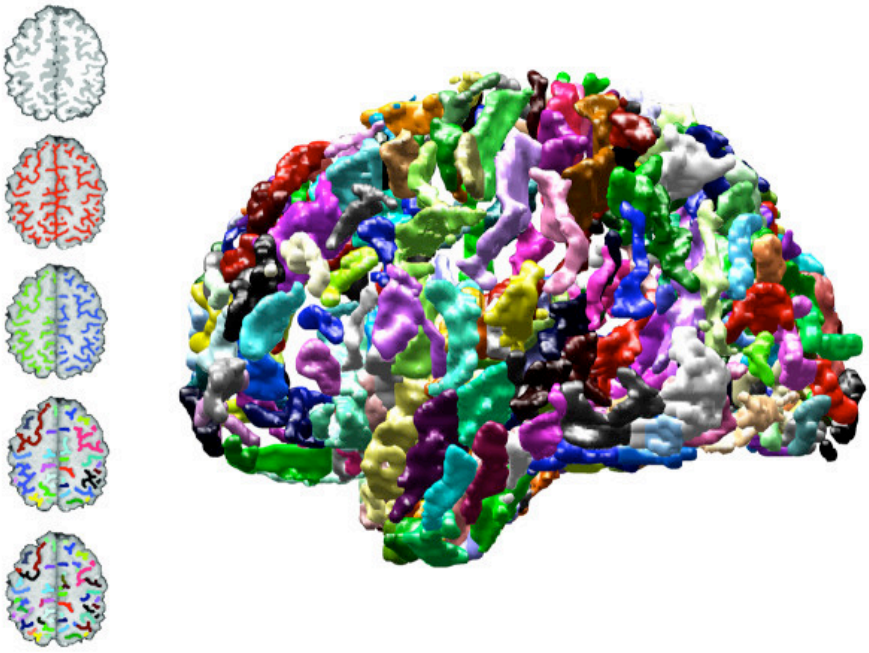

Figure 1: Piece construction
Mindboggle's five steps to construct sulcus pieces (left, top to bottom): Darker pixels (non-white matter) of a (1) tissue-segmented horizontal slice are thinned to a skeleton (2), which is split into left/right hemispheres (3). Adjacent pixels are grouped into pieces (4). These 2-D pieces are used to make 3-D pieces, shown in cross-section (5) and in 3-D (right figure; left side of the brain, front facing left, isosurfaced pieces of arbitrary color). The pieces do not have to correspond accurately to anatomical divisions because Mindboggle uses them not to label a brain directly, but to compute a transformation to deform atlas label boundaries. (Figure from [3]).

4.A.3. Match combinations of pieces from a brain atlas with pieces from a target brain
Once a set of (sulcus) pieces is extracted from each hemisphere of a labeled (atlas) and unlabeled (target) brain image, then Mindboggle presently establishes structural correspondences across the two brains by finding matches between combinations of features in the following manner (Fig. 2):

1. Find all nearby combinations of (1, 2, or 3) pieces from the atlas for each target piece.
2. Measure the similarity of these pairs by applying a cost function.
3. Sort costs across all atlas piece combinations for all target pieces.
4. Select matches to maximize similarity and minimize redundant piece selection (Eq. 1).

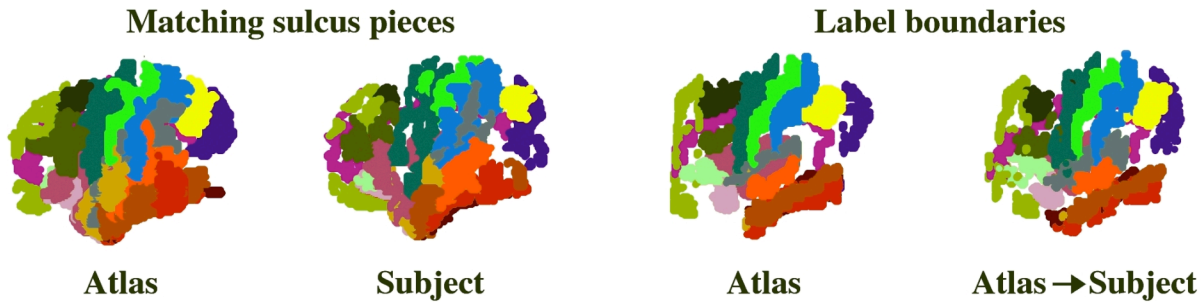

Figure 2: Matching sulcus pieces and transforming label boundaries
Left pair: Each atlas piece is matched with a combination of subject pieces by minimizing a cost function. Each colored region represents multiple pieces with a single sulcus label.
Right pair: Atlas label boundaries are transformed according to displacements between matching pieces, resulting in piece-wise linearly deformed atlas label boundaries in the subject brain.
All images are left views with the front of the brain facing left.

To select the best set of matches for each atlas piece, Mindboggle orders the tentative matches by a sum of four weighted differences serving as a cost function (Eq. 1):

Eq. 1: $EM = w_N N + w_V V + w_P P + w_O O$

Three of the four costs are differences between quantities computed from subject and atlas skeleton pieces: the number of points (N), number of subvolumes (V), and mean position (P). The fourth cost is the degree of non-overlap (O) of subvolumes occupied by the pieces. Number of subvolumes refers to the number of cubic subdivisions of the coregistration space that are occupied by at least one point from the piece(s) of interest. Non-overlap of two pieces, P1 and P2, is equal to the fraction of subvolumes of P1 that do not overlap P2 added to the fraction of subvolumes of P2 that do not overlap P1. The default weights $w_N N$, $w_V V$, $w_P P$, and $w_O O$ were empirically determined and were found to be robust.

4.A.4. Transform and fill atlas label boundaries

After identifying matching structures across an atlas and target brain, Mindboggle transforms the atlas label boundaries in the target brain space, according to the following steps:

1. Translate patches of atlas label boundaries according to the difference in position for each atlas-target match (Fig. 2).
2. Warp the original atlas label boundaries to these patch-wise translated boundaries, using a modified Self-Organizing Map algorithm.
3. Replace certain label boundaries with planes placed at atlas-appointed positions.
4. Propagate labels (using neighborhood majority labels) to fill the label boundaries within non-white matter (Fig. 3).
5. Each voxel is assigned multiple labels (by running Mindboggle multiple times with different individual atlases).

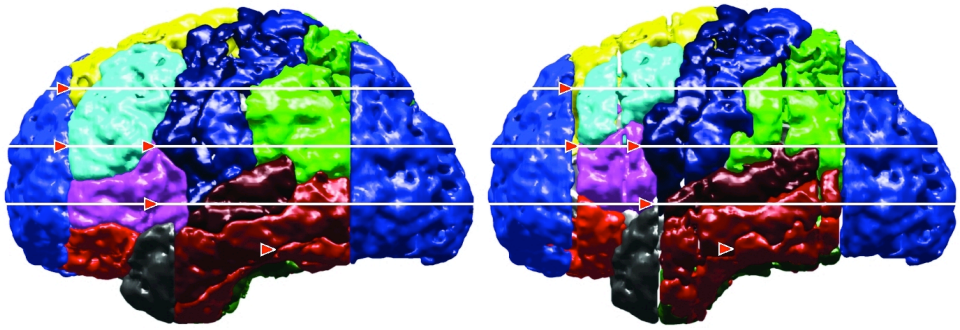

Figure 3: Mindboggle results
Atlas labels are warped to their transformed boundaries, then propagated to fill the brain's gray matter mask (left view). *Left:* an atlas after label warping and propagation to a target brain. *Right:* the manually labeled target brain. Red tags denote the largest discrepancies between Mindboggle-assigned and manually assigned labels for this individual. (Figure from [2].)

4.B. Evaluation of Mindboggle

When Mindboggle was created, software dedicated to labeling human brain gyral regions in a fully automated manner was not available. Klein used popular nonlinear registration methods (SPM2 [65], AIR [69], and ANIMAL [68]) to warp an atlas brain to an unlabeled target brain, applied the resulting transform matrix to the atlas labels, then propagated the labels through the target brain.

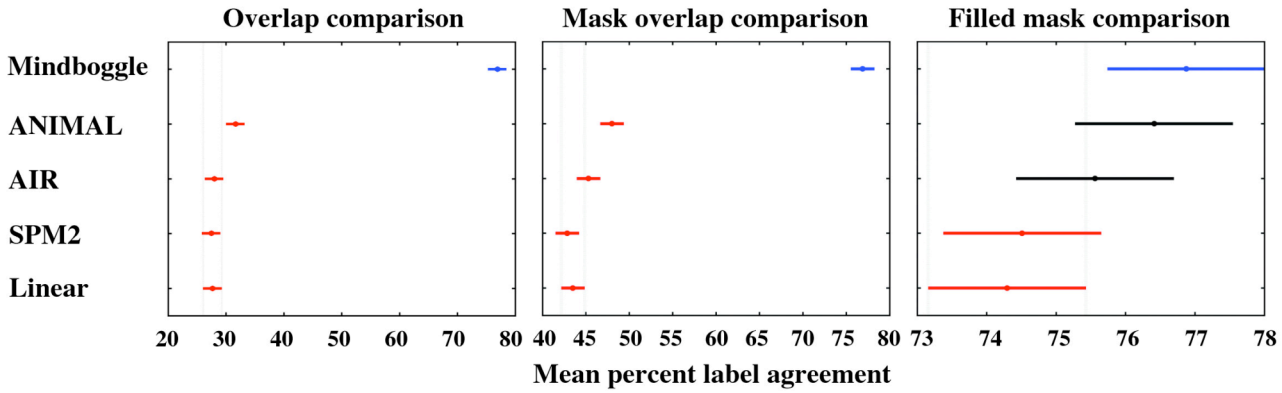

Figure 4: A one-way ANOVA was performed to test if the means are the same for the overlap label agreements obtained by different methods (linear, SPM2, AIR, ANIMAL, Mindboggle). A multiple comparison test was then performed using Tukey's honestly significant difference criterion to determine which pairs of means are significantly different. The graphs display the mean for each method with a 95% confidence interval around the mean, based on the Studentized range distribution. If intervals are disjoint, their means are considered significantly different. The same test was conducted again for the mask overlap and filled mask overlap label agreements (see text). Mindboggle's result is in blue and any significantly different result is in red. Mindboggle obtained a significantly higher mean union label agreement than any other method ($p<0.05$). After applying Mindboggle's label propagation procedure, only Mindboggle obtained a significantly higher mean label agreement than did linear registration or SPM2. (Figure from [2].)

