Spectral Computations on Nontrivial Line Bundles

Alexander Vais, Benjamin Berger, Franz-Erich Wolter

Welfenlab, Division of Computer Graphics, Leibniz University of Hannover, 30167 Hannover, Germany

Abstract

Computing the spectral decomposition of the Laplace-Beltrami Operator on a manifold M has proven useful for applications such as shape retrieval and geometry processing. The standard operator acts on scalar functions which can be identified with sections of the trivial line bundle $M \times \mathbb{R}$. In this work we propose to extend the discussion to Laplacians on nontrivial real line bundles. These line bundles are in one-to-one correspondence with elements of the first cohomology group of the manifold with $\mathbb{Z}_2$ coefficients. While we focus on the case of two-dimensional closed surfaces, we show that our method also applies to surfaces with boundaries. Denoting by β the rank of the first cohomology group, there are 2^β different line bundles to consider and each of these has a naturally associated Laplacian that possesses a spectral decomposition. Using our new method it is possible for the first time to compute the spectra of these Laplacians by a simple modification of the finite element basis functions used in the standard trivial bundle case. Our method is robust and efficient. We illustrate some properties of the modified spectra and eigenfunctions and indicate possible applications for shape processing. As an example, using our method, we are able to create spectral shape descriptors with increased sensitivity in the eigenvalues with respect to geometric deformations and to compute cycles aligned to object symmetries in a chosen homology class.

Keywords: Spectral geometry processing, vector bundles, computational topology, Laplace operator, finite elements

1. Introduction

Curves, surfaces and solids are commonly used in computer graphics, computer vision and computer aided geometric design, where they serve as basic building blocks or data elements. Several continuous and discrete representations exist and many algorithms have been proposed to operate on these representations for a variety of tasks or to convert between them. It is often useful to view these objects within the framework of differential geometry where they become instances of manifolds, with or without boundary. Loosely speaking, a manifold M is a space that locally looks like the Euclidean space. In this setting it becomes possible to transfer the tools and techniques of multivariate calculus onto M and to develop algorithms that benefit from the rich arsenal of techniques available in this mathematical fundament.

In this paper we deal with the Laplace-Beltrami operator which is the generalization of the Euclidean Laplacian. As in the Euclidean case it can be defined by $\Delta f := -\operatorname{div} \nabla f$, where ∇f is the gradient of the real-valued function f defined on the manifold M and div is the divergence operator. We follow the above sign convention, making Δ a positive definite operator. The Laplace operator is abundant throughout mathematics, physics and engineering. It plays an important role in describing physical phenomena such as heat diffusion and wave propagation. Moreover, in the context of shape analysis, the Laplace-Beltrami operator has proven useful for a variety of applications due to the fact that it captures important geometric and topological information about the shape represented by the manifold in an isometrically invariant way.

Solving the eigenvalue problem $\Delta f = \lambda f$ results in a sequence of eigenvalues $\lambda_1, \lambda_2, \ldots$, called *spectrum* of Δ and a sequence of eigenfunctions $f_1, f_2, \ldots$ corresponding to the eigenvalues. The eigenfunctions have the useful property that they are orthonormal with respect to the L^2 inner product

$$(f, h) = \int_M f(x)h(x)dM$$

and form a basis for the corresponding Hilbert space of functions defined on M. For rectangular or spherical manifolds M we obtain the well-known Fourier bases and spherical harmonics, respectively. From this functional-analytic point of view, knowledge of the eigenvalues and eigenfunctions leads to an improved understanding of the Laplacian itself, since it becomes diagonal in the basis of its eigenfunctions. Intuitively this means that the action of Δ on a function f can be described by the action of an infinite diagonal matrix, the diagonal entries being the eigenvalues, on the infinite vector of coefficients describing f with respect to the eigenfunction basis, see e.g. [1]. The aforementioned decomposition of f provides a natural frequency decomposition for dealing with data on the manifold at different scales.

1.1. Related Work and Motivation

Laplacians in several computational discretizations, see e.g. [2–5], and their spectral decomposition have received much

Email addresses: vais@welfenlab.de (Alexander Vais), bberger@welfenlab.de (Benjamin Berger), few@welfenlab.de (Franz-Erich Wolter)

attention in the geometric processing community.

Among the applications are shape and image retrieval using a prefix of the spectrum as a fingerprint or "Shape-DNA" [6–12], geometric signal processing operations [13–15], surface remeshing and parametrization [16–19], creating descriptors for shape matching [20–26], shape segmentation and registration [27, 28], statistical shape analysis for medical studies [29–31] and symmetry detection[32], just to mention a few. A survey of some applications can be found in [33]. According to our knowledge the earliest research works employing Laplace (Beltrami) spectra as finger prints for classification, retrieval and matching of shapes and images appears to have been done at the Welfenlab in the years 1999-2001, cf. [11], for details see [12]. Theoretical background is provided by a branch of mathematics known as spectral geometry, in which several results have been obtained, relating e.g. asymptotic expansions of fundamental solutions of the heat and wave equation with the geometry of the manifold, see e.g. [34–38].

In this paper we introduce numerical computations for Laplacian operators on nontrivial Euclidean line bundles over two-dimensional surfaces M. To motivate and explain the idea, consider first the one-dimensional simple case where M is the unit circle S^1 or equivalently the interval $[0, 2\pi]$ with end points identified. Any real-valued function f on M can be considered as a 2π-periodic function on $\mathbb{R}$. The Laplace-Beltrami Operator becomes $\Delta = -\frac{d^2}{dx^2}$ and the eigenvalue problem reads:

$$-f''(x) = \lambda f(x), \quad f(0) = f(2\pi), \quad f'(0) = f'(2\pi)\,.$$

The solutions in this case are easily obtained: The eigenvalues are $\lambda_k := k^2$ for any integer k and the corresponding eigenfunctions are linear combinations of $\sin(kx)$ and $\cos(kx)$. At this point, consider the question:

What happens if f is chosen 2π-antiperiodic?

This means the boundary conditions become

$$f(0) = -f(2\pi), \quad f'(0) = -f'(2\pi)\,.$$

It is easy to check that the eigenvalues are $\lambda_k = (k + \frac{1}{2})^2$ and the corresponding eigenfunctions are linear combinations of the functions $\sin((k + \frac{1}{2})x)$ and $\cos((k + \frac{1}{2})x)$ for integers k.

For higher-dimensional manifolds the spectral decomposition is not easily obtained since only very special metrics give rise to closed form solutions via techniques such as separation of variables [39]. Hence, numerical approximation methods, typically based on the finite element or finite difference schemes have to be used.

Our goal is to study the generalization of the question above in the situation where M is a two-dimensional manifold, focusing on a numerical approach based on the finite element method. We will show that our method also covers manifolds with boundaries. Intuitively we have the freedom to choose if a function f on M has a sign-flip discontinuity when crossing a set of specific loops or paths on M. In order to discuss these matters on solid mathematical ground, it is appropriate to interpret the situation within the framework of vector bundles, where f becomes a smooth section of a certain vector bundle over M that incorporates the sign flips in its transition functions, ultimately leading to modifications in the discretized problems under consideration. Similar discontinuities arise for example in the context of modeling singularities occurring in global parametrization [17–19].

We note that analogous situations also commonly occur in physics in the context of describing the wave function of fermionic particles, i.e. half-integer spin particles such as electrons, giving rise to the notion of *spin structure*. The choice of spin structure is essential in defining the spinor bundle, the spinor Laplacian and the Dirac operator which is a square root of the spinor Laplacian. Recently, related concepts have made their way into applications in computer graphics and geometric modeling. In [40] an algorithm is presented that employs a quaternionic discrete Dirac operator for constructing conformal mappings of surfaces. This has applications such as distortion-minimizing texture mapping and curvature-based shape editing. A key step in the aforementioned algorithm is the computation of an eigenfunction of a modification of this discrete operator. The authors show that their discrete Dirac operator is related to a Dirac operator in the continuous setting. For example they show that it is locally equivalent to the standard Dirac operator for a spin $\frac{1}{2}$ particle in the plane.

While objects with simple topology like the Euclidean plane or compact oriented surfaces of genus zero have essentially only one spin structure, it is known that there are several possible spin structures for surfaces with more complicated topology and each one induces a different spinor bundle, different Dirac operators and associated spectra. More precisely the possible spin structures are in correspondence with elements of the first cohomology group $H^1(M, \mathbb{Z}_2)$, see [41].

In our case, we consider real line bundles, which are also in one-to-one correspondence with elements of $H^1(M, \mathbb{Z}_2)$. However, while the fiber of a spinor bundle is some representation of the spin group, typically $\mathbb{C}^2$, the fiber of our bundles is $\mathbb{R}$. This similarity suggests that the ideas presented here are applicable to more general settings, for example in order to compute the spectral decomposition of Dirac operators for different spin structures.

1.2. Outline

In the following section we recall several notions and facts of differential geometry and topology, focusing especially on the concept of a vector bundle. The Laplace-Beltrami operator will be identified with a Laplacian operator associated to the trivial line bundle over M. However, in general the trivial line bundle over M is just one of several different line bundles and for each of them an associated Laplacian exists.

In the third section we describe an algorithm that computes the spectra of Laplacian operators associated naturally to any of these nontrivial vector bundles.

