

Published in final edited form as:

IEEE Trans Med Imaging. 2007 April ; 26(4): 471–478. doi:10.1109/TMI.2003.819294.

Simultaneous Registration and Parcellation of Bilateral Hippocampal Surface Pairs for Local Asymmetry Quantification

Nicholas A. Lord*, Jeffrey Ho, Baba C. Vemuri, and Stephan Eisenschenk
University of Florida, Gainesville, FL 32611 USA

Abstract

In clinical applications where structural asymmetries between homologous shapes have been correlated with pathology, the questions of definition and quantification of “asymmetry” arise naturally. When not only the degree but the position of deformity is thought relevant, asymmetry localization must also be addressed. Asymmetries between paired shapes have already been formulated in terms of (nonrigid) diffeomorphisms between the shapes. For the infinity of such maps possible for a given pair, we define optimality as the minimization of deviation from isometry under the constraint of piecewise deformation homogeneity. We propose a novel variational formulation for segmenting asymmetric regions from surface pairs based on the minimization of a functional of both the deformation map and the segmentation boundary, which defines the regions within which the homogeneity constraint is to be enforced. The functional minimization is achieved via a quasi-simultaneous evolution of the map and the segmenting curve, conducted on and between two-dimensional surface parametric domains. We present examples using both synthetic data and pairs of left and right hippocampal structures and demonstrate the relevance of the extracted features through a clinical epilepsy classification analysis.

I. Introduction

The temporal lobe's role in memory and learning makes it a natural point of focus in the study of neurodegenerative disease progression. In particular, the medical imaging literature abounds with attempts to identify and correlate abnormalities of the hippocampus with Alzheimer's disease, schizophrenia, and epilepsy. A precise analysis of such abnormalities could unlock the ability not only to track the progression of existent disease, but to identify at-risk individuals for preventative treatment. While early studies typically characterized hippocampal shape in terms of such simple global measures as volume, length, and surface area, it was shown as early as [19] that analysis of regional asymmetries could improve disease classification capability. Several methods for fine-grained regional hippocampal shape analysis have since been suggested. Gerig *et al.* [9] included a medial shape representation with age and drug treatment data in an exploratory statistical analysis of the hippocampus's link to schizophrenia. Shen *et al.* [13] conducted a statistical analysis based on the spherical harmonic (SPHARM) representation method. Styner *et al.* [14] tested the power of a SPHARM-based medial representation to separate monozygotic from dizygotic twins through lateral ventricular structure, and schizophrenics from normals through hippocampal and hippocampus-amygdalan structures. Davies *et al.* [7] devised a minimum description length framework for statistical shape modeling and extracted modes of variation between normal and schizophrenic populations. Bouix *et al.* [2] employed medial surfaces in

a local width analysis. Using a viscous fluid flow model, Csernansky *et al.* [6] computed diffeomorphic maps of patient hippocampi onto a reference, producing a dense inward/outward deformation field over each hippocampal surface. The surface itself was additionally manually segmented to allow for the regional comparison of the deformation fields as part of an attempt to separate healthy individuals from those exhibiting dementia of the Alzheimer's type. All reported results have indicated the benefit of incorporating regional information into the analysis. This is not surprising, as simple scalar measures of entire structures necessarily discard a wealth of information concerning the full characterization of those structures. However, statistical results thus far obtained are preliminary, and shape analysis results can actually appear contradictory across different studies (for instance, Styner notes in [15] the contrast between the primary abnormality localization in the hippocampal tail found in that work and the localization in the head reported in [5]). Furthermore, while statistically significant differences of hippocampal shape have been identified between diseased and normal sample populations, reliable classification of a sizeable number of patients with respect to those categories has not yet occurred. The problem thus remains very much open.

The need to identify and possibly isolate subregions of interest on the hippocampal surface suggests segmentation (henceforth synonymous with “parcellation”) and the need to compare hippocampal surfaces between and within (i.e., detection of asymmetry between halves of the same structure) patients represents an inherent requirement for registration. A unified segmentation and registration scheme is thus a natural approach to the problem at hand. The first such scheme known to us was given by Yezzi *et al.* [1], who unified mutual information (MI)-based flat 2-D rigid image registration with a level set implementation of a piecewise constant (Chan–Vese) segmentation scheme through a variational principle. Wyatt *et al.* [21] accomplished rather the same thing, solving instead a maximum *a posteriori* (MAP) problem in a Markov random field (MRF) framework. Xiaohua *et al.* [22] extended this to the nonrigid registration case, and Unal *et al.* [16] addressed the same problem using coupled partial differential equations (PDEs). In [18], Wang *et al.* presented a simultaneous nonrigid registration and segmentation technique for multimodal data sets. Jin *et al.* [10] simultaneously evolved a surface segmentation and a radiance-discontinuity-detecting segmentation of the evolving surface itself. Young *et al.* [23] combined joint segmentation and (scaled rigid) registration with morphing active contours for segmentation of groups of CT images. Pohl *et al.* [11] produced an expectation maximization (EM)-based algorithm for simultaneous affine registration and subcortical segmentation of MRIs using a labelled cortical atlas.

