SIAM J. IMAGING SCIENCES
Vol. 7, No. 3, pp. 1645–1668

Statistical Shape Analysis of Neuroanatomical Structures via Level-Set–based Shape Morphing*

Tammy Riklin Raviv†, Yi Gao‡, James J. Levitt§, and Sylvain Bouix§

Abstract. Groupwise statistical analysis of the morphometry of brain structures plays an important role in neuroimaging studies. Nevertheless, most morphometric measurements are often limited to volume and surface area, as further morphological characterization of anatomical structures poses a significant challenge. In this paper, we present a method that allows the detection, localization, and quantification of statistically significant morphological differences in complex brain structures between populations. This is accomplished by a novel level-set framework for shape morphing and a *multishape* dissimilarity-measure derived by a modified version of the Hausdorff distance. The proposed method does not require explicit one-to-one point correspondences and is fast, robust, and easy to implement regardless of the topological complexity of the anatomical surface under study. The proposed model has been applied to well-defined regions of interest using both synthetic and real data sets. This includes the corpus callosum, striatum, caudate, amygdala-hippocampal complex, and superior temporal gyrus. These structures were selected for their importance with respect to brain regions implicated in a variety of neurological disorders. The synthetic databases allowed quantitative evaluations of the method. Results obtained with real clinical data of Williams syndrome and schizophrenia patients agree with published findings in the psychiatry literature.

Key words. statistical shape analysis, modified Hausdorff distance, level-sets, brain morphology

AMS subject classifications. 35A15, 57N25, 62B99, 92C99

DOI. 10.1137/13093978X

1. Introduction.

The detection of groupwise neuroanatomical shape differences sheds new light on neurodevelopmental studies and on how neurodegenerative diseases impact brain morphology. Recent advances in brain magnetic resonance (MR) imaging and image analysis have led to numerous morphometric studies of brain anatomy in normal neurodevelopment [5, 41], the effect of handedness [68], and the presence of pathologies including Alzheimer's disease [9, 14], schizophrenia [30, 40, 61], and fetal alcohol exposure [29], to mention just a few.

Due to the inherent difficulty of properly representing the morphology of the structures of

*Received by the editors October 3, 2013; accepted for publication (in revised form) March 11, 2014; published electronically August 19, 2014. A preliminary version of this paper appeared as *Statistical shape analysis for population studies via level-set based shape morphing* in the proceedings *Computer Vision—ECCV* 2012. *Workshops and Demonstrations*, Springer-Verlag, Berlin, Heidelberg, 2012, pp. 42–51 [58]. This work was supported by NIH grant R01MH082918.

http://www.siam.org/journals/siims/7-3/93978.html

†Electrical and Computer Engineering Department, Ben-Gurion University, Beer-Sheva, 84105, Israel (rrtammy@ee.bgu.ac.il).

‡Department of Electrical and Computer Engineering, The University of Alabama at Birmingham, Birmingham, AL 35205 (gaoyi.cn@gmail.com).

§Psychiatry and Neuroimaging Laboratory, Brigham and Women's Hospital, Harvard Medical School, Boston, MA 02215 (james_levitt@hms.harvard.edu, sylvain@bwh.harvard.edu).

interest, the morphometric measures are often limited to volume and surface area, as well as other quantitative measures such as curvature, smoothness, and thickness. Yet, these features provide only a partial description of the anatomy and are often insufficient to define a clear distinction between populations.

Deformation-based morphometry (DBM) methods, e.g., tensor-based morphometry (TBM), and voxel-based morphometry (VBM) are image analysis tools for identifying regional structural differences from the gradients (Jacobian matrices) of deformation fields [1, 16]. In its classical form, VBM does not refer explicitly to specific anatomical structures but considers images as continuous scalar measurements and tests for local differences at a predefined spatial scale. The art of setting the smoothness kernel, which determines the scale at which spatial anatomical differences can be expressed, is a major limitation in VBM [77, 78]. Nevertheless, VBM methods gain their popularity in morphological studies, since they work directly on the grayscale images without having to extract the boundaries of particular structures of interest.

The shape of a structure, however, often provides rich anatomical characteristics, which may allow one to distinguish between groups of subjects, and could be related to differences in cognitive functioning, development, or neurological symptoms, for example, in cortical folding studies [79, 18]. In the last decades, there has been significant progress in the development of algorithms for use in computational anatomy based on statistical analysis of shapes.

There exist several general approaches to shape analysis. A common class of methods makes use of medial representations, which are based on a mathematical model for shape called the medial surface. This structure provides a compact description of a volumetric object while reflecting its local symmetries. However, establishing a one-to-one correspondence to compare different subjects is necessary, whether it is accomplished through registration [5] or modeling [57, 65, 69, 76].

A different shape analysis approach, which also relies on compact shape representation, uses global feature vectors determined by eigenfunctions as in the "Shape DNA" paradigm of [56, 50], spherical harmonics [6, 33], or invariant moment representations [43]. Such methods are usually numerically stable and allow for the computation of relevant statistics. However, the resulting set of feature vectors is rarely intuitive; therefore the interpretation of the results in terms of anatomical changes can be difficult.

The main body of shape analysis literature is based on the representation of an object's surface or interior, along with a study of the mechanical deformations required to transform one object into another [4, 13, 15, 70]. While providing intuitively interpretable results, this popular paradigm relies on nonlinear registration techniques to establish one-to-one correspondences between subjects. For an excellent survey on medical image registration, see [63] and references therein.

Surface representation approaches can be based either on preselected landmarks as in [19, 13] or, more commonly, on dense correspondences via shape parameterization. However, the hypothesis that a one-to-one correspondence must be established is not obvious unless the objects under study have similar shapes. The calculation of statistical significance from the recovered deformation fields also poses a challenge, as the number of tests is usually much larger than the number of subjects under study [26, 66]. These difficulties can be partially addressed by mapping different shape surfaces into a common parameterized surface. For example, the commonly used SPHARM-PDM (*spherical harmonics–point distribution*

model) [25, 67] solves the correspondence problem by defining a spherical parameterization for each of the objects in a shape population, assuming correspondence between equivalently parameterized points. In a similar fashion, mapping surfaces into cylinders was suggested in [72]. While these techniques are assumed to be more robust, they pose significant constraints on the topology and the overall structure of the shapes to be analyzed. In order to overcome this limitation, organs with complex structures are partitioned into segments such that the mapping for each shape segment is performed separately [72, 53]. Recently, an alternative solution to the uniform spherical parameterization constraint was suggested [52]. It is based on combining the SPHARM-PDM with the entropy-based particle system correspondence model [8].

A different thread of dense-correspondence algorithms is based on path-straightening to minimize geodesics on shape spaces [64]. In [36, 37] a generalization of this technique to three dimensions was used to solve reparameterization and registration problems simultaneously. While this approach seems to be promising, it involves an interleaved optimization procedure which requires a certain degree of structural simplicity and similarity between shapes to assure convergence.

The objective of the proposed study is the detection of morphometric differences in anatomical structures between different populations. In contrast to the above methods, the suggested framework exclusively relies on shape boundaries and boundary-based dissimilarities. Specifically, we use a modified and continuous version of the Hausdorff distance, which in its classical form defines a metric between point clouds of possibly different cardinality. This allows us to circumvent one-to-one point correspondences. Moreover, we show that defining a shape distance based on boundaries, rather than area (or volume) overlap, is more suitable for handling convoluted structures such as cortical brain regions.

The Hausdorff distance and the related Gromov–Hausdorff distance have been used before, e.g., in [7, 11, 21, 46], to define distances between objects. For example, in [7] the Gromov–Hausdorff distance was used for calculating the diffusion distance between points on a surface, allowing the comparison of pairs of nonrigid shapes with different topology. In contrast, here shapes are represented by their signed distance functions (SDFs) or, equivalently, by level-sets [51] in a manner similar to that of the seminal paper [42]. Moreover, the proposed method is not limited to a pair of shapes but considers a shape population and constructs, via level-set evolution, its center by minimizing the corresponding groupwise modified Hausdorff distance.

The level-set framework was originally suggested for image segmentation; see [51, 10, 54]. Usually, an image likelihood term, which is based on the piecewise smoothness model of the image features (gray-levels, colors, gradients, texture, etc.), dominates the evolution of the segmenting contour or three-dimensional (3D) surface. Here, unlike in the common approaches, the level-set paradigm is used for shape morphing and in particular for the construction of a shape that represents "the mean" of a given shape ensemble, represented by binary images. In this sense, the proposed approach is also different from the level-set–based shape averaging in [3], which considers the image gray-levels while minimizing the sum of square differences between shapes.

The main contribution of the proposed methodology is a variational framework, based on a novel level-set functional, for a groupwise shape morphing. A boundary-based shape distance term governs the evolution of the mean of a given set of shapes. Spatial statistical analysis is

obtained by calculating the minimal distances between the surface of the constructed mean shape and each of the input shapes. We then compute two-sample t-tests at every location on the mean shape surface to look for statistically significant differences (SSDs) between populations. The resulting raw p-values are adjusted for multiple comparison using the false discovery rate [49]. The suggested model allows for the extraction of statistics that are sensitive enough to detect subtle changes between populations, yet robust enough to avoid common statistical errors.

We test the proposed method on well-defined regions of interest using both synthetic and real data sets. The structures were selected for their importance with respect to brain regions implicated in schizophrenia and other neurological disorders. These include synthetic sets of the striatum and the amygdala-hippocampal complex (AHC) and real data of the corpus collasum (CC) in Williams syndrome (WS) patients, the superior temporal gyrus (STG) in first-episode schizophrenia patients, and the caudate nucleus in women with schizotypal personality disorder (SPD). We show that our algorithm can accurately detect, locate, and quantify known morphological changes. The results obtained for the real clinical data agree with published findings based on volumetric comparisons of the structures of interest. Moreover, smaller and more localized differences that could not be captured by volume measurements were detected.

This paper extends a previous report [58] by the inclusion of additional experimental results, examples, implementation details, and theoretical discussions.

The rest of the paper is organized as follows. In section 2 we introduce our novel mathematical shape model and the related unique level-set framework. Section 2.1 presents the continuous modified Hausdorff distance (MHD) and its use for pairwise shape dissimilarities. Sections 2.2 and 2.3 discuss the parametric transformation and elastic shape morphing, respectively. Section 2.4 introduces the mean shape construction algorithm. Implementation details are given in section 2.5. Population analysis is discussed in section 2.6. Experimental results are presented in section 3. We conclude in section 4.