Afterwards we present results obtained using our method and give examples, focusing on the differences that arise from the non-trivial bundles. We also validate our results against known spectra and spectra computed by a different method for a restricted class of shapes, namely flat tori and tori of revolution.

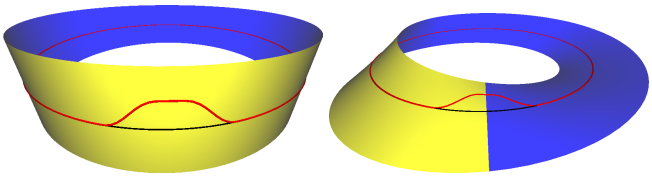

Figure 1: Line bundles over the circle.

Finally we give a conclusion and indicate possible applications and directions for future work.

2. Mathematical Background

We review some of the terminology and notation that we will need, avoiding technical details. For a comprehensive treatment, including formal definitions, we refer to textbooks on differential geometry and topology such as [42–44].

A vector bundle E over a manifold M is obtained by assigning to each point $p \in M$ a vector space E_p in a continuous way. The dimension of the vector spaces is called *rank* of the bundle and vector bundles of rank one are called *line bundles*.

The simplest line bundle is the Cartesian product $M \times \mathbb{R}$ which assigns to each point a copy of the real numbers. This is the so-called *trivial* line bundle. However, not every line bundle is trivial: It is possible to assign to each point a copy of $\mathbb{R}$ such that the resulting bundle is not homeomorphic to $M \times \mathbb{R}$. Figure 1 visualizes two possible line bundles over the circle $M := S^1$. The first bundle is the trivial bundle which is topologically an infinite cylinder, although only a finite portion is shown in the figure. The second bundle in the figure looks locally like a Cartesian product, while globally it is twisted like a Möbius strip of infinite width. Again, only a finite portion is shown.

A map $s : M \to E$ with the property that $s(p) \in E_p$ is called a *section* of E. The space of all smooth sections is denoted by $\Gamma(E)$. Sections of real line bundles are easy to visualize: Locally they can be imagined as the graphs of real-valued functions. The red lines in Fig. 1 show sections of the trivial and non-trivial real line bundles over the circle.

Besides the line bundles, two very important vector bundles associated to a manifold are the tangent bundle TM which is the collection of all tangent spaces, and the cotangent bundle T^*M which is the collection of vector spaces being dual to the tangent spaces. Sections of these bundles are the familiar vector fields and differential one-forms.

A useful concept for computations in vector bundles is the notion of repers: A local reper or frame of E over $U \subset M$ is a collection of sections $e_1, \dots, e_k$ such that for all $p \in U$ the vectors $e_1(p), \dots, e_k(p)$ form a basis of E_p. A reper is called global if this property extends to all $p \in M$.

Every vector bundle admits local repers. A vector bundle admitting a global reper is trivial. Notice that for line bundles, a reper consists of only one non-vanishing section.

Often, vector bundles come equipped with additional structures, such as a fiber metric or a connection.

A *fiber metric* is a scalar product $\langle \cdot, \cdot \rangle : E_p \times E_p \mapsto \mathbb{R}$ depending smoothly on p. Vector bundles with such a fiber metric are called *Euclidean vector bundles*. We will assume throughout this paper that the tangent bundle is Euclidean. In this case, its fiber metric is called a *Riemannian metric* or *metric tensor* which is often denoted by g or by a matrix g_{ij} with respect to a local reper of TM. This allows us to measure lengths, angles and volumes on the manifold M.

A *connection* ∇ on a vector bundle E is a differential operator $\nabla : \Gamma(E) \to \Gamma(E \otimes T^*M)$ that is $\mathbb{R}$-linear and satisfies the Leibniz rule $\nabla(f\phi) = \phi \otimes df + f\nabla\phi$ for any function f and any section ϕ. Here d denotes the exterior derivative, mapping a real-valued function to its differential, which is a real-valued one-form. A connection can be considered as an extension of this concept, mapping E-sections to E-valued one-forms.

The space of sections $\Gamma(E)$ of an Euclidean vector bundle becomes a Hilbert space by introducing the natural inner product

$$(\phi, \psi) := \int_M \langle \phi, \psi \rangle \, dM$$

The Riemannian metric on TM induces a fiber metric on T^*M. The fiber metric of E and the fiber metric of T^*M induce a fiber metric in the bundle $E \otimes T^*M$ and therefore a natural inner product on the space of E-valued one-forms. These inner products are needed in order to define the adjoint to ∇ which is denoted by ∇^*. The adjointness property means that $(\nabla\phi, \alpha) = (\phi, \nabla^*\alpha)$ for all sections ϕ and all E-valued one-forms α. It is essential for deriving the weak variational formulation in the finite element method.

Using the connection and its adjoint, the *Connection Laplacian* associated to ∇ is defined by $\Delta := \nabla^*\nabla$.

The Laplace-Beltrami operator is defined by $\Delta = d^*d$ where d the differential and d^* is its adjoint. It is a special case of a connection Laplacian since d is a connection on the trivial Euclidean line bundle $M \times \mathbb{R}$.

2.1. Classification of Euclidean Line Bundles

Having identified the classical Laplace-Beltrami operator with the connection Laplacian d^*d of the trivial Euclidean line bundle, we can now proceed to consider other line bundles. To characterize these, we will make use of some tools of algebraic topology, see Appendix A for a quick review and [44] for details.

According to the following classification theorem, which we cite from [45], the following objects are equivalent via appropriate isomorphisms:

- An Euclidean line bundle E over M.
- An element ξ of the first cohomology group $H^1(M, \mathbb{Z}_2)$.
- A two-sheeted cover $\rho : \widetilde{M} \to M$.
- A homomorphism $\phi : \pi_1(M) \to \mathbb{Z}_2$.

Any element $\xi \in H^1(M, \mathbb{Z}_2)$ of the cohomology group is completely characterized by its values on a basis. For a compact closed oriented surface of genus g, the rank of the first

(co-)homology group is $2g$. Therefore we choose a basis $\Gamma = \{\gamma_1, \ldots, \gamma_{2g}\}$ of cycles whose equivalence classes generate $H_1(M, \mathbb{Z}_2)$.

Intuitively we can identify a section of an Euclidean line bundle E with a real-valued function on M that does or does not flip sign when crossing a loop γ_i. We will encode this information by a tuple (Γ, χ) where χ is a $\mathbb{Z}_2$-valued vector of dimension $2g$, setting $\chi(i) = 0$ if no sign flip occurs across γ_i and setting $\chi(i) = 1$ if a sign flip occurs. In the latter case, we will call γ_i an *active generator*. We will denote the corresponding Euclidean line bundle by $E_{\Gamma,\chi}$.

More formally, we can construct the bundle as follows: Let $\phi : \pi_1(M) \to \mathbb{Z}_2$ be a homomorphism. Let $\widetilde{M}$ be the universal covering of M. Using ϕ and interpreting the fundamental group as the group of deck transformations, we define an action of $\pi_1(M)$ on the trivial line bundle $\widetilde{M} \times \mathbb{R}$ via:

$$\gamma(\widetilde{q}, v) = \left(\gamma\widetilde{q}, (-1)^{\phi(\gamma)} v\right) \quad \text{for } \gamma \in \pi_1(M), \widetilde{q} \in \widetilde{M}, v \in \mathbb{R}$$

The bundle $E := (\widetilde{M} \times \mathbb{R})/\pi_1(M)$ is an Euclidean line bundle over M. The connection $\nabla = d$ on $M \times \mathbb{R}$ lifts to a connection on $\widetilde{M} \times \mathbb{R}$ and induces a connection on E. This connection in turn induces a connection Laplacian that acts locally just like the Laplace-Beltrami operator.

Note that a smooth section of E can be identified with a smooth real-valued function on a (two-sheeted) covering space of M, provided that it satisfies the chosen symmetry or anti-symmetry conditions. Therefore not every function on the covering space corresponds to a smooth section of the bundle.

3. Description of our Algorithm

Our goal is the following: Given an Euclidean line bundle $E_{\Gamma,\chi}$ over M we want to compute the spectral decomposition of the associated connection Laplacian.

In order to carry out the numerical computation we prescribe sign-flips across the loops in Γ similar to the sign flip that arose in the boundary conditions in the motivating introductory example employing the Möbius strip. The sign-flips are incorporated into the finite element basis functions. The resulting spectrum does not depend on the specific choice of loops, but only on the element $\xi \in H^1(M, \mathbb{Z}_2)$ that the tuple (Γ, χ) represents.

In the first subsection we determine the system of loops Γ, which is a necessary preliminary step for the subsequent steps. Then we describe the classical finite element approach. Finally we introduce the sign-flip modifications necessary in order to account for the nontrivial bundles.

