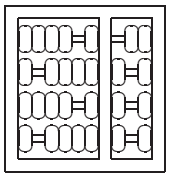

INSTITUTO DE COMPUTAÇÃO
UNIVERSIDADE ESTADUAL DE CAMPINAS

Intrinsic Mesh Segmentation

Fernando de Goes
Siome Goldenstein
Luiz Velho

Technical Report - IC-07-017 - Relatório Técnico

May - 2007 - Maio

Intrinsic Mesh Segmentation

Fernando de Goes* Siome Goldenstein† Luiz Velho‡

Abstract

Mesh segmentation offers a desirable divide-and-conquer strategy for many graphics applications. In this paper, we present a novel, efficient, and intrinsic method to segment meshes following the minima rule. The eigenfunctions of the Laplace-Beltrami operator define locality and volume-shape preserving functions over a surface. Inspired on Manifold learning theory, we use these functions as the basis of an embedding space for mesh vertices and group them using *k*-means clustering. We also present a new kind of segmentation hierarchy built from the analysis of the Laplace-Beltrami operator spectrum.

1 Introduction

Mesh segmentation offers a desirable divide-and-conquer strategy for many graphics applications. Essentially, there are two main types of mesh segmentation algorithms: the ones that partition meshes into patches and the ones that decompose meshes into parts. Patch-type segmentation creates disk-like sub-meshes and is used for texture mapping, building charts, mesh simplification, and geometry-image creation. Part-type segmentation recovers meaningful components and can be used for morphing, compression, skeleton extraction, animation, and modeling. See [Sha04, AKM*06, Sha06] for recent surveys on mesh segmentation.

In this work, we focus on part-type segmentation following the minima rule [HS97]. In the last years, several approaches have been proposed to this problem. Many of them are based on clustering methods using mesh attributes such as curvature, dihedral angles, and geodesic distances. In [PKA03], a fast marching watershed algorithm was used with curvature values to assist shape matching. In [STK02], *k*-means clustering with geodesic distances decomposes meshes into parts for morphing purposes. In [KT03], a fuzzy clustering with graph cuts finds hierarchical segments extracting mesh skeleton. In [LZ04], Liu and Zhang explored the Polarization theory to achieve segmentation using a spectral clustering on the eigenvectors of an affinity matrix based on geodesic distances and dihedral angles. In [ZL05], Zhang and Liu improved previous work computing eigenvectors through Nyström approximation and finding segments by a linear search with salient cut metrics. A boundary approach to mesh segmentation is proposed in [LLS*05]. This work used curvature estimation and geometric-snakes to scissor mesh through salient contours.

*Institute of Computing - UNICAMP - Brazil
†Institute of Computing - UNICAMP - Brazil
‡VISGRAF Project - IMPA - Brazil

Despite of good results, the wide variety of mesh attributes may induce undesirable problems to previous techniques. Curvature attributes tend to create over-segmentation around local concavities and make the algorithms sensitive to pose. Metrics based on geodesic distances are also sensitive to pose and to topologic changes. Besides, irregular meshes complicate the extraction of small features and hierarchical segmentations. Recent works treat above problems through pose invariant tools. In [KLT05], a Multi-dimensional scaling tranforms meshes into canonical poses enabling detection of mesh features and hierarchical components. In [SCOS07], it is presented a volume-shape function to examine the diameter of the neighborhood of surface points, defining a pose oblivous shape signature for meshes.

This paper presents a novel part-type mesh segmentation based on Manifold learning theory and the Laplace-Beltrami operator. The algorithm is close to linear on mesh size and on number of segments. It uses only intrinsic properties of the surface, resulting in a new and more natural kind of hierarchical segmentation, extracting small features, and intrinsically following the minima rule. The main idea of our approach is to explore the locality preserving property of the Laplace-Beltrami eigenfunctions to define a new embedding for mesh vertices, as described in Manifold learning theory. Then mesh vertices are grouped by a *k*-means clustering with Euclidian distances. To decide the desired number of segments, we analyse the growth of the Laplace-Beltrami spectrum, which follows surface geometry and topology.

Our work has the following contributions:

- Introduces an efficient and intrinsic mesh segmentation algorithm.
- Proposes a new interpretation of the Laplace-Beltrami eigenfunctions based on Manifold learning theory.
- Explores the Laplace-Beltrami spectrum to identify hierarchical and meaningful mesh segments.
- Presents a new kind of hierarchy of segments aligned to the minima rule.

2 Laplace-Beltrami Operator

The Laplace-Beltrami operator (LBop) is a generalization of the Laplacian to *nD*-manifolds. Intuitively, it measures the difference of a function f between each point of the manifold and its neighborhood.