Mindboggle was shown to give results for intact (Fig. 4) and artificially lesioned (Fig. 5) brain images that compared favorably against linear and the other nonlinear methods [2]. The artificial lesion case demonstrates that a feature-matching approach does not fail if features are removed, whereas nonlinear deformation approaches tend to erroneously fill the lesioned regions with labels.

The evaluation measures we used were overlap (intersection over union), mask overlap (intersection over manually labeled regions), and filled mask overlap (same as mask overlap, after Mindboggle's label propagation procedure was applied to completely fill unlabeled gray matter masks with labels).

For each subject, automatically assigned atlas labels, A_i, are compared against manually assigned labels, M_i, for each label, i. The total number of voxels (volume) with a given label assigned either automatically or manually is $|A_i \cup M_i|$, corresponding to the union of label sets A_i and M_i. The number of voxels that have been assigned this label both automatically and manually is $|A_i \cap M_i|$, corresponding to the intersection of label sets A_i and M_i. One measure of overlap between label sets A_i and M_i, called "overlap ratio" is defined as the volume of intersection divided by the average of the two volumes under comparison. This is a very conservative measure of overlap, so we defined overlap as the volume of intersection divided by the volume of union (also called the Jaccard coefficient), which may be extended to multiple labeled regions by summing over a set of labels of index i:

Eq. 2: $$\text{Overlap} = \frac{\sum_i |A_i \cap M_i|}{\sum_i |A_i \cup M_i|}$$

The mask overlap penalizes for discrepancies in overlap but not in size by comparing label sets A_i and M_i only within the set M_i, the manually labeled voxels:

Eq. 3: $$\text{Mask overlap} = \frac{\sum_i |A_i \cap M_i|}{\sum_i |M_i|}$$

The filled mask overlap is the same as mask overlap, but does not penalize for missing labels, by comparing manual with automated labels after the automated labels have been propagated to fill all unlabeled voxels:

Eq. 4: $$\text{Filled mask overlap} = \frac{\sum_i |A_i^F \cap M_i|}{\sum_i |M_i|}$$

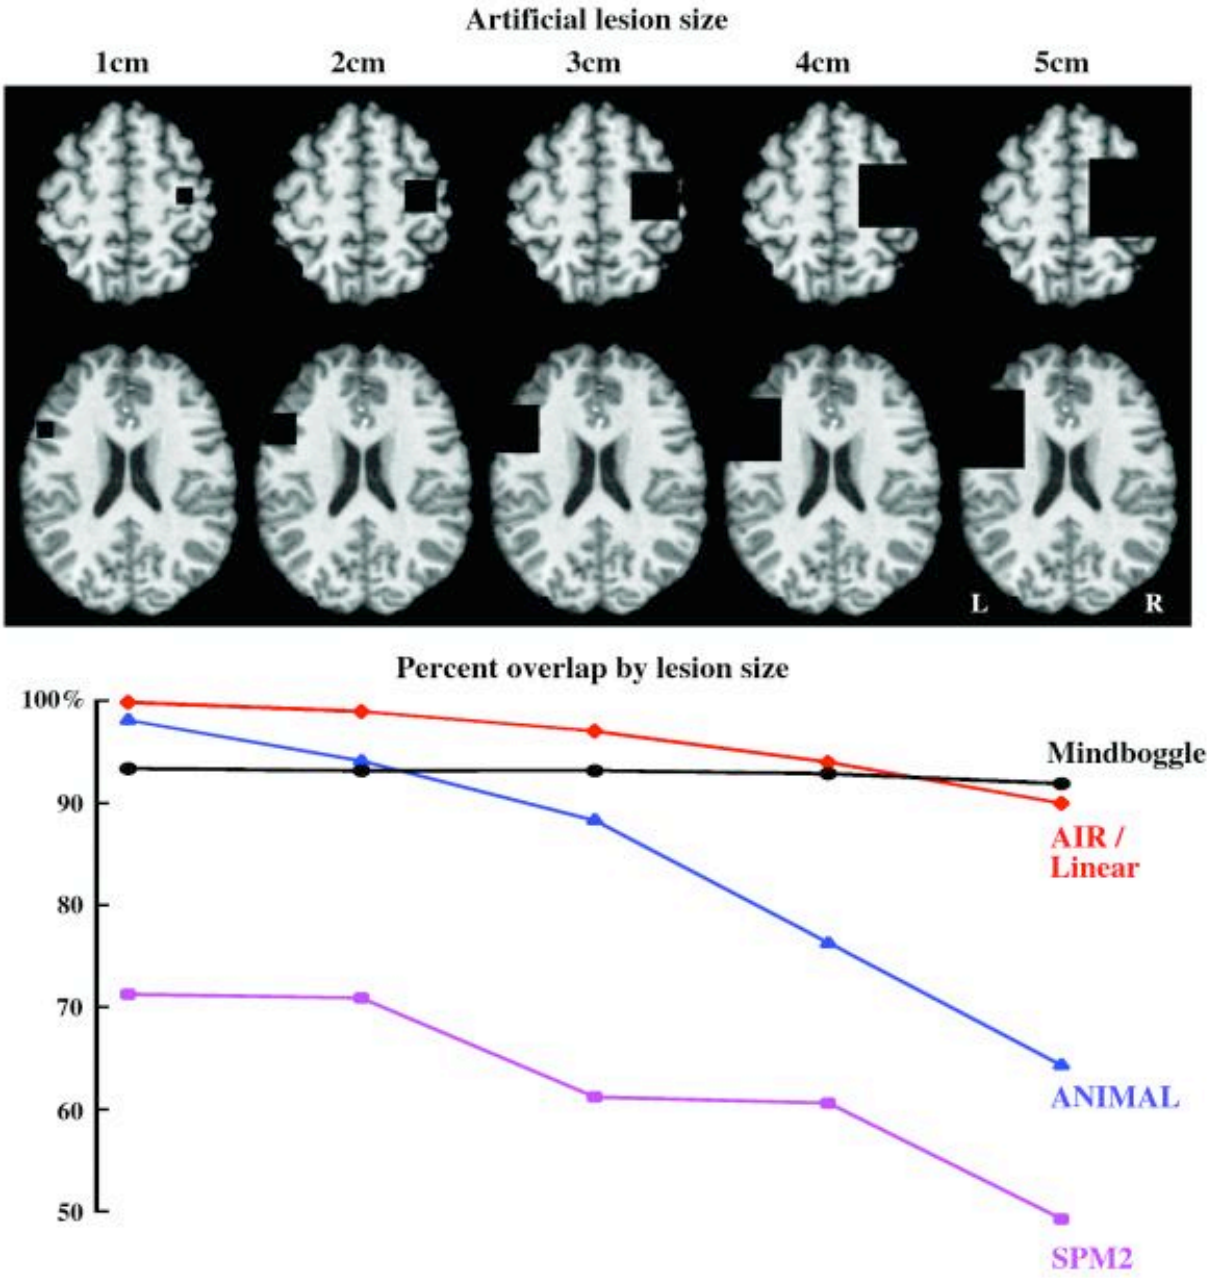

Fig. 5: Registering to artificially lesioned brain image data
Warping methods were used to register an atlas to an artificially lesioned version of itself. One of two lesions was centered on Broca's region (top row), and the other on the right motor strip (second row). The shape of the artificial lesion was a cube and was increased from 1 to 5 cm^3 (left to right in the figure). Unlike other methods that warp labels within lesioned areas, Mindboggle's labeling performance hardly degrades as the lesion size increases. (Figure from [2].)

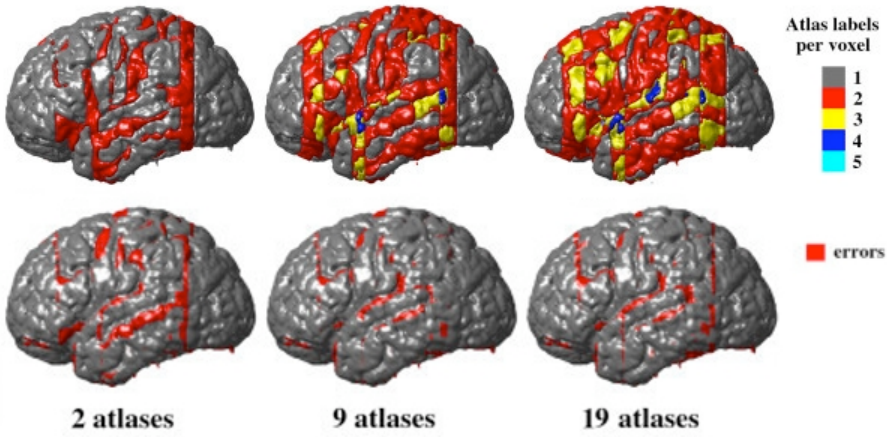

Fig. 6: Number of labels, number of errors per voxel
Top: Colors indicate the number of different labels assigned to each voxel by multiple atlases (2, 9, and 19 from left to right). Gray indicates one label, when all atlases agree. Increasing the number of atlases increased the number of different labels assigned to each voxel, but a voxel rarely had more than 3 labels.
Bottom: Red indicates voxels where the majority label of those assigned by different atlases disagrees with the voxels' manually assigned labels. Increasing the number of atlases reduces labeling errors (see Fig. 8). (Figure after Figs. 4 and 6 in [3].)