3.1. Computing Homology Generators

We assume that our manifold is furnished with a triangulation (V, E, F) where V is the set of vertices, $E \subset V^2$ is the set of edges and $F \subset V^3$ is the set of faces. The set of vertices together with the edges form a graph (V, E). The dual of this graph is obtained as the graph (F, E^*) where $E^* \subset F^2$ is given by all pairs of faces sharing an edge in E. In the following, we identify E and E^* with each other. The problem of computing cycles on surfaces has been researched before, see e.g. [46–49] and the references therein. For our purposes we apply the algorithm by Erickson and Whittlesey [47]:

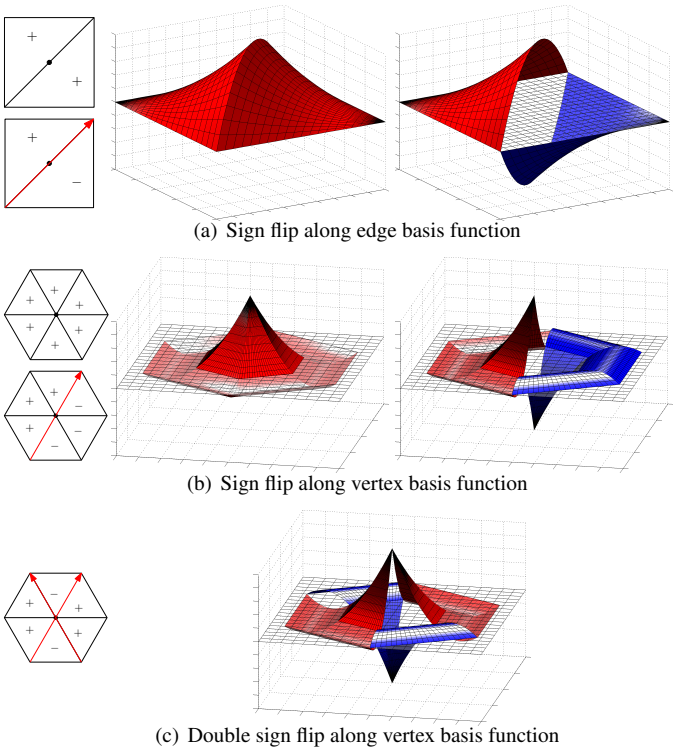

(a) Sign flip along edge basis function

(b) Sign flip along vertex basis function

(c) Double sign flip along vertex basis function

Figure 2: Modified quadratic finite element basis functions.

1. Compute a spanning tree $T \subset E$ of the graph (V, E).
2. Compute a spanning tree $T^* \subset E^*$ of the dual graph (F, E^*) using only edges not occurring in T
3. Compute the set L of all edges not occurring in T or T^*. Each edge $e \in L$ induces a cycle in T.

For a closed manifold of genus g, this algorithm yields a set of $2g$ cycles that generate $H_1(M, \mathbb{Z}_2)$. It can be run with different spanning trees to obtain different generators, though any set of cycles generating the first homology group is sufficient for our later calculations.

3.2. Finite Elements Formulation

The general outline for applying a finite element computation to the the Laplacian eigenvalue problem $d^*df = \lambda f$ is obtained in two steps: First, taking the inner product with an arbitrary test function φ we obtain the equation:

$$(d^*df, \varphi) = (df, d\varphi) = \lambda(f, \varphi) \quad \forall \varphi$$

Here we exploited the adjointness property of the operators d and d^* with respect to the inner products on functions and one-forms. This weak variational formulation is discretized by writing the unknown function f as a linear combination $f = f^1\varphi_1 + \cdots f^N\varphi_N$ of a collection (φ_k) of suitable basis functions and solving the discrete generalized eigenvalue problem

$$Af = \lambda Bf \tag{1}$$

where A and B are $N \times N$ matrices and $f = (f^k)$ is a vector of dimension N. The entries of the matrices are computed by evaluating the inner products

$$A_{ij} := \int_M \langle d\varphi_i, d\varphi_j \rangle dM, \quad B_{ij} = \int_M \varphi_i \varphi_j dM.$$

According to standard finite element constructions, see e.g. [50], the basis functions are constructed by piecing together polynomials over the individual triangles in a triangulation of M to yield functions with local support that satisfy the interpolation conditions $\varphi_i(q_j) = \delta_{ij}$ on a set of N nodes (q_k) spaced regularly at the vertices, on the edges and in the interior of the triangles. This establishes a one-to-one correspondence between every node q_k and the basis function φ_k evaluating to one precisely at that node and to zero at all other nodes. Depending on where q_k is located, the corresponding basis function is called a vertex, edge or bubble function, respectively.

For the classical Laplace-Beltrami eigenvalue problem, the basis functions are continuous. In our method we modify some of the basis functions to have a sign flip across some of the homology generators computed in the previous section. Any basis function whose support is crossed by one or more active generators is affected. Depending on the type of basis function we have three cases to consider:

- Bubble functions are not affected since their support is constrained to one triangle and they vanish on the edges of the triangle.
- An edge basis function has a support consisting of two triangles. Assuming that an active generator γ_k passes through the common edge and denoting by T_R the triangle on the right-hand side of γ_k, we flip the sign of the basis function over T_R.
- A vertex basis function has a support consisting of all triangles incident to a vertex v. Assuming that a single active generator γ_k passes through v, it divides the triangle fan into two sets T_L and T_R corresponding to the triangles to the left and to the right of γ_k. We apply a sign flip to the part of the basis function over T_R. If more than one active generator passes through the vertex of a vertex basis function, the modifications are aggregated.

Note that our computations work well with linear finite elements, in which case there are only vertex basis functions, as well as with higher order elements that require edge and bubble functions. The latter can be employed if high accuracy is desired. Figure. 2 illustrates how modified quadratic edge and basis functions look like. For more details regarding the implementation we refer to Appendix B.

The resulting matrices A and B define a generalized symmetric eigenvalue problem, in which A and B are symmetric and B is positive definite. In our implementation we used the SLEPc library [51] employing the Krylow-Shur algorithm and using the shift-and-invert transform for solving the resulting sparse problem.

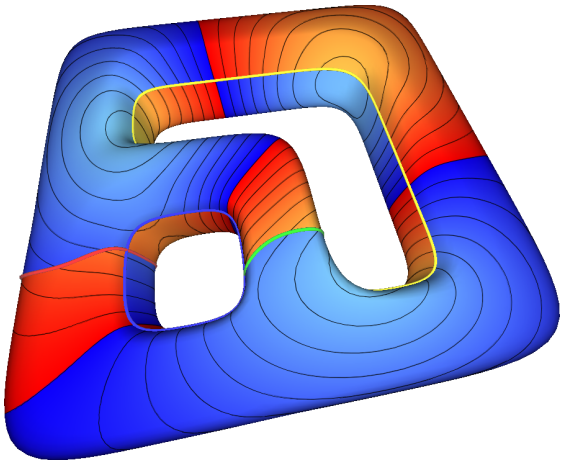

(a) Plot of the 8th eigenfunction on a genus two object. All loops are active.

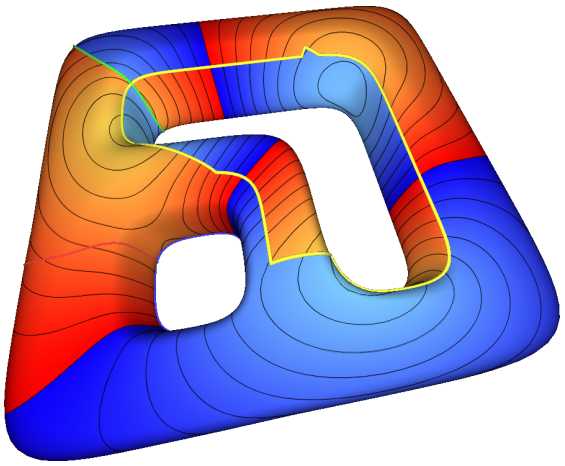

(b) Same computation using a different set of loops. All loops except the red one are active. Notice that the contour lines of absolute value are unaffected.

Figure 3: Example: Two views of the same eigenfunction.

4. Results

4.1. Experimental Results and Analysis

The described algorithm was implemented in C++ and ran on a collection of test models with positive genus. We used the SLEPc library [51] for solving the generalized matrix eigenvalue problem and OpenGL for visualization.

4.1.1. Invariance with respect to the Homology Generators

A typical result of our method is shown in fig. 3(a): The depicted model has genus two. The chosen basis of the first homology group consists of the four cycles colored in red, green, blue and yellow. All loops have been chosen active by setting $\chi = 1111$. The spectral decomposition of the corresponding Laplacian was performed with quadratic finite elements. The mesh is colored by the resulting 8th eigenfunction, clearly exhibiting the sign-flip discontinuities across the loops as the color palette maps positive values to blue colors and negative values to red colors. The picture also shows contour lines of constant *absolute* value. As expected, these lines are smooth,

unaffected by the shape of the loops or the underlying mesh. It is useful to remember that the discontinuous function that we plot on the surface of the object is just a visualization of a *smooth* section of a nontrivial Euclidean line bundle.

We point out, that it is not necessary to choose a nice basis of loops, as any homology basis would suffice. The loops don't have to be short or straight and they may even degenerate in the sense of being locally parallel to other loops or to themselves, as long as the sign flip modifications are correctly applied in the FEM basis functions. The spectrum and the contour lines of absolute value will not be affected by such changes. Any changes are purely due to the way the active loops are positioned in relation to each other with respect to $\mathbb{Z}_2$ homology and there are precisely 2^{2g} different configurations. For example, fig. 3(b) shows a different configuration of loops: The yellow loop has been deformed, the blue loop has been moved down and the green loop switched from one handle to another. This configuration represents a different basis for the homology group. It would have been the same if the green loop had only been deformed. In order to account for the change in the homology basis, the red loop has been deactivated and the spectral decomposition was performed with active yellow, green and blue loops. Despite all these changes, we obtain exactly the same spectrum and the same 8th eigenfunction. Of course we would have obtained a different 8th eigenfunction if we had chosen a different bundle.