With the relevant background established, we now present a novel scheme for shape comparison that rests on the fundamental assumptions that: 1) the deformity of homologous anatomical structures can be quantified as the deviation from isometry of the deformation map between their surfaces and 2) the evolution of the global correspondence must allow for a partial disconnection between the normal and abnormal regions, as these are by definition not expected to exhibit the same deformation patterns. As such, surface segmentation rests at the heart of our approach, which can be viewed both as a registration tool and an asymmetry localizer. For the topologically spherical surfaces considered here, our intrinsic approach leads naturally to an elegant 2-D parametric representation of both the segmentation and registration. Our use of Riemannian surface structure in the matching criterion is similar in spirit to work done by Wang *et al.* (e.g., [20]) in which the authors presented a technique to nonrigidly register surfaces using a 2-D (parametric) diffeomorphic map constructed from Riemannian surface structure information. The integrated manifold segmentation is enough to distinguish our approach from this work, but additionally, our definition of system energy in terms of piecewise isometry (i.e., using first fundamental forms in the matching criterion) inherently balances the consideration that the map should accurately match Riemannian

surface characteristics with the idea that the map should incur as little deformation as possible in doing so. As such, we eliminate the need for additional regularization of the map gradient. In fact, the real novelty of the algorithm is more fundamental: this is the first proposal for simultaneous nonrigid surface registration and segmentation wherein the segmentation is driven by the evolving registration map itself.

The rest of the paper is organized as follows: Section II contains the mathematical formulation and algorithm conceptualization. In Section III, we delve into specifics of the implementation. Section IV contains experimental results on both synthetic and real data sets. Finally, in Section V, we draw conclusions.

II. Simultaneous Segmentation and Registration Algorithm

A. Differential Geometric Framework

To begin the explanation of our framework, we first turn our attention to the already-studied problem of 2-D image pair segmentation and nonrigid registration. A solution to this problem (regardless as to how one obtains it) clearly takes the form of: 1) a differentiable one-to-one and onto mapping function which shares its domain with one of the images and its range with the other and 2) a segmenting curve which lies within one image domain and is carried into the other through the mapping function. With images, there is a natural and obvious basis for computing this solution: the image intensity values. Intensity agreement represents the matching criterion by which algorithms such as mutual information maximization estimate pairwise image registration. Statistics derived from intensity values form the driving force of region-based segmentation methods (in Chan–Vese [3], for instance, the statistic is the mean of each curve-defined region). Further, differential qualities (e.g., gradient, Laplacian) of the intensity profiles can prove useful in identifying edges, corners, and other useful features. When conducting pairwise segmentation and registration of shapes rather than images, however, the appropriate measures are perhaps less obvious. It remains to us to define them.

Before defining a matching criterion, we first show how it is possible to cast the shape registration and segmentation problem into a computational framework similar to the image framework described above (save for the matching criterion itself) through appropriate parametrization. To accomplish this, one need note only the following simple topological facts: 1) hippocampi can be thought of as 2-D Riemannian manifolds embedded in \mathbb{R}^3 ; 2) these manifolds are topologically equivalent to the sphere (closed and genus zero); 3) the exclusion of a pair of poles from the sphere produces a cylindrical topology; and 4) any surface topologically equivalent to the cylinder can be parametrized by a class of functions whose common domain is a single rectangular patch with periodic boundary conditions at a pair of opposite sides. Note that this last fact directly implies that any function defined on the surface in question can be fully represented as a function over a closed rectangular domain with the stated boundary conditions. In particular, a smooth map between two hippocampi can be represented as a smooth map (deformation field) between two parametric rectangular domains satisfying the boundary conditions. Through this simple and natural representation of the surfaces, our solution format has become quite similar to that of the 2-D image segmentation and registration problem in that all of the computations are performed on and between rectangular domains. In contrast, when computation takes place in the embedding (3-D) space before being restricted to an approximation to the surface within that space, one cannot draw nearly so direct a parallel.

Clearly, the matching criterion which we define in place of image intensity must communicate information about the shapes being registered and segmented. But a well-known theorem due to Bonnet (whose formal statement we omit for the sake of brevity: see

for instance [8]) tells us that the first and second fundamental forms (FFF and SFF) contain *all* surface shape information. For the 2-D manifold case, these quantities are defined as follows:

$$FFF = \begin{pmatrix} X_u \cdot X_u & X_u \cdot X_v \\ X_u \cdot X_v & X_v \cdot X_v \end{pmatrix} \quad (1)$$

$$SFF = \begin{pmatrix} N \cdot X_{uu} & N \cdot X_{uv} \\ N \cdot X_{uv} & N \cdot X_{vv} \end{pmatrix} \quad (2)$$

where \mathbf{X} is the 3-D surface map, and u, v are the map parameters (equivalently, the coordinate system of the parametric domain), N is the surface normal (defined as $(X_u \times X_v) / |X_u \times X_v|$), and subscripts denote differentiation. There are different ways to conceptually distinguish the two quantities. For one, the geometric information contained in the FFF is of first order, while that of the SFF is second order in nature. Alternately, the FFF encodes the intrinsic geometry of the parameterized surface while the SFF encodes the extrinsic geometry, or, the FFF encodes local length information, the SFF surface curvature (its trace, the mean curvature, is used as a matching criterion in [20]). Because of its lower differential order (and thus greater relative stability) and intrinsic nature, we choose to base our approach on comparison between the FFF tensors of the two surface parametric domains. Given topological equivalence, FFF equivalence everywhere through a bijective deformation map defines that map as a global isometry: this indicates that the extent of failure to match this characteristic serves well as a measure of shape dissimilarity.

