Eurographics Symposium on Geometry Processing (2007)
Alexander Belyaev, Michael Garland (Editors)

Laplace-Beltrami Eigenfunctions for Deformation Invariant Shape Representation

Raif M. Rustamov

Purdue University, West Lafayette, IN

Abstract

A deformation invariant representation of surfaces, the GPS embedding, is introduced using the eigenvalues and eigenfunctions of the Laplace-Beltrami differential operator. Notably, since the definition of the GPS embedding completely avoids the use of geodesic distances, and is based on objects of global character, the obtained representation is robust to local topology changes. The GPS embedding captures enough information to handle various shape processing tasks as shape classification, segmentation, and correspondence. To demonstrate the practical relevance of the GPS embedding, we introduce a deformation invariant shape descriptor called G2-distributions, and demonstrate their discriminative power, invariance under natural deformations, and robustness.

1. Introduction

Of crucial importance in computer graphics, shape modeling, medical imaging and 3D face recognition is matching, retrieval, correspondence, and segmentation of non-rigid, deformable shapes. An interesting problem, then, is to obtain a shape representation that is invariant under natural deformations, and, at the same time, contains enough information to perform these shape processing tasks.

Since the natural articulations of shapes usually leave unchanged the geodesic distances between the surface points, such deformations correspond to various *isometric*– the metric tensor stays unchanged – embeddings of the surface into Euclidean space. Thus, it is most natural to base deformation invariant representations on geodesic distances. One such representation, the canonical forms of [EK03], have been successfully used for such tasks as deformable shape classification [EK03], and pose invariant segmentation [KLT05].

Unfortunately, geodesic distances are sensitive to local topology changes. As a result, the representations based on them will have limited robustness. Can we avoid using the geodesic distances completely?

Our positive answer to this question is inspired by Lévy's beautiful paper [Lév06], where drawing on an elegant analogy with Chladni plates, Lévy convincingly argues that the *eigenfunctions* of the Laplace-Beltrami differential operator "understand the geometry" – in some sense, they capture the global properties of the surface. Potential applications of these eigenfunctions, as exemplified in [Lév06], include signal processing on surfaces, geometry processing, pose transfer, and parametrization. Another source of inspiration is [RWP05], where the *eigenvalues* of the same operator were used as a shape descriptor.

Our main contribution is to introduce a deformation invariant representation of surfaces, namely the *GPS* embedding (Section 4), which is based on combining the Laplace-Beltrami eigenvalues and eigenfunctions. The *GPS* embedding is itself a surface in the infinite-dimensional space, where the inner product and distance are related to the Green's function. Similar to canonical forms, the *GPS* embedding is invariant under natural deformations of the original surface, and can be used for deformable shape processing – its potential applications are as wide as that of canonical forms. We believe that the *GPS* embedding is the first representation to achieve such a scope without using geodesic distances at all.

We describe our framework for computing the *GPS* embedding in Section 5. It is motivated by the Finite Element approach of [VL07], but our explanations carry more geometric flavor. We make several remarks about the discrete Laplace-Beltrami operator that we think are novel.

In Section 6 we demonstrate how our framework can be employed for non-rigid shape classification. To this end we introduce a deformation invariant shape descriptor – *G*2-*distributions*. The idea is simple: for a given surface compute its *GPS* embedding; then, find the $D2$ shape distribution [OFCD02] of the *GPS* embedding (remember, the *GPS* embedding is also a surface). What renders the resulting descriptor useful for non-rigid shape retrieval is the deformation invariance of the *GPS* embedding, and, thereby, of the descriptor. Curiously, these $G2$-distributions turn out to be related to the distribution of the Green's function's values on the surface.

Our initial experiments show that, first, $G2$-distributions are insensitive to isometric deformations; second, they are robust to local topology changes; third, they show promise to be discriminating among different object classes. These observations provide a practical confirmation of some of the theoretical properties of the *GPS* embedding, and reinforce our belief that the *GPS* embedding captures enough information and is robust enough to provide a practically useful framework for deformable shape processing.

2. Related work

Deformation invariant shape representation: Most similar to our approach in scope are methods based on spectral embedding [EK03, JZ07]. One considers the matrix of pairwise geodesic distances between points on the surface. Spectral embedding, for example Multidimensional Scaling, is used to "flatten" this structure – to get an embedding of these points into the Euclidean space such that Euclidean distances differ from the original geodesic ones as little as possible. Since object articulations change geodesic distances little, this approach yields an isometry invariant representation. Such representations were used for shape classification [EK03, JZ07], part correspondence [JZ06], and segmentation [KLT05].

Unfortunately, these methods can be very sensitive to local changes in the topology – a "short circuit" can affect many geodesic distances by rendering canonical forms of two similar objects very different. A solution to circumvent this problem tries to combine both intrinsic (deformation invariant, e.g. geodesic) and extrinsic (not deformation invariant, e.g. Euclidean) distance measures as in [BBK07]. Notice that when extrinsic features are incorporated, the resulting shape representation loses its isometry invariance.

