



SMI 2013

Laplacians on flat line bundles over 3-manifolds

Alexander Vais*, Daniel Brandes, Hannes Thielhelm, Franz-Erich Wolter

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



ARTICLE INFO

Article history:

Received 18 March 2013

Received in revised form

29 May 2013

Accepted 29 May 2013

Available online 12 June 2013

Keywords:

Spectral geometry

Vector bundles

Computational topology

Laplace operator

Knots

Seifert surfaces

FEM

ABSTRACT

The well-known Laplace–Beltrami operator, established as a basic tool in shape processing, builds on a long history of mathematical investigations that have induced several numerical models for computational purposes. However, the Laplace–Beltrami operator is only one special case of many possible generalizations that have been researched theoretically. Thereby it is natural to supplement some of those extensions with concrete computational frameworks. In this work we study a particularly interesting class of extended Laplacians acting on sections of flat line bundles over compact Riemannian manifolds. Numerical computations for these operators have recently been accomplished on two-dimensional surfaces. Using the notions of line bundles and differential forms, we follow up on that work giving a more general theoretical and computational account of the underlying ideas and their relationships. Building on this we describe how the modified Laplacians and the corresponding computations can be extended to three-dimensional Riemannian manifolds, yielding a method that is able to deal robustly with volumetric objects of intricate shape and topology. We investigate and visualize the two-dimensional zero sets of the first eigenfunctions of the modified Laplacians, yielding an approach for constructing characteristic well-behaving, particularly robust homology generators invariant under isometric deformation. The latter include nicely embedded Seifert surfaces and their non-orientable counterparts for knot complements.

© 2013 Elsevier Ltd. All rights reserved.

1. Introduction and related work

From a physics and engineering perspective the well-known Laplacian acting on functions that are defined over some space M is essential for modeling common phenomena such as heat diffusion and wave propagation on M . In the corresponding mathematical models it often arises from variational methods applied to some energy minimization principle. Its properties make it a versatile tool for obtaining well-behaved functions or studying the underlying space.

The physical relevance and mathematical properties of the Laplacian have motivated several generalizations in various directions, such as the extension from scalar functions to vector or tensor fields. For example the vector Laplacian is relevant in the study of electromagnetics whereas analogous differential operators are used in linear elasticity. Furthermore, by going from Euclidean spaces to curved Riemannian spaces, the Laplace–Beltrami operator acting on functions and the Hodge–de Rham Laplacian acting on differential forms provide natural

generalizations of the Laplacian or vector Laplacian, respectively. These and other more abstract generalizations are studied in a branch of mathematics known as spectral geometry.

Although many fundamental theoretical questions are still unsolved, the field of spectral geometry has established remarkable results that show the Laplacians to capture various geometric and topological information about the underlying space. However, most of these results are not directly amenable to computational methods and are rather given in terms of asymptotic relations or curvature-dependent bounds on the eigenvalues, see e.g. [1]. There is a large gap between the abstract constructions in theory and concrete computational methods applicable to given shapes. In particular, the transition from two to three dimensions is more challenging in practice than indicated by the general theory.

With the increasing availability of computing power, the Laplace operator has attracted considerable interest in computational geometry and shape processing, driven by the desire to exploit it for practical algorithms and based on a variety of discretizations, see e.g. [2–6]. Among applications employing numerically computed Laplacian invariants are shape and image retrieval using spectral prefixes [7–9] based on early research [10,11] and patented in [12] with a retrospective discussed in [13]. The Laplacian has also been successfully used in signal processing operations [14], surface remeshing [15], parametrization [16–18], mesh deformation [19], descriptors for shape matching [20–24],

* Corresponding author. Tel.: +49 511 762 2910.

E-mail addresses: vais@gdv.uni-hannover.de (A. Vais),
dbrandes@gdv.uni-hannover.de (D. Brandes),
thielhel@gdv.uni-hannover.de (H. Thielhelm),
few@gdv.uni-hannover.de (F.-E. Wolter).

segmentation and registration [25] and statistical and topological shape analysis [26–28], just to mention a few. A survey of some of these and other applications can be found in [29].

The route taken in most of these applications is to start directly with a discretization defined in terms of concrete equations that are valid for a point cloud, graph, or mesh representation. Therefore, it is common to discuss the Laplace–Beltrami operator in a specific discretization, most notably as the so-called Cotangent–Laplacian [3]. However, comparatively few computational attempts have gone beyond modeling the classical Laplace–Beltrami operator by considering for example the spectrum of operators derived from different energy functionals [30,31], the Hodge–de Rham Laplacian [32] or quaternionic-valued operators [33].

The contribution of this paper fits into dealing with a class of operators going beyond the usual Laplace–Beltrami operator. We extend the work in [34,35], proposing a general method to explicitly construct a flat line bundle over a compact three-dimensional manifold M represented by a simplicial complex K and to perform a spectral decomposition for the associated connection Laplacian. Note that the general concept of connection Laplacians has recently also been investigated in the context of so-called vector diffusion maps for analyzing point-based data sets, see [36].

The case where M is equipped with non-Euclidean geometry and the trivial connection has been an object of study within physics, see for example [37,38] considering two-dimensional hyperbolic surfaces. Specific constantly curved three-dimensional settings have recently attracted attention, too, see e.g. [39]. Our method applies in these settings as well as in the general arbitrarily curved case.

As we will use knot complements to construct three-dimensional bounded manifolds, some of our result relate to so-called Seifert surfaces [40]. While these are topologically easily constructed, obtaining nice geometrical embeddings is challenging, see [41]. This topic has also been researched in the context of electromagnetic computations to deal with the multi-valuedness of scalar potentials by introducing cuts, see the work of Kotiuga [42,60]. As we will illustrate, our method yields well-behaved embeddings of Seifert surfaces or their non-orientable counterparts.

2. Contribution

Combining ideas from spectral geometry and algebraic topology, the aim of this paper is to investigate the so-called connection Laplacians on flat line bundles from a computational point of view. These operators generalize the well-known Laplace–Beltrami operator which has become ubiquitous in shape processing. One can interpret most of those Laplacians as perturbations of the ordinary Laplacian d^*d by a first-order differential expression, namely

$$\Delta_\omega f = d^* df + 2(df, \omega) + (d^* \omega + |\omega|^2) f \quad (1)$$

where ω is an imaginary-valued closed differential one-form. Employing the notion of a connection, Δ_ω is often called the Bochner or connection Laplacian associated to the flat connection $d_\omega = d + \omega$. One way of understanding connection Laplacians is in terms of introducing certain sign flips or phase shifts across embedded hypersurfaces representing closed chains, i.e. so-called cycles within relative homology. Focusing on two-dimensional manifolds, an approach for obtaining the spectral decomposition of such Laplacians has been recently introduced in [34,35].

We follow up and extend those approaches by describing a general method that is able to deal with three-dimensional volumetric objects of complex topology. While in two dimensions it is comparatively easy to find a suitable 1-cycle resembling a curve and to perform the flips/phase shifts across this curve, the

corresponding situation in three dimensions is more difficult. Obtaining suitable generators in this case requires more sophisticated algorithms which typically produce quite cluttered outputs. These generators can exhibit complex self-intersections or even be non-orientable, thereby obscuring how and where precisely to apply the required sign flips or phase shifts consistently.

In this paper we investigate topologically complex three-dimensional manifolds M by computing the spectral decompositions of the generalized Laplacians. Our approach applies to compact manifolds that may be unbordered and even equipped with a non-Euclidean geometry.

We describe how to overcome the above-mentioned difficulties by constructing complex line bundles over simplicial complexes representing M based on a formal approach. Following the classical definition of a bundle we define an atlas and associated bundle transition functions in terms of a discrete one-form on the dual mesh or — in other words — in terms of a discrete flat connection using the terminology from [43]. We show that the resulting atlas is well-defined in case the one-form is closed. This is necessary to ensure the correctness of the computations building upon this atlas.

As an application we compute smooth well-behaved embeddings of two-dimensional homology generators for any considered homology class. Those are invariant under isometric transformations and robust to noise and discretization.

Outline: In Sections 3–5 we discuss essential mathematical preliminaries in the smooth and discrete settings. Sections 6–8 describe the core of our approach. Section 9 summarizes the algorithm used for obtaining the results discussed in Section 10.

3. Basics

An appropriate mathematical setting for our discussion is provided by differential geometry, see e.g. [44], starting with a given Riemannian manifold M , possibly with boundary. For shape processing this manifold is typically, but not necessarily, embedded in an Euclidean space and can be pictured as a curve, surface, or volume. A differentiable manifold is usually defined in terms of an atlas, being a collection of open sets U_i covering M , together with chart homeomorphisms $U_i \rightarrow \mathbb{R}^n$ that induce differentiable chart transitions. The metric tensor, denoted by g or g_{ij} in local coordinates, allows for measuring metric properties such as lengths, angles and volumes on M . This tensor can be assumed to be given a-priori or to be induced by the embedding.