2.1 Definition and Properties

For *nD*-manifolds, the Laplace-Beltrami operator is defined as:

$$\Delta = \nabla \circ \nabla = \delta d + d\delta = \sum_i \frac{1}{\sqrt{|g|}} \frac{\partial}{\partial x_i} \sqrt{|g|} \frac{\partial}{\partial x_i}, \tag{1}$$

where δ and d are exterior calculus operators, and $|g|$ is the determinant of the manifold metric tensor.

As the definition suggests, the LBop contains information about the locality of each region on the manifold. Specifically to surfaces, this feature ensures that the operator follows principal curvatures and directions. Therefore, we can write the LBop of a point p as a linear combination of second derivatives of the surface around p parameterized by the arc lengths in the principal directions of p.

The LBop is a fundamental tool in many mathematics and physics problems. In particular, LBop is used to solve the heat diffusion process and its eigenfunctions solve the Helmholtz equation

$$\Delta f = -\lambda f, \tag{2}$$

which solutions represent natural vibrations over manifolds.

The eigenfunctions of the LBop present interesting properties. First of all, they represent a generalization of the function basis of Fourier analysis for *nD*-manifold. In [VL07], these basis are called *Manifold Harmonic Basis* (MHB). As a consequence of this property, the first MHB represent low frequencies surface informations and recover the overall surface shape, while high frequencies basis indicate finer details. Second, MHB functions are smooth, minimizing the Dirichlet energy. Third, the zero set (also called *nodal lines*) of the *n*th MHB subdivides the manifold into maximal *n* subdomains with intersections at constant angles. For more results about the geometry of eigenfunctions, see [JNT01]. Based on above properties, we conclude that the MHB indicate locality and volume-shape preserving surface functions.

The spectrum of the LBop is composed by the sequence of eigenvalues $0 \leq \lambda_1 \leq \lambda_2 \leq \ldots \uparrow \infty$ corresponding to squared spatial frequencies. It is well known that the eigenvalues are strongly related to the surface geometry and topology, determining its area, Betti numbers, and Euler characterist. In [RWP06], the spectrum of the LBop was also used to define meshes signatures for shape matching.

Research on heat equation proved that the asymptotic behavior of the heat trace $Z(t)$ is expressed by

$$Z(t) = \sum_i e^{-\lambda_i t}. \tag{3}$$

Consequently, the error approximation of a heat solution can be controlled by the growth of the spectrum [CL06]. Geometrically, this fact means that the maxima gaps of the spectrum indicate the number of meaningful regions on the surface, providing a hierarchical way to segment meshes into parts (Figure 1).

2.2 Discrete Case

In graphics, a surface is represented by a mesh $M = (V,E,F)$ and then LBop must have a discrete version. The first approach given to the LBop of meshes was following Spectral Graph theory [Chu97] and it defines the LBop as a combinatorial matrix $L = D - A$, where D is a diagonal matrix of vertices valences and A is the graph adjacency matrix. However, this version of the LBop does not take geometry into acount.

A geometric based version of the LBop was first presented by [PP93], defining *cotangent weights* . In [MDSB02, VL07], a FEM (Finite Element Method) was used to complete the derivation

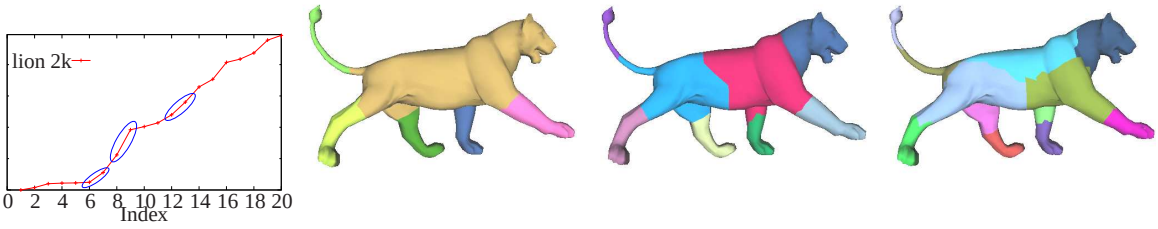

Figure 1: Analysis of LBop spectrum and the hierarchy of segments. Blue ellipses indicate the large spectrum gaps.

of the discrete LBop. Following [VL07], we define a discrete LBop of a mesh by the matrix L

$$\begin{array}{rcl} L(i,j) & = & -(\cot(\alpha_{i,j})+\cot(\beta_{i,j}))/2Area(i) \\ L(i,i) & = & \sum_{j\in N(i)} L(i,j), \end{array} \tag{4}$$

where $N(i)$ is the set of neighbor vertices of i, $Area(i)$ is the Voronoi area around i, and $\alpha_{i,j}$ and $\beta_{i,j}$ are the angles opposite to the edge (i,j).

