Laplace-Spectra as Fingerprints for Shape Matching

Martin Reuter[*]
University of Hannover
Welfenlab - Division of CG

Franz-Erich Wolter[†]
University of Hannover
Welfenlab - Division of CG

Niklas Peinecke[‡]
University of Hannover
Welfenlab - Division of CG

Abstract

This paper introduces a method to extract fingerprints of any surface or solid object by taking the eigenvalues of its respective Laplace-Beltrami operator. Using an object's spectrum (i.e. the family of its eigenvalues) as a fingerprint for its shape is motivated by the fact that the related eigenvalues are isometry invariants of the object. Employing the Laplace-Beltrami spectra (not the spectra of the mesh Laplacian) as fingerprints of surfaces and solids is a novel approach in the field of geometric modeling and computer graphics. Those spectra can be calculated for any representation of the geometric object (e.g. NURBS or any parametrized or implicitly represented surface or even for polyhedra). Since the spectrum is an isometry invariant of the respective object this fingerprint is also independent of the spatial position. Additionally the eigenvalues can be normalized so that scaling factors for the geometric object can be obtained easily. Therefore checking if two objects are isometric needs no prior alignment (registration/localization) of the objects, but only a comparison of their spectra. With the help of such fingerprints it is possible to support copyright protection, database retrieval and quality assessment of digital data representing surfaces and solids.

CR Categories: I.5.3 [Computing Methodologies]: Pattern Recognition—Similarity measures; J.6 [Computer Applications]: Computer aided Engineering—Computer aided design (CAD); K.5.1 [Legal aspects of computing]: Hardware/Software Protection—Copyrights

Keywords: Laplace-Beltrami operator, shape invariants, fingerprints, shape matching, database retrieval, copyright protection, NURBS, parameterized surfaces and bodies, polyhedra

1 Introduction

The characterization and design of the shape of 3D-objects are central problems in computer graphics and geometric modeling. The development of software and hardware tools useful to design and visualize the shape of 3D-objects has advanced rapidly during the past twenty years. Nonetheless fundamental problems pertaining to the characterization of shape are still widely unresolved. It is for example a basic question to quickly identify and retrieve a given object stored in a huge database or to find all similarly shaped objects. During the past forty years a vast number of shape matching and searching techniques have been developed (e.g. using moments,

[*]e-mail:reuter@gdv.uni-hannover.de
[†]e-mail:few@gdv.uni-hannover.de
[‡]e-mail:peinecke@gdv.uni-hannover.de

spherical harmonics or Reeb graphs - a recent survey can be found in [Iyer et al. 2005], see also [Funkhouser et al. 2003]). It should be pointed out that most approaches dealing with shape matching describe procedures to realign the geometric objects (usually called localization or registration cf. [Prokop and Reeves 1992], [Tucker and Kurfess 2003]), and work only on a specific representation (mainly 3D polygonal meshes) of the object.

The point-set of a solid 3D-object with smooth boundary may be described in very different ways (cf. [Wolter et al. 2004]), for example in boundary representation (B-Rep) using NURBS surface patches. Even when restricted to NURBS surfaces, it is not easy to decide if the given objects are similar in their shape. A simple comparison of the control points used to represent the boundary surfaces does not help at all. The problem becomes even more complicated if we consider the possibility that the solid's boundary surfaces might be represented in various other ways e.g. by trigonometric or by implicitly defined functions.

In some of these cases the problem of identifying shapes (for example to protect the copyright of the designer) has been approached with the help of watermarks. This is of special interest when dealing with delicate high precision material e.g. turbine blades, whose design needs major research effort and expensive investments. Even though NURBS patches are very popular today, most watermark techniques deal with polygonal meshes only. Often the watermark data is embedded into these meshes by slightly modifying the vertex location, the connectivity of the mesh or the frequency domain (cf. [Benedens 1999], [Ohbuchi et al. 2002]). For NURBS surfaces watermarking is more difficult and only very few algorithms exist. An algorithm proposed by Ohbuchi et al. [1999] does not change the surface, but is not very robust. Generally, watermarks can be destroyed by a representation change or by a reparametrization of the object, if they are not embedded into the geometry. On the other hand, embedding data into geometry rather than into the representation changes the shape of the object which is unacceptable in many cases. Therefore, the shape of an object has to be identified by geometric invariants that we will call fingerprints or signatures. An example for a fingerprint of shape intrinsic information used to identify shape via registration / alignment of umbilics can be found in [Ko et al. 2003]. Shape intrinsic information does not depend on the given representation of the shape and can be understood as a fingerprint of the shape (if enough information is contained). Many fingerprints (e.g. eigenvalues of the inertia tensor) have strong limitations with respect to the amount of completeness up to which these invariants determine the shape of the object they are related to. It would be helpful to find a vector of numbers associated with the given object having the following properties:

