Computers & Graphics 46 (2015) 264–274

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

SMI 2014

Mesh saliency via ranking unsalient patches in a descriptor space

Pingping Tao [a], Junjie Cao [a,b], Shuhua Li [a], Xiuping Liu [a,*], Ligang Liu [c]

[a] School of Mathematical Sciences, Dalian University of Technology, China
[b] School of Mathematics and Information Science, Nanchang Hangkong University, China
[c] School of Mathematical Sciences, University of Science and Technology of China, China

ARTICLE INFO

Article history:
Received 2 July 2014
Received in revised form
26 September 2014
Accepted 27 September 2014
Available online 7 October 2014

Keywords:
Mesh saliency
Manifold ranking

ABSTRACT

This paper presents a novel mesh saliency detection approach based on manifold ranking in a descriptor space. Starting from the over-segmented patches of a mesh, we compute a descriptor vector for each patch based on Zernike coefficients, and the local distinctness of each patch by a center-surround operator. Patches with small or high local distinctness are named as background or foreground patches, respectively. Unlike existing mesh saliency methods which focus on local or global contrast, we estimate the saliency of patches based on their relevances to some of the most unsalient background patches, i.e. background patches with the smallest local distinctness, via manifold ranking. Compared with ranking with some of the most salient foreground patches as queries, this improves the robustness of our method and contributes to make our method insensitive to the queries estimated. The ranking is performed in the descriptor space of the patches by incorporating the manifold structure of the shape descriptors, which therefore is more applicable for mesh saliency since the salient regions of a mesh are often scattered in spatial domain. Finally, a Laplacian smoothing procedure is applied to spread the patch saliency to each vertex. Comparisons with the state-of-the-art methods on a wide range of models show the effectiveness and robustness of our approach.

1. Introduction

Visual saliency is an important and fundamental research problem in neuroscience and psychology to investigate the mechanism of human visual systems. It has also been an attractive topic in computer vision and computer graphics in recent years. Mesh saliency reflects perceptually important points or regions of polygonal meshes. While mesh saliency may not outperform some differential geometry measures, such as curvature, as a mesh feature in all the applications, it has great value for human centered visual computing applications, such as abstraction [1], simplification [2], smoothing [3], illumination [4], shape matching [5], rendering [4] and viewpoint selection [6–9], especially with the rapid growth of number and size of 3D models.

Although there has been significant progress in mesh saliency detection, most previous works depend on a local center-surround operator [10–12] and multi-scale computation [2,3,5,13]. They tend to be susceptible to noise. For example, Refs. [2,10–12] simply select regions where the curvature of a surface vertex or patch is different from its immediate neighbors. The computation of curvature employed by these methods is sensitive to noise. Ref. [3] proposes a novel approach for mesh saliency estimation considering both local contrast and global rarity, which is robust against noise. However, it is not easy to tune the parameters to obtain faithful results. Leifman et al. [13] choose 20% of the most distinct vertices and extreme vertices as focus points. Regions that are close to the focus points catch more attention than faraway regions. But some salient regions, the distinctness of all vertices within which is less than 20%, will be missed. Furthermore, the multi-scale operation employed in the above methods may only alleviate the influence of noise to some degree.

To handle the aforementioned problems, we propose a novel, simple and robust method for detecting regions of interest on surfaces. Visual forms may possess one or several centers of gravity about which the form is organized [13]. Human attention is firstly attracted by the most representative salient elements (we name them as saliency queries) and then the visual attention will be transferred to other regions [14]. We employ a semi-supervised algorithm, manifold ranking, to imitate the process. First, we oversegment a mesh into patches and compute a descriptor for each of them via Zernike coefficients, which is more informative and robust than curvature. Then the distinctness of each patch is estimated locally using these descriptors. Instead of selecting some foreground patches such as saliency queries, patches with the smallest distinctness values, i.e. some of the most unsalient background patches, are chosen as queries to improve the robustness

* Corresponding author.
E-mail addresses: taoping85@gmail.com (P. Tao), jjcao1231@gmail.com, jjcao@dlut.edu.cn (J. Cao), sue142857@gmail.com (S. Li), xpliu@dlut.edu.cn (X. Liu), lgliu@ustc.edu.cn (L. Liu).

http://dx.doi.org/10.1016/j.cag.2014.09.023
0097-8493/

of our method. With these queries and a self-adapting graph defined in the descriptor space of patches, the saliency of all the patches is determined using manifold ranking. Ranking in the descriptor space is more applicable for mesh saliency since the salient regions of a mesh are usually scattered in the spatial domain. Finally, we use simple Laplacian smoothing to spread the patch saliency to vertex saliency. The patch descriptor and the strategy to generate the queries together contribute to make our

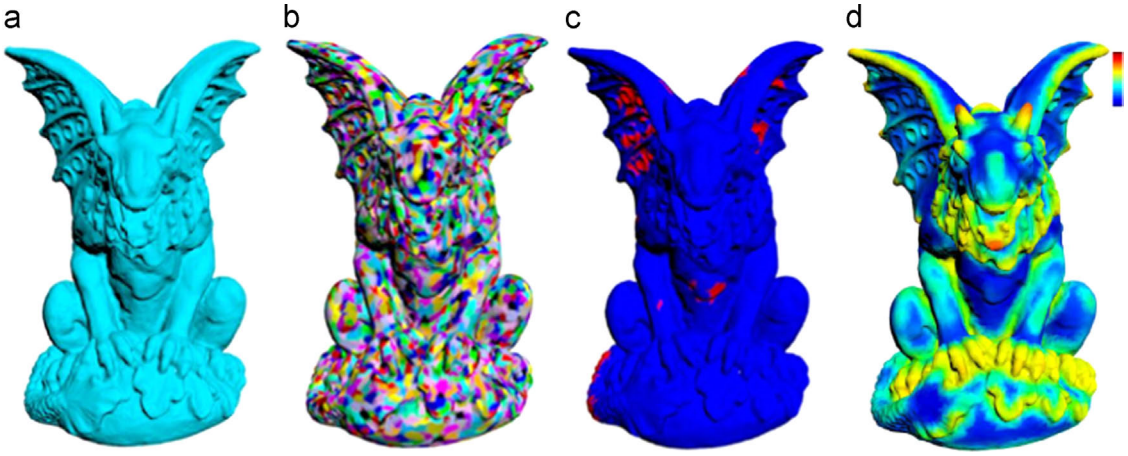

Fig. 1. Method overview. (a) Input mesh. (b) Over-segmentation. (c) Some of the most unsalient background patches (red). (d) Mesh saliency. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

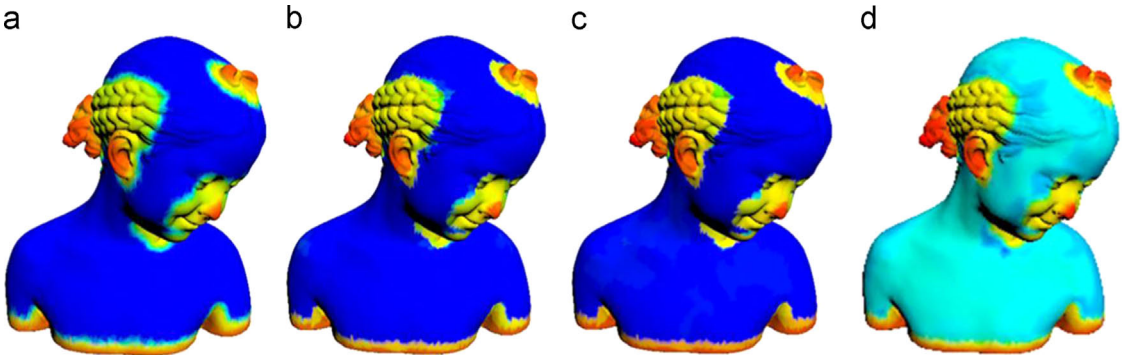

Fig. 2. Our method is robust to patch numbers. From left to right, the patch numbers are 3924, 1970, 902 and 772.