Note that the matrix L is sparse, so efficient eigen-problem solvers can be used. In this paper, we use the ARPACK solver (see `www.caam.rice.edu/software/ARPACK`). To make computations faster, we use $0.5(L+L^t)$, a symmetric version of L [Lev06]. Figure 2 compares eigenvectors of the LBop computed by the combinatorial version and the geometric one.

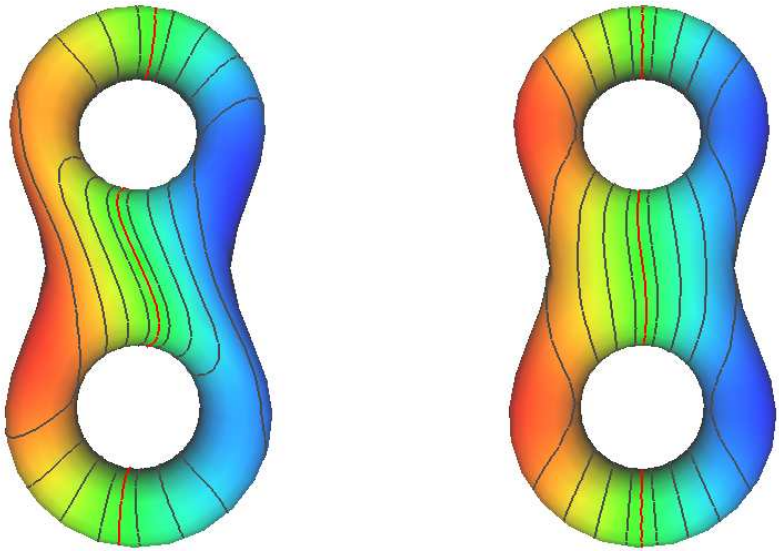

Figure 2: The 3rd eigenfunction of the combinatorial LBop (left) and of the discrete LBop (right). The color indicates values increasing from blue to red. The contours show isolines and nodal lines are the red ones.

Besides the preservation of continuous properties, the discrete LBop gives coherent results through multiresolution analysis. We verify that the initial band of the spectrum (low frequencies) is preserved after successive decimations of a mesh (Figure 3). A similiar result is shown in [DBG*06], but Dong et al. used the LBop just with *cotangent weights*, what requires a mass-adjustment into its spectrum. Considering the vertices area into LBop, we achieve a band spectrum preserving with no required adjustment. In the same way, the MHB of low frequencies are also preserved, but it is worth noting that, as the frequency increases, there are more oscilations on the MHB and higher mesh resolutions are required to achieve good approximations.

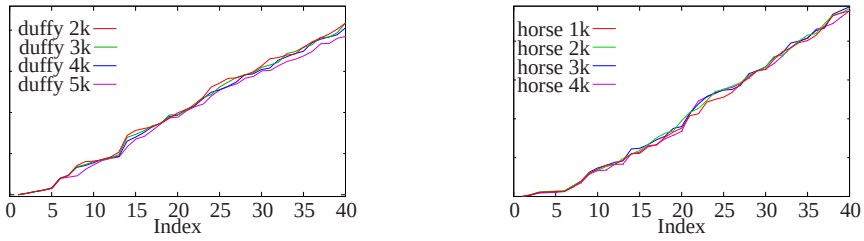

Figure 3: Examples of the initial band of the LBop spectrum preserved after sucessive decimations.

2.3 Applications

In the recent years, the LBop has been used in a diverse set of graphics applications. Among them: smoothing [NISA06], conformal parametrization [GY03], remeshing [DBG*06, TACSD06], pose transfer [Lev06], registration [RWP06], and spectral filtering [VL07]. For more, see [Tau00, Sor06].

3 Manifold Learning

Manifold learning is an area of machine learning that aims to recover intrinsic information about a set of points of a mD-manifold embedded in R^n ($m < n$). Usually, it is desirable to have new representations for the points that preserve manifold characteristics. These new embeddings can be applied to non-linear dimensionality reduction problem, data classification, data organization, and data clustering.

Many of the Manifold learning methods are based on spectral analysis. A classical approach is the PCA (Principal Component Analysis) where the most relevant scatter directions of data set are recovered computing the eigenvectors of high eigenvalues of a covariance matrix. To work with non-linear spaces, Kernel-PCA modifies the distance metric between points by kernel functions. In Isomap [TdSL00], the distance metric used is geodesic lengths over the manifold and, using Multi-dimensional scaling, a new point representation can be computed approximating geodesic distances by Euclidian distances. Isomap is also used in graphics for surface parameterization [ZSGS04].

