Global Medical Shape Analysis Using the Volumetric Laplace Spectrum

Martin Reuter*, Marc Niethammer†, Franz-Erich Wolter‡, Sylvain Bouix†, Martha Shenton†

*Dept. of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139
reuter@mit.edu

†Brigham and Women's Hospital
Harvard Medical School
Boston, MA 02215
{marc,sylvain,shenton}@bwh.harvard.edu

‡Inst. für Mensch-Maschine-Kommunikation
Leibniz Universität Hannover
30167 Hannover, Germany
fwolter@mit.edu

Abstract

This paper proposes to use the volumetric Laplace spectrum as a global shape descriptor for medical shape analysis. The approach allows for shape comparisons using minimal shape preprocessing. In particular, no registration, mapping, or remeshing is necessary. All computations can be performed directly on the voxel representations of the shapes. The discriminatory power of the method is tested on a population of female caudate shapes (brain structure) of normal control subjects and of subjects with schizotypal personality disorder. The behavior and properties of the volumetric Laplace spectrum are discussed extensively for both the Dirichlet and Neumann boundary condition showing advantages of the Neumann spectra. Both, the computations of spectra on 3D voxel data for shape matching as well as the use of the Neumann spectrum for shape analysis are completely new.

1. Introduction

Morphometric studies of brain structures have classically been based on volume measurements. More recently, shape studies of gray matter brain structures have become popular. Methodologies for shape comparison may be divided into global and local shape analysis approaches. While local shape comparisons [1], [2], [3] yield powerful, spatially localized results that are relatively straightforward to interpret, they usually rely on a number of preprocessing steps. In particular, one-to-one correspondences between surfaces need to be established, shapes need to be registered and resampled, possibly influencing shape comparisons. While global shape comparison cannot spatially localize shape changes, global approaches may be formulated with a significantly reduced number of assumptions and preprocessing steps, staying as true as possible to the original data.

This paper describes a methodology for global shape comparison based on the Laplace-Beltrami spectrum [4], [5], [6], [7]. In particular, the paper focuses on the *volumetric* Laplace spectrum of three-dimensional solids. Previous approaches for global shape analysis in medical imaging include the use of invariant moments [8], the shape index [9], and global shape descriptors based on spherical harmonics [10]. The proposed methodology based on the Laplace-Beltrami spectrum differs in the following ways from these previous approaches:

- It works directly for any Riemannian manifold, whereas spherical harmonics based methods require spherical representations, and invariant moments do not easily generalize to arbitrary Riemannian manifolds. It may be used to analyze surface, solids, non-spherical objects, etc. in different representations.

- Only minimal preprocessing of the data is required. Three dimensional volume data may be represented by its boundary surface, separating the object interior from its exterior or by the volume itself (a volumetric, region-based approach). In the former case, the extraction of a surface approximation from a (possibly manually segmented) binary image volume is the only preprocessing step required. In the volumetric case, which is the focus of this paper, even this preprocessing step can be avoided and computations may be performed directly on the voxels of a given binary segmentation. This is in sharp contrast to other shape comparison methods, requiring object registration, remeshing, etc. The Laplace-Beltrami spectrum is invariant to rigid transformations, isometries, and to object discretization (as long as the discretization is sufficiently accurate).

This paper supplements previous surface-based Laplace-Beltrami shape analysis work on medical data [11] as well as work on pixel data (images) [7]. Furthermore, it extends the computation of volumetric

0-7695-3005-2/07 $25.00 © 2007 IEEE
DOI 10.1109/CW.2007.42

spectra (i.e., spectra of 3D solids) for the purpose of shape analysis to voxel data (described initially in [6] for tetrahedrizations). For the first time, Neumann spectra are used as shape descriptors, with powerful properties for coarse object discretizations. Furthermore, in addition to the eigenvalues, eigenfunction analyses are introduced using 2D examples to shed some light on the behavior of the spectra. A detailed implementation description is provided based on the general concept of the Laplace-Beltrami operator.

Section 2 describes the theoretical background of the Laplace-Beltrami spectrum and the numerical computation of the eigenvalues. Normalizations of the spectra, properties of the Neumann spectrum as well as the influence of noise and of the discretization are investigated. Section 3 explains the statistical methods used for the analysis of groups of Laplace-Beltrami spectra. Results for two groups of female caudate shapes are given in Section 4. The paper concludes with a summary and outlook in Section 5.

2. Shape-DNA: The Laplace-Beltrami Spectrum

In this section we introduce the necessary background for the computation of the Laplace-Beltrami spectrum beginning sequence (also called "Shape-DNA"). The "Shape-DNA" is a fingerprint or signature computed only from the intrinsic geometry of an object. More precisely: it is the beginning sequence of the spectrum of the Laplace-Beltrami operator defined for real valued functions on Riemannian manifolds. It can be used to identify and compare objects like surfaces and solids independently of their representation, position and (if desired) independently of their size. This methodology was first introduced in [4] though a sketchy description of basic ideas and goals of this methodology is already contained in [12]. The Laplace-Beltrami spectrum can be regarded as the set of squared frequencies (the so called natural or resonant frequencies) that are associated to the eigenmodes of an oscillating membrane defined on the manifold. Let us review the basic theory in the general case (for more details refer to [6] and especially [5]).

