



Laplace–Beltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis

Martin Reuter^{a,b,*}, Franz-Erich Wolter^c, Martha Shenton^d, Marc Niethammer^{e,f}

^a Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, United States

^b A.A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Harvard Medical School, Boston, United States

^c Inst. für Mensch-Maschine-Kommunikation, Leibniz Universität Hannover, Germany

^d Brigham and Women's Hospital, Harvard Medical School, Boston, United States

^e Department of Computer Science, UNC-Chapel Hill, United States

^f Biomedical Research Imaging Center, School of Medicine, UNC - Chapel Hill, United States

ARTICLE INFO

Article history:

Received 29 October 2008

Accepted 19 February 2009

Keywords:

Laplace–Beltrami spectra

Eigenvalues

Eigenfunctions

Nodal domains

Morse–Smale complex

Reeb graph

Brain structure

Caudate nucleus

Schizotypal personality disorder

ABSTRACT

This paper proposes the use of the *surface-based Laplace–Beltrami* and the *volumetric Laplace eigenvalues and eigenfunctions* as shape descriptors for the comparison and analysis of shapes. These spectral measures are isometry invariant and therefore allow for shape comparisons with minimal shape pre-processing. In particular, no registration, mapping, or remeshing is necessary. The discriminatory power of the 2D surface and 3D solid methods is demonstrated on a population of female caudate nuclei (a subcortical gray matter structure of the brain, involved in memory function, emotion processing, and learning) of normal control subjects and of subjects with schizotypal personality disorder. The behavior and properties of the Laplace–Beltrami eigenvalues and eigenfunctions are discussed extensively for both the Dirichlet and Neumann boundary condition showing advantages of the Neumann vs. the Dirichlet spectra in 3D. Furthermore, topological analyses employing the Morse–Smale complex (on the surfaces) and the Reeb graph (in the solids) are performed on selected eigenfunctions, yielding shape descriptors, that are capable of localizing geometric properties and detecting shape differences by indirectly registering topological features such as critical points, level sets and integral lines of the gradient field across subjects. The use of these topological features of the Laplace–Beltrami eigenfunctions in 2D and 3D for statistical shape analysis is novel.

© 2009 Elsevier Ltd. All rights reserved.

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 [1–3] yield powerful, spatially localized results that are relatively straightforward to interpret, they usually rely on a number of pre-processing 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 pre-processing steps, staying as true as possible to the original data.

This paper describes a methodology for global shape comparison based on the Laplace–Beltrami eigenvalues and for local comparison based on selected eigenfunctions (without the need to register the shapes). The Laplace–Beltrami operator for non-rigid shape analysis of surfaces and solids was first introduced in [4–6] together with a description of the background and up to cubic finite element computations on different representations (triangle meshes, tetrahedra, NURBS patches). In [7,8] the eigenvalues of the (mass density) Laplace operator were used to analyze pixel images. This article focuses on statistical analyses of the Laplace–Beltrami operator on triangulated surfaces and of the *volumetric* Laplace operator on 3D solids and extends earlier works [9,10] by additionally analyzing eigenfunctions and their topological features to localize shape differences. [9] introduces the analysis of eigenvalues of the 2D surface to medical applications. Especially [10] can be seen as a preliminary study to this work, already involving eigenvalues and eigenfunctions for shape analysis. Related work in anatomical shape processing that uses eigenfunctions of the

* Corresponding author at: Martinos Center for Biomedical Imaging, Mass. General Hospital, Harvard Medical, Boston, MA, United States.

E-mail address: reuter@mit.edu (M. Reuter).

URL: <http://reuter.mit.edu> (M. Reuter).

Laplace–Beltrami operator computed via standard linear FEM on triangle meshes includes [11,12] who employ the eigenfunctions as an orthogonal basis for smoothing and the nodal domains of the first eigenfunction for partitioning of brain structures. In [13] a Reeb graph is constructed for the first eigenfunction of a modified Laplace–Beltrami operator on 2D surface representations to be used as a skeletal shape representation. The modified operator gives more weight to points located on the geodesic medial axis (also called cut locus [14]) which originated in computational geometry (see [15,16] for its computation) and has become useful in biomedical imaging. In [17] the Laplace–Beltrami operator is employed for surface parametrization but without computing eigenfunctions or eigenvalues.