So far, the Manifold learning methods presented here work globally, trying to preserve distance values between all pair of points, even between distant ones. In [RS00], a new method called LLE (Locally Linear Embedding) proposes to treat points as linear combinations of their neighbors and then it shows that the important information to be preserved in the embedding space is the point locality, i.e., the relation of proximity among points and not distance values.

Extending the idea of locality, eigenfunctions of the manifold LBop (MHB) can be explored over the set of points. MHB result in a new embedding where Euclidian distances measure the proximity between points on the manifold. The use of MHB was first introduced in [BN02] by the name of *Laplacian Eigenmaps*. A recent extension of this approach is the *Diffusion maps* ([CL06]), where the distortion caused by the distribution of samples over the manifold is minimized.

To show that the MHB define embedding spaces with Euclidian distance metric preserving proximity, let's review how the MHB of a set of points are computed. Given a set of points $\{x_i\}_{i=1}^n$, first we connect close points by edges and then we compute the LBop L weighting an edge (x_i, x_j) by the function $w(x_i, x_j)$ – note that, for our purposes, both informations are given intrinsically by the mesh. Now, to find an embedding that preserves the manifold locality, we compute a vector y such that

$$y = \text{argmin}_z z^t L z = \text{argmin}_z \sum_{i,j} w(i,j)(z_i - z_j)^2. \tag{5}$$

Generalizing the above idea to a m dimensional embedding, we use the first m solutions $\{y^i\}_{i=1}^m$ of Equation 5 and define the embedding of a point x_j as the vector $Y_j = [y^i_j]_{i=1}^m$. Finally, it is easy to verify that

$$Y = \text{argmin}_Z Z^t L Z = \text{argmin}_Z \sum_{i,j} w(i,j)\|Z_i - Z_j\|^2. \tag{6}$$

In [MS01], Meila and Shi presented a similiar MHB embedding but in a stochastic fashion. By building a Markov chain over the data set and defining its LBop by a transition matrix, they proved that an exact k-partition of the data set is only possible when the first k MHB are piecewise constant. They also noted that the correct number of partitions k can be determined automatically analysing gaps between the eigenvalues of the LBop.

4 Our Method

The main motivation of our approach is to associate the properties of the MHB of surfaces (Section 2) with the advantages of embedding points by MHB as exposed by Manifold learning theory (Section 3).

4.1 Main Techniques

We now explain the three techniques necessary to accomplish the desired association between MHB with Manifold learning theory and describe our segmentation method.

Embedding and clustering vertices

The MHB compose an embedding space preserving the manifold locality and measuring proximity by Euclidian distances. At the same time, MHB also represent volume-shape functions on surfaces. Therefore, embedding mesh vertices by MHB gives us a way to group vertices into meaningful volumes.

Meaningful volumes are parts of a 3D object delimited by short boundaries and in concavity regions. Such features also state the minima rule. Hence, MHB embedding naturally results in a part-type segmentation following the minima rule.

To group vertices we use a *k*-means clustering [Llo82] on the embedding space. The *k*-means procedure acts as an *EM* optimization. On the *E*xpectation phase, the data set is partitioned into *k* groups minimizing the distance from *k* seed points. On the *M*aximization phase, each seed point is updated by the average of points inside its group. The quality of results depends directly on the choice of the initial seeds. In our implementation, we use the mutually furthest apart points as initial seeds.

The dimension of the embedding space

In Section 2, we saw that the *n*th MHB and, consequently, all lower MHB divide the surface up to *n* subdomains. As MHB are smooth functions, we can use this property as a relaxation of the optimal case of piecewise constant functions described by Meila and Shi [MS01]. From this fact, we set the dimension of the embedding space equal to the desired number of clusters.

The number of segments

The growth of the spectrum of the LBop depends on surface geometry and topology such that similiar eigenvalues correspond to related MHB. In Manifold learning, we have that the *k* first piecewise constant MHB are separated from the others MHB by a maximum gap in the spectrum. Therefore, looking for large gaps in the spectrum of the LBop, we can choose different values of *k* to segment the mesh. In practice, we note that, when we choose *k* not in a gap, the segmentation usually generates non-compact parts.

As the value of *k* increases, more features of the mesh are segmented given us a new kind of hierarchy among segments. In previous works, the hierarchy of segments determines a tree where one parent segment is split into others parts independently. In our approach, for each level of the hierarchy, current segments can be recombined and reorganized, leading to a more natural and coherent segmentation.

We show the structure of our method in Algorithm 1.

Algorithm 1 Intrinsic Mesh Segmentation

Require: A mesh $M = (V, E, F)$.
Ensure: k sub-meshes $M_i = (V_i, E_i, F_i)$.