P1. Isometry:
Congruent solids (or isometric surfaces) should have the same fingerprint being independent of the solid's given representation. For some applications it is necessary that the fingerprint is independent of the object's size.

P2. Similarity:
Similar shaped solids should have similar fingerprints. The fingerprint should depend continuously on the shape deformation.

P3. Efficiency:
  The effort needed to compute those fingerprints should be reasonable.

P4. Completeness:
  Ideally, those fingerprints should give a complete characterization of the shape. One step further it would be desirable that those fingerprints could be used to reconstruct the solid.

P5. Compression:
  In addition it would also be desirable that the fingerprint data should not be redundant, i.e. a part of it could not be computed from the rest of the data.

P6. Physicality:
  Furthermore, it would be nice if an intuitive geometric or physical interpretation of the meaning of the fingerprints would be available.

Concerning property **P1** it should be noted that congruent objects, i.e. objects that can be transformed into each other by a proper or non-proper Euclidean motion, are always isometric. Furthermore, isometric solid bodies are always congruent, which is also true for planar domains. Moreover, the property **P1** is important when almost isometric objects like hands with different finger positions or faces with different expressions shall be compared and identified. See e.g. [Elad and Kimmel 2003] for a method using discrete geodesic distances and multidimensional scaling to generate similar signature surfaces, and see [Bronstein et al. 2003] for an application to face recognition.

This paper proposes to use the eigenvalues of the Laplace operator of a domain or solid or the Laplace-Beltrami operator of a surface in the Euclidean space as a fingerprint. This fingerprint (or signature) can be calculated for different object representations (mentioned above) and can even be calculated for grayscale or color images by understanding such an image as height functions or higher dimensional manifolds. This fingerprint fulfills the above properties (with the only exception of **P4**). It is independent of the spatial position and (if desired) even of the object's size, making registration or localization completely unnecessary. It consists of a family of positive numbers that can be compared easily and fast, permitting this approach to be used in time-critical applications like database retrieval.

To avoid any missunderstanding, note that our Laplace-Beltrami operator does not operate on any mesh vertices, but rather on the underlying manifold itself. *It is therefore different from discrete Laplacians on graphs or meshes.* Even though these discrete Laplacians have been used for e.g. dimensionality reduction [Belkin and Niyogi 2003] or mesh compression [Karni and Gotsman 2000], our introduction of our computation of the Laplace-Beltrami spectra of the underlying manifolds in the areas of geometric modeling - CAD in particular and in computer graphics in general - appears to be completely new with the only exception being [Wolter and Friese 2000] containing a sketchy description of some basic ideas and goals discussed in this paper. Moreover the application of the Laplace-Beltrami spectra as fingerprints in order to discriminate and search objects in geometric databases appears to be new. Although there has been done a considerable amount of theoretical research in geometry on the Laplace-Beltrami operator, there only exists very little work dealing with computational research (see e.g. [Descloux and Tolley 1983]).

2 Theoretical Background

In this section we will explain the theoretical background that is needed to understand the spectrum of the Laplace operator and its computation. Let f be a real-valued function, with $f \in C^2$, defined on a Riemannian manifold M (differentiable manifold with Riemannian metric). We define the **Laplace-Beltrami Operator** Δ of a function f as $\Delta f := \mathrm{div}(\mathrm{grad}\, f)$ with grad f being the **gradient** of f and div the **divergence** on the manifold. The Laplace-Beltrami operator can also be stated in local coordinates (see [Chavel 1984]). In this work we restrict ourselves to surfaces in Euclidean space. The **Helmholtz equation** (also known as the **Laplacian Eigenvalue Problem**) is stated as