2.1. Definitions

Let f be a real-valued function, with $f \in C^2$, defined on a Riemannian manifold M (differentiable manifold with Riemannian metric). The **Laplace-Beltrami Operator** Δ is:

$$\Delta f := \operatorname{div}(\operatorname{grad} f) \tag{1}$$

with $\operatorname{grad} f$ the **gradient** of f and div the **divergence** on the manifold (Chavel [13]). The Laplace-Beltrami operator is a linear differential operator. It can be calculated in local coordinates. Given a local parametrization $\psi : \mathbb{R}^n \to \mathbb{R}^{n+k}$ of a submanifold M of $\mathbb{R}^{n+k}$ with

$$\begin{aligned} g_{ij} &:= \langle \partial_i \psi, \partial_j \psi \rangle, \quad G &&:= (g_{ij}), \\ W &:= \sqrt{\det G}, &(g^{ij}) &:= G^{-1}, \end{aligned} \tag{2}$$

(where $i, j = 1, \ldots, n$ and $\det$ denotes the determinant) the Laplace-Beltrami operator becomes:

$$\Delta f = \tfrac{1}{W} \textstyle\sum_{i,j} \partial_i (g^{ij} W \partial_j f). \tag{3}$$

If M is a domain in the Euclidean plane $M \subset \mathbb{R}^2$, the Laplace-Beltrami operator reduces to the well known Laplacian:

$$\Delta f = \frac{\partial^2 f}{(\partial x)^2} + \frac{\partial^2 f}{(\partial y)^2}. \tag{4}$$

The wave equation $\Delta u = u_{tt}$ may be decomposed into its time-dependent and its spatially dependent parts $u(\boldsymbol{x}, t) = f(\boldsymbol{x}) a(t)$. Separating variables in the wave equation yields [14]

$$\frac{\Delta f}{f} = \frac{a_{tt}}{a} = -\lambda, \quad \lambda = const.$$

Thus, the vibrational modes may be obtained through the **Helmholtz equation** (also known as the **Laplacian eigenvalue problem**) on manifold M with or without boundary

$$\Delta f = -\lambda f. \tag{5}$$

The solutions of this equation represent the spatial part of the solutions of the wave equation (with an infinite number of eigenvalue λ and eigenfunction f pairs). In the case of M being a planar region, $f(u, v)$ in equation (5) can be understood as the natural vibration form (also **eigenfunction**) of a homogeneous membrane with the **eigenvalue** λ. Any properties of the material are ignored. The standard boundary condition of a fixed membrane is the **Dirichlet boundary condition** where $f \equiv 0$ on the boundary of the domain (see Figure 1 for two eigenfunctions of the disk). In this work we also use the **Neumann boundary condition** where the derivative in the normal direction of the boundary $\frac{\partial f}{\partial n} \equiv 0$ is zero along the boundary. Here the normal direction n of the boundary should not be confused with a normal of the embedded Riemannian manifold (e.g., surface normal). n is normal to the boundary and tangential to the manifold. We will speak of the Dirichlet or Neumann spectrum depending on the used boundary condition.

The **spectrum** is defined to be the family of eigenvalues of the Helmholtz equation (eq. 5), consisting of

a diverging sequence $0 \leq \lambda_1 \leq \lambda_2 \leq \cdots \uparrow +\infty$, with each eigenvalue repeated according to its multiplicity and with each associated finite dimensional eigenspace. In the case of the Neumann boundary condition the first eigenvalue λ_1 is always equal to zero, because in this case the constant functions are non trivial solutions of the Helmholtz equation.

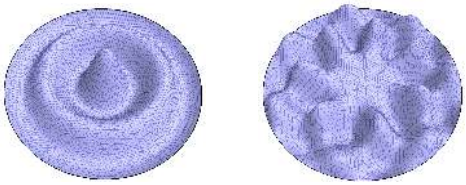

Figure 1. Eigenfunction 30 and 50 of the disk.

Because of the rather simple Euclidean nature of the voxel representations used later, the more general definitions given above are not necessarily needed to understand the computation and results of this paper. Nevertheless, we described the general Riemannian background to align this paper with previous work on parametrized surfaces and solids and with future work. Furthermore, this approach clarifies that the eigenvalues are indeed isometric invariants with respect to the Riemannian manifold. Note that two solid bodies embedded in $\mathbb{R}^3$ are isometric if and only if they are congruent (translated, rotated and mirrored). In the surface case (e.g., discussed in [11]) this is not true, since non-congruent but isometric surfaces exist resulting in a large class of non-congruent surfaces having the same spectrum.

2.2. Properties

The following paragraphs describe well known results on the Laplace-Beltrami operator and its spectrum. The spectrum is an **isometric invariant** as it only depends on the gradient and divergence which in turn are defined to be dependent only on the Riemannian structure of the manifold (eq. 3), i.e., the intrinsic geometry. Furthermore, we know that **scaling** an n-dimensional manifold by the factor a results in eigenvalues scaled by the factor $\frac{1}{a^2}$. Therefore, by normalizing the eigenvalues, shape can be compared regardless of the object's scale (and position as mentioned earlier). The spectrum **depends continuously** on the shape of the membrane [14]. The spectrum does **not characterize** the shape completely, since some non-isometric manifolds with the same spectrum exist (for example see [15]). Nevertheless these artificially constructed cases appear to be very rare cf. [6] (e.g., in the plane they have to be concave with corners and until now only isospectral pairs could be found). A substantial amount of **geometrical and topological** information is known to be contained in the spectrum [16] (Dirichlet as well as Neumann). Even though we cannot crop a spectrum without loosing information, we showed in [5] that it is possible to extract important information just from the first few Dirichlet eigenvalues (approx. 500). The spectra has more **discrimination power** than simple measures like surface area, volume or even the shape index (the normalized ratio between surface area and volume, $SI = A^3/(36\pi V^2) - 1$). See Figure 2 for simple shapes with same shape index, that the Laplace-Beltrami spectrum can distinguish. For more properties see [6], [5].