1: Compute the LBop L of M.
2: Compute the initial band B of the spectrum of L.
3: Find gaps k_i in B. We identify a gap by the k_i-th eigenvalue where it starts (See 4.2).
4: Choose k as the desired number of segments among k_i's.
5: Compute the first k MHB $[y^j]_{j=1}^k$of L.
6: Set $Y_i = [y_i^j]_{j=1}^k$ as the embedding representation for each vertex $i \in V$.
7: Apply *k-means* over Y to find k sub-meshes.
8: [OPTIONAL] Fit sub-meshes boundaries (See 4.3).

4.2 Numerical Solution

The time complexity of our approach to segment a mesh $M = (V,E,F)$ into k parts depends essentially on steps 2, 5, and 7 of Algorithm 1. The cost of k-means (step 7) is $O(ck|V|)$, where c indicates the number of k-means iterations. In our experiments, the value of c is negligible and, as $k \ll |V|$, the cost of this step becomes linear on mesh size. In turn, the cost of steps 2 and 5 depends on eigen-problem solver. To compute the first k MHB, the ARPACK solver used in our implementation is superlinear on k and linear on the number $|E|$ of non-zero elements of the LBop.

It is well known that eigen-problem solvers are very expensive to compute MHB for surfaces with more than a few thousand vertices. The reason for that is the poor condition number to compute lower frequencies of a matrix, requiring high number of iterations to converge to a solution. To compensate for this problem, we invert the spectrum of the LBop. For this, first we compute the maximum eigenvalue $\bar{\lambda}$ of the LBop matrix L and then we compute the highest k eigenvectors of $\bar{\lambda}I - L$. See [VL07] for an alternative solution.

In step 2, we compute a band B of the spectrum of the LBop. The length of B used is 40, but different lengths can be set. As the desired number of segments k is very small compared to the length of the band B, step 2 turns to the dominant computation of our algorithm. However, we can accelerate this computation through a multiresolution approach as described in Section 2.2. So we decimate the orignal mesh to compute the band B and after that we return to the full mesh.

4.3 Fitting Boundary

Fitting boundary is a commom problem on part-type segmentation algorithms. Usually, it is related with imprecisions on segmentation criteria. Hierarchical approaches as [KT03, ZL05, KLT05] explore the tree structure of the hierarchy to refine boundaries. Each node of the three represents a sub-mesh and it is partitioned independently from others nodes. As the number of parts to divide a node is small, the boundaries involve only pairs of segments and then can be refined by a graph min-cut algorithm.

In our method, the boundary smoothing depends on the quality of the approximation of the MHB used in the embedding space. As discussed in Section 2.2, higher MHB may present numerical precision problems. So mesh segmentation into many parts (e.g., more than 15) can result in jagged boundaries, erroneously clustering boundary vertices. Usually, when the mesh is decomposed into many parts, the boundaries involve more than two segments and, for this reason, graph min-cut is not enough to correct them. Nevertheless, for small number of segments, we achieve good boundaries with no adjustment.

5 Results

We have applied our method into different meshes up to 5000 vertices. All experiments were carried on an Intel Pentium 4 3.2 GHz machine with 2 GB RAM. Table 5 reports time statistics. We have only used decimation to compute the inital band of the LBop spectrum for meshes with more than 3000 vertices. It is also important to note that we do not fit boundaries. The segments colors are chosen at random.

Table 1: Statistics of our algorithm. For each mesh, we give respectively the number of vertices, edges, and faces, the decimation resolution (if necessary), computing time for the 40th first eigenvalues, the number of segments k, computing time for the k first MHB, and computing time for k-means. All times are in seconds.

Mesh	**Size**	**Decimation**	**Band**	k	**MHB**	k**-means**
lion	(2002,6000,4000)	-	38	12	6.5	0.8
dino-pet	(1681,4970,3281)	-	20	16	15.3	0.7
duffy	(4590,13761,9171)	(2004,6003,3999)	49.5	13	14.4	4.1
mouse	(1872,5610,3740)	-	50.9	11	10.3	0.8
neptune	(2996,9000,6000)	-	29.3	13	12.4	1.8
horse	(4009,12008,8000)	(2009,6008,4000)	53.4	7	27.5	2.1
julius	(2107,6213,4107)	-	13.1	6	4.4	0.7
camel	(1900,5694,3796)	-	55.5	6	24.8	0.4
rabbit	(1238,3678,2442)	-	7.8	9	2.7	0.3
hand	(1453,4326,2874)	-	7.4	7	3	0.3
bird	(567,1695,1130)	-	16.5	4	4.7	0.05
scarrow	(1802,5350,3548)	-	155	5	28.8	0.4
ET	(429,1281,854)	-	1.4	5	1.3	0.02
chess piece	(250,744,496)	-	0.93	3	0.57	0.01

Figures 1, 4, and 5 show examples of the hierarchical segmentation generated by our approach. In the dino-pet example of Figure 4(c)–(d), the body, the neck, and the head are recombined and split in different ways between the third and fourth level of the hierarchy, featuring the new kind of hierarchy of segments. It is also immediate to verify that the segments represent natural parts of the objects with small boundaries in concavity regions, following the minima rule.

