



Shape Matching Based on Diffusion Embedding and on Mutual Isometric Consistency

Avinash Sharma, Radu Horaud

► To cite this version:

Avinash Sharma, Radu Horaud. Shape Matching Based on Diffusion Embedding and on Mutual Isometric Consistency. NORDIA 2010 - Workshop on Nonrigid Shape Analysis and Deformable Image Alignment, Jun 2010, San Francisco, United States. IEEE, pp.29-36, 2010, <10.1109/CVPRW.2010.5543278>. <inria-00549406>

HAL Id: inria-00549406

<https://hal.inria.fr/inria-00549406>

Submitted on 21 Dec 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Shape Matching Based on Diffusion Embedding and on Mutual Isometric Consistency

Avinash Sharma
INRIA Rhône-Alpes
FRANCE

Avinash.Sharma@inria.fr

Radu Horaud
INRIA Rhône-Alpes
FRANCE

Radu.Horaud@inria.fr

Abstract

We address the problem of matching two 3D shapes by representing them using the eigenvalues and eigenvectors of the discrete diffusion operator. This provides a representation framework useful for both scale-space shape descriptors and shape comparisons. We formally introduce a canonical diffusion embedding based on the combinatorial Laplacian; we reveal some interesting properties and we propose a unit hypersphere normalization of this embedding. We also propose a practical algorithm that seeks the largest set of mutually consistent point-to-point matches between two shapes based on isometric consistency between the two embeddings. We illustrate our method with several examples of matching shapes at various scales.

1. Introduction

Recent advancement in technology has enabled easy acquisition and modeling of visual data, leading to a vast amount of 3D shape data. A large part of the existing databases uses *mesh* representations which model a shape as a graph structure where vertices correspond to points uniformly sampled from a continuous surface lying on a Riemannian manifold, and edges describe the local topology of the surface. Hence shape analysis heavily relies on local and global properties associated with the underlying manifold.

Analyzing 3D data in a single framework is still a challenging task considering the large variability of the data gathered with different acquisition devices. Matching two shapes is one such challenging shape analysis task. Most of the previous attempts of shape matching can be broadly categorized as *extrinsic* or *intrinsic* approaches depending on how they analyze

the properties of the underlying manifold. Intrinsic approaches are a natural choice for finding dense correspondences between articulated shapes, as they embed the shape in some canonical domain which preserves some important properties of the manifold, *e.g.*, geodesics and angles [4, 29, 16, 31, 22, 24, 18, 10].

Interestingly, mesh representation also enables adaptation of well established graph matching algorithms that allow to combine intrinsic and extrinsic (appearance) features and were previously successfully applied for images feature matching [17, 11, 27, 30, 20]. Some more recent approaches apply hierarchical matching to find dense correspondences [15, 32, 25]. But many of them suffer with the problem of either computational intractability or a lack of proper distance metric as Euclidean metric is not directly applicable to non-rigid shapes. A recent benchmarking of shape matching methods was performed [7]. According to the results summarized in [7] one of the best performing shape matching algorithms is [22] that uses Laplacian embedding and point registration based on a variant of the expectation-maximization (EM) algorithm [14] to register the two sets of points associated with the Laplacian embeddings of the two shapes. The two embedded representations need to be correctly aligned in order to properly initialize the EM algorithm. This alignment is critical for the spectral matching methods because the two isometric embeddings are defined up to the signs and orderings of the eigenvalues of the Laplacian. The eigenvector histogram proposed in [22] works well, provided that the embeddings of the two shapes have approximately the same statistical properties. Therefore, it would be useful to investigate alternative solutions based on some other intrinsic shape properties that are more robust to discrepancies in the statistical properties of the two embedded shapes.