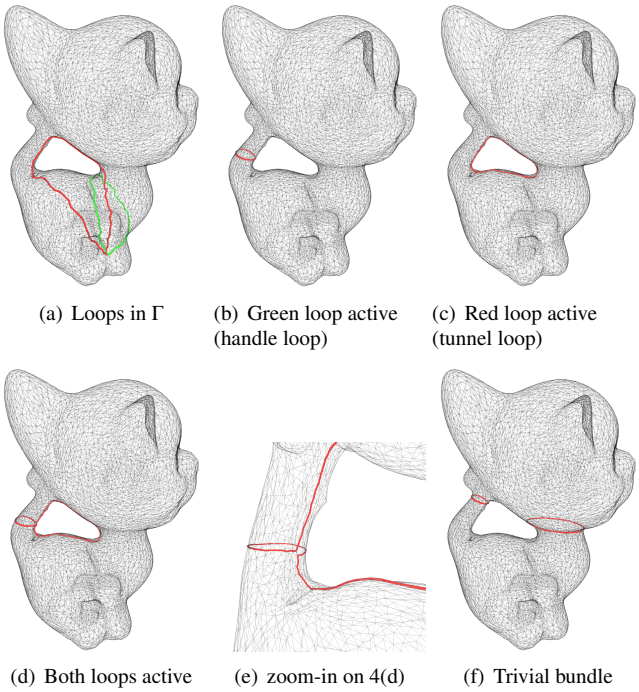

(a) Loops in Γ (b) Green loop active (handle loop) (c) Red loop active (tunnel loop)

(d) Both loops active (e) zoom-in on 4(d) (f) Trivial bundle

Figure 4: Zero sets of the first nontrivial eigenfunction for a genus 1 object.

4.1.2. Zero Sets of Nontrivial Eigenfunctions

The zero sets of the eigenfunctions of the Laplace-Beltrami operator, sometimes also called nodal sets, have been used for shape segmentation since they often identify privileged directions related to the symmetries of the objects or capture surface protrusions that are often well aligned with perceptual features, see [4, 32]. Among all eigenfunctions, the first one corresponding to the lowest non-zero eigenvalue has been proven to be especially useful, since it varies smoothly along the principal direction of elongated objects [52],[53] and in many cases its zero set aligns with a central symmetry of the object. This property has been exploited for example in medical studies to detect the Corpus Callosum, an important structure in brain anatomy, located in the middle between the left and right hemispheres of the white matter surface [54].

Regarding the zero set of a section s on a nontrivial line bundle $E_{\Gamma,\chi}$, we observe that this set, when interpreted as a chain, is closed and homologous to the collection of active loops in Γ. This is seen as follows: Denote by δ the zero-set of s, denote by γ the union of active loops, assuming without loss of generality that all loops have been slightly deformed so they intersect only transversally. Identify s with a function $\widetilde{s} : M \to \mathbb{R}$ that has the sign-flip discontinuities when crossing the set γ and let D be the subset of M where $\widetilde{s}$ is positive. Interpreting γ and δ as 1-cycles and D as a 2-chain, we have that D is bounded by $\gamma + \delta$. Therefore $\gamma = \delta$ with respect to $\mathbb{Z}_2$ homology.

Applying this observation to the eigenfunctions of the nontrivial bundle Laplacians, we can compute *cycles that are geometry aware in the above mentioned sense within any chosen homology class*. The example in fig. 4 shows the zero-sets of the first eigenfunctions on all four bundles associated to a genus one kitten model. The computations were performed using the two homology generators depicted in fig. 4(a). The nontrivial bundle cases are shown in figs. 4(b), 4(c), 4(d), while the classical trivial bundle case is shown in 4(f) for comparison. Note that in general, the zero set need not be connected or can have degeneracies. However the latter typically appears only in perfectly symmetric situations that easily break up to yield paths as shown in fig. 4(e). It is also possible to obtain long cycles that wrap around the surface in complicated ways within their respective homology class.

Concentrating on short cycles, we performed several experiments with the zero sets of the first eigenfunctions and obtained results similar to [49] who compute cycles aligned with the principal curvature direction fields on the surface or [48] who compute short cycles that wrap around the handles and tunnels of the surface. In our approach, we compute the zero sets of all $2^{2g} - 1$ first eigenfunctions on the nontrivial bundles of an object with genus g and sort these by total length. Then we pick out a set of $2g$ cycles that are linearly independent. Figure 5 shows the output of this method for objects of genus one to four. *We note that the resulting loops, while not necessarily the shortest possible, do align well with symmetries of the object.*

The obtained zero sets form cycles that are characteristic for their respective homology class and depend on the intrinsic metric of the surface. Therefore they can be used in the context of shape analysis for objects with non-trivial topology. An example application emphasizing the use of characteristic cycles in medical shape analysis can be found in [55].

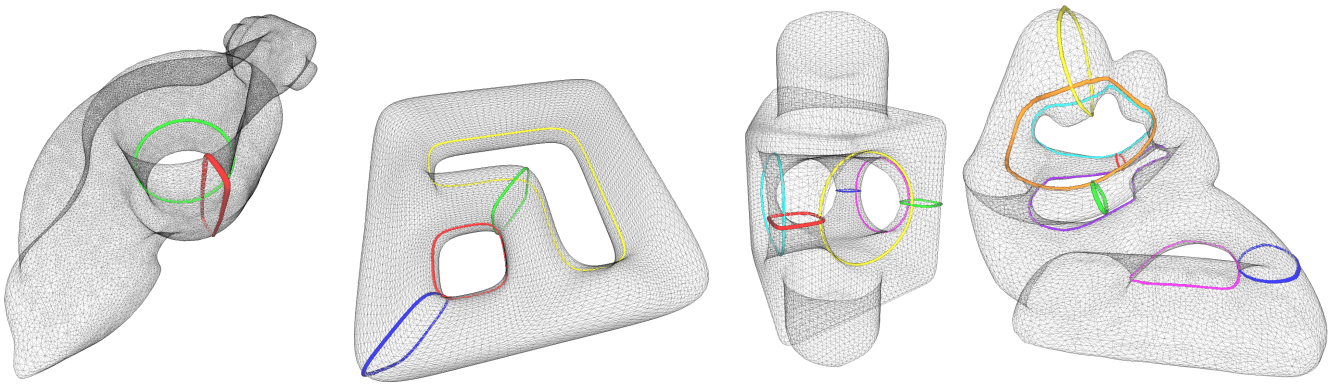

Figure 5: Homology generators calculated from the zero sets of the first eigenfunctions.

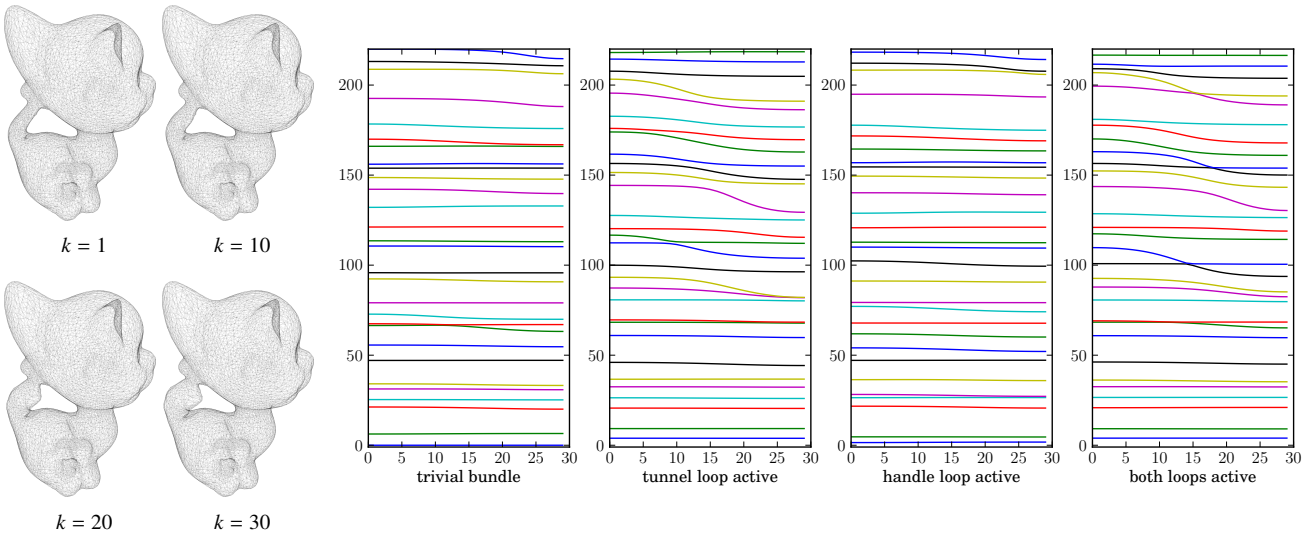

Figure 6: Evolution of all four spectra for a deformation of the kitten model.

