3D Laplacian Pyramid Signature

Kaimo Hu and Yi Fang

Electrical and Computer Engineering
New York University in Abu Dhabi, Abu Dhabi,UAE

Abstract. We introduce a simple and effective point descriptor, called 3D Laplacian Pyramid Signature (3DLPS), by extending and adapting the Laplacian Pyramid defined in 2D images to 3D shapes. The signature is represented as a high-dimensional feature vector recording the magnitudes of mean curvatures, which are captured through sequentially applying Laplacian of Gaussian (LOG) operators on each vertex of 3D shapes. We show that 3DLPS organizes the intrinsic geometry information concisely, while possessing high sensitivity and specificity. Compared with existing point signatures, 3DLPS is robust and easy to compute, yet captures enough information embedded in the shape. We describe how 3DLPS may potentially benefit the applications involved in shape analysis, and especially demonstrate how to incorporate it in point correspondence detection, best view selection and automatic mesh segmentation. Experiments across a collection of shapes have verified its effectiveness.

1 Introduction

Signatures play an important role in shape analysis. On one hand, they are concise presentations of the model and are easily commeasurable; on the other hand, a good signature preserves much geometric information of the shape.

Current signatures defined on 3D shapes can be classified into *shape signature* (also referred to as shape descriptor) and *point signature*. Shape signatures are usually defined on a whole 3D model or a partial model, while point signatures are defined on each vertex of the mesh and capture the local or global shape characteristics from the perspective of that vertex. Since point signatures possess more local information, and sometimes can be easily converted to shape signature by calculating their statistic distributions [1], they have wider applications than shape signatures, such as point clouds registration, mesh saliency detection and mesh segmentation.

Early work on designing point signatures mainly focused on characterising the spatial geometric information of a shape [2–5]. While performing well in capturing the spacial geometrical attributes of models ranging from local to global, they are incompetent to describe the intrinsic frequency domain information embedded in a 3D model, which is sometimes essential for model analysis and understanding. Recently, some new point signatures have been proposed to reflect the frequency domain information [6–9]. However, they usually suffer from heavy calculations (see the related work in Sec. 2 for detail).

Inspired by the successful applications of Laplacian Pyramid operator applied in image processing [10], we introduce a simple yet effective point descriptor: 3D Laplacian Pyramid Signature (3DLPS), based on the evaluation of Laplacian Pyramid after applying a series of Laplacian of Gaussian (LOG) operators at each vertex on a 3D model. Every vertex on the meshed surface is then characterized with a high-dimensional feature vector that measures the magnitudes of mean curvatures, which are captured from these sequentially applied LOG operators. We show that 3DLPS has the following desirable properties: 1) it organizes the intrinsic frequency information of a shape in an efficient, multi-scale way; 2) it shows high sensitivity and specificity; and 3) it is insensitive to articulated objects. Due to these nice properties, 3DLPS has the potential to benefit many applications involved in shape analysis, including point correspondence detection, best view selection and automatic mesh segmentation. We have illustrated these across a collection of shapes. In summary, our main contributions are:

- The 3D Laplacian Pyramid Signature is introduced to characterize the geometric information embedded in the shape in a concise way.
- The desirable properties of 3DLPS are explained, which ensure the effectiveness of potential applications in shape analysis.
- The applications of point correspondence detection, best view selection and automatic mesh segmentation based on 3DLPS are evaluated.

2 Related Work

Signatures have been extensively studied across areas as diverse as computer vision, structure biology and others. In this paper, we mainly review the related works that are designed as point signatures in the context of shape analysis. Furthermore, as is related to 3DLPS, we also briefly review the applications of the Laplacian operator and pyramid defined on 2D images and 3D shapes.

Point signatures on 3D shapes Early work on designing point signatures is based on spatial domain. To make signatures robust against rigid transformation, a common strategy is to summarize the shape distribution in neighborhoods of a point [2]. For instance, the spin image method [3] constructs a 2D histogram that encodes the density of oriented points, and the shape context [4, 5, 11] captures the distribution over distances and angles of all other points on the shape according to the current point. While these signatures are widely used in spacial domain, they do not catch the frequency domain information, which are sometimes essential for shape analysis and understanding.

Recent efforts to track the problem of a robust shape signature are diffusion based approaches [6, 7, 12, 13]. This process provides a natural notion of scale to describe the shape around a point. Rustamov [6] proposed the global point signature (GPS), which shows that the eigenfunctions of the Laplace-Beltrami nicely characterize the geometric features for points. Sun et al. [7] defined the heat kernel signature (HKS) by restricting the well-known heat kernel to the temporal domain. Dey et al. [8] merged the concept of HKS with the persistent homology, and designed a pose-oblivious algorithm for partial shape matching.

Aubry et al. [9] proposed the Wave Kernel Signature (WKS), which represents the average probability of measuring a quantum mechanical particle at a specific location. All these methods somewhat catches the frequency domain information through a multi-scale way. However, they suffer from a computational challenge of estimating an eigen-decomposition of a huge Laplacian matrix. On the contrary, 3DLPS does not require the calculation of eigenvalues and eigenvectors, and thus is more robust and easier to compute, while still capturing the frequency domain information for shape analysis.

Another category of approaches for defining multi-scale point signatures is to capture the features of points at shapes that are resulted from geometry processing operations. For example, Li and Guskov [14] first obtained a series of increasingly smoothed version of a given shape, and then constructed point signatures for features found at each smoothed version of the shape. Manay et al.'s integral invariant signature [15] also falls into this category. 3DLPS is somewhat similar to this class of signatures. However, we use the Laplacian Pyramid operators, which concentrate more on the frequency geometric information.