Commonly, vector bundles are introduced to equip the manifold with additional structure, see e.g. [44,45]. Intuitively, a rank k vector bundle E over a manifold M is obtained by assigning to each point $p \in M$ a k -dimensional vector space E_p in a continuous way. The vector space E_p is called fiber over p . Vector bundles of rank one are called line bundles. The prototypical example of a vector bundle is the tangent bundle TM which is the collection of all tangent spaces of M . Its dual is the cotangent bundle T^*M . Applying k times the exterior product to T^*M one obtains the bundles $\wedge^k T^*M$. A section of a bundle E is a differentiable map $s : M \rightarrow E$ with the property $s(p) \in E_p$ for all p . The space of sections of E is denoted by $\Gamma(E)$. These constructions are quite natural and familiar as for example $\Gamma(TM)$ is the space of vector fields and $\Gamma(\wedge^k T^*M)$, usually denoted by $\Omega^k(M)$, is the space of differential k -forms. The space of complex-valued functions or differential zero-forms on M can also be considered as the space of sections of the trivial line bundle $M \times \mathbb{C}$.

The most important operation on differential forms is the exterior derivative $d : \Omega^k \rightarrow \Omega^{k+1}$. Forms in the kernel of d are called closed, those in the image of d are called exact. Since $d^2 = 0$, the exterior derivative gives rise to the de Rham cohomology groups $H_{dR}^k(M)$ as the quotient groups of closed forms modulo

exact forms. Two forms are said to be co-homologous if their difference is exact. Stokes theorem $\int_M d\alpha = \int_{\partial M} \alpha$ relates the integral of a k -form over a k -dimensional chain to an integral over its $(k-1)$ boundary chain ∂M .

Assuming a manifold M to be represented appropriately by a simplicial complex K , chains are defined to be formal linear combinations of simplices in K and a boundary operator ∂ is defined, satisfying $\partial^2 = 0$. Thus it gives rise to the (simplicial) homology groups $H_k(M, R)$ by calling two chains homologous if their difference forms the boundary of a higher-dimensional chain. Cohomology is introduced in a similar fashion, giving rise to the cohomology groups $H_k(M, R)$. It turns out that these groups are isomorphic to the de Rham cohomology groups, establishing a link between global analysis and topology. For an introduction to homology theory, we refer the interested reader to the vast literature on algebraic topology, and in particular [46,45]. See also [47,48] for presentations especially tailored to geometry processing applications.

The metric tensor induces inner products (\cdot, \cdot) on the spaces of sections of the various bundles mentioned above and therefore the adjoints $d^* : \Omega^k \rightarrow \Omega^{k-1}$, often also called co-differentials, satisfying the important relation $(d\alpha, \beta) = (\alpha, d^*\beta)$ for all α, β . With these preliminaries we have the Laplace–Beltrami operator $\Delta : \Omega^0 \rightarrow \Omega^0$, defined by $\Delta = d^*d$.

4. Connections and holonomy

Generally a connection ∇ on a vector bundle E is a linear differential operator that maps sections of E to E -valued one-forms, i.e. a map $\nabla : \Gamma(E) \rightarrow \Gamma(E \otimes T^*M)$ satisfying the Leibniz rule

$$\nabla(fs) = s \otimes df + f \nabla s, \quad f \in C^\infty(M), \quad s \in \Gamma(E).$$

An intuitive approach to connections can be given in terms of parallel transport and holonomy, see e.g. [45]. Note that we are not necessarily thinking of the tangent bundle and the canonical metric-induced Levi-Civita connection at this point. Rather, we focus on the connection $d_\omega = d + \omega$ for some one-form ω in the trivial line bundle $E = M \times \mathbb{C}$. The parallel transport of an initial vector $f_0 \in E_p$ along a curve $\gamma : [0, 1] \rightarrow M$ starting in $\gamma(0) = p$ is defined by solving the differential equation $(d_\omega f)(\dot{\gamma}) = 0$ or more explicitly

$$\dot{f}(t) + \omega(\dot{\gamma}(t))f(t) = 0$$

for a vector field $f(t)$ along γ satisfying the initial condition $f(0) = f_0$. The vector $f_1 = f(1)$ at $q = \gamma(1)$ is said to be the parallel transport of f_0 along γ . The parallel transported vectors at q are related to the initial vectors at p by a linear map $E_p \rightarrow E_q$. See Fig. 1a for an illustration in the smooth setting, where we have depicted the complex numbers with their a phase and amplitude as arrows with a certain direction and length respectively. In case the one-form ω is restricted to have imaginary values only, the amplitude is invariant under parallel transport.

Note that the concept of a connection can be discretized in the spirit of discrete exterior calculus as discussed for example in [43]: given a simplicial complex K of dimension m and a curve through K , one can imagine the effect of parallel transport to be concentrated at the transitions between adjacent m -simplices. This amounts to viewing the smooth one-form ω as a discrete one-form on the dual mesh $*K$, see Fig. 1b.

If γ is closed, then both f_0 and f_1 are based at the same point p and the linear map mentioned above is actually an endomorphism on E_p . This map is called the holonomy of the connection d_ω around the curve γ based at p . If ω is closed, i.e. if $d\omega = 0$, the connection is called flat. For flat connections, homologous loops induce the same holonomy. In this case the holonomy of d_ω around

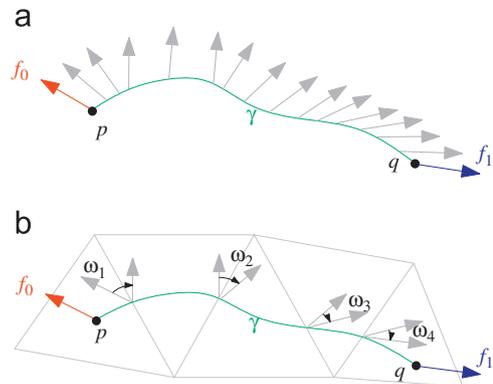


Fig. 1. Illustration of parallel transport. (a) Smooth case. (b) Discrete case.

contractible loops is the identity, while for non-contractible loops this is not necessarily the case. If ω is not only closed but also exact, i.e. if $\omega = d\xi$ for some function ξ , then the holonomy around any loop is the identity. It is precisely the case of flat connections which have trivial local holonomy but non-trivial global holonomy that is the subject of interest we are addressing in this paper.

5. Vector bundles

Assume M to be a given manifold covered by a set of charts $\{U, V, \dots\}$. Let E be a complex line bundle over M . For each chart U the set of all fibers $E_p, p \in U$ is assumed to be homeomorphic to the Cartesian product $U \times \mathbb{C}$ and one speaks of a trivialization of E over U , which we will denote by $E_U = U \times \mathbb{C}$. It is well-known that a bundle on M can be constructed completely by choosing a collection of bundle transition functions, one for each pair of overlapping charts. That is, given any two charts $U, V \subset M$ with $U \cap V \neq \emptyset$ one prescribes a continuous function $\psi(U, V) : U \cap V \rightarrow \mathbb{C}^* = \mathbb{C} \setminus \{0\}$. If $E_U = U \times \mathbb{C}$ and $E_V = V \times \mathbb{C}$ are two trivializations of E and $X \in E_p$ (where $p \in U \cap V$) is some vector represented by the pair $(p, x_U) \in U \times \mathbb{C}$ and $(p, x_V) \in V \times \mathbb{C}$, then x_U and x_V are related by

$$x_U = \psi(U, V)(p)x_V.$$

We denote this relationship by the diagram $V \xrightarrow{\psi(U, V)} U$. The bundle transition functions have to satisfy the so-called co-cycle condition

$$\psi(U, V) \psi(V, W) = \psi(U, W)$$

for any three overlapping charts U, V, W in order to make the construction well-defined. A collection of transition functions $\psi = \{\psi(U, V)\}$ satisfying these conditions is a Čech 1-cocycle. In short we can say that a Čech 1-cocycle completely determines the vector bundle.

Assume now that E is a complex line bundle over M . A connection in E can be given by a collection of differential one-forms ω_U , one for each chart U in an atlas of M , subject to the following transition condition, cf. [45]: if U, V are overlapping charts, then

$$\omega_U = \psi(V, U)^{-1} \omega_V \psi(V, U) + d\psi(U, V) = \omega_V + d\psi(U, V).$$

If all bundle transition functions are constant, then the above condition for the connection one-forms simplifies to $\omega_U = \omega_V$. In this case, setting $\omega_U = 0$ for all U in the atlas of M defines a valid connection in E . This connection is flat, yet it can have non-trivial holonomy that is encoded in the bundle transition functions. To see this, imagine a loop $\gamma(t)$ covered by a sequence of charts U_1, U_2, \dots, U_n and starting and ending in a point p . Beginning with a vector $f \in E_p$ the parallel transport can be computed in terms of

the local trivializations as follows: represent $f(t)$ over U_i by a complex-valued function $f_i(t)$. Since $\omega_{U_i} = 0$ it follows that $f_i(t)$ is in fact constant. A chart transition $U_i \rightarrow U_{i+1}$ induces the transition from f_i to f_{i+1} by $f_{i+1} = \psi(U_{i+1}, U_i)f_i$. After a full trip around γ as we arrive at $U_n = U_1$, the vector f_n differs from f_1 by the complex number given by the product $\prod_i \psi(U_{i+1}, U_i)$ which is in general not $1 \in \mathbb{C}$.