Another important criterion for shape matching is

the proper detection of interest points on the two shapes. A benchmark for detection and description of features on 3D shapes has been recently presented in [6]. The benchmark concludes that the heat (or diffusion) kernel based features are robust and repeatably detected. The diffusion kernel is the fundamental solution to the diffusion equation and can be represented in terms of the eigenvalues and eigenfunctions of the Laplacian operator. This leads to an important intrinsic shape representation known as *diffusion-kernel embedding* [1]. The diffusion kernel also allows the definition of the diffusion distance which is an interesting alternative to the geodesic distance and which can be used in various ways such as: building invariant shape signatures, i.e., distance distributions [21, 8], computing differential properties of meshes and of point clouds [19], or comparing shapes [5, 23]. One motivation for using diffusion signatures and diffusion metrics is that they are relatively robust to changes in topology. Nevertheless, we remark that they are quite sensitive to the choice of the scale (or time) parameter. At small scales only local properties are captured while at large scales all points on the shape are at equal distances. When the scale parameter tends to infinity the heat kernel collapses to zero, unless it is properly normalized [1, 26]. Also it is important to use some kind of normalized representation for a proper comparison of two shapes such as to avoid issues with sampling and changes in size.

In this paper we propose to find a sparse set of matches between two shapes using the intrinsic properties of the diffusion embeddings of the shapes. First we analyze in detail the *canonical* diffusion embedding that is based on the combinatorial graph Laplacian, we show that it has interesting statistical properties, and we propose a unit hypersphere intrinsic representation. This allows us to devise a simple method to estimate a global scale factor between the two embeddings whenever the two shapes to be matched correspond to different samplings. Then we present a method to detect feature (or interest) points on shapes at different scales using these diffusion embedding representations. Next, we describe a sparse shape matching method. We propose a method that finds the largest set of mutually consistent features between the two shapes. Finally, we present sparse shape matching results along with a discussion on the choice of the scale parameter for the feature detection and matching.

2. Diffusion Embeddings

In this section we will introduce different intrinsic representations of shapes using the diffusion (heat) op-

erator. Analysis of diffusion processes on graphs will provide many time-dependent and time-invariant intrinsic representations of shapes as well as different spectral distance metrics applicable to them.

2.1. Heat Diffusion on an Undirected Graph

In the case of a undirected graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ where $\mathcal{V}(\mathcal{G}) = \{v_1, \dots, v_n\}$ is the vertex set, $\mathcal{E}(\mathcal{G}) = \{e_{ij}\}$ is the edge set, the heat or diffusion equation writes:

$$\left(\frac{\partial}{\partial t} + \mathbf{L}\right) \mathbf{f}(t) = 0 \quad (1)$$

where $\mathbf{L}$ denotes the combinatorial graph Laplacian (section 2.2) [9, 28] and $\mathbf{f}(t) = (f(v_1; t) \dots f(v_n; t))^T$ is a $n \times 1$ vector of real-valued functions: Each such function associates a real number to a graph vertex $f : \mathcal{V} \rightarrow \mathbb{R}$ which can be interpreted as a set of *heat values* at each graph vertex at time $t > 0$. We denote by $\mathbf{f}(0)$ the initial heat distribution. Notice that both the Laplacian and heat operators map real valued functions, i.e. f , defined on the graph onto other such functions. The fundamental solution of the discrete diffusion equation is given by:

$$\mathbf{H}(t) = e^{-t\mathbf{L}} \quad (2)$$

where the exponential of a matrix is defined by:

$$e^{\mathbf{A}} = \sum_{k=0}^{\infty} \frac{\mathbf{A}^k}{k!} \quad (3)$$

2.2. The Combinatorial Graph Laplacian Operator

A combinatorial Laplacian (also known as *un-normalized* Laplacian) of a undirected graph $\mathcal{G}$ can be written as :

$$\mathbf{L} = \mathbf{D} - \mathbf{\Omega} \quad (4)$$

where $\mathbf{\Omega}$ is the $n \times n$ *weighted adjacency matrix* of $\mathcal{G}$ with positive entries and $\mathbf{D}$ is the *degree matrix* $\mathbf{D} = \text{Diag}[d_1 \dots d_i \dots d_n]$ where $d_i = \sum_{i \sim j} \Omega_{ij}$ ($i \sim j$ denotes the set of vertices v_j which are adjacent to v_i). The spectral decomposition of the combinatorial Laplacian is $\mathbf{L} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$. The eigenvector $\mathbf{u}_k = (u_{1k} \dots u_{ik} \dots u_{nk})^T$ is a column of $\mathbf{U}$. From (4) we obtain $\lambda_1 = 0$ and $\mathbf{u}_1 = \mathbf{1}$, hence $\mathbf{u}_{k \geq 2}^T \mathbf{1} = 0$. Using $\mathbf{u}_k^T \mathbf{u}_k = 1$ we obtain the following properties:

$$\sum_{i=1}^n u_{ik} = 0, \quad 2 \leq k \leq n \quad (5)$$

$$|u_{ik}| < 1, \quad 1 \leq i \leq n, 2 \leq k \leq n \quad (6)$$

2.3. The Combinatorial Diffusion Operator

Using the spectral decomposition of the Laplacian matrix, we further obtain the spectral decomposition of the graph diffusion operator:

$$\mathbf{H}(t) = \mathbf{U}e^{-t\mathbf{\Lambda}}\mathbf{U}^\top \quad (7)$$

with $e^{-t\mathbf{\Lambda}} = \text{Diag}[e^{-t\lambda_1} \dots e^{-t\lambda_n}]$. Notice that from the properties of the eigenvalues of the Laplacian we obtain

$$1 = \phi_1 > \phi_2 \geq \dots \phi_n > 0$$

Each entry $h(i, j; t)$ of the diffusion operator is a *diffusion kernel function* that estimates the amount of heat at vertex v_i and at time t , starting from a heat distribution located at a single vertex v_j . From above and by omitting the constant eigenvector associated $\lambda_1 = 0$, we obtain:

$$h(i, j; t) = \sum_{k=2}^n e^{-t\lambda_k} u_{ik} u_{jk} \quad (8)$$

The diagonal terms of the heat operator correspond to the *auto-diffusion function* defined at each vertex, i.e., $h(i, i; t)$.

2.4. The Diffusion-Kernel Distance

By interpreting (8) as an inner product we obtain a close-form expression for the distance between two graph vertices, which can be interpreted as an Euclidean metric:

$$\begin{aligned} d_t^2(i, j) &= h(i, i; t) + h(j, j; t) - 2h(i, j; t) \\ &= \sum_{k=2}^n \left(e^{-\frac{1}{2}t\lambda_k} (u_{ik} - u_{jk}) \right)^2 \end{aligned} \quad (9)$$

2.5. Canonical Diffusion Embedding

We start by defining the *canonical embedding* associated with the combinatorial diffusion operator which can be decomposed as:

$$\mathbf{H}(t) = \left(\mathbf{U}e^{-\frac{1}{2}t\mathbf{\Lambda}} \right) \left(\mathbf{U}e^{-\frac{1}{2}t\mathbf{\Lambda}} \right)^\top \quad (10)$$

The rows of the $n \times n$ matrix $\mathbf{U}e^{-t\mathbf{\Lambda}/2}$ can be viewed as the coordinates of the graph's vertices in some *feature space*. Indeed, one can define a mapping $\mathbf{x} : \mathcal{V} \rightarrow \mathbb{R}^{n-1}$ (omitting the constant vector corresponding to $\lambda_1 = 0$) to obtain:

$$\mathbf{x}_i = \left(e^{-\frac{1}{2}t\lambda_2} u_{i2} \quad \dots \quad e^{-\frac{1}{2}t\lambda_n} u_{in} \right)^\top \quad (11)$$

From (6) we obtain bounds on the coordinate values for all $k > 1$:

$$-1 < -e^{-\frac{1}{2}t\lambda_k} < x_{ik} < e^{-\frac{1}{2}t\lambda_k} < 1 \quad (12)$$

It is practical to use only a subset of the operator's eigenvectors, namely the eigenvectors corresponding to the K smallest non null eigenvalues of the Laplacian, $K \leq n-1$. Therefore we obtain the following representation, where $\mathbf{U}_K$ is a $n \times K$ sub-matrix of $\mathbf{U}$ formed with K column vectors $\mathbf{u}_2, \dots, \mathbf{u}_{K+1}$, $\mathbf{U}_K^\top \mathbf{U}_K = \mathbf{I}_K$, and $\mathbf{\Lambda}_K = \text{Diag}[\lambda_2, \dots, \lambda_{K+1}]$ is a $K \times K$ matrix:

$$\mathbf{X} = [\mathbf{x}_1 \dots \mathbf{x}_i \dots \mathbf{x}_n] = \left(\mathbf{U}_K e^{-\frac{1}{2}t\mathbf{\Lambda}_K} \right)^\top \quad (13)$$