Applications of Laplacian-beltrami and Pyramid The discrete versions of Laplace and Laplace-Beltrami operators, both referred to as Laplacians, are widely used in image processing [16] and geometry processing [17]. To mention a few, Taubin [18] used the graph Laplacian for surface fairing; Ni et al. [19] used different weight of Laplacians to control the number of critical points; Dong et al. [20] described an approach to the quadrangulation of manifold meshes using Laplacian eigenfunctions; Reuter et al. [21] investigated the discrete Laplace-Beltrami operators for shape segmentation.

Pyramid is a type of multi-scale signal representation developed by the computer vision and image processing communities, in which a signal or an image is subject to repeated smoothing and subsampling [22]. Burt and Adelson [10] applied the Laplacian operator and Gaussian filter in the smoothing step, and presented the Laplacian Pyramid operator for compact image coding and edge enhancement. Kobbelt et al. [23] and Guskov et al. [24] introduced the ideas of pyramid to 3D shape processing for mesh modeling, smoothing and sampling.

Though both the Laplacian operator and the ideas of Pyramid have been successfully applied in image and mesh processing, we found little work has been addressed on deriving point signatures based on Laplacian Pyramid. Since Laplacian Pyramid can be regarded as a spectral decomposition and compression of 3D shapes, it is promising to apply it for capturing the frequency domain information embedded in the shape, which would provide an alternative perspective for shape analysis. This is the main motivation of our work.

3 3D Laplacian Pyramid Signature

Similar to real-world digital images, 3D shapes are in general both scale-variant and highly nonstationary in space. Inspired by the successful applications of Laplacian Pyramid decomposition in images, we define the Laplacian Pyramid operator on 3D surface shapes. Since the Laplacian operator filters the high

frequency information in frequency domain, and the Pyramid decomposition tracks 3D shapes in different frequencies, 3DLPS is supposed to be capable of capturing the intrinsic properties of a point on different frequencies with respect to the shape. In this section, we first give the definition and construction algorithm of 3DLPS, and then explain its desirable attributes for shape analysis.

3.1 Definition and Construction

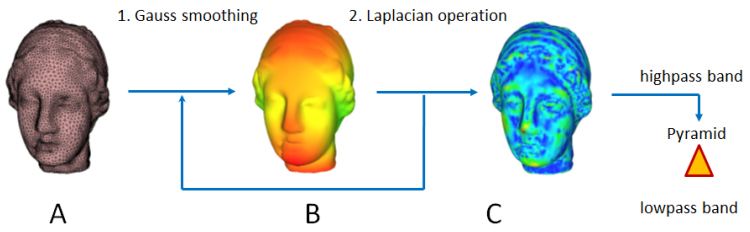

Fig. 1. Pipeline of 3DLPS extraction. The color ramps in B) and C) indicate the magnitudes of δ-coordinates od the models before and after the corresponding operations.

Laplacian Matrix on Mesh Let $\mathscr{M} = (V, E, F)$ be a triangular mesh with n vertices. For each vertex v_i, the δ-coordinates is defined as the difference between the absolute coordinates of v_i and the weighted average of its immediate neighbors in $\mathscr{M}$ [17],

$$\delta_i = (\delta_i^{(x)}, \delta_i^{(y)}, \delta_i^{(z)}) = v_i - \frac{\sum_{j \in N(i)} w_{ij} v_j}{\sum_{j \in N(i)} w_{ij}}, \tag{1}$$

where $N(i) = \{j | (i,j) \in E\}$. to better approximate the mean-curvature normals using δ-coordinates, we employ the "cotangent weights" $w_{ij} = (\cot \alpha_{ij} + \cot \beta_{ij})/(2 * |\Omega_i|)$ as proposed in [25], where $|\Omega_i|$ is the size of the Voronoi cell of v_i, and α_{ij}, β_{ij} are the two angles opposite to edge (i,j). Finally, the normalized Laplacian Matrix L is defined as

$$L_{ij} = \begin{cases} 1 & \text{if } i = j \\ -w_{ij} / \sum_{j \in N(i)} w_{ij} & \text{if } (i,j) \in E \\ 0 & \text{otherwise} \end{cases}. \tag{2}$$

Laplacian Pyramid on Vertices Similar to the Laplacian Pyramid decomposition applied in images, we decompose the Cartesian coordinates on 3D shapes into "highpass" bands and a "lowpass" band. This is achieved by iteratively applying the Laplacian Matrix on the Cartesian coordinates of $\mathscr{M}$. Let

$A = D - L$, where D is the unit diagonal matrix, then the decomposition of the Cartesian coordinates on 3D shapes can be represented as

$$D \cdot \mathbf{C} = A \cdot \mathbf{C} + L \cdot \mathbf{C}, \tag{3}$$

where $\mathbf{C}$ is a $n \times 3$ matrix representing the Cartesian coordinates of the vertices. Here $A \cdot \mathbf{C}$ is the "lowpass" band, and $L \cdot \mathbf{C}$ is the "highpass" band at the current iteration. At each iteration, the "highpass" band $L \cdot \mathbf{C}$ is concatenated to form the vertex coordinates differences, while the "lowpass" band $A \cdot \mathbf{C}$ is passed down as the input of the next iteration (as shown in Fig. 1).