$$\Delta f = -\lambda f. \qquad (1)$$

The solutions of this equation represent the spatial part of the solutions of the wave equation. In the surface case ($n = 2$) $f(u,v)$ can be understood as the natural vibration form (also eigenfunction) of a homogeneous membrane with the eigenvalue λ. The solutions of the general vibration problem are the solutions $f(u,v)$ of this differential equation on the surface. Any constants of the material are ignored. The boundary condition of a fixed membrane is $f \equiv 0$ on the boundary of the surface domain (Dirichlet boundary condition). Because of this physical interpretation the question if the eigenvalues of the Laplace operator determine the shape of a planar domain has been rephrased by the late mathematician L. Bers in a terse, impressively concise and pictorial way: "**Can one hear the shape of a drum?**" (cf. [Protter 1987] for a historic account).

2.1 Properties of the Spectrum

The following results on the Laplace-Beltrami operator are well known:

1. The **spectrum** is defined to be the family of eigenvalues to the Laplace eigenvalue problem, consisting of a sequence $0 \leq \lambda_1 \leq \lambda_2 \leq \cdots \uparrow +\infty$, with each eigenvalue repeated according to its multiplicity and with each associated eigenspace being finite dimensional ([Berger et al. 1971], p.142). In the case of a closed manifold without boundary the first eigenvalue λ_1 is equal to zero, because here the constant functions are a non trivial solution of the Helmholtz equation. If a boundary exists, the first eigenvalue is always greater than zero, since the only constant solution is trivial (because of the boundary condition).

2. The spectrum is an isometric invariant as it only depends on the gradient and divergence which in turn are defined to be dependent only on the Riemannian structure of the manifold (this implies property **P1.Invariance**).

3. The spectrum depends continuously on the shape of the membrane ([Courant and Hilbert 1993], p.366), thus complying with property **P2.Similarity**. Moreover, it can be shown with similar arguments that the spectrum depends continuously on the Riemannian metric of the Riemannian manifold in general.

4. Furthermore, we know that scaling a surface by the factor a results in scaled eigenvalues by the factor $\frac{1}{a^2}$. Therefore, by normalizing the eigenvalues, shape can be compared regardless of the objects scale.

5. Sadly, the spectrum does not completely determine the shape, even though geometrical data is contained in the eigenvalues. Finally, in 1992 it could be shown [Gordon et al. 1992] that pairs of isospectral but not congruent planar domains exist. Therefore, it will not be possible to satisfy property **P4.Completeness**. Nevertheless no three pairwise isospectral but non-isometric manifolds have been constructed so far and all known pairs of isospectral planar domains have been shown to be non-convex having non-smooth boundaries. It is not sure if triples or isospectral continuous deformations exist in lower dimensions at all. The counter-examples that were constructed were always somewhat artificial and appear to be exceptional. Furthermore some shapes can completely be characterized by their spectrum (e.g. simple analytic surfaces of revolution, cf. [Zelditch 1998] or the planar disks, cf. [Kac 1966]). Therefore, but also for other reasons partly based on experimental studies, we feel that the spectra of the Laplace-Beltrami operator have a significant discrimination power.

6. A substantial amount of geometrical and topological information is known to be contained in the spectrum (for example the dimension and the volume of the manifold). Beyond that, McKean and Singer [1967] showed that in case of a compact manifold with boundary the surface area of the boundary can be obtained from the spectrum. On a surface (dim = 2) the Riemannian volume is of course the surface area and the Riemannian surface area of the boundary is its length. In case of a closed surface and of a planar domain with smooth boundary McKean and Singer deduced the possibility to "hear" the Euler characteristic (and therefore for planar domains the number of holes). Until today even more geometrical information contained in the spectrum has been discovered by analyzing the trace of the heat kernel and its asymptotic expansion (cf. [Protter 1987]). These results endorse property **P6.Physicality** and even **P4.Completeness**, since all geometrical and topological properties determined by the spectrum have to be identical for isospectral objects. In a follow-up paper we will give more details on this subject and show how it is possible to extract these data from our calculated eigenvalues numerically (this has not been done so far). For the proofs of convergence and description of our algorithm the available space in this paper is too limited.