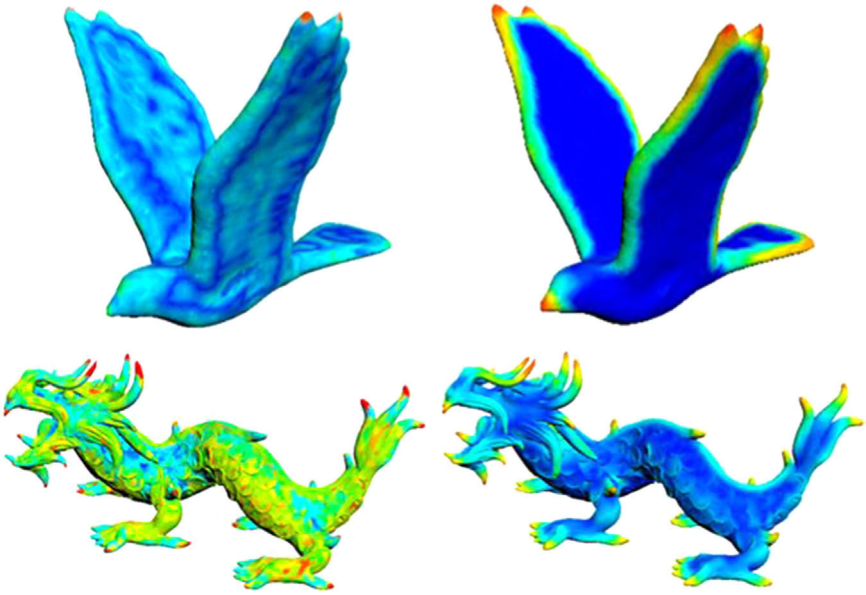

Fig. 3. Mesh saliency via ranking with curvature-based patch descriptor (the left column) is sensitive to local geometry, while using Zernike-based patch descriptor (the right column) is more tidier and faithful.

method robust to noise. Experiments demonstrate that our single scale method outperforms many state-of-the-art methods, especially when the models are contaminated by noise. The contributions of our method are as follows:

- Manifold ranking is introduced to imitate the transformation of human attention for mesh saliency detection. It is easy to compute and efficient.
- Instead of transferring saliency from unstable foreground patches, some background patches are selected as queries, which makes our method not sensitive to input parameters and robust against noise.
- Considering that salient regions of a mesh are usually scattered in the spatial domain, a self-adapting graph is defined in the descriptor space of patches, which benefits to reveal the saliency patches independent of their locations and cardinality.

2. Related work

Existing mesh saliency methods fall into two categories: salient points detectors [15–17] and salient regions detectors. Our method belongs to the latter which generates a saliency map, i.e. a saliency value for each vertex.

Early mesh saliency detection methods estimate saliency of a 3D shape by computing saliency in its 2D projection. Guy and Medioni

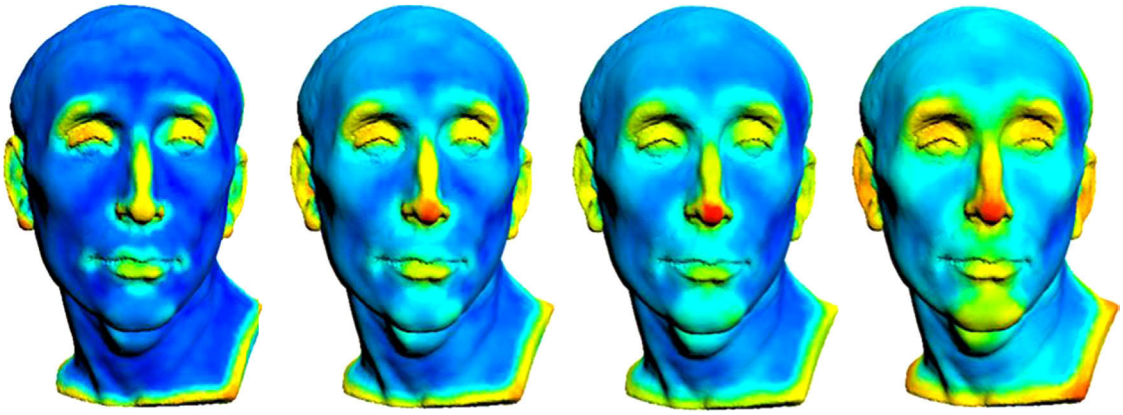

Fig. 4. Our mesh saliency results are insensitive with different descriptor scales. The results shown from left to right are obtained with the scale of 3.0l, 4.0l, 5.0l, 6.0l.

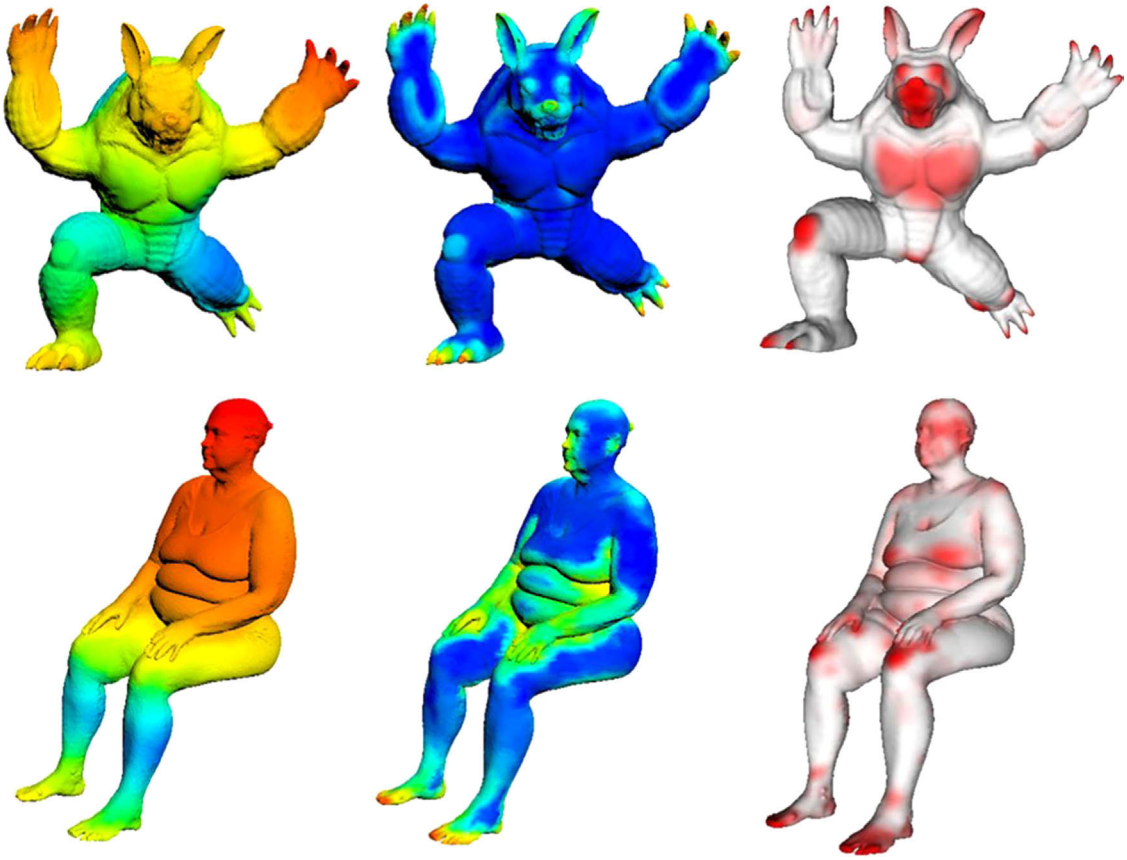

Fig. 5. The saliency maps generated using a two-loop graph in the spatial domain and those generated using our self-adaption graph G in the descriptor space, and the pseudo-ground truth [16] (from left to right). Taking the top row as an example, the saliency of left and right thigh is different when ranking in the spatial domain and is the same when ranking in the descriptor space.

[18] propose a scheme for computing a saliency map based on edges in a 2D image, and apply it to 3D meshes. Yee et al. [4] use the method in [19] to calculate a saliency result of a coarsely rendered 2D projection of a 3D dynamic scene. Mantiuk et al. [20] present a 2D saliency algorithm to guide MPEG compression of a 3D scene animation.

Recently, some works compute saliency on meshes directly. Shilane et al. [5] analyze distinctive regions based on performing a shape-based search. However, the results undesirably change with training database. Leifman et al. [13] develop an approach to detect the interest regions of surfaces for viewpoint selection.