Figure 6 shows that our approach is consistent to pose variations, segmenting the different poses of a horse into the same parts: e.g., head, legs, body, and tail. All horse models have the same resolution. More segmentation examples are in Figure 7.

In Figure 4(a)–(b), we compare our results to those presented in [LZ04], which also use k-means and eigenvectors of a matrix to segment meshes. Despite of similiar results, our method extracts more natural components and takes only 36 seconds to compute four levels of the hierarchy, while Liu and Zhang algorithm spends several minutes to compute only six segments. This discrepancy occurs because the matrix computed in [LZ04] is dense and measures the affinity between each pair of faces of the mesh, which costs $O(|F|^2 \log |F|)$ using Dijkstra's algorithm and makes the dimension of their matrix almost twice of ours.

6 Conclusion and Future Works

In this paper, we have presented a novel hierarchical part-type mesh segmentation algorithm. Our method is simple and totally intrinsic with respect to the properties of surfaces, without any constraint or salient metric. We have also introduced a new kind of hierarchy more natural to human perception, recombining and splitting segments in each level of the hierarchy.

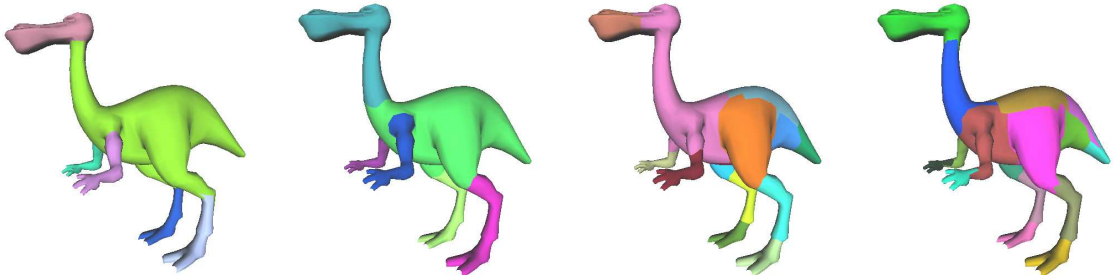

(a) [LZ04] (b) First level. (c) Third level. (d) Fourth level.

Figure 4: First row: comparision between [LZ04] and the first level of our method for the dino-pet model. Second row: example of our hierarchy of segments. Note how the body, the neck, and the head of the model are recombined and split in different parts.

Our approach is based on a new interpretation of the Laplace-Beltrami eigenfunctions (MHB) using Manifold learning theory. We define the MHB as a new embedding space for mesh vertices where Euclidian distances measure vertices proximity. We also use the Laplace-Beltrami spectrum to identify the correct number of segments that decomposes meshes hierarchically.

The complexity of our method is dominated by the eigen-problem solver. In our current implementation, we use ARPACK solver, which is superlinear on the number of segments and linear on the number of nonzeros elements of the matrix. Even though it has a close to linear asymptotic behavior, its cost increases drastically for meshes with more than a few thousands of vertices. As future work, we intend to use out-of-core computation as presented in [VL07] and to introduce multiresolution solvers over meshes. Then we will be able to segment dense meshes into long hierarchies.

Another direction for future work is the design of a new boundary refinement strategy to treat complex boundaries among many segments. We also plan to exploit user-defined priors to produce different types of segmentation, and adapt our method to segment point clouds from 3D laser scanners.

References

[AKM*06] ATTENE M., KATZ S., MORTARA M., PATANE G., SPAGNUOLO M., TAL A.: Mesh segmentation - a comparative study. In *Proc. of the IEEE Intl. Conf. on Shape Modeling and Applications* (2006), p. 7.

[BN02] BELKIN M., NIYOGI P.: Laplacian eigenmaps and spectral techniques for embedding and clustering. *Advances in Neural Information Processing Systems 1* (2002), 585–592.

[Chu97] CHUNG F. R. K.: *Spectral Graph Theory*. CBMS Regional Conference Series in Mathematics. 1997.