Therefore, the *canonical embedded* coordinates of a graph vertex v_i correspond to the i -th column of the $K \times n$ matrix $\mathbf{X}$. It is worthwhile to notice that in this case the embedded (or feature space) is an Euclidean space. Hence, there is a straightforward interpretation of the heat-kernel signature (HKS) [26] (or of the auto-diffusion function (ADF) [13]), namely:

$$\|\mathbf{x}_i\|^2 = h(i, i; t) \quad (14)$$

2.6. Maximum-variance embedding

In this paragraph we show a useful property of the embedding of a graph based on the spectrum of the combinatorial heat operator, namely that the eigenvectors of $\mathbf{H}$ are the directions of maximum variance of the the embedded points $\mathbf{X}$.

Let's compute the $K \times K$ covariance matrix of $\mathbf{X}$:

$$\mathbf{S}_X = \frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^\top \quad (15)$$

Using (13) and (5), i.e., the entries of each row of $\mathbf{U}_K^\top$ sum to zero, we obtain:

$$\begin{aligned} \bar{\mathbf{x}} &= \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \\ &= \frac{1}{n} e^{-\frac{1}{2}t\mathbf{\Lambda}_K} \begin{pmatrix} \sum_{i=1}^n u_{i2} \\ \vdots \\ \sum_{i=1}^n u_{i, K+1} \end{pmatrix} \\ &= \begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix} \end{aligned} \quad (16)$$

Hence, the covariance:

$$\begin{aligned}
\mathbf{S}_X &= \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top \\
&= \frac{1}{n} \mathbf{X} \mathbf{X}^\top \\
&= \frac{1}{n} \left(\mathbf{U}_K e^{-\frac{1}{2} t \mathbf{\Lambda}_K} \right)^\top \left(\mathbf{U}_K e^{-\frac{1}{2} t \mathbf{\Lambda}_K} \right) \\
&= \frac{1}{n} e^{-t \mathbf{\Lambda}_K}
\end{aligned} \tag{17}$$

is a diagonal form with entries $e^{-t \lambda_2}/n \geq \dots \geq e^{-t \lambda_{K+1}}/n$. Therefore, the non null eigenvectors of the combinatorial heat-kernel matrix correspond to the directions of maximum variance [2]. Moreover, using this result in conjunction with (12) we obtain the following proposition (see [12] for an alternative proof):

Proposition 1. *The set of points $\mathbf{X}$, corresponding to the spectral embedding of the combinatorial heat-kernel matrix, and defined by (13) admits the eigenvectors of $\mathbf{H}$ defined by (7) as the axes of maximum variance. The embedded representation is strictly bounded by an ellipsoid whose center coincides with the origin of the axes and with eccentricities $e^{-t \lambda_2}, \dots, e^{-t \lambda_{K+1}}$.*

2.7. Alignment of Two Different Embeddings

Consider the task of aligning the embedded representations of two different shapes. The covariance matrix of an embedding, that we just characterized, provides a principled framework to align two distributions of points drawn from two different manifolds. Indeed, one prerequisite to align two distributions of points is to ensure that they occupy the same volume in space. This can be easily done using the covariance matrix. Let $\mathbf{S}_X$ and $\mathbf{S}_{X'}$ be the covariance matrices of two different embedding $\mathbf{X}$ and $\mathbf{X}'$ with respectively n and n' points. The volume of $\mathbf{X}$ can be computed using the determinant of the covariance matrix. Hence, the following constraint:

$$\det(\mathbf{S}_X) = \det(\mathbf{S}_{X'}) \tag{18}$$

Using (17) one can easily derive the following relationship:

$$t' \text{Trace}(\mathbf{\Lambda}'_K) = t \text{Trace}(\mathbf{\Lambda}_K) + K \log n/n' \tag{19}$$

This allows to choose the parameter t' associated with $\mathbf{X}'$ as a function of the free parameter t . Notice that if the two sets of points have very different cardinalities, namely if $n \ll n'$, t' may take negative values. In practice one should choose t such that $t' > 0$.