Using the filled mask overlap measure, Mindboggle's accuracy was shown to improve significantly by using multiple, individual atlases rather than a single, individual atlas [3] (Figs. 6 and 7) using a simple majority voting rule to decide on a single label for each voxel from the labels contributed by all of the atlases.

4.C. Algorithmic improvements

Since a description and evaluation of Mindboggle was published, we have made a number of algorithmic improvements that we intend to further test before release in the next version of Mindboggle:

- Subcortical labeling: Mindboggle no longer crops any part of the brain's gray matter, allowing for more valid volume and shape measurements and more ROIs.
- Interhemispheric division: Mindboggle now splits the brain much more quickly and, in preliminary tests, more accurately, by median filtering a surface extracted from a vertical slab of the skeleton.
- Distance function: Filling label boundaries using a distance transform runs about ten times faster and appears to be as accurate as the slow and iterative filling algorithm Mindboggle presently uses.
- A brain may now be labeled in a little over 5 minutes using Matlab code that has not been optimized, as compared against an hour for BrainVisa and over 16 hours for Freesurfer (on a Macintosh Pro, 3 GHz, OSX 10.4.10) (Fig. 1).

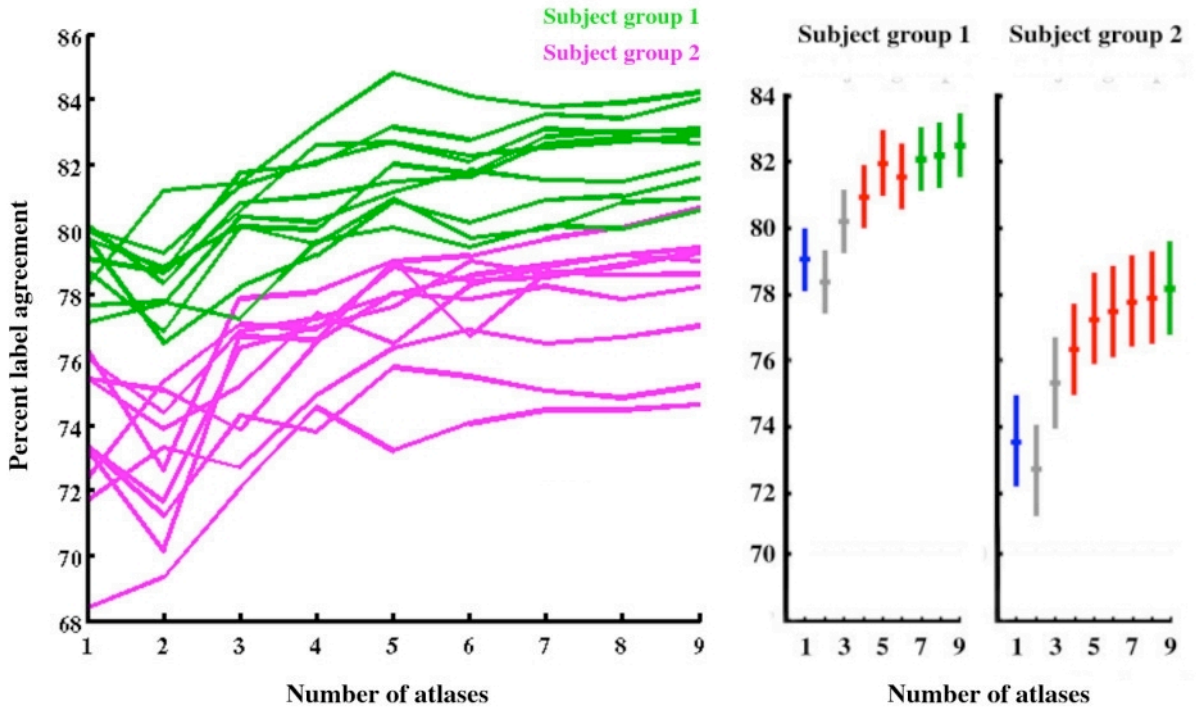

Fig. 7a (left): **Multiple label agreement by subject pool**
Two subject groups of 10 subjects each were manually labeled by different labelers. Each member of group 1 was labeled using Mindboggle with 1 atlas from the same group, then with 2, 3, up to 9 atlases, each atlas selected at random from the remaining atlases. This was repeated for each member from group 2. The results are clearly separable between the two groups. Group 1 consisted of similar subjects (right-handed, young, single race); group 2 consisted of a broad range of ages of multiple races and unknown handedness. Therefore, labeling accuracy is sensitive to variance in the subject population, and to different manual labeling methods. **Fig. 7b** (right): The same comparison was conducted as in Fig. 5 between the means of the label agreements obtained by different numbers of atlases. Red and green results are significantly higher than blue (single atlas) results and green results are significantly higher than gray results (*e.g.*, using 7 or 9 atlases resulted in significantly higher label agreements than using up to 3 atlases). (Figure after Figs. 7 and 9 in [3].)

4.D. Shape analysis using parametric active surfaces

In addition to the geometric features Mindboggle employs to compare shapes, such as mean distance between two coregistered shapes, their volumes, degree of overlap, *etc*., parametric shape characteristics can also be used to determine corresponding brain structures across brains for the purpose of brain structure identification and brain parcellation. Such shape analysis usually requires extraction and modeling of surfaces for mathematical and quantitative measurements to be used as criteria for shape comparison.

In one of our previous studies on 4-D live ultrasound of the human heart [7], a 3-D active surface model was proposed and used for left ventricular shape analysis in order to compare automated segmentation with manual tracing. The specific surface model used in that application was based on a finite-element model using cubic Hermite polynomials as surface descriptors. As presented in the paper, this framework can very efficiently represent convoluted surfaces with very few parameters. Moreover, after the polynomials are automatically fitted to a segmented surface, a continuous surface model can be automatically constructed. Based on the model, quantitative shape comparisons beyond volume similarity and overlap measures can be easily computed. In the paper [7], root-mean squared errors (RMSE) were used as quantitative criteria. In addition to phantom data, the segmentation algorithm was tested on each ultrasound frame of a clinical patient data set (87 frames in total). Quantitative evaluation of the automated segmentation compared against manual tracings yielded an RMSE of 2.58±1.58 mm, which corresponds to approximately 3 pixels. Additional measures like a quantitative error mapping through the entire surface can be generated. Representative results of a segmentation comparison are shown in Figure 8.

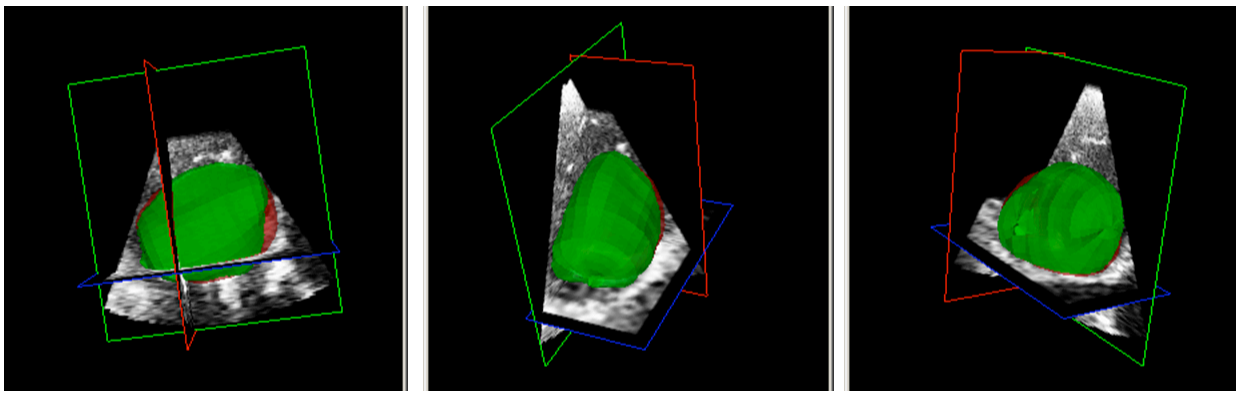

(a) (b) (c)

Fig. 8: Comparison of clinical data between an automated method (red) and manual tracing (green) on one ultrasound frame of the left ventricle of a human heart at three different view angles (a-c).

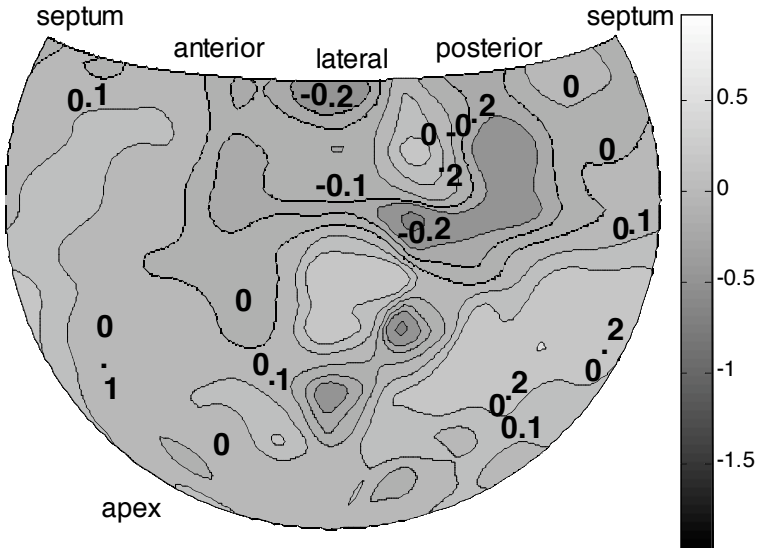

Fig. 9: Hammer mapping of relative errors between two left (cardiac) ventricular surfaces.