6. Construction of a line bundle

Assume the m -dimensional manifold M to be topologically described by a simplicial complex K . Let K^j , $0 \leq j \leq m$, be the set of j -simplices in K . To each simplex $\sigma \in K^j$ we consider the neighborhood U_σ , consisting of all m -simplices in K that are incident to σ . Fig. 2a shows how these domains can look like for $j = 0, \dots, m$ in two and three dimensions. We can think of each U_σ as a chart of M . The union of all those charts covers M to yield an atlas of M , which in the following will be used to construct a line bundle. For example, Fig. 2b shows part of a two-dimensional simplicial decomposition of a manifold M with two charts U_σ and U_μ corresponding to vertices σ, μ . These two charts overlap in the chart $U_{\sigma\mu}$ where $\sigma\mu$ is the 1-simplex joining σ and μ .

Now assume that ω is a given \mathbb{C}^* -valued closed discrete one-form on the dual complex $*K$. It assigns to each dual edge $*\eta$, η being an $(m-1)$ -simplex in K , a value $\omega(*\eta) \in \mathbb{C}^*$. We can define bundle transition functions as follows: for a fixed j -simplex σ , let τ_0, \dots, τ_n be the m -simplices incident to σ . For two neighboring τ_i, τ_j let η_{ij} be their common $(m-1)$ simplex. We first define the bundle-transition functions between U_σ and U_{τ_i} for each i by imposing the conditions

$$\psi(U_{\tau_i}, U_\sigma) = \omega(*\eta_{ij})\psi(U_{\tau_i}, U_\sigma).$$

In the following we will prove that the above conditions can be fulfilled, since ω is closed by definition. There is one degree of freedom that is fixed by imposing e.g. $\psi(U_{\tau_0}, U_\sigma) = 1$. Thus, we effectively construct locally a discrete dual 0-form as a potential for ω .

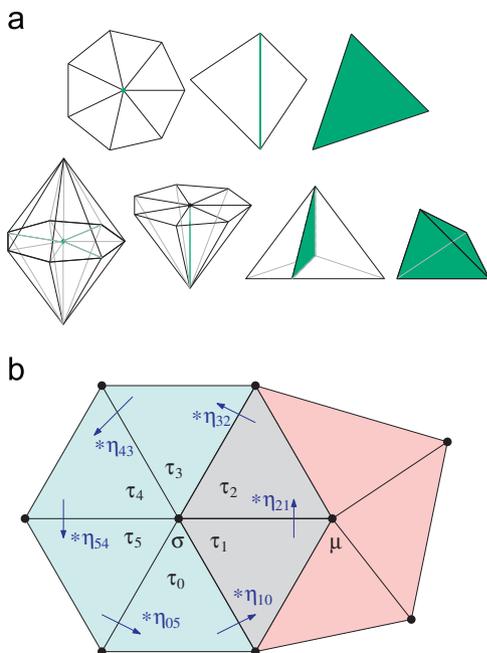
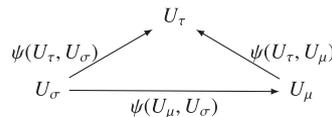


Fig. 2. Constructing a discrete vector bundle. (a) Charts in 2D and 3D. (b) Two overlapping charts.

Proof. In order to see that $\psi(U_\tau, U_\sigma)$ is well-defined, consider the two-dimensional case $m=2$ first. If $\dim \sigma = 2$, then $U_\tau = U_\sigma$ is the only triangle incident to U_σ and there is nothing to check. If $\dim \sigma = 1$, there are at most two triangles incident to σ and at most one condition, which can be obviously satisfied. The remaining case of σ being a vertex is the only non-trivial case. In this case we have a whole triangle fan with σ in its center and only half of a fan if σ is on the boundary. Starting by setting $\psi(U_{\tau_0}, U_\sigma) = 1$ for an initial triangle τ_0 in the fan, we can visit its neighbor τ_1 and determine $\psi(U_{\tau_1}, U_\sigma)$ uniquely. Continuing in the same direction, we visit the neighbor τ_2 of τ_1 and determine $\psi(U_{\tau_2}, U_\sigma)$ uniquely. After a full trip, we arrive back at τ_0 and need to check that no conflict arises for the value of $\psi(U_{\tau_0}, U_\sigma)$. However, this follows directly from ω being closed, i.e. from the vanishing holonomy of ω around σ .

In the three-dimensional setting $m=3$, the cases $\dim \sigma \in \{3, 2, 1\}$ are similar to those discussed above, since we have exactly one tetrahedron incident to itself, at most two tetrahedra incident to a triangle, or at most a cycle of tetrahedra incident to an edge, respectively. The remaining case $\dim \sigma = 0$ amounts to a ball (or half ball) of tetrahedra incident to the vertex σ . These tetrahedra are in one-to-one correspondence with triangles on the surface of the (half-)ball. Every condition between two adjacent tetrahedra can be viewed as a condition between their associated adjacent triangles. Therefore, ω induces a discrete one-form $\tilde{\omega}$ on the dual mesh of the surface triangulation. The problem of finding a valid assignment $\tau \mapsto \psi(U_\tau, U_\sigma)$ amounts to determining a discrete dual 0-form \tilde{f} on the aforementioned triangulated (hemi-)sphere S which satisfies $d\tilde{f} = \tilde{\omega}$. This problem has a solution, since S is simply-connected. Moreover, the solution is unique up to a constant. \square

Having defined the transition functions between charts U_σ and U_τ for j -simplices σ and m -simplices τ , we can proceed to define the transition functions between the charts associated to two arbitrary simplices. Thus, let σ, μ be two simplices such that the overlap of U_σ and U_μ contains the m -simplex U_τ . We obtain the following diagram:



and the cocycle-condition suggests to define

$$\psi(U_\mu, U_\sigma) = \psi(U_\tau, U_\mu)^{-1}\psi(U_\tau, U_\sigma).$$

We have to prove that this is well-defined, i.e. independent of the simplex τ used in the construction.

Proof. Let σ and μ be two simplices in K with $U_\sigma \cap U_\mu \neq \emptyset$. For $U_\tau, \tau \in K$ to be contained in $U_\sigma \cap U_\mu$, we must have $\tau = \sigma \cup \mu \cup \kappa$ for some κ being disjoint from μ and σ (the simplices in the last equation are to be understood in the sense of sets of vertices with \cup being the usual union of sets.). Let now T be the set of all such τ . If $\tau_1, \tau_2 \in T$ are neighboring m -simplices sharing a common $(m-1)$ -simplex η , then we have

$$\begin{aligned} \psi(U_{\tau_1}, U_\mu)^{-1}\psi(U_{\tau_1}, U_\sigma) &= (\omega(*\eta)\psi(U_{\tau_2}, U_\mu))^{-1}\omega(*\eta)\psi(U_{\tau_2}, U_\sigma) \\ &= \psi(U_{\tau_2}, U_\mu)^{-1}\psi(U_{\tau_2}, U_\sigma) \end{aligned}$$

Since any two elements in T can be connected by a chain of neighboring elements of T , the above equality extends to all $\tau_1, \tau_2 \in T$. \square

Thus, starting with a closed discrete \mathbb{C}^* -valued dual one-form ω , we obtain a valid trivialization of a vector bundle. It consists of

the collection of all charts $U_\sigma, \sigma \in K$ covering M together with the specification of all transition functions $\psi(U_\sigma, U_\mu)$ which are constant on $U_\sigma \cap U_\mu$. As discussed at the end of the previous section, the bundle can be equipped with a flat connection having locally trivial holonomy around each simplex. Nevertheless, non-trivial global holonomy can arise.

7. Finite element discretization

Let ∇ be a connection such as d or d_ω . The general outline for applying a finite element computation to the Laplacian eigenvalue problem $\nabla^* \nabla f = \lambda f$ is obtained in two steps: first, taking the L^2 inner product with an arbitrary test function φ one obtains the equation

$$(\nabla f, \nabla \varphi) = \lambda (f, \varphi) \quad \forall \varphi.$$

This weak variational formulation is discretized by writing the unknown function f as a linear combination $f = f^1 \varphi_1 + \dots + f^N \varphi_N$ of a collection (φ_k) of suitable basis functions and solving the discrete generalized eigenvalue problem $Af = \lambda Bf$, 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

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

Using the connection $\nabla = d$, we obtain the classical finite element discretization of the Laplace–Beltrami operator.