2.8. Unit Hyper-Sphere Normalization

One disadvantage of the standard embeddings using diffusion is that, when t tends to infinity, the embedded representation collapses to 0. In order to avoid the collapsing on $(0 \dots 0)$ we can "re-normalize" the embedding such that the embedded vertex coordinates lie on a unit sphere of dimension K , which yields:

$$\tilde{\mathbf{x}}_i = \frac{\mathbf{x}_i}{\|\mathbf{x}_i\|} \tag{20}$$

In more detail, the k -th coordinate of $\tilde{\mathbf{x}}_i$ writes as:

$$\tilde{x}_{ik}(t) = \frac{(e^{-t \lambda_k} u_{ik}^2)^{1/2}}{\left(\sum_{l=2}^{K+1} e^{-t \lambda_l} u_{il}^2 \right)^{1/2}} \tag{21}$$

The diffusion distance is equivalent to the geodesic distance on a unit sphere:

$$d_S(i, j; t) = \arccos \frac{h(i, j; t)}{h(i, i; t)^{1/2} h(j, j; t)^{1/2}} \tag{22}$$

2.9. Scale-Space Detection of Features Points

Here we present a method that uses diffusion embeddings to obtain a scale-space representation of 3D shapes in order to detect salient feature points. We detect salient features on shapes by finding the local extrema of (14) using the two ring neighborhood at each vertex. When analyzing two shapes, we choose a time scale t and compute t' using the method proposed in section 2.7.

In figure 1 first row, we show the modulo of (14) computed at different time scales. We use 130 eigenvalues and corresponding eigenvectors of the combinatorial Laplacian to approximate the heat kernel at each vertex. In the second row of figure 1, we show the detected local extrema of the modulo of (14) (with the colored spheres where red corresponds to maxima and blue corresponds to minima) as the function of t . At small value of t we capture local geometry and hence large number of extrema while at large value of t we capture global geometry and hence small number of extrema. We observe that the detection of maxima is more stable in case of real meshes and so we propose to use these maxima to select a subset of points onto the given shape. For two different shapes with embedded representation $\mathbf{X}$ and $\mathbf{X}'$, we detect two subsets of points $\mathbf{Y} \subset \mathbf{X}$ and $\mathbf{Y}' \subset \mathbf{X}'$.

3. Mutually Consistent Matches

In this section we address the problem of matching two sets of points detected onto two different shapes in




Figure 1. *Modulo* of (14) plotted over the shape (*first row*) and corresponding local extrema (*second row*) for different choice of t parameter. (Red and blue colors corresponds to higher and lower function values.)

section 2.9. We formulate the point-to-point matching problem in the framework of *association-graph*. The *association-graph* models the pairwise consistency of all possible point-to-point matches between two set of points [3]. Each node of the *association-graph* represents a point-to-point match and every (positive weighted) edge represents the mutual isometric consistency between two point-to-point matches. Two mutually consistent graph nodes are strongly connected (i.e. have high edge weights) if the corresponding points onto the two shapes have similar spectral distances.

Thus, a set of mutually consistent point-to-point matches can be obtained by finding a subset of strongly




Figure 2. Association graph structure: Each node represent a point-to-point matching and every edge represents the mutual consistency between two nodes measured as positive weight. The edges with low score are depicted with dashed line as opposed to the ones with high consistency score.

connected nodes of the *association-graph*. This is an instance of largest maximal clique problem which is a *NP hard* problem. There are many algorithms that find approximate solutions by enforcing different constraints. In [17], a spectral relaxation and one-to-one matching constraint were employed to find such solution.

We propose to find an approximate solution for the point-to-point matching problem by iteratively searching the graph node with the largest degree (i.e. a node strongly connected to other nodes). This corresponds to selecting a point-to-point match which is the most consistent with the other matches. In each iteration, after selecting a node, we eliminate all the other nodes from our search space that violates one-to-one matching constraint *w.r.t.* the selected node. In the end, every node is either selected or eliminated and we obtain a subset of mutually consistent point-to-point matches as the set of selected nodes. It can be easily shown that our algorithm is equivalent to finding a solution similar to [17] but using the principal eigenvector corresponding to the largest eigenvalue (in this case 1) of the *normalized* weighted adjacency matrix of the *association-graph*. Our method is efficient in the sense that we do not require an eigen decomposition step.