4.1.3. Comparison of the Different Spectra

By considering all combinations of loops that are active, we obtain prefixes of all 2^{2g} different spectra for a model with genus g. For the model in fig. 3 we have $g = 2$ and the resulting sixteen spectra are shown in fig. 7 indexed by the vectors χ. Of these spectra, only the bottom row, corresponding to the trivial bundle could be computed before. The other fifteen spectra are new and we will call them nontrivial spectra.

We note that the first eigenvalue is zero if and only if the bundle is trivial. This is easily seen as follows: If E is trivial, then $f := 1$ gives an eigenfunction for $\lambda = 0$. Now, assume E is not trivial, and f is an eigenfunction corresponding to $\lambda = 0$. Because $\Delta f = d^* df = 0$, it follows that $df = 0$ and therefore f is locally constant. If the constant were not zero, then f would be a globally non-vanishing section which is a contradiction to E being nontrivial. So f must be zero everywhere, but this is a contradiction to f being an eigenfunction.

The existence of the nontrivial spectra suggests extensions of concepts such as Shape-DNA[6, 7]. Considering the basic example of the Laplacian on a circle of circumference L, the eigenvalues given by $\lambda_k = \omega_k^2$ for $\omega_k = 2\pi k/L$, $k \in \mathbb{Z}$ indicate a dependence on the circumference L. Intuitively, it is the necessity of an eigenfunction to be periodic on the circle that selects these particular values for ω. As discussed in the introduction, anitperiodicity or equivalently the choice of the nontrivial bundle results in a shift of the eigenvalues. Similar behavior manifests itself when considering surfaces. As an example we created a deformation of the kitten model by thickening its tail over thirty time frames. The left part of fig. 6 shows frame $k \in \{1, 10, 20, 30\}$ of this sequence. For each frame we computed the lowest thirty eigenvalues of the classical Laplace-Beltrami spectrum and the other three nontrivial spectra as defined by activating all combinations of the handle and tunnel loop, see fig. 4(b) and 4(c). The resulting spectral curves are shown in the right part of fig. 6. It is interesting to note that some nontrivial spectra are more susceptible to the deformation. More precisely, the influence of thickening a handle is reflected in the eigenvalues when introducing sign-flips across the corresponding tunnel loop. *This property can be exploited to generate more sensitive fingerprints in the context of Shape-DNA by combining information from multiple spectra. It can also be used to generate shape descriptors that are more sensi-*

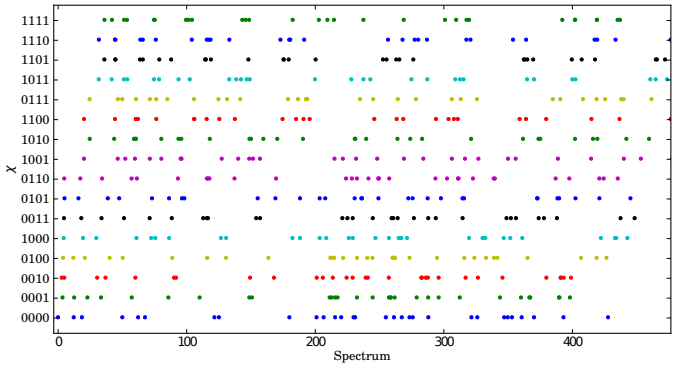

Figure 7: Spectra associated to the different bundles on the object in fig. 3.

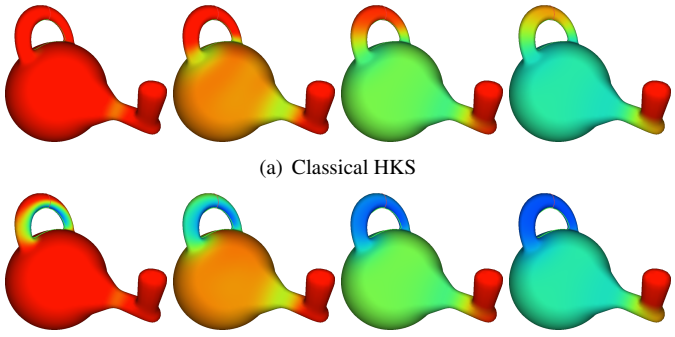

(a) Classical HKS

(b) Nontrivial bundle HKS

Figure 8: Comparison of the heat kernel signature function on the trivial and a nontrivial bundle.

tive to deformation affecting a chosen homology class.

4.1.4. Behavior of the Heat Kernel Signature

Considering the heat diffusion equation $\frac{\partial f}{\partial t} = -\Delta f$ with given initial data $f(t,x)|_{t=0}$, it can be shown that the eigenvalues λ_k and the eigenfunctions f_k of the Laplace-Beltrami operator Δ determine the heat kernel

$$H_t(x,y) := e^{-t\Delta} = \sum_{k=0}^{\infty} e^{-t\lambda_k} f_k(x) f_k(y) \quad t > 0; x, y \in M,$$

also known as the fundamental solution of the heat equation. Knowing the heat kernel allows us to compute how the initial heat distribution evolves with time, see e.g. [1, 38]. Intuitively $H_t(x,y)$ describes the amount of heat obtained after time t at the point y if unit heat was initially concentrated at the point x for $t = 0$. The heat kernel can be used to define the so called heat kernel signature $S(t,x) := H_t(x,x)$, which was introduced in [21] as a multi-scale shape descriptor. The multi-scale nature is owed to the fact that for small t the heat diffusion process is mostly influenced by the local geometry of M while for large t it is influenced by large geometric features and the topology of the manifold. For example $S(t,x)$ becomes constant as $t \to \infty$ since heat eventually distributes uniformly over M.

Considering the heat diffusion equation on nontrivial bundles has interesting implications: Imagine a point p located on a handle of M. After placing a heat source at p at the time $t = 0$, the heat begins to diffuse around the handle. For a time parameter $t > 0$ that is relevant in scale with respect to the diameter of the handle, the heat has enough time to diffuse to the points opposite to p on the surface of the handle. If the non-trivial bundle E is defined by an active tunnel loop γ on the handle, then, loosely speaking, heat cancels itself due to the sign flip across γ and the handle cools down rapidly. An illustration of this phenomenon is shown in fig. 8 which depicts an object colored by the HKS function $S(t,x)$ at several instants of increasing time t, comparing the classical HKS with the HKS derived from the Laplacian on a nontrivial bundle.

4.1.5. Extension to Manifolds with Boundary

In the discussion so far we assumed our manifold M to be closed, i.e. $\partial M = \emptyset$. Our method generalizes easily for manifolds with boundary. Recall that Euclidean line bundles are classified by elements of the first cohomology group $H^1(M, \mathbb{Z}_2)$. The Lefschetz duality theorem (see for example [56]) implies that there is an isomorphism between $H^1(M, \mathbb{Z}_2)$ and the first relative homology group $H_1(M, \partial M, \mathbb{Z}_2)$. In this relative homology, two 1-chains are homologous if their difference is the boundary of a 2-chain plus some 1-chain on ∂M.

In order to compute a set of paths that generate $H_1(M, \partial M, \mathbb{Z}_2)$, we first apply the algorithm of section 3.1, since, following a remark in [49], the algorithm remains correct by considering every boundary loop to enclose a hypothetical face whose geometry need not be determined. If ∂M consists of b boundary loops, we additionally compute $b - 1$ paths that connect different boundary loops. Again we can choose any paths, but for visualization we prefer short connections that are easily obtained with a modified Dijkstra algorithm.

Having obtained $2g + b - 1$ generators of $H_1(M, \partial M, \mathbb{Z}_2)$ we can prescribe sign flips across each of them to define 2^{2g+b-1} different bundles. The Laplacian eigenvalue problem needs to be supplemented with appropriate boundary conditions, such as Dirichlet or Neumann, on each boundary loop.

An example is shown in fig. 9. The object has genus three and three boundary loops, yielding $2 \cdot 3$=6 loops and $3 - 1 = 2$ paths which are all relative cycles that generate the first relative homology group. Choosing all cycles active and applying Dirichlet boundary conditions to the boundary loops we obtain eigenfunctions which typically look as shown in fig. 9(b). Note that the boundary loops constitute contour lines of value zero in this case. For Neumann boundary conditions the contour lines run orthogonal to the boundary loops as shown in fig. 9(c).

4.2. Validation against Known Results

To demonstrate the correctness of our approach, we compare the numerical output of our implementation to well-known results in special cases where the eigenvalue spectra are explicitly computable or at least computable using another approach. We discuss two cases: The flat torus and tori of revolution.

