Orienting Raw Point Sets by Global Contraction and Visibility Voting

Junjie Cao[⋆1] Ying He[†] Zhiyang Li[⋆] Xiuping Liu[⋆] Zhixun Su[⋆]
[⋆]Department of Applied Mathematics, Dalian University of Technology, China
[†]School of Computer Engineering, Nanyang Technological University, Singapore

Abstract

We present a global method for consistently orienting a defective raw point set with noise, non-uniformities and thin sharp features. Our method seamlessly combines two simple but effective techniques—constrained Laplacian smoothing and visibility voting—to tackle this challenge. First, we apply a Laplacian contraction to the given point cloud, which shrinks the shape a little bit. Each shrunk point corresponds to an input point and shares a visibility confidence assigned by voting from multiple viewpoints. The confidence is increased (resp. decreased) if the input point (resp. its corresponded shrunk point) is visible. Then, the initial normals estimated by principal component analysis are flipped according to the contraction vectors from shrunk points to the corresponding input points and the visibility confidence. Finally, we apply a Laplacian smoothing twice to correct the orientation of points with zero or low confidence. Our method is conceptually simple and easy to implement, without resorting to any complicated data structures and advanced solvers. Numerous experiments demonstrate that our method can orient the defective raw point clouds in a consistent manner. By taking advantage of our orientation information, the classical implicit surface reconstruction algorithms can faithfully generate the surface.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.5]: Computational Geometry And Object Modeling—

Keywords: orientation; raw points; surface reconstruction; constrained Laplacian smoothing

1. Introduction

The last decade has witnessed the rapid development of 3D scanning devices. Although these devices are capable of generating highly dense samples with improved precision, the acquired point sets are inevitably defect-ridden [18]. The typical defects include noise, outliers, non-uniformities and holes. Furthermore, raw point sets without inside/outside information, or even surface normals, have drawn increasing attention due to the ever-broadening range of general digitizing devices that do not provide such properties. Reconstructing surfaces from raw points thus remains critical concerns [5, 20, 8, 1, 21, 22, 18]. Consistently oriented normals are critical for surface reconstruction, since they define the reconstructed surface to the first order and identify the inside/outside information. The state-of-the-art reconstruction algorithms [5, 19, 20, 12] may produce poor quality results without consistent orientation [9, 14]. This issue can be solved if reliable oriented normals are available.

As pointed out in [1, 18], robust orientation is as difficult as reconstructing the whole surface itself. The existing orientation approaches can be roughly classified into three classes: the photometric stereo approach, the geometric approach and the volumetric approach. The photometric stereo approaches [15, 24] are delicate processes suffering from a non-ideal acquisition environment, such as specular reflections and shadows. The volumetric approaches [28, 17, 26, 10, 4, 18] are robust against noise and outliers at the expense of high computation cost. The large holes also pose challenges to most of these approaches. In general, the geometric approaches [7, 27, 6, 9, 14] usually work locally and hence are more efficient than photometric stereo and volumetric approaches. For example, the most frequently used method [7] and its variants [27, 6] try to propagate local unreliable information to deduce the global property - the surface orientation. Hence they suffer from noise, outliers, sparsity and thin sharp features. Recently, Huang *et al.* [9] consolidated the defective point set to produce a set of de-noised and uniformly distributed points over the original point set first. Then they conducted a tangential-direction-favored normal propagation scheme to deal with nearby surface sheets. However, the down-sampling strategy they employed may further drop limited information of highly sparse regions, thus giving rise to incorrect orientation (see (b) of Fig. 1). Liu and Wang's approach (ORT) [14] orients sparse point clouds by making use of a coarse surface reconstruction with a modified adaptive spherical cover scheme (ASC). Compared to other geometric approaches, it is more robust against the above mentioned defects. However, the ORT method may generate inconsistent orientations for some isolated, small regions (see (c) of Fig. 1) due to its local and coarse reconstruction.