B. System Energy Functional

We now formalize the problem and fully specify its solution as follows. Let \mathbf{S}_1 and \mathbf{S}_2 be two surfaces in \mathbb{R}^3 . The Euclidean metric in \mathbb{R}^3 induces Riemannian metrics g_1 and g_2 on \mathbf{S}_1 and \mathbf{S}_2 respectively. The goal is to segment regions in \mathbf{S}_1 and their corresponding regions in \mathbf{S}_2 while mapping those correspondents so as to optimally match their metric structures. Specifically, the algorithm must compute a homeomorphism f between \mathbf{S}_1 and \mathbf{S}_2 and a set of closed curves γ_1 on \mathbf{S}_1 and γ_2 on \mathbf{S}_2 . The restriction of f to the complement $\mathbf{S}_1 \setminus \gamma_1$ is a diffeomorphism between $\mathbf{S}_1 \setminus \gamma_1$ and $\mathbf{S}_2 \setminus \gamma_2$. If U_1, \dots, U_j and V_1, \dots, V_n denote the collections of open components (topological discs) of $\mathbf{S}_1 \setminus \gamma_1$ and $\mathbf{S}_2 \setminus \gamma_2$, respectively, then f maps U_i diffeomorphically onto V_i for each i while matching the Riemannian structures of each U_i and $V_i = f(U_i)$.

We solve the simultaneous segmentation and registration problem outlined above using a variational framework. The energy functional \mathcal{E} is defined as a functional of a pair (f, γ_1) , where γ_1 is a set of closed curves on \mathbf{S}_1 and $f: \mathbf{S}_1 \rightarrow \mathbf{S}_2$ is a homeomorphism which is C^∞ on $\mathbf{S}_1 \setminus \gamma_1$. Let $f^* g_2$ denote the pull-back metric on \mathbf{S}_1 . $\mathcal{E}(f, \gamma_1)$ is then defined as

$$\mathcal{E}(f, \gamma_1) = \int_{\mathbf{S}_1 \setminus \gamma_1} e dA + \alpha \int |d\gamma_1(t)/dt|_{g_1} dt + \beta \left(\int_{\mathbf{S}_{1in}} |e - \bar{e}_{in}|_{g_1}^2 dA + \int_{\mathbf{S}_{1out}} |e - \bar{e}_{out}|_{g_1}^2 dA \right) \quad (3)$$

where

$$e = |f^* g_2 - g_1|_{g_1}^2 \quad (4)$$

$$\bar{e}_{\text{in}} = \frac{\int_{\mathbf{S}_{\text{lin}}} e \, dA}{\int_{\mathbf{S}_{\text{lin}}} dA}, \quad \bar{e}_{\text{out}} = \frac{\int_{\mathbf{S}_{\text{out}}} e \, dA}{\int_{\mathbf{S}_{\text{out}}} dA}. \quad (5)$$

\mathbf{S}_{lin} and \mathbf{S}_{out} are the regions of \mathbf{S}_1 inside and outside of γ_1 , respectively, and α, β are positive weighting constants.

The quantity e defined in (4) provides, at each point on \mathbf{S}_1 , a measure of similarity between the Riemannian structures g_1 and g_2 as they correspond under the current estimate of f . As such, it represents the local deviation from isometry, and its integral over the domain (we have adopted the convention of using \mathbf{P}_1 , the parametric domain of \mathbf{S}_1) provides a global measure of how far the given deformation f is from an isometry. This is the first term in the system energy represented in (3). In local (para-metric) coordinates, it is given by the distance between the two matrices G_1 and $J^T G_2 J$

$$|J^T G_2 J - G_1|_{g_1}^2 \quad (6)$$

where J is the Jacobian of f when f is expressed from the local coordinate system of \mathbf{P}_1 onto that of \mathbf{P}_2 , and G_1 and G_2 are 2×2 positive definite matrices expressing the metrics g_1 and g_2 in the coordinate systems of \mathbf{P}_1 and \mathbf{P}_2 , respectively, as g_2 in (1). The second term in \mathcal{E} restricts the length of the segmenting curve, so as to achieve the minimal necessary segmentation (shorter length and/or fewer open components) under the other energy constrains. Finally, the third and fourth terms place a homogeneity constraint on e within each connected component defined by the segmenting curve. This corresponds to our notion that the curve should partition the surfaces in such a way as to isolate regions exhibiting one sort of deformation characteristic from those exhibiting another. Note that these terms are defined in terms of both f and f_1 , and as such the minimization of their energy will drive both segmentation and registration. This coupling represents the fundamental difference between a simultaneous framework and one in which the segmentation follows the registration (in series).