4.2.1. Flat Torus

The flat torus of dimensions a and b is obtained by identifying the opposite sides of the rectangle $[0,a] \times [0,b]$. Applying

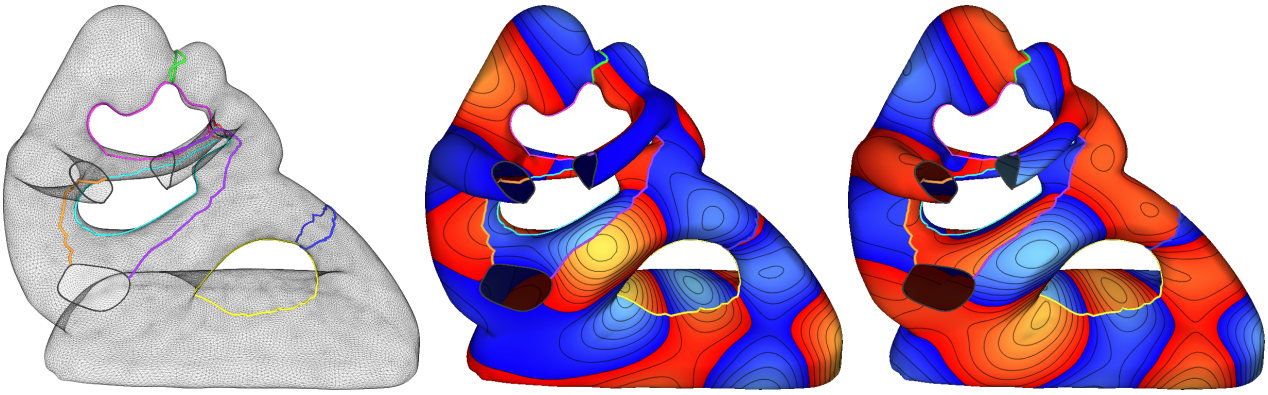

(a) Cut loops and paths (b) An eigenfunction with Dirichlet boundary conditions (c) An eigenfunction with Neumann boundary conditions

Figure 9: Manifold with genus $g = 3$ and $b = 3$ boundaries.

separation of variables the problem can be reduced to two independent instances of the one-dimensional circle case. The eigenvalues of Δ for $\chi = (0, 0)$ are given by the expression

$$\lambda_{m,n} := 4\pi^2 \left(\frac{m^2}{a^2} + \frac{n^2}{b^2} \right)$$

for $m, n \in \mathbb{N}$. The multiplicities are one for $m = n = 0$, two if either $m = 0$ or $n = 0$ but not both and four or higher in case $m > 0, n > 0$.

For $\chi \neq (0, 0)$ the eigenvalue 0 disappears and the integers get shifted by $\frac{1}{2}$ in case of a sign flip. In the case $\chi = (1, 1)$ we obtain

$$\lambda_{m,n} := 4\pi^2 \left(\frac{\left(m + \frac{1}{2}\right)^2}{a^2} + \frac{\left(n + \frac{1}{2}\right)^2}{b^2} \right)$$

We created a flat rectangle mesh with vertices on the borders identified according to the topology of the torus. The two-dimensional FEM computations were carried out using the sign-flipped basis functions. The resulting eigenvalues were found to coincide within numerical tolerances with the eigenvalues obtained using the above equations.

4.2.2. *Surfaces of Revolution*

A surface of revolution generated by a progenitor curve $(x, z) : I \to \mathbb{R}^2$ is defined by the parametrization

$$X(t, \varphi) = x(t) \begin{pmatrix} \cos\varphi \\ \sin\varphi \\ 0 \end{pmatrix} + z(t) \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

We assume the progenitor curve to be parametrized by arc-length and defined on the interval $I = [0, L]$. In this case, the metric tensor associated with the parametrization is

$$(g_{ij}(t, \varphi)) = \begin{pmatrix} 1 & 0 \\ 0 & x^2(t) \end{pmatrix}$$

The metric is therefore captured by the volume element $w(t) := \sqrt{\det g} = x(t)$ which is a function of t alone. From the local description of the Laplace-Beltrami operator we obtain

$$\Delta f = -\partial_t^2 f - \frac{1}{w^2} \partial_\varphi^2 f - \frac{\partial_t w}{w} \partial_t f$$

By separation of variables, the substitution $f(t, \varphi) = \sin(\beta\varphi) u(t)$ leads to the following differential equation for u:

$$-\partial_t \left(w \partial_t u \right) + \frac{\beta^2}{w} u = \lambda w u$$

In order to compute the eigenvalues of a genus one surface of revolution, we need to solve a sequence of Sturm-Liouville problems, one for each possible value of β. The possible values of β and the boundary conditions are given by

$$\beta = k + \frac{\chi(1)}{2}, \; u(0) = (-1)^{\chi(2)} u(L), \; \partial_t u(0) = (-1)^{\chi(2)} \partial_t u(L)$$

for non-negative integers k. We created a torus of revolution and carried out the computations for all four possible bundles using this approach. The Sturm-Liouville problems were solved using a one dimensional finite element method in which the interval $[0, L]$ is subdivided by a sequence of nodes $t_0 = 0 < t_1 < \cdots < t_n = L$. We identify the nodes t_0 and t_1. For the antiperiodic boundary condition $\chi(2) = 1$ the sign flip was easily introduced into the basis function associated to this node.

Similar to the flat case, we created a mesh for the torus of revolution and computed its eigenvalues using our modified two-dimensional finite elements, comparing the results with the eigenvalues obtained from the sequence of one-dimensional finite element computations. Again the results agree within numerical tolerances.

5. Conclusion and Outlook

The proposed method generalizes the numerical computation of the spectral decomposition of the Laplace-Beltrami operator from the trivial line bundle case to other nontrivial line bundles. This is achieved by adapting the finite element basis functions to the topological structure of the bundles using appropriate

sign flips. Our method is applicable to models with nontrivial topology as measured by a non-vanishing first cohomology group and is robust and efficient. We gave experimental evidence and theoretical substantiation of several interesting differences that arise due to the choice of different bundles and indicated how these can be exploited for finding collections of cycles aligned with object symmetries or used to extend classical shape descriptors such as the Shape-DNA or the heat kernel signature.

Based on the previous success of algorithms using differential-geometric notions, including the Laplace-Beltrami operator, we suggest that our method can be considered as a first attempt to make the additional machinery of bundles accessible to numerical computations in the context of shape processing.

Following the path of ideas outlined in this paper, research could proceed in several directions: First, application studies on the behavior of the modified Laplacians should be conducted in order to exploit the additional flexibility offered by them and to explore other ways in which they can serve as useful building blocks in shape processing algorithms.

Second, it is possible to study the extension of the proposed method to three-dimensional manifolds, i.e. solids that are bounded by two-dimensional surfaces.

Third, it is possible to consider other differential operators aside from the Laplacians discussed here. Finally, it is interesting to generalize the computational discussion from real rank one bundles to other vector bundles and differential operators that act on their sections. Consider for example the set of isomorphism classes of complex line bundles over a manifold M. It is known that these bundles are characterized by elements of the second cohomology group $H^2(M, \mathbb{Z})$ with integer coefficients. Therefore we obtain a different vector bundle for each $n \in \mathbb{Z}$, identifying the integer zero with the trivial bundle $M \times \mathbb{C}$ and the Euler characteristic of M with the tangent bundle TM, fitting vector fields nicely into this unifying view.

We reckon that the exploration of these issues from a computational point of view will have several fruitful applications.

Appendix A. Algebraic Topology

Let M be a path-connected topological space. A continuous map $\gamma : [0,1] \to M$ with $\gamma(0) = \gamma(1) = p$ is called a *loop* with base point p. Two loops are called homotopic if one can be gradually deformed into the other. The set of equivalence classes of loops with a fixed base point is called the *first fundamental group* of M, denoted by $\pi_1(M, p)$.

A *covering* of M is a space $\widetilde{M}$ with a surjective local homeomorphism $\rho : \widetilde{M} \to M$. If $\widetilde{M}$ is simply connected, then it is called a universal covering of M. A deck transformation is a homeomorphism $h : \widetilde{M} \to \widetilde{M}$ such that $\rho \circ h = \rho$. Using these notions, another geometric interpretation of the first fundamental group is given by the fact that it is isomorphic to the deck transformation group of the universal cover of M.

To visualize this construction, consider the torus $M = S^1 \times S^1$ as shown in fig. A.10. Its fundamental group is generated by two loops, denoted by a and b. Cutting the torus along these

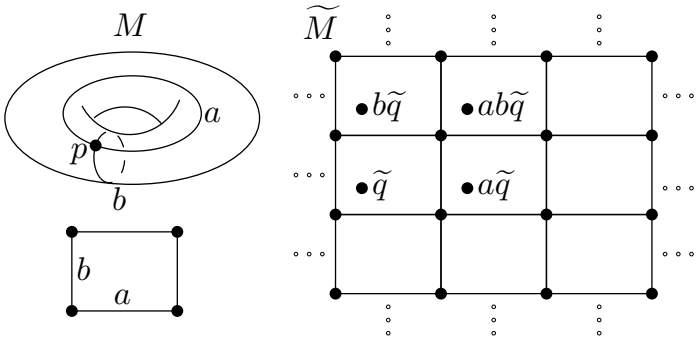

Figure A.10: Fundamental group of the torus acting on its universal covering space by deck transformations.

loops creates a rectangle, the so-called fundamental domain. The universal covering is the Euclidean plane $\widetilde{M} = \mathbb{R}^2$, which is tiled with copies of the fundamental domain. The fundamental group acts on the points $\widetilde{q} \in \widetilde{M}$ by translation, sending $\widetilde{q}$ to the corresponding point in another copy of the fundamental domain.