To address the above issues, we present a global geometric approach for consistent orientation from defective raw points. Our method seamlessly combines two effective techniques—

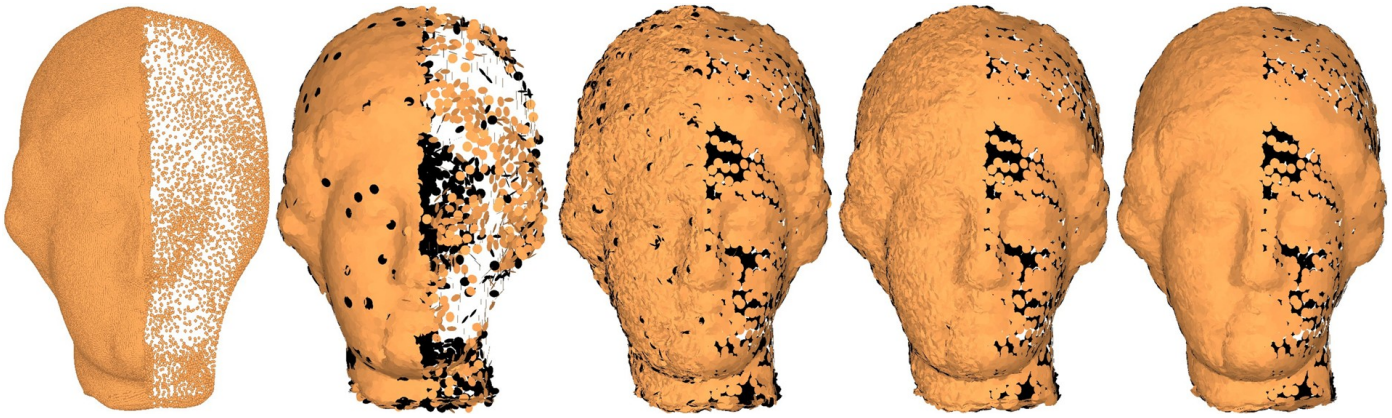

(a) 72.5K (b) Cons (c) ORT (d) BOT (e) Ours

Figure 1: Orientation results of the noisy, non-uniformly distributed point clouds of Venus model (see (a)) via different approaches. (b) Cons: the point clouds consolidation approach [9]. (c) ORT: the adaptive spherical cover approach [14]. (d) BOT: the binary orientation tree [4]. (e) Our approach by constrained Laplacian smoothing and visibility voting. The back-facing points are rendered in black splats.

constrained Laplacian smoothing and visibility voting of point set—and involves four main steps (see Fig. 2). Taking a raw point set as input, we estimate its initial normals by principal component analysis (PCA), or other more advanced algorithms (such as [13]). The input points are shrunk via constrained Laplacian smoothing. We compute the visibility of both the shrunk and input points via the hidden point removal (HPR) operator [11]. The smoothness of the shrunk points increases the chance of being viewed, and they compensate for sparse regions in some degree to avoid uncovering the model's inner region. To handle holes and occlusions, multiple viewpoints are leveraged to vote to assign a confidence level to each point. Finally, a visibility confidence-weighted smoothing and a majority-determined confidence-weighted smoothing are leveraged respectively to orient the normals globally. Our approach is effective and simple without resorting to any complicated data structures and advanced solvers. The experimental results show that our method is robust aganist noise, non-uniformities, sparsity and thin sharp features by inferring the consistent orientation globally. Our method can also handle models with large holes that are challenging to most of the volumetric approaches. We demonstrate the efficacy of our method by reconstructing high-quality surfaces using our consistent orientation information.

2. Related Work

It is technically challenging to accurately orient normals from a defective raw point set without any prior knowledge of the surface. A large body of literature is devoted to addressing this challenge. Due to the page limit, we review only the volumetric and the geometric methods, which are closely related to our approach.

Volumetric methods tag in/out information or sign an unsigned distance field based on some volumetric representations. Generally, volumetric methods handle noise and outliers very well, at the expense of high computational cost. Many methods [28, 17, 26, 10] work for the regular volumetric grid enclosing the point set. To improve the efficiency of the above methods, Chen *et al.* applied a visibility checking approach to tag an octree [4]. The in/out tags of the corners are determined according to their visibility relative to the input point set from a set of pre-defined viewpoints. However, a corner is tagged as out with only one view, which restricts the algorithm to complete point clouds. Our visibility confidence contributes to this. Recently, Mullen *et al.* [16] signed a robust unsigned distance function by utilizing stochastic ray tracing and global smoothing. Although this method is insensitive to noise and outliers, it is complicated, with many heuristic parameters.