III. Implementation of Energy Functional Minimization

The minimization of the system energy defined by (3) is achieved by alternating between the deformation map and segmentation estimation processes. When optimization steps are constrained to be sufficiently small, this alternation approximates minimization of \mathcal{E} simultaneously with respect to both f and γ_1 . The Euler-Lagrange equation for (3) can be derived through standard calculus of variations. The equation for f contains a fourth-order expression involving f and its derivatives. Because of the high differential order of this analytic expression, we instead choose to minimize the functional \mathcal{E} with respect to f directly through a constrained optimization process (more detail in Sections III-B and III-C). Minimization with respect to γ_1 , on the other hand, is implemented via a level set segmentation as described by Chan and Vese in [3]. The only difference between this module of our method and that described therein is our generalization of the process to genus zero 2-manifolds, which entails respecting nonuniform surface length and area elements in accordance with the metric information.

A. Hippocampal Surface Data Acquisition and Formatting

To produce the surface data used in this study, MRI scans were first acquired with a 1.5 T MRI scanner (Siemens Magnetom 42SPA, Siemens Medical Systems, Iselin, NJ) using 3-D magnetization-prepared rapid gradient echo (MPRAGE) sequences. The gradient echo

images were transferred for postprocessing. The sequence parameters included the following: obtained in sagittal plane, 250-mm field of view, repetition/echo times of 10/4 ms, T1 weighted sequence of 300 ms, 10° flip angle, 130 × 256 matrix, and 180-mm slab with 128 partitions producing 1.25-mm gapless sections. The segmentation of the hippocampal surfaces from the MRI data was done manually by a trained neuroanatomist, and a smooth surface was fit to the marked points by using a deformable pedal surface as described by Vemuri and Guo in [17]. Each segmented hippocampal surface obtained from the application of this technique was represented by a 40 × 21 mesh of points on that surface, periodic in one direction. (This mesh is the instantiation of the surface parametric domain as discussed in Section II-A.) Homologous hippocampi were brought into rigid alignment through application of the iterative closest points (ICP) algorithm following the segmentation. The rigid alignment step included extraction of volume data and normalization of the shapes with respect to this data.

B. Discrete Representation of Deformation Map and Segmentation Curve

An intrinsic map between two structures that are topologically equivalent to cylinders can be stored (in discretized form) as a two-vector function over an $m \times n$ grid, where m and n are the grid point counts in the u and v parameter directions. The two-vector at each grid point on the domain grid represents the displacement vector between the domain grid point and the corresponding target grid point. Since we are representing a cylinder, which is periodic in one direction, with a grid, we must enforce this periodicity through a modulo arithmetic implemented with respect to one of the grid coordinate directions (u). The grid must then be bounded in the remaining (v) parametric direction, with the two bounded ends corresponding to the “top” and “bottom” of the cylinder being parameterized. The coordinates at the poles (columns 1 and n of the grid) are excluded from the map evolution process, and so we conduct energy minimization with respect to mn displacement u coordinates and $m(n - 2)$ displacement v coordinates, which can be stored as a single variable vector of size $2m(n - 1)$. Note that this optimization is constrained in that the v component must everywhere with respect to the boundaries corresponding to the top and bottom of the cylinder. We choose also to restrict the modular family of solutions for the u components (corresponding to 360° rotations around the cylindrical axis) to falling within a modulus above or below the initial conditions. We evolve the system under these constraints using the LargeScale version of Matlab's “fmincon” function, a subspace trust region method based on the interior-reflective Newton method described in [4]. We exploit the inherent sparsity structure of the Hessian (as described in Section III-C). The Riemannian characteristics themselves, whose values of course drive the energy minimization process, are calculated through analytical differentiation of a bicubic surface fit to the given points.

The segmenting curve is simply stored and evolved according to well established level set curve evolution techniques (see [12] for extensive discussion). The one modification present in our case is that length and area computations respect the distances between grid points by using the appropriate Riemannian metric at each point (as opposed to the uniform Euclidean metric used for flat domains and targets).

C. Sparsity of Hessian

The trust region method which conducts the energy minimization with respect to the map (for a fixed segmentation) can exploit the fact that local changes in the map can only effect local changes in the induced metric structure. To make this point clear, let us consider the fact that the FFF is defined as in (1), where X is the surface spatial location 3-vector and u and v are the 2-D parametric coordinates. The FFF at a point depends only on the derivatives of X with respect to u and v . Furthermore, for a fixed segmentation, the energy itself is a function only of the FFF structure on the target manifold imposed by the evolving map. But

consider each component of the $2m(n-1)$ vector which we are evolving: infinitesimal variation in one of these components can only perturb the metric structure on the target manifold locally, as the FFF definition shows that it depends only on X_u and X_v at the given (u, v) , and these partial derivatives themselves are inherently local. It should then be immediately clear that mixed second partials can only be nonzero when both variables being differentiated against lie within the same neighborhood (as defined by the differentiation scheme): this leads to the Hessian structure depicted in Fig. 2, wherein the solid diagonals represent the nonzero matrix elements (and the map variable vector structure was explained in Section III-B). Exploitation of this sparsity structure is necessary for a second-order optimization scheme to achieve computational feasibility.