Now, assume M is a manifold represented by a singular simplicial complex and let R be an arbitrary ring. A k-chain is a formal linear combination of oriented k-simplices with coefficients in the ring R. The set of k-chains form a group C_k under addition. The boundary operator $\partial_k : C_k \to C_{k-1}$ is a linear operator that maps any oriented simplex to the chain consisting of its appropriately signed oriented boundary simplices. A chain $\alpha \in C_k$ is called closed, or cycle, if $\partial\alpha = 0$ and it is called exact if it can be written as $\alpha = \partial\gamma$ for some $\gamma \in C_{k+1}$. Any exact chain is closed as a consequence of the fact that the boundary of a boundary is empty, i.e. $\partial^2 = 0$. Therefore, the sequence of chain groups C_k with the boundary operators in between form a chain complex.

Now let Z_k be the group of closed k-chains and let B_k be the group of exact k-cycles. Two cycles $\alpha, \beta \in Z_k$ are called homologous, if $\alpha - \beta \in B_k$. The k-th homology groups are the quotient groups $H_k(M, R) := Z_k / B_k$ induced by this equivalence relation.

The set of homomorphisms from H_k to R form the k-th cohomology group $H^k(M, R)$. If R is a field, these homology and cohomology groups are finite-dimensional vector spaces and for our purposes it suffices to consider $R = \mathbb{Z}_2$.

Appendix B. Implementation Details

Appendix B.1. Construction of the FEM Basis Functions

The construction of the basis functions is carried out using FEM techniques by piecing together polynomial functions on individual triangles of the finite element mesh. The mesh can be the original triangulation defining the object or a refined version of it for higher accuracy computations. We summarize the typical procedure for the so-called nodal elements and refer to [50] for details.

Given a reference triangle $R \subset \mathbb{R}^2$, spanned by the vertices $(u, v) \in \{(0,0), (1,0), (0,1)\}$, any real-valued polynomial function $\psi(u, v)$ has the form of a finite linear combination

of monomials $u^i v^j$ up to some degree. For the construction of a finite element space, one typically defines a set of polynomials $\psi_1, \ldots, \psi_K$ by prescribing the interpolation conditions $\psi_i(p_j) = \delta_{ij}$ on a set of local nodes $p_1, \ldots, p_K \in R$. These nodes can be chosen in an equidistant manner or they can be chosen as Fekete points for better numerical conditioning.

Every triangle T in a triangulation of M can be parametrized over the reference triangle R by a map $\xi_T : R \to T$. Projecting the set of local nodes on R onto M using ξ_T for all T, eliminating duplicates and enumerating the resulting set of nodes, we obtain a set of global nodes $q_1, \ldots, q_N \in M$.

Figure B.11 shows the situation for linear and quadratic finite elements. Each figure shows the reference element R and the location of the local nodes as well as two triangles $T_1, T_2 \subset M$ that are parametrized over R via the maps ξ_{T_1},ξ_{T_2} and the locations of the global nodes. Linear polynomial basis functions on R, defined by the above mentioned interpolation conditions, are given by the expressions

$$\psi_1 = 1 - u - v \qquad \psi_2 = u \qquad \psi_3 = v,$$

while quadratic polynomial basis functions on R are given by

$$\begin{aligned}
\psi_1 &= 1 - 3u + 2u^2 - 3v + 4uv + 2v^2 \\
\psi_2 &= 4u - 4u^2 - 4uv \\
\psi_3 &= -u + 2u^2 \\
\psi_4 &= 4uv \\
\psi_5 &= -v + 2v^2 \\
\psi_6 &= 4v - 4uv - 4v^2
\end{aligned}$$

Higher order finite element spaces are constructed in the same way for higher accuracy computations, if desired.

In order to construct the global basis functions on M from these polynomials, we denote by $\alpha(T, k)$ the index i such that $q_k = \xi_T(p_i)$ and assign to each global node q_k a global basis function φ_k defined piecewise by

$$\varphi_k|_T := \pm\psi_{\alpha(T,k)} \circ \xi_T^{-1} \tag{B.1}$$

for all triangles T that contain the node q_k and by $\varphi_k|_T := 0$ for all other triangles. The sign choice depends on whether the basis function is flipped on this triangle or not, as described in section 3.2.

Everything discussed so far in the construction of the basis functions boils down to computing the entries of the matrices

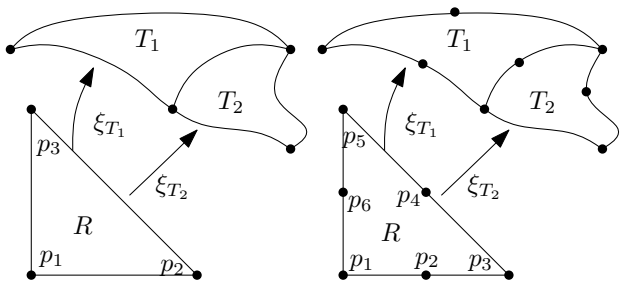

Figure B.11: Node numbering for a linear and for a quadratic element.

A and B of the generalized eigenvalue problem in eq. (1) by evaluating the integrals on the reference domain R. Using the connectivity information given by α, the basis polynomials ψ_i on R and the metric coefficients $g_{ij} := \langle \partial_i \xi_T, \partial_j \xi_T \rangle$ we compute the inverse metric coefficients $(g^{ij}) = (g_{ij})^{-1}$, the volume element $w := \sqrt{\det(g_{ij})}$ and obtain by summing over the relevant triangles:

$$A_{kl} = \sum_T \int_R \sum_{i,j} g^{ij} \cdot \pm\partial_i \psi_{\alpha(T,k)} \cdot \pm\partial_j \psi_{\alpha(T,l)} \cdot w \, dR$$

$$B_{kl} = \sum_T \int_R \pm\psi_{\alpha(T,k)} \cdot \pm\psi_{\alpha(T,l)} \cdot w \, dR$$

The integrals can be computed by numerical integration such as Gauss-Legendre quadrature by computing a linear combination of the integrand evaluated at a set of sample points in R with appropriate weights [57]. If the embedding of the triangulation is planar, the maps ξ_T are affine and their derivatives are constant. In this case the integrals can be evaluated exactly.

Appendix B.2. Determining the Sign Flips

In order to determine the sign flips, it is algorithmically convenient to proceed as follows: we store with each triangle T an array `T.dof` of length K that maps the local node numbers to the global node numbers. More precisely we set `T.dof[i]=k` if $\alpha(T, k) = i$. With each triangle T we also store a bit-valued array `T.flip` of length K that indicates which sign flips occur. For $i = \alpha(T, k)$ the entry `T.flip[i]` determines the sign of $\varphi_k|T$ in eq. (B.1).

Initially, we set all flip bits to zero. Then, for each active loop γ we walk along its edges, looking at every triangle T adjacent to the right of γ, and flip `T.flip[i]` for all local node numbers i that correspond to the global node number of a node lying on γ. Note that the described implementation handles cases with more than one loop passing through a vertex as shown in fig. 2(c). It also handles cases in which more than one active loop passes through the same edge and the effects of the sign flips cancel.

Acknowledgments

The authors would like to thank Dr. L. Habermann and H. Thielhelm for valuable comments. We also thank the AIM@SHAPE Shape Repository for making accessible the 3D models used in this paper.

References

[1] S. Rosenberg, The Laplacian on a Riemannian Manifold: An Introduction to Analysis on Manifolds, Cambridge University Press, 1997.

[2] G. Xu, Discrete Laplace-Beltrami Operators and their Convergence, Computer Aided Geometric Design 21 (8) (2004) 767–784.

[3] M. Wardetzky, S. Mathur, F. Kälberer, E. Grinspun, Discrete Laplace Operators: No Free Lunch, in: Proc. Eurographics Symposium on Geometry Processing, 2007, pp. 33–37.

[4] M. Reuter, S. Biasotti, D. Giorgi, G. Patanè, M. Spagnuolo, Discrete Laplace-Beltrami Operators for Shape Analysis and Segmentation, Computers & Graphics 33 (3) (2009) 381–390.

[5] M. Alexa, M. Wardetzky, Discrete Laplacians on General Polygonal Meshes, ACM Transactions on Graphics 30 (4) (2011) 102.

[6] M. Reuter, F.-E. Wolter, N. Peinecke, Laplace-Spectra as Fingerprints for Shape Matching , in: Proc. ACM Symposium on Solid and Physical Modeling, 2005, pp. 101–106.

[7] M. Reuter, F.-E. Wolter, N. Peinecke, Laplace-Beltrami Spectra as Shape DNA of Surfaces and Solids, Computer-Aided Design 38 (4) (2006) 342–366.

[8] N. Peinecke, F.-E. Wolter, M. Reuter, Laplace-Spectra as Fingerprints for Image Recognition , Computer-Aided Design 39 (6) (2007) 460–476.

[9] N. Peinecke, F.-E. Wolter, Mass Density Laplace-Spectra for Image Recognition, in: Proc. Conference on Cyberworlds, IEEE, 2007, pp. 409–416.

[10] F.-E. Wolter, P. Blanke, H. Thielhelm, A. Vais, Computational Differential Geometry Contributions of the Welfenlab to GRK 615, in: Modelling, Simulation and Software Concepts for Scientific-Technological Problems, Vol. 57 of LNACM, Springer, 2011, pp. 211–236.

[11] F.-E. Wolter, K.-I. Friese, Local and Global Geometric Methods for Analysis Interrogation, Reconstruction, Modification and Design of Shape, in: Proc. Computer Graphics International, IEEE, 2000, pp. 137–151, also available as Welfenlab report No. 3. ISSN 1866-7996.