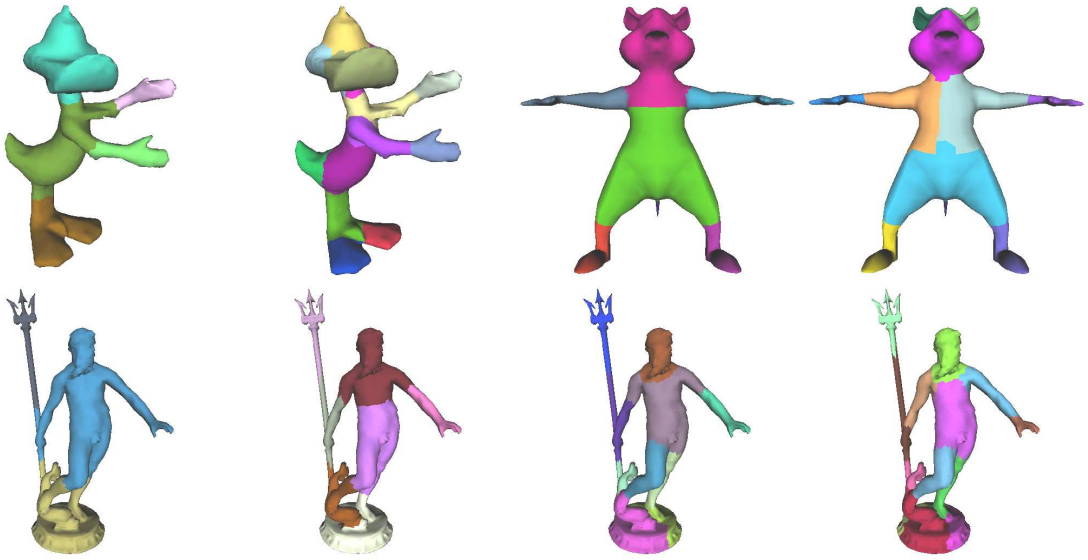

Figure 5: More examples of hierarchical segmentations.

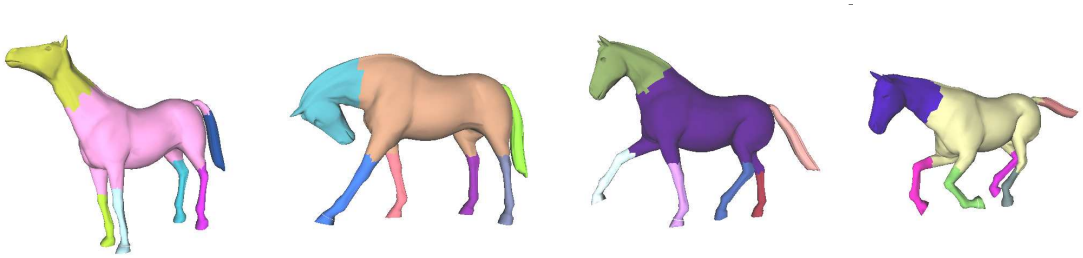

Figure 6: Consistent segmentations of various poses of a horse model.

[CL06] COIFMAN R., LAFON S.: Diffusion maps. *Applied and Comp. Harmonic Analysis: Special issue on Diffusion Maps and Wavelets 21* (2006), 5–30.

[DBG*06] DONG S., BREMER P.-T., GARLAND M., PASCUCCI V., HART J. C.: Spectral surface quadrangulation. *ACM Transactions on Graphics 25*, 3 (2006), 1057–1066.

[GY03] GU X., YAU S.-T.: Global conformal surface parameterization. In *Proc. of the Eurographics/ACM SIGGRAPH Symp. on Geometry processing* (2003), pp. 127–137.

[HS97] HOFFMAN D. D., SIGNH M.: Salience of visual parts. *Cognition 63* (1997), 29–78.

[JNT01] JAKOBSON D., NADIRASHVILI N., TOTH J.: Geometric properties of eigenfunctions. *Russian Math. Surveys 56*, 6 (2001), 1085–1105.