Figure 9 shows an example of the Hammer mapping technique, which is an area-ratio preserving flattened mapping, in this case of the left ventricle of the human heart. In this figure, relative errors between automated segmentation and manual tracing were mapped. This mapping technique can also be applied to brain structures and parcellations; after using a finite-element model combined with an active surface model, we can map any quantitative shape analysis metrics throughout the whole cortical surface of the brain.

By generalizing this demonstrated active shape framework to brain structures for structure identification and to ROIs for parcellation, we will be able to compare shapes and conduct quantitative evaluations in a more refined manner than geometric shape analysis alone can provide.

Research Design and Methods

5.A. Summarized Aims

This is a proposal to perform a large-scale shape analysis of brain structures for automatic identification. This procedure will be used to facilitate automatic whole-brain parcellation (see Figure 11):

5.A.1. Optimize the automated extraction of brain structures from brain image data.
5.A.2. Analyze the shapes of brain structures.
5.A.3. Manually label brains for training, testing, and brain atlas construction.
5.A.4. Construct a Bayesian framework to identify brain structures using shape analysis and manual labels.
5.A.5. Facilitate brain parcellation using these identified structures.
5.A.6. Evaluate the automated identification of structures and ROIs.

5.A.1. Optimize the automated extraction of brain structures

Presently, skeletal construction is performed on each slice independently, and the resulting slice-stack is further processed to construct a 3-D skeleton. A major problem with this method is that skeletal construction is affected by the orientation of the brain, or the axis along which the image is sliced. We will therefore implement a new "exact medial axis" algorithm for creating a 3-D skeleton [4,5,6] to replace steps 3, 4, 6, and 7 in 4.A.2 (C code is available online). This will avoid having to choose along which axis a skeletal slice-stack should be constructed, avoid skeletal artifacts such as those arising from folds lying along a slice, increase accuracy and reduce inconsistencies across acquisitions and orientations, and result in much faster processing times. An exact algorithm for the discrete bisector function and a thinning algorithm which produces homotopic discrete Euclidean skeletons [5] will be applied to an exact Euclidean medial axis method [4]. To refine the thinness characteristics of the skeletons for better shape characterization, we will transform the skeletons to higher resolution for further homotopic thinning [6].

5.A.2. Analyze the shapes of brain structures

We will adopt and create a variety of shape comparison metrics for the purpose of characterizing and matching brain structures by their shapes and for characterizing parcellated ROIs. Shape analysis will take two forms: parametric shape analysis using an active surface model and geometric shape analysis. Parametric and geometric descriptors are used to create shape features that characterize local and global shape attributes.

5.A.2.a. Parametric shape analysis

For parametric shape analysis we plan to use and extend on the Active Geometric Functions (AGF) framework [7] outlined in section 4.E to exploit its properties for efficient shape characterization. In order to efficiently capture the shape of the surface of an individual brain structure or the whole cortex we propose to use 3rd-order Hermite polynomials as surface descriptors instead of using a traditional triangulated mesh with a linear surface within each patch. This approach allows for dimensionality reduction by solving an N-1 dimensional variational problem, i.e. by using a 2-D function we can represent 3-D surfaces such as individual brain structures. Consider the following example.

In 1-D, there are four cubic Hermite basis functions:

Eq. 5:

$$H_0^0(\xi) = 1-3\xi^2+2\xi^3$$
$$H_0^1(\xi) = \xi\ (\xi-1)^2$$
$$H_1^0(\xi) = \xi^2\ (3-2\xi)$$
$$H_1^1(\xi) = \xi^2\ (\xi-1)$$

On a 2-D finite element patch, as shown in Figure 10, there will be 4 basis functions associated with each local coordinate direction (ξ_1 or ξ_2), which generates a total of 16 2-D basis functions. Figure 10b

shows an example of using 3rd-order Hermite polynomials to represent a convoluted 3-D surface in 3-D space using a single 2-D finite element patch. For comparison, if linear quadrilateral patches are used, about 100 linear patches are required to represent the same surface with comparable accuracy.

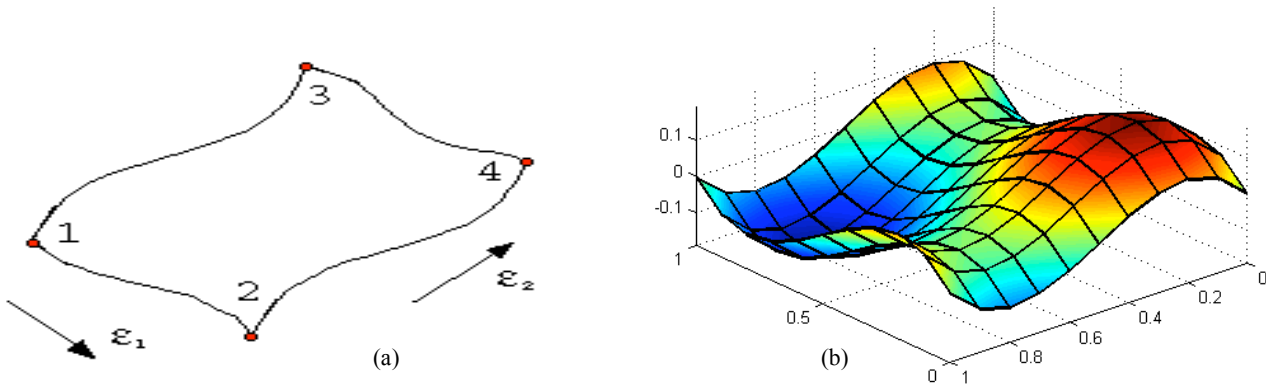

Fig. 10: Illustration of a cubic Hermite surface representation: (a) a 2-D surface finite element patch used in our model; (b) a single patch with cubic Hermite representation can efficiently characterize a convoluted surface as shown in color, whereas it takes about 100 linear quadrilateral patches to achieve similar accuracy.

We propose to derive measures for minimal shape parameterization to sufficiently represent the geometry of individual brain structures in terms of accuracy and efficacy. For each brain structure we plan to compare and measure the minimal node representation of a 2-D finite element patch with cubic Hermite representation by means of quantitative surface analysis to hand labeled structures.

For each brain structure or ROI, a subject-specific surface model will be built by a finite element model using a cubic Hermite patch. As shown in Figure 10 we can represent arbitrary complex 3-D surfaces with this approach. The model will be fitted to the surface via energy minimization similar to the method presented by Song *et al.* [8]. The 3-D active mesh model is deformed under external forces derived from a minimization of the Mumford-Shah energy functional:

Eq. 6: $$E(f,\vec{C}) = \beta\int_{\Omega}(f-g)^2\,dV + \alpha\int_{\Omega\setminus\vec{C}}|\nabla f|^2\,dV + \gamma\oint_{\vec{C}}ds$$

in which $\vec{C}$ denotes the smoothed and closed segmentation, g represents the observed image data, f is a piecewise smoothed approximation to g with discontinuities only along $\vec{C}$, and Ω denotes the image domain. The first integral enforces the similarity between f and g, which drives the model to fit the parcellated brain surface; the second integral controls the smoothness of f; and the last integral is actually the length of the segmented boundary, which acts as internal elasticity to prevent leaking at weak boundaries to overcome imperfect segmentation.

A surface model with exactly the same structure could be fitted to the corresponding structures or ROIs in a brain atlas or a labeled subject brain. Once the fitting step is completed, quantitative shape comparisons can be performed based on the metrics, such as root-mean squared errors, localized error mappings, *etc.*, derived from the model. We note that the AGF framework also allows for real-time segmentation of large volume data often found in brain image analysis. This parameterization will allow us to build smooth surface- and volume representations of the shapes of interest to derive local and global geometric shape features.

5.A.2.b. Geometric shape analysis

Brain structures will be treated as geometric shapes (discrete voxels or point distributions) in their native 3-D Euclidean space. Geometric analysis of shapes will include gross shape descriptors such as volume, extent, orientation, etc. In addition, we intend to build up a rich feature set of local and global

shape features using Wavelet analysis [110-111]. Optionally, alternative parameterizations and shape descriptors could be used to further enhance our parametric and geometric feature set. From this large feature set we intend to perform feature selection to obtain the minimal set of shape features that most contribute to maximized matching performance. Standard feature selection methods [112] could be applied for this task.

5.A.3. Manually label brains for training, testing, and brain atlas construction

We will require many manually labeled brain images for training, validation, feature selection, Bayesian prior estimation (5.A.4), and atlas construction. We will therefore manually edit Freesurfer parcellated brains ourselves for this research and for use by the biomedical community. The neuroanatomist Jason Tourville (see letter of support) will perform the editing of parcellations of 300 brains from the OASIS brain image repository over the course of the first year. The work will consist of manually editing the Freesurfer cortical and subcortical regions of interest to conform to underlying individual anatomy. The manually labeled brain image data will be made available as a resource for the imaging community via NITRC. All parcellations will be spatially aligned to a common reference frame using existing registration techniques.

5.A.4. Construct a Bayesian framework to identify brain structures

Mindboggle presently combines pieces of brain features and matches these combinations across brains in a rule-based manner, without drawing upon earlier experience such as information gained from a population of manually parcellated brains. We intend to statistically describe the variations and covariations in the morphology of our population of manually parcellated (5.A.3) ROIs, and use this information in a Bayesian framework to guide feature combination and matching.