2.2 Numerical Computation

For the numerical computation the first step is to translate the Laplacian eigenvalue problem into a **variational problem** (using Greens formula)

$$\iint \varphi \Delta f \, d\sigma = -\lambda \iint \varphi f \, d\sigma \qquad (2)$$

(with $d\sigma = W d\vec{x} = W \, du \, dv$ being the surface element in the surface case). We then discretize this problem using the Galerkin technique and **Finite Element Method** with linear, quadratic and cubic polynomial form functions ([Strang 1986]) defined on a triangle mesh to obtain a general eigenvalue problem of the form $AU = \lambda BU$ with A, B being sparse positive (semi-)definite symmetric band matrices. The solution vectors U (eigenfunctions) with corresponding eigenvalues λ can then be calculated (e.g. with direct solvers or the Lanczos method). It should be noted that the above integrals are computed on the surface (not on vertices of a given mesh) and are therefore relatively independent of the given mesh (as long as the mesh fulfills some refinement and condition standards). Beyond that, this method is completely independent of the given parametrization.

In the present state we can use the following object representations as input: triangulations of 2d-parameter space together with any given parametrized surface (including B-Splines and NURBS), polygonal surfaces, tetrahedrized polyhedra or 3d-parameter spaces. It is possible to glue parameter spaces with each other or with themselves in order to construct closed or more complex objects. Furthermore, a surface sensitive triangulation technique based on [Chen and Bishop 1997] has been implemented for the creation of high quality meshes on the surfaces. More details on the implementation will be presented in a follow-up paper.

2.3 Convergence and Accuracy

It is well known that the convergence rate of the FEM method with degree p behaves asymptotically (with an error of order $O(h^{p+1})$) as the element size h tends to zero and if the exact solution contains no singularities (see e.g. [Zienkiewicz and Taylor 2000]). This is not true for many discrete Laplace-Beltrami operators defined on meshes and used for computer graphic applications that are not convergent in general (cf. [Xu 2004]). Furthermore, our approximations are very accurate as can be verified when comparing the approximated results with the exact ones known from theory in some cases (rectangle, circle, sphere, cube, ball, cylinder etc.). Additionally, the fact that we are able to extract the correct geometrical data (volume, boundary length, Euler characteristic) from the heat trace expansion indicates that the calculated eigenvalues are very precise in these cases.

3 Applications

With the help of our shape fingerprints several applications can be realized. As mentioned before one can use these fingerprints to identify objects for the purpose of copyright protection, even when they are given in different representations. Similar objects having similar fingerprints can be detected fast by comparing only the fingerprints, enabling the use of our technique for database retrieval and shape matching. Furthermore these fingerprints can be used for quality assessment, e.g. when converting objects to prove that the shape of the converted object has not been modified.

In order to apply our technique in these areas three requirements need to be fulfilled. Firstly the fingerprint data needs to be very **accurate**, ensuring that the same fingerprint is calculated for an object when using different meshes or even when it is calculated analytically. This requirement can be fulfilled, by using dense enough meshes, since the convergence of the FEM is known. Secondly similar but different objects need to have different fingerprints, making it possible to **discriminate** these objects. This property is only violated by the artificially constructed isospectral but not isometric objects described earlier. Thirdly similar objects need to have similar fingerprints in order to detect the **similarity**. This requirement is completely fulfilled by the spectra.

First we want to corroborate the requirement "accuracy" by looking at a few example calculations where the exact eigenvalues are known. The eigenvalues of the unit disk are the squares of the roots of the Bessel functions J_n ([Courant and Hilbert 1993], p.261). By approximating the area of the disk with triangles we will always get a discretization error at the boundary. Therefore the resulting eigenvalues are not very accurate (the first value having an error of 0.067% with cubic elements and 4021 vertices). These results can be improved by utilizing a different parametrization (not the identity) defined on a polygonal parameter space glued with itself at the appropriate edges as can be seen in figure 1. This way the eigen-
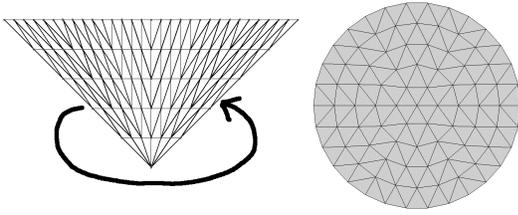
Figure 1: Parameter space of the disk

value λ_{99} is approximated with a relative error of only 0.005%. Generally, we observe an advantage of the Laplace-Beltrami operator over the simple Laplacian even for planar domains: using the free choice of a parametrization we have utilized something like "free form finite elements". By this means any surface can be triangulated **without** discretization error at the boundary. The computations of eigenvalues based on these parametrizations are very accurate.