We choose the finite element basis functions φ_i to be polynomial basis functions with supports given by the charts U_σ , cf. Section 5. Any such domain decomposes into m -simplices U_τ which are diffeomorphic to an m -dimensional reference simplex R . For concreteness, let R be given in the (u^1, \dots, u^m) -plane with vertices $v_0 = 0$ and $v_i = e_i$ for $1 \leq i \leq m$, where e_i is the i th Cartesian basis vector. This is just the simplex enclosed by the coordinate planes and the plane $u^1 + \dots + u^m = 1$. Then linear scalar basis functions on R are given by

$$\phi_0 = 1 - \sum_i u^i, \quad \phi_1 = u^1, \quad \dots \quad \phi_m = u^m.$$

Similarly higher-order polynomials are specified for high accuracy computations, see e.g. [49]. There are different types of basis functions, depending on the support. The most familiar are the so-called vertex basis functions. They are associated with a support of the form U_σ for a vertex σ . These are the only kind of basis functions arising if one uses linear finite elements. Typically they take the value one at σ and fall off to vanish on the boundary ∂U_σ . In the higher-order case, one additionally considers functions naturally associated to supports U_σ for j -simplices σ with $1 \leq j \leq m$, see Fig. 2a.

In order to evaluate the inner products for the matrices A and B , we need to perform integration. This can be done by summing contributions of individual m -simplices U_τ contained in the intersection of the supports of the two basis functions involved. For example let φ_i and φ_j have supports U_σ and U_μ respectively. Then

$$A_{ij} = \sum_{U_\tau \subset U_\sigma \cap U_\mu} \int_{U_\tau} \langle \psi(U_\tau, U_\sigma) \nabla \varphi_i, \psi(U_\tau, U_\mu) \nabla \varphi_j \rangle dU_\tau,$$

$$B_{ij} = \sum_{U_\tau \subset U_\sigma \cap U_\mu} \int_{U_\tau} \langle \psi(U_\tau, U_\sigma) \varphi_i, \psi(U_\tau, U_\mu) \varphi_j \rangle dU_\tau.$$

Any U_τ can be parametrized over the reference simplex R , giving rise to concrete expressions for the basis functions φ_i, φ_j and their differentials in terms of the polynomials ϕ_k , the volume element dU_τ and the metric tensor needed to evaluate the scalar products. Therefore, all computations are ultimately reduced to a numerical integration problem over R .

8. Determining the possible line bundles

In order to determine how many different flat complex line bundles are available one needs to know the cohomology group $H^1(M, G)$ with coefficients in the Abelian group $G = U(1) \subset \mathbb{C}^*$. If one restricts G to the subgroup $\{\pm 1\}$, then the group $H^1(M, \mathbb{Z}_2)$ counts the real line bundles, see e.g. [50] and the examples in [34].

Within literature there are several algorithms for computing homology groups, see e.g. [51,52], with co-homology being considered more difficult [53]. However $H^1(M, G)$ is not easily obtained in practice, since most currently available algorithms and software packages are not designed for this task. For example the package Chomp [54], which is state of the art within publicly available codes, focuses on computing homology (not co-homology) only using \mathbb{Z} or \mathbb{Z}_n coefficients. It is based on elementary reductions and collapses, see [55] for more background.

Theory suggests to apply the universal coefficient theorem for cohomology [46] to obtain the exact sequence

$$0 \rightarrow \text{Ext}(H_0(M, \mathbb{Z}), G) \rightarrow H^1(M, G) \rightarrow \text{Hom}(H_1(M, \mathbb{Z}), G) \rightarrow 0$$

As $\text{Ext}(H_0(M, \mathbb{Z}), G) = \text{Ext}(\mathbb{Z}, G) = 0$, it follows that $H^1(M, G)$ is isomorphic to $\text{Hom}(H_1(M, \mathbb{Z}), G)$. Furthermore the classification theorem of finitely generated Abelian groups implies that $H_1(M, \mathbb{Z})$ decomposes into $H_1(M, \mathbb{Z}) = \mathbb{Z}^n \oplus \mathbb{Z}_{n_1} \oplus \dots \oplus \mathbb{Z}_{n_r}$, i.e. into a free cyclic subgroup of rank n (where n is commonly known as the first Betti number) and a torsion subgroup in which every element has finite order. To determine $H^1(M, G)$ one has to count the possibilities of mapping each generator of $H_1(M, \mathbb{Z})$ to an element of G while respecting the order relations: while a generator $1 \in \mathbb{Z}$ can be mapped arbitrarily, leading to essentially G possibilities, a generator $1 \in \mathbb{Z}_k$ has to be mapped to an element $\alpha \in G$ with $\alpha^k = 1$, i.e. to one of the k th roots of unity. These considerations give a computable answer to the question of how many possibilities there are to construct a line bundle with a flat unitary connection. For example, if M is an orientable closed surface of genus g then $H_1(M, \mathbb{Z}) = \mathbb{Z}^{2g}$ and the flat bundles can be parametrized by elements of $H^1(M, G) = G^{2g}$.

In order to actually compute a discrete G -valued dual one-form ω consider the following sequence of group homomorphisms:

$$H_{m-1}(K, \partial K, C) \rightarrow H^1(*K, C) \rightarrow H^1(*K, G)$$

where we take C to be \mathbb{Z} or \mathbb{Z}_n for some n . Note that the first group is the $(m-1)$ -st relative homology group that can be computed via [54]. The first map is the isomorphism implied by Poincaré–Lefschetz duality, cf. [46]. This map is available for any Abelian group C if M is orientable. If M is not orientable, the isomorphism still holds for $C = \mathbb{Z}_2$. The second map in the above sequence is induced by a given group homomorphism $C \rightarrow G$.

In the examples discussed later in this paper, we focus on real line bundles, taking G to be $\{\pm 1\}$, $C = \mathbb{Z}$ and the group homomorphism $C \rightarrow G$ to be given by $z \mapsto (-1)^z$.

9. Algorithm summary

Let K be a given simplicial complex describing an m -dimensional manifold. Our algorithm can be summarized as follows:

Stage 1: Compute a set of generators γ_i for $H_{m-1}(K, \partial K, C)$ where we take C to be \mathbb{Z}_2 if the manifold is not orientable. If it is orientable we choose C to be \mathbb{Z} or another finite cyclic group \mathbb{Z}_p . The γ_i can be interpreted as generators of $H^1(*K, C)$.

Stage 2: Use a combination of the γ_i and a group homomorphism $C \rightarrow G$ to create a G -valued discrete dual one-form ω .

- Stage 3: Use ω to construct a line bundle structure by computing the transition maps $\psi = \{\psi(U_\tau, U_\sigma)\}$ as described in Section 6.
- Stage 4: Solve the Laplacian eigenvalue problem on the constructed bundle using the finite element approach described in Section 7.
- Stage 5: Extract information from the spectrum and eigenfunctions. In this paper we have focused on extracting the zero set of the eigenfunction corresponding to the lowest eigenvalue.

10. Computational examples and discussion

We have illuminated the constructive nature of our approach by implementing the complete process in C++ and applying it to various models of different topological complexity. Some of those models have been taken from the AIM@Shape Repository. The tetrahedrizations have been computed using Tetgen [56]. For solving the sparse systems we have relied upon the SLEPc library [57].

10.1. Knot and link complements

To create examples of topologically non-trivial three-manifolds with boundary we construct knot and link complements. A knot is

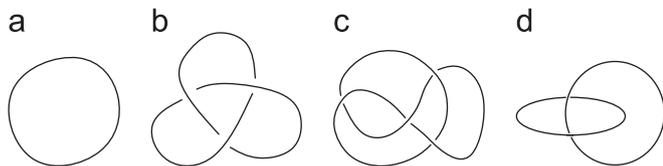


Fig. 3. Representation of knots in terms of knot projections. (a) Unknot, (b) trefoil, (c) figure eight, (d) Hopf link.

an embedding of a circle into space, whereas a link is an embedding of multiple circles. Pictorially knots and links are often represented in terms of a planar projection, such as those shown in Fig. 3. While any knot itself is homeomorphic to the standard circle, the complement of the knot can be topologically quite complicated.

For our purposes we consider knots and links to be thickened by some amount by sweeping a small circle around them, yielding tubular objects O . We construct complements of our tubular knots and links within some larger, yet bounded object H , typically a box. In each example the resulting bounded three-manifold $M = H \setminus O$ is the base domain to which we apply our computations.

An interesting property of knots is that the first fundamental group of the complement can be given in terms of a so-called Wirtinger presentation deduced from a knot-projection [58]. From this presentation it is easy to show that the abelianization of the first fundamental group, which is the first homology group, has rank one. The discussion of Section 8 implies the first cohomology group $H^1(M, \mathbb{Z}_2)$ of our knot complements to be spanned by one generator γ . We apply our algorithm to the bundle constructed from the discrete dual one-form $\omega = \exp(i\pi\gamma) = (-1)^\gamma$, computing the first eigenfunction of the Laplacian on this bundle. Let us denote by S its zero-set. In general S is a regular surface, which we will assume in the following discussion. It can be shown that S represents a class of the relative homology group $H_{m-1}(M, \partial M, \mathbb{Z}_2)$, see [34].