Our approach is also based on embedding a surface into a higher-dimensional Euclidean space. However, *GPS* embedding does not rely on extrinsic features at all, yet it is robust to local topology changes. Moreover, together with the eigenvalues, *GPS* embedding is a complete isometry invariant of a surface – given the *GPS* embedding and the spectrum of the Laplace-Beltrami operator there is a unique corresponding surface *up to an isometric deformation*.

It is worthwhile to emphasize our differences from Jain and Zhang [JZ06, JZ07] once more, because of the common theme of eigenvalues and eigenvectors. Notice that Jain and Zhang use the eigenvalues and eigenvectors of the geodesic distance matrix after application of some kernel; they do not use the Laplace-Beltrami operator. We, on the other hand, do not use geodesic distances or any variation of them at all, but use the eigenvalues and eigenfunctions of the Laplace-Beltrami operator.

Laplace-Beltrami differential operator also appears in the work of Reuter et al. [RWP05]. They propose to use the set of Laplace-Beltrami eigenvalues – the spectrum – as a shape signature. They show that the spectrum contains enough information to discriminate shapes. However, it should be noted that the spectrum does not determine the surface uniquely up to isometry; there are so called isospectral shapes – non-isometric surfaces that have coinciding spectra [Cip93]. More importantly, the possible applications of the *GPS* embedding are wider than just of the spectrum alone.

Applications of Laplace-Beltrami: Discrete versions of Laplace and Laplace-Beltrami operators, usually both referred to as Laplacians, found many applications in geometry processing [Sor06]. To mention a few, Taubin's seminal paper [Tau95] proposes graph Laplacian with Tutte weights for surface fairing. In [NGH04] Laplacians with different weights are used to control the number of critical points of a function on a surface. Dong et al. [DBG*06] use the eigenvectors of a discrete Laplacian for surface quadrangulation. Discrete Laplacian shows up in [XPB06], where Xu et al. handle surface blending, N-sided hole filling and free-form surface fitting using partial differential equations.

In manifold learning, eigenvalues and eigenvectors of Laplace-Beltrami operator were used to define eigenmaps [BN03] and an infinite collection of so called diffusion maps [CLL*05]; this collection includes the map we use to define the *GPS* embedding. Inadvertently, our formulas in Section 4 are similar to [CLL*05]. However, both our motivations and justifications are different; and it is important that we single out only one map among all possible. Diffusion maps were proposed for dimensionality reduction, whereas the *GPS* embedding does exactly the opposite – embeds a surface into a higher dimensional space. It is worth noting that in manifold learning when one passes to the discrete setting, the discrete Laplacians used are weighted graph Laplacians, while for our approach it is absolutely essential to use one of the "faithful" Laplacians – the ones based on discrete differential geometry, because otherwise the representation becomes dependent on the particular triangulation of the surface.

3. Laplace-Beltrami framework

For a closed compact manifold surface S, let Δ denote its Laplace-Beltrami differential operator. Consider the equation

$$\Delta\phi = \lambda\phi.$$

A scalar λ for which the equation has a nontrivial solution is called an eigenvalue of Δ; the solution ϕ is called an eigenfunction corresponding to λ. Note that $\lambda = 0$ is always an eigenvalue – the corresponding eigenfunctions are constant functions.

The eigenvalues of the Laplace-Beltrami operator are non-negative and constitute a discrete set. We will assume that the eigenvalues are distinct, so we can put them into ascending order

$$\lambda_0 = 0 < \lambda_1 < \lambda_2 < \ldots < \lambda_i < \ldots$$

The appropriately normalized eigenfunction corresponding to λ_i will be denoted by ϕ_i. The normalization is achieved using the L_2 inner product. Given two functions f and g on the surface, their inner product is denoted by $\langle f, g\rangle$, and is defined as the surface integral

$$\langle f, g\rangle = \int_S fg.$$

Thus, we require that $\langle \phi_i, \phi_i\rangle = 1$.

Since the Laplace-Beltrami operator is Hermitian, the eigenfunctions corresponding to its different eigenvalues are orthogonal:

$$\langle \phi_i, \phi_j\rangle = \int_S \phi_i \phi_j = 0,$$

whenever $i \neq j$. Given a function f on the surface, one can expand it in terms of the eigenfunctions

$$f = c_0\phi_0 + c_1\phi_1 + c_2\phi_2 + \cdots,$$

where the coefficients are

$$c_i = \langle f, \phi_i\rangle = \int_S f\phi_i.$$

Thus, eigenfunctions of the continuous Laplace-Beltrami operator give an orthogonal basis for the space of functions defined on the surface. Lévy's main point in [Lév06] is that this basis is *the one*; the expansion coefficients provide a *canonical* parametrization of functions defined on the surface, and all geometry processing should be carried out in this coefficient domain. For example, to smooth a function one should simply discard the coefficients corresponding to the larger eigenvalues, i.e. truncate the infinite expansion above.