There are also many mesh saliency algorithms extended from image saliency methods. Itti et al. [19] use the center-surround mechanism at different scales to compute a saliency map. Inspired by [19], Lee et al. [2] define mesh saliency in a scale-dependent manner using a center-surround operator on Gaussian-weighted mean curvatures. They also show that mesh saliency can visually enhance mesh simplification and viewpoint selection. Instead of using only the local geometric cues, Wu et al. [3] extend Cheng et al. [21] to compute mesh saliency. They consider both local contrast and global rarity. Spectral analysis is once an effective approach for image saliency detection [22]. Similar to the spectral residual analysis in [22], Song et al. [23] analyze the spectral attributes of the log-Laplacian spectrum of a mesh. The spectral residual in the spectral domain is transformed back to the spatial domain to obtain the mesh saliency map. Although there are many image saliency methods [21,24,25] in spatial domain achieving far faithful results than spectral based methods recently, they are not extended to 3D meshes and the spectral approach [23] is the best mesh saliency method as far as we known. We generalize the idea of Yang et al. [24] to mesh saliency detection. The approach is expected to work for 3D mesh, since we can treat image and mesh in a united way using a graph consisting of oversegmented patches. However similar to the above, there are still many problems needed to be conquered during the generalization, such as the choice of proper descriptor and construction of the graph used in manifold ranking, since image and mesh have different generation mechanisms and the nature between image saliency and mesh saliency is different.

3. Approach

3.1. Algorithm overview

Our method involves three main steps (see Fig. 1). Taking a mesh as input, we oversegment it into patches and compute a descriptor for each of them based on Zernike coefficients [28] (Section 3.3). Then a rough distinctness of each patch is estimated locally (Section 3.4). Among the most indistinct patches, we chose some of them as queries. After constructing a self-adaption graph in the descriptor space of patches (Section 3.5), the saliency of all the patches is computed based on their relevances to the given queries via manifold ranking (Section 3.6). Finally, a smooth vertex saliency map is achieved by splatting the patch saliency. The pseudocode of our method is shown in Algorithm 1. Before detailing the main steps, we introduce manifold ranking briefly in Section 3.2.

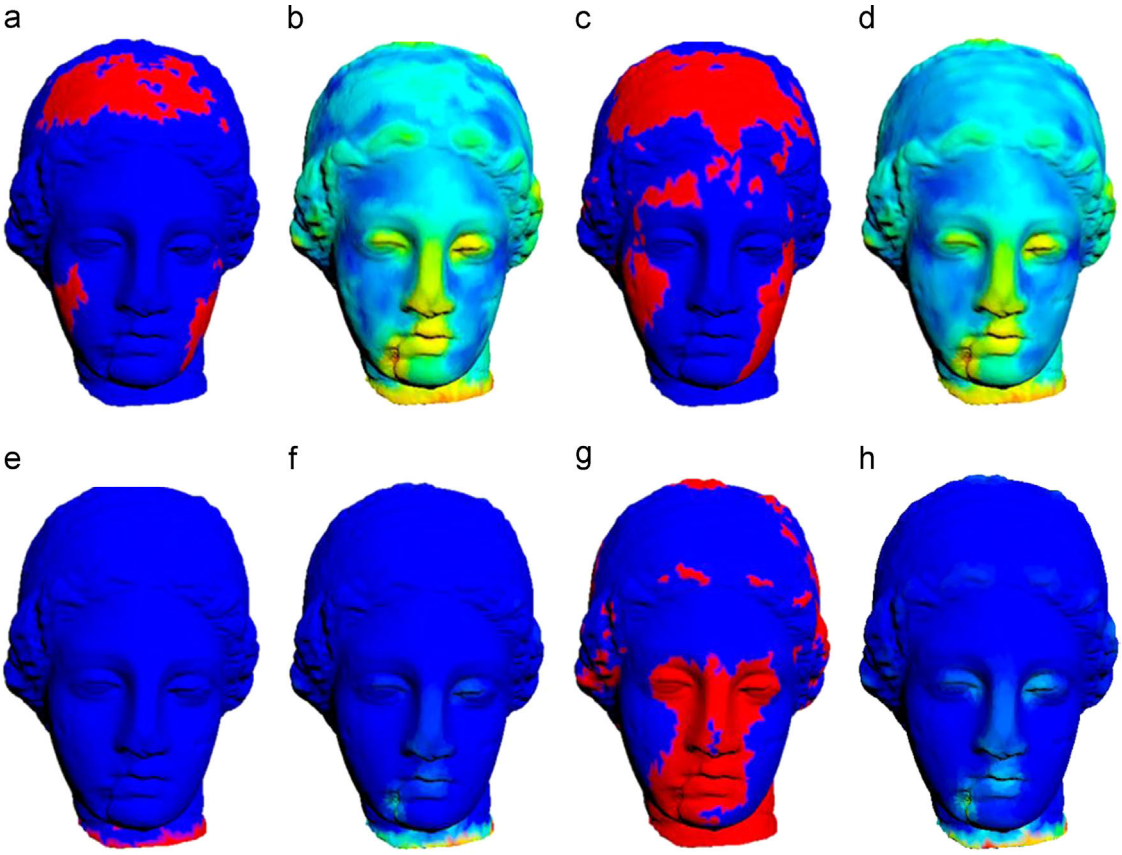

Fig. 6. Saliency generated by ranking with background queries is more stable than with foreground queries. The red patches in (a) and (c) denote the background queries with local contrast smaller than 0.05 and 0.1. (b) and (d) are the corresponding saliency maps of (a) and (c). (e) and (g) are the foreground queries with local contrast greater than 0.6 and 0.2, respectively. (f) and (h) are the corresponding saliency maps of (e) and (g), respectively. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

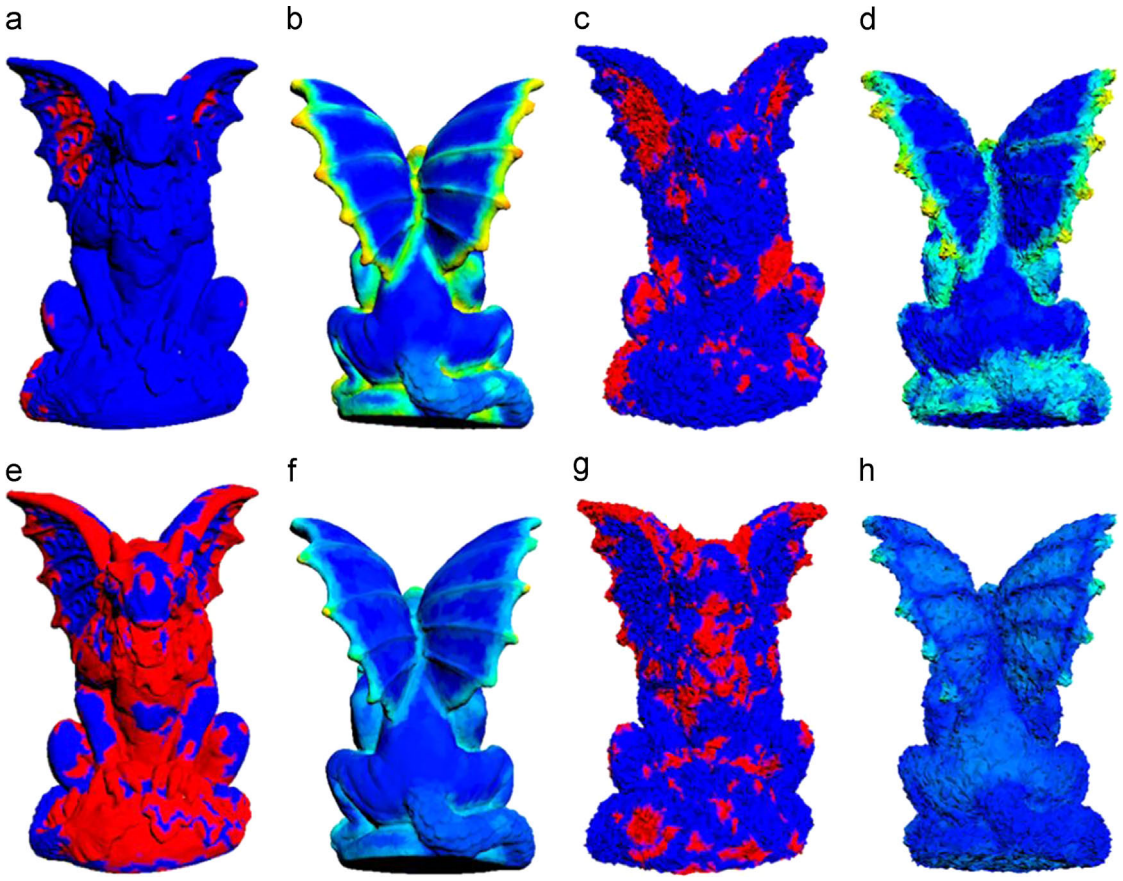

Fig. 7. Ranking with unsalient background queries is more robust than with salient foreground queries. The red patches in (a) and (c) denote the background queries with local contrast smaller than 0.1 on the original model and the model contaminated by 40% noise. (b) and (d) are the corresponding saliency maps of (a) and (c), respectively. (e) and (g) are the salient foreground patches with local contrast greater than 0.2 on the original model and the model contaminated by 40% noise, respectively. (f) and (h) the corresponding saliency maps of (e) and (g), respectively.