2. Methods.

2.1. Shape representation and distance. Shape representation plays a major role in morphological analysis. A shape S_i is commonly represented by a point cloud $\{\mathbf{x}_i \in \mathbb{R}^N\}$, where N is the shape dimension. For example, $\mathbf{x}_i$ may stand for the coordinates $\{(x_i^v, y_i^v, z_i^v)\}_{v=1,\ldots,V_i}$ of the surface voxels of the ith object. Let S_i and S_j be two distinctive point sets which may have a different cardinality, i.e., $|S_i| \neq |S_j|$. The classical definition of the Hausdorff distance is as follows:

$$D_{\mathrm{H}}(S_i, S_j) = \max\left\{ \sup_{\mathbf{x}_i \in S_i} \inf_{\mathbf{x}_j \in S_j} d(\mathbf{x}_i, \mathbf{x}_j), \sup_{\mathbf{x}_j \in S_j} \inf_{\mathbf{x}_i \in S_i} d(\mathbf{x}_i, \mathbf{x}_j) \right\}, \tag{2.1}$$

where $d(x_i, x_j)$ is the Euclidean distance between a pair of points $x_i \in S_i$ and $x_j \in S_j$. In the proposed framework, we use a shape dissimilarity measure which is a modified version of the Hausdorff distance, replacing the maximum of the left- and the right-hand side terms in (2.1) and the supremum in each term by summation:

$$D_{\mathrm{MH}}(S_i, S_j) = \sum_{\mathbf{x}_i \in S_i} \inf_{\mathbf{x}_j \in S_j} d(\mathbf{x}_i, \mathbf{x}_j) + \sum_{\mathbf{x}_j \in S_j} \inf_{\mathbf{x}_i \in S_i} d(\mathbf{x}_i, \mathbf{x}_j). \tag{2.2}$$

While the MHD is not a metric distance as it violates the triangle inequality, it appears to be more robust than the classical Hausdorff distance in the presence of noise and irregularities [20]. Direct implementation of (2.2) requires curve or surface parameterization. While this is reasonably doable (at least for two-dimensional (2D) shapes) via active contour techniques such as the *snakes* [32], it is subject to significant limitations. For example, topological changes cannot be handled in an elegant manner, self-intersections and shocks cannot be avoided, and the extension to three dimensions is difficult. To avoid these obstacles, we choose to represent shape boundaries by their SDFs and utilize the level-set framework [51].

We next reformulate (2.2) using a continuous form. Let $\omega_i \in \Omega$ denote the image region which corresponds to the ith object, where $\Omega \in \mathbb{R}^3$ is the image domain. The boundary of ω_i is denoted by $\partial\omega_i$. Rather than using a discrete point set, we represent an object's shape by the SDF of its boundary: $\phi_{S_i}: \Omega \to \mathbb{R}$. Specifically,

$$\phi_i(\mathbf{x}) = \begin{cases} d_E(\mathbf{x}, \partial\omega_i) & \text{if} \quad \mathbf{x} \in \omega_i, \\ -d_E(\mathbf{x}, \partial\omega_i) & \text{if} \quad \mathbf{x} \in \Omega \setminus \omega_i, \end{cases} \tag{2.3}$$

where $d_E(.)$ denotes a Euclidean distance. We assume that $\phi(\mathbf{x})$ is differentiable almost everywhere, and its gradient satisfies the eikonal equation $|\nabla\phi(\mathbf{x})| = 1$. We define the distance between the shapes of objects i and j as the MHD between their boundaries, using the continuous form of (2.2):

$$D_{\mathrm{CMHD}}(i,j) = \int_{\partial\omega_i} |\phi_{S_j}(\mathbf{x})| ds + \int_{\partial\omega_j} |\phi_{S_i}(\mathbf{x})| ds, \tag{2.4}$$

where ds is an infinitesimal part of a curve (in two dimensions) or a surface (in three dimensions) which defines the shape's boundary. As ϕ_{S_j} is an SDF, the value of $|\phi_j(\mathbf{x})|$ represents the minimal Euclidean distance from $\mathbf{x}$ to the boundary of ω_j. Now, integrating along $\mathbf{x} \in \partial\omega_i$ provides the left-hand side of (2.4). Similarly, the right-hand side of (2.4) is obtained by integrating the values of $|\phi_{S_i}(\mathbf{x})|$ for $\mathbf{x} \in \partial\omega_j$. Figure 1 visually demonstrates the main concept introduced in (2.4). Suppose that both ω_i and ω_j are defined on the same image domain Ω. In Figure 1(a) the shape boundary $\partial\omega_i$ (green) is plotted over ϕ_{S_j} (color-coded), while in Figure 1(b) the boundary $\partial\omega_j$ (blue) is plotted over ϕ_{S_i} (color-coded). Now, "traveling" along the green contour, for example, and integrating the absolute values of ϕ_{S_j} expresses the left-hand side of (2.4). The same applies to the right-hand side of (2.4), represented by Figure 1(b).

2.2. Parametric transformations. As in [1], we define "shape" as a set of geometric features of an object that is invariant to 12-parameter (6-parameter in two dimensions) affine transformation. We therefore decouple shape deformation into a parametric and a nonparametric transformation. The latter, termed *morphing* here, involves local (usually small) perturbations of the shape boundaries and will be discussed in section 2.3.

In the proposed framework, a shape S_i is aligned to another shape S_j by a parametric transformation $T_{j\to i}$ that minimizes their MHD:

$$T_{i\to j} = \arg\min_{T_{i\to j}} D_{\mathrm{CMHD}}(S_j, T_{i\to j}(S_i)). \tag{2.5}$$

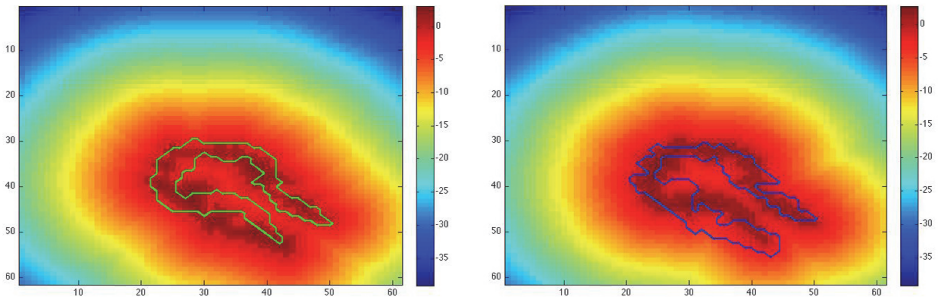

(a) (b)

Figure 1. Obtaining the MHD via signed distance functions. *The boundaries $\partial\omega_i$ and $\partial\omega_j$ are represented by the green* (a) *and blue* (b) *contours, while the values of ϕ_{S_j} and ϕ_{S_i} (shown, respectively, in* (a) *and* (b)*) are color-coded. Each of the terms in* (2.4) *is obtained by integrating the Euclidean distances of the image pixels (voxels in three dimensions) to the zero-level of the respective shape (i.e., $|\phi_{S_j}|$ and $|\phi_{S_i}|$) along the contour of the other shape ($\partial\omega_i$ and $\partial\omega_j$, respectively).*

Since an affine transformation of an SDF is not an SDF, the transformation $T_{i\to j}$ is directly applied to $\mathbf{x} \in S_i$ such that the MHD between the i-shape and the j-shape is minimized (2.4). For the sake of notation brevity and simplicity we hereafter denote the SDF corresponding to $T_{i\to j}(S_i)$ by $\phi_{S_i}^{T^j}$. Note that shape alignment could be alternatively addressed by maximizing/minimizing other similarity/dissimilarity measures. For example, a pairwise comparison of shapes represented as binary images χ_i and χ_j (indicator functions) is commonly done by calculating their Dice score [17]:

$$D_{\text{DICE}} = \frac{2|\chi_i \cap \chi_j|}{|\chi_i| + |\chi_j|}.$$

Alternatively, referring to the SDFs of shapes i and j, the least square error (LSE) could be calculated as follows:

$$D_{\text{LSE}} = \int_\Omega |\phi_{S_i} - \phi_{S_j}|^2 d\mathbf{x}.$$

The advantage of aligning convoluted shapes (typical to anatomical structures) by minimizing the MHD, rather than by either *maximizing* the overlapping shape regions via the Dice score or *minimizing* their LSE measure, is demonstrated in Figure 2.[1] Figure 2(b) shows that when shape registration is performed such that the Dice measure is maximized (or, alternatively, the LSE is minimized), the tails of the tadpoles that cannot significantly overlap are "ignored." In contrast, the tails interleave at the expense of the alignment of the heads and the necks, when the optimization criterion is based on the proximity of the boundaries, as is the case with the MHD measure (see Figure 2(c)).

The values of the Dice scores, LSE, and the proposed MHD measures, which were obtained

[1]We deliberately chose 2D examples to visually demonstrate the main concepts introduced in section 2. The proposed method is evaluated on 3D data, as shown in section 3.

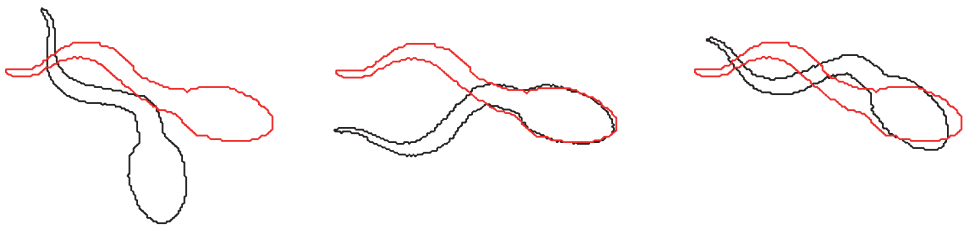

(a) Before alignment (b) Max. Dice/Min. LSE (c) Min. MHD

Figure 2. Shape alignment: An illustrative 2D toy example. (a) *The contours before alignment. Shape alignment is obtained by an affine registration of the black contour toward the red one either by* (b) *maximizing the shape overlap calculated by their Dice score/minimizing the LSE measure, or by* (c) *minimizing their MHD as suggested by our method. The corresponding values of the Dice scores, LSE, and the proposed MHD dissimilarity measures for each of the optimization criteria are given in Table 1 to reject the possibility that the differences are due to the fact that either of the registration processes was trapped in a local minimum.*

Table 1

A comparison of the Dice, LSE, and MHD scores obtained for each of the optimization criteria that were used to align the tadpole contours shown in Figure 2.

Score \ Opt. criterion	Min. LSE	Max. Dice	Min. MHD
LSE	827	833	1175
Dice	0.678	0.683	0.525
MHD	0.492	0.518	0.204

for shape registration using each of the optimization criteria, are given in Table 1,[2] to show that the different registration results were obtained for different global optima. Obviously, the dice score is relatively low, and the LSE measure is relatively high, while the MHD is the lowest when the optimization criterion for registration is minimizing the MHD.

The affine transformation parameters are recovered by a multigrid search in the parameter space [2] followed by a downhill simplex optimization [48], using a MATLAB optimization toolbox [44].

2.3. Level-set–based shape morphing. The SDFs $\{\phi_i\}$ can be viewed as level-set functions, where their zero-levels define the boundaries of the respective shapes. We use the sigmoidal "logistic" function of ϕ as a regularized form of the Heaviside function:[3]

$$H_\epsilon(\phi) = \frac{1}{2}\left(1 + \tanh\left(\frac{\phi}{2\epsilon}\right)\right) = \frac{1}{1 + e^{-\phi/\epsilon}}, \tag{2.6}$$

[2]In practice, we ran the registration algorithm three times for each of the three optimization criteria. However, the registration result obtained for maximizing the Dice score is not shown here as it is visually very similar to the result obtained by minimizing the LSE measure.