IV. Results

To illustrate the method's performance, we first present a visual display of its output and follow by demonstrating the classification power of the extracted features. The input dataset consisted of 60 L/R hippocampus pairs, collected and formatted as described in Section III-A. Clinicians provided a trinary classification of this dataset: LATL (left atrophied temporal lobe) epileptic, RATL (right atrophied temporal lobe) epileptic, and Control. In this step, six samples were discarded due to ambiguity of their clinical class membership, leaving the final set at 54 patient samples (15 LATL, 16 RATL, 23 Control).

A. Output Visualization

Our result visualization format can best be understood through reference to Fig. 1, in which we present a holistic schematic. For each sample, we illustrate: 1) the deformation map f as the warp necessary to carry points in \mathbf{P}_1 to their correspondents in \mathbf{P}_2 ; 2) the segmentation γ_1 and the deformation energy function e as they exist in \mathbf{P}_1 ; and 3) the segmentation γ_1 as it exists on \mathbf{S}_1 through the parametrizing function from \mathbf{P}_1 (there is a corresponding region on \mathbf{S}_2 as well, not depicted here but fully specified for given f and γ_1 through the parametrizing functions: γ_2 is specified by f and γ_1). As noted in Fig. 1, the segmentation is depicted on through the shading of its interior and exterior with red and blue, respectively; the curve(s) itself is then of course the boundary between the shaded regions.

We first consider the case of a pair of cylinders, where one member of the pair has had its surface distorted according to an outward normal vector field of magnitude dictated by a Gaussian function confined to a known support. Since the support is known, we can use its segmentation as one quantification of accuracy. Fig. 3 illustrates the obtained results. Here, 86% of the distortion area is included in the segmentation and 0% of the undistorted area is included in the segmentation. This means that the segmenting curves lie slightly within the distorted regions, as can be seen in the figure. Even the 14% "misclassification" cannot necessarily be seen completely as error in this context: since the protrusions are far less severe towards the boundaries of their support (the feet of the hills), they are accordingly far more similar in deformation level to the undistorted regions than to the heavily distorted regions towards the support centers. As such, the segmentation arrived at is in compliance with the two-mean framework (which produces a high/low segmentation). While we do later note that it is possible to refine this framework (see conclusion), it is important to understand that this manner of grouping is not inherently a limitation, particularly considering the intended application. We follow by demonstrating robustness of the segmentation in the face of reparametrization of the surface of comparison and random noise in the surface point cloud (see caption for parameter levels). When repeated under these conditions, the segmentation remains 100% consistent at the pixel resolution with that obtained prior.

For the real cases described previously, we present the results obtained on one member of each class (RATL, LATL, and control), as columns of Fig. 4. Each figure depicts a mirrored left hippocampus on the left side of each panel and the corresponding right hippocampus to which it is compared on the right side. This aspect of the study (inherently) lacks ground truth, and a “correct” result is thus one in which the segmented regions indeed appear to be the portions of one hippocampus which are shaped least like the corresponding portions of the hippocampus to which they are being compared. This is readily evident in the presented figures. The provision of an input device to allow for expert manual segmentation (closed curves confined to the surface) for validation is itself a research problem. This does not mean, however, that we cannot evaluate the quality of the data extracted from the real set. We present a classification analysis as evidence.

B. Classification Analysis

As it is clear that volume does indeed contain information relevant in the shape comparison of homologues, it is not surprising that it is a significant feature in distinguishing epileptics from controls. However, it is equally evident that the complex process of neurodegeneration associated with epilepsy cannot be captured in full so simply. The utility of volume as a sole basis for classification is thus inherently limited. We have provided a fine-grained shape difference characterization with a mind to overcoming precisely this limitation. The test performance of a classifier should thus increase when our extracted features are included as input. Note that we have worked with volume-normalized data to deliberately orthogonalize the extracted features with respect to volume. Since volume itself is assumed to be a relevant feature, it is possible that the best test classification accuracy will be obtained by using volume and local shape difference information in conjunction.

Since we produce a segmented energy function over the entire surface (parametric) domain, we have options as to how to present features to the classifier algorithm. We can collect summary statistics that profit from both segmentation and deformation quantification, such as mean deformation energy [e as defined in (4)] inside and outside of the segmenting curves. Alternately, once the shapes have been volume-normalized and mutually registered, we can compare the (quasi-)continuous e functions across patients by forming feature vectors from the e values at corresponding sample locations. There is also an array of supervised learning methods from which to select a preferred classification algorithm; we choose the ones observed to give best test performance on all feature sets examined, including volume alone (as further detailed in the following).