We will adopt the general approach and slightly modify the nomenclature of a successful Bayesian framework for face identification [104]. In that work, the authors formulate a probabilistic similarity measure which is based on the probability that the differences between two images, Δ, are within the variations (*e.g.*, expressions) of the same face. In their formulation, they are solving a binary pattern classification problem between intraclass (same-face) variation, Ω_I, and extraclass (different-face) variation, Ω_E. The similarity measure is the posterior probability that two images are within intraclass variation given a difference Δ between the images, according to Bayes' Theorem:

Eq. 7: $$S(I_1,I_2) = P(\Omega_I \mid \Delta) = P(\Delta \mid \Omega_I)\, P(\Omega_I)/P(\Delta)$$

where $$P(\Delta) = P(\Delta \mid \Omega_I)P(\Omega_I)+P(\Delta \mid \Omega_E)P(\Omega_E)$$

The maximum *a posteriori* (MAP) rule is then used to solve the classification problem. Two images are considered to correspond to the same class if $P(\Omega_I \mid \Delta) > P(\Omega_E \mid \Delta)$, or equivalently, if $S(I_1,I_2) > 1/2$.

This is analogous to our feature-matching problem; just as they need to decide whether a candidate face corresponds to a given (identified) face, we need to decide whether a candidate feature (combination of pieces) corresponds to a given (labeled) feature (*e.g.*, the central sulcus). Our images I_1 and I_2 will correspond to images not of faces but of these features, and each class refers to the set of instances of a given feature (*e.g.*, central sulci in a set of brains).

Because we can automatically extract features from an atlas brain and assign each feature to a labeled region, we can assemble feature classes from our population of atlases (5.A.3). We will employ a similarity measure that will be based on parametric as well as geometric analysis of the shapes capturing local and global shape morphology. By quantifying the pair-wise differences between each member of a feature class (*e.g.*, the central sulcus) and any other member across the population, we can statistically describe intraclass variations Ω_I. If we perform a similar comparison between a member of a feature class and any member of any other feature class across the population, we can statistically describe extraclass variations Ω_E. We will estimate the likelihoods $P(\Delta \mid \Omega_I)$ and $P(\Delta \mid \Omega_E)$ and priors $P(\Omega_I)$ and $P(\Omega_E)$ from these variations (with the assumption that their distributions are Gaussian).

Calculating Bayesian priors should drastically reduce the number of likely combinations of pieces and the number of candidate matches to consider. This would make it computationally feasible to remove the stringent constraints in the present version of Mindboggle, namely, that only a maximum of three pieces are allowed to combine for matching with each single atlas piece.

5.A.4.a. MAP Estimation

To compute the a posteriori probability $P(\Omega_I | \Delta)$ and $P(\Omega_E | \Delta)$ we make use of our hand labeled brain atlases from section 5.A.3. For each individual brain structure we select a training subset to model intraclass and extraclass variations of each individual ROI. To reduce computational cost we plan to use feature dimensionality reduction using either PCA or other available dimensionality reduction techniques on the feature set we obtained from section 5.A.2. In order to compute $P(\Omega_I | \Delta)$ and $P(\Omega_E | \Delta)$ we need to estimate the likelihood and prior probability terms in equation 7. We begin by computing pair-wise differences between the feature class and the individual brain structures following a projection onto intra- and extraclass Gaussian subspace models. We model the intra- and extraclass variations as high dimensional Gaussians on the reduced feature space. Gaussian densities are modeled as:

$$P(\Delta|\Omega_E) = \frac{e^{-\frac{1}{2}\Delta^T \Sigma_E^{-1} \Delta}}{(2\pi)^{D/2}|\Sigma_E|^{1/2}}$$

Eq. 8:
$$P(\Delta|\Omega_I) = \frac{e^{-\frac{1}{2}\Delta^T \Sigma_I^{-1} \Delta}}{(2\pi)^{D/2}|\Sigma_I|^{1/2}}.$$

The likelihood term can be easily computed by Euclidean distances on the reduced feature space representation by directly working in the eigen-space of the feature set. To reduce computational complexity we aim to pre-compute the feature vectors of our population of individual brain structures. The maximum likelihood estimate for the intra- and extraclass variations can be simply computed as:

$$P(\Delta|\Omega_E) = \frac{e^{-\frac{1}{2}||\mathbf{e_j}-\mathbf{e_k}||^2}}{(2\pi)^{D/2}|\Sigma_E|^{1/2}}$$

Eq. 9:
$$P(\Delta|\Omega_I) = \frac{e^{-\frac{1}{2}||\mathbf{i_j}-\mathbf{i_k}||^2}}{(2\pi)^{D/2}|\Sigma_I|^{1/2}}$$

Here **i** and **e** represent the reduced eigen-feature space of the ROI and other brain structures in the database. For the prior terms we will assume equal probability between both classes and plan to investigate different prior distributions to study the effect and performance of different priors for robust identification and matching of brain structures.

After computing $P(\Delta | \Omega_I)$ and $P(\Delta | \Omega_E)$ we plan to use the probabilistic similarity measure in eq. 7 to query and match individual brain structures for identification and brain parcellation. We believe that our Gaussian assumption on $P(\Delta | \Omega_I)$ and $P(\Delta | \Omega_E)$ and feature dimensionality reduction approach allows for efficient and robust brain structure identification and matching. To further increase computational performance we could adopt a maximum likelihood similarity matching requiring us to only compute the a posteriori probability of the intraclass similarity. Potential difficulties could lie in the independent assumption of our feature space as well as the Gaussian assumption on our feature classes. In case where performance of our initial approach might show non-satisfying performance we would investigate other distributions and sampling strategies to compute the MAP for identifying and matching brain structures. Further concerns lie in the eigen feature space representation of the individual brain structures. Non-linearity within the data could result in unsatisfying results when

using linear subspace projection methods. One potential workaround is to consider a generalized PCA approach with a data representation that allows for permutational invariants [113].

5.A.5. Facilitate brain parcellation using these identified structures

In the original version of Mindboggle, skeletal structures are not anatomically identified, but are merely used to find structural correspondences across brains. With the Bayesian framework of 5.A.4, the skeletal structures will be identified in a probabilistic fashion; moreover, they will be constructed with guidance from Bayesian priors as opposed to finding "all nearby combinations of ... pieces" (4.A.3). Once these structures have been identified, atlas label boundaries can be transformed as per 4.A.4. By identifying the anatomy of structures rather than simply combining them in a blind fashion before matching with structures from an atlas brain, we are confident that parcellation results will be more accurate and precise.

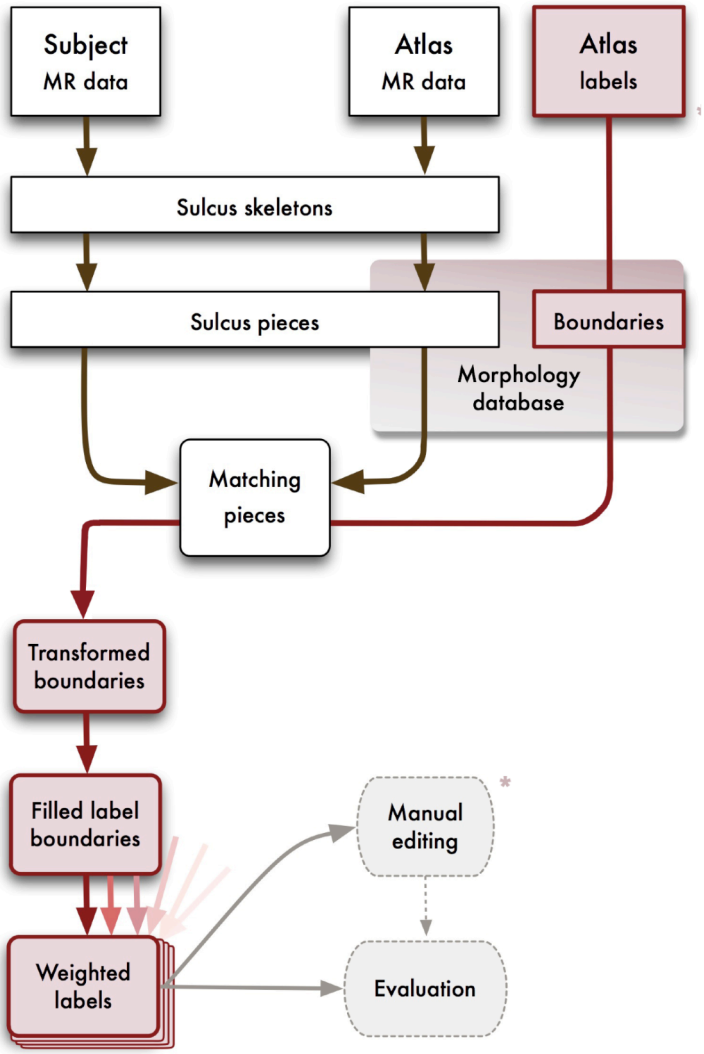

Fig. 11: Schematic overview with references to sections within Research Design and Methods. A morphology database (middle right) is populated with atlas label boundaries and structural data and, within a Bayesian framework, guides sulcus piece construction and matching. Manual editing of the final, automatically generated label volumes can provide more Atlas labels* and in turn enrich the morphology database.

5.A.6. Evaluate the automated identification of structures and ROIs

We will use the dataset of 300 manually edited parcellated brains of section 5.A.3 for use as atlases, as well as training and test data to evaluate our work. We will evaluate both the automated identification of brain structures (skeletons) as well as the automated parcellation into ROIs, as compared against manual

parcellations. Comparisons will be made between each and every pair of brains. For 300 brains, that would mean each brain is compared with every other brain, resulting in 44,850 comparisons.