[3]Representation of the regularized form of a Heaviside function by the logistic function has been suggested before in [59] to assign a probabilistic interpretation to the values of $H_\epsilon(\phi(\mathbf{x}))$.

where ϵ defines the slope of the sigmoid and is arbitrarily set to one here. The proposed sigmoid has a similar behavior yet is not identical to the one used by [10]. Note that by simple manipulation of (2.6) we can express ϕ by $H_\epsilon(\phi)$:

$$\phi = \epsilon \log\left(\frac{H_\epsilon(\phi)}{1 - H_\epsilon(\phi)}\right), \tag{2.7}$$

where the values of $H_\epsilon(\phi(\mathbf{x}))$ are in the open set $(0, 1)$ for every $\phi(\mathbf{x})$.

The boundary of a shape S_i can be approximated by $\partial\omega_{S_i} = |\nabla H_\epsilon(\phi_{S_i})|$, defining the continuous MHD between S_i and S_j as follows:

$$E_{\text{CMHD}}(S_i, S_j) = \int_\Omega \left[|\phi_{S_j}||\nabla H_\epsilon(\phi_{S_i})| + |\phi_{S_i}||\nabla H_\epsilon(\phi_{S_j})|\right] d\mathbf{x}. \tag{2.8}$$

Morphing of a shape S_j toward S_i can be obtained if the respective level-set function ϕ_{S_j} evolves to satisfy the following minimization scheme:

$$\phi_{S_j}^{\rightarrow i} = \arg\min_{\phi_{S_j}} E_{\text{CMHD}}(S_i, S_j). \tag{2.9}$$

In practice, we add an *area preserving term* (*volume preserving term* in three dimensions) and a *length preserving term* (*surface area preserving term* in three dimensions) which regularize the expression in (2.8). The length (surface area) and the area (volume) preserving terms limit the extent of possible changes of these significant shape properties (due to, for example, excessive shrinkage or expansion) throughout the morphing process. Let A_j^0 be the initial area (volume) of S_j obtained by pixel (or voxel) count. The area embedded in ϕ_{S_j} throughout its evolution process is estimated as $A_j^t = \int_\Omega H_\epsilon(\phi_{S_j}(\mathbf{x}; t))d\mathbf{x}$.

We therefore define the *area preserving term* as

$$E_A(\phi_{S_j}; A_j^0) = \left[\int_\Omega H_\epsilon(\phi_{S_j}(\mathbf{x}))d\mathbf{x} - A_j^0\right]^2. \tag{2.10}$$

In analogy to the area/volume preserving term we derived a *length/surface area preserving term*, which takes the following form:

$$E_L(\phi_{S_j}; L_j^0) = \left[\int_\Omega |\nabla H_\epsilon(\phi_{S_j}(\mathbf{x}))|d\mathbf{x} - L_j^0\right]^2, \tag{2.11}$$

where L_j^0 is the length (surface area) of S_j calculated by integrating $|\nabla H_\epsilon(\phi_{S_j}(\mathbf{x}))|$ at $t = 0$.

Combining (2.8), (2.10), and (2.11), the unified level-set functional which defines the Hausdorff-based shape morphing takes the form

$$E(\phi_{S_j}, \phi_{S_i}) = W_{\text{MHD}} E_{\text{CMHD}}(\phi_{S_i}, \phi_{S_j}) + W_L E_L(\phi_{S_j}; L_j^0) + W_A E_A(\phi_{S_j}; A_j^0). \tag{2.12}$$

The gradient descent equation that determines the evolution of ϕ_{S_j} is derived from the first

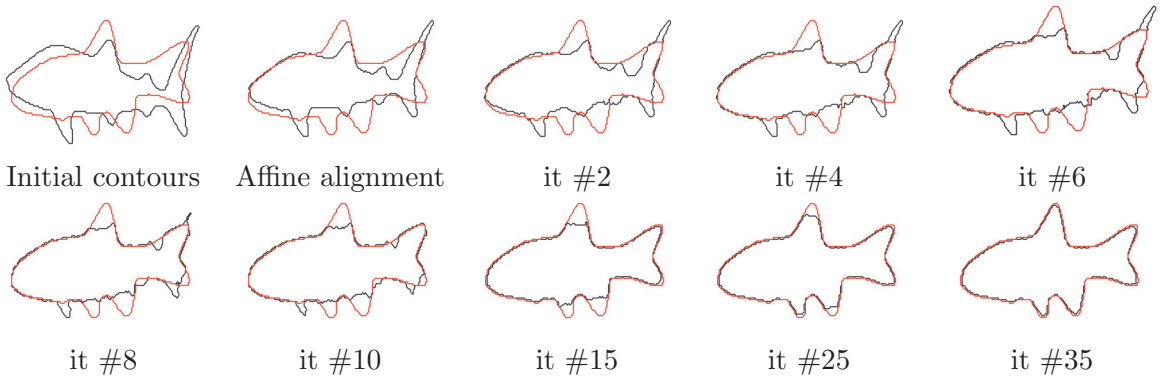

Figure 3. Shape morphing. *The black contour is morphed toward the red contour by minimizing the MHD level-set functional in (2.12) via the gradient descent process introduced in (2.16). Prior to the morphing process, the contours are aligned by an affine transformation, as described in section 2.2.*

variation of the functional in (2.12):

$$\begin{aligned}\frac{\partial \phi_{S_j}}{\partial t} &= W_{\mathrm{MHD}}\left[\operatorname{sign}(\phi_{S_j})|\nabla H_\epsilon(\phi_{S_i})| + \delta_\epsilon(\phi_{S_j})\operatorname{div}\left(\frac{\nabla\phi_{S_j}}{|\nabla\phi_{S_j}|}|\phi_{S_i}|\right)\right] \\ &\quad + W_A\delta_\epsilon(\phi_{S_j})\left[\int_\Omega H_\epsilon(\phi_{S_j}(\mathbf{x}))d\mathbf{x} - A_j^0\right] \\ &\quad + W_L\delta_\epsilon(\phi_{S_j})\operatorname{div}\left(\frac{\nabla\phi_{S_j}}{|\nabla\phi_{S_j}|}\right)\left[\int_\Omega |\nabla H_\epsilon(\phi_{S_j}(\mathbf{x}))|d\mathbf{x} - L_j^0\right],\end{aligned} \tag{2.13}$$

where sign$(\cdot)$ is the sign function, div is the divergence operator, and $\delta_\epsilon(\phi)$ is the derivative of $H_\epsilon(\phi)$ with respect to ϕ. We note that the derivative of the absolute value function $|\phi_{S_j}|$ (left-hand side term of (2.8)) is not defined at zero. A more stable alternative to the sign function used here would be based on the left- and right-side partial derivatives as suggested by [38].

The volume preserving term generalizes the commonly used balloon term [12], where the scalar $A^{\mathrm{diff}} \triangleq A_j - A_j^0$ determines both the size and the direction (expansion or contraction) of the evolution of ϕ_{S_j} and its respective boundary (zero-level). In a similar manner, the surface area preserving term generalizes the commonly used smoothness (or regularization term) [32]. The difference $L^{\mathrm{diff}} \triangleq L_j - L_j^0$ affects both the size and the sign of the contour's evolution. Note that while the smoothness term always acts to shorten the length of the active contour (reducing the curvature), here when the difference L^{diff} is negative, the boundary evolves in the opposite direction, increasing the curve length (surface area).

We use a toy example (fish contours shown in Figure 3) to demonstrate the shape morphing process obtained by applying iteratively the above gradient descent equation (2.13). The target shape S_i is represented by the red contour. The black contour, which illustrates S_j in the equation, dynamically evolves until it is aligned with S_i. Prior to the morphing process, the black contour is registered to the red one by an affine transformation such that the MHD between the shapes is minimized (section 2.2).

2.4. Evolution of the mean shape. "Shape ensemble" is obtained by 3D manual delineations of the anatomical structure under study (e.g., left hemisphere amygdala) from MR images of different subjects of the same population (e.g., schizophrenia patients).

We define the mean S^M of a shape ensemble $\{S_1 \dots S_N\}$ as the shape that minimizes the sum of the distances from all the shapes in the set:

$$\hat{S}^M = \arg\min_{S^M} \sum_{i=1}^{N} D_{\mathrm{CMHD}}(T_{i\to M}(S_i), S^M), \tag{2.14}$$

where $T_{i\to M}$ is the affine transformation (obtained throughout the optimization process) that aligns a shape S_i to the population center as defined in section 2.2.

The MHD level-set functional defined in (2.12) is now extended to take the following form:

$$\begin{aligned} E(S^M, \{S_i\}) = \frac{W_{\mathrm{MHD}}}{N} \sum_{i=1}^{N} \int_\Omega \left[|\phi_{S^M}||\nabla H_\epsilon(\phi_{S_i}^{T^M})| + |\phi_{S_i}^{T^M}||\nabla H_\epsilon(\phi_{S^M})| \right] d\mathbf{x} \\ + W_A \left[\int_\Omega H_\epsilon(\phi_{S^M}(\mathbf{x}))d\mathbf{x} - A^0 \right]^2 \\ + W_L \left[\int_\Omega |\nabla H_\epsilon(\phi_{S^M}(\mathbf{x}))|d\mathbf{x} - L^0 \right]^2. \end{aligned} \tag{2.15}$$

Here, A^0 and L^0 are the averages of the ensemble's areas (volumes) and lengths (surface areas), respectively.[4]

The incorporation of the length (surface area) and the area (volume) preserving terms in (2.15) implies that S^M can no longer be considered as the Fréchet mean (or Karcher mean) of the ensemble. We, however, keep this terminology though, in fact, we should have used the term *canonical representation of the shape ensemble* instead.

As neither S^M nor the affine parameters $T_{i\to M}$ that align each shape S_i to S^M are known, we use an alternating minimization technique in which (2.5) (for each shape S_i in the ensemble) and (2.14) are jointly solved. While the affine transformation parameters are inferred by using a global optimization method, the mean shape is generated via gradient descent optimization of a level-set functional:

$$\begin{aligned} \frac{\partial \phi_{S^M}}{\partial t} = \frac{W_{\mathrm{MHD}}}{N} \sum_{i=1}^{N} \left[\mathrm{sign}(\phi_{S^M})|\nabla H_\epsilon(\phi_{S_i}^{T^M})| + \delta_\epsilon(\phi_{S^M})\mathrm{div}\left(\frac{\nabla \phi_{S^M}}{|\nabla \phi_{S^M}|} |\phi_{S_i}^{T^M}| \right) \right] \\ + W_A \delta_\epsilon(\phi_{S^M}) \left[\int_\Omega H_\epsilon(\phi_{S^M}(\mathbf{x}))d\mathbf{x} - A^0 \right] \\ + W_L \delta_\epsilon(\phi_{S^M})\mathrm{div}\left(\frac{\nabla \phi_{S^M}}{|\nabla \phi_{S^M}|} \right) \left[\int_\Omega |\nabla H_\epsilon(\phi_{S^M}(\mathbf{x}))|d\mathbf{x} - L^0 \right], \end{aligned} \tag{2.16}$$

where $A^{\mathrm{diff}} \triangleq \int_\Omega H_\epsilon(\phi_{S^M}(\mathbf{x}))d\mathbf{x} - A^0$ and $L^{\mathrm{diff}} \triangleq \int_\Omega |\nabla H_\epsilon(\phi_{S^M}(\mathbf{x}))|d\mathbf{x} - L^0$ are scalars, calculated before each gradient descent iteration.