Table 1
Run times for computing mesh saliency.

Model	Patch and descriptor computing (s)	Queries selecting (s)	Saliency detecting (s)
Armadillo (34K)	161.93	3.52	3.32
Bimba (30K)	138.57	2.69	2.62
Dragon (30K)	161.33	2.78	3.92
Gargoyle (25K)	124.69	2.24	2.68
Bunny (14K)	70.50	1.02	0.48
Human (15K)	113.33	1.18	0.93

Algorithm 1. Manifold ranking based mesh saliency.

Input: A mesh M.
Output: A saliency map S.
Step 1: Oversegment M into patches and compute patch descriptors in Section 3.3.
Step 2: Compute the distinctness of each patch using Eq. (2).
Step 3: Construct the similarity matrix W using Eq. (3).
Step 4: Choose some of the most unsalient background patches as queries and determine the indicator vector y.
Step 5: Compute the patch saliency via manifold ranking using Eq. (1) with the above queries.
Step 6: Achieve a smooth saliency across vertices by splatting the patch saliency.

3.2. Saliency via manifold ranking

We describe the problem of detecting saliency via manifold ranking [26,27] as follows: given a number of patches as queries, all the patches are ranked based on their relevances to the given queries respecting the manifold structure of the descriptor space of patches. Specially, given a set of patches $X=\{x_1,\ldots,x_l,x_{l+1},\ldots,x_n\}$ where $\{x_i|i\le l\}$ are queries, the goal is to compute a ranking function $f=[f_1,\ldots,f_n]^T$ which assigns a ranking score f_i to each patch x_i. Let $y=[y_1,y_2,\ldots,y_n]^T$ denote an indicator vector, in which $y_i=1$ indicates x_i is a query and $y_i=0$ otherwise. Then a graph G is defined to depict the manifold structure of the patch descriptors. Let $W=[w_{ij}]_{n\times n}$ denote the similarity matrix of the graph and $D=diag\{d_{11},\ldots,d_{nn}\}$, where $d_{ii}=\sum_j w_{ij}$. The following cost function is defined:

$$Q(f)=\frac{1}{2}\left(\sum_{i,j=1}^{n} w_{ij}\left\|\frac{1}{\sqrt{d_{ii}}}f_i-\frac{1}{\sqrt{d_{jj}}}f_j\right\|^2+\mu\sum_{i=1}^{n}\|f_i-y_i\|^2\right). \quad (1)$$

The first term is a smoothing term which means the saliency of similar patches should be close, and the second term is a fitting term which means that the ranked saliency should not deviate too much from the queries. The two terms are balanced by the parameter μ. The resulted ranking function is the solution of $f^*=\arg\min Q(f)$, which can be written as $f^*=(I-\alpha S)^{-1}y$, where I is an identity matrix, $\alpha=1/(1+\mu), S=D^{-1/2}WD^{-1/2}$. Similar to [24], $f^*=(D-\alpha W)^{-1}y$ is used. Define $A=(D-\alpha W)^{-1}$ and set the

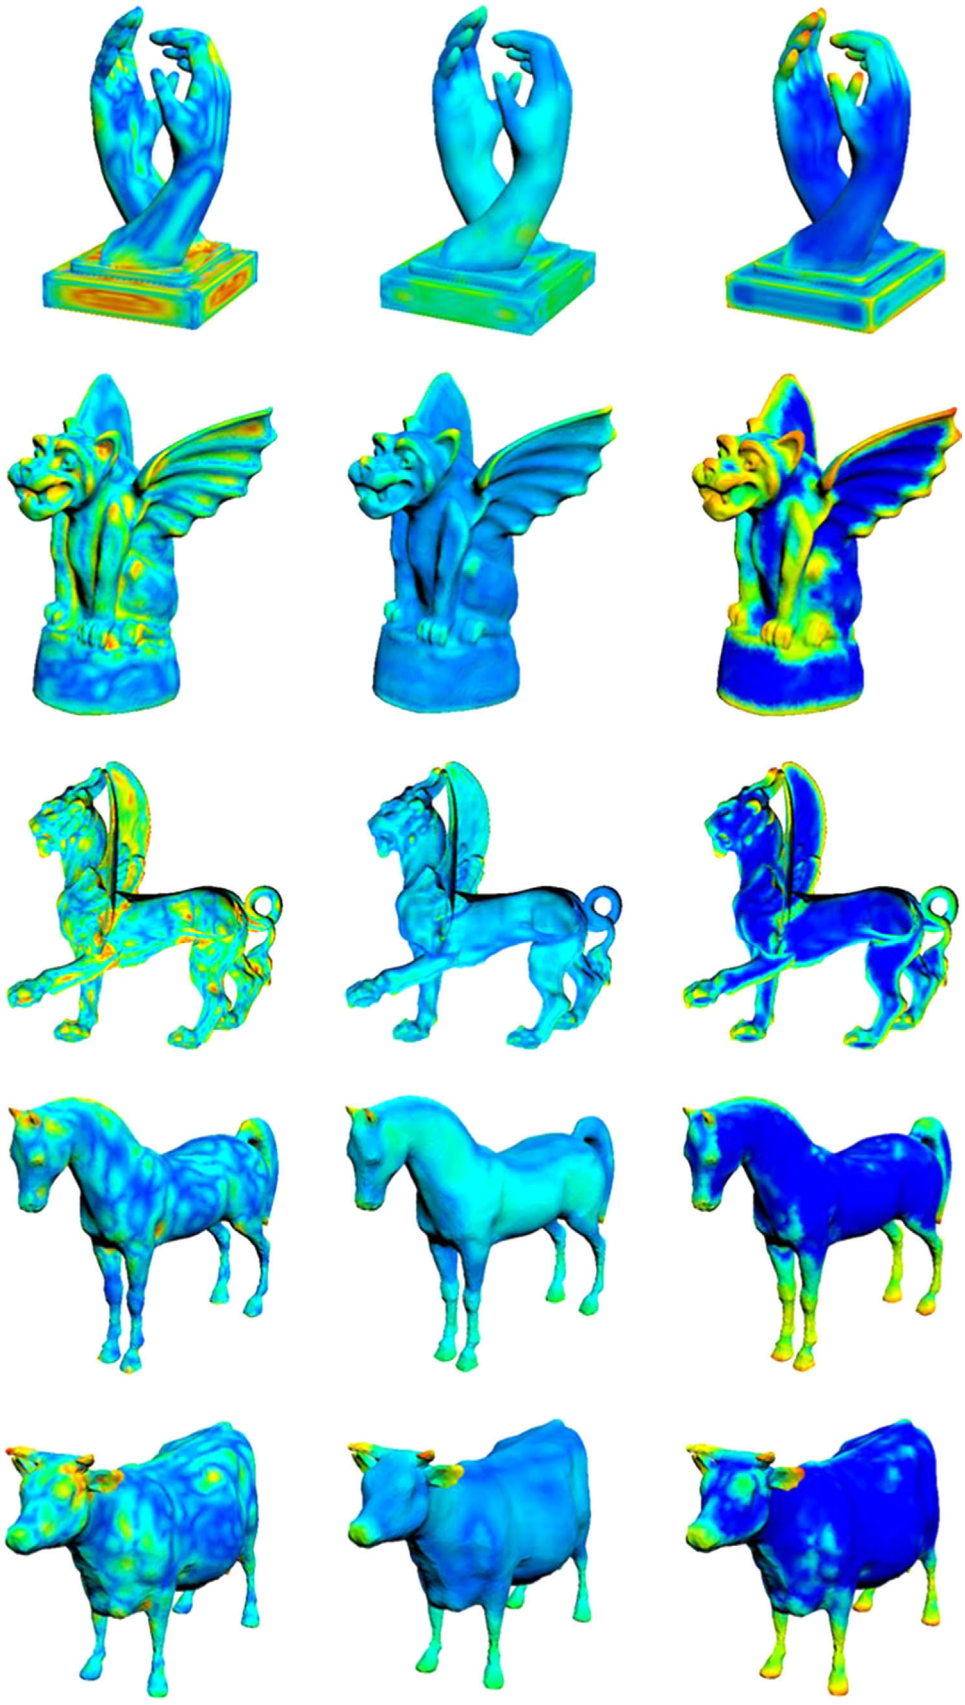

Fig. 8. Mesh saliency results of Lee et al. [2] (the left column), Wu et al. [3] (the middle column) and our method (the right column).

diagonal elements of A to 0, which means that the ranking score of a query is ranked by the other queries (except itself). This operation contributes to make our result more robust since a few unfaithfully estimated queries may be corrected by other queries.