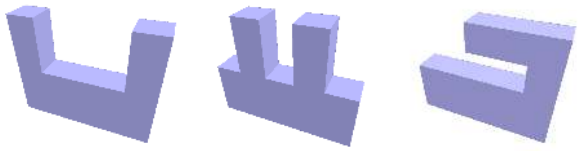

Figure 2. Objects with same shape index but different spectra

2.3. Variational Problem

For the numerical computation, the first step is to translate the Helmholtz equation into a **variational problem**. To accomplish this, we need Green's formula

$$\iint \varphi \Delta f d\sigma = -\oint \varphi \frac{\partial f}{\partial n} ds - \iint \nabla(f,\varphi) d\sigma \quad (6)$$

(Blaschke [17] p.227) with the Nabla operator

$$\nabla(f,\varphi) := Df\ G^{-1}\ (D\varphi)^T = \sum (\partial_i f\ g^{ij}\ \partial_j \varphi) \quad (7)$$

where $Df = (\partial_1 f, \partial_2 f, ...)$. With the Dirichlet ($f, \varphi \equiv 0$) or the Neumann ($\frac{\partial f}{\partial n} \equiv 0$) boundary condition Green's formula simply looses the boundary term. We multiply the Helmholtz equation with test functions $\varphi \in C^2$, complying with the boundary condition. By integrating over the area and using Green's formula we obtain:

$$\begin{aligned} & \varphi \Delta f && = -\lambda \varphi f \\ \Leftrightarrow\ & \iint Df\ G^{-1}\ (D\varphi)^T\ d\sigma && = \lambda \iint \varphi f\ d\sigma \end{aligned} \quad (8)$$

(with $d\sigma = W\ du\ dv$ being the surface element in the 2D case or the volume element $d\sigma = W\ du\ dv\ dw$ in the 3D case). Every function $f \in C^2$ on the open domain and continuous on the boundary solving the variational equation for all test functions φ is a solution to the Laplace eigenvalue problem (Braess [18], p.35). We will use the variational formulation to obtain a system of equations to construct an approximation of the solution.

2.4. Implementation

To solve the Helmholtz equation on any Riemannian manifold we employ the **Finite Element Method** [19]. We choose a tessellation of the manifold into so called elements (e.g., triangles or in this work cuboid voxel). We then choose linearly independent test functions with up to cubic degree (the so called form functions F_i) defined on the cuboid voxel elements. These high degree functions lead to a better approximation and consequently to better results, but because of their higher degree of freedom more node points have to be inserted into the elements. See [5] or [6] for a more detailed description of the **discretization** used in the FE method that finally leads to the following general eigenvalue problem $AU = \lambda BU$ with the matrices

$$\begin{aligned} A &= (a_{lm}) &&:= \left(\iint DF_l \; G^{-1} \; (DF_m)^T d\sigma\right), \\ B &= (b_{lm}) &&:= \left(\iint F_l F_m d\sigma\right). \end{aligned} \tag{9}$$

Where F_l is the form function being one at node l and zero at all other nodes. Here U is the vector $(U_1, \ldots, U_n)$ containing the unknown values of the solution at each node and A, B are **sparse positive (semi-) definite symmetric matrices**. The solution vectors U (eigenvectors) with corresponding eigenvalues λ can then be calculated. The eigenfunctions are approximated by $\sum U_i F_i$. In case of the Dirichlet boundary condition, the boundary nodes do not get a number assigned to them and do not show up in this system. In case of a Neumann boundary condition, every node is treated exactly the same, no matter if it is a boundary node or an inner node. Since only a small number of eigenvalues is needed, a Lanczos algorithm can be employed to solve this large symmetric eigenvalue problem much faster than with a direct method. In this work we use the ARPACK package to compute the eigenfunctions and -values starting from the smallest eigenvalue in increasing eigenvalue order. It should be noted that the integrals mentioned above are basically independent of the mesh (as long as the mesh fulfills some refinement and condition standards). Beyond that, this method is completely independent of the given parametrization.

2.5. Form Functions

In order to compute the entries of the two matrices A and B (equation 9) we need the form functions F_i and their partial derivatives $(\partial_k F_i)$ in addition to the metric values from equation (2). The form functions are a basis of functions representing the solution space.

Any piecewise polynomial function F of degree d can easily be linearly combined by a base of global form functions F_i (of same degree d) having the value one at a specific nodal point i and zero at the others. For linear functions it is sufficient to use only the vertices of the mesh as nodal points. In case of a voxel the values at the 8 vertices are sufficient to define a tri-linear function in the inside $c_1+c_2u+c_3v+c_4w+c_5uv+c_6uw+c_7vw+c_8uvw$, for higher degrees further nodal points have to be inserted. If we look at a 2D example (a single triangle of a triangulation), a linear function above the triangle can be linearly combined by the three form functions at the corners. These local functions can be defined on the unit triangle (leg length one) and mapped to an arbitrary triangle. Figure 3 shows examples of a linear and a quadratic local form function for triangles. It

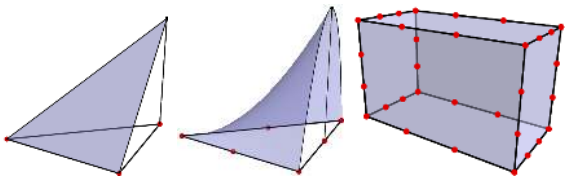

Figure 3. A linear and a quadratic form function and location of 32 nodes for cubic serendipity FEM voxel.

can be seen that the form function has the value one at exactly one node and zero at all the others. Note that in the case of the quadratic form function new nodes have to be introduced at the midpoint of each edge, because quadratic functions in two variables have six degrees of freedom. On each element containing n nodes exactly n local form functions will be constructed this way. The form functions and their derivatives can be defined explicitly on the unit triangle or unit cube. Since high order approximations lead to much better results, we mainly use cubic form functions of the serendipity family for the computation of the spectra in this paper. To set up these functions over a cuboid domain new nodes have to be inserted (two nodes along each edge makes 32 nodes together with the vertices, see figure 3). A cubic function of the serendipity family with three variables has 32 degrees of freedom, that can be fixed by giving the function values at these 32 locations. More details on the construction of these local functions can be found in most FEM books (e.g. Zienkiewicz [19]). For each element the results of the integrals (9) are calculated for every combination m, l of nodes in the element and added to the corresponding entry in the matrix A or B. Since this entry differs only from zero when the associated global form functions F_i overlap (i.e. the associated nodes share the same element) the matrices A and B will be sparse.