As a first example consider the complement of the unknot within a box. The upper part of Fig. 4 shows the tubular neighborhood of the unknot, which is an ordinary torus, colored in black, while the enclosing box is rendered in a semi-transparent way. The resulting surface S is shown in red. Depending on the relative size of the torus within the box, S can switch from being an inner membrane to an outer one. Note however, that both are homologous modulo the boundary of M and valid representatives of the single non-trivial class in $H_2(M, \partial M, \mathbb{Z}_2)$.



Fig. 4. Top: complement of the unknot within a box. Bottom: two different trefoil complements.

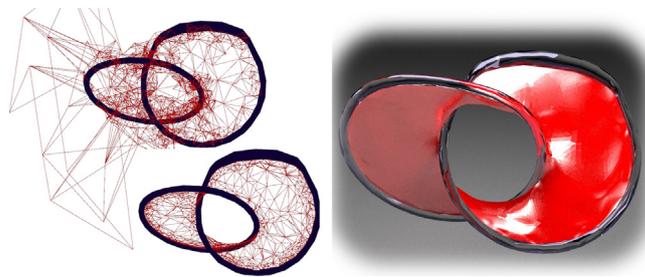


Fig. 5. Complement of a Hopf link within an enclosing ball (not shown).

The lower part of Fig. 4 shows the result for two different variants of the trefoil knot complement. The two cases differ as one trefoil has been elongated in the vertical direction. Although the two resulting zero surfaces are homologous modulo the boundary of M using \mathbb{Z}_2 coefficients, they differ in a topologically remarkable way. The surface shown in the lower right picture is an example of a surface commonly known as a Seifert surface of the trefoil knot. Remarkably Seifert surfaces do exist in general and can be constructed topologically directly from any knot projection using an algorithm given by Seifert [40]. It is however much more difficult to explicitly construct a nice geometrical embedding. Seifert surfaces play an important role in knot theory, giving for example rise to the notion of genus g_K of a knot K , which is the smallest number that is the genus of a Seifert surface for K . In contrast, the surface on the lower left in Fig. 4 is not-orientable. It is topologically a band with three half-turns that is homeomorphic to a Möbius strip.

Note that such non-orientable zero sets could never have been obtained from considering ordinary functions, i.e. sections of the trivial real line bundle. This follows from the fact that ordinary functions yield two-sided zero sets $S \subset M$ and two-sidedness together with the orientability of M imply the orientability of S , cf. [45]. To summarize the situation: if S is orientable it can be thought of as a Seifert surface for the knot. However, our approach can yield non-orientable analogs of Seifert surfaces, too.

Fig. 5 illustrates how our method is applied to the complement of the Hopf link within a larger ball. The cohomology group $H^1(M, \mathbb{Z}_2)$ has rank two and contains four different classes. The sub-figure on the upper left shows a discrete dual one-form, which represents one of those four classes. The first eigenfunction on the line bundle induced by that class yields the surface shown in the lower left and also rendered in the right sub-figure.

Further examples for zero surfaces generated by our method applied to knot and link complements are shown in Fig. 6.

Our method also deals with more commonly considered models such as those shown in Fig. 7. The solid cat model shown on the left has the topology of a solid torus, thus possessing a single non-trivial real line bundle. The zero set S is a surface cutting the tail depicted in the leftmost sub-figure. For the complement of the solid cat model within a larger box one obtains a membrane, similar to the knot complement examples. The chair and statue models, also depicted in Fig. 7, have a more complicated topology. For the chair model we applied our method to the complement within a larger box using a line bundle whose connection has a holonomy of -1 for those loops linking one of the two large holes in the backrest. This results in S consisting of two red membranes appearing as a kind of padding in the backrest. Finally, we also applied our method to the solid statue on the right which appears in a wire-frame rendering. For this complex high-genus model we have chosen one of the several real line bundles and visualized the resulting surface S .

Similarly to the ordinary Laplacian, the low-frequency eigenfunctions of the connection Laplacians are insensitive to noise. This is illustrated in Fig. 8 where we have applied random noise to

a solid model. Note that the resulting cut surface remains stable. The same can be said with respect to varying the resolution of the volumetric discretization.

10.2. Mean curvature of the zero-sets

The zero surfaces computed by our method generally align well to symmetries of M and are reminiscent of minimal surfaces. Of course these are not minimal surfaces in general. In fact, the mean curvature of a level set S for a function f can be expressed as $H = -\text{div } N$, where $N = \nabla f / |\nabla f|$ is the normal vector field along S , see e.g. [44]. Since in our case S is the zero set of the first eigenfunction f of Δ , we have for any $x \in S$ that $\Delta f(x) = \lambda f(x) = 0$, yielding

$$H = \frac{\partial_N |\nabla f|}{|\nabla f|}.$$

The length of ∇f , with f being the first eigenfunction, usually varies slowly in a neighborhood of the surface. Thus, according to the above equation, the mean curvature of S tends to be small.

10.3. Spanning surfaces induced by harmonic one-forms

In this section we compare our work to related considerations concerning level sets of harmonic one-forms. These have been subject of research in the context of electromagnetic modeling, see e.g. [59]. As a simple example, consider a loop of wire in space, possibly knotted, carrying a constant current I . Since the magnetic field H in the space exterior to the wire is known to be irrotational in this situation, i.e. satisfying $\nabla \times H = 0$, one can write locally $H = \nabla \phi$ for some scalar potential ϕ . However, there is no global potential function ϕ satisfying this relation, since for a loop L enclosing the wire we have by Amperè's law

$$\oint_L H \, dl = I.$$

In order to make ϕ a well-defined single-valued function, one is forced to introduce a so-called cut surface into the computation, such that any loop linking the current has to intersect it.

In [42,60], Kotiuga rigorously defined cuts in this context and proposed a method for calculating those using a finite element approach. This approach basically amounts to solving a scalar Poisson equation yielding a harmonic real-valued function f that has a discontinuity of 2π along a previously computed discrete relative homology generator. The differential df is closed (since $d(df) = 0$) and co-closed (since $d^*df = \Delta f = 0$) and therefore a harmonic one-form. The function $g := \exp(if)$ is a continuous well-defined function that maps into the circle S^1 . For almost all $\varphi \in S^1$ the preimage $S_\varphi = g^{-1}(\varphi)$ yields a suitable smooth cut.

We have implemented this method in order to compare the cuts S_φ with our spanning surfaces S . An example showing S_φ for various φ is depicted in Fig. 9, where a solid double torus has been carved out from a cube.

Both methods can be understood in the context of providing well-behaving surfaces contained in the homology classes of the