[4] We assume that the standard deviations of the shape ensemble's areas (volumes) and lengths (surface areas) are relatively small.

2.5. Mean shape construction in practice.

2.5.1. SDF construction. The proposed shape morphing algorithm has to be performed in the space of SDFs to ensure that the SDFs of the shape contours are generated in the initial stage of the algorithm and after any arithmetic operation or coordinate transformation (but translation), including gradient descent updates. SDF construction is implemented via the fast marching algorithm [35, 27].

2.5.2. Shape alignment. Shape analysis algorithms require shape alignment. However, the question of what is the "appropriate" transformation space for medical imaging applications is arguable. One can remove pose, pose and scale (uniform or nonuniform) (e.g., SPHARM [52]), or full affine (TBM registration–based techniques [1]). Sometimes it goes as far as computing shape invariants to isometries (e.g., using a Laplace–Beltrami operator [50]). Although here we look only at deformations after affine registration, we do compute affine parameters and can easily go back and extract the shearing parameters from the affine matrix if desired.

2.5.3. Initialization. Since the arithmetic mean of the set of SDFs $\phi_1, \ldots, \phi_N$ is not an SDF, we use the inverse-logit of the SDFs as in [55] and average the regularized Heaviside functions of the SDFs, i.e., $H_\epsilon(\phi_{S_i})$. The initial estimation of the level-set function, representing the mean shape, i.e., $\phi_{S^M}^{\text{INIT}}$, is determined by the contour $\partial\omega_M$ defined by $\{\mathbf{x} \in \Omega \mid \frac{1}{N}\sum_{i=1}^{N} H_\epsilon(\phi_{S_i}(\mathbf{x})) = \frac{1}{2}\}$. This preliminary step to the proposed framework for mean shape generation by shape morphing is in the spirit of current approaches to groupwise image registration, in which a set of images is simultaneously aligned to the latent population center. Rather than selecting any individual image as the template, which may introduce a bias, the process is carried out in an iterative manner, starting from a rough estimate of the unknown mean [73, 74] or median [22]. For a more detailed review and alternative approaches, see also [75] and references therein.

In practice, the ensemble shapes are first aligned by translation such that the center of mass of each shape coincides with the mean of the centers of mass of all the shapes. We then average the shape, defining an "approximated" mean, and calculate the affine transformations of each shape to that mean. Then a better approximation to the mean shape is calculated by averaging the affine-transformed shapes. We use the approximated mean shape to initialize the level-set–based gradient descent process.

2.5.4. Stopping criteria and parameter setting. The algorithm stops when the integration of the absolute differences between the Heaviside functions of the evolving level-set function in two consecutive iterations is negligible, i.e., smaller than a predefined threshold η:

$$\int_\Omega |H_\epsilon(\phi^M(t+1)) - H_\epsilon(\phi^M(t))| < \eta, \tag{2.17}$$

where t denotes the iteration number. In our experience, upon convergence, the left-hand side term in the above inequality (2.17) decreases to zero. We, however, set η to a positive and small scalar, taking into account cases in which there are small fluctuations around the minimum.

We tune the weights W_{MHD}, W_L, and W_A of the energy functionals (2.12) and (2.15) to balance the respective gradient descent terms (2.13) and (2.16) while preserving the proportion to the changes in length/surface area and area/volume. Specifically, let $W_L = w_L/L^0$ and $W_A = w_A/A^0$. The parameters w_L, w_A, and W_{MHD} are scaled such that the MHD term, the length term divided by L^{diff}, and the area term divided by A^{diff} have the same order of magnitude, and the absolute value of their sum is bounded by one.

We set both the step size of the level-set evolution (dt) and the slope of the regularized Heaviside (ϵ) to one. Although adaptive modification of dt using a line search algorithm [47] may facilitate convergence, the results obtained are sufficiently good, as the weighted sum of gradient descent terms is normalized. The parameter ϵ determines the number of levels around the zero-level (the shape boundary) that are updated at each gradient descent iteration. Its size is usually related to the scale of the shape being analyzed.

2.5.5. Toy example. The initial approximation of the mean shape ϕ_{SM}^{INIT} is constructed by the shape "averaging" process, which smooths out structural details that may characterize a given shape set. The evolution process described in the previous section, in which ϕ_{SM}^{INIT} is iteratively morphed toward each of the shapes $\{S_i\}$, is therefore necessary for generating a shape that faithfully represents the ensemble.

Figure 4 illustrates the process of mean shape construction. In this simplified example the shape ensemble consists of six 2D slices of STG structures, shown in Figures 4(a)–4(f). Figure 4(g) shows an initial approximation of the mean shape after alignment of the input shapes by an affine transformation. It is worth noting that while in three dimensions the initial approximation of the mean shape does not undergo significant topological changes (e.g., holes or disassembly) as in two dimensions, its appearance may not be a meaningful representation of the ensemble shapes. The final mean shape estimate obtained by the gradient descent process in (2.16) is shown in Figure 4(h). Figure 5 shows the pairwise MHD distances (color-coded) of the planar STG shapes and their mean that are shown in Figure 4. Note that the MHD between each of the input shapes and the constructed (mean) shape is much smaller than the pairwise MHD between that shape and any of the other input shapes.

(a) (b) (c) (d) (e) (f) (g) (h)

Figure 4. Mean shape construction. (a)–(f) *Input shapes.* (g) *Initial approximation of the mean shape.* (h) *Final estimation of the mean of shapes* (a)–(f) *obtained by the suggested method.*

2.6. Localization of shape differences between populations. We now present how our model can effectively detect local shape deformations within a population. Given a point on the mean shape boundary, $\mathbf{x} \in \partial\omega_{\text{Mean}}$, we can directly obtain its signed distance to each of the affine aligned shapes by looking up the distance in the corresponding SDF $\phi_n(\mathbf{x})$. For each of those voxels, statistics on the $\phi_n(\mathbf{x})$ can capture shape deformation.

Let $\phi_1^{\text{NC}}, \ldots, \phi_N^{\text{NC}}$ and $\phi_1^{\text{AB}}, \ldots, \phi_M^{\text{AB}}$ be the SDFs representing shapes of a particular anatomical structure in two populations. Let $d^{\text{NC}}(\mathbf{x})$ and $d^{\text{AB}}(\mathbf{x})$ be two vectors of lengths

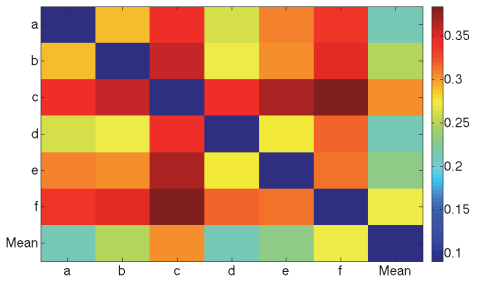

Figure 5. *Pairwise MHD distances (color-coded) of the planar STG shapes and their mean shown in Figure 4. Note that the MHD between each of the input shapes and the constructed (mean) shape is much smaller than the pairwise MHD between that shape and any of the other input shapes.*

N and M, respectively, of the (signed) distances of the corresponding shape ensembles to $\mathbf{x}$. We then apply two-sample t-tests with significance level of 0.05 to each of the mean shape surface voxels. The assumption that $d^{\mathrm{NC}}(\mathbf{x})$ and $d^{\mathrm{AB}}(\mathbf{x})$ come from a normal distribution was checked using the Lilliefors test (MATLAB command *lillietest* [45]). We assume different variances. Having p-values below the significance level allows us to locate SSDs between the two populations. In the following experiments, we also use the false discovery rate (FDR) to correct for multiple comparisons [49].

3. Experiments.

We evaluated the proposed algorithm using synthetic and real data sets: synthetic sets of the amygdala-hippocampal complex (AHC) (section 3.1) and of the striatum (section 3.2), and real data of the corpus collasum (CC) of Williams syndrome patients (section 3.3), superior temporal gyrus (STG) in first-episode schizophrenia patients (section 3.4), and the caudate nucleus in women with schizotypal personality disorder (SPD) (section 3.5). Results were corrected for multiple comparisons using the FDR approach in [49]. We applied an FDR of 5%.

3.1. Synthetic AHC.

Manual segmentations of the left AHC in 40 normal controls were taken from the laboratory database. An unbiased atlas of the AHC was created from the 40 samples [60]. The resulting atlas was then warped back to the subject space using 20 randomly selected inverse warps obtained in the previous step. The resulting samples compose the normal control (NC) group. The remaining 20 were manipulated by adding (or removing) a specific number of voxels using a hemisphere such that either a bump (or dimple) would be created and labeled "abnormal" (AB). Eight pairs of NC/AB data sets were generated. Each AB set had a bump (or a dimple) located in the head of the AHC and with a radius of $3, 4, 5$, and 6 voxels, respectively.

For each AB/NC data set, we generated the mean AHC and performed a t-test comparing the NC and AB distances to the mean at each point on the mean shape's surface. Successful results of these eight experiments ($\{\text{bump, dimple}\} \times \{3, 4, 5, 6\}$) are shown in Figure 6. We also evaluated the method by looking at the ratio of the statistically significant voxels over the total number of surface voxels for the bump (dimple) as the size of the deformation increases (Figure 7).