3.1. Association Graph Construction

We find point-to-point matching between two sets of points $\mathbf{Y}$ and $\mathbf{Y}'$ in section 2.9 For the simplicity of notation, let's refer these subsets using indexes and we can write : $\mathbf{Y} = \{1, \dots, i, \dots, j, \dots, a\}$ and $\mathbf{Y}' = \{1, \dots, i', \dots, j', \dots, b\}$. The *association-graph* will have a set of nodes as $\{(i, j')\} \forall i \in \mathbf{Y}$ and $\forall j' \in \mathbf{Y}'$. Each node (i, j') represents a point-to-point match between $i \in \mathbf{Y}$ and $j' \in \mathbf{Y}'$.

Let $p = (i, i')$ and $q = (j, j')$ be two *association-graph* nodes. Let $\mathcal{W} = \{w_{pq}\}$ be the weighted adjacency matrix of the *association-graph* where w_{pq} be the edge weight between node p and q and $w_{pp} = 0$. If the two nodes violated the one-to-one matching constraint (i.e. a single point on a shape matches to multiple points on another shape), we set corresponding edge weight to a very small positive value ϵ . For all the other edges we compute a positive weight (w_{pq}) defined strictly between 0 and 1. We call it mutual consistency score between two nodes and this score is high (close to 1) when the spectral distance between i and j on the first shape is similar to the spectral distance between i' and j' on the second shape.

In order to compute the edge weights, we need to compare two distance measures computed on different shapes. We propose to align the two different shapes by representing them with unit hyper-sphere normalized embeddings (section 2.8) and computing different time scales t and t' for each of them (section 2.7). The edge weights are given by:

$$w_{pq} = \exp\left(-\frac{1}{\sigma^2} \left(\frac{d_{\mathcal{S}}(i, j; t) - d_{\mathcal{S}}(i', j'; t')}{d_{\mathcal{S}}(i, j; t) + d_{\mathcal{S}}(i', j'; t')}\right)^2\right) \quad (23)$$

where $d_{\mathcal{S}}$ is computed by (22) and t and t' are obtained by (19). The sum of distance term in the denominator ensures that the score is not biased for the small distances and the parameter σ controls the sensitivity of the consistency score.

Figure 2 shows the construction of *association-graph*. We consider *association-graph* as a complete graph but for the simplicity of visualization, we have shown only the strong edges and some weak edges. A subset of strongly connected nodes shown in the figure represent mutually consistent set of point-to-point matches that are weakly connected to the other nodes.

3.2. Mutually Consistent Matching Algorithm

Let $\mathcal{D}$ be the degree matrix of the *association-graph*. Then according to [28], the principal eigenvector of the normalized weighted adjacency matrix $\mathcal{D}^{-\frac{1}{2}}\mathcal{W}\mathcal{D}^{-\frac{1}{2}}$ is the vector $\mathbf{z} = \mathcal{D}^{\frac{1}{2}}\mathbf{1}$, which implies $\mathbf{z}(p) = \sqrt{d_p}$ for each node p . The vector $\mathbf{z}$ is the null space of *normalized* Laplacian of the *association-graph*.

Now, let's define a vector $\mathbf{y}$ as the $n \times 1$ indicator vector such that $\mathbf{y}(p) = 1$ if p is labeled as a consistent point-to-point match and $\mathbf{y}(p) = -1$ if p is labeled as an inconsistent point-to-point match and $\mathbf{y}(p) = 0$ otherwise. We initialize $\mathbf{y}$ with the zero vector. We summarize our algorithm as follows:

1. Find $p = \arg \max_l \mathbf{z}(l)$ while considering only those nodes for which $\mathbf{y}(l) = 0$ (i.e. find unlabeled node with the largest degree).
2. Label p as a consistent point-to-point match by setting $\mathbf{y}(p) = 1$.
3. Set $\mathbf{y}(q) = -1$ for all the q 's that violate one-to-one matching constraint *w.r.t.* the selected match p .
4. If $\mathbf{y}$ has more zero elements then go to step 1 else stop and return $\mathbf{y}$.

However, our method does not perform well if the *association-graph* does not have a single strongly connected cluster. This can happen if the detected feature points (i.e. extrema of 14) are uniformly distributed on shapes causing similar consistency scores for all the *association-graph* edges or due to existence of large number of symmetry in two shapes.