Fig. 5(a) demonstrates the success rates in classifying controls versus epileptics (LATL and RATL groups combined into a single set) and Fig. 5(b) demonstrates analogous results for the problem of separating RATL and LATL members. The training and test statistics were collected through a standard leave-one-out cross validation procedure: test results should be regarded as the reliable performance indicators (and high discrepancy between training and test percentages as evidence of overfitting). The feature vectors being compared are: 1) volume alone; 2) volume with a set of six summary statistics derived from the algorithm output (total area within the curve(s), total area outside of the curve(s), total energy e within the curve(s), total energy e outside of the curve(s), mean energy e within the curve, mean energy e outside of the curve); 3) the function e at 600 sampled locations; and 4) a concatenation of feature vectors (1), (2), and (3). The chosen abbreviations for these feature vectors are “VOL,” “VOL+6SUM,” “E,” and “E+VOL+6SUM,” respectively. The classifiers used are: 1) support vector machine with polynomial basis (SVM w/PB); 2) support vector machine with radial basis (SVM w/RB); and 3) kernel Fisher discriminant with polynomial basis (KFD w/PB).

The results of both studies make the method's success clear. In the case of distinguishing epileptics from controls, all experimental feature vectors demonstrated superior test classification accuracy, with a maximum performance of 88.89% on “E VOL+6SUM” as classified by SVM w/PB, as compared to 79.56% for “VOL” as classified by SVM w/RB. Note in particular that feature vector “E” outperforms “VOL” as well, despite the fact that “E” contains *no* volume information. This portion of the experiment confirms our notion that local distortion information is an important complement to volume data and suggests that it can in some cases outweigh volume in independent relevance. While we do not observe this latter phenomenon in the LATL versus RATL case, we do still observe substantial outperformance of “VOL” by “VOL+6SUM” (90.32% versus 80.65%), again demonstrating the algorithm's extraction of clinically relevant shape information not represented by volume. The large difference between training and test accuracies for the other two feature vectors indicates overfitting, which is enabled by their higher dimensionalities. This indicates that an overly fine-grained approach can confound classification on this set, but that subregion identification and characterization remains nonetheless crucial. We have not come across clinical test classification accuracies of this level in any of the literature we have surveyed.

V. Conclusion

We have presented a novel scheme for simultaneous nonrigid registration and segmentation of homologous shapes, wherein the registration and segmentation processes are driven by intrinsic geometric characteristics of the shapes themselves. As such, we are able not only to identify surface pairs representing large deformations, but also to specify which subregions of the shapes appear most likely to be involved in the deformation, to produce an estimate of the implicit deformation field, and to quantify the deformation energy of the segmented subregions separately from that of the remaining surface patches. We have noted that the method's results conform to ground truth in cases where it can be meaningfully defined and to expectation in cases where it cannot. The most important validation comes in the form of our substantial improvement of disease classification accuracy relative to standard volumetric analysis. In the case of epileptics versus controls, we have provided evidence that volume-independent local shape information is actually a more significant predictor than volume itself.

It should be noted that the two-region Chan–Vese-inspired framework is merely a good “first draft” of a segmentation scheme for this sort of problem. Extending the functional to the multiple mean case, or to that of generalized Mumford-Shah (of which the Chan-Vese approach is a special case), would potentially enable the identification of intermediate deformation levels currently being grouped with their nearest matches in the high/low system. The FFFs can also be compared using a norm other than Frobenius. Finally, all evolution-based methods (apart from the most trivial cases) suffer from some amount of initialization dependence. Global optimization schemes could be expected to further boost the algorithm's performance.

Acknowledgments

Data sets were partially provided by Dr. C. Leonard of the UF McKnight Brain Institute. Thanks to S. Kodipaka for his invaluable assistance in the data analysis portion of this work.

The work of B. Vemuri was supported in part by the NIH under Grant R01-NS046812. The work of N. A. Lord was supported by a UF Alumni Fellowship.