For evaluating automated brain structure identification within the Bayesian framework of 5.A.4, we will compute the fraction of skeleton voxels in a brain that are both (1) closest to a manually labeled anatomical boundary in that brain and are also (2) matched with atlas skeleton voxels that are closest to the same anatomical boundary in the atlas brain. In other words, a skeleton voxel closest to the central sulcus in one brain is considered correctly identified (with respect to a given atlas) if it is matched with an atlas structure that is closest to the central sulcus in the atlas brain. This atlas comparison is repeated for the remaining 298 atlases.

For evaluating automated parcellations, we will employ the evaluation protocol that we successfully implemented and about which we have submitted a manuscript for publication (described in 3.D). Briefly, the volume and the degree of overlap for each automatically labeled ROI will be compared quantitatively with its corresponding manually labeled ROI according to the overlap measures of Equations 2, 3, and 4, and the overlap ratio (4.B) as well as label boundary distance measures. We will conduct these evaluations across different brains as above, and also across different acquisitions of the same brain to test the consistency and robustness of our method. For this within-subject evaluation, we will use the multiple acquisition data included with the OASIS dataset, again coupled with our manually edited parcellations of these data (5.A.3).

After the above evaluations, we have multiple eager and willing beta testing laboratories at Columbia University, Cold Spring Harbor Laboratory (Brain Architecture Project), MIT, and elsewhere.

5.B. Distribute software and resources

To facilitate the widespread adoption of the algorithms developed as part of this research, we will make the codebase, formats, and distribution as mainstream and forward-thinking as possible:

5.B.1. Codebase

Our algorithms will be implemented in cross-platform, free, open source, and easily maintainable and adaptable code. The software will be written in C and in Python, a freely available (almost universally preinstalled), open source, stable programming language with strong community support that is extremely well-suited to scientific computing, particularly since the recent development of NumPy and SciPy [107]. Mindboggle will take advantage of Python's object-oriented features to create easily maintainable, modular code with built-in, automatic documentation generation capabilities (pydoc). Python emphasizes simplicity and consistency: "There should be one -- and preferably only one -- obvious way to do it." This emphasis facilitates future maintainability of Python code by those other than the original programmer. We will make extensive use of NumPy for dealing with multidimensional arrays, Matplotlib [108] for 2-D plotting of (publication quality) interim results, iPython [109] for interactive programming, and existing open source Python scripts for basic utilities. Our programs will be distributed with either a GNU license (http://www.gnu.org/copyleft/gpl.html) as it presently has, or a modified BSD license, like that of Slicer, VTK, ITK, Python, NumPy, SciPy, etc. (*e.g.*, http://www.slicer.org/cgi-bin/License/SlicerLicenseForm.pl).

5.B.1. Formats

Our programs will adhere to standard brain image formats (e.g., NIfTI, Analyze, and Dicom) and parcellation methods [101] to characterize ROIs for use with different software packages and a range of research applications such as ROI analysis, diffusion tensor imaging, and morphometry.

5.B.2. Documentation

In addition to writing detailed and clear documentation within each and every software program (automatically assembled and formatted by pydoc), we will create online documentation detailing the use of Mindboggle and its individual components. We will also distribute a manual with the software, and create online tutorials.

5.B.3. Distribution

We will develop Mindboggle under the scrutiny of NITRC [96], where the original version of Mindboggle is also registered and downloadable, with dedicated forum, mailing list, and Subversion version control.

Feedback through NITRC will help us to create new versions that are faithful to the needs, desires, and habits of the neuroimaging community. The codebase, documentation, and related files will be made available through NITRC's website and Mindboggle's website.

There are very few publicly available data in a consistent form regarding morphological variability of brain regions. We will be constructing such a database while establishing Bayesian priors for ROI feature matching (5.A.4). We recognize that this information could be of great value to others interested in morphological co/variations within and across brains, so we will make these data publicly accessible from a website as well.

5.C. Software design

We will be roughly following the agile process referred to as feature-driven development. We will break each major feature down into a hierarchical list of smaller features that may be implemented in workable code in a shorter time. We will plan by feature, design by feature, and build by feature in an iterative and incremental manner. We will work in the general spirit of agile development in that we will remain adaptive to change, develop and test features for workable code models, and test accuracy against a training dataset at each stage of development.

Reference Cited

1. "Mindboggle" http://www.binarybottle.com/mindboggle.html
2. Klein, Arno, and Joy Hirsch. "Mindboggle: a scatterbrained approach to automate brain labeling." *NeuroImage* 24, no. 2 (January 15, 2005): 261-280.
3. Klein, Arno, Brett Mensh, Satrajit Ghosh, Jason Tourville, and Joy Hirsch. "Mindboggle: Automated brain labeling with multiple atlases." *BMC Medical Imaging* 5, no. 7. http://www.biomedcentral.com/1471-2342/5/7.
4. Remy, E., and E. Thiel. "Exact medial axis with euclidean distance." *Image and Vision Computing* 23, no. 2 (February 1, 2005): 167-175.
5. Couprie, Michel, David Coeurjolly, and Rita Zrour. "Discrete bisector function and Euclidean skeleton in 2D and 3D." *Image Vision Comput* 25, no. 10 (2007): 1543-1556.
6. Saúde, Couprie, and Lotufo. "Exact Euclidean Medial Axis in Higher Resolution." *Discrete Geometry for Computer Imagery*, 2006. http://dx.doi.org/10.1007/11907350_51
7. Q. Duan, G. Shechter, L. F. Gutierrez, D. Stanton, L. Zagorchev, A. F. Laine, and D. Elgort, "Augmenting CT Cardiac Roadmaps with Segmented Streaming Ultrasound," in *SPIE International Symposium, Medical Imaging 2007* San Diego, CA, USA, 2007.
8. T. Song, E. D. Angelini, B. D. Mensh, and A. F. Laine, "Comparison Study of Clinical 3D MRI Brain Segmentation Evaluation," in *26th Annual International Conference IEEE Engineering in Medicine and Biology Society (EMBS)* SanFrancisco, California, 2004.
9. "OASIS: Home Page." http://www.oasis-brains.org
10. McCarley, R., C. Wible, M. Frumin, Y. Hirayasu, J. Levitt, I. Fischer, et al. *MRI anatomy of schizophrenia*. 1999.
11. Duncan, John S. "Seizure-induced neuronal injury: Human data." *Neurology* 59, no. 90095 (November 12, 2002): 15S-20.
12. Villarreal, Gerardo, Derek A Hamilton, Helen Petropoulos, Ira Driscoll, Laura M Rowland, Jaqueline A Griego, et al. "Reduced hippocampal volume and total white matter volume in posttraumatic stress disorder." *Biological psychiatry* 52, no. 2 (July 15, 2002): 119-25.
13. Kates, Wendy R., Bradley S. Folley, Diane C. Lanham, George T. Capone, and Walter E. Kaufmann. "Cerebral growth in Fragile X syndrome: Review and comparison with Down syndrome." *Microscopy Research and Technique* 57, no. 3 (2002): 159-167.
14. Pinter, J D, S Eliez, J E Schmitt, G T Capone, and A L Reiss. "Neuroanatomy of Down's syndrome: a high-resolution MRI study." *The American journal of psychiatry* 158, no. 10 (October 2001): 1659-65.
15. Cook, I.A., M.L. Morgan, J.J. Dunkin, S. David, E. Witte, R. Lufkin, et al. "Estrogen replacement therapy is associated with less progression of subclinical structural brain disease in normal elderly women: a pilot study." *International Journal of Geriatric Psychiatry* 17, no. 7 (2002): 610-618.
16. Busatto G. F, Garrido G. E.J, Almeida O. P, Castro C. C, Camargo C. H.P, Cid C. G, et al. "A voxel-based morphometry study of temporal lobe gray matter reductions in Alzheimer's disease." *Neurobiology of Aging* 24, no. 2 (March 2003): 221-231.
17. Soares, J.C., and J.J. Mann. " The Anatomy of Mood Disorders--Review of Structural Neuroimaging Studies." *Biological Psychiatry* 41 (1997): 86-106.
18. Beyer, John L, and K Ranga R Krishnan. "Volumetric brain imaging findings in mood disorders." *Bipolar Disorders* 4, no. 2 (2002): 89-104.
19. Kruger, G., and G. H. Glover. "The physiological noise in oxygen-sensitive magnetic resonance imaging." *Magnenetic Resonance in Medicine* 46 (2001): 631-637.
20. Lowe, M.J., and J. A. Sorensen. "Spatially filtering functional magnetic resonance imaging data." *Magnetic Resonance in Medicine* 37, no. 5 (1997): 723-729.
21. Amunts, Katrin, Axel Schleicher, Uli Burgel, Hartmut Mohlberg, Harry B. Uylings, Karl Zilles, et al. "Broca's region revisited: cytoarchitecture and intersubject variability." *Journal of Comparative Neurology* 412, no. 2 (September 1999): 319-341.
22. Nieto-Castanon A, Ghosh S. S, Tourville J. A, and Guenther F. H. "Region of interest based analysis of functional imaging data." *NeuroImage* 19, no. 4 (2003): 1303-1316.
23. Maddah, Mewes, Haker, Grimson, and Warfield. "Automated Atlas-Based Clustering of White Matter Fiber Tracts from DTMRI." *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2005*, 2005.