Laplacian of Gaussian on 3D Surfaces Triangular meshes are discrete approximations of continuous surfaces defined on $\mathbb{R}^3$. Hence, directly applying the Laplacian Pyramid operator may cause high noise due to the mesh irregularity. To reduce its sensibility to the discrete approximation as well as the noise, we apply the Gaussian smoothing filter each time before performing Laplacian operator. For Vertex v_i, the Cartesian coordinate after Gaussian smoothing is

$$C(v_i) = \sum_{v_j \in V, ||v_i - v_j|| < 3\sigma} \frac{1}{2\pi\sigma^2} \exp(-\frac{||v_i - v_j||^2}{2\sigma^2}), \tag{4}$$

where σ is the standard deviation of Gaussian distribution, and $||v_i - v_j||$ represents the distance between v_i and v_j. We use the Euclidean distance in our experiments because it gives better results than geodesic distance and is easy to compute [26]. Note that if $||v_i - v_j|| \geq 3\sigma$, the values become very small in Gaussian distribution, we only consider a small neighborhood in which all the vertices has Euclidean distances smaller than 3σ with v_i.

Algorithm 1: Laplacian on 3D surface mesh

Input: m: the maximum iterations;
$\mathscr{M}$: the input surface mesh.
Output: the concatenated "highpass" bands $\mathscr{C}$.

Initialize $\mathscr{C}$ as a null matrix;
for $i \leftarrow 1$ *to* m **do**
  Apply Gaussian smoothing on $\mathscr{M}$ using Equ. 4;
  Decompose the Cartesian coordinates $\mathbf{C}$ on $\mathscr{M}$ into "highpass" band $L \cdot \mathbf{C}$ and "lowpass" band $A \cdot \mathbf{C}$, as formulated in Equ. 3;
  $\mathscr{C} \leftarrow [\mathscr{C},\ L \cdot \mathbf{C}]$; ▷ Concatenate "highpass" band
  $\mathscr{M} \leftarrow A \cdot \mathbf{C}$. ▷ Update $\mathscr{M}$ as "lowpass" band
end

The Laplacian Pyramid operation is finally shown in Algorithm 1. Given a 3D mesh $\mathscr{M}$, we repeatedly apply Gaussian smoothing and Laplacian decomposition. After each decomposition, we concatenate the "highpass" band to the output, and update the 3D surface mesh as the "lowpass" band. The final output is the concatenation of v's δ-coordinates after all iterations.

Construction of 3D Point Signature The concatenation of δ-coordinates itself is able to characterize the mean curvatures of vertices on the surfaces. However, it contains redundant information. To make the 3DLPS efficient for shape analysis, we covert the concatenation of δ-coordinates into a compact and concise representation. Suppose the concatenation of differential vertex coordinates $\mathscr{C}$ has the form $\mathscr{C} = \{\delta_1, \delta_2, ...\delta_m\}$, then we define the final 3DLPS as $S = \{|\delta_1|, |\delta_2|, ...|\delta_m|\}$, where $|\delta|$ is the length of the vector δ.

3.2 Attributes of 3D Laplacian Pyramid Signature

Inherited from the properties of Laplacian Pyramid operator defined on 2D images, 3DLPS exhibits some nice attributes that are desirable for 3D shape analysis. In the following, we will explain them in detail.

Intrinsic and multi-frequency. By intrinsic and multi-frequency, we mean that 3DLPS captures the intrinsic information, and represents them in a multi-frequency way.

Laplacian operator is a smoothing schema, where the high frequency information tends to be smoothed out first, and the low frequency information will be retained for longer time. It is well known from differential geometry that [17]

$$\lim_{|\gamma|\to 0} \frac{1}{|\gamma|} \int_{v\in\gamma} (v_i - v) dl(v) = -H(v_i)\mathbf{n}_i, \tag{5}$$

where $H(v_i)$ is the mean curvature at v_i and $\mathbf{n}_i$ is the surface normal at v_i. Therefore, the magnitude approximates a quantity proportional to the local mean curvature [18]. This means that at each iteration in Algorithm 1, the "highpass" band $L \cdot \mathbf{C}$ encapsulates the local mean curvature normal at each vertex.

We apply the "cotangent weights" schema to approximate the mean-curvature normal. These geometry-dependent weights lead to $L\cdot\mathbf{C}$ with normal components only, rather than encoding the tangential components, which may be non-zero on planar 1-rings. This makes our 3DLPS approximately captures the $\delta-$differences along the normal component, and only captures the local frequency geometry information rather than mesh regularities.

With the increase of iterations, the high frequencies will be smoothed out sequentially, and the lower frequencies will be decomposed into $L \cdot \mathbf{C}$. Therefore, our 3DLPS represents the decreasing frequency information of the vertex. Fig. 2 illustrates the affections of Laplacian Pyramid operations on the homer model. We see that at the beginning, the high frequency parts (nose, fingers and mouth) are smoothed out. With the increase of iterations, the arms and legs shrink tenderly. Since the belly is in low frequency, it changes placidly during the Laplacian operations. Our 3DlPS captures all these information concisely.

Sensitivity and specificity. By sensitivity and specificity, we mean 3DLPS shows high discrimination among different points, and share high similarity among analogous points with respect to the frequencies embedded in the shape.

The sensitivity and specificity of the 3DLPS can be deduced from the attributes of intrinsic and multi-frequency. Since 3DLPS captures the features at

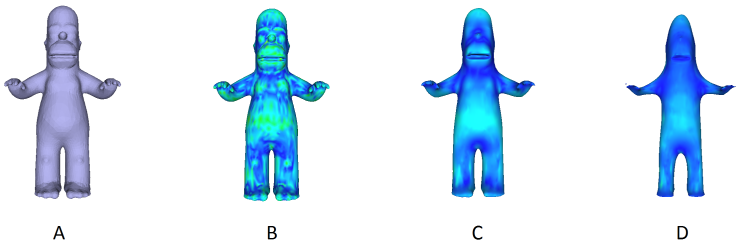

Fig. 2. Affections of Laplacian Pyramid operations on homer. A) The oiginal model; B) magnitudes of the 1st δ-coordinates; C): magnitudes of the 15th δ-coordinates; D) magnitudes of the 30th δ-coordinates.

different frequencies, if two points are embedded with different frequency information (e.g. the points on nose and the points on belly in Fig. 2), their 3DLPSs will be distinctive from each other. On the contrary, if two points are embedded with similar frequency information (e.g. the points on the left leg and the corresponding points on the right leg in Fig. 2), their 3DLPS will be similar. We can intuitively verify that from the color ramp depicted in Fig. 2.