4. Results

We show the point-to-point matching results obtained with our algorithm on synthetic and real meshes. We use the diffusion embedding to perform scale-scale detection of feature points on two shapes (section 2.9). We compute the unit hyper-sphere normalized embedding representation (section 2.8) for every shape using the first 130 eigenvalues and eigenvectors of the diffusion operator. We align two embeddings to find appropriate time scale t and t' for each shape (see section 2.9), while matching two different shapes. We set parameter σ to 0.5 while computing the *association-graph* edge weights (23) for all the experiments.

We perform two sets of experiments. In the first experiment we show the sensitivity of our method to the choice of time scale parameter. Figure 3 shows matching results for two synthetic horse meshes (approximately 19,000 vertices) and two real human meshes (approximately 1,500 vertices). In figure 3(a-c) we choose $t = t' = 200$ for the feature point detection. While in figure 3(d-f) we choose $t = 30$ and t' is almost same as t the two shapes have similar sampling and almost the same number of vertices. Next we choose several different scales for matching. In case of synthetic meshes, where two shapes are strictly isometric (same eigenvalues), t and t' are same (see figure 3(a-c)). For the second row of figure 3, we get almost the same time scale t and t' as the two shapes have similar size and sampling.




Figure 3. Mutually consistent matches obtained at different time scales. Feature points were detected at fixed time scale for each row. Here t_{mat} and t'_{mat} is the time scale at which feature points were matched.

Results in figure 3(c,f) can be understood by the fact that at a very large time scale features points have a similar representation, so the *association-graph* is not able to capture the correct mutual consistency between matching pairs, resulting in incorrect matching. In case of smaller scales, such as $t = 0$ in figure 3(d) the diffusion embedding reduces to Laplacian embedding. In this case diffusion distance between two feature points detected at large scale is dominated by local surface noise (in real meshes) instead of global shape structure, and is not useful to get correct consistency score. This leads to wrong matching results. However, at small scale our method performs well if the two shape are strictly isometric (see figure 3(a)).

In the second experiment we show matching results between two shapes with different sampling and size. In figure 4(a,b), we show matching result on two different real meshes with approximately 16,000 and 1,500 vertices. As we can see, we obtain significantly different time scales for the two shapes. In figure 4(c), we show robustness of our method to significant topological noise in the real mesh.

5. Conclusion

We presented discussion on shape representation based on discrete diffusion operator, a normalized embedding and a practical way to find appropriate times scale for two different shapes with their application for

sparse shape registration. As part of future work it will be interesting to combine matching obtained at multiple scales. The problem of intrinsic shape symmetry can be addressed by including complementary information like texture and other extrinsic features while defining the *association-graph*.

References

- [1] P. Bérard, G. Besson, and S. Gallot. Embedding Riemannian manifolds by their heat kernel. *Geometric and Functional Analysis*, 4(4):373–398, 1994. 2
- [2] C. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006. 4
- [3] R. C. Bolles and R. A. Cain. Recognizing and locating partially visible objects: The local-feature-focus method. *The International Journal of Robotics Research*, 1:57–82, 1982. 5
- [4] A. Bronstein, M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proc. of National Academy of Sciences*, 2006. 1
- [5] A. Bronstein, M. Bronstein, M. Mahmoudi, R. Kimmel, and G. Sapiro. A Gromov-Hausdorff framework with diffusion geometry for topologically robust non-rigid shape alignment. *IJCV*, 2009. 2
- [6] A. M. Bronstein, M. M. Bronstein, B. Bustos, U. Castellani, M. Crisani, B. Falcidieno, L. J. Guibas, I. Kokkinos, V. Murino, M. M. Ovsjanikov, G. Patané, I. Sipiran, M. Spagnuolo, and J. Sun. Shrec 2010: robust feature detection and description benchmark. In




Figure 4. Mutually consistent matches. Feature detection and matching was performed at the same scale. We first chose t for the first shape and compute the t' for the other shape using (19)