24. Xia, Turken, Whitfield-Gabrieli, and Gabrieli. “Knowledge-Based Classification of Neuronal Fibers in Entire Brain.” *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2005*, 2005.
25. Sheline, Yvette I., Mark A. Mintun, Stephen M. Moerlein, and Abraham Z. Snyder. “Greater Loss of 5-HT2A Receptors in Midlife Than in Late Life.” *Am J Psychiatry* 159, no. 3 (March 1, 2002): 430-435.
26. Parsey, Ramin V., Lawrence S. Kegeles, Dah-Ren Hwang, Norman Simpson, Anissa Abi-Dargham, Osama Mawlawi, et al. “In Vivo Quantification of Brain Serotonin Transporters in Humans Using [11C]McN 5652.” *J Nucl Med* 41, no. 9 (September 1, 2000): 1465-1477.
27. Drevets, W C, E Frank, J C Price, D J Kupfer, P J Greer, C Mathis, et al. “Serotonin type-1A receptor imaging in depression.” *Nuclear medicine and biology* 27, no. 5 (July 2000): 499-507.
28. Aston, John A D, Vincent J Cunningham, Marie-Claude Asselin, Alexander Hammers, Alan C Evans, Roger N Gunn, et al. “Positron emission tomography partial volume correction: estimation and algorithms.” *Journal of cerebral blood flow and metabolism : official journal of the International Society of Cerebral Blood Flow and Metabolism* 22, no. 8 (August 2002): 1019-34.
29. Rousset, O.G., P. Deep, H. Kuwabara, A.C. Evans, A.H. Gjedde, P. Cumming, et al. “ Effect of partial volume correction on estimates of the influx and cerebral metabolism of 6-[(18)F]fluoro-L-dopa studied with PET in normal control and Parkinson's disease subjects.” *Synapse* 37, no. 2 (2000): 81-89.
30. Mawlawi, Osama, Diana Martinez, Mark Slifstein, Allegra Broft, Rano Chatterjee, Dah-Ren Hwang, et al. “Imaging Human Mesolimbic Dopamine Transmission With Positron Emission Tomography: I. Accuracy and Precision of D2 Receptor Parameter Measurements in Ventral Striatum.” *J Cereb Blood Flow Metab* 21, no. 9 (2001): 1034-1057.
31. M. Kass, A. Witkin and D. Terzopoulos, "Snakes: Active contour models." International Journal of Computer Vision, Vol. 1, pp. 321-331, 1987.
32. D. Terzopoulos and D. Metaxas, "Dynamic 3D models with local and global deformations: Deformable superquadratics." IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 13, No. 7, pp. 703-714, 1991.
33. S. Osher and J. A. Sethian, "Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations." Journal of Computational Physics, Vol. 79, No. 1, pp. 12-49, 1988.
34. V. Caselles, R. Kimmel and G. Sapiro, "Geodesic active contours." International Journal of Computer Vision, Vol. 22, pp. 61-79, 1997.
35. D. Mumford and J. Shah, "Boundary detection by minimizing functional", International Conference on Computer Vision and Pattern Recognition, pp. 22-26, San Francisco, CA, USA, 1985.
36. T. F. Chan and L. A. Vese, "Active contours without edges." IEEE Transactions on Image Processing, Vol. 10, No. 2, pp. 266 - 277, 2001.
37. Jouandet, M., and M. Deck. “The prenatal growth of the human cerebral cortex: – a brainprint analysis.” *Radiology* 188 (1993): 765-774.
38. Le Goualher, G., A. M Argenti, M. Duyme, W. F. C Baar\'e, H. E Hulshoff Pol, D. I Boomsma, et al. “Statistical sulcal shape comparisons: application to the detection of genetic encoding of the central sulcus shape.” *NeuroImage* 11 (2000): 564-574.
39. Ono, Michio, Stefan Kubik, and Chad D. Abernathey. *Atlas of the Cerebral Sulci*. Thieme, 1990.
40. Rademacher, J., V. S. Caviness, H. Steinmetz, and A. M. Galaburda. “Topographical Variation of the Human Primary Cortices: Implications for Neuroimaging, Brain Mapping, and Neurobiology.” *Cereb. Cortex* 3, no. 4 (July 1, 1993): 313-329.
41. Steinmetz, H, A Herzog, Y Huang, and T Hackländer. “Discordant brain-surface anatomy in monozygotic twins.” *The New England journal of medicine* 331, no. 14 (October 6, 1994): 951-2.
42. Evans, Alan C., D. L. Collins, S. R. Mills, E. D. Brown, R. L. Kelly, T. M. Peters, et al. “3D statistical neuroanatomical models from 305 MRI volumes.” Edited by Lowell Klaisner. *Nuclear Science Symposium & Medical Imaging Conference: 1993 IEEE Conference Record*, October 1993.24.
43. Thompson, P M, R P Woods, M S Mega, and A W Toga. “Mathematical/computational challenges in creating deformable and probabilistic atlases of the human brain.” *Human brain mapping* 9, no. 2 (February 2000): 81-92.
44. Wright, I C, P Sham, R M Murray, D R Weinberger, and E T Bullmore. “Genetic contributions to regional variability in human brain structure: methods and preliminary results.” *NeuroImage* 17, no. 1 (September 2002): 256-71.
45. Zilles, Karl, Axel Schleicher, Christian Langemann, Katrin Amunts, Patricia Morosan, Nicola Palomero-Gallagher, et al. “Quantitative analysis of sulci in the human cerebral cortex: Development, regional heterogeneity, gender difference, asymmetry, intersubject variability and cortical architecture.” *Human Brain Mapping* 5, no. 4 (1997): 218-221.

46. Zilles, K, R Kawashima, A Dabringhaus, H Fukuda, and T Schormann. “Hemispheric shape of European and Japanese brains: 3-D MRI analysis of intersubject variability, ethnical, and gender differences.” *NeuroImage* 13, no. 2 (February 2001): 262-71.
47. Caviness Jr., V.S., J. Meyer, N. Makris, and D.N. Kennedy. “MRI-based topographic parcellation of human neocortex: an anatomically specified method with estimate of reliability.” *Journal of Cognitive Neuroscience* 8 (1996): 566-587.
48. Fiez, J. A, H. Damasio, and T. J Grabowski. “Lesion segmentation and manual warping to a reference brain: intra- and interobserver reliability.” *Human Brain Mapping* 9 (2000): 192-211.
49. Lancaster, J. L, M. G Woldorff, L. M Parsons, M. Liotti, C. S Freitas, L. Rainey, et al. “Automated Talairach Atlas labels for functional brain mapping.” *Human Brain Mapping* 10 (2000): 120-131.
50. Towle, Vernon L, Leila Khorasani, Stephen Uftring, Charles Pelizzari, Robert K Erickson, Jean-Paul Spire, et al. “Noninvasive identification of human central sulcus: a comparison of gyral morphology, functional MRI, dipole localization, and direct cortical mapping.” *NeuroImage* 19, no. 3 (July 2003): 684-97.
51. Talairach, J., and G. Szikla. *Atlas d'anatomie stereotaxique du telencephale: etudes anatomo-radiologiques*. Paris: Masson & Cie, 1967.
52. Talairach, J., and P. Tournoux. *Co-Planar Stereotaxic Atlas of the Human Brain*. New York: Thieme Medical Publishers, 1988.
53. Maintz, J.B. Antoine, and Max A. Viergever. “An Overview of Medical Image Registration Methods.” 1996.
54. Maintz, J., and M. Viergever. “A survey of medical image registration.” *Medical Image Analysis* 2, no. 1 (1998): 1-36.
55. Toga, Arthur W. *Brain Warping*. Academic Press, 1999.
56. Gholipour, A., N. Kehtarnavaz, R. Briggs, M. Devous, and K. Gopinath. “Brain Functional Localization: A Survey of Image Registration Techniques.” *Medical Imaging, IEEE Transactions on* 26, no. 4 (2007): 427-451.
57. Crum, W. R, L. D Griffin, D. L. G Hill, and D. J Hawkes. “Zen and the art of medical image registration: correspondence, homology, and quality.” *NeuroImage* 20 (2003): 1425-1437.
58. Rogelj, P., S. Kovacic, and J. Gee. *Validation of a nonrigid registration algorithm for multimodal data*. 2002. http://citeseer.ist.psu.edu/rogelj02validation.html
59. Bookstein, F. L. “Principal warps: thin plate splines and the decomposition of deformations.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11 (1989): 567-585.
60. Chui, H., L. Win, R. Schultz, J. Duncan, and A. Rangarajan. “A unified feature registration method for brain mapping.” Edited by M. F Insana and R. M Leahy. *Proc. 17th Information Processing in Medical Imaging, Davis, CA, USA, Lecture Notes in Computer Science*, 2001.
61. Evans, Alan C., Weiqian Dai, D. Louis Collins, Peter Neelin, and Sean Marett. “Warping of a computerized 3-D atlas to match brain image volumes for quantitative neuroanatomical and functional analysis.” Edited by Murray H. Loew. *Medical Imaging V: Image Processing*, June 1991.
62. Hartkens, T., D. L. G Hill, A. D Castellano-Smith, D. J Hawkes, C. R Maurer Jr. , A. J Martin, et al. “Using points and surfaces to improve voxel-based non-rigid registration.” Edited by T. Dohi and R. Kikinis. *Proc. 5th Medical Image Computing and Computer-Assisted Intervention, Tokyo, Japan, Lecture Notes in Computer Science*, 2002.
63. Johnson, H. J, and G. E Christensen. “Landmark and intensity-based, consistent thin-plate spline image registration.” Edited by M. F Insana and R. M Leahy. *Proc. 17th Information Processing in Medical Imaging, Davis, CA, USA, Lecture Notes in Computer Science*, 2001.
64. Magnotta, V. A, H. J Bockholt, H. J Johnson, G. E Christensen, and N. C Andreasen. “Subcortical, cerebellar, and magnetic resonance based consistent brain image registration.” *NeuroImage* 19 (2003): 233-245.
65. “SPM - Statistical Parametric Mapping.” http://www.fil.ion.ucl.ac.uk/spm/
66. Friston, K. J, J. Ashburner, J. B Poline, C. D Frith, J. D Heather, R. S. J Frackowiak, et al. “Spatial Registration and Normalization of Images.” *Human Brain Mapping* 2 (1995): 165-189.
67. “The Brain Imaging Software Toolbox.” http://www.bic.mni.mcgill.ca/software/
68. Collins, D. L, and A. C Evans. “Animal: automatic non-linear image matching and anatomical labeling.” In *Brain Warping*, edited by A. W Toga, 123-142. San Diego: Academic Press, 1999.
69. “Automated Image Registration.” http://bishopw.loni.ucla.edu/AIR5/index.html
70. Woods, Roger P., Scott T. Grafton, John D.G. Watson, Nancy L. Sicotte, and John C. Mazziotta. “Automated Image Registration: II. Intersubject Validation of Linear and Nonlinear Models.” *Journal of Computer Assisted Tomography* 22, no. 1 (February 1998): 153-165.