2.6. Cuboid Voxel Elements

For piecewise flat objects the computation described above can be simplified, thus speeding up the construction of the two matrices A and B significantly. If the local geometry is flat we do not need to integrate numerically on the manifold since the metric G (see equation 2) is constant throughout each element. The integrals can be computed once for the unit element explicitly and then mapped linearly to the corresponding element. This makes the time consuming numerical integration process needed for curved surfaces or solids completely unnecessary.

As opposed to the case of a surface triangulation with a piecewise flat triangle mesh (with possibly different types of triangles), the uniform decomposition of a 3D solid into cuboid voxels leads to very simple Finite Elements. A parametrization over the unit cube of a cuboid with side length s_1, s_2, s_3 (and volume V) yields a diagonal first fundamental matrix $G = diag\left((s_1)^2, (s_2)^2, (s_3)^2\right)$ where $W = \sqrt{det(G)} = V$ and $G^{-1} = diag\left(\frac{1}{(s_1)^2}, \frac{1}{(s_2)^2}, \frac{1}{(s_3)^2}\right)$. These values are not only constant for an entire voxel, they are identical for every voxel (since the voxels are identical). Therefore we can pre-compute the contribution of every voxel to the matrices A and B once for the whole problem after setting up the form functions F_l as described above:

$$\begin{aligned} a_{l(i),m(j)} &+= V \iiint_0^1 \left(\sum_{k=1}^{3} \frac{\partial_k F_i \ \partial_k F_j}{(s_k)^2}\right) du\ dv\ dw \\ b_{l(i),m(j)} &+= V \iiint_0^1 F_i F_j \ du\ dv\ dw. \end{aligned} \tag{10}$$

The local indices i, j label the (e.g. 32) nodes of the cuboid voxel element and thus the corresponding local form functions and their partial derivatives. These integrals can be pre-computed for every combination i, j. In order to add $(+=)$ these local results into the large matrices A and B only a lookup of the global vertex indices $l(i), m(j)$ for each voxel is necessary. Therefore the construction of the matrices A and B can be accomplished in $O(n)$ time for n elements.

2.7. Normalizing the Spectrum

As mentioned above, the Laplace-Beltrami spectrum is a diverging sequence. Analytic solutions for the spectrum and the eigenfunctions are only known for a limited number of shapes (e.g., the cuboid, the cylinder, the solid ball). The eigenvalues for the cuboid

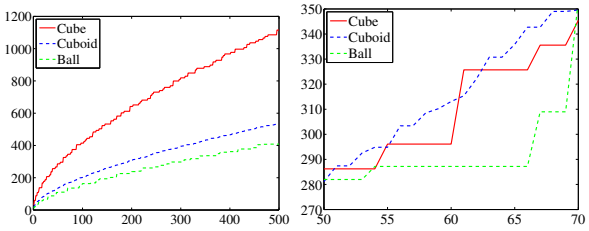

Figure 4. Unnormalized and normalized (zoom-in) exact Spectra of Cube, Cuboid, Ball.

with side length s_1, s_2 and s_3 for example are

$$\lambda_{M,N,O} = \pi^2 \left(\frac{M^2}{(s_1)^2} + \frac{N^2}{(s_2)^2} + \frac{O^2}{(s_3)^2}\right)$$

with $M, N, O \in \mathbb{N}^+$ for the Dirichlet case and $M, N, O \in \mathbb{N}$ for the Neumann case. In general the Dirichlet and Neumann eigenvalues of a 3D solid asymptotically tend to a curve dependent on the volume of the 3D manifold M:

$$\lambda_n \sim \left(\frac{6\pi^2 n}{\text{vol}(M)}\right)^{\frac{2}{3}}, \quad \text{as} \quad n \uparrow \infty. \tag{11}$$

Figure 4 (left) shows the discrete Dirichlet spectra of a unit cube ($V = 1$), a cuboid with side length $1, 1.5, 2$ ($V = 3$) and a unit ball ($V = \frac{4}{3}\pi$). It can be seen how the difference in volume manifests itself in different scalings of the eigenvalue asymptotes. A statistical method able to distinguish shapes needs to account for this diverging behavior so not to limit the analysis to an analysis of volume. Therefore the Laplace-Beltrami spectra should be normalized by volume. The zoom-in of the volume normalized spectra (Figure 4 right) shows that shape differences are preserved in the spectra after volume normalization.

2.8. Exactness of the Spectrum

To verify the accuracy of the numerically computed spectra, we compare the eigenvalues of a cuboid with side length $(1, 1.5, 2)$ and of a ball with radius one to the known exact values. In the case of the cuboid we computed the first 200 eigenvalues. The maximum absolute difference occurring in the Dirichlet spectra was less than 0.044 (which is less than 0.015 % relative error). This is due to the fact that the voxels represent the cuboid exactly without any approximation error at the boundary. The Neumann spectra have only a maximum absolute difference of less than 0.01 (which is less than 0.005% relative error), due to the higher resolution.

In case of the ball an exact representation by voxels is not possible, therefore the numerical results differ more strongly from the analytical ones especially

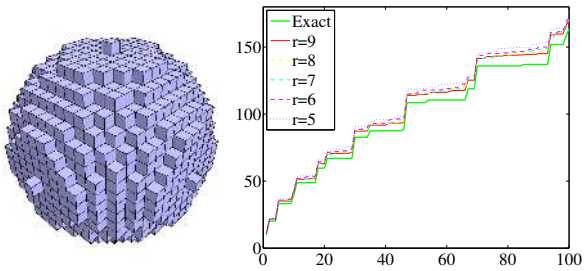

Figure 5. Approximation of the Ball.