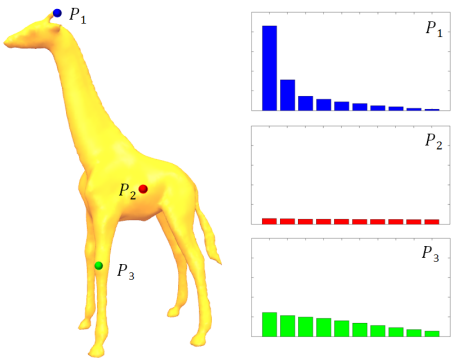

Fig. 3. 3DLPS of different points on the giraffe model.

Fig. 3 gives a more intuitive illustration. P_1 lies on the ear, where there are plenty of high frequency information, thus its 3DLPS has large values in high frequency components. Since P_2 lies on the flat areas, all the components in the 3DLPS are stationary and low; the 3DLPS of P_3 are between those of P_1 and P_2, since the legs of the giraffe have higher frequency than the belly.

Insensitive to articulated deformation. By insensitive to articulated deformation, we mean that for two models, in which one is articulated deformed from another, their 3DLPS signatures are differs only in the joint parts, but remains similar in other regions.

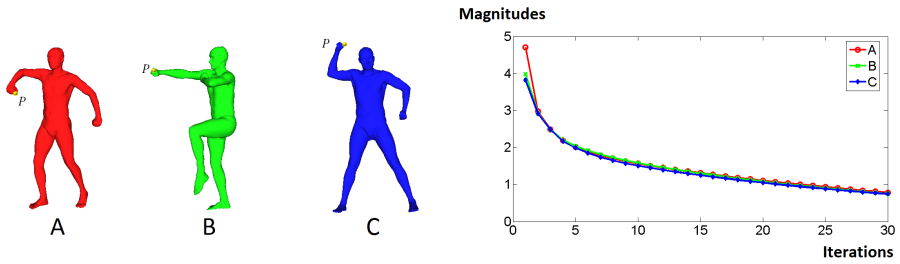

Fig. 4. 3DLPSs of P on different poses of similar men. Due to the different local geometries, the first few components of the 3DLPSs are a little different.

Intuitively, the articulated deformation only changes the geometries near the joint parts, so our 3DLPSs on the two models only differ from each other in these joint parts. Though in theory, the 3DLPS of a point P is affected by its k-ring neighbors, where k is the iteration times, in practice the far distance neighbors affect little of the 3DLPSs in large scales. Fig. 4 shows the corresponding points of similar men with different poses and their according 3DLPSs. Since P is far away from joints, their 3DLPSs are pretty similar in high frequency components.

4 Applications

Since 3DLPS concisely captures the frequency information of points, it has the potential to benefit applications involved in shape analysis from the perspective of frequency domain. As a brief illustration, we will show how 3DLPS is successfully applied in point correspondence detection, best view selection and mesh segmentation. Other applications may also potentially benefit from 3DLPS.

4.1 Point Correspondence Detection

The general solution of correspondence detection is to define the descriptors on surfaces or points, and then match the descriptors between them. Since 3DLPS reflects the multi-frequency information embedded in the shape, it can be incorporated into the general solution for point correspondence detection.

However, the original 3DLPS is sensitive to local geometries in high frequencies, whose magnitudes usually dominate the feature vectors (See the curves in Fig. 4). To solve this problem, we define a new signature based on 3DLPS, called *Context 3DLPS*, which is more robust and adaptive, yet rotation-invariant. We apply it as the new tool for robust point correspondence detection.

Definition 1 (Context 3DLPS) Given a vertex v, we average over all vertices in its i-th neighbor ring, and denote it as $\bar{S}_i$ (the 0th neighbor vertex is v itself), then the Context 3DLPS of v is defined as $S(v, nr) = (\bar{S}_0, \bar{S}_1, ..., \bar{S}_{nr})$, where nr is the number of neighbor rings.

The similarities of Context 3DLPSs are measured by χ^2 distance in our experiments, which is defined as

$$Sim(h_1, h_2) = \sum_{i=1}^{m} \frac{(h_1(i) - h_2(i))^2}{h_1(i) + h_2(i)}, \tag{6}$$

where h_1 and h_2 are the Context 3DLPS of v_1 and v_2 respectively, and m is the dimension of the Context 3DLPS.

For Context 3DLPS, the larger nr is, the larger context of information around v we considered, yet the larger memory storage and computation are required. In our experiments, we found $nr = 10$ is a good tradeoff between context information amount and computational cost, as shown in Fig. 5.

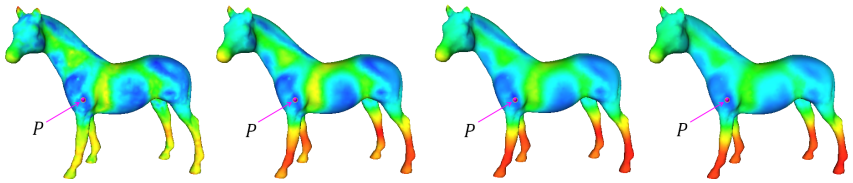

(a) $nr = 0$. (b) $nr = 5$. (c) $nr = 10$. (d) $nr = 15$.