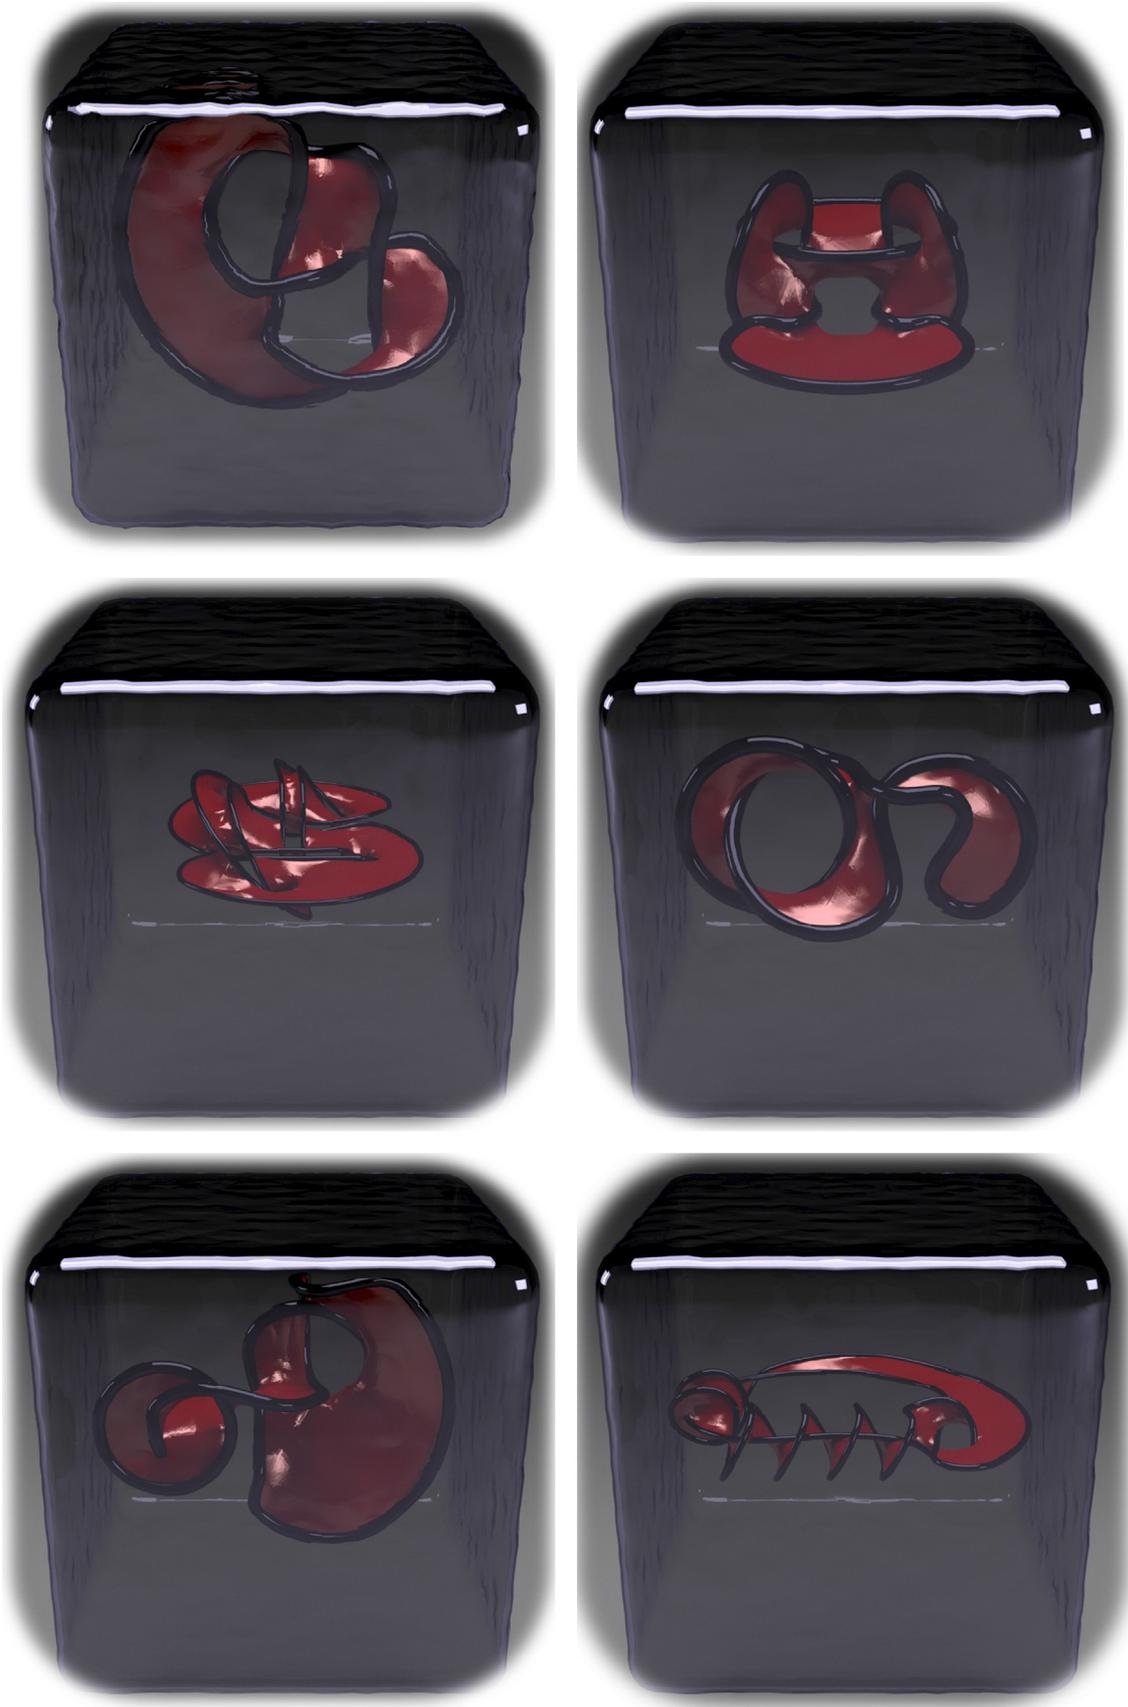


Fig. 6. More examples of knot and link complements.

initial generators required for the computation. While our method yields a unique characteristic representative for the considered homology class based on \mathbb{Z}_2 coefficients, the one-form based approach yields a one-parameter family of orientable representatives for the considered class using \mathbb{Z} coefficients. In this family no unique member is distinguished canonically.

10.4. Complements in closed manifolds

In the context of shape analysis one is interested in characteristic features of objects that are invariant or at least depend as little as possible on issues not pertaining the geometry or topology of the shape. As discussed, the Laplace–Beltrami operator and the generalized Bochner Laplacians are natural tools in this context, yielding robust features invariant under isometries.

In some of our previous examples we have constructed objects M by carving out some object O from a larger object H sitting inside three-dimensional Euclidean space, i.e. $M = H \setminus \phi(O)$, where $\phi : O \rightarrow H \subset \mathbb{R}^3$ is a map embedding O into H . Our investigations in this context should be understood as examining the object M and not O and thereby yielding features characteristic for M .

However, one might be interested in studying a given object O indirectly by applying a spectral analysis to the complement object M . In this case one should care to minimize the effect of the embedding. More generally, consider two different embeddings $\phi_1, \phi_2 : O \rightarrow H$ of an object O into a space H . If $\phi_2 \circ \phi_1^{-1} : \phi_1(O) \rightarrow \phi_2(O)$ can be extended to an isometry of H , then the complements $M_i = H \setminus \phi(O_i)$, are isometric and a spectral analysis for M_1 and M_2 yields the same result.

As an example for the influence of embedding, consider again a tubular unknot O placed within a larger box H at two different positions, cf. Fig. 10. The second is placed closer to the boundary of the box. Computing the zero set of the first eigenfunction as before yields the two red membranes. Note that the membrane for the second embedding is bent a little towards the upper boundary of the box. This is due to the interaction of the outer boundary with O . This example suggests to consider an embedding within a compact three-dimensional manifold H without boundary.

A classical approach to construct such a compact manifold is to choose U to be some larger space, while choosing Γ to be a subgroup of the group $\text{Iso}(U)$ of isometries of U . With certain restrictions on Γ , the quotient space $H = U/\Gamma$ becomes a manifold that inherits the geometry of U . Furthermore U is the universal covering space of H and Γ is isomorphic to the first fundamental group of H . The isometries of H are given by the centralizer of Γ in $\text{Iso}(U)$.

Resuming the previous example, imagine identifying opposite faces of the cube as to obtain the three-torus $H = T^3 = \mathbb{R}^3/\Gamma$, where Γ is the group of translations generating a cubical lattice in \mathbb{R}^3 . This is a compact homogeneous three-manifold with an Euclidean metric. We have extended our method to deal with such manifolds

being described by appropriate gluing rules. Repeating the computation above yields the membrane shown on the right of Fig. 10. The result appears flat, as expected. In fact positioning the unknot anywhere within T^3 yields such a flat membrane as there is no interfering boundary affecting the result. Moreover we know from the above discussion, that the isometries of T^3 consist of all translations of \mathbb{R}^3 modulo the lattice Γ . Therefore, theory predicts that any two embeddings of an object into T^3 differing by an Euclidean motion result in isometric manifolds. This is confirmed

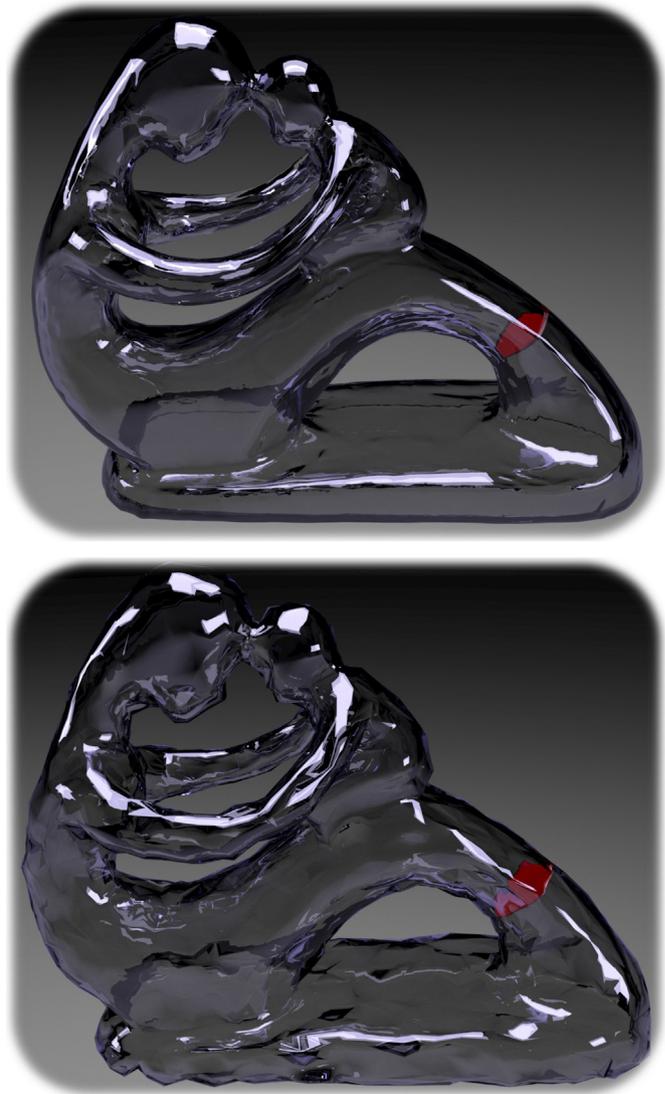


Fig. 8. Influence of noise.