Whether the Laplace-Beltrami eigenbasis is *the one* or not, it still would be relevant for deformable shape matching – this is because of Laplace-Beltrami's isometry invariance. Perhaps, another factor to single out the Laplace-Beltrami operator among the infinitude of differential operators would be its "simplicity" and well-studiedness.

4. Global Point Signatures

Can we intrinsically characterize a point on a surface – describe its location without referring to an external coordinate system? Let us remind the reader that the Laplace-Beltrami operator and its eigenfunctions are intrinsic in that sense: the values of eigenfunctions can be thought as numbers attached to the points on the surface, these numbers do not depend on how the surface is located in Cartesian coordinates. Thus, it is natural to try to characterize the points by the values of the eigenfunctions.

Given a point $\mathbf{p}$ on the surface, we define its *Global Point Signature*, $GPS(\mathbf{p})$, as the infinite-dimensional vector

$$GPS(\mathbf{p}) = \left(\frac{1}{\sqrt{\lambda_1}}\phi_1(\mathbf{p}), \frac{1}{\sqrt{\lambda_2}}\phi_2(\mathbf{p}), \frac{1}{\sqrt{\lambda_3}}\phi_3(\mathbf{p}), \ldots\right),$$

where $\phi_i(\mathbf{p})$ is the value of the eigenfunction ϕ_i at the point $\mathbf{p}$. Notice that ϕ_0 is left out because it is void of information. The name comes from the intuition that this infinite-dimensional vector is a signature, a characterization of the point within the global "context" of the surface. Our motivation for normalizing by the inverse root of eigenvalues will be provided shortly.

GPS can be further considered as a mapping of the surface into infinite dimensional space. The image of this map will be called the *GPS embedding* of the surface. We will refer to the infinite dimensional ambient space of this embedding as the *GPS* domain. Let us list some of the properties of the *GPS* embedding.

First, the GPS embedding of a surface without self-intersections has no self-intersections either. We need to prove that distinct points have distinct images under the *GPS*. To this end, suppose that for two surface points $\mathbf{p} \neq \mathbf{q}$ we have $GPS(\mathbf{p}) = GPS(\mathbf{q})$. This means that the eigenfunctions of the Laplace-Beltrami operator satisfy the equality $\phi_i(\mathbf{p}) = \phi_i(\mathbf{q})$. Given any function f on the surface, consider its expansion in terms of the eigenfunctions. Under some mild conditions, the expansion will converge to f pointwise; consequently, the equality $f(\mathbf{p}) = f(\mathbf{q})$ will hold. However, one can easily imagine a "nice" function that takes distinct values at those two points – a contradiction meaning that the *GPS*s must have been different.

Second, GPS embedding is an isometry invariant. This means that two isometric surfaces will have the same image under the *GPS* mapping. Indeed, the Laplace-Beltrami operator is defined completely in terms of the metric tensor, which is itself an isometry invariant. Consequently, the eigenvalues and eigenfunctions of isometric surfaces coincide, i.e. their *GPS* embeddings also coincide.

Third, given the GPS embedding and the eigenvalues, one can recover the surface up to isometry. In fact, eigenvalues and eigenvectors of the Laplace-Beltrami operator uniquely determine the metric tensor. This stems from completeness of eigenfunctions, which implies the knowledge of Laplace-Beltrami, from which one immediately recovers the metric tensor [Ber03], and so, the isometry class of the surface.

Fourth, the GPS embedding is absolute, it is not subject to *rotations or translations of the ambient infinite-dimensional space.* To explain, consider the result, say, of geodesic MDS embedding. This embedding is determined only up to translations and rotations, there is no uniquely determined positional normalization relative to the embedding domain. Thereby, for example, in order to compare two shapes, one still needs to find the appropriate rotations and translations to align the MDS embeddings of the shapes. On the other hand, the *GPS* embedding is uniquely determined – two isometric surfaces will have exactly the same *GPS* embedding (except for reflections, because the signs of eigenfunctions are not fixed), no rotation or translation in the ambient infinite-dimensional space will be involved. For example, the center of mass of the *GPS* embedding will automatically coincide with the origin – this follows from the orthogonality of the eigenfunctions, namely from the equality $\langle \phi_i, \phi_0 \rangle = 0$, if one remembers that ϕ_0 is constant.

Fifth, the inner product and, thereby, the Euclidean distance in the GPS domain have a meaningful interpretation. Here we will need a quick digression about the Green's function, $G(\mathbf{x}, \mathbf{x}')$. One way to solve the differential equation

$$\Delta u = g,$$

where g is some function on the surface, is through the formula

$$u(\mathbf{x}) = \int_S G(\mathbf{x}, \mathbf{x}')g(\mathbf{x}'),$$

where the surface integral is taken with respect to $\mathbf{x}'$. Notice, that the absolute value of $G(\mathbf{x}, \mathbf{x}')$ measures how much $u(\mathbf{x})$ is influenced by the value of $g(\mathbf{x}')$ – how relevant is the input at $\mathbf{x}'$ for the output at $\mathbf{x}$.