Fig. 5. Illustration of Context 3DLPS with various parameter nr. The color ramp depicts the differences of Context 3DLPS between their according locations and P, in which hot color means high differences. When nr is too small, the Context 3DLPS only captures the very local context information, which makes its specificity low (dissimilar points are wrongly regarded as similar).

To detect the corresponding points of P, we first calculate the Context 3DLPS for each vertex, and then measure the similarities between P and all other vertices using Equ. 6. The top k points with the smallest χ^2 distances are regarded as the candidates. To make the candidates distinct from each other and reduce the number of candidates, we merge the candidates within the top k points as long as their distances are smaller than a user-specified threshold ϵ (in experiments, we set ϵ as 10 times the average edge length of the mesh), and set the new merged candidate as the one whose χ^2 distance is the smallest to P. Finally, Users can manually select the best correspondence point, or apply some algorithms to further refine the candidates [27–29]. As our main goal is to demonstrate the power of 3DLPS, we do not involve the matching algorithms here.

Fig. 6 shows the results of corresponding point candidates detected by our algorithm. In this experiment, we first designate points P_1, P_2 and P_3 on model 1, and then calculate the top 3 most similar points on model 1 and model 2 respectively. In Fig 6 (a), we verify that our method is capable of detecting the symmetric points. For instance, Q_{21} and Q_{31} are detected as the most similar points to P_2 and P_3 respectively. In Fig. 6 (b), we verify that our method is insensitive to articulated deformation. Though model 2 is a deformation of model

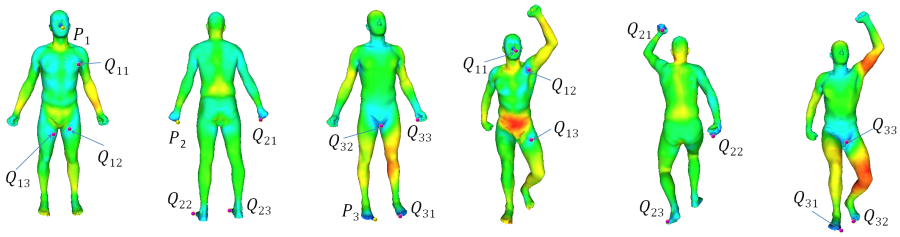

(a) P_1, P_2, P_3 and their most similar points on model 1.

(b) The most similar points of P_1, P_2 and P_3 on model 2.

Fig. 6. Illustration of our methods to detect the corresponding points. P_i is the ith user designated point, and Q_{ij} indicates the jth similar point to P_i. We designate P_1, P_2 and P_3 on model 1, and detect the top 3 similar points on both model 1 and model 2. Note that model 2 is an articulated deformation of model 1.

1, our algorithm can still detect Q_{21} and Q_{31} as the most similar points to P_2 and P_3, followed by their symmetric points Q_{22} and Q_{32}. We use color ramp to illustrate the χ^2 distance between the designated points and other vertices, in which cool colors indicate small distances.

4.2 Best View Selection

Automatic generation of best views for 3D models has drawn much attention [30], due to its applications in 3D model browsing, automatic camera replacement, 3D scene generation and view-based 3D object recognition [31, 32].

Among all the existing algorithms, one category aims at maximizing the amount of features visible from the viewpoint. These methods usually associate a goodness measure to a number of candidate views, where the goodness measure is a function of some objectives related to the geometrical properties of the object [26]. 3DLPS can be naturally incorporated as the geometrical properties for best view selection. In addition, since it exhibits various frequency information in 3DLPS, it provides an adaptive mechanism for visualizing different frequency information according to user specifications.

The framework of our algorithm is the same as the mesh saliency [26], in which the only difference is we replace the mesh saliency of v by the function $f(v)$ defined on v based on 3DLPS. For a viewpoint P, let $F(P)$ be the set of vertices visible from P, we compute the visible information from P as

$$U(P) = \sum_{v \in F(P)} f(v). \tag{7}$$

Then the viewpoint with maximum visible information P_m is

$$P_m = \underset{P}{\operatorname{argmax}}\ U(P). \tag{8}$$

Similar to [26], we use the gradient-descent-based optimization heuristic to select good viewpoints for efficiency.

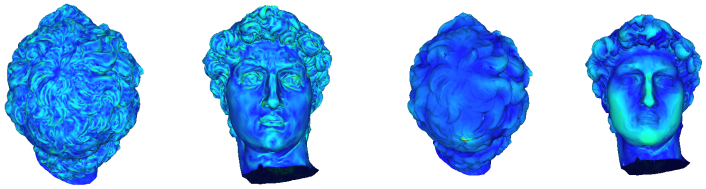

(a) Color ramp of the magnitudes of the 1st component.

(b) Color ramp of the magnitudes of the 30th component.

Fig. 7. a): different views of the magnitudes of the 1st component of the 3DLPSs. The left view is determined as the best view by our algorithm; b) different views of the magnitudes of the 30th component of the 3DLPSs. the right view is determined as the best view by our algorithm.

Fig. 7 intuitively shows how magnitudes change with the variations of frequency in 3DLPSs. If we define $f(v)$ as the magnitude of the first component of v's 3DLPS, then the best view is selected as the left view in Fig. 7(a). Note that this view is similar to the view based on maximizing the pure curvatures, since both of them tend to exhibit the high frequency information. If $f(v)$ is defined as the 30th component of v's 3DLPS, our algorithm tend to select the views that exhibit the low frequency information, as shown in the right view of Fig. 7(b).

By defining different $f(v)$ based on 3DLPS, our algorithm is adaptive to generate different best views that maximize different information. For instance, if we define $f(v)$ as the 5th component of v's 3DLPS, the best view selected by our algorithm is as Fig 8(e).