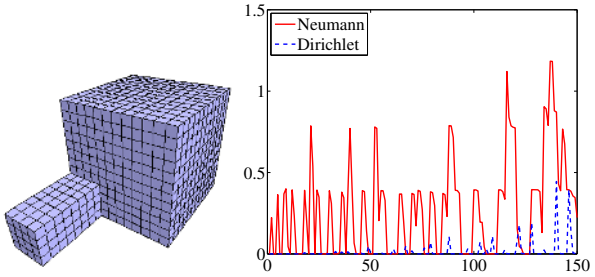

Figure 6. The first 150 eigenvalues of the cube with tail subtracted from the eigenvalues of the cube for the Dirichlet and Neumann case.

for high eigenvalues (up to 6% relative error for the first 100 Dirichlet eigenvalues). Since the exact values of the object represented by the voxelization are unknown, a fair analysis of the accuracy of the computation is difficult. Nevertheless, it is interesting to see that the numerical values closely approximate the exact ones of the ball the more voxels are used (see Figure 5, the value r describes the number of voxels used in the direction of the radius).

2.9. Neumann Spectrum

To demonstrate that Neumann spectra can be used to pick up significant geometric features much faster than Dirichlet spectra the eigenvalues of the cube with a tail (see Figure 6, left) were computed for the Neumann and the Dirichlet boundary condition and compared to the values of the cube. Figure 6 (right) shows the differences of the first 150 eigenvalues for the two different boundary conditions. Since the cube with tail has a larger volume, its eigenvalues are expected to be smaller than the values of the cube (see Section 2.2 2.). This fact is reflected in the graph (Figure 6 right) where the differences are always positive. It can be clearly seen that the Neumann spectrum picks up the differences much earlier than the Dirichlet spectrum. This is due to the fact that the Neumann boundary condition allows the solutions to swing at the boundary whereas the Dirichlet condition forces them to be

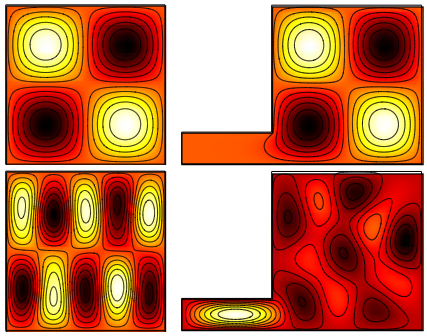

Figure 7. Dirichlet Eigenfunctions 4 (top) and 18 (bottom)

zero on the boundary, strongly reducing their freedom especially in the region of the tail.

A 2D example of a square with a tail (S_T) illustrates the different behaviors of Neumann and Dirichlet boundary conditions. Figures 7 and 8 depict a comparison of a few eigenfunctions of the square with tail (S_T) and of the unit square (S_1) for both the Dirichlet and the Neumann case. For the Dirichlet case (Figure 7), the lower eigenfunctions do not detect the attached tail (7 top). For higher frequencies the nodal domains shrink (7 bottom) until they are finally able to slip into the smaller features. Because of the strict Dirichlet boundary condition, this only happens starting with the 18th eigenfunction. From a signal processing point of view it is sensible that functions with higher frequencies can be used to analyze smaller features.

The Neumann spectrum behaves differently (Figure 8). Because of the higher degree of freedom (with respect to the free vibration of the eigenfunctions at the boundary), small features like the tail influence the eigenfunctions already very early. It is unnecessary to compare the smallest eigenvalue which is always zero with constant eigenfunctions. But already the eigenfunctions with the smallest frequency above zero (8 top) are very different from each other since one of the extrema is shifted into the tail. This is reflected in a change of more than 50% of the corresponding eigenvalue. The third eigenfunction (8 bottom) of S_T on the other hand is zero in the tail region and therefore almost identical with S. The corresponding eigenvalues are almost the same.

2.10. Influence of Noise

As demonstrated in Section 2.7, volume normalizations can lead to good spectral alignments. Nevertheless, having identical noise levels for the populations under investigation is essential, since different noise levels will affect surface areas differently (high noise

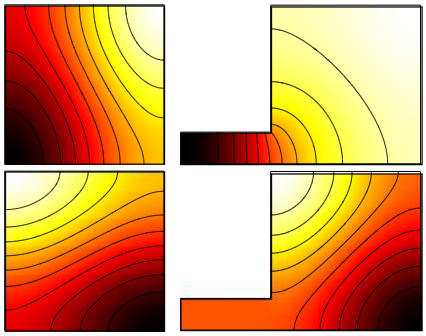

Figure 8. Neumann Eigenfunctions 2 (top) and 3 (bottom)

levels yield highly irregular bounding surfaces especially if the voxel resolution is low). Because surface area is contained in the spectrum this has an influence on the eigenvalues. Violating the assumption of similar noise levels therefore leads to the detection of noise level differences as opposed to shape differences, as demonstrated in Figure 9 where the spectra of the ball are depicted with different levels of added noise. A fixed probability for adding a voxel to or removing a voxel from the object boundary was chosen for each experiment. Only voxels maintaining 6-connectivity were added or removed. Increasing the noise level moves the corresponding spectra further apart. (For the analysis of identically acquired and processed shapes – e.g., obtained through manual segmentations of MRI data – a similar noise level is a reasonable assumption; we also assume the accuracy of the spectra calculations to stay the same for the whole population.)

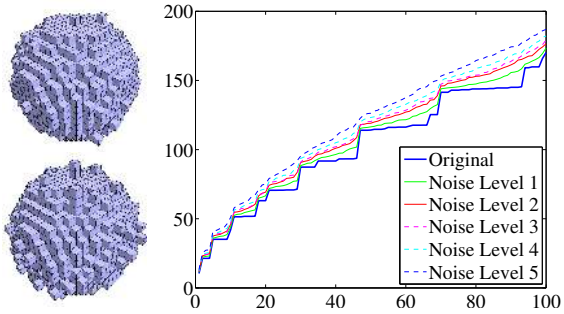

Figure 9. Influence of noise on the ball spectrum (Dirichlet). Noise-levels influence the spectral signature of a shape.