Geometric methods work directly on point clouds. A frequently used geometric orienter was proposed by Hoppe *et al.* [7]. Starting from a seed point, it flips inconsistent normals of its neighbors along a minimum spanning tree. Xie *et al.* [27], Guennebaud *et al.* [6], and Huang *et al.* [9] improved [7] by modifying the seeds and propagation criterion. Although the method of Huang *et al.* [9] is generally effective, serious sparsity gives rise to incorrect propagation since the down-sampling strategy they employed further drops the limited information of highly sparse regions [4, 14]. Apart from the above propagation approaches, Liu and Wang [14] orient sparse point clouds by making use of a coarse surface reconstruction with a modified adaptive spherical cover scheme (ASC). The sphere detection and splitting operation make the method robust to non-uniform sparsity. However, the cleaning process they used may lead to degenerated reconstruction, which gives rise to inconsistent orientation in turn. Although all the above orienters can deal with raw point sets with sparsity, incompleteness and noise to some degree, they are limited by the local property, i.e., no global information is involved to determine the orientation. On the contrary, our algorithm, composed of constrained implicit Laplacian smoothing and HPR, infers the orientation globally to achieve a considerable improvement over the previous methods.

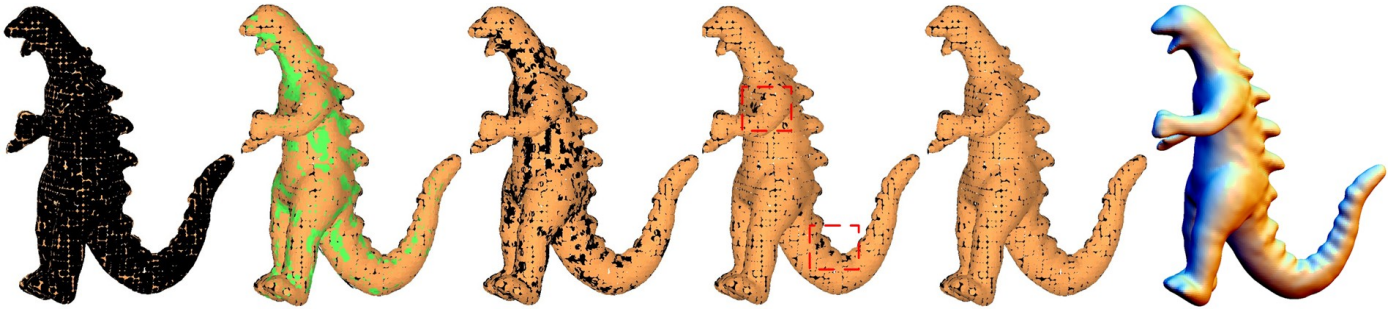

(a) (b) (c) (d) (e) (f)

Figure 2: Overview of our algorithm. (a) Input model with wrong initial orientation. (b) The input points (orange) and shrunk points (green). (c) Initial orientation estimation by visibility checking of the input and shrunk points. (d) Orientation after the visibility confidence weighted smoothing (Locally inconsistent orientations are highlighted in red rectangles). (e) Consistent orientation after the majority-determined confidence weighted smoothing. (f) The RBF-based reconstruction [19] of the model with oriented normals by our method.

3. Approach

Let $P = (X, N)$ denote a point set such that a point is given by the pair $p_i = (x_i, n_i)$, where $x_i \in X \subseteq \mathbb{R}^3$ is the point location and $n_i \in N \subseteq \mathbb{R}^3$ is the normal. The goal of our work is to compute a consistent orientation, *i.e.* making all the normals N outward-facing. In the following subsections, we distinguish among an arbitrary point set $P = (X, N)$, the input point set $P = (X, N^{in})$, the output point set $P = (X, N^{out})$, the point set with estimated orientation by visibility voting $P = (X, N^{vv})$, the point set with orientations computed by the visibility confidence weighted smoothing $P = (X, N^{sv})$, and the point set with orientations computed by the majority-determined confidence weighted smoothing $P = (X, N^{sm})$. Furthermore, we define a normal flipping operator $Flip(N^1, N^2) = \{sign(n_i^1 \cdot n_i^2) n_i^1\}$, which flips normal set N^1 according to N^2.