A more reasonable strategy is to exhibit different frequency information approximately equally in the best view. To achieve this goal, we define

$$f(v) = \sum_{i=1}^{m} rank(v, i), \tag{9}$$

where m is the length of 3DLPS, and $rank(v, i)$ indicates the index of v sorted according to the ith components of all vertices in ascending order. This function normalizes the low frequency information and high frequency information into a commeasurable metric space. Thus, by maximizing $f(v)$s across all the vertices with all the specified frequencies, the best view is supposed to exhibit various frequency information equally, as shown in Fig. 8(f).

Fig. 8 also gives a brief comparison with some of the state-of-the-art approaches. Note that the result of [26] is similar to our result in Fig. 8(f). However, our result exhibits more information on different frequencies (refer to the color ramp in Fig. 7 for comparison). Actually, the method proposed in [26] is exactly the Difference of Gaussian (DOG) defined on 3D, which is an approximation of Laplacian of Gaussian (LOG) utilized in our method. Since we separate different frequencies and define the signature as a vector, rather than summing them up to form a scale value [26], our method is more adaptive to exhibit different frequency information.

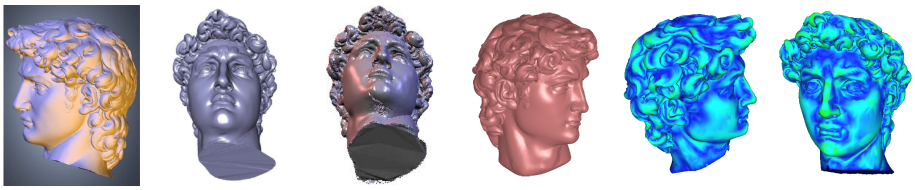

(a) (b) (c) (d) (e) (f)

Fig. 8. Comparison of the best views of David head. a): Result of [26]; b): Result of [33]; c): Results of [34]; d): Result of [35]; e): Our result by maximizing the magnitudes of the 15th component in 3DLPSs; f): Our result by defining $f(v)$ as Equ. 9. Note in e) and f), the pictures are mapped to color ramps, in which cool color means small value. To make our best views visibly pleasant, we manually rotated our best views to the correct upright direction.

4.3 Automatic Mesh Segmentation

By observation, we found in many models, functional parts are usually different in frequency domain, and the boundaries between them often coincide with the regions where the frequency magnitudes change rapidly (see Fig. 9(c)). Since 3DLPS provides a good measure for frequency differences, it provides a natural way for segmentation of this kind of objects.

Inspired by this observation, we propose a label-swop based mesh segmentation method, as shown in Algorithm 2. In this algorithm, we first initialize the vertices with different labels, and then segment the mesh using a label-swop mechanism, as shown in lines 5-17 in Algorithm 2. We first estimate the probability density function of each vertex (shown in line 8-9) using Gaussian weights, and then swop the label of each vertex to its exemplar that possesses the locally highest estimated probability density function value. Finally, the segments are merged sequentially until the user specified segments number is achieved.

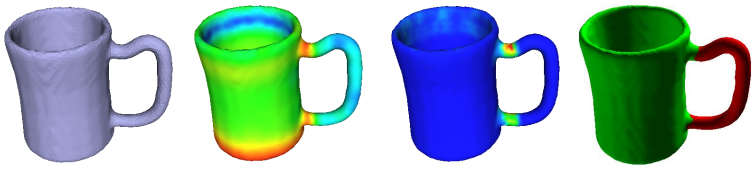

(a) (b) (c) (d)

Fig. 9. Illustration of our label-swop based mesh segmentation algorithm. a): the original model; b): the color ramp mappings of 3DLPSs ($h = 15$); c): the local variations of low frequency components in b) ($nr = 3$); d): the final segmentation result ($n = 2$).

Fig. 9 illustrates how our label-swop algorithm performs on the cup model. In this experiments, we set $n = 2$, $h = 15$ and $nr = 3$. As shown in Fig. 9(c), the

Algorithm 2: Label-swop based mesh segmentation

Input: n: the number of segments specified by users;
h: the to be used index of component in 3DLPS;
nr: the parameter to determine local regions.
Output: L: label array that represents the result.

for $i \leftarrow 1$ *to* $size(V)$ **do**
  $L(v_i) \leftarrow i$; ▷ initialize vertices with different labels
  $M(v_i) \leftarrow$ the hth component of the 3DLPS of v_i;
end
for $i \leftarrow 1$ *to* $size(V)$ **do**
  $N(v_i) \leftarrow \max_{v_j \in N(v_i, nr)} ||M(v_i) - M(v_j)||$; ▷ $N(v_i, nr)$ is v_i's nr-ring neighbors
end
for $i \leftarrow 1$ *to* $size(V)$ **do**
  $D(v_i) \leftarrow \sum_{j \in N(v_i, nr)} N(v_j) \exp(-\frac{||v_i - v_j||^2}{2\sigma^2})$;
end
foreach *Vertex* v *in the mesh* **do**
  $v_{final} \leftarrow v$;
  while $\exists v_j \in N(v_i, 1)$ *such that* $D(v_j) > D(v_{final})$ **do**
    $v_{final} \leftarrow v_j$; ▷ swop v_j as the new exemplar
  end
  $L(v) \leftarrow L(v_{final})$;
end
$m \leftarrow$ the number of different labels in L;
while $m > n$ **do**
  foreach *adjacent segments* S_i *and* S_j **do**
    $v_{min}(i, j) \leftarrow$ the vertex v whose $D(v)$ is the minimum along the boundaries of S_i and S_j;
  end
  $v_{join} \leftarrow \operatorname{argmax}_v D(v_{min}(i, j))(1 \le i, j \le m, i \ne j)$;
  join the segments adjacent to v_{join} by updating L;
end

local variations of the specified components for all the vertices clearly marked the boundaries between the body and the handle of the cup. This leads to consistent segmentations with human intuition.