Previous approaches for global shape analysis in medical imaging describe the use of invariant moments [18], the shape index [19], and global shape descriptors based on spherical harmonics [20]. The proposed methodology based on the Laplace–Beltrami spectrum differs in the following ways from such approaches.

1. It may be used to analyze *surfaces or solids independently of their isometric embedding* whereas methods based on spherical harmonics or invariant moments are not isometry invariant (finding large shape differences in bendable near-isometric shapes that might only be located differently but otherwise the same, e.g. a person in different body postures). Furthermore, some spherical harmonics-based methods require spherical representations and invariant moments do not easily generalize to arbitrary Riemannian manifolds.

2. Only *minimal pre-processing* of the data is required, in particular no registration is needed. 3D volume data may be represented by its 2D boundary surface, separating the object interior from its exterior or by the 3D volume itself (a volumetric, region-based approach). In the former case, the extraction of a surface approximation from a binary image volume is the only pre-processing step required. In the volumetric case even this pre-processing step can be avoided and computations may be performed directly on the voxels of a given binary segmentation.¹ This is in sharp contrast to other shape comparison methods, requiring additional object registration, remeshing, etc. The presented Laplace–Beltrami eigenvalues and eigenfunctions are invariant to rigid transformations, isometries, and to grid/mesh discretization (as long as the discretization is sufficiently accurate) [6] and fairly robust with respect to noise.

This article summarizes and significantly extends previous Laplace–Beltrami shape analysis work on subcortical brain structures [9,10]. Results are presented both for the 2D surface case (triangle mesh), as well as for 3D solids consisting of non-uniform voxel data. Neumann spectra are used as shape descriptors in 3D, with powerful discrimination properties for coarse geometry discretizations. In addition to the eigenvalues (allowing only global shape comparisons), new eigenfunction analyses are introduced employing the Morse–Smale complex and Reeb graph to shed light on the behavior of the spectra as well as on local shape differences. This can be done by automatically defining local geometric features described by topological features of the eigenfunctions (e.g. critical points, nodal domains, level sets and integral curves of the gradient field). The first eigenfunctions indirectly register these features robustly across shapes, therefore an explicit mesh registration is not necessary. In this paper we are mainly interested in the statistical analysis of populations of shapes. We use a study of differences in

a subcortical structure (the caudate nucleus) as a real world example to demonstrate the applicability of the presented methods. The presented topological study of eigenfunctions is a novel approach for statistical shape analyses.

Section 2 describes the theoretical background of the Laplace–Beltrami operator and the numerical computation of its eigenvalues and eigenfunctions. Normalizations of the spectra, properties of the Neumann spectrum as well as the influence of noise and of the discretization are investigated. Section 3 gives an overview of the used topological structures, namely the Morse–Smale complex and the Reeb graph while Section 4 explains the statistical methods used for the analysis of populations of Laplace–Beltrami spectra. Results for two populations of female caudate shapes are given in Section 5. This section is subdivided into the 2D and 3D analyses. Within each of these subsections, we start with a global analysis on the eigenvalues and continue with local shape measures derived from a selection of eigenfunctions. The paper concludes with a summary and outlook in Section 6.

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. 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 [21]. 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 a generalized oscillating membrane defined on the manifold. We will 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:

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

with $\operatorname{grad} f$ the *gradient* of f and div the *divergence* on the manifold (Chavel [22]). The Laplace–Beltrami operator is a linear differential operator. It can be calculated in local coordinates. Given a local parametrization

$$\psi : \mathbb{R}^n \rightarrow \mathbb{R}^{n+k} \quad (2)$$

of a submanifold M of \mathbb{R}^{n+k} with

$$\begin{aligned} g_{ij} &:= \langle \partial_i \psi, \partial_j \psi \rangle, & G &:= (g_{ij}), \\ W &:= \sqrt{\det G}, & (g^{ij}) &:= G^{-1}, \end{aligned} \quad (3)$$

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

$$\Delta f = \frac{1}{W} \sum_{ij} \partial_i (g^{ij} W \partial_j f). \quad (4)$$

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

$$\Delta f = \frac{\partial^2 f}{(\partial x)^2} + \frac{\partial^2 f}{(\partial y)^2}. \quad (5)$$

¹ Note that of course other pre-processing steps might be necessary to initially obtain the geometric data, such as scanning, manual or automatic segmentation of the image. For the purpose of shape analysis, the shape has to be given in a standard representation, which is usually 3D voxel data or 2D triangular meshes.