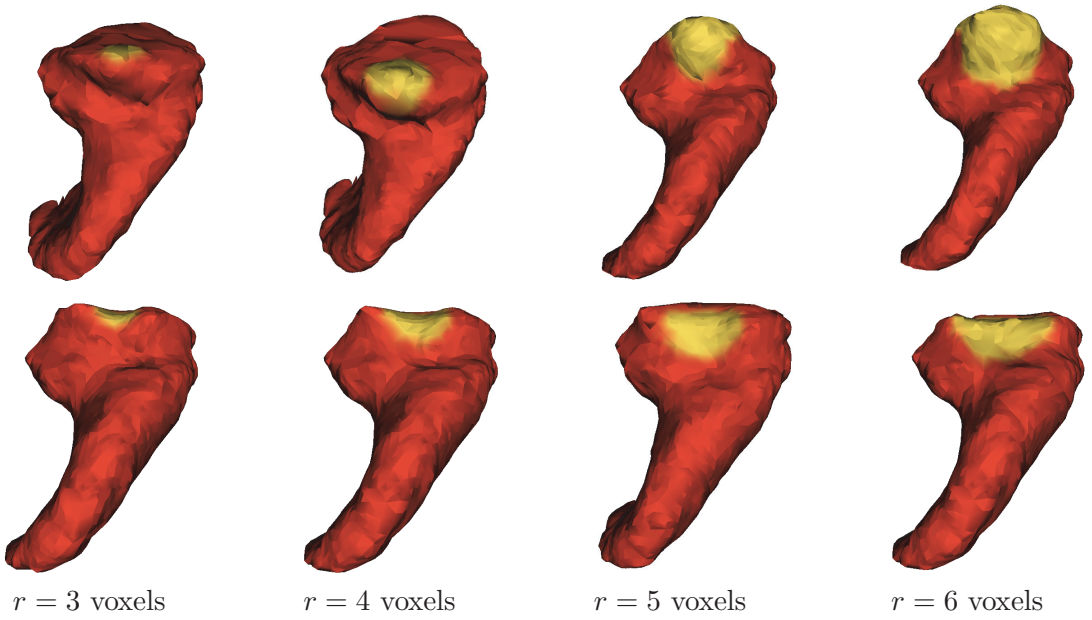

Figure 6. AHC data set. *p-value maps displayed on the mean shapes of NC/AB data sets with a bump (top row) or dimple (bottom row) of radius (from left to right)* 3, 4, 5, *and* 6. *Red indicates nonsignificant p-values, while the yellow colors present a scale of FDR corrected p-values (below the threshold).*

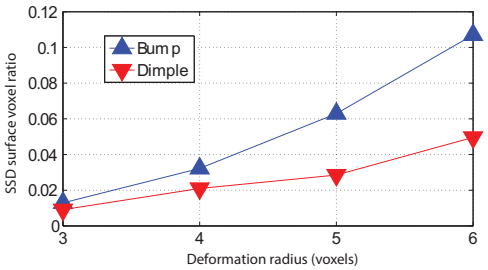

Figure 7. Method evaluation for the synthetic AHC data sets. *The Y-axis displays the ratio between the number of surface voxels which have p-values that indicate SSD between the populations and the total number of surface voxels. This ratio increases as the size of the synthetic distortion (bump or dimple) increases.*

3.2. Synthetic striatum. Synthetic striatum shapes were generated through manifold learning based on a training set of 27 real samples [23, 24]. Abnormal examples were generated via random processes of either thinning or thickening of specific, well-defined regions of the striatum (see Figures 8(a)–8(c), 8(e)–8(g)). Two sets of examples for the right and the left striatum, each containing 50 normal and 50 abnormal examples, were tested. Results are shown in Figure 8. Note that the distorted regions (highlighted in red), corresponding to voxels with significant (FDR corrected) p-values (Figures 8(d) and 8(h)), were precisely detected.

Next we compared three additional striatum data sets each containing 100 samples that were synthetically distorted in three distinct locations. Each of these sets was compared with a fourth set containing 100 nondeformed striatum examples. Figure 9 shows the mean shapes of the deformed and the nondeformed population for each pair of sets. The extent of detected

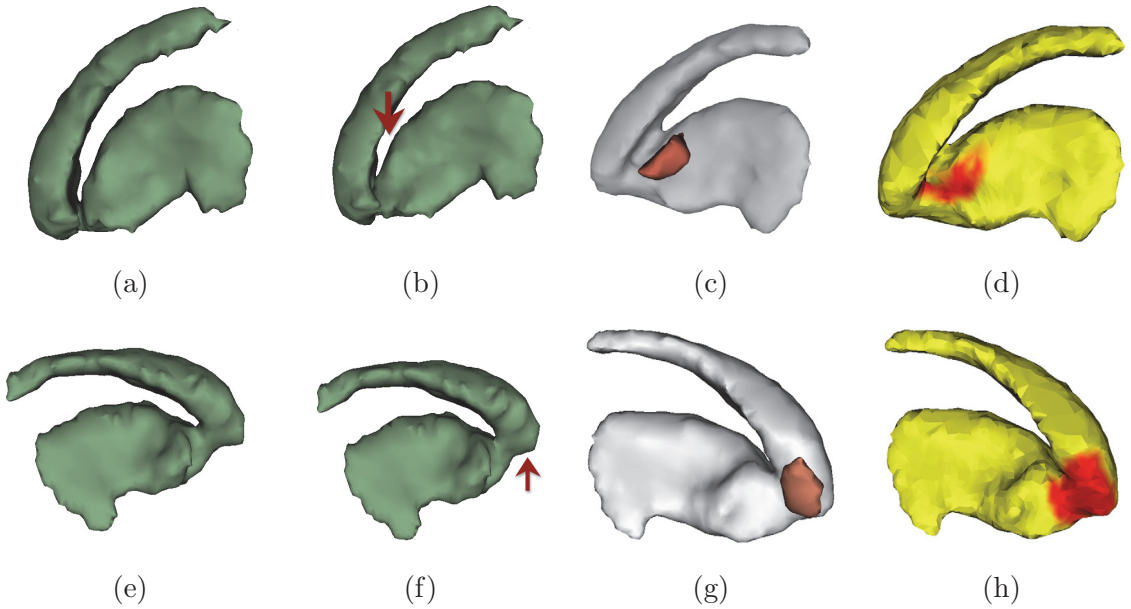

Figure 8. Left (top row) and right (bottom row) striatum data sets. (a), (e): *Randomly selected examples of the left* (a) *and the right* (e) *striatum.* (b), (f): *Left* (b) *and right* (f) *striatum shapes after applying shape deformations to the respective shapes shown in* (a) and (e)*, respectively.* (c), (g): *Mean (left and right) striatum shapes along with the averaged artificial deformation (red).* (d), (h): *Mean (left and right) striatum shapes along with the respective p-value maps comparing distorted and undistorted data sets. Yellow indicates nonsignificant p-values. Red colors present a scale of FDR corrected p-values (below the threshold). Note that although the deformations and the extent of the deformed regions (indicated by arrows) are subtle, they were successfully detected by our algorithm.*

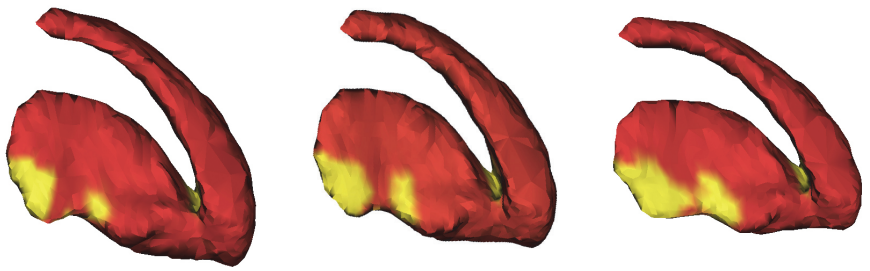

Figure 9. Striatum data sets. *p-value maps displayed on the mean shapes of NC/AB data sets with increasing (from* (a) *to* (c)*) synthetic deformations in three locations upon the striatum surface. Red indicates nonsignificant p-values, while the yellow colors present a scale of FDR corrected p-values (below the threshold).*

distorted regions (highlighted in red), corresponding to voxels with significant (FDR corrected) p-values, increases as the level of deformation increases. Quantitative results are shown in Figure 10. Here, the X-axis displays the extent of the synthetic distortion (deformation level). The left-hand Y-axis displays the ratio between the number of surface voxels which have p-values that indicate SSD between the populations and the total number of surface voxels.

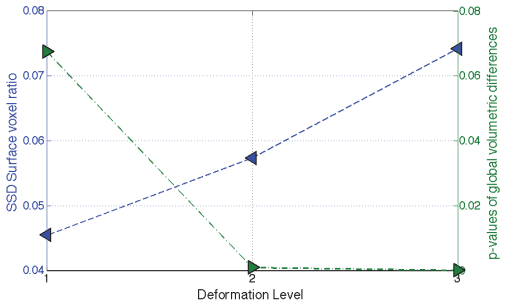

Figure 10. Method evaluation for the synthetic striatum data sets. *The left-hand Y-axis displays the ratio between the number of surface voxels which have p-values that indicate SSD between the populations and the total number of surface voxels. This ratio increases as the extent of the synthetic distortion (deformation level) increases (blue). The right-hand Y-axis displays the p-values obtained by applying the two-sample t-test to the volumes of the deformed and nondeformed shape sets. Note that the calculated p-value for the set with the smallest deformation level is above* 0.05. *The calculated p-value decreases as the deformation level increases (green).*

This ratio increases as the extent of the synthetic distortion (deformation level) increases. The right-hand Y-axis displays the p-values obtained by applying the two-sample t-test to the volumes of the deformed and nondeformed shape sets. As the deformation level increases (becomes more significant) the corresponding p-values for the volume measurements decrease. Note that the calculated p-value for the set with the smallest deformation level is above 0.05. This result demonstrates the capability of the proposed method to detect and locate small deformations that cannot be captured by comparing global shape properties such as shape volume.

3.3. 2D CC data of Williams syndrome patients. Williams syndrome (WS) patients display a distinctive cognitive and behavioral profile which has been associated with several neuroanatomical changes, including volumetric and morphological changes in both gray and white matter (e.g., [71]). In our study, a T1 coronal acquisition (256×192 matrix, 1.5mm thickness) was acquired for 16 WS patients and 17 controls from the Genetic Medical Institute (Portugal) and the Genomic Foundation in Galicia (Spain). A 2D cross-section of the CC was manually segmented in the mid-sagittal slice of each image. In Figure 12 the mean shape of both the patients and the NC subjects has been constructed using our level-set–based MHD shape morphing method. Then, the MHD between each shape in the two sets and the mean shape were calculated. We then performed a population analysis based on these measures using a two-sample t-test. We found an SSD between WS and NC based on a global MHD measure (2.4): $p = 0.03$, $t = 2.22$. For a comparison, the results for a two-sample t-test on the perimeters of the aligned CC shapes were $p = 0.041$, $t = 3.0983$. There was no SSD between the CC areas of the WS and NC populations (we did not normalize based on the intracranial volumes, as this information was not available to us). We next calculated the closest (in terms of magnitude) signed distances between each of the 220 pixels along the mean CC contour and each of the 33 CCs of both the NC and WS populations. The results are shown in Figure 11. In this figure, the signed distances are color-coded. The X-axis displays the pixel's index along the mean CC contour. The Y-axis displays the subject's index, where NC subjects'

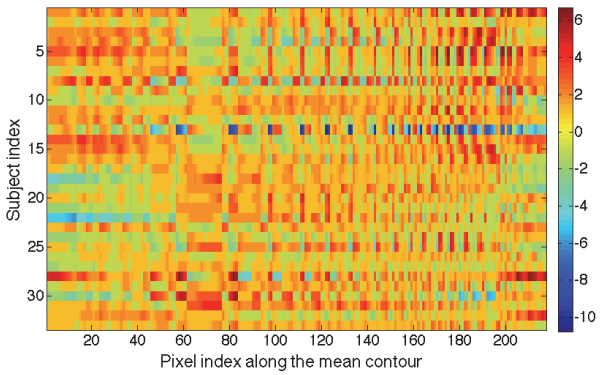

Figure 11. *Color-coded signed distances of each pixel along the mean shape contour to the closest pixel along the CC contour of each subject. X-axis: Pixel's index along the mean CC contour. Y-axis: Subject's index, where NC subjects' indices are* 1–17 *and WS subjects' indices are* 18–33.