Fig. 7. Examples for characteristic surfaces (shown in red) for different objects. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)



Fig. 9. Examples of iso surfaces induced by harmonic one-forms.

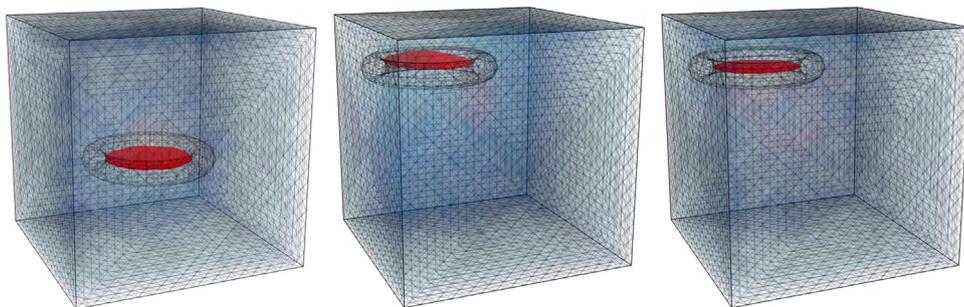


Fig. 10. Complements of tubular unknots. Left and middle: within a cube. Right: within T^3 . (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

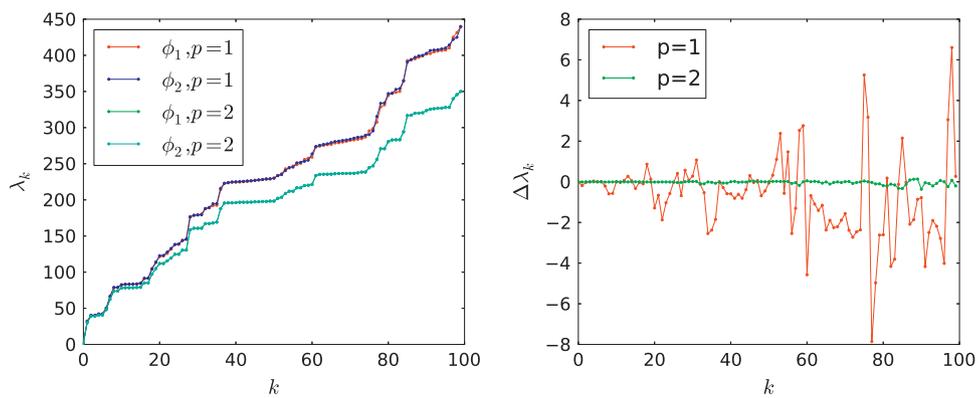


Fig. 11. Comparison of eigenvalue spectra for two different embeddings ϕ_1, ϕ_2 within T^3 using linear and quadratic FEM.

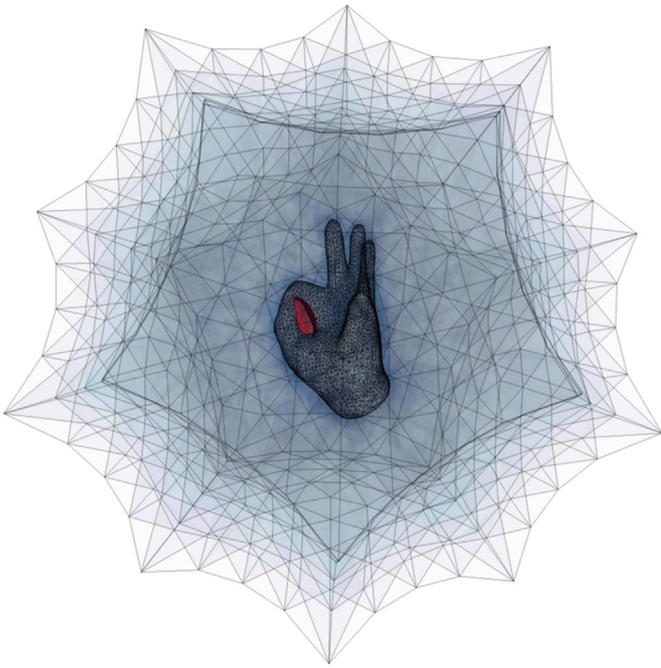


Fig. 12. Model of a human hand carved out from the Seifert–Weber space. Our method yields a spanning surface clamped in between the thumb and forefinger.

in practice as follows. Consider two embeddings (ϕ_1 and ϕ_2) of O within T^3 analogous to the embeddings shown in Fig 10. Fig. 11 compares the resulting eigenvalue spectra. The left part shows a plot of the first 100 eigenvalues for both embeddings using linear ($p=1$) and quadratic ($p=2$) finite elements. Note that the resulting spectra for $p=1$ barely differ, as expected. Yet, a closer look at their difference as shown on the right reveals discrepancies that increase with the eigenvalue index. These disappear for quadratic elements, where the resulting difference curve nearly vanishes.

Besides the three-dimensional torus, other Euclidean compact closed manifolds may be constructed in a similar fashion. There are essentially ten different classes, see [61] for a discussion. However, there are also gluing rules giving rise to manifolds that cannot be equipped with an Euclidean metric. For example gluing opposite faces of a dodecahedron with a $5/10$ twist yields the so-called Seifert–Weber space having a hyperbolic metric. Using instead a $3/10$ twist yields the Poincaré Dodecahedral space (PDS) having an elliptic metric. The SWS has homology group $H_1(M, \mathbb{Z}) = \mathbb{Z}_5^2$ while the PDS is a so-called homology sphere, which means that it has the same homology groups as the sphere S^3 , in particular $H_1=0$. These constructions have been introduced in [62]. More recently they have attracted interest of cosmologists seeking to determine the shape of our universe by studying the cosmic microwave background radiation. See [63] for an introduction into this topic.

Without going into details, our method can deal with such curved background geometries. Fig. 12 shows an example involving the complement of an object within the Seifert–Weber space which is depicted within the Poincaré model of hyperbolic space.

11. Conclusion

Based on the ideas introduced in [34,35] we extended the theory and algorithms to construct a line bundle structure for a manifold represented by a simplicial complex. This structure is explicitly described in terms of charts that are naturally used as supports for finite element basis functions and corresponding

bundle transition functions. This allows us to solve partial differential equations on flat line bundles.

Our approach allows to compute the spectral decomposition of Laplacians acting on sections of flat line bundles over three-dimensional manifolds of complex shape and topology. These include compact manifolds with boundary equipped with an arbitrary Riemannian metric. We have computed smooth well-behaved isometry-invariant homology generators, that are robust to noise and mesh discretization. Our examples show spanning surfaces for various geometries including Seifert surfaces and their non-orientable counterparts for knot complements.

Focusing on the latter application, the considered examples have involved flat real line bundles, whereas the presented approach easily allows for complex flat line bundles, too. Those allow for a larger space of flat connections and for extracting features in an analogous way as demonstrated in [35] for the two-dimensional case. The approach pursued in this paper suggests to explore this direction further by investigating and visualizing complex eigenfunctions on three-manifolds.

Considering the three-dimensional nature of the results we have presented, one might look into adapting the methods to deal more directly with voxel-based data common in applications. Furthermore theory suggests that one might consider bundles of higher rank. These allow for non-Abelian structure groups, such as the group $SO(3)$ of rotations acting on \mathbb{R}^3 or $SU(2)$ of unitary matrices operating on \mathbb{C}^2 being related to quaternions. This is an interesting topic for further research.