2.11. Influence of the Discretization

Domain discretizations can significantly affect computational results. Figure 10 shows a two-dimensional example for a domain consisting of a small and a large square connected by a thin rectangle. Insufficient degrees of freedom, due to a coarse discretization, can lead to insufficiently resolved vibrational modes. In

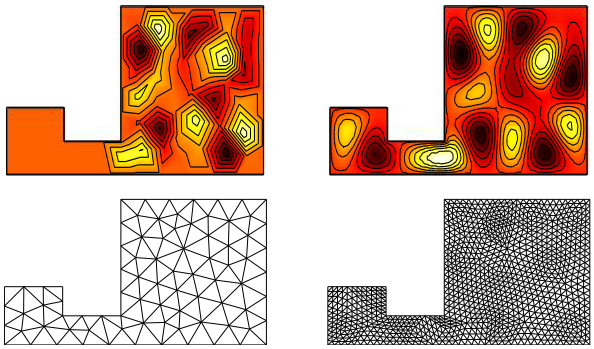

Figure 10. Dirichlet Eigenmode 19 for coarse discretizations (left) and fine (right).

particular, thin structures may simply be overlooked. The Neumann spectra are not influenced as strongly by the discretization.

3. Statistical analysis of LB spectra

The (possibly normalized) beginning sequence of the Laplace-Beltrami spectrum (called ShapeDNA) can be interpreted as a point $\boldsymbol{v} \in \mathbb{R}^n_{\geq 0}$ in the n-dimensional positive Euclidean space. Given the ShapeDNA $\boldsymbol{v}_i$ of many individual objects divided into two populations A and B we use permutation tests to compare group features to each other (200,000 permutations were used for all tests). We call a set of objects the object population. Permutation testing is a nonparametric, computationally simple way of establishing group differences by randomly permuting group labels. Let $S_A = \{\boldsymbol{v}_i\}$ and $S_B = \{\boldsymbol{v}_j\}$ denote two sets of ShapeDNA associated to individuals for group A and for group B respectively. Assume for example that we want to investigate if elements in S_A have on average a larger Euclidean norm than the elements in S_B (due to some external influences). A possible test statistic $stat$ would be the sum of the lengths of the elements in S_A ($stat := \sum_{\boldsymbol{v}_i \in S_A} \|\boldsymbol{v}_i\|$). We define the p-value to be the fraction of permutations (of labels A and B; keeping the number of elements per group fixed) having a greater or equal sum $stat$ than the original set S_A. The values of S_A will be considered significantly larger than the ones of S_B at a prespecified significance level α if $p \leq \alpha$ (taken as $\alpha = 5\%$ here). Note, that rejecting the null hypothesis of two populations being equal given a significance level α only implies that the probability of making a type I error (i.e., the probability of detecting false positives; "detecting a difference when there is none in reality") is α, but does not exclude the possibility of making such an error.

We use a two-sided, nonparametric, multivariate permutation test based on the maximum t-statistic to

analyze the high-dimensional spectral feature vectors (ShapeDNA). All shapes are initially corrected for brain volume differences. See [20] for details on permutation testing. The maximum t-statistic is chosen due to the usually small number of available samples in medical image analysis, compared to the dimensionality of the ShapeDNA feature vectors (preventing the use of the Hotelling T^2 statistic). It is defined as

$$stat = t_{max} := \max_{1 \le j \le N} \frac{|\bar{\boldsymbol{v}}_{A,j} - \bar{\boldsymbol{v}}_{B,j}|}{SE_j}. \quad (12)$$

Here, N is the vector dimension, $\bar{\boldsymbol{v}}_{A,j}$ indicates the mean of the j-th vector component of group A, and SE_j is the pooled standard error estimate of the j-th vector component, defined as

$$SE_j = \frac{\sqrt{(n_A - 1)\sigma_{A,j}^2 + (n_B - 1)\sigma_{B,j}^2}}{\sqrt{\frac{1}{n_A} + \frac{1}{n_B}}}, \quad (13)$$

where n_i is the number of subjects in group i (with $i \in A, B$) and $\sigma_{i,j}$ is the standard deviation of vector component j of group i. The maximum t-statistic is particularly sensitive to differences in at least one of the components of the feature vector [21]. It is a summary statistic, which allows for the detection of differences between feature vectors across populations. However, it does not determine which components show statistically significant differences.

Nevertheless, testing the individual statistical significance of vector components is possible. Such testing needs to be performed over a whole set of components, since it is usually not known beforehand which component of a ShapeDNA vector will be a good candidate for statistical testing. (I.e., we can in general not simply pick *one* individual vector component (eigenvalue) for statistical testing.) To account for multiple comparisons when testing over a whole set of vector components, the significance level needs to be adjusted (since "the chance of finding differences that are purely random in nature increases with the number of tests performed"). See [22], [20] for background on schemes for multiple comparison corrections or [11] for an application to population studies using the ShapeDNA.

4. Results

Volume measurements are the simplest means of morphometric analysis. While volume analysis results are easy to interpret, they only characterize one morphometric aspect of a structure. The following Sections introduce the volumetric Laplace spectrum as a method for a more complete global structural description using the analysis of a caudate as an exemplary brain structure involved in memory function, emotion processing and learning. For brevity all results are presented for the right caudate only.

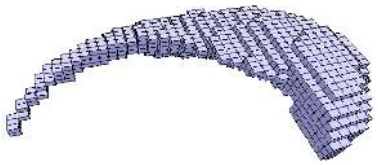

Figure 11. Example of a caudate shape consisting of voxels.

Magnetic Resonance Images (MRI) of the brains of thirty-two neuroleptic-naïve female subjects diagnosed with Schizotypal Personality Disorder (SPD) and of 29 female normal control subjects were acquired on a 1.5-T General Electric MR scanner. Spoiled-gradient recalled acquisition (SPGR) images (voxel dimensions 0.9375 x 0.9375 x 1.5 mm) were obtained coronally. The images were used to delineate the caudate nucleus (see Figure 11) and to estimate the intracranial content (ICC). For details see [23].

4.1. Laplace-Beltrami Spectrum Results

The Laplace-Beltrami spectrum was computed for two different shape normalizations:

(1) The shapes were volume normalized to **unit brain volume** (using the ICC measurements).
(2) The shapes were normalized to **unit caudate volume**.