We first contract P a bit to a shrunk point set $S = (X')$, which has a one-to-one correspondence with P by their indices. Then we define an orientation indication vector set $V = \{v_i = x_i - x_i'\}$ from the S to P. If s_i is inside the input shape, v_i is most likely outward. By checking the visibility of P and S from multi-viewpoints, we assign each pair of points, p_i and s_i, a visibility confidence c_i, and flip a portion of initial normals N^{in}, according to V and C to generate a coarse orientation estimation N^{vv}. Next, we conduct a visibility confidence-weighted global smoothing to orient the undetermined orientations of points with zero-confidence, and generate N^{sv}. To reduce the interference of noise, sparsity and positions and directions of observation, we identify the largest region with consistent orientation, increase the confidences of all points in that region, and apply the smoothing with the new confidence again to produce N^{sm}. Finally, we compute the final consistent orientation $N^{out} = Flip(N^1, N^{sm})$. The following subsections detail the above steps.

3.1. Contracting

The volume-reducing characteristic of Laplacian smoothing [25] has been exploited to perform a robust curve skeleton extraction for triangular meshes [2] and point sets [3]. As stated in [2], the input geometry can be shrunk into a nearly zero-volume shape by iteratively solving the linear system:

$$\begin{bmatrix} W_L L \\ W_H \end{bmatrix} X' = \begin{bmatrix} 0 \\ W_H X \end{bmatrix}, \tag{1}$$

where L is an $n \times n$ Laplacian matrix, and W_L and W_H are the diagonal weight matrices balancing the contraction and attraction constraints, the i-th diagonal element of W_L (resp. W_H) is denoted $W_{L,i}$ (resp. $W_{H,i}$). The solution of this system minimizes the quadratic energy [23]:

$$\|W_L L X'\|^2 + \sum_i W_{H,i}^2 \|x_i' - x_i\|^2, \tag{2}$$

where the first term removes geometric details along the normal directions using implicit Laplacian smoothing, and the second preserves shape geometry during contractions.

We expect the orientation indication vector v_i to point outward if s_i is located inside the input shape. Although more shrunk points are likely to be located inside with more iterations, there is no guarantee that all shrunk points S, are inside the input shape. Excessive contraction is prone to make v_i inward, and even s_i is inside, near the relatively small and concave features. Hence we set $W_H = 2I$, where I is an $n \times n$ identity matrix, and contract $P = (X, N^{in})$ just once to avoid excessive contraction of S. The scale of the Laplacian coordinate is proportional to a point's neighborhood extent (under the same local neighborhood shape); the contraction forces from the Laplacian equations are smaller for denser models. Thus to handle models of different sizes and resolutions, we set $W_L = 1/(5E)$ for all experiments in this paper, where E is the mean neighborhood extent of the point cloud.

To handle nearby surface sheets, the k-nearest neighbors (kNNs), with default value $k = 10$, are first employed to construct the Laplacian operator. Then, the distance measure (3) defined in [9] is applied to filter out $k/2$ furthest neighbors from the kNNs.

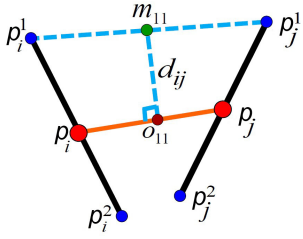

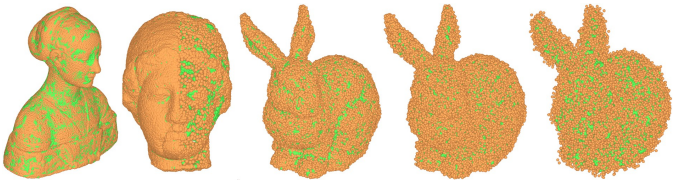

Figure 3: The implicit global Laplacian contraction is robust against various defects. From left to right are contraction results of models with open boundaries, non-uniform sampling, and 0.5%, 1% and 3% noises respectively. The contracted shapes are rendered in green and the original shapes in orange.