- Eurographics Workshop on 3D Object Retrieval*, 2010. 2
- [7] A. M. Bronstein, M. M. Bronstein, U. Castellani, A. Dubrovina, L. J. Guibas, R. P. Horaud, R. Kimmel, D. Knossow, E. v. Lavante, M. D., M. Ovsjanikov, and A. Sharma. Shrec 2010: robust correspondence benchmark. In *Eurographics Workshop on 3D Object Retrieval*, 2010. 1
- [8] M. Bronstein and I. Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *CVPR*, 2010. 2
- [9] F. Chung. *Spectral Graph Theory*. American Mathematical Society, 1997. 2
- [10] A. Dubrovina and R. Kimmel. Matching shapes by eigendecomposition of the Laplace-Beltrami operator. In *3DPVT*, 2010. 1
- [11] O. Duchenne, F. Bach, I. Kweon, and J. Ponce. A tensor based algorithm for high order graph matching. In *CVPR*, 2009. 1
- [12] F. Fouss, A. Pirotte, J. Renders, and M. Saerens. Random-walk computation of similarities between nodes of a graph, with application to collaborative recommendation. *IEEE KDE*, 19(3):355–369, 2007. 4
- [13] K. Gebal, J. A. Baerentzen, H. Aanaes, and R. Larsen. Shape analysis using the auto diffusion function. In *SGP*, 2009. 3
- [14] R. Horaud, F. Forbes, M. Yguel, G. Dewaele, and J. Zhang. Rigid and articulated point registration with expectation conditional maximization. *PAMI*, 2010. 1
- [15] Q. Huang, B. Adams, M. Wicke, and L. J. Guibas. Non-rigid registration under isometric deformations. *Computer Graphics Forum*, 27(5):1449–1457, 2008. 1
- [16] V. Jain, H. Zhang, and O. van Kaick. Non-rigid spectral correspondence of triangle meshes. *International Journal of Shape Modeling*, 2007. 1
- [17] M. Leordeanu and M. Hebert. A spectral technique for correspondence problems using pairwise constraints. In *ICCV*, 2005. 1, 5
- [18] Y. Lipman and T. Funkhouser. Mobius voting for surface correspondence. *SIGGRAPH*, 28(3), 2009. 1
- [19] C. Luo, I. Safa, and Y. Wang. Approximating gradients for meshes and point clouds via diffusion metric. In *SGP*, 2009. 2
- [20] J. Maciel and J. P. Costeira. A global solution to sparse correspondence problems. *PAMI*, 25:187–199, 2003. 1
- [21] M. Mahmoudi and G. Sapiro. Three-dimensional point cloud recognition via distributions of geometric distances. *Graphical Models*, 71:22–31, 2009. 2
- [22] D. Mateus, R. Horaud, D. Knossow, F. Cuzzolin, and E. Boyer. Articulated shape matching using Laplacian eigenfunctions and unsupervised point registration. In *CVPR*, 2008. 1
- [23] F. Memoli. Spectral Gromov-Wasserstein distances for shape matching. In *Proc. of NORDIA Workshop*, 2009. 2
- [24] M. R. Ruggeri, G. Patané, M. Spagnuolo, and D. Saupé. Spectral-driven isometry-invariant matching of 3d shapes. *IJCV*, 2009. 1
- [25] Y. Sahillioglu and Y. Yemez. 3d shape correspondence by isometry-driven greedy optimization. In *CVPR*, 2010. 1
- [26] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *SGP*, 2009. 2, 3
- [27] L. Torresani, V. Kolmogorov, and C. Rother. Feature correspondence via graph matching : Models and global optimization. In *ECCV*, 2008. 1
- [28] U. von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007. 2, 6
- [29] S. Wang, Y. Wang, M. Jin, X. Gu, D. Samaras, and P. Huang. Conformal geometry and its application on 3d shape matching. *PAMI*, 29(7):1209–1220, 2007. 1
- [30] R. Zass and A. Shashua. Probabilistic graph and hypergraph matching. In *CVPR*, 2008. 1
- [31] W. Zeng, Y. Zeng, Y. Wang, X. Yin, X. Gu, and D. Samras. 3d non-rigid surface matching and registration based on holomorphic differentials. In *ECCV*, 2008. 1
- [32] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samras, and N. Paragios. Dense non-rigid surface registration using high order graph matching. In *CVPR*, 2010. 1