3.3. Patch generation and patch descriptor

We employ the method of Wu et al. [3] to generate a set of patches P for any input mesh. Three parameters are used to control the number of patches. As can be seen in Fig. 2, our method is not sensitive to the patch number. The saliency maps in (a)–(c) of Fig. 2 are similar. Hence we use unified parameters in all the experiments.

Given the patches P, a proper patch descriptor is needed to compare different patches. We define the patch descriptor as the mean of the descriptors of the vertices within the patch. As known, the curvature computation is fairly sensitive to noise. Alternatively, Zernike coefficients [28] are adopted as the vertex descriptor, which is rotationally invariant and stable, as shown in Fig. 3.

For each vertex, the square sub-region of its tangent plane with radius r is considered. The sub-region is divided into a $16 * 16$ grid. A heightmap is calculated as the Euclidean distance from grid points to the mesh surface along the normal direction. It encodes the local shape surrounding the vertex. In order to achieve rotation

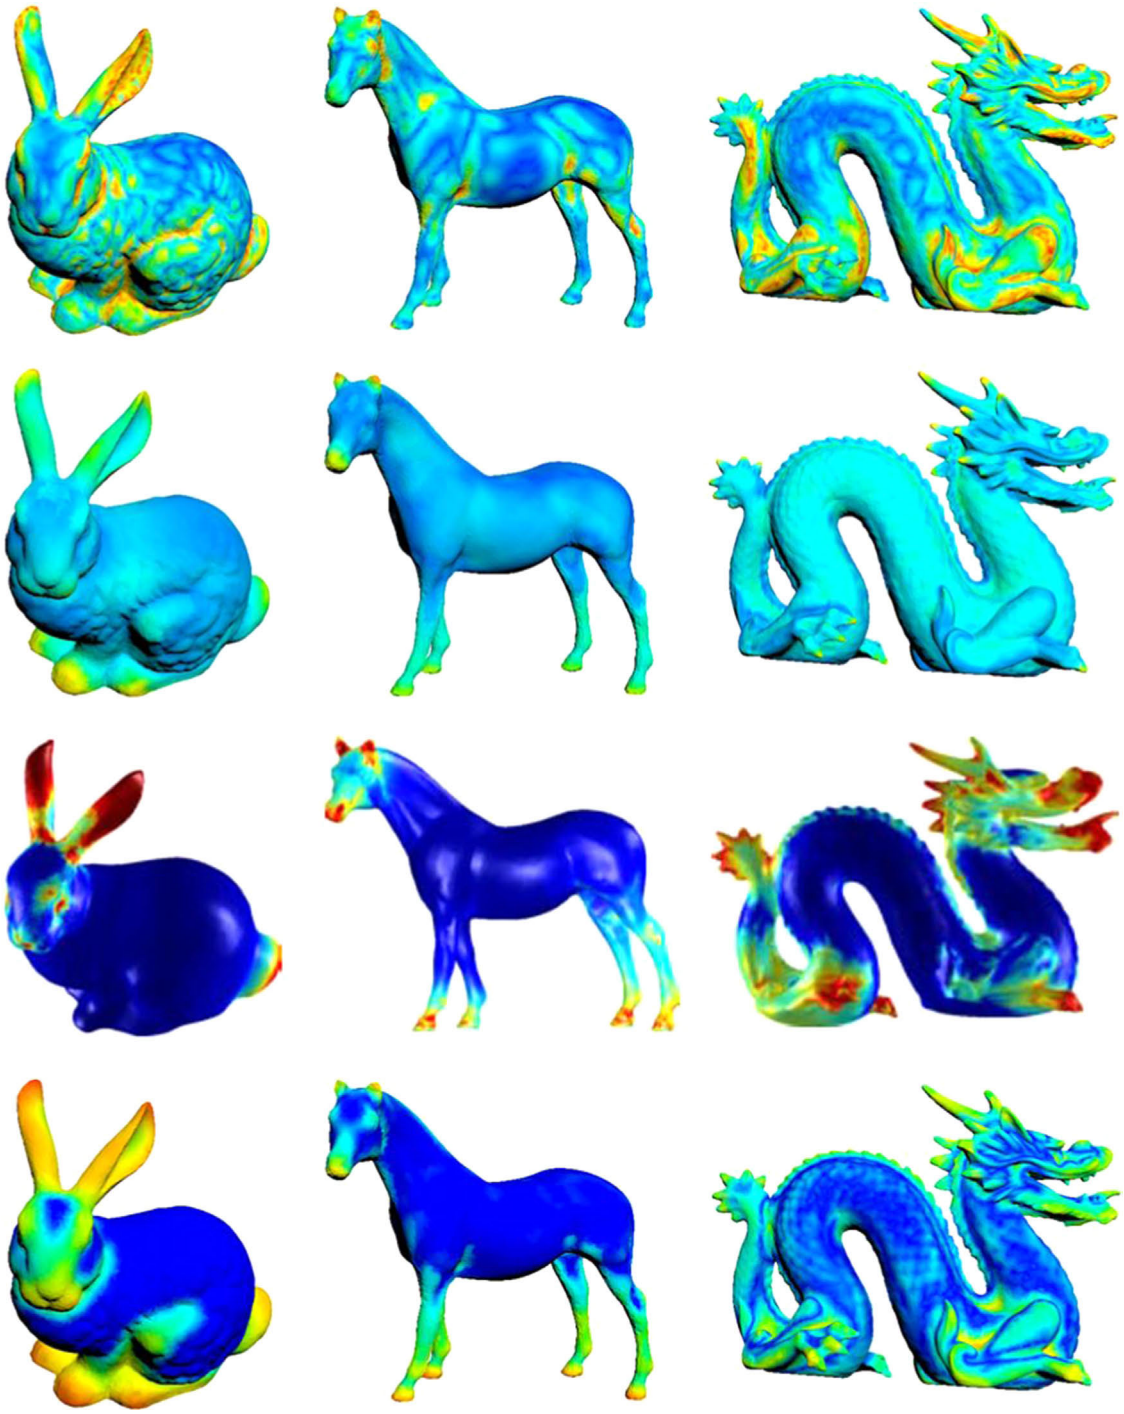

Fig. 9. Mesh saliency results of Lee et al. [2] (the top row), Wu et al. [3] (the second row), Leifman et al. [13] (the third row), and our method (the bottom row).

invariance, Maximo et al. [28] compute the Zernike-basis expansion of the heightmap. The descriptor depicts a wider range of surface shape as the radius r becomes bigger. Compared to multi-scale descriptors in [5,3], a single scale descriptor is enough to generate faithful results using our method. Furthermore, our method is insensitive to the scale, i.e. the radius r for computing the Zernike descriptor. As shown in Fig. 4, the mesh saliency is not influenced by the scale greatly. We set $r=3.0\,l$ in all experiments, where $l=0.5\%$ of the longest diagonal of the mesh's bounding box.

3.4. Local patch distinctness

To generate queries for the ranking process, we estimate local distinctness by computing the local contrast among the patches:

$$C(p)=\frac{1}{n}\sum_{q\in N_p} d(z_p,z_q), \tag{2}$$

where N_p denotes 3-ring neighborhood of patch p, n is the number of N_p, z_p and z_q are the Zernike descriptors of patches p and q,