We provide a comparison of our method with some of the state-of-the-art methods. We conclude that for the models whose ground-truth segmentations coincide with the low frequency variations, such as ants, octopuses, tables and so on, our algorithm performs consistently better than existing methods (see Fig. 10 for details).

However, our segmentation algorithm suffers from two main limitations: 1) users have to specify the number of segments in advance; 2) for the models whose semantic segmentations do not coincide with low frequency variations, such as human faces, our algorithm may fail. As demonstrated in [1, 29, 42–44] that no one automatic segmentation algorithm is better than the others for all types of

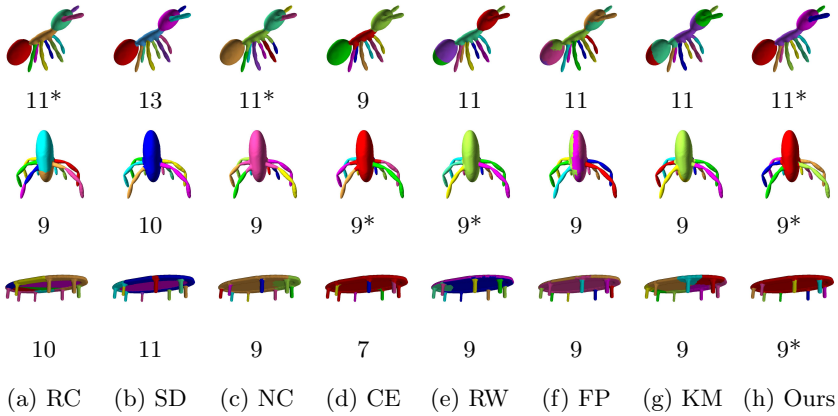

Fig. 10. Comparison with other segmentation methods. The number below each image indicates the number of the according segments, and the "*" following it means the according segmentation approximately coincides with the ground-truth manual segmentation. The methods used from left to right columns are a): Randomized cuts (RC) [36]; b): Shape diameter (SD) [37]; c) Normalized cuts (NC) [36]; d): Core extraction (CE) [38]; e) Random walks (RW) [39]; f) Fitting primitives (FP) [40]; g) K-Means (KM) [41]; h) Our method (Ours), respectively.

objects, we are confident that our algorithm provides an alternative means of segmentation for some types of models.

5 Discussion and Conclusion

In this paper, we have proposed a simple yet efficient point descriptor, called 3D Laplacian Pyramid signature, which reveals the frequency domain information of points embedded in the shape. Our 3DLPS is a multi-scale representation of the points and thus can be applied to plenty of applications involved in shape analysis, such as point correspondence detection, best view selection and automatic mesh segmentation. Compared with existing point signatures such as GPS, HKS and WKS, our proposed 3DLPS is more simpler to calculate, yet efficient enough for plenty of applications, since it inherits the nice properties of Laplacian Pyramid defined in 2D images.

However, 3DLPS still suffers from some limitations. For example, given two isometric deformed models, the 3DLPSs of vertices on the joint regions may be a little different, since 3DLPS simply catches the magnitudes of changes induced by Laplacian operations. This means that 3DLPS is not completely invariant under isometric transformations. Another limitation is that 3DLPS may suffer from noises of high frequency information (see Fig. 5). In this case, we have to introduce Context 3DLPS for some applications such as robust point correspondence detection. We will further investigate how to solve these problems.

References

1. Fang, Y., Sun, M., Ramani, K.: Temperature distribution descriptor for robust 3d shape retrieval. 4th Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment (NORDIA'11) (2011)
2. Chua, C.S., Jarvis, R.: Point signatures: A new representation for 3d object recognition. International Journal of Computer Vision **25** (1997) 63–85
3. Johnson, A.: Spin-Images: A Representation for 3-D Surface Matching. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA (1997)
4. Belongie, S., Malik, J., Puzicha, J.: Shape context: A new descriptor for shape matching and object recognition. In: NIPS. (2000) 831–837
5. Kokkinos, I., Bronstein, M.M., Litman, R., Bronstein, A.M.: Intrinsic shape context descriptors for deformable shapes. In: CVPR. (2012) 159–166
6. Rustamov, R.M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In: Symposium on Geometry Processing. (2007) 225–233
7. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. SGP '09: Proceedings of the Symposium on Geometry Processing (2009) 1383–1392
8. Dey, T.K., Li, K., Luo, C., Ranjan, P., Safa, I., Wang, Y.: Persistent heat signature for pose-oblivious matching of incomplete models. Comput. Graph. Forum **29** (2010) 1545–1554
9. Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In: ICCV Workshops. (2011) 1626–1633
10. Burt, P.J., Edward, Adelson, E.H.: The laplacian pyramid as a compact image code. IEEE Transactions on Communications **31** (1983) 532–540
11. Gal, R., Shamir, A., Cohen-Or, D.: Pose-oblivious shape signature. Visualization and Computer Graphics, IEEE Transactions on **13** (2007) 261–271
12. Bronstein, A.M., Bronstein, M.M., Kimmel, R., Mahmoudi, M., Sapiro, G.: A gromov-hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. Int. J. Comput. Vision **89** (2010) 266–286
13. Bronstein, A.M., Bronstein, M.M., Guibas, L.J., Ovsjanikov, M.: Shape google: Geometric words and expressions for invariant shape retrieval. ACM Trans. Graph. **30** (2011) 1:1–1:20
14. Li, X., Guskov, I.: Multiscale features for approximate alignment of point-based surfaces. In: Symposium on Geometry Processing. Volume 255 of ACM International Conference Proceeding Series., Eurographics Association (2005) 217–226
15. Manay, S., Hong, B.W., Yezzi, A.J., Soatto, S.: Integral invariant signatures. In: ECCV (4). (2004) 87–99
16. Paris, S., Hasinoff, S.W., Kautz, J.: Local laplacian filters: edge-aware image processing with a laplacian pyramid. ACM Trans. Graph. **30** (2011) 68
17. Sorkine, O.: Laplacian Mesh Processing. PhD thesis, School of Computer Science, Tel Aviv University (2006)
18. Taubin, G.: A signal processing approach to fair surface design. In: Proceedings of the 22Nd Annual Conference on Computer Graphics and Interactive Techniques. SIGGRAPH '95, New York, NY, USA, ACM (1995) 351–358
19. Ni, X., Garland, M., Hart, J.C.: Fair morse functions for extracting the topological structure of a surface mesh. ACM Trans. Graph. **23** (2004) 613–622
20. Dong, S., Bremer, P.T., Garland, M., Pascucci, V., Hart, J.C.: Spectral surface quadrangulation. ACM Trans. Graph. **25** (2006) 1057–1066