REFERENCES

1. Yezzi A, Zollei L, Kapur T. A variational framework for joint segmentation and registration. *IEEE-MMBIA*. 2001:44–51.
2. Bouix S, Pruessner JC, Collins DL, Siddiqi K. Hippocampal shape analysis using medial surfaces. *NeuroImage*. 2005; 25(4):1077–1089. [PubMed: 15850726]
3. Chan TF, Vese LA. Active contours without edges. *IEEE Trans. Image Processing*. Feb.; 2001 10(2):266–277.
4. Coleman TF, Li Y. An interior, trust region approach for nonlinear minimization subject to bounds. *SIAM J. Optimization*. 1996; 6:418–445.
5. Csernansky JG, Wang L, Jones D, Rastogi-Cruz D, Posener JA, Heydebrand G, Miller JP, Miller MI. Hippocampal deformities in schizophrenia characterized by high dimensional brain mapping. *Amer. J. Psychiatry*. 2002; 159(12):2000–2006. [PubMed: 12450948]
6. Csernansky JG, Wang L, Swank J, Miller JP, Gado M, McKeel D, Miller MI, Morris JC. Preclinical detection of alzheimer's disease: Hippocampal shape and volume predict dementia onset in the elderly. *NeuroImage*. 2005; 25(3):783–792. [PubMed: 15808979]
7. Davies RH, Twining CJ, Allen PD, Cootes TF, Taylor CJ. Shape discrimination in the hippocampus using an mdl model. *Proc. IPMI*. 2003:38–50.
8. do Carmo, MP. *Differential Geometry of Curves and Surfaces*. Prentice-Hall; Englewood Cliffs, NJ: 1976.
9. Gerig G, Muller KE, Kistner EO, Chi Y-Y, Chakos M, Styner M, Lieberman JA. Age and treatment related local hippocampal changes in schizophrenia explained by a novel shape analysis method. *Proc. MICCAI*. 2003:653–660.
10. Jin H, Yezzi AJ, Soatto S. Region-based segmentation on evolving surfaces with application to 3D reconstruction of shape and piecewise constant radiance. *Proc. ECCV*. 2004; 2:114–125.
11. Pohl KM, Fisher J, Grimson WEL, Kikinis R, Wells WM. A Bayesian model for joint segmentation and registration. *NeuroImage*. 2006; 31(1):228–239. [PubMed: 16466677]
12. Sethian, JA. *Level Set Methods and Fast Marching Methods*. Cambridge Univ. Press; Cambridge, MA: 1996.
13. Shen L, Ford J, Makedon F, Saykin A. Hippocampal shape analysis: Surface-based representation and classification. *SPIE Proc. Med. Imag. 2003: Image Process*. 2003:253–264.
14. Styner M, Gerig G, Lieberman J, Jones D, Weinberger D. Statistical shape analysis of neuroanatomical structures based on medial models. *Med. Image Anal.* 2003; 7(3):207–220. [PubMed: 12946464]
15. Styner M, Lieberman JA, Gerig G. Boundary and medial shape analysis of the hippocampus in schizophrenia. *Proc. MICCAI*. 2003:464–471.
16. Unal GB, Slabaugh GG. Coupled PDEs for non-rigid registration and segmentation. *Proc. CVPR*. 2005:168–175.
17. Vemuri BC, Guo Y. Snake pedals: Compact and versatile geometric models with physics-based control. *IEEE Trans. Pattern Anal. Mach. Intell.* 2000; 22:445–459.
18. Wang F, Vemuri BC. Simultaneous registration and segmentation of anatomical structures from brain MRI. *Proc. MICCAI*. 2005:17–25.
19. Wang L, Joshi SC, Miller MI, Csernansky JG. Statistical analysis of hippocampal asymmetry in schizophrenia. *NeuroImage*. 2001; 14(3):531–545. [PubMed: 11506528]
20. Wang Y, Chiang M-C, Thompson PM. Mutual information-based 3D surface matching with applications to face recognition and brain mapping. *Proc. ICCV*. 2005:527–534.
21. Wyatt PP, Noble JA. Map MRF joint segmentation and registration of medical images. *Med. Image Anal.* 2003; 7(4):539–552. [PubMed: 14561557]
22. Xiaohua C, Brady JM, Rueckert D. Simultaneous segmentation and registration of medical images. *Proc. MICCAI*. 2004:663–670.
23. Young Y-N, Levy D. Registration-based morphing of active contours for segmentation of CT scans. *Math. Biosci. Eng.* 2005; 2:79–96. [PubMed: 20369913]

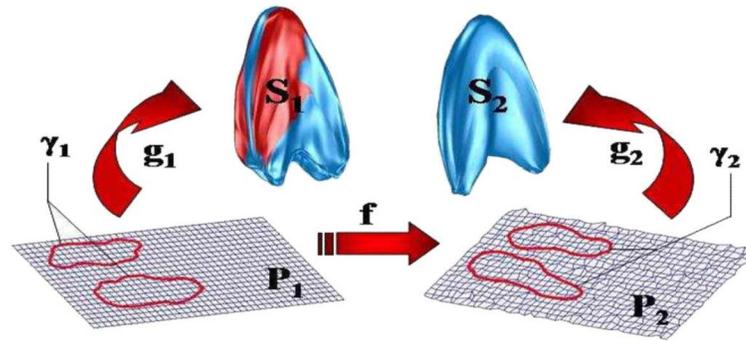


Fig. 1.

Illustration of framework. Homologous shapes S_1 and S_2 are associated with flat rectangular patches (with appropriate BCs) P_1 and P_2 respectively. Patches store all Riemannian metric information g_1 and g_2 of their respective surfaces in terms of FFF matrices G_1 and G_2 . Metric information is then matched between P_1 and P_2 through a homeomorphic map f of P_1 onto P_2 (diffeomorphic except possibly on the curves γ) and segmenting curve γ_1 in P_1 which is carried by f into P_2 as γ_2 . S_1 and S_2 are thus registered and segmented via parametric correspondence. Note that f can be visualized as deformation of P_1 and that segmentation can be represented on S_1 by shading interiors red and exteriors blue.