Green's function can be written in terms of the eigenfunctions as follows

$$G(\mathbf{x}, \mathbf{x}') = \sum_{i=1}^{\infty} \frac{\phi_i(\mathbf{x})\phi_i(\mathbf{x}')}{\lambda_i}.$$

Clearly, in our setting, $G(\mathbf{p}, \mathbf{q}) = GPS(\mathbf{p}) \cdot GPS(\mathbf{q})$ –the dot product of two infinite-dimensional vectors – which shows that the inner product in the *GPS* domain corresponds to nothing but the Green's function.

Why Green's function is important? It would not be an exaggeration to say that every successful application of Laplace-Beltrami operator points to the relevance of Green's function in shape processing. Let us give an example from mesh editing. To modify a mesh, [YZX*04] consider guidance fields for each mesh vertex – these represent the sought modifications. However, directly applying such a modification would tear the mesh. Instead, the best possible (in the least squares sense) modification is found that will keep the mesh intact. Such a modification is provided by the solution of the equation $\Delta u = g$, where g is the divergence of the guidance field. Consequently, within this mesh editing framework, the magnitude of the Green's function $|G(\mathbf{p}, \mathbf{q})|$ is a measure of how much the points $\mathbf{p}$ and $\mathbf{q}$ are bound by modifications of each other – thus, Green's function in some sense measures the extent to which two points are geometrically "bundled" together. Thus, the inner product in the *GPS* domain is a measure of "togetherness" of two points. Our segmentation example in Figure 1 depends on this fact.

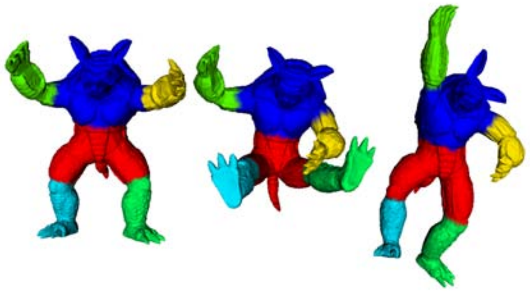

Figure 1: *The k-means clustering on the GPS coordinates results in a pose invariant segmentation.*

To conclude, eigenvalues and eigenfunctions of the Laplace-Beltrami operator are isotopy invariants of a surface. The *GPS* embedding is based on a combination of these, a combination specifically designed to yield a meaningful inner product and, thereby, distance in the *GPS* domain. In contrast to geodesic distances, eigenfunctions and eigenvalues carry more global character, which leads to more stable representations. As a result, at least theoretically, the *GPS* embedding provides an ideal tool for processing of non-rigid shapes – matching, segmentation, and correspondence.

5. Discrete setting

As a terminological note let us mention that it is customary to call Laplace-Beltrami operator as a Laplacian. In addition, one should differentiate between a *discrete* Laplacian and a *combinatorial* one, e.g. Graph Laplacian: the former is specifically designed to keep many of the properties of its continuous counterpart and to faithfully capture geometric and topological properties of the underlying surface, while the latter is sensitive to the peculiarities of the particular triangulation. For our purposes it is crucial that mesh dependence is as minimal as possible – we have to use a discrete Laplacian.

Interestingly, constructing a discrete Laplace-Beltrami operator is a highly non-trivial task. Perhaps [PP93] was the first paper to consider an approach different from one based on the central difference formula. Afterwards, several versions were proposed in [DMSB99], [GSS99], [Tau00], [MDSB02], [Xu04b]. The comparative study carried out in [Xu04a] singles out the versions described in Desbrun et al. [DMSB99] and Meyer et al. [MDSB02] – these two were the only ones to converge in important cases to the continuous counterpart. Moreover, Xu [Xu06] proposes another convergent discrete Laplacian by modifying the one by Meyer et al.; his theoretical analysis shows that, at least on the sphere, this modification leads to better convergence properties. Thus, we decided to base our computations on the discrete Laplacian of Xu, which we will shortly review.

Before embarking, let us explain our notation. We use arrows to distinguish (column) vectors. For example $\vec{v}$ is a vector, and its i-th entry is denoted by v_i, without an arrow; on the other hand, $\vec{v}_i$ is the i-th vector within some indexed set of vectors. Capital letters are used for matrices. Thus, M is a matrix, and its ij entry would be denoted by M_{ij}. The vertices of the triangle mesh representing the surface are denoted by p_i. Given a function f on the surface, its discrete version is the vector $\vec{f}$ with $f_i = f(p_i)$.

5.1. The generalized eigenvalue problem

We shortly review the definition of the discrete Laplace-Beltrami differential operator. For a function f defined on the surface, the value of Δf is approximated as

$$\Delta f(p_i) \approx \frac{1}{s_i} \sum_{j \in N(i)} \frac{\cot\alpha_{ij} + \cot\beta_{ij}}{2} [f(p_j) - f(p_i)].$$

The angles appearing in this formula are depicted in the Figure 2; s_i is the area of the shaded region in the same figure. The summation is over all vertex indices j adjacent with vertex p_i. Let us denote

$$m_{ij} = \frac{\cot\alpha_{ij} + \cot\beta_{ij}}{2}$$

when i and j are adjacent, and $m_{ij} = 0$ otherwise.