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

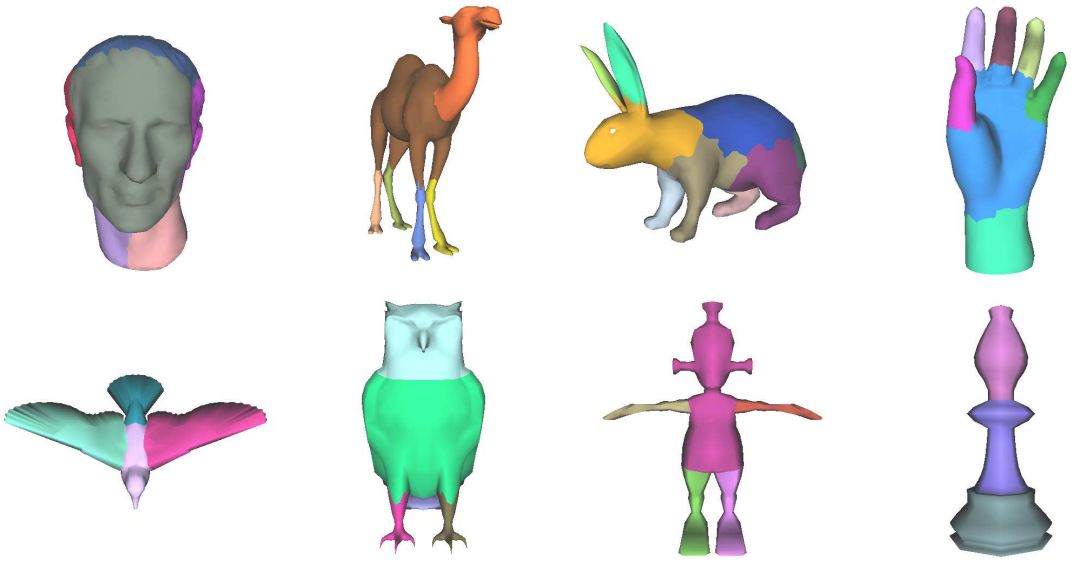

Figure 7: More segmentation results.

[KT03] KATZ S., TAL A.: Hierarchical mesh decomposition using fuzzy clustering and cuts. In *SIGGRAPH* (2003), pp. 954–961.

[Lev06] LEVY B.: Laplace-beltrami eigenfunctions towards an algorithm that "understands" geometry. In *Proc. of the IEEE Intl. Conf. on Shape Modeling and Applications* (2006), p. 13.

[Llo82] LLOYD S.: Least square quantization in pcm. *IEEE Transactions on Information Theory 28* (1982), 129–137.

[LLS*05] LEE Y., LEE S., SHAMIR A., COHEN-OR D., SEIDEL H.-P.: Mesh scissoring with minima rule and part salience. *Computer Aided Geometric Design 22*, 5 (2005), 444–465.

[LZ04] LIU R., ZHANG H.: Segmentation of 3D meshes through spectral clustering. In *Proc. of the Computer Graphics and Applications, 12th Pacific Conference* (2004), pp. 298–305.

[MDSB02] MEYER M., DESBRUN M., SCHRODER P., BARR A. H.: Discrete differential-geometry operators for triangulated 2-manifolds. In *VisMath* (2002).

[MS01] MAILA M., SHI J.: A random walks view of spectral segmentation. In *AI and STATISTICS* (2001).

[NISA06] NEALEN A., IGARASHI T., SORKINE O., ALEXA M.: Laplacian mesh optimization. In *Proc. of ACM GRAPHITE* (2006), pp. 381–389.

[PKA03] PAGE D. L., KOSCHAN A., ABIDI M.: Perception-based 3D triangle mesh segmentation using fast marching watersheds. In *Proc. of the Computer Vision and Pattern Recognition* (2003), vol. II, pp. 27–32.

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

[RS00] ROWEIS S. T., SAUL L. K.: Nonlinear dimensionality reduction by locally linear embedding. *Science 290*, 5500 (2000), 2323–2326.

[RWP06] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace-beltrami spectra as "Shape-DNA" of surfaces and solids. *Computer-Aided Design 38*, 4 (2006), 342–366.

[SCOS07] SHAMIR A., COHEN-OR D., , SHAPIRA L.: Consistent partitioning of meshes. 2007.

[Sha04] SHAMIR A.: A formulation of boundary mesh segmentation. In *Proc. of the 3D Data Processing, Visualization, and Transmission, 2nd Intl. Symp.* (2004), pp. 82–89.

[Sha06] SHAMIR A.: Segmentation and shape extraction of 3D boundary meshes. In *State-of-art Report: Eurographics* (2006), pp. 137–149.

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

[STK02] SHLAFMAN S., TAL A., KATZ S.: Metamorphosis of polyhedral surfaces using decomposition. *Computer Graphics Forum 21*, 3 (2002), 219–228.

[TACSD06] TONG Y., ALLIEZ P., COHEN-STEINER D., DESBRUN M.: Designing quadrangulations with discrete harmonic forms. In *Proc. of Eurographics/ACM SIGGRAPH Symp. on Geometry processing* (2006), pp. 201–210.

[Tau00] TAUBIN G.: Geometric signal processing on polygonal meshes. In *State-of-art Report: Eurographics* (2000).

[TdSL00] TENENBAUM J. B., DE SILVA V., LANGFORD J. C.: A global geometric framework for nonlinear dimensionality reduction. *Science 290*, 5500 (2000), 2319–2323.

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

[ZL05] ZHANG H., LIU R.: Mesh segmentation via recursive and visually salient spectral cuts. In *Proc. of Vision, Modeling, and Visualization* (2005), pp. 429–436.

[ZSGS04] ZHOU K., SYNDER J., GUO B., SHUM H.-Y.: Iso-charts: stretch-driven mesh parameterization using spectral analysis. In *Proc. of Eurographics/ACM SIGGRAPH Symp. on Geometry processing* (2004), pp. 45–54.