Fig. 1. Eigenfunction 30 and 50 of the disk.

The wave equation

$$\Delta u = u_{tt}, \quad (6)$$

may be decomposed into its time dependent and its spatially dependent parts

$$u(x, t) = f(x)a(t). \quad (7)$$

Separating variables in the wave equation yields [23]

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

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

$$\Delta f = -\lambda f. \quad (8)$$

The solutions of this equation represent the spatial part of the solutions of the wave equation (with an infinite number of eigenvalue λ_i and eigenfunction f_i pairs). In the case of M being a planar region, $f(u, v)$ in Eq. (8) can be understood as the natural vibration form (also *eigenfunction*) of a homogeneous membrane with the *eigenvalue* λ . The square roots of the eigenvalues are the resonant or natural frequencies ($\omega_i = \sqrt{\lambda_i}$). If a periodic external driving force is applied at one of these frequencies, an unbounded response will be generated in the medium (important, for example, for the construction of bridges). In this work the material properties are assumed to be uniform. The standard boundary condition of a fixed membrane is the *Dirichlet boundary condition* where $f \equiv 0$ on the boundary of the domain (see Fig. 1 for two eigenfunctions of the disk). In some cases we also apply 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 boundary condition used.

The *spectrum* is defined to be the family of eigenvalues of the Helmholtz equation (Eq. (8)), consisting of a diverging sequence $0 \leq \lambda_1 \leq \lambda_2 \leq \dots \uparrow +\infty$, with each eigenvalue repeated according to its multiplicity and with each associated finite-dimensional eigenspace (represented by the corresponding base of eigenfunctions). In the case of the Neumann boundary condition and for closed surfaces without boundary the first eigenvalue λ_1 is always equal to zero, because in this case the constant functions are solutions of the Helmholtz equation. We then omit the first eigenvalue so that λ_1 will be the first non-zero eigenvalue.

Because of the rather simple Euclidean nature of the voxel representations used later, the more general (Riemannian) definitions given above are not necessarily needed to understand the computation in the 3D voxel case. Nevertheless, the metric terms are helpful when dealing with cuboid voxels (as we do) and of course for analyzing the 2D boundary surfaces of the shapes. Furthermore,




Fig. 2. Objects with same shape index but different spectra.

this approach clarifies that the eigenvalues are indeed isometry 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 this is not true, since non-congruent but isometric surfaces exist.

2.2. Properties

The following paragraphs describe the well-known results on the Laplace–Beltrami operator and its spectrum.

- (i) The spectrum is *isometry 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. (4)), i.e., the intrinsic geometry.
- (ii) Furthermore, *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).
- (iii) Changes of the membrane's shape result in *continuous* changes of its spectrum [23].
- (iv) The spectrum does *not characterize* the shape completely, since some non-isometric manifolds with the same spectrum exist (for example see [24]). 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).
- (v) A substantial amount of *geometrical and topological* information is known to be contained in the spectrum [25] (Dirichlet as well as Neumann). Even though we cannot crop a spectrum without losing information, we showed in [5] that it is possible to extract important information just from the first few Dirichlet eigenvalues (approx. 500).
- (vi) The nodal lines (or nodal surfaces in 3D) are the zero level sets of the eigenfunctions. When the eigenfunctions are ordered by the size of their eigenvalues, then the nodes of the n th eigenfunction divide the domain into maximal n sub-domains, called the nodal domains [23]. Usually the number of nodal domains stays far below n .
- (vii) The spectra have more *discrimination power* than simple measures like surface area, volume or the shape index (the normalized ratio between surface area and volume, $SI = A^3/(36\pi V^2) - 1$) [19]. See Fig. 2 for simple shapes with identical shape index, that can be distinguished by their Laplace–Beltrami spectrum.² Furthermore, as opposed to the spectrum, a moment-based method did not detect significant shape differences in the medical application presented in Section 5. The discrimination power of the spectra can be increased when employing both the spectra of the 2D boundary surface and the 3D solid body (cf. isospectral GWW prisms in [6]).