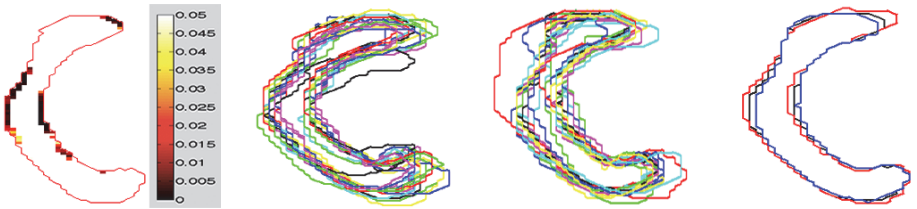

Figure 12. CC data of WS patients. (a) *p-values (color-coded) along the mean CC contour spatially calculated for the NC and WS populations.* (b)–(c) *Superposition of the CC contours for the NC* (b) *and WS* (c) *populations. Contours' colors are random.* (d) *Superposition of the mean contours of the NC (red), the WS (blue), and the entire set of* 33 *subjects (black).*

indices are 1–17 and WS subjects' indices are 18–33. We then calculated the p-values for each pixel along the mean CC contour for the NC and WS populations. The results are shown in Figure 12(a), in which color-coded, statistically significant p-values (below 0.05) are plotted upon the mean CC contour (red). This figure allows us to detect the (planar) CC segments that distinguish (statistically) between WS patients and NC subjects. A visual assessment of this result is shown in Figures 12(b)–12(d). Figures 12(b)–12(c) show superposition of the CC contours for the NC and WS populations, respectively. The mean contours of the NC (red), the WS (blue), and the entire set of 33 subjects (black) are superimposed in Figure 12(d).

3.4. STG in first-episode schizophrenic patients. We used manual segmentation of the left and right STGs in 19 patients diagnosed with first-episode schizophrenia and 14 matched NCs originally acquired for a brain volumetric study [28]. Examples of the left STG of some of the subjects are shown in Figure 13.

We generated the mean shapes of the patient and NC data sets for the left and right STGs. We computed a t-test comparing the two populations at each point on the mean shape's surface. We then applied a threshold at an FDR of 0.05 to the resulting p-values. The color-coded p-value maps displayed upon the mean STG surface of the left hemisphere (Figure 14(b))

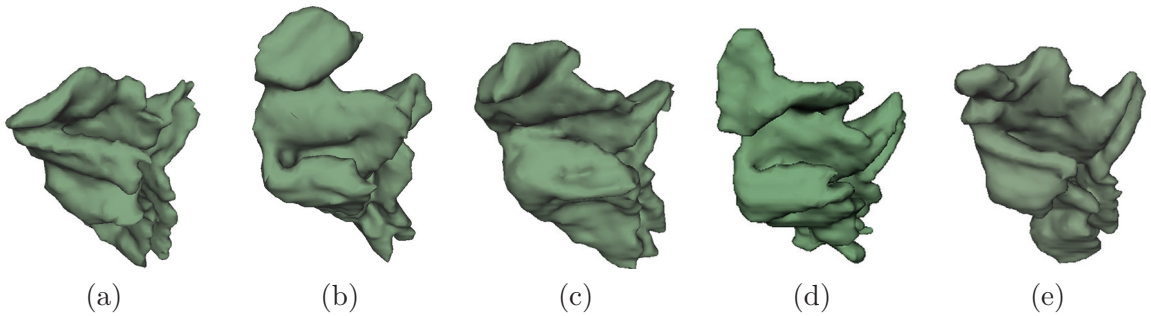

Figure 13. Left STG of first-episode schizophrenics. *A few examples demonstrating the complexity of this structure and its variability among patients.*

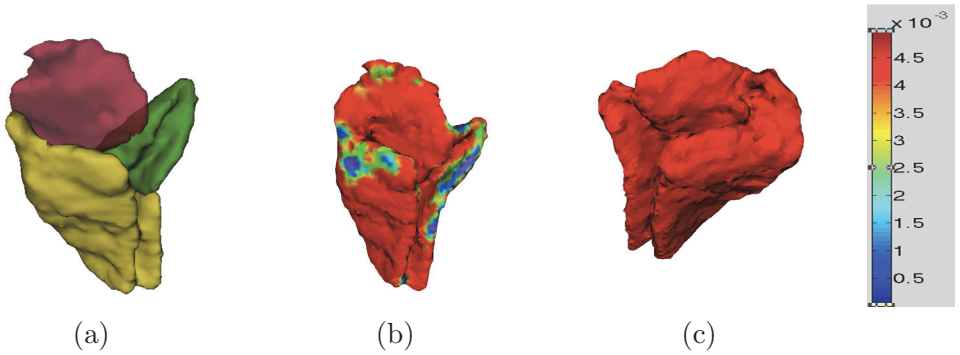

Figure 14. STG in first-episode schizophrenia. (a) *Left STG composed of its substructures: planum temporale (red), Heschl's gyrus (green), rest of the STG (yellow);* (b) *p-value map of the left STG comparing schizophrenia patients and NCs;* (c) *p-value map of the right STG for the same populations. Red indicates nonsignificant p-values; green-to-purple colors present a scale of p-vales below the* 0.05 *FDR threshold. Note that the right STG shows no differences.*

localize the STG regions which demonstrate SSDs between the populations (green-to-purple spots). No SSDs were detected for the right STG (homogeneous red color). To evaluate these findings we partitioned the STG into three regions: the planum temporale (Figure 14(a), in red), the Heschl gyrus (Figure 14(a), in green), and the rest of the STG (Figure 14(a), in yellow). We then performed volume and surface area measurements on each of these subregions and the entire STGs (left and right). This is, in fact, the common practice in psychiatry and neuroanatomy research. We received SSDs for the left Heschl gyrus measurements: p-values of 0.005 for the volume and 0.035 for the surface area measurements. Borderline p-values were obtained for the left planum temporale. Specifically, the p-values for the volume and surface area measurements were 0.053 and 0.058, respectively. Measurements on the right STG substructures and on the entire STGs did not show any statistical significance. Note that, except from the left part of the left STG, the results obtained by the proposed method are in line with the results presented above for the volume and surface area measurements as well as with published findings on gray matter reduction in these substructures in schizophrenia patients [62]. We suspect that the focal differences in the left part of the left STG that were captured by our p-value map were too small (with respect to the size of this substructure) to be detected otherwise.

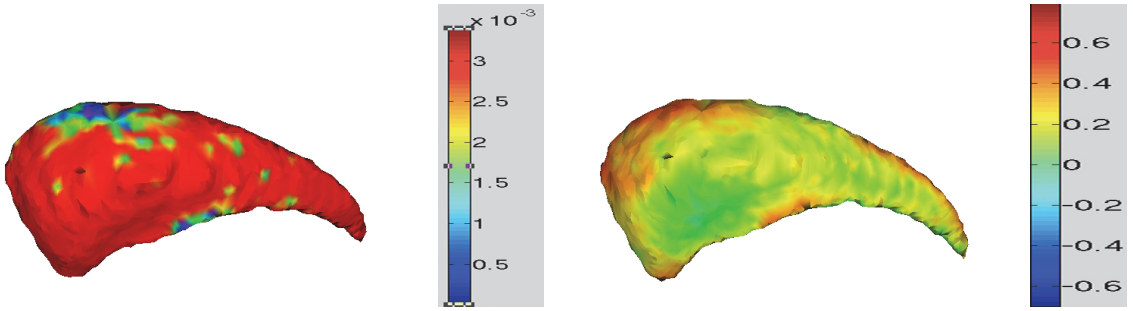

(a) (b)

Figure 15. Caudate nucleus in SPD patients. (a) *Mean right caudate along with the respective p-value maps. Red indicates nonsignificant p-values. Yellow-to-blue indicates p-value below a* 0.05 *FDR.* (b) *Signed differences between the NC subjects and the SPD patients. Positive values are in yellow-to-red, negative values in blue.*

3.5. Caudate nucleus in SPD patients. MR brain scans of 61 women—32 SPD patients and 29 NC subjects— were manually segmented to extract the caudate nucleus in the left and right hemispheres [34]. Statistically significant morphological differences have been detected in the right caudate using the proposed algorithm (Figure 15). Figure 15(b) presents the sign of the differences (inflation or deflation) between the NC subjects and the SPD patients. A comparison between Figure 15(a) and Figure 15(b) shows that the SSDs between the NC and SPD populations are due to deflation, i.e., gray matter reduction in the patients' caudate. These results are comparable to the manual volumetric measures as well as to the shape statistics obtained from using SPHARM-PDM methodology [39].

4. Discussion and conclusions. We presented a novel framework for performing shape analysis for population studies. The core of the method is a level-set algorithm for shape morphing that is based on the modified Hausdorff distance (MHD). The proposed algorithm was applied to a variety of neuroanatomical structures, including synthetic sets of the striatum and the amygdala-hippocampal complex (AHC) and real data of the corpus collasum (CC) in Williams syndrome (WS) patients, the superior temporal gyrus (STG) in first-episode schizophrenia patients, and the caudate nucleus in women with schizotypal personality disorder (SPD). We were able to locate statistically significant shape differences between NC and patient populations which are consistent with prior findings based on volumetric measurements of manually segmented structures. Moreover, our experiments on both synthetic and real clinical data show that the method is sensitive enough to detect deformations that are too small to be captured by volume measurements.

In contrast to most existing methods, here shape alignment by affine transformations and nonrigid shape morphing is done by minimizing the MHD. We believe that using a boundary-based dissimilarity measure rather than a metric based on volume overlap might be more suitable when dealing with twisted or convoluted shapes, such as cortical brain regions. This assumption is supported by the examples shown in Figure 2. The MHD term dominates the evolution of the proposed level-set functional. As we use a gradient descent strategy to solve a nonconvex problem, the optimization process may be trapped in (undesired) local

minima, which may result, for example, in topological changes or an excessive shrinkage of the evolving mean shape. We therefore incorporate two additional soft constraints: the length (surface area) and the area (volume) preserving term. These terms generalize the traditional balloon and smoothness terms. Their signs determine the directions of the normals to the evolving contour (or surface). Their magnitudes are proportional to the unsigned differences in volume or surface area between the ensemble shapes and the constructed mean shape in the current iteration. Therefore, when the volume or surface area of an evolving shape does not vary significantly during a morphing process, the contribution of these terms to the entire cost functional is negligible.