21. Reuter, M., Biasotti, S., Giorgi, D., Patan, G., Spagnuolo, M.: Discrete laplacec-beltrami operators for shape analysis and segmentation. Computers & Graphics **33** (2009) 381 – 390 IEEE International Conference on Shape Modelling and Applications 2009.
22. Meer, P., Baugher, E.S., Rosenfeld, A.: Frequency domain analysis and synthesis of image pyramid generating kernels. IEEE Trans. Pattern Anal. Mach. Intell. **9** (1987) 512–522
23. Kobbelt, L., Campagna, S., Vorsatz, J., Seidel, H.P.: Interactive multi-resolution modeling on arbitrary meshes. In: Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques. SIGGRAPH '98, New York, NY, USA, ACM (1998) 105–114
24. Guskov, I., Sweldens, W., Schröder, P.: Multiresolution signal processing for meshes. Computer Graphics Proceedings (SIGGRAPH 99) (1999) 325–334
25. Pinkall, U., Polthier, K.: Computing discrete minimal surfaces and their conjugates. Experimental Mathematics **2** (1993) 15–36
26. Lee, C.H., Varshney, A., Jacobs, D.W.: Mesh saliency. ACM Trans. Graph. **24** (2005) 659–666
27. Dubrovina, A., Kimmel, R.: Matching shapes by eigendecomposition of the Laplace-Beltrami operator. In: 3DPVT. (2010)
28. Funkhouser, T., Shilane, P.: Partial matching of 3D shapes with priority-driven search. In: Symposium on Geometry Processing. (2006)
29. Sun, M., Fang, Y., Ramani, K.: Center-shift: An approach towards automatic robust mesh segmentation (arms). In: CVPR. (2012) 2145–2152
30. Hu, K.M., Wang, B., Yuan, B., Yong, J.H.: Automatic generation of canonical views for cad models. In: CAD/Graphics. (2011) 17–24
31. Dutagaci, H., Cheung, C.P., Godil, A.: A benchmark for best view selection of 3d objects. In: Proceedings of the ACM Workshop on 3D Object Retrieval. 3DOR '10, New York, NY, USA (2010) 45–50
32. Secord, A., Lu, J., Finkelstein, A., Singh, M., Nealen, A.: Perceptual models of viewpoint preference. ACM Transactions on Graphics **30** (2011)
33. Mortara, M., Spagnuolo, M.: Semantics-driven best view of 3d shapes. Computers & Graphics **33** (2009) 280–290
34. Yamauchi, H., Saleem, W., Yoshizawa, S., Karni, Z., Belyaev, A.G., Seidel, H.P.: Towards stable and salient multi-view representation of 3d shapes. In: SMI, IEEE Computer Society (2006) 40
35. Leifman, G., Shtrom, E., Tal, A.: Surface regions of interest for viewpoint selection. In: CVPR, IEEE (2012) 414–421
36. Golovinskiy, A., Funkhouser, T.: Randomized cuts for 3D mesh analysis. ACM Transactions on Graphics (Proc. SIGGRAPH ASIA) **27** (2008)
37. Shapira, L., Shamir, A., Cohen-Or, D.: Consistent mesh partitioning and skeletonisation using the shape diameter function. Vis. Comput. **24** (2008) 249–259
38. Katz, S., Leifman, G., Tal, A.: Mesh segmentation using feature point and core extraction. The Visual Computer **21** (2005) 649–658
39. kun Lai, Y., min Hu, S., Martin, R.R., Rosin, P.L.: Fast mesh segmentation using random walks. In: In ACM Symposium on Solid and Physical Modeling. (2008)
40. Attene, M., Falcidieno, B., Spagnuolo, M.: Hierarchical mesh segmentation based on fitting primitives. THE VISUAL COMPUTER **22** (2006) 181–193
41. Shlafman, S., Tal, A., Katz, S.: Metamorphosis of polyhedral surfaces using decomposition. In: Computer Graphics Forum. (2002) 219–228
42. Chen, X., Golovinskiy, A., Funkhouser, T.: A benchmark for 3D mesh segmentation. ACM Transactions on Graphics (Proc. SIGGRAPH) **28** (2009)

43. Fang, Y., Liu, Y.S., Ramani, K.: Three dimensional shape comparison of flexible proteins using the local-diameter descriptor. BMC Structural Biology **9** (2009) 29
44. Fang, Y., Sun, M., Kim, M., Ramani, K.: Heat-mapping: A robust approach toward perceptually consistent mesh segmentation. In: CVPR. (2011) 2145–2152