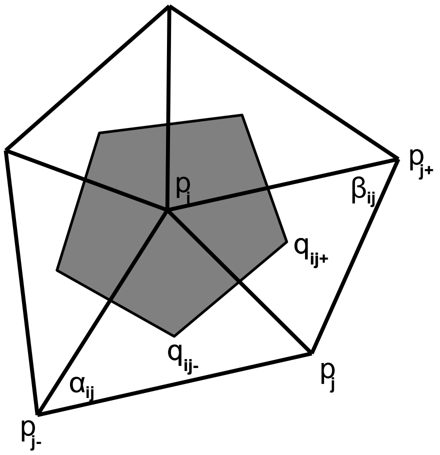

Figure 2: *Definitions of the angles and the area appearing in the discrete Laplace-Beltrami operator.*

Using the column-vector $\vec{f}$, the formula above can be written as a matrix-vector multiplication $\Delta f \approx L\vec{f}$. The involved matrix L – the discrete Laplacian– has the entries as follows

$$L_{ij} = \begin{cases} \sum_k m_{ik}/s_i & \text{if } i = j, \\ -m_{ij}/s_i & \text{if } i \text{ and } j \text{ adjacent}, \\ 0 & \text{otherwise}. \end{cases}$$

Since the areas s_i associated with mesh vertices can vary from vertex to vertex, the discrete Laplacian matrix L is not symmetric. Finding the discrete counterpart of Laplace-Beltrami eigenvalues and eigenfunctions is equivalent to the standard eigenvalue problem for the matrix L

$$L\vec{v} = \lambda\vec{v}.$$

The non-symmetry of L causes problems – both numerical and theoretical. First, we do not have a guarantee that the eigenvalues and eigenvectors of a nonsymmetric matrix will be real; even if they were real, the numerical procedures would sometimes yield complex results. Second, it is not clear how to normalize the eigenvectors – using the usual dot product of vectors causes inconsistency. Indeed, the eigenfunctions of the continuous Laplace-Beltrami operator are orthogonal, while the eigenvectors of the discrete version are not (if one uses the usual dot product). Vallet and Levy [VL07] use the Finite Element Method to explain the root of this inconsistency. Essentially, the following explanation is equivalent to theirs, yet it has more geometric falvor.

Let us rewrite the eigenvalue problem above as a *generalized eigenvalue problem*. Consider the diagonal matrix S with $S_{ii} = s_i$. Denote by M the matrix whose entries are given by $M_{ij} = m_{ij}$, the cotangent weights above. Notice that $L = S^{-1}M$. The equation $L\vec{v} = \lambda\vec{v}$ can be rewritten as $S^{-1}M\vec{v} = \lambda\vec{v}$, or

$$M\vec{v} = \lambda S\vec{v}. \tag{1}$$

Although this formulation is equivalent to the standard one – we would get the same eigenvalues and eigenvectors as in the standard case – this one goes under the name of generalized eigenvalue problem.

Let us remind that two matrices appear in a generalized eigenvalue problem, say A and B: the equality to satisfy is $A\vec{v} = \lambda B\vec{v}$. If matrix A is symmetric, and matrix B is symmetric positive-definite, then the generalized eigenvectors $\vec{v}$ corresponding to different generalized eigenvalues λ are orthogonal. However, the orthogonality here is in terms of B-inner (dot) product:

$$\langle \vec{u}, \vec{w} \rangle_B = \vec{u}^T B\vec{w}.$$

Moreover, all of the generalized eigenvalues/eigenvectors are real. We see that if $B = I$, the identity matrix, we are back to the standard eigenvalue problem, and the statements above are well-known facts about symmetric matrices. Also notice that the I-inner product would be the standard dot product of vectors.

Going back to the equation 1, note that both of S and M are symmetric, and, obviously, S is positive-definite. By the facts mentioned about the generalized eigenvalue problem, the eigenvalues and eigenvectors of discrete Laplacian are to be real. Of course, this fact is not new – the news is that numerical methods for such generalized eigenvalue problems supposedly "know" that the result must be real and behave accordingly. We used the Arnoldi method of ARPACK (this is how MATLAB solves eigenvalue problems) in our experiments, never complex eigenvalues or eigenvectors were obtained.

5.2. Geometric interpretation

One more piece of information from the generalized formulation is about the orthogonality of eigenvectors. We see that the generalized eigenvectors of the problem 1, which are same as the standard eigenvectors of the Laplacian L, are orthogonal with respect to the S-inner product

$$\langle \vec{v}_i, \vec{v}_j \rangle_S = \vec{v}_i^T S \vec{v}_j = 0,$$

when $i \neq j$. We would like to interpret this geometrically.