While the proposed method is general and may be applied to a variety of shapes and shape classes, it is in particular usable in neurology and psychiatry research, in which the connection between focal brain atrophies (i.e., loss of neurons and synapses in specific brain regions) and neurological disorders is investigated. A common practice in psychiatric studies is to detect statistically significant volumetric differences of suspected brain structures between NCs and patients. Those volume measurements are usually based on segmentation of brain structures and substructures. For example, the STG contains the primary auditory cortex, which is responsible for processing sounds. To test the hypothesis that auditory hallucinations (a common positive symptom of schizophrenia) are related to gray matter volume loss in the STG, statistics on volumetric measurements of patients and NCs were conducted. Nevertheless, when an entire structure is considered (such as the STG), the comparison (based on statistical tools) may not be sensitive enough to detect small, localized brain atrophies, as was demonstrated by our measurements (section 3.4) and others [62]. This premise is also supported by our synthetic striatum experiment (section 3.2 and Figure 10). Moreover, even when the deformations are detected, they cannot be located. It is therefore common to calculate volume differences on the manually segmented substructures: the planum temporale, the Heschl gyrus, and the rest of the STG [62]. For first-episode schizophrenia patients statistically significant volume differences were detected in the left hemisphere planum temporale and Heschl gyrus [31, 62]. Similar findings were obtained in the STG experiment performed by the proposed statistical shape analysis technique (section 3.4 and Figure 14). We note that the partitioning of brain structures, such as the STG, is difficult due to the complexity and heterogeneity of sulcal patterns in cortical regions. To address this challenge, heuristics and thumb rules are often used. This may lead to biased and error-prone segmentations. The proposed fully automatic shape analysis method is therefore a more accurate, repeatable, and objective alternative.

The strength of the proposed method is its capability to detect local morphological differences between populations. Nevertheless, scattered and spatially inconsistent shape dissimilarities cannot be captured via voxelwise comparison as we do here. The alternative is therefore either using global MHD measures of the entire structure or, for the detection of more focal differences, calculating the MHD measures of surface patches of growing sizes (in a scale-space manner) similar to the spatial smoothing in VBM. This possible extension is a subject for a future study.

Acknowledgments. We thank Adriana Sampaio for the corpus collasum datasets and Mert Rory Sabuncu and Boon Thye (Thomas) Yeo for the fruitful discussions which triggered this work.

REFERENCES

[1] J. Ashburner and K.J. Friston, *Voxel-based morphometry—The methods*, NeuroImage, 11 (2000), pp. 805–821.

[2] J.R. Bergen, P. Anandan, K.J. Hanna, and R. Hingorani, *Hierarchical model-based motion estimation*, in Proceedings of the Second European Conference on Computer Vision (ECCV '92), Springer-Verlag, London, 1992, pp. 237–252.

[3] B. Berkels, G. Linkmann, and M. Rumpf, *An SL(2) invariant shape median*, J. Math. Imaging Vis., 37 (2010), pp. 85–97.

[4] F.L. Bookstein, *Shape and the information in medical images*, Comput. Vision Image Underst., 66 (1997), pp. 79–118.

[5] S. Bouix, J.C. Pruessner, D.L. Collins, and K. Siddiqi, *Hippocampal shape analysis using medial surfaces*, NeuroImage, 25 (2005), pp. 1077–1089.

[6] Ch. Brechbuhler, G. Gerig, and O. Kübler, *Parametrization of closed surfaces for 3-D shape description*, Comput. Vision Image Underst., 61 (1995), pp. 154–170.

[7] A.M. Bronstein, M.M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro, *A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching*, Int. J. Comput. Vis., 89 (2010), pp. 266–286.

[8] J. Cates, P.T. Fletcher, M. Styner, H.C. Hazlett, and R. Whitaker, *Particle-based shape analysis of multi-object complexes*, in Medical Image Computing and Computer-Assisted Intervention—MICCAI 2008, Springer-Verlag, Berlin, Heidelberg, 2008, pp. 477–485.

[9] E. Ceyhan, M.F. Beg, C. Ceritoglu, L. Wang, J.C. Morris, J.G. Csernansky, M.I. Miller, and J.T. Ratnanather, *Quantization and analysis of hippocampal morphometric changes due to dementia of Alzheimer type using metric distances based on large deformation diffeomorphic metric mapping*, Comput. Med. Imaging Graph., 35 (2011), pp. 275–293.

[10] T.F. Chan and L.A. Vese, *Active contours without edges*, IEEE Trans. Image Process., 10 (2001), pp. 266–277.

[11] G. Charpiat, O. Faugeras, R. Keriven, and P. Maurel, *Approximations of shape metrics and application to shape warping and empirical shape statistics*, in Statistics and Analysis of Shapes, H. Krim and A. Yezzi, Jr., eds., Birkhäuser Boston, Boston, 2006, pp. 363–395.

[12] L.D. Cohen and I. Cohen, *Finite-element methods for active contour models and balloons for 2-D and 3-D images*, IEEE Trans. Pattern Anal. Mach. Intell., 15 (1993), pp. 1131–1147.

[13] T.F. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham, *Active shape models–Their training and application*, Comput. Vision Image Underst., 61 (1995), pp. 38–59.

[14] S.G. Costafreda, I.D. Dinov, Z. Tu, Y. Shi, C.-Y. Liu, I. Kloszewska, P. Mecocci, et al., *Automated hippocampal shape analysis predicts the onset of dementia in mild cognitive impairment*, NeuroImage, 56 (2011), pp. 212–219.

[15] J.G. Csernansky, S. Joshi, L. Wang, J.W. Haller, M. Gado, J.P. Miller, U. Grenander, and M.I. Miller, *Hippocampal morphometry in schizophrenia by high dimensional brain mapping*, Proc. Natl. Acad. Sci. USA, 95 (1998), pp. 11406–11411.

[16] C. Davatzikos, A. Genc, D. Xu, and S.M. Resnick, *Voxel-based morphometry using the RAVENS maps: Methods and validation using simulated longitudinal atrophy*, NeuroImage, 14 (2001), pp. 1361–1369.

[17] L. Dice, *Measure of the amount of ecological association between species*, Ecology, 26 (1945), pp. 297–302.

[18] D.L. Dierker, E. Feczko, J.R. Pruett, S.E. Petersen, B.L. Schlaggar, J.N. Constantino, J.W. Harwell, T.S. Coalson, and D.C. Van Essen, *Analysis of cortical shape in children with simplex autism*, Cereb. Cortex, to appear.

[19] I.L. Dryden and K.V. Mardia, *Statistical Shape Analysis*, John Wiley & Son, Chichester, UK, 1998.

[20] M.-P. Dubuisson and A.K. Jain, *A modified Hausdorff distance for object matching*, in Proceedings of the 12th IAPR International Conference on Pattern Recognition (ICPR), Vol. 1, IEEE Computer Society Press, Los Alamitos, CA, 1994, pp. 566–568.

[21] I. Eckstein, A.A. Joshi, C.-C.J. Kuo, R. Leahy, and M. Desbrun, *Generalized surface flows for deformable registration and cortical matching*, in Medical Image Computing and Computer-Assisted Intervention—MICCAI 2007, Springer-Verlag, Berlin, Heidelberg, 2007, pp. 692–700.

[22] T. Fletcher, T. Venkatasubramanian, and S. Joshi, *The geometric median on Riemannian manifolds with application to robust atlas estimation*, NeuroImage, 45 (2009), pp. S143–S152.
[23] Y. Gao and S. Bouix, *Synthesis of realistic subcortical anatomy with known surface deformations*, in Mesh Processing in Medical Image Analysis 2012, Springer-Verlag, Berlin, Heidelberg, 2012, pp. 80–88.
[24] Y. Gao, T. Riklin Raviv, and S. Bouix, *Shape analysis, a field in need of careful validation*, Hum. Brain Mapp., to appear; available online at http://onlinelibrary.wiley.com/doi/10.1002/hbm.22525/full.
[25] G. Gerig, M. Styner, M. Shenton, and J. Lieberman, *Shape versus size: Improved understanding of the morphology of brain structures*, in Medical Image Computing and Computer-Assisted Intervention—MICCAI 2001, Springer-Verlag, Berlin, Heidelberg, pp. 24–32.
[26] P. Golland, W.E. Grimson, M.E. Shenton, and R. Kikinis, *Detection and analysis of statistical differences in anatomical shape*, Med. Image Anal., 9 (2005), pp. 69–86.
[27] M.S. Hassouna and A.A. Farag, *Multistencils fast marching methods: A highly accurate solution to the Eikonal equation on Cartesian domains*, IEEE Trans. Pattern Anal. Mach. Intell., 29 (2007), pp. 1563–1574.
[28] Y. Hirayasu, M.E. Shenton, D.F. Salisbury, C.C. Dickey, I.A. Fischer, P. Mazzoni, T. Kisler, et al., *Lower left temporal lobe MRI volumes in patients with first-episode schizophrenia compared with psychotic patients with first-episode affective disorder and normal subjects*, Amer. J. Psychiatry, 155 (1998), pp. 1384–1391.
[29] J. Joseph, C. Warton, S.W. Jacobson, J.L Jacobson, C.D. Molteno, A. Eicher, P. Marais, O.R. Phillips, K.L. Narr, and E.M. Meintjes, *Three-dimensional surface deformation-based shape analysis of hippocampus and caudate nucleus in children with fetal alcohol spectrum disorders*, Hum. Brain Mapp., 35 (2014), pp. 659–672.
[30] S.H. Joshi, K.L. Narr, O.R. Philips, K.H. Nuechterlein, R.F. Asarnow, A.W. Toga, and R.P. Woods, *Statistical shape analysis of the corpus callosum in schizophrenia*, NeuroImage, 64 (2013), pp. 547–559.
[31] K. Kasai, M.E. Shenton, D.F. Salisbury, Y. Hirayasu, T. Onitsuka, M.H. Spencer, D.A. Yurgelun-Todd, R. Kikinis, F.A. Jolesz, and R.W. McCarley, *Progressive decrease of left Heschl gyrus and planum temporale gray matter volume in first-episode schizophrenia: A longitudinal magnetic resonance imaging study*, Arch. Gen. Psychiatry, 60 (2003), pp. 766–775.
[32] M. Kass, A.P. Witkin, and D. Terzopoulos, *Snakes: Active contour models*, Int. J. Comput. Vis., 1 (1988), pp. 321–331.
[33] A. Kelemen, G. Szekely, and G. Gerig, *Elastic model-based segmentation of 3D neuroradiological data sets*, IEEE Trans. Med. Imaging, 10 (1999), pp. 828–839.
[34] M.S. Koo, J.J. Levitt, R.W. McCarley, L.J. Seidman, C.C. Dickey, M.A. Niznikiewicz, M.M. Voglmaier, et al., *Reduction of caudate volume in neuroleptic-naive female subjects with schizotypal personality disorder*, Biol. Psychiatry, 60 (2006), pp. 40–48.
[35] D.-J. Kroon, *Accurate Fast Marching*, MATLAB Central File Exchange, http://www.mathworks.com/matlabcentral/fileexchange/24531-accurate-fast-marching/ (6 February 2011).
[36] S. Kurtek, E. Klassen, Z. Ding, S. Jacobson, J. Jacobson, M. Avison, and A. Srivastava, *Parameterization-invariant shape comparisons of anatomical surfaces*, IEEE Trans. Med. Imaging, 30 (2011), pp. 849–858.
[37] S. Kurtek, E. Klassen, J.C. Gore, Z. Ding, and A. Srivastava, *Elastic geodesic paths in shape space of parameterized surfaces*, IEEE Trans. Pattern Anal. Mach. Intell., 34 (2012), pp. 1717–1730.
[38] I. Ladrón de Guevara, J. Muñoz, O.D. de Cózar, and E.B Blázquez, *Robust fitting of circle arcs*, J. Math. Imaging Vision, 40 (2011), pp. 147–161.
[39] J.J. Levitt, M. Styner, M. Niethammer, S. Bouix, M.S. Koo, M.M. Voglmaier, C.C. Dickey, et al., *Shape abnormalities of caudate nucleus in schizotypal personality disorder*, Schizophr. Res., 110 (2009), pp. 127–139.
[40] J.J. Levitt, C.-F. Westin, P.G. Nestor, R.S.J. Estepar, C.C. Dickey, M.M. Voglmaier, L.J. Seidman, et al., *Shape of the caudate nucleus and its cognitive correlates in neuroleptic-naive schizotypal personality disorder*, Biol. Psychiatry, 55 (2004), pp. 177–184.
[41] M. Lin, P.T. Fwu, C. Buss, E.P. Davis, K. Head, L.T. Muftuler, C.A. Sandman, and M.-Y. Su, *Developmental changes in hippocampal shape among preadolescent children*, Int. J. Dev. Neurosci., 31 (2013), pp. 473–481.