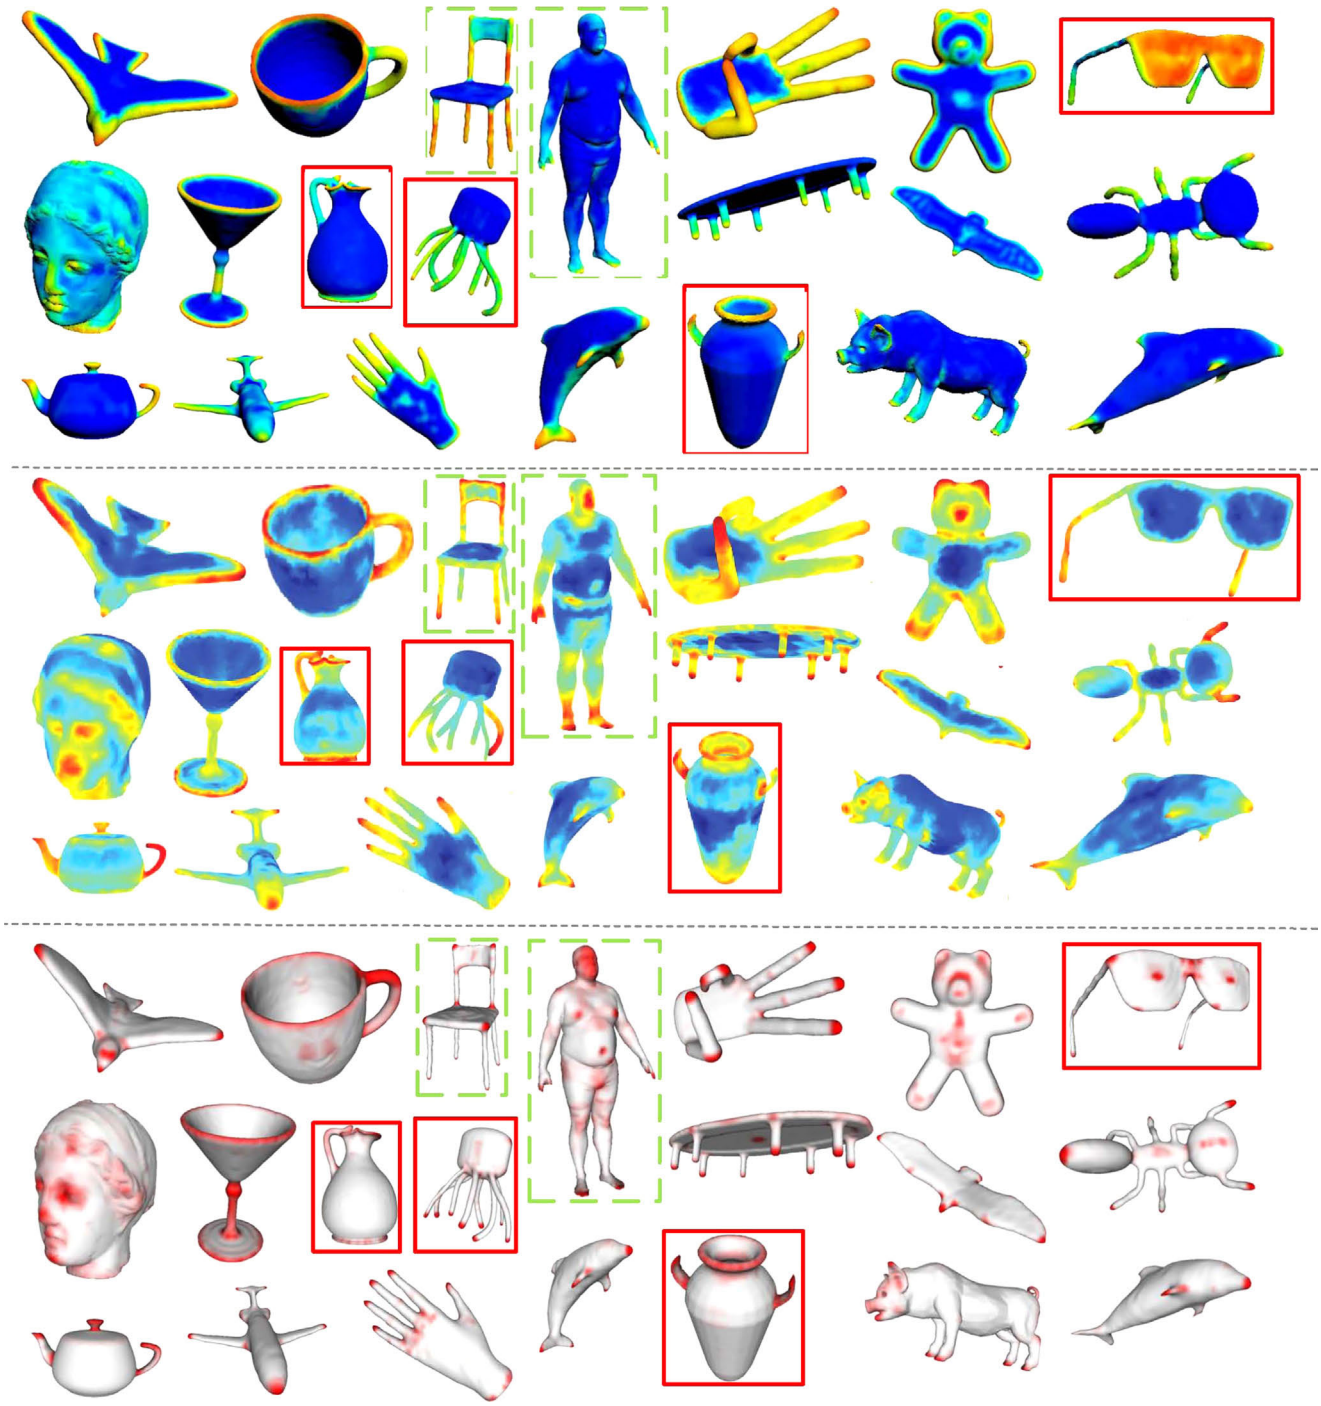

Fig. 10. Comparisons of our results (the top part) with spectral mesh saliency [23] (the middle part) and pseudo-ground truth [16] (the bottom part). If our results are more faithful, compared with pseudo-ground truth, the corresponding modes are marked in solid red boxes, otherwise in dashed green boxes. Our results are comparable with spectral mesh saliency for the rest of models in this figure. All models are courtesy of the Watertight Track of the 2007 SHREC Shape-based Retrieval Contest. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

respectively, and $d(\cdot,\cdot)$ is the Euclidean distance between the patch descriptors. More faithful local contrast can be obtained by using more rings at the cost of more computation time. In order to take a balance, numerical experiments suggest that 3-ring neighborhood is a good choice.

3.5. Graph construction

To get satisfactory results with less running time, we define the ranking graph $G=(P,E)$ on patch level, where P is a set of patches and E is the edge set. The choice of E and weights defined on it affect the final results to a great extent [24,25]. Yang et al. [24] define a two-ring graph first and introduce more edges by connecting any pair of patches on the four sides of a image for image saliency detection. They assume that the patches on the image sides tend to be background.

However, there is no corresponding assumption for 3D shapes. Furthermore, the nature between image saliency and mesh saliency is different. The salient regions of an image are usually concentrated, while those of a mesh are often dispersed. Different from the graph construction method in [24], G is built in a descriptor space in a self-adaption way. For each patch, we find those patches within the radius r_z and connect an edge between the current patch and any neighboring patch. r_z is estimated as follows. We compute feature distances between all patches and the most salient foreground patch, and sort them in ascending order. The distance corresponding to the maximal distance increment from the fifth distance is selected as r_z. Thus, even the most salient patches will not have a few edges. The similarity between two patches is defined by

$$w_{ij}=e^{-\|z_i-z_j\|/\sigma^2},\quad i,j\in P, \tag{3}$$

where z_i and z_j denote the descriptors of patches i and j, respectively. The descriptor distance is normalized to the range [0,1] and we set $\sigma=10^{-2}*r_z/d_z$, where d_z denotes the maximum descriptor distance between all patches.

Sidi et al. [29] incorporate the manifold structure of the shape descriptors when constructing the graph of a collection of mesh patches too. The diffusion in the descriptor space contributes to relate the parts independent of their poses, locations and cardinality. As demonstrated in the experiments, our ranking in the descriptor space reveals the saliency patches independent of their locations and cardinality too. As illustrated in Fig. 5, the saliency map via the self-adaption graph in the descriptor space is more faithful and close to the pseudo-ground truth than that via the two-loop graph in spatial domain.

3.6. Ranking with background queries

Similar to [2,3], vertices or patches with high local contrast are considered as saliency points or patches. Once the local distinctness C is computed using Eq. (2) and normalized, we can select some patches with the highest local distinctness as the most salient foreground queries, i.e. patches with local distinctness larger than some threshold. However a proper threshold is hard to select and the foreground queries are sensitive to noise, as illustrated in Figs. 6 and 7. Hence we select some of the most unsalient background patches as the queries, which are patches with local distinctness lower than 0.1, and all the patches will be ranked according to these queries. Unlike using the foreground queries, the ranking results are more robust for noise and a wide range of thresholds can be used to determine the background queries.