In the continuous case, the inner product contains a surface integral. Let us investigate how a surface integral can be discretized. Let f be a continuous function defined on the surface. The reader can easily see that the approximation

$$\int_S f \;\approx\; \sum_i f_i s_i$$

is appropriate. Indeed, the sum of s_i's is equal to the total area of the surface mesh – the corresponding regions constitute a complete covering of the surface, and our approximation corresponds to assuming the function to be constant within these regions. Now let us approximate the continuous inner product of two functions f and g:

$$\langle f, g \rangle = \int_S fg \approx \sum_i f_i g_i s_i = \vec{f}^T S \vec{g} = \langle \vec{f}, \vec{g} \rangle_S.$$

The discrete version of the continuous inner product is the S-inner product!

To make the point stronger, consider computing the center of mass of a surface. Of course, the naive averaging formula $C_M = \frac{1}{n}\sum_i p_i$ would not work unless the vertices are distributed evenly on the surface. However, the following one would give a good approximation $C_M = \frac{1}{area}\sum_i s_i p_i$. In the same manner, the formula for the inner product must contain the the areas s_i to weight the vertex contributions appropriately.

Clearly, the S-orthogonality of the eigenvectors corresponds to the orthogonality of the eigenfunctions in the continuous case. The generalized formulation has immediately revealed the root of our yearning for symmetric Laplacians – the hidden expectation that eigenvectors should be orthogonal with respect to the *standard* dot product (I-inner product). Moreover, since the eigenvectors of symmetric Laplacians are orthogonal with respect to the standard dot product – a product that by the above does not have any clear geometric meaning, unless the mesh is uniform – one should not expect symmetric Laplacians to be faithful to the continuous Laplace-Beltrami differential operator. We call this the **dual faithfulness criterion**: if a Laplacian is to be faithful, then its eigenvectors should be orthogonal with respect to a meaningful inner product, a product that is an appropriate discretization or approximation of the continuous counterpart. Consequently, *for a given triangulated surface, a symmetric discrete Laplacian can be faithful only if the mesh vertices are distributed uniformly over the surface area.* Let us emphasize that this result is not about the particular form of Laplacians we are considering (in that case it would have been a trivial fact), but is a statement expressing impossibility in general.

5.3. Computations

All of the computations were performed in MATLAB. The eigenvectors were normalized to have unit length. Here length is defined in terms of the S-inner product by the formula

$$||\vec{v}||^2 = \langle \vec{v}, \vec{v} \rangle_S = \vec{v}^T S \vec{v}.$$

Notice that this normalization is crucial for computations of distances in the *GPS* domain.

For shape classification we used decimated models – number of vertices ranged from a few thousands to twenty five thousands. Due to the $1/\sqrt{\lambda}$ dependence of the *GPS* coordinates, the distances in the *GPS* domain were dominated by the low frequency eigenvectors (small eigenvalues). The number of eigenvectors needed was no more than 25. As a result, the computations took only 5-7 seconds on an Intel T7200 2GHz laptop. Still, to give an idea, computing 250 eigenvectors took about 70-100 seconds.

6. Shape classification using the *GPS* embedding

The *GPS* embedding gives a deformation independent embedding of a shape into the infinite dimensional space. To achieve fast comparison of models, from this embedding we extract a concise descriptor. Of course, we do not work with infinite dimensional space, but consider the projection onto the first d dimensions.

For simplicity, we use a modification of $d2$ distributions introduced in [OFCD02]. Essentially, $d2$ is the histogram of pairwise distances between the points uniformly sampled from the surface. To capture more information we modify $d2$ somewhat: instead of using just one histogram, we construct m^2 histograms(actually $m(m+1)/2$, because of symmetry), where m is some integer. First, consider m equally spaced spherical shells centered at the origin of d-dimensional space. These subdivide the image of the *GPS* embedding into patches contained between successive shells. Each of our m^2 histograms captures the distribution of distances between points one of which belongs to one inter-shell patch, the other to another patch.

Let us repeat again that these histograms are not computed on the object itself, but on its *GPS* embedding. Clearly, our descriptor is related to the distribution of Green's function $G(\mathbf{p},\mathbf{q})$ for uniformly sampled points $\mathbf{p}$ and $\mathbf{q}$ on the surface. This is why we call this descriptor as the object's $G2$-distribution.

Results: For successful shape classification local modifications even if they include topology changes should affect the shape signature only slightly. This is where, for example, the geodesic multidimensional scaling experiences problems [BBK07]. To test the stability of the *GPS* embedding in this context we compare the $G2$-distributions of the original Homer model with its modification. The modification adds a "short-circuit" at the feet by welding the mesh within the red circle in Figure 3. The corresponding $G2$-distribution histograms are shown – red for the original and blue for the modified. Clearly, there is only small difference between the signatures.

How well does *GPS* tolerate large natural shape deformations? Seven deformations of Armadillo obtained by methods of [YBS06], see Figure 4, provide us with a testing ground. Figure 5 clearly reveals that isometric deformations, however large they might be, influence the $G2$-distributions only slightly. All of the Armadillo models were clustered together very tightly. The same figure also reveals that $G2$-distributions can distinguish objects belonging to different categories well.