$$\mathcal{D}_{ij} = 1 - |n_i \cdot n_j| \frac{\max_{r,s\in\{1,2\}} \|m_{rs} - o_{rs}\|}{1 + \|p_i - p_j\|}, \tag{3}$$

where m_{rs} is the midpoint of the line segment $\overline{p_i^r p_j^s}$, $r, s \in \{1, 2\}$; points p_i^1, p_i^2, p_j^1 and p_j^2 are a unit distance away from p_i and p_j along n_i and n_j; and o_{rs} is the perpendicular projection of m_{rs} onto the estimated tangent line $\overline{p_i p_j}$ or its extensions.

Empirically, the type of Laplacian has no obvious influence on the final orientation. Hence we use the combinational weights Laplacian for simplicity, without projecting the neighbors to a tangential plane and computing a local Delaunay triangulation. It is still robust against various defects as illustrated in Fig. 3. The choice of other parameters follows the description in [3].

3.2. Tagging Confidence by Visibility Checking

After the contraction, point sets P and S are interweaved (see (b) of Fig. 2). Whether s_i is inside/outside the input shape can be examined by visibility voting via the hidden point removal (HPR) operator [11, 16]. Given a viewpoint outside of the input shape, if s_i is visible, v_i is most likely inward. Conversely, if p_i is visible, v_i is most likely outward. However, this property does not hold when a hole exists. Also, if the region is very sparse, both p_i and s_i can be visible, and we cannot discern which one is more reliable.

To reduce the interference of the above two cases, we observe the model from various viewpoints and each viewpoint votes to assign a confidence to each visible point p_i or s_i. We first set $c_i = 0, i \in [1, n]$. If p_i (resp. s_i) is visible from a viewpoint, let $c_i = c_i + 1$ (resp. $c_i = c_i - 1$). Thus if p_i and s_i are visible simultaneously, the two times of voting cancel each other out. After checking all the pre-defined viewpoints, the confidence c_i tells us whether s_i is inside/outside and how reliable it is.

As illustrated in Fig. 4, the more viewpoints are employed to vote, the more faithful an orientation is deduced. In our experiments, we normalize the input model into a bounding box centered at the origin, with the length of the main diagonal equal to 1.6. Hence the viewpoints can be uniformly distributed on the surface of a cube (noted as viewcube) centered at the origin, with the Cartesian coordinates of the vertices, which are $(\pm 1, \pm 1, \pm 1)$. Obviously, the more viewpoints (excluding those facing a hole or open boundary) that are leveraged, the more qualified confidences (more $|c_i \neq 0|$ and the more reliable sign of each c_i) we will obtain. Since our method does not depend on a volumetric structure, it is impossible to find proper viewpoints as much as was done in [4]. Actually, resorting to our final global smoothing procedure, distributing the viewpoints uniformly on the viewcube can produce faithful orientation in our experiments. Empirically, 14 viewpoints are taken, which are the centers of the six faces and the eight corners of the viewcube.

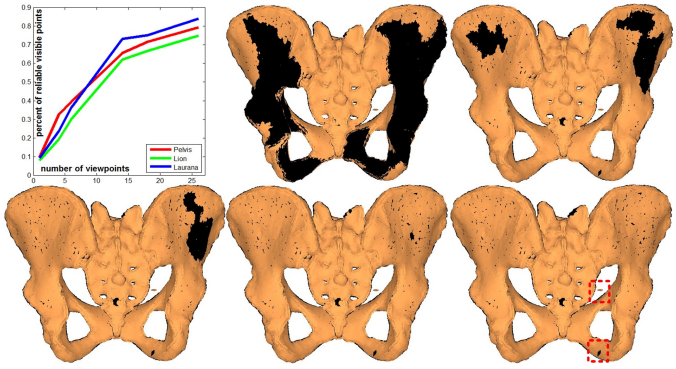

Figure 4: With the increasing number of the viewpoints, more visible points can be located in a reliable manner, which leads to more faithful orientation. The curve plot in the top left shows the relation between the number of view points and visible points of Pelvis, Lion and Laurana. The subfigures from left to right, top to