A similar eigenvalue calculation as in the disk case can be done for the sphere. Even though the sphere is a surface with curvature, the eigenvalues computed with this parametrization are very accurate. $\lambda_{101} = 110$ for example is approximated with 5447 nodal points to be 110.0036093 with a relative error of 0.0033% and with 11522 nodal points to be 110.0003989 with a relative error of only 0.0004%. To show that we can use the surface of a polyhedron (planar triangular facets) instead of a parametrized surface as input, we calculated the eigenvalues of a polyhedron approximating the unit sphere with 1282 vertices and 2560 triangles. The eigenvalues ($\lambda_1 \approx 0$, $\lambda_{2-4} \approx 2.005$, $\lambda_{4-5} \approx 6.014$ ...) are very close to the eigenvalues of the sphere (0,2 and 6), as expected, since the polyhedron is a good approximation of it and because of the continuity property of the spectrum (cf. section 2.1). Still it should be noted that this faceted sphere is of course only an approximation and not the sphere itself, therefore the spectra have to differ slightly from each other.

For the purpose of shape comparison and identification we need to apply a distance calculation to the spectra. In most cases shape comparison is applied to objects independent of their size. Therefore the fingerprints (spectra) need to be normalized first. Four normalization methods are possible: The fingerprint of a d-dimensional Riemannian manifold can be

1. ... divided by its first non-zero eigenvalue.

2. ... divided by the factor c of its fitting curve $f(x) = cx^{\frac{2}{d}}$ (being the slope of its fitting line for $d = 2$).

3. ... multiplied by $V^{\frac{2}{d}}$, where V is the volume of the manifold extracted by extrapolation from the spectrum (for $d = 2$ this is simply a multiplication by the surface area).

4. ... multiplied by $V^{\frac{2}{d}}$, where V is the real volume that has been calculated externally via a pre-process.

The first normalization method is sufficient when trying to identify an object since the fingerprints will be exactly the same after normalization. The methods two to four can be used if similar shapes are to be detected. It should be noted that the slope of the fitting line is a rough approximation to $4\pi/A$ where A is the area of the two dimensional manifold (cf. [Minakshisundaram and Pleijel 1949]). After normalization the fingerprints are cropped to lower the dimension n (to 10-100) and the mutual Euclidean distances of the resulting n-dimensional vectors can be calculated. It is possible to use a different distance (another p-Norm, a symmetric Hausdorff

distance or the correlation) but our results show that the Euclidean distance leads to good results.

For the application of copyright protection, database retrieval and quality assessment it is necessary that shapes can be identified even if the object is given in a different parametrization, with a different spatial position and size. Therefore we computed the fingerprints of the three B-Spline patches $B1, B1'$ and $B2$ (backs of display dummies, see table 2). The patch $B1'$ is a scaled, translated and rotated version of patch $B1$ where additionally the degree has been raised. Patch $B2$ has a different leaner shape. As expected, the fingerprints of $B1$ and $B1'$ are almost identical (distance 0.079) while the one of $B2$ differs from the others by a distance of 45.6 (when scaling with the fitting line and using $n = 50$ values). When using different meshes for the calculation, a different representation as a height function or even when introducing small errors (at the control points), the fingerprint comparison leads to the same clear results.