[42] R. Malladi, J.A. Sethian, and B.C. Vemuri, *Shape modeling with front propagation: A level set approach*, IEEE Trans. Pattern Anal. Mach. Intell., 17 (1995), pp. 158–175.
[43] J.F. Mangin, F. Poupon, E. Duchesnay, D. Rivière, A. Cachia, D.L. Collins, A.C. Evans, and J. Régis, *Brain morphometry using 3D moment invariants*, Med. Image Anal., 8 (2004), pp. 187–196.
[44] MATLAB, *Optimization Toolbox*, The MathWorks, http://www.mathworks.com/help/optim/index.html (1994–2014).
[45] MATLAB, *Statistics Toolbox*, The MathWorks, https://www.mathworks.com/help/stats/index.html (1994–2014).
[46] F. Memoli and G. Sapiro, *A theoretical and computational framework for isometry invariant recognition of point cloud data*, Found. Comput. Math., 5 (2005), pp. 313–347.
[47] J.J. Moré and D.J. Thuente, *Line search algorithms with guaranteed sufficient decrease*, ACM Trans. Math. Software, 20 (1994), pp. 286–307.
[48] J.A. Nelder and R. Mead, *A simplex method for function minimization*, Comput. J., 7 (1965), pp. 308–313.
[49] T. Nichols and S. Hayasaka, *Controlling the familywise error rate in functional neuroimaging: A comparative review*, Stat. Methods Med. Res., 12 (2003), pp. 419–446.
[50] M. Niethhammer, M. Reuter, F.-E. Wolter, S. Bouix, N. Peinecke, M.-S. Koo, and M.E. Shenton, *Global medical shape analysis using the Laplace-Beltrami spectrum*, in Medical Image Computing and Computer-Assisted Intervention—MICCAI 2007, Springer-Verlag, Berlin, Heidelberg, 2007, pp. 850–857.
[51] S. Osher and J.A. Sethian, *Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations*, J. Comput. Phys., 79 (1988), pp. 12–49.
[52] B. Paniagua, L. Bompard, J. Cates, R. Whitaker, M. Datar, C. Vacheta, and M. Styner, *Combined SPHARM-PDM and entropy-based particle systems shape analysis framework*, in Medical Imaging 2012: Biomedical Applications in Molecular, Structural, and Functional Imaging, Proc. SPIE 8317, SPIE, Bellingham, WA, 2012, 83170L.
[53] B. Paniagua, L. Cevidanes, H. Zhu, and M. Styner, *Outcome quantification using SPHARM-PDM toolbox in orthognathic surgery*, Int. J. Comput. Assist. Radiol. Surg., 6 (2011), pp. 617–626.
[54] N. Paragios and R. Deriche, *Geodesic active regions: A new paradigm to deal with frame partition problems in computer vision*, J. Vis. Commun. Image Rep., 13 (2002), pp. 249–268.
[55] K.M. Pohl, J. Fisher, S. Bouix, M.E. Shenton, R.W. McCarley, W.E. Grimson, R. Kikinis, and W.M. Wells, *Using the logarithm of odds to define a vector space on probabilistic atlases*, Med. Image Anal., 11 (2007), pp. 465–477.
[56] M. Reuter, F.-E. Wolter, and N. Peinecke, *Laplace-Beltrami spectra as 'Shape-DNA' of surfaces and solids*, Comput. Aided Des., 38 (2006), pp. 342–366.
[57] M. Rezanejad and K. Siddiqi, *Flux graphs for 2D shape analysis*, in Shape Perception in Human and Computer Vision, Springer-Verlag, London, 2013, pp. 41–54.
[58] T. Riklin Raviv, Y. Gao, J.J. Levitt, and S. Bouix, *Statistical shape analysis for population studies via level-set based shape morphing*, in Computer Vision—ECCV 2012. Workshops and Demonstrations, Lecture Notes in Comput. Sci. 7583, Springer-Verlag, Berlin, Heidelberg, 2012, pp. 42–51.
[59] T. Riklin Raviv, K. Van Leemput, B.H. Menze, W.M. Wells, and P. Golland, *Segmentation of image ensembles via latent atlases*, Med. Image Anal., 14 (2010), pp. 654–665.
[60] M.R. Sabuncu, B.T. Yeo, K. Van Leemput, T. Vercauteren, and P. Golland, *Asymmetric image-template registration*, in Medical Image Computing and Computer-Assisted Intervention—MICCAI 2009, Springer-Verlag, Berlin, Heidelberg, 2009, pp. 565–573.
[61] C.C. Schultz, G. Wagner, K. Koch, C. Gaser, M. Roebel, C. Schachtzabel, I. Nenadic, J.R. Reichenbach, H. Sauer, and R.G. Schlösser, *The visual cortex in schizophrenia: Alterations of gyrification rather than cortical thickness—A combined cortical shape analysis*, Brain Struct. Funct., 218 (2013), pp. 51–58.
[62] M.E. Shenton, C.C. Dickey, M. Frumin, and R.W. McCarley, *A review of MRI findings in schizophrenia*, Schizophr. Res., 49 (2001), pp. 1–52.
[63] A. Sotiras, C. Davatzikos, and N. Paragios, *Deformable medical image registration: A survey*, IEEE Trans. Med. Imaging, 32 (2013), pp. 1153–1190.
[64] A. Srivastava, E. Klassen, S.H. Joshi, and I.H. Jermyn, *Shape analysis of elastic curves in euclidean spaces*, IEEE Trans. Pattern Anal. Mach. Intell., 33 (2011), pp. 1415–1428.

[65] M. Styner, G. Gerig, J. Lieberman, D. Jones, and D. Weinberger, *Statistical shape analysis of neuroanatomical structures based on medial models*, Med. Image Anal., 7 (2003), pp. 207–220.
[66] M. Styner, J.A. Lieberman, D. Pantazis, and G. Gerig, *Boundary and medial shape analysis of the hippocampus in schizophrenia*, Med. Image Anal., 8 (2004), pp. 197–203.
[67] M. Styner, I. Oguz, S. Xu, C. Brechbühler, D. Pantazis, J. Levitt, M. Shenton, and G. Gerig, *Framework for the statistical shape analysis of brain structures using SPHARM-PDM*, Insight J., 1071 (2006), pp. 242–250.
[68] Z.Y. Sun, S. Klöppel, D. Rivière, M. Perrot, R. Frackowiak, H. Siebner, and J.F. Mangin, *The effect of handedness on the shape of the central sulcus*, NeuroImage, 60 (2012), pp. 332–339.
[69] T.B. Terriberry, J.N. Damon, S.M. Pizer, S.C. Joshi, and G. Gerig, *Population-based fitting of medial shape models with correspondence optimization*, in Information Processing in Medical Imaging, Lecture Notes in Comput. Sci. 4584, Springer-Verlag, Berlin, Heidelberg, 2007, pp. 700–712.
[70] P.M. Thompson, K.M. Hayashi, G.I. De Zubicaray, A.L. Janke, S.E. Rose, J. Semple, M.S. Hong, et al., *Mapping hippocampal and ventricular change in Alzheimer disease*, NeuroImage, 22 (2004), pp. 1754–1766.
[71] F. Tomaiuolo, M. Di Paola, B. Caravale, S. Vicari, M. Petrides, C. Caltagirone, *Morphology and morphometry of the corpus callosum in Williams syndrome: A T1-weighted MRI study*, Neuroreport, 13 (2002), pp. 2281–2284.
[72] Y. Wang, J. Zhang, B. Gutman, T.F. Chan, J.T. Becker, H.J. Aizenstein, O.L. Lopez, R.J. Tamburo, A.W. Toga, and P.M. Thompson, *Multivariate tensor-based morphometry on surfaces: Application to mapping ventricular abnormalities in HIV/AIDS*, NeuroImage, 49 (2010), pp. 2141–2157.
[73] G. Wu, H. Jia, Q. Wang, and D. Shen, *Sharpmean: Groupwise registration guided by sharp mean image and tree-based registration*, NeuroImage, 56 (2011), pp. 1968–1981.
[74] Y. Xie, J. Ho, and B.C. Vemuri, *Image atlas construction via intrinsic averaging on the manifold of images*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2010, pp. 2933–2939.
[75] S. Ying, G. Wu, Q. Wang, and D. Shen, *Groupwise registration via graph shrinkage on the image manifold*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2013, pp. 2323–2330.
[76] P.A. Yushkevich, H. Zhang, and J.C. Gee, *Continuous medial representation for anatomical structures*, Med. Image Anal., 25 (2006), pp. 1547–1564.
[77] T. Zhang and C. Davatzikos, *ODVBA: Optimally-discriminative voxel-based analysis*, IEEE Trans. Med. Imaging, 30 (2011), pp. 1441–1454.
[78] T. Zhang and C. Davatzikos, *Optimally-discriminative voxel-based morphometry significantly increases the ability to detect group differences in schizophrenia, mild cognitive impairment, and Alzheimer's disease*, NeuroImage, 79 (2013), pp. 94–110.
[79] K. Zilles, N. Palomero-Gallagher, and K. Amunts, *Development of cortical folding during evolution and ontogeny*, Trends Neurosci., 36 (2013), pp. 275–284.