The ranking function f is calculated using Eq. (1), which means the relevances to the background queries. From this perspective,

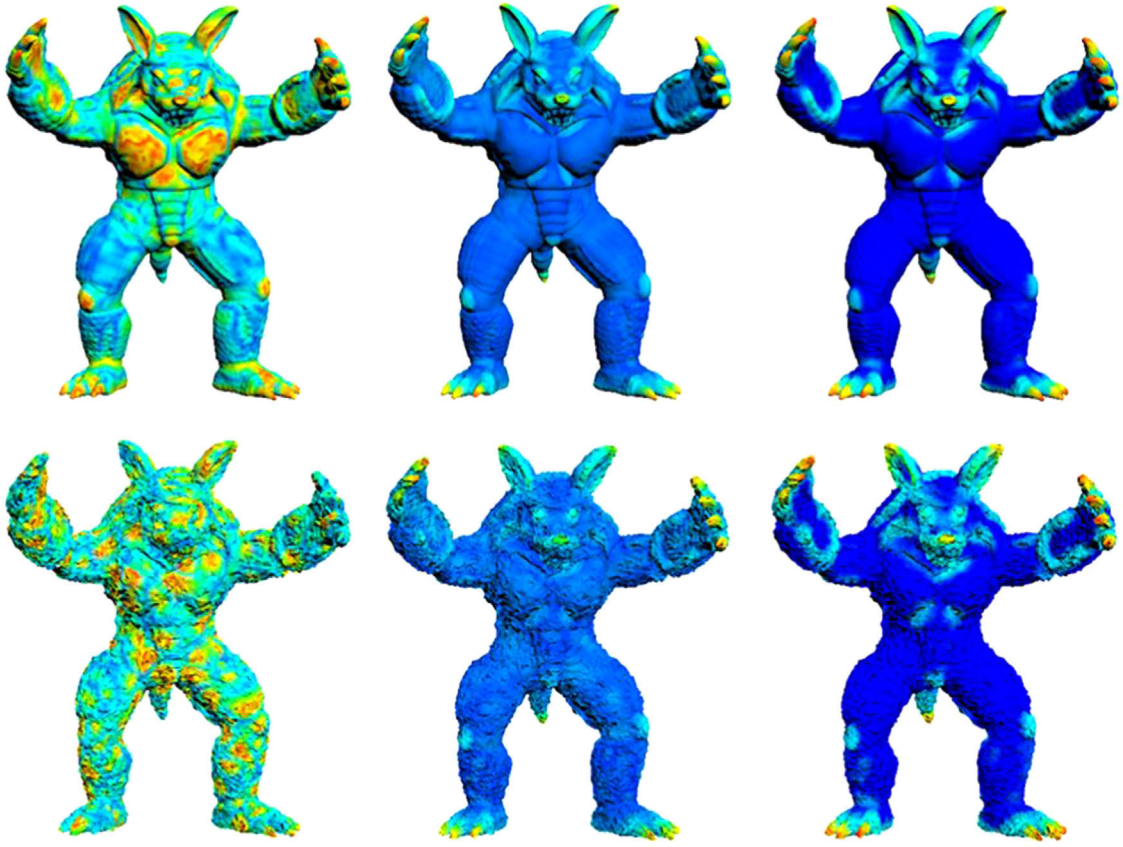

Fig. 11. Our method is more robust against noise. The top row shows mesh saliency results of Lee et al. [2], Wu et al. [3] and our method for the original Armadillo model. The second row shows the results of them for the model with 30% random noise relative to the average edge length.

we compute the saliency map as

$$S(p) = 1 - f(p). \quad (4)$$

Finally, a Laplacian smoothing procedure is employed to spread the patch saliency to each vertex.

4. Experimental results

In this section, we evaluate our method on a broad set of object categories and compare it with state-of-the-art methods, including Lee et al. [2], Wu et al. [3], Leifman et al. [13], Song et al. [23] and pseudo-ground truth [16]. The pseudo-ground truth is from data collected in a user study using a regression model trained based on meshes of the same class. All the experiments are run on a Intel Xeon E5630 2.53 GHz CPU with 12 G memory. The performance of our method is shown in Table 1. The saliency computation involves three steps, namely patches and descriptors computing, queries selecting and saliency detecting. For a model with 30K vertices, the first step requires about 146.63 s. The queries selecting and saliency detecting takes about 2.8 s and 3.13 s, respectively.

As we can see from Figs. 8 and 9, the results of Lee et al. [2] are greatly influenced by local changes of the curvature. Taking the Feline as an example, the result is disorganized. Wu et al. [3] generate better results. However, it is not easy to tune the parameters to get faithful results. Hence some salient regions may be not identified evidently, such as the eyes of the Cow and Horse in Figs. 8 and 9. Leifman et al. [13] choose 20% of the most distinct vertices and the extreme vertices as focus points. They

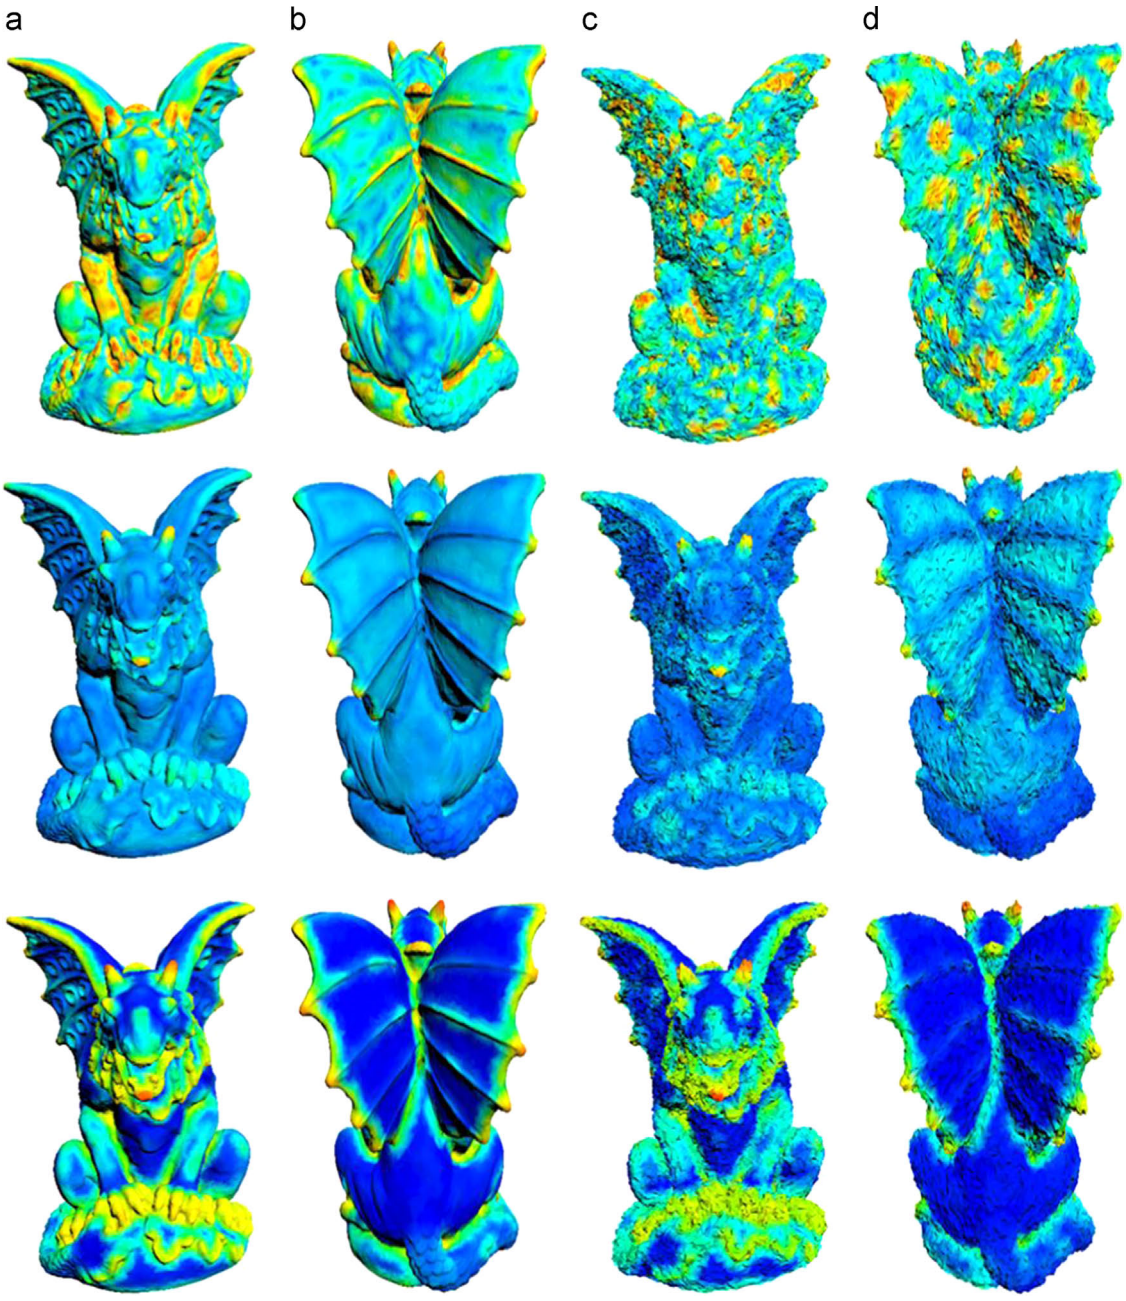

Fig. 12. Our method (bottom row) is more robust against noise than Lee et al. [2] (middle row) and Wu et al. [3] (bottom row). The left two columns are front and back view of saliency results for the original Gargoyle model. The right two columns are front and back view of saliency results for the Gargoyle model with 30% random noise relative to the average edge length.