In order to get more inner nodes especially into the very thin tail region of the caudate shapes we use the **dual** of the voxel graph. The dual voxel graph is the voxel image where each voxel of the original image (**regular** voxel image) is considered to be a vertex of the dual graph. Voxels inside the domain become inner vertices while voxel just outside of the domain's boundary now become boundary vertices. Figure 12 depicts a 2D case (using pixels) with the original regular domain on the left and the dual on the right. The dual graph enlarges the domain by half a layer thus creating new inner nodes needed especially inside the thinner features (tail). The use of the dual graph is necessary since a global refinement of the voxel leads to large FEM problems that may cause memory problems on a standard computer. The two populations show significant differences in volume, surface area and in their 2D surface spectra [11]. We expect to find significant differences also in their 3D volume spectra. The following paragraphs look first at the Dirichlet, then at the Neumann spectrum for both the regular and dual voxel graph.

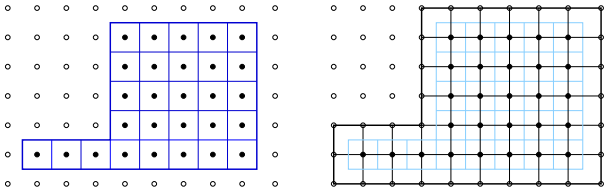

Figure 12. Regular pixel domain (left) and its dual (right).

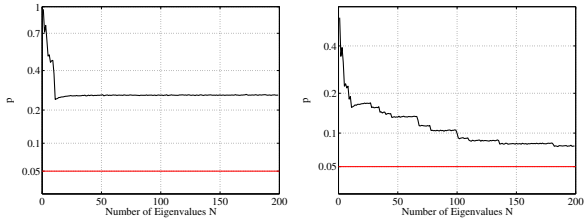

Figure 13. Accumulative maximum t-statistic results (brain volume normalized Dirichlet) for the regular (left) and the dual (right) voxel graphs.

Dirichlet Spectra. Figure 13 shows the statistical results for the regular voxel graph and its dual for the caudate populations. The graph shows the corresponding p value (see Section 3) when using the first N brain volume normalized eigenvalues of the spectra for the statistical analysis. We want to stress that not just a single eigenvalue is used but the whole beginning sequence of the first N values (therefore we call these plots accumulated statistic plots[1]). A p value below the 5% horizontal line is considered to be statistically significant. In such a case the beginning sequence of the spectrum is considered to be able to distinguish between the two caudate populations (NC and SPD). Figure 13 shows that the beginning sequence of the Dirichlet spectrum does not yield any statistically significant results. Employing the dual graph yields lower p-values when higher eigenvalues get involved. This observation is sensible, considering the fact that especially in the thin tail part of the caudates only very few inner nodes exist (see Section 2.11 for the effects of low resolution). The dual graph has more degrees of freedom, introducing inner nodes in the thin tail area, and therefore improves the result. For a detailed analysis employing the Dirichlet spectrum higher voxel resolutions would be necessary.

Neumann Spectra. As demonstrated in section 2.9, the Neumann spectrum can help to identify shape differences much earlier than the Dirichlet spectrum.

1. For a statistical study, the number of used eigenvalues *needs to be chosen and fixed beforehand.* The presented accumulated statistic plots are merely meant to illustrate the behavior of population statistics based on the Laplace-Beltrami spectrum for varying lengths of the spectral feature vectors.

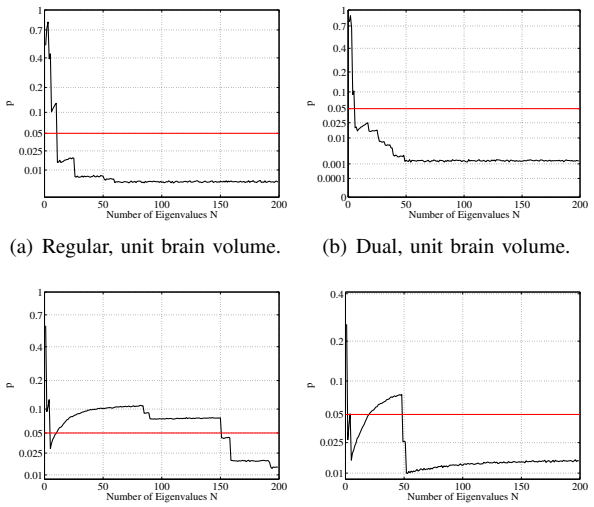

(a) Regular, unit brain volume. (b) Dual, unit brain volume.

(c) Regular, unit caudate volume. (d) Dual, unit caudate volume.

Figure 14. Accumulated statistic, Neumann boundary conditions.

Figure 14 shows the accumulated statistics plots for different normalizations of the Neumann spectra for the regular and the dual case. In both cases (regular and dual) the unit brain normalized spectra show very similar behavior (14 a,b): already very early, significant differences are detected. Because these differences might simply reflect the different caudate volumes (known to be significant) a normalization to unit caudate volume (14 c,d) is applied to reveal volume-independent shape differences. The eigenvalues in the regular case (14 c) do not show statistically significant shape differences until about 150 eigenvalues are involved, however, the dual case (14 d) shows a significant p value already for 50 or more eigenvalues used. In this case the frequencies become large enough (and therefore the nodal domains small enough) to detect the smaller features, presumably in the tail area. This assumption aligns with the Dirichlet case (Figure 13 dual) where better p-values are obtained when employing higher frequencies. Finally it should be noted that computing the twelve invariant moments (invariant with respect to similarity group actions) proposed by Mangin et al. [8] for the caudate shape population did not yield statistically significant shape differences.