To show how the spectrum can distinguish between different shapes with the same area, perimeter, incircle, circumcircle and sum of angles, the eigenvalues of the two polygons Isovol $I1$ and $I2$ (cf. table 2) have been calculated. It can be seen in table 1 that al-

EV	patch 1	patch 2	diff.
λ_1	1.4303697	1.3619184	0.0684513
λ_2	2.9252157	2.4698985	0.4553171
λ_3	3.4078884	3.4220482	0.0141597
λ_4	4.0994027	3.8857960	0.2136067
λ_5	4.6141383	5.3698187	0.7556803

Table 1: Eigenvalues of the iso-domains

ready the first few eigenvalues of the two patches are quite different from each other. The Euclidean distance of the two normalized 50-dimensional fingerprints here is 53.5. Even when transforming $I1$ slowly into $I2$ by moving the north bay to the right (and keeping the above specified geometric properties constant), the corresponding continuous movement of the eigenvalues can be detected.

We also applied our technique to a database of 1000 randomly generated B-Spline surface patches (with 3×3 up to 6×6 coefficients). By using the Euclidean distance of the normalized 11-dimensional vectors of eigenvalues, each patch could be uniquely identified even with deliberately different (not optimal) meshes introducing distinct calculation errors. Still, these inaccurate eigenvalues yielded distances of less than 0.02. Furthermore, from the 500.000 possible pairs of different patches only 300 had a distance of less than 0.3 to each other, none was closer than 0.15.

As a final example we will calculate the fingerprints of all the objects depicted in table 2. We have already talked about $B1, B1', B2, I1$ and $I2$. The objects $H1$ and $H2$ (parts of a car hood) are deformations of each other. The fingerprint of the unit square has been calculated with our FEM method ($S1$) and by analytic methods (SR the real spectrum). The disk fingerprint has been computed with a fine ($D1$) and a coarser ($D2$) mesh (both having smooth boundary) and by using a polygonal disk approximation (linearly bordered $D3$). All of these objects shall now be compared with each other in a way, that not only identifies the identical shapes, but also detects similarities. For this purpose the fingerprints were first normalized by the slope of the fitting line and then cropped to contain only the first 50 eigenvalues. To visualize their position, these vectors were embedded linearly into the two dimensional space with the help of classical multidimensional scaling (MDS), a method performing a principal coordinate analysis (PCA). When plotting only the first two dimensions the display error of the real mutual distances is kept as small as possible [Cox and Cox 2001]. The
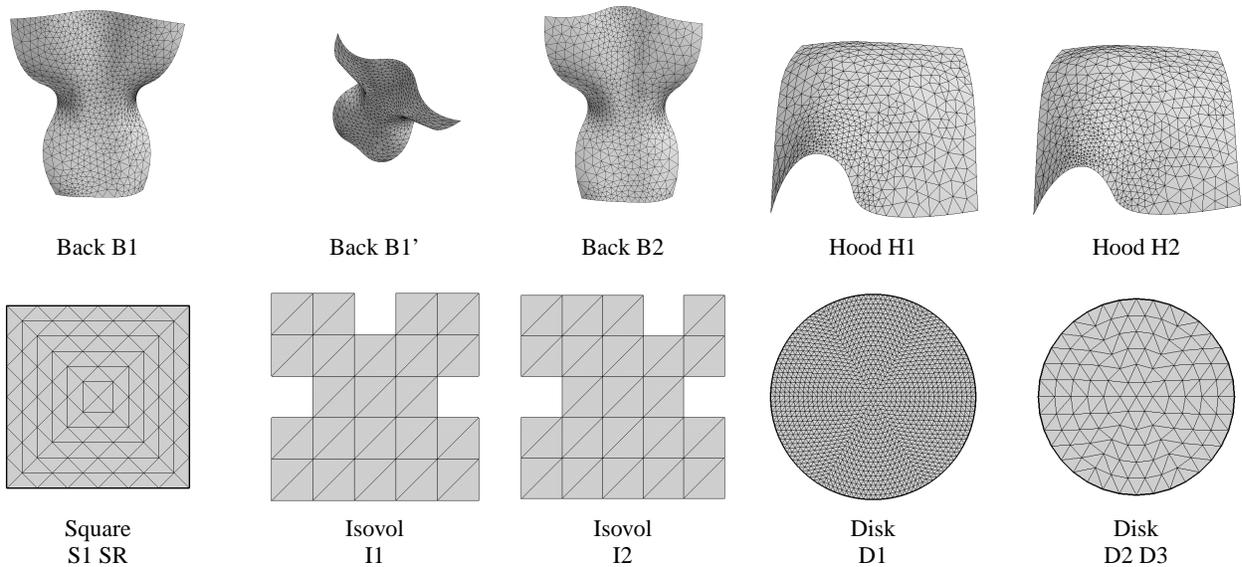
Table 2: 2d test shapes