may lose some saliency regions that do not contain focus points. As shown in Fig. 9, the feet of the Bunny are not obtained. Our method distinguishes some small salient regions better than the other three methods, such as the regions in the head of the Feline in Fig. 8 and the Horse in Fig. 9.

Comparisons of our results with Song et al. [23] and Chen et al. [16] are shown in Fig. 10. For the models in solid red boxes, such as the Vase, the Jar and the Glass, our results are more consistent with the pseudo-ground truth. While Song et al. [23] achieve better results for those models in dashed green boxes, such as the Human and the Chair. In general, our saliency results are largely comparable with them.

Figs. 11 and 12 show the saliency results when the input meshes contain noise. We add 30% (of the average edge length) random noise to perturb vertex coordinates. Method of Lee et al. [2] is not robust because curvature is highly affected by the noise. The saliency of the Armadillo and the noise contaminated Armadillo are dramatically different, see Fig. 11. Both Wu et al. [3] and our method compute mesh saliency in a global view. They are more robust against noise. However, as shown in Fig. 12, our method achieves more stable results especially for Gargoyle's wings.

5. Conclusion and limitation

This paper presents a novel method to compute mesh saliency via manifold ranking. It is robust against noise since it takes some of the most unsalient background patches as queries. By ranking in the descriptor space composed of the patch descriptors based on Zernike coefficients, we can reveal the salient regions independent of their locations and cardinality. Compared with various state-of-the-art multi-scale approaches, our simple single scale method generates comparable even more faithful results on a wide range of shapes, especially when the models are contaminated by noise.

One limitation of our method is that it does not incorporate any high-level priors. Taking the Desk in Fig. 10 for example, we fail to detect the center region of the desktop as salient region without the guiding from advanced semantic information. In the future, we plan to investigate semantic priors for 3D shapes and incorporate these priors into mesh saliency detection.

Acknowledgement

The authors sincerely thank all the anonymous reviewers for their valuable suggestions. Junjie Cao is supported by the NSFC Fund (No. 61363048) and the Fundamental Research Funds for the Central Universities (Nos. DUT13RC206, DUT13JS04). Xiuping Liu is supported by the NSFC Fund (Nos. 61173102, 61370143). Ligang Liu is supported by the National Natural Science Foundation of China (61222206) and One Hundred Talent Project of the Chinese Academy of Sciences.

References

[1] Yang Y-B, Lu T, Lin J-J. Saliency regions for 3D mesh abstraction. In: The 10th pacific rim conference on advances in multimedia information processing; 2009. p. 292–9.
[2] Lee CH, Varshney A, Jacobs DW. Mesh saliency. ACM Trans Graph 2005;24(3):659–66.
[3] Wu J, Shen X, Zhu W, Liu L. Mesh saliency with global rarity. Graph Models 2013;75(5):255–64.
[4] Yee H, Pattanaik S, Greenberg DP. Spatiotemporal sensitivity and visual attention for efficient rendering of dynamic environments. ACM Trans Graph 2001;20(1):39–65.
[5] Shilane P, Funkhouser T. Distinctive regions of 3D surfaces. ACM Trans Graph 2007;26(2):7.
[6] Kim Y, Varshney A. Persuading visual attention through geometry. IEEE Trans Vis Comput Graph 2008;14(4):772–82.
[7] Feixas M, Sbert M, González F. A unified information-theoretic framework for viewpoint selection and mesh saliency. ACM Trans Appl Percept 2009;6(1):1.
[8] Kim Y, Varshney A. Saliency-guided enhancement for volume visualization. IEEE Trans Vis Comput Graph 2006;12(5):925–32.
[9] Yamauchi H, Saleem W, Yoshizawa S, Karni Z, Belyaev A, Seidel H-P. Towards stable and salient multi-view representation of 3D shapes. In: IEEE international conference on shape modeling and applications; 2006. p.40.
[10] Gal R, Cohen-Or D. Salient geometric features for partial shape matching and similarity. ACM Trans Graph 2006;25(1):130–50.
[11] Li X, Guskov I. Multiscale features for approximate alignment of point-based surfaces. In: Symposium on geometry processing; 2005. p. 217–26.
[12] Gelfand N, Mitra NJ, Guibas LJ, Pottmann H. Robust global registration. In: Symposium on geometry processing; 2005. p. 197–6.
[13] Leifman G, Shtrom E, Tal A. Surface regions of interest for viewpoint selection. In: IEEE conference on computer vision and pattern recognition; 2012. p. 414–21.
[14] Itti L, Koch C, Niebur E. A model of saliency-based visual attention for rapid scene analysis. IEEE Trans Pattern Anal Mach Intell 1998;20(11):1254–9.
[15] Castellani U, Cristani M, Fantoni S, Murino V. Sparse points matching by combining 3D mesh saliency with statistical descriptors. Comput Graph Forum 2008;27(2):643–52.
[16] Chen X, Saparov A, Pang B, Funkhouser T. Schelling points on 3D surface meshes. ACM Trans Graph 2012;31(4):29.
[17] Song R, Liu Y, Martin RR, Rosin PL. 3D point of interest detection via spectral irregularity diffusion. Vis Comput 2013;29(6–8):695–705.
[18] Guy G, Medioni G. Inference of surfaces, 3D curves, and junctions from sparse, noisy, 3d data. IEEE Trans Pattern Anal Mach Intell 1997;19(11):1265–77.
[19] Itti L, Koch C, Niebur E. A model of saliency-based visual attention for rapid scene analysis. IEEE Trans Pattern Anal Mach Intell 1998;20(11):1254–9.
[20] Mantiuk R, Myszkowski K, Pattanaik S. Attention guided mpeg compression for computer animations. In: The 19th spring conference on computer graphics; 2003. p. 239–44.
[21] Cheng M-M, Zhang G-X, Mitra NJ, Huang X, Hu S-M. Global contrast based salient region detection. In: IEEE conference on computer vision and pattern recognition; 2011. p. 409–16.
[22] Hou X, Zhang L. Saliency detection: a spectral residual approach. In: IEEE conference on computer vision and pattern recognition; 2007. p. 1–8.
[23] Song R, Liu Y, Martin RR, Rosin PL. Mesh saliency via spectral processing. ACM Trans Graph 2014;33(1):6.
[24] Yang C, Zhang L, Lu H, Ruan X, Yang M-H. Saliency detection via graph-based manifold ranking. In: IEEE conference on computer vision and pattern recognition; 2013. p. 3166–73.
[25] Liu R, Cao J, Lin Z, Shan S. Adaptive partial differential equation learning for visual saliency detection. In: IEEE conference on computer vision and pattern recognition; 2014. p. 3866–73.
[26] Zhou D, Weston J, Gretton A, Bousquet O, Schölkopf B. Ranking on data manifolds. In: Advances in neural information processing systems, vol. 16; 2004. p. 169–76.
[27] Zhou D, Bousquet O, Lal TN, Weston J, Schölkopf B. Learning with local and global consistency. In: Advances in neural information processing systems, vol. 16; 2004. p. 321–8.
[28] Maximo A, Patro R, Varshney A, Farias R. A robust and rotationally invariant local surface descriptor with applications to non-local mesh processing. Graph Models 2011;73(5):231–42.
[29] Sidi O, van Kaick O, Kleiman Y, Zhang H, Cohen-Or D. Unsupervised co-segmentation of a set of shapes via descriptor-space spectral clustering. ACM Trans Graph 2011;30(6):1–10.