References

- [1] Rosenberg S. The Laplacian on a Riemannian manifold: an introduction to analysis on manifolds. Cambridge University Press; 1997.
- [2] Bobenko AI, Springborn BA. A discrete Laplace–Beltrami operator for simplicial surfaces. *Dis Comput Geom* 2007;38(4):740–56.
- [3] Pinkall U, Juni SD, Polthier K. Computing discrete minimal surfaces and their conjugates. *Exp Math* 1993;2:15–36.
- [4] Wardetzky M, Mathur S, Kälberer F, Grinspun E. Discrete laplace operators: no free lunch. In: *Proceedings of the eurographics symposium on geometry processing*; 2007. p. 33–7.
- [5] Reuter M, Biasotti S, Giorgi D, Patanè G, Spagnuolo M. Discrete Laplace Beltrami operators for shape analysis and segmentation. *Comput Graph* 2009;33(3):381–90.
- [6] Dey TK, Ranjan P, Wang Y. Convergence, stability, and discrete approximation of Laplace spectra. In: *ACM-SIAM symposium on discrete algorithms*; 2010. p. 650–63.
- [7] Reuter M, Wolter F-E, Peinecke N. Laplace–Beltrami spectra as shape DNA of surfaces and solids. *Computer-Aided Des* 2006;38(4):342–66.
- [8] Peinecke N, Wolter F-E, Reuter M. Laplace-spectra as fingerprints for image recognition. *Computer-Aided Des* 2007;39(6):460–76.
- [9] Peinecke N, Wolter F-E. Mass density Laplace-spectra for image recognition. In: *Cyberworlds, IEEE*; 2007. p. 409–16.
- [10] Wolter F-E, Friehe K-I. Local and global geometric methods for analysis interrogation, reconstruction, modification and design of shape. In: *Computer graphics international*; 2000. p. 137–51. Also available as Welfenlab report No. 3. ISSN 1866-7996.
- [11] Wolter F-E, Howind T, Altschaffel T, Reuter M, Peinecke N. Laplace-Spektren - Anwendungen in Gestalt- und Bildkognition, available as Welfenlab Report No. 7. ISSN 1866-7996.
- [12] Wolter F-E, Peinecke N, Reuter M. Verfahren zur Charakterisierung von Objekten/a method for the characterization of objects (surfaces, solids and images), German Patent Application; June 2005 (pending). US Patent US2009/0169050 A1; July 2, 2009, 2006.
- [13] Wolter F-E, Blanke P, Thielhelm H, Vais A. Computational differential geometry contributions of the Welfenlab to GRK 615. In: *Modelling, simulation and software concepts for scientific-technological problems, LNACM*, vol. 57; 2011. Springer, p. 211–36.
- [14] Vallet B, Lévy B. Spectral geometry processing with manifold harmonics. *Comput Graph Forum* 2008;27(2):251–60.
- [15] Dong S, Bremer PT, Garland M, Pascucci V, Hart JC. Spectral surface quadrangulation. *ACM TOG* 2006;25(3):1057–66.
- [16] Tong Y, Alliez P, Cohen-Steiner D, Desbrun M. Designing quadrangulations with discrete harmonic forms. In: *Proceedings of the eurographics symposium on geometry processing*; 2006. p. 201–10.
- [17] Kälberer F, Nieser M, Polthier K. QuadCover-surface parameterization using branched coverings. *Comput Graph Forum* 2007;26(3):375–84.

- [18] Bommers D, Zimmer H, Kobbelt L. Mixed-integer quadrangulation. *ACM TOG* 2009;28(3):77.
- [19] Botsch M, Sorkine O. On linear variational surface deformation methods. *IEEE Trans Visual Comput Graph* 2008;14(1):213–30.
- [20] Rustamov RM. Laplace–Beltrami eigenfunctions for deformation invariant shape representation. In: Belyaev A, Garland M, editors. *Proceedings of the eurographics symposium on geometry processing*; 2007. p. 225–33.
- [21] Sun J, Ovsjanikov M, Guibas L. A concise and provably informative multi-scale signature based on heat diffusion. *Comput Graph Forum* 2009;28(5):1383–92.
- [22] Bronstein MM, Kokkinos I. Scale-invariant heat kernel signatures for non-rigid shape recognition. In: *Conference on computer vision and pattern recognition, IEEE*; 2010. p. 1704–11.
- [23] Ruggeri M, Patane G, Spagnuolo M, Saube D. Spectral-driven isometry-invariant matching of 3D shapes. *J Comput Vis* 2010;89(2):248–65.
- [24] Mémoli F. Spectral Gromov–Wasserstein distances for shape matching. In: *International conference on computer vision*; 2009. p. 256–63.
- [25] Reuter M. Hierarchical shape segmentation and registration via topological features of Laplace–Beltrami eigenfunctions. *J Comput Vis* 2010;89(2):287–308.
- [26] De Silva V, Morozov D, Vejdemo-Johansson M. Persistent cohomology and circular coordinates. *Dis Comput Geomet* 2011;45(4):737–59.
- [27] Reuter M, Wolter F-E, Shenton M, Niethammer M. Laplace–Beltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis. *Comput Aided Des* 2009;41(10):739–55.
- [28] Niethammer M, Reuter M, Wolter F-E, Bouix S, Peinecke N, Koo M-S, et al. Global medical shape analysis using the Laplace–Beltrami spectrum. In: *MICCAI*; 2007. Springer. p. 850–7.
- [29] Zhang H, Van Kaick O, Dyer R. Spectral mesh processing. *Comput Graph Forum* 2010;29(6):1865–94.
- [30] Wardetzky M, Bergou M, Harmon D, Zorin D, Grinspun E. Discrete quadratic curvature energies. *Comput Aided Geomet Des* 2007;24(8):499–518.
- [31] Hildebrandt K, Schulz C, von Tycowicz C, Polthier K. Eigenmodes of surface energies for shape analysis. In: Mourrain B, Schaefer S, Xu G, editors. *Proceedings of the geometric modeling and processing. Lecture notes in computer science*, vol. 6130; 2010. Springer. p. 296–314.
- [32] Zobel V, Reininghaus J, Hotz I. Generalized heat Kernel signatures. In: *International conference on computer graphics, visualization and computer vision*; 2011. p. 93–100.
- [33] Crane K, Pinkall U, Schröder P. Spin transformations of discrete surfaces. *ACM TOG* 2011;30:104.
- [34] Vais A, Berger B, Wolter F-E. Spectral computations on nontrivial line bundles. *Comput Graph* 2012;36(5):398–409.
- [35] Vais A, Berger B, Wolter F-E. Complex line bundle Laplacians. *Vis Comput* 2012;28(8):1–13.
- [36] Singer A, Wu H-T. Vector diffusion maps and the connection Laplacian. *Commun Pure Appl Math* 2012;65(8):1067–144.
- [37] Aurich R, Steiner F. Periodic-orbit sum rules for the Hadamard–Gutzwiller model. *Phys D: Nonlin Phenom* 1989;39(2–3):169–93.
- [38] Maintrot M. Finite element method on Riemann surfaces and applications to the Laplacian spectrum. PhD thesis. EPFL; 2012.
- [39] Pansart JP. Numerical calculation of the lowest eigenmodes of the Laplacian in compact orientable 3-dimensional hyperbolic spaces; 2008. p. 27 arXiv0809.0591.
- [40] Seifert H. Über das Geschlecht von Knoten. *Math Ann* 1935;110(1):571–92.
- [41] van Wijk J, Cohen A. Visualization of Seifert surfaces. *IEEE Trans Visual Comput Graph* 2006;12(4):485–96.
- [42] Kotiuga PR. An algorithm to make cuts for magnetic scalar potentials in tetrahedral meshes based on the finite element method. *IEEE Trans Magn* 1989;25(5):4129–31.
- [43] Crane K, Desbrun M, Schröder P. Trivial connections on discrete surfaces. *Comput Graph Forum (SGP)* 2010;29(5):1525–33.
- [44] do Carmo MP. *Riemannian geometry*. Boston: Birkhäuser; 1992.
- [45] Frankel T. *Geometry of physics*. Cambridge University Press; 2003.
- [46] Hatcher A. *Algebraic topology*. Cambridge University Press; 2002.
- [47] Desbrun M, Kanso E, Tong Y. Discrete differential forms for computational modeling. *Dis Differ Geomet: Appl Introd* 2008:287–324.
- [48] Edelsbrunner H, Harer J. *Computational topology. An introduction*. American Mathematical Society; 2010.
- [49] Solin P. *Partial differential equations and the finite element method*. John Wiley & Sons; 2005.
- [50] Murdoch TA. Twisted-calibrations and the cone on the veronese surface. PhD thesis. Rice University; 1988.
- [51] Dey TK, Guha S. Computing homology groups of simplicial complexes in \mathbb{R}^3 . *JACM* 1998;45(2):266–87.
- [52] Boltcheva D, Canino D, Merino Aceituno S, Léon J-C, De Florian L, Héty F. An iterative algorithm for homology computation on simplicial shapes. *Computer-Aided Des* 2011;43(11):1457–67.
- [53] Dlotko P, Specogna R. Physics inspired algorithms for (co)homology computation; December 2012. p. 1–18 arXiv1212.1360.
- [54] CHOMP—computational homology project; 2012. URL (<http://chomp.rutgers.edu/>).
- [55] Kaczynski T, Mischaikow K, Mrozek M. *Computational homology. Applied mathematical sciences*, vol. 157. Springer-Verlag; 2004.
- [56] TetGen: a quality tetrahedral mesh generator and three-dimensional Delaunay triangulator; 2011. URL (<http://wias-berlin.de/software/tetgen/>).
- [57] Hernandez V, Roman JE, Vidal V. SLEPC: a scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Trans Math Software* 2005;31(3):351–62.
- [58] Stillwell J. *Classical topology and combinatorial group theory*. Springer; 1993.
- [59] Gross PW, Kotiuga PR. *Electromagnetic theory and computation: a topological approach*. Cambridge University Press; 2004.
- [60] Kotiuga PR. On making cuts for magnetic scalar potentials in multiply connected regions. *J Appl Phys* 1987;61(8):3916.
- [61] Conway JH, Rossetti JP. Describing the Platycosms arXivmath/0311476.
- [62] Weber C, Seifert H. Die beiden Dodekaederräume. *Math Zeitschrift* 1933;237–53.
- [63] Weeks JR. *The shape of space. second ed.* CRC Press; 2001.