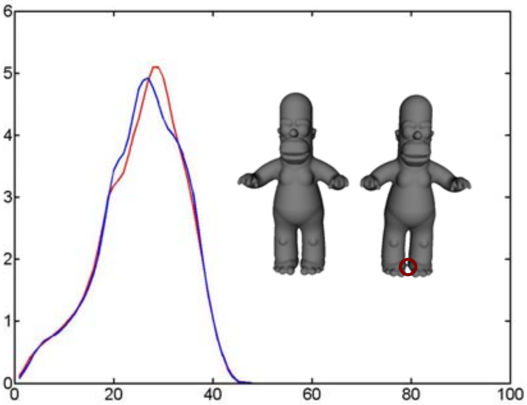

Figure 3: *The G2-distribution histograms for the original Homer versus topology modified version of the model are shown. For clarity $m = 1$ distribution is shown; using more refined G2-distribution gives similar results, e.g. for $m = 8$ see Figure 5. The horizontal axis shows the histogram bin numbers.*

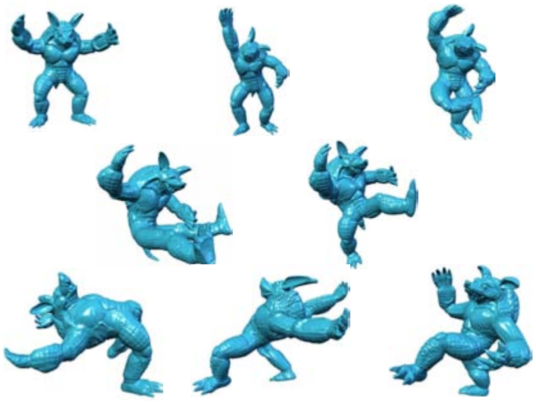

Figure 4: *Armadillo and its deformations.*

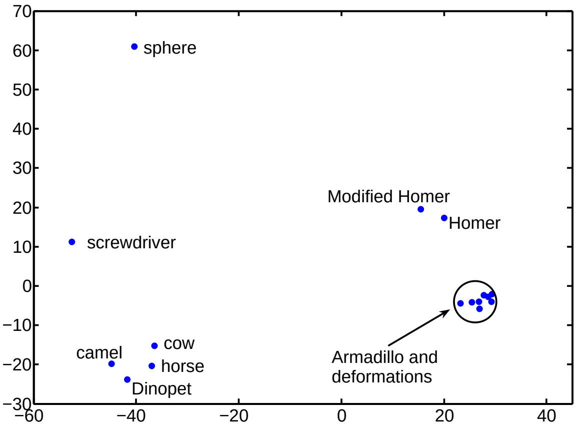

Figure 5: *The classical MDS projection of shape similarities as computed using the G2-distributions. The dimension of the GPS embedding is $d = 15$; the number of shells used is $m = 8$, which generates 36 histograms. The sum of L2 distances between respective histograms is used to compare two objects.*

The last concern we address is whether our discrete framework is sensitive to the underlying triangulation. The answer is again provided by Armadillos: the original Armadillo has been independently simplified from a denser model, and it has some five times more vertices than the deformations, yet it has been clustered together with the other Armadillos.

7. Summary and future work

We have described a new framework to represent non-rigid shapes. Our main contribution is to introduce the *GPS* embedding as a means of ripping a surface from its "transient", Euclidean embedding related properties, to keep its essence – features that are isometry invariant.

To demonstrate the practical relevance of the *GPS* embedding we introduced $G2$-distributions as shape descriptors, and have conducted initial studies of their discriminative power, robustness to local shape changes, including topology modifications, and deformation independence. We plan to perform large-scale experiments to further understand the properties of the *GPS* embedding based signatures.

The main drawback of the Laplace-Beltrami framework is its inability to deal with degenerate meshes. We did not mention surfaces with boundaries neither. However, we think that one should be able to handle them by imposing appropriate boundary conditions.

We should also mention that in practice there are two problems while working with eigenvalues and eigenvectors in general [JZ06]: the signs of eigenvectors are undefined, and two eigenvectors may be swapped. Using d2 distributions indirectly addresses both of these issues. Further analysis is needed to clarify the consequences of these factors for shape processing when the *GPS* embedding is used directly.

Apart from shape classification, we expect the *GPS* embedding to be relevant in the context of shape correspondence and segmentation. Figure 1 shows a preliminary result from our segmentation experiments using the *GPS*. Simple k-means clustering based on distances in the *GPS* domain was performed to segment the Armadillos into six patches, no further optimization has been done. The figure demonstrates pose-oblivious nature of such segmentation. These results and applications described in [Lév06, VL07] raise hope that Laplace-Beltrami eigenfunctions will provide to geometry processing what Fourier basis has provided to signal processing.

8. Acknowledgements

All of the models except the Dinopet and the sphere were downloaded from AIM@SHAPE Shape Repository. Seven deformations of Armadillo obtained by methods of [YBS06] are courtesy of Shin Yoshizawa. The rest of the models are courtesy of INRIA.