71. “Image Registration Toolkit.” http://wwwhomes.doc.ic.ac.uk/~dr/software/
72. Rueckert, D., L.I. Sonoda, C. Hayes, D.L.G. Hill, M.O. Leach, D.J. Hawkes, et al. “Nonrigid registration using free-form deformations: application to breast MR images.” *Medical Imaging, IEEE Transactions on* 18, no. 8 (1999): 712-721.
73. “FreeSurfer.” http://surfer.nmr.mgh.harvard.edu/
74. Dale, Anders M., Bruce Fischl, and M. I. Sereno. “Cortical Surface-Based Analysis I: Segmentation and Surface Reconstruction.” *NeuroImage* 9, no. 2 (February 1999): 179-194. ftp://ftp.nmr.mgh.harvard.edu/pub/flat/freesurfer/docs/recon1.ps.
75. “BrainVISA/Anatomist Home Page.” http://brainvisa.info/
76. Cointepas, Yann, Jean-François Mangin, Line Garnero, Jean-Baptiste Poline, and Habib Benali. “BrainVISA: Software Platform for Visualization and Analysis of Multi-modality Brain Data.” *NeuroImage* 13, no. 6 (June 2001): 98. http://www.academicpress.com/www/journal/hbm2001/11140.html.
77. Mangin, J.-F., D. Riviere, A. Cachia, E. Duchesnay, Y. Cointepas, D. Papadopoulos-Orfanos, et al. “A framework to study the cortical folding patterns.” *NeuroImage* 23, no. Supplement 1 (2004): S129-S138.
78. Fischl, B., D. H Salat, E. Busa, M. Albert, M. Dieterich, C. Haselgrove, et al. “Whole brain segmentation: Automated labeling of neuroanatomical structures in the human brain.” *Neuron* 33 (2002): 341-355.
79. Fischl, B., A. van der Kouwe, C. Destrieux, E. Halgren, F. S\'egonne, D. H Salat, et al. “Automatically parcellating the human cerebral cortex.” *Cerebral Cortex* 14 (2004): 11-22.
80. Desikan, Rahul S., Florent Segonne, Bruce Fischl, Brian T. Quinn, Bradford C. Dickerson, Deborah Blacker, et al. “An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest.” *NeuroImage* 31, no. 3 (July 1, 2006): 968-980.
81. Crivello, F., T. Schormann, N. Tzourio-Mazoyer, P. E Roland, K. Zilles, B. M Mazoyer, et al. “Comparison of spatial normalization procedures and their impact on functional maps.” *Human Brain Mapping* 16 (2002): 228-250.
82. Dinov, I. D, M. S Mega, P. M Thompson, R. Woods, D. W Sumners, E. R Sowell, et al. “Quantitative comparison and analysis of brain image registration using frequency-adaptive wavelet shrinkage.” *IEEE Transactions on Information Technology in Biomedicine* 6 (2002): 73-85.
83. Gee, J. C, D. C Alsop, and G. K Aguirre. “Effect of spatial normalization on analysis of functional data.” Edited by K. M Hanson. *Medical Imaging 1997: Image Processing, Proc. SPIE, Bellingham, WA*, 1997.
84. Hellier, P., C. Barillot, I. Corouge, B. Gibaud, G. Le Goualher, L. Collins, et al. “Retrospective Evaluation of Inter-subject Brain Registration.” *Lecture Notes in Computer Science* 2208 (2001): 258-??
85. Hellier, P., J. Ashburner, I. Corouge, C. Barillot, and K. J Friston. “Inter subject registration of functional and anatomical data using SPM.” Edited by T. Dohi and R. Kikinis. *Proc. 5th Medical Image Computing and Computer-Assisted Intervention, Tokyo, Japan, Lecture Notes in Computer Science*, 2002.
86. Hellier, P, C Barillot, I Corouge, B Gibaud, G Le Goualher, D L Collins, et al. “Retrospective evaluation of intersubject brain registration.” *IEEE transactions on medical imaging* 22, no. 9 (September 2003): 1120-30.
87. Crum, William R, Oscar Camara, Daniel Rueckert, Kanwal K Bhatia, Mark Jenkinson, Derek L G Hill, et al. “Generalised overlap measures for assessment of pairwise and groupwise image registration and segmentation.” *Medical image computing and computer-assisted intervention : MICCAI* 8, no. Pt 1 (2005): 99-106.
88. Tzourio-Mazoyer, N., B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, et al. “Automated Anatomical Labeling of Activations in SPM Using a Macroscopic Anatomical Parcellation of the MNI MRI Single-Subject Brain.” *NeuroImage* 15, no. 1 (January 2002): 273-289.
89. Hammers, Alexander, Ioannis Gousias, Laila Ahsan, Spyridion Vossos, Chi-Hua Chen, Rolf A. Heckermann, et al. “ A frequency-based atlas of the human brain: Extension to 30 brains / 83 regions and validation with SPM5.” June 10, 2007.
90. Eickhoff, Simon B., Klaas E. Stephan, Hartmut Mohlberg, Christian Grefkes, Gereon R. Fink, Katrin Amunts, et al. “A new SPM toolbox for combining probabilistic cytoarchitectonic maps and functional imaging data.” *NeuroImage* 25, no. 4 (May 1, 2005): 1325-1335.
91. “LONI | Atlases.” http://www.loni.ucla.edu/Atlases/index.jsp
92. “Brain Atlases included with FSL.” http://www.fmrib.ox.ac.uk/fsl/fslview/atlas-descriptions.html

93. Heckemann, Rolf A., Joseph V. Hajnal, Paul Aljabar, Daniel Rueckert, and Alexander Hammers. "Automatic anatomical brain MRI segmentation combining label propagation and decision fusion." *NeuroImage* 33, no. 1 (October 15, 2006): 115-126.
94. Rohlfing, Torsten, Daniel B. Russakoff, and Calvin R. Maurer Jr. "Performance-Based Classifier Combination in Atlas-Based Image Segmentation Using Expectation-Maximization Parameter Estimation." *IEEE Trans. on Medical Imaging* 23, no. 8 (2004): 983+.
95. "IBSR (Internet Brain Segmentation Repository)." http://www.cma.mgh.harvard.edu/ibsr/
96. "The Neuroimaging Informatics Tools and Resources Clearinghouse (NITRC)." http://www.nitrc.org
97. "NIPY: Neuroimaging tools for Python." http://neuroimaging.scipy.org/
98. Millman, K. Jarrod, and Matthew Brett. "Analysis of Functional Magnetic Resonance Imaging in Python." *IEEE Computing in Science & Engineering* 9, no. 3 (June 2007): 52-55.
99. "The MathWorks - MATLAB and Simulink for Technical Computing." http://www.mathworks.com/
100. Smith, Stephen M., Mark Jenkinson, Mark W. Woolrich, Christian F. Beckmann, Timothy E.J. Behrens, Heidi Johansen-Berg, et al. "Advances in functional and structural MR image analysis and implementation as FSL." *NeuroImage* 23, no. Supplement 1 (2004): S208-S219.
101. Kohonen, T. *Self-organizing maps*. New York: Springer-Verlag, 1997.
102. P.J. Besl, N.D. McKay, A method for registration of 3-D shapes, IEEE Transactions on Pattern Analysis and Machine Intelligence 14 (1992) 239–256.
103. Cordon, O., S. Damas, and J. Santamaria. "Feature-based image registration by means of the CHC evolutionary algorithm." *Image and Vision Computing* 24, no. 5 (May 1, 2006): 525-533.
104. Moghaddam, Baback, Tony Jebara, and Alex Pentland. "Bayesian Modeling of Facial Similarity." *Advances in Neural Information Processing Systems (NIPS 11)*, 1998.
105. Warfield, S.K., K.H. Zou, and W.M. Wells. "Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation." *Medical Imaging, IEEE Transactions on* 23, no. 7 (2004): 903-921.
106. Martin-Fernandez, Bouix, Ungar, McCarley, and Shenton. "Two Methods for Validating Brain Tissue Classifiers." *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2005* (2005).
107. "SciPy." http://www.scipy.org/
108. Hunter, John D. "Matplotlib: A 2D Graphics Environment." *IEEE Computing in Science & Engineering* 9, no. 3 (June 2007): 90-90.
109. Pérez, Fernando, and Brian E. Granger. "IPython: A System for Interactive Scientific Computing." *IEEE Computing in Science & Engineering* 9, no. 3 (June 2007): 21-99.
110. A. F. Laine, Wavelets in temporal and spatial processing of biomedical images *Annu. Rev. Biomed. Eng.* pp. 511–50, 2000.
111. Unser, M., A. Aldroubi, and A.F. Laine, Wavelets in Medical Imaging *IEEE Tran. Med. Imag.*, vol. 22, no. 3 (2003) 285-288.
112. Jebara, T. and T. Jaakkola. "Feature Selection and Dualities in Maximum Entropy Discrimination" in *the 16th Conference on Uncertainty in Artificial Intelligence*, UAI 2000.
113. Jebara, T. Images as Bags of Pixels. ICCV 2003.