[12] F.-E. Wolter, T. Howind, T. Altschaffel, M. Reuter, N. Peinecke, Laplace-Spektra - Anwendungen in Gestalt- und Bildkognition, available as Welfenlab Report No. 7. ISSN 1866-7996.

[13] G. Taubin, A Signal Processing Approach to Fair Surface Design, in: Proc. Computer Graphics and Interactive Techniques, ACM, 1995, pp. 351–358.

[14] B. Lévy, Laplace-Beltrami Eigenfunctions: Towards an Algorithm that Understands Geometry, International Conference on Shape Modeling and Applications.

[15] B. Vallet, B. Lévy, Spectral Geometry Processing with Manifold Harmonics, Computer Graphics Forum 27 (2) (2008) 251–260.

[16] S. Dong, P. Bremer, M. Garland, V. Pascucci, J. Hart, Spectral Surface Quadrangulation, ACM Transactions on Graphics 25 (3) (2006) 1057–1066.

[17] Y. Tong, P. Alliez, D. Cohen-Steiner, M. Desbrun, Designing Quadrangulations with Discrete Harmonic Forms, in: Proc. Eurographics Symposium on Geometry Processing, 2006, pp. 201–210.

[18] F. Kälberer, M. Nieser, K. Polthier, QuadCover-Surface Parameterization using Branched Coverings, Computer Graphics Forum 26 (3) (2007) 375–384.

[19] D. Bommes, H. Zimmer, L. Kobbelt, Mixed-Integer Quadrangulation, ACM Transactions on Graphics 28 (3) (2009) 77.

[20] R. Rustamov, Laplace-Beltrami Eigenfunctions for Deformation Invariant Shape Representation, in: A. Belyaev, M. Garland (Eds.), Proc. Eurographics Symposium on Geometry Processing, 2007, pp. 225–233.

[21] J. Sun, M. Ovsjanikov, L. Guibas, A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion, Computer Graphics Forum 28 (5) (2009) 1383–1392.

[22] M. M. Bronstein, I. Kokkinos, Scale-Invariant Heat Kernel Signatures for Non-Rigid Shape Recognition, in: Conference on Computer Vision and Pattern Recognition, IEEE, 2010, pp. 1704–1711.

[23] M. Ruggeri, G. Patan, M. Spagnuolo, D. Saupe, Spectral-Driven Isometry-Invariant Matching of 3D Shapes, Journal of Computer Vision 89 (2010) 248–265.

[24] K. Hildebrandt, C. Schulz, C. von Tycowicz, K. Polthier, Eigenmodes of Surface Energies for Shape Analysis, in: B. Mourrain, S. Schaefer, G. Xu (Eds.), Advances in Geometric Modeling and Processing, Vol. 6130 of Lecture Notes in Computer Science, Springer, 2010, pp. 296–314.

[25] F. Mémoli, Spectral Gromov-Wasserstein Distances for Shape Matching, in: International Conference on Computer Vision, 2009, pp. 256–263.

[26] A. Bronstein, M. Bronstein, R. Kimmel, M. Mahmoudi, G. Sapiro, A Gromov-Hausdorff Framework with Diffusion Geometry for Topologically Robust Non-Rigid Shape Matching, Journal of Computer Vision 89 (2) (2010) 266–286.

[27] M. Reuter, Hierarchical Shape Segmentation and Registration via Topological Features of Laplace-Beltrami Eigenfunctions, Journal of Computer Vision 89 (2010) 287–308.

[28] T. Hou, H. Qin, Robust Dense Registration of Partial Nonrigid Shapes, IEEE Transactions on Visualization and Computer Graphics, preprint (2011).

[29] M. Reuter, F.-E. Wolter, M. Shenton, M. Niethammer, Laplace-Beltrami Eigenvalues and Topological Features of Eigenfunctions for Statistical Shape Analysis, Computer Aided Design 41 (10) (2009) 739 – 755.

[30] M. Niethammer, M. Reuter, F.-E. Wolter, S. Bouix, N. Peinecke, M.-S. Ko, M. E. Shenton, Global Medical Shape Analysis using the Laplace-Beltrami-Spectrum, Conference on Medical Image Computing and Computer Assisted Intervention.

[31] M. Reuter, M. Niethammer, F.-E. Wolter, S. Bouix, M. Shenton, Global Medical Shape Analysis Using the Volumetric Laplace Spectrum, in: Proc. Conference on Cyberworlds, IEEE, 2007, pp. 417–426.

[32] M. Ovsjanikov, J. Sun, L. Guibas, Global Intrinsic Symmetries of Shapes, Computer Graphics Forum 27 (5) (2008) 1341–1348.

[33] H. Zhang, O. Van Kaick, R. Dyer, Spectral Mesh Processing, Computer Graphics Forum 29 (6) (2010) 1865–1894.

[34] S. Minakshisundaram, A. Pleijel, Some Properties of the Eigenfunctions of the Laplace-Operator on Riemannian Manifolds, Canad. J. Math 1 (1949) 242–256.

[35] S. R. S. Varadhan, On the Behavior of the Fundamental Solution of the Heat Equation with Variable Coefficients, Communications on Pure and Applied Mathematics 20 (2) (1967) 431–455.

[36] Y. Colin de Verdière, Spectre du Laplacien et longueurs des géodésiques périodiques I, Compositio Math 27 (1973) 83–106.

[37] J. Duistermaat, On the Spectrum of Positive Elliptic Operators and Periodic Bicharacteristics, in: J. Chazarain (Ed.), Fourier Integral Operators and Partial Differential Equations, Vol. 459 of Lecture Notes in Mathematics, Springer Berlin / Heidelberg, 1975, pp. 15–22.

[38] S. A. Molchanov, Diffusion Processes and Riemannian Geometry, Russian Mathematical Surveys 30 (1975) 1–63.

[39] S. J. Farlow, Partial Differential Equations for Scientists and Engineers, Dover Publishing Inc., 1993.

[40] K. Crane, U. Pinkall, P. Schröder, Spin Transformations of Discrete Surfaces, in: ACM Transactions on Graphics, Vol. 30, ACM, 2011, p. 104.

[41] H. B. Lawson, M.-L. Michelsohn, Spin Geometry, Princeton University Press, 1994.

[42] M. P. do Carmo, Riemannian Geometry, Birkhäuser Boston, 1992.

[43] T. Frankel, The Geometry of Physics: An Introduction, Cambridge University Press, 2011.

[44] A. Hatcher, Algebraic Topology, Cambridge University Press, 2002.

[45] T. A. Murdoch, Twisted-Calibrations and the Cone on the Veronese Surface, Ph.D. thesis, Rice University (1988).

[46] É. C. de Verdière, F. Lazarus, Optimal System of Loops on an Orientable Surface, Discrete and Computational Geometry 33 (3) (2005) 507–534.

[47] J. Erickson, K. Whittlesey, Greedy Optimal Homotopy and Homology Generators, in: Proc. ACM-SIAM Symposium on Discrete Algorithms, SODA '05, SIAM, Philadelphia, PA, USA, 2005, pp. 1038–1046.

[48] T. Dey, K. Li, J. Sun, D. Cohen-Steiner, Computing Geometry-aware Handle and Tunnel Loops in 3D Models, ACM Transactions on Graphics 27 (3) (2008) 1–9.

[49] P. Diaz-Gutierrez, D. Eppstein, M. Gopi, Curvature Aware Fundamental Cycles, Computer Graphics Forum 28 (7) (2009) 2015–2024.

[50] P. Solin, Partial Differential Equations and the Finite Element Method, John Wiley & Sons, 2005.

[51] V. Hernandez, J. E. Roman, V. Vidal, SLEPc: A Scalable and Flexible Toolkit for the Solution of Eigenvalue Problems, ACM Transactions on Mathematical Software 31 (3) (2005) 351–362.

[52] Y. Shi, R. Lai, S. Krishna, N. Sicotte, I. Dinov, A. W. Toga, Anisotropic Laplace-Beltrami Eigenmaps: Bridging Reeb Graphs and Skeletons, Computer Vision and Pattern Recognition Workshop (2008) 1–7.

[53] Y. Shi, J. Morra, P. Thompson, A. Toga, Inverse-Consistent Surface Mapping with Laplace-Beltrami Eigen-Features, in: Information Processing in Medical Imaging, Springer, 2009, pp. 467–478.

[54] R. Lai, Y. Shi, N. Sicotte, A. Toga, Automated Corpus Callosum Extraction via Laplace-Beltrami Nodal Parcellation and Intrinsic Geodesic Curvature Flows on Surfaces, Proc. International Conference on Computer Vision.

[55] S. Xin, Y. He, C. Fu, D. Wang, S. Lin, W. Chu, J. Cheng, X. Gu, L. Lui, Euclidean Geodesic Loops on High-Genus Surfaces Applied to the Morphometry of Vestibular Systems, Medical Image Computing and Computer-Assisted Intervention–MICCAI 2011 (2011) 384–392.

[56] H. Edelsbrunner, J. Harer, Computational Topology, An Introduction, American Mathematical Society, 2010.

[57] S. Wandzura, H. Xiao, Symmetric Quadrature Rules on a Triangle, Computers and Mathematics with Applications 45 (12) (2003) 1829–1840.