I am deeply grateful to the anonymous reviewers for their detailed and useful comments that helped improve the article immensely.

References

[BBK07] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Joint intrinsic and extrinsic similarity for recognition of non-rigid shapes*. Tech. Rep. CIS-2007-01-2007, Computer Science Department, Technion, March 2007.

[Ber03] BERGER M.: *A panoramic view of Riemannian geometry*. Springer-Verlag, Berlin, 2003.

[BN03] BELKIN M., NIYOGI P.: Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput. 15*, 6 (2003), 1373–1396.

[Cip93] CIPRA B.: You can't always hear the shape of a drum. *What's happening in the Mathematical Sciences*, 1 (1993).

[CLL*05] COIFMAN R. R., LAFON S., LEE A. B., MAGGIONI M., NADLER B., WARNER F., ZUCKER S. W.: Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *PNAS 102*, 21 (2005), 7426–7431.

[DBG*06] DONG S., BREMER P.-T., GARLAND M., PASCUCCI V., HART J. C.: Spectral surface quadrangulation. In *TOG(SIGGRAPH)* (2006), pp. 1057–1066.

[DMSB99] DESBRUN M., MEYER M., SCHRÖDER P., BARR A. H.: Implicit fairing of irregular meshes using diffusion and curvature flow. In *SIGGRAPH* (1999), pp. 317–324.

[EK03] ELAD A., KIMMEL R.: On bending invariant signatures for surfaces. *IEEE Trans. Pattern Analysis and Machine Intelligence 25*, 10 (2003), 1285–1295.

[GSS99] GUSKOV I., SWELDENS W., SCHRÖDER P.: Multiresolution signal processing for meshes. In *SIGGRAPH* (1999), pp. 325–334.

[JZ06] JAIN V., ZHANG H.: Robust 3D shape correspondence in the spectral domain. In *Shape Modeling International* (2006).

[JZ07] JAIN V., ZHANG H.: A spectral approach to shape-based retrieval of articulated 3d models. *Computer Aided Design 39* (2007), 398–407.

[KLT05] KATZ S., LEIFMAN G., TAL A.: Mesh segmentation using feature point and core extraction. *The Visual Computer 21*, 8-10 (2005), 649–658.

[Lév06] LÉVY B.: Laplace-Beltrami eigenfunctions: Towards an algorithm that understands geometry. In *Shape Modeling International* (2006).

[MDSB02] MEYER M., DESBRUN M., SCHRÖDER P., BARR A.: Discrete differential geometry operators for triangulated 2-manifolds. In *Proceedings of Visual Mathematics* (2002).

[NGH04] NI X., GARLAND M., HART J. C.: Fair Morse functions for extracting the topological structure of a surface mesh. In *TOG(SIGGRAPH)* (2004), pp. 613–622.

[OFCD02] OSADA R., FUNKHOUSER T., CHAZELLE B., DOBKIN D.: Shape distributions. *TOG 21*, 4 (2002), 807–832.

[PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experiment. Math. 2*, 1 (1993), 15–36.

[RWP05] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace-spectra as fingerprints for shape matching. In *Solid and Physical Modeling* (2005), pp. 101–106.

[Sor06] SORKINE O.: Differential representations for mesh processing. *Computer Graphics Forum 25*, 4 (2006), 789–807.

[Tau95] TAUBIN G.: A signal processing approach to fair surface design. In *SIGGRAPH* (1995), pp. 351–358.

[Tau00] TAUBIN G.: Geometric signal processing on polygonal meshes. In *EUROGRAPHICS* (2000).

[VL07] VALLET B., LEVY B.: *Spectral Geometry Processing with Manifold Harmonics*. Tech. rep., April 2007.

[XPB06] XU G., PAN Q., BAJAJ C. L.: Discrete surface modelling using partial differential equations. *Comput. Aided Geom. Des. 23*, 2 (2006), 125–145.

[Xu04a] XU G.: Convergence of discrete Laplace-Beltrami operators over surfaces. *Comput. Math. Appl. 48*, 3-4 (2004), 347–360.

[Xu04b] XU G.: Discrete Laplace-Beltrami operators and their convergence. *Comput. Aided Geom. Des. 21*, 8 (2004), 767–784.

[Xu06] XU G.: Discrete Laplace-Beltrami operator on sphere and optimal spherical triangulations. *Int. J. Comput. Geometry Appl. 16*, 1 (2006), 75–93.

[YBS06] YOSHIZAWA S., BELYAEV A., SEIDEL H.-P.: *Skeleton-driven Laplacian mesh deformations*. Research Report MPI-I-2006-4-005, Max Planck Institut Informatik, Saarbruecken, October 2006.

[YZX*04] YU Y., ZHOU K., XU D., SHI X., BAO H., GUO B., SHUM H.-Y.: Mesh editing with Poisson-based gradient field manipulation. In *TOG(SIGGRAPH)* (2004), pp. 644–651.