5. Conclusion

This paper shows the applicability and utility of the volumetric Laplace spectrum with Dirichlet and Neumann boundary conditions for 3D shape analysis in medical imaging. In the context of shape matching the use of the Neumann spectra as well as the compu-

tations in 3D are completely new. It could be demonstrated that the Neumann spectra can detect statistically significant differences of the two populations (using the maximum t-statistic) not only in size, but also in shape, hinting at caudate shape differences in Schizotypal Personality Disorder. Computing the Dirichlet and Neumann spectra directly on the 3D voxel data is feasible on a standard desktop computer. The Neumann spectra are of interest, since they recognize shape differences much earlier than the Dirichlet spectra and also work much better if the voxel resolution is very low. Especially the higher eigenvalues yield statistically significant results, indicating true shape differences mainly in areas with smaller features (e.g., the tail area of the caudates). This result is promising since it demonstrates the possibility to compare shape based on a specific size of the features of interest (multiresolution shape matching). Future work will focus on this possibility.

Acknowledgment

This work was partially funded by a Humboldt Foundation postdoctoral fellowship to the first author and was supported in part by a Department of Veteran Affairs Merit Award (SB,MS), a Research Enhancement Award Program (MS), and National Institute of Health grants R01 MH50747 (SB,MS), K05 MH070047 (MS), and U54 EB005149 (MN,SB,MS).

References

[1] M. Styner, I. Oguz, S. Xu, C. Brechbühler, D. Pantazis, J. J. Levitt, M. E. Shenton, and G. Gerig, "Framework for the statistical shape analysis of brain structures using SPHARM-PDM," in *Proceedings of the Open Science Workshop at MICCAI*, 2006.

[2] D. Nain, M. Styner, M. Niethammer, J. J. Levitt, M. E. Shenton, G. Gerig, A. Bobick, and A. Tannenbaum, "Statistical shape analysis of brain structures using spherical wavelets," in *Proceedings of the International Symposium of Biomedical Imaging (ISBI)*, 2007.

[3] D. Mamah, L. Wang, D. Barch, G. de Erausquin, M. Gado, and J. Csernansky, "Structural analysis of the basal ganglia in schizophrenia," *Schizophrenia Research*, vol. 89, no. 1–3, pp. 59–71, 2007.

[4] M. Reuter, F.-E. Wolter, and N. Peinecke, "Laplace-spectra as fingerprints for shape matching," in *SPM '05: Proceedings of the 2005 ACM Symposium on Solid and Physical Modeling*. ACM Press, 2005, pp. 101–106.

[5] M. Reuter, *Laplace Spectra for Shape Recognition*. Books on Demand, ISBN 3-8334-5071-1, 2006.

[6] M. Reuter, F.-E. Wolter, and N. Peinecke, "Laplace-Beltrami spectra as "Shape-DNA" of surfaces and solids," *Computer-Aided Design*, vol. 38, no. 4, pp. 342–366, 2006.

[7] N. Peinecke, F.-E. Wolter, and M. Reuter, "Laplace spectra as fingerprints for image recognition," *Computer-Aided Design*, vol. 39, no. 6, pp. 460–476, 2007.

[8] J.-F. Mangin, F. Poupon, E. Duchesnay, D. Riviere, A. Cachia, D. L. Collins, A. C. Evans, and J. Regis, "Brain morphometry using 3D moment invariants," *Medical Image Analysis*, vol. 8, pp. 187–196, 2004.

[9] J. J. Levitt, C.-F. Westin, P. G. Nestor, R. S.-J. Estepar, C. C. Dickey, M. M. Voglmaier, L. J. Seidman, R. Kikinis, F. A. Jolesz, R. W. McCarley, and M. E. Shenton, "Shape of the caudate nucleus and its cognitive correlates in neuroleptic-naive schizotypal personality disorder," *Biological Psychiatry*, vol. 55, pp. 177–184, 2004.

[10] G. Gerig, M. Styner, D. Jones, D. Weinberger, and J. Lieberman, "Shape analysis of brain ventricles using SPHARM," in *Proceedings of the Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA)*, 2001, pp. 171–178.

[11] M. Niethammer, 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 *Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2007.

[12] F.-E. Wolter and K. Friese, "Local and global geometric methods for analysis, interrogation, reconstruction, modification and design of shape," in *Proceedings of CGI'00*, 2000, pp. 137–151.

[13] I. Chavel, *Eigenvalues in Riemannian Geometry*. Academic Press, 1984.

[14] R. Courant and D. Hilbert, *Methods of Mathematical Physics, Vol I*. Interscience, 1953.

[15] C. Gordon, D. Webb, and S. Wolpert, "Isospectral plane domains and surfaces via Riemannian orbifolds," *Inventiones Mathematicae*, vol. 110, pp. 1–22, 1992.

[16] H. McKean and I. Singer, "Curvature and the eigenvalues of the Laplacian," *Journal of Differential Geometry*, vol. 1, pp. 43–69, 1967.

[17] W. Blaschke and K. Leichtweiß, *Elementare Differentialgeometrie*. Springer Verlag, 1973.

[18] D. Braess, *Finite Elemente*. Springer Verlag, 1997.

[19] O. Zienkiewicz and R. Taylor, *The Finite Element Method - Volume 1: The Basis*. Butterworth Heinemann, 2000.

[20] P. Good, *Permutation, Parametric, and Bootstrap Tests of Hypotheses*, 3rd ed. Springer, 2005.

[21] J. M. Boyett and J. J. Shuster, "Nonparametric one-sided tests in multivariate analysis with medical applications," *Journal of the American Statistical Association*, vol. 72, no. 359, pp. 665–668, 1977.

[22] C. R. Genovese, N. A. Lazar, and T. Nichols, "Thresholding of statistical maps in functional neuroimaging using the false discovery rate," *Neuroimage*, vol. 15, pp. 870–878, 2002.

[23] K. Min-Seong, J. Levitt, R. W. McCarley, L. J. Seidman, C. C. Dickey, M. A. Niznikiewicz, M. M. Voglmaier, P. Zamani, K. R. Long, S. S. Kim, and M. E. Shenton, "Reduction of caudate nucleus volumes in neuroleptic-naive female subjects with schizotypal personality disorder," *Biological Psychiatry*, vol. 60, pp. 40–48, 2006.