For more properties see [6,5].

² In fact, Riemannian volume and volume of the boundary are spectrally determined (see also [6] where these values were numerically extracted from the beginning sequence of the spectrum in several 2D and 3D cases).

2.3. Variational formulation

For the numerical computation, the first step is to translate the Helmholtz equation into a *variational formulation*. This is accomplished using Green's formula

$$\int \int \varphi \Delta f d\sigma = - \oint \varphi \frac{\partial f}{\partial n} ds - \int \int \nabla(f, \varphi) d\sigma. \quad (9)$$

(Blaschke [26] p.227) with the Nabra operator defined as

$$\nabla(f, \varphi) := Df G^{-1} (D\varphi)^T = \sum (\partial_i f g^{ij} \partial_j \varphi) \quad (10)$$

with the vector $Df = (\partial_1 f, \partial_2 f, \dots)$. Employing the Dirichlet ($f, \varphi \equiv 0$) or the Neumann ($\frac{\partial f}{\partial n} \equiv 0$) boundary condition (Eq. (9)) simplifies to

$$\int \int \varphi \Delta f d\sigma = - \int \int \nabla(f, \varphi) d\sigma. \quad (11)$$

The Helmholtz equation (8) is multiplied with test functions $\varphi \in C^2$, complying with the boundary condition. By integrating over the area and using (11) one obtains:

$$\begin{aligned} \varphi \Delta f &= -\lambda \varphi f \\ \Leftrightarrow \int \int \varphi \Delta f d\sigma &= -\lambda \int \int \varphi f d\sigma \\ \Leftrightarrow \int \int \nabla(\varphi, f) d\sigma &= \lambda \int \int \varphi f d\sigma \\ \Leftrightarrow \int \int Df G^{-1} (D\varphi)^T d\sigma &= \lambda \int \int \varphi f d\sigma \end{aligned} \quad (12)$$

(with $d\sigma = Wdudv$ being the surface element in the 2D case or the volume element $d\sigma = Wdudvdw$ 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 [27], p.35). This variational formulation is used to obtain a system of equations constructing an approximation of the solution.

2.4. Implementation

To solve the Helmholtz equation on any Riemannian manifold the *Finite Element Method* (FEM) [28] can be employed. We choose a tessellation of the manifold into the so-called elements (e.g., triangles or cuboid voxel). Then linearly independent test functions with up to cubic degree (the form functions F_i) can be defined on the triangles or cuboid voxel elements (explained in the next section). The 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 detailed description of the *discretization* used in FEM that finally leads to the following general eigenvalue problem

$$AU = \lambda BU \quad (13)$$

with the matrices

$$\begin{aligned} A &= (a_{lm}) := \left(\int \int DF_l G^{-1} (DF_m)^T d\sigma \right), \\ B &= (b_{lm}) := \left(\int \int F_l F_m d\sigma \right), \end{aligned} \quad (14)$$

where F_l is a piecewise polynomial form function with value one at node l and zero at all other nodes. Here U is the vector (U_1, \dots, 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 the 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 the 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 [29] 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 [30] together with SuperLU [31] and a shift-invert method, to compute the eigenfunctions and eigenvalues starting from the smallest eigenvalue in increasing eigenvalue order. The sparse solver implemented in Matlab uses a very similar indirect method.

It should be noted that the integrals mentioned above are independent of the mesh (as long as the mesh fulfills some refinement and condition standards). Since the solution of the sparse generalized eigenvalue problem can be done efficiently with external libraries, we will now focus on the construction of the matrices A and B .

2.5. Form functions

In order to compute the entries of the two matrices A and B (Eq. (14)) we need the form functions F_i and their partial derivatives ($\partial_k F_i$) in addition to the metric values from Eq. (3). 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 node i and zero at the others. For linear functions it is sufficient to use only the vertices of the triangle mesh as nodes. In the case of a voxel the values at the 8 vertices are sufficient to define a trilinear function in the inside $c_1 + c_2 u + c_3 v + c_4 w + c_5 uv + c_6 uw + c_7 vw + c_8 uvw$. For higher degree approximations further nodes have to be inserted. When applying a Dirichlet boundary condition with zero values at the boundary, we only need a form function for each node in the interior of the domain. 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. Fig. 3 shows examples of a linear and a quadratic local form function for triangles. It can be seen that the form function has the value 1 at exactly one node and 0 at all the others. Note that in the case of the quadratic form function new nodes were 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 Fig. 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. A full tricubic approach of the Lagrange family needs 64 nodes (32 along the edges, 24 inside the faces, and 8 inside the cuboid) and increases the total degree of freedom tremendously without adding much accuracy to the solution. More details on the construction of these local functions can be found in most FEM books (e.g. Zienkiewicz [28]). For each element the results of the integrals (14) 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 0 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.