2d-MDS plot (Figure 2) can be understood as a projection of the 50 dimensional vectors onto their best fitting plane. It shows very well how identical objects are mapped to the same spot and how similar objects form groups. These are very good results, considering that only the first two most important dimensions are plotted and that further information is contained in higher (less important) dimensions.
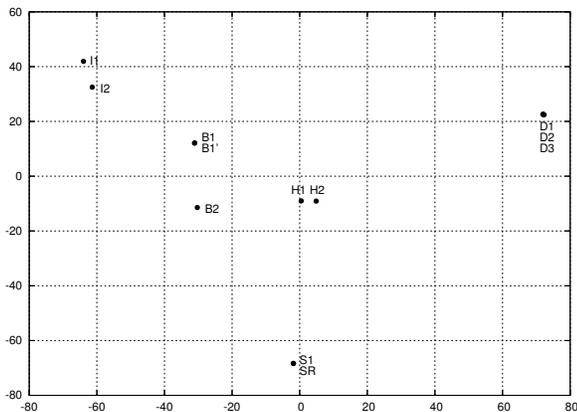

Figure 2: 2d MDS plot of fingerprints

4 Conclusions and Outlook

We have introduced a novel method using the eigenvalues of the Laplace-Beltrami operator as a fingerprint for a given surface or solid and showed the possibility of numerically calculating such fingerprints for different types of geometric objects (parametrized surfaces including NURBS and polyhedra) independent of the given spatial location and scaled size. We demonstrated that shape identification and comparison can be done by using only a few eigenvalues, making it possible to locate objects rapidly within huge databases. As shown, it is even possible to use the Euclidean distance of normalized spectra to detect similar objects. Via comparison of these spectral fingerprints it is possible to compare a suspicious object with a copyrighted one, to detect if the object might be an illegal copy. These fingerprints can also be used as a quality measure when converting surfaces or solids into different representations.

Moreover we succeeded in numerically extracting the volume, boundary length and Euler characteristic from the computed spectrum. Those numerical computations appear to be new and will be presented in a follow-up paper. Further research will extend this work to implicitly defined surfaces and solids employing e.g. new concepts ([Boissonnat et al. 2004], [Ohtake and Belyaev 2002]) useful to mesh implicit surfaces. As another extension parallel processing and multigrid methods could be used to speed up and improve our eigenvalue computation making it possible to use higher resolutions and to compute spectra of even more complex surfaces and especially of 3d-solids.

References

BELKIN, M., AND NIYOGI, P. 2003. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computations 15*, 1373–1396.

BENEDENS, O. 1999. Geometry-based watermarking of 3D models. *IEEE Computer Graphics and Applications 19*, 1, 46–55.

BERGER, M., GAUDUCHON, P., AND MAZET, E. 1971. *Le Spectre d'une Variété Riemannienne*. Lecture Notes in Math. 194, Springer Verlag.

BOISSONNAT, J.-D., COHEN-STEINER, D., AND VEGTER, G. 2004. Isotopic implicit surface meshing. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, ACM Press, 301–309.

BRONSTEIN, A. M., BRONSTEIN, M. M., AND KIMMEL, R. 2003. Expression-invariant 3d face recognition. In *Proc. AVBPA*, Lecture Notes in Comp. Science No. 2688, Springer, 62–69.

CHAVEL, I. 1984. *Eigenvalues in Riemannian Geometry*. Academic Press.

CHEN, H., AND BISHOP, J. 1997. Delaunay triangulation for curved surfaces. In *6th International Meshing Roundtable Proceedings*, 115–127.

COURANT, R., AND HILBERT, D. 1993. *Methoden der mathematischen Physik*. Springer Verlag.

COX, T. F., AND COX, M. A. A. 2001. *Multidimensional Scaling*, 2nd ed. Chapman and Hall.

DESCLOUX, J., AND TOLLEY, M. 1983. An accurate algorithm for computing the eigenvalues of a polygonal membrane. *Comp. Methods Appl. Mech. Engrg. 39*, 37–53.