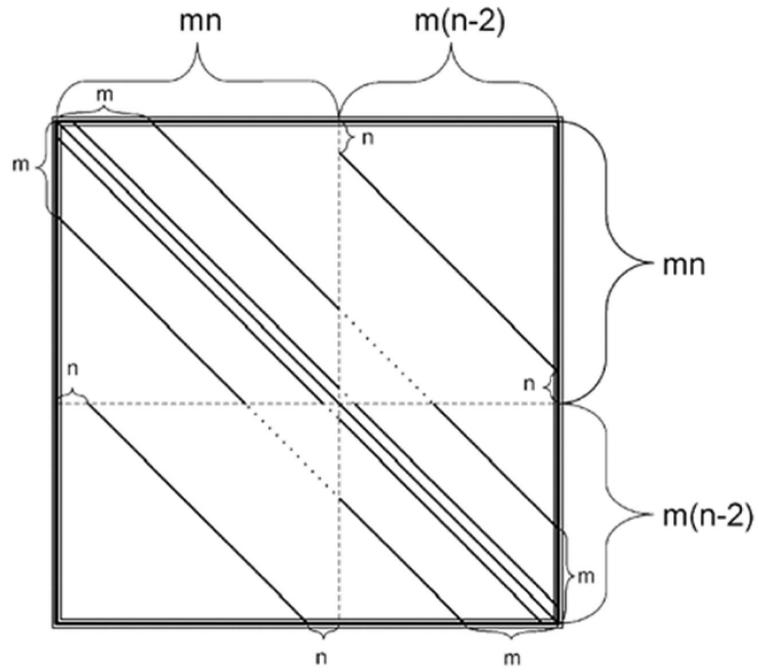


Fig. 2.
Hessian sparsity structure.

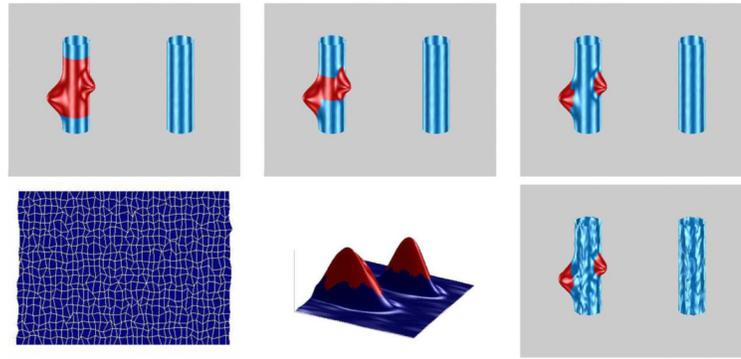


Fig. 3.

Synthetic results. Top: segmentation of prominent distortions from cylindrical surface (left: init.; middle: 5 its.; right: 20 its.). Bottom: performance under surface noise ($\sigma = 0.01$ for unit cylinder) and reparametrization (left: warp field, $\sigma = 0.2$ in each component; middle: segmented energy field in parametric domain; right: converged result, at 20 its.).

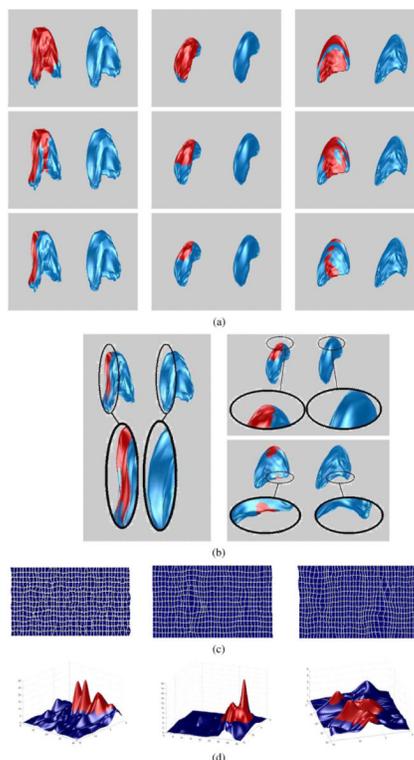


Fig. 4. Real results (left column: LATL member; middle column: RATL member; right column: control member). (a) Segmentation evolution on sample patients (top to bottom: initialization, three, ten iterations). (b) Converged segmentations (left: LATL member, top right: RATL member, bottom right: control member). (c) Warp of parametric domain induced by evolution process. (d) Final segmented distortion energy e graph as function over parametric domain.

	SVM w/ PB		SVM w/ RB		KFD w/ PB	
	Training	Test	Training	Test	Training	Test
VOL	79.59	77.78	79.63	79.56	77.81	77.78
VOL+6SUM	92.31	81.48	88.68	81.48	80.43	77.78
E	100.00	85.19	94.44	81.48	94.51	87.04
E+VOL+6SUM	100.00	88.89	96.26	81.48	99.97	85.19

(a)

	SVM w/ PB		SVM w/ RB		KFD w/ PB	
	Training	Test	Training	Test	Training	Test
VOL	83.25	80.65	80.65	80.65	79.78	77.42
VOL+6SUM	93.55	90.32	93.76	87.10	93.87	80.65
E	100.00	70.97	100.00	67.74	99.89	67.74
E+VOL+6SUM	100.00	74.19	100.00	70.97	93.23	67.74

(b)

Fig. 5.

Epilepsy classification results, optimal test accuracies over all classifiers shown in bold. (a) Control versus epilepsy. (b) LATL versus RATL.