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

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 Eq. (3)) 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 even simpler 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 :

$$G = \text{diag}((s_1)^2, (s_2)^2, (s_3)^2) \quad (15)$$

$$W = \sqrt{\det(G)} = (s_1)^2(s_2)^2(s_3)^2 = V \quad (16)$$

$$G^{-1} = \text{diag}\left(\frac{1}{(s_1)^2}, \frac{1}{(s_2)^2}, \frac{1}{(s_3)^2}\right). \quad (17)$$

These values are not only constant for an entire voxel, they are identical for each 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_i as described above:

$$a_{l(i),m(j)} + = V \int \int \int_0^1 \int \left(\sum_{k=1}^3 \frac{\partial_k F_i \partial_k F_j}{(s_k)^2} \right) dudvdw \quad (18)$$

$$b_{l(i),m(j)} + = V \int \int \int_0^1 \int F_i F_j dudvdw.$$

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 sphere, the cuboid, the cylinder, the solid ball). The eigenvalues for the unit 2-sphere for example are $\lambda_i = i(i + 1)$, $i \in \mathbb{N}_0$ with multiplicity $2i + 1$. In general the eigenvalues asymptotically tend to a line with a slope dependent on the surface area of the 2D manifold M

$$\lambda_n \sim \frac{4\pi n}{\text{area}(M)}, \quad \text{as } n \uparrow \infty. \quad (19)$$

Therefore a difference in surface area manifests itself in different slopes of the eigenvalue asymptotes. Fig. 4 shows the behavior of the spectra of a population of spheres and a population of ellipsoids respectively. The sphere population is based on a unit sphere where Gaussian noise is added in the direction normal to the surface of the noise-free sphere. Gaussian noise is added in the same way to the ellipsoid population. Since the two basic shapes (sphere and ellipsoid) differ in surface area, their unnormalized spectra diverge (Fig. 4a), so larger eigenvalues lead to a better discrimination of groups. Surface area normalization greatly improves the spectral alignment (Fig. 4b). Fig. 4c and d show zoom-ins of the spectra for small eigenvalues. Even for the surface area normalized case, the spectra of the two populations clearly differ. Therefore the spectra can be used to pick up the difference in shape in addition to the size differences.

A similar analysis can be done for 3D solids. The eigenvalues for the cuboid (3D solid) 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 :

$$\lambda_n \sim \left(\frac{6\pi^2 n}{\text{vol}(M)} \right)^{\frac{2}{3}}, \quad \text{as } n \uparrow \infty. \quad (20)$$

Fig. 5 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 surface area or volume. Therefore the Laplace–Beltrami spectra should be normalized. Fig. 6 shows the spectra of the volume normalized solids. The zoom-in shows that shape differences are preserved in the spectra after volume normalization.

2.8. Exactness of the spectrum

When using an FEM with p -order form functions, the order of convergence is known. For decreasing mesh size h it is $p + 1$ for eigenfunctions and $2p$ for eigenvalues [32]. This is the reason, why it makes sense to use higher order elements (we use up to cubic) instead of a global mesh refinement.

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 is 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




Fig. 5. Unnormalized exact spectra of cube, cuboid, ball.




Fig. 6. Volume normalized spectra and zoom-in.

Fig. 4. Spectral behavior from top to bottom: (a) unnormalized, (b) Area normalized, (c) unnormalized (zoom), (d) Area normalized (zoom).

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 at the boundary.

In the case of the ball an exact voxel representation is not possible, therefore the numerical results differ more strongly from the analytical ones especially 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 Fig. 7, the value r describes the number of voxels used in the direction of the radius).




Fig. 7. Approximation of the ball.