ELAD, A., AND KIMMEL, R. 2003. On bending invariant signatures for surfaces. *IEEE Transactions on PAMI 25*, 10, 1285–1295.

FUNKHOUSER, T., MIN, P., KAZHDAN, M., CHEN, J., HALDERMAN, A., DOBKIN, D., AND JACOBS, D. 2003. A search engine for 3d models. *ACM Transactions on Graphics 22*, 1, 83–105.

GORDON, C., WEBB, D., AND WOLPERT, S. 1992. Isospectral plane domains and surfaces via Riemannian orbifolds. *Invent. Math. 110*, 1–22.

IYER, N., JAYANTI, S., LOU, K., KALYANARAMAN, Y., AND RAMANI, K. 2005. Three-dimensional shape searching: state-of-the-art review and future trends. *Computer Aided Design 37*, 5, 509–530.

KAC, M. 1966. Can one hear the shape of a drum? *The Amer. Math. Monthly 73*, 4 (Apr.), 1–23.

KARNI, Z., AND GOTSMAN, C. 2000. Spectral compression of mesh geometry. In *Siggraph 2000, Computer Graphics Proceedings*, ACM Press / ACM SIGGRAPH / Addison Wesley Longman, K. Akeley, Ed., 279–286.

KO, K. H., MAEKAWA, T., PATRIKALAKIS, N. M., MASUDA, H., AND WOLTER, F.-E. 2003. Shape intrinsic properties for free-form object matching. *ASME J. of Comp. and Inf. Sci. in Eng. 3*, 4 (Dec.), 325–333.

MCKEAN, H., AND SINGER, I. 1967. Curvature and the eigenvalues of the Laplacian. *J. Differential Geometry 1*, 43–69.

MINAKSHISUNDARAM, S., AND PLEIJEL, A. 1949. Some properties of the eigenfunctions of the Laplace-operator on Riemannian manifolds. *Canadian Journal of Mathematics 1*, 242–256.

OHBUCHI, R., MASUDA, H., AND AONO, M. 1999. A shape-preserving data embedding algorithm for nurbs curves and surfaces. In *Proceedings of the International Conference on Computer Graphics*, IEEE Computer Society, 180–187.

OHBUCHI, R., MUKAIYAMA, A., AND TAKAHASHI, S. 2002. A frequency-domain approach to watermarking 3d shapes. *Eurographics 21*, 3.

OHTAKE, Y., AND BELYAEV, A. G. 2002. Dual/primal mesh optimization for polygonized implicit surfaces. In *Proceedings of the seventh ACM symposium on Solid modeling and applications*, ACM Press, 171–178.

PROKOP, R. J., AND REEVES, A. P. 1992. A survey of moment-based techniques for unoccluded object representation and recognition. *CVGIP: Graph. Models Image Process. 54*, 5, 438–460.

PROTTER, M. H. 1987. Can one hear the shape of a drum? Revisited. *SIAM Review 29*, 2 (June).

STRANG, G. 1986. *Introduction to Applied Mathematics*. Wesley-Cambridge Press.

TUCKER, T. M., AND KURFESS, T. R. 2003. Newton methods for parametric surface registration. Part I. Theory. *Computer-Aided Design 35*, 1, 107–114.

WOLTER, F.-E., AND FRIESE, K. 2000. Local and global geometric methods for analysis, interrogation, reconstruction, modification and design of shape. In *Proc. Computer Graphics International*.

WOLTER, F.-E., PEINECKE, N., AND REUTER, M. 2004. Geometric modeling of complex shapes and engineering artifacts. In *Encyclopedia of Computational Mechanics*, E. Stein, R. de Borst, and T. Hugues, Eds., vol. 1. John Wiley and Sons.

XU, G. 2004. Discrete Laplace-Beltrami operators and their convergence. *Computer Aided Geometric Design 21*, 767–784.

ZELDITCH, S. 1998. The inverse spectral problem for surfaces of revolution. *Journal of Differential Geometry 49*, 2, 207–264.

ZIENKIEWICZ, O., AND TAYLOR, R. 2000. *The Finite Element Method - Volume 1: The Basis*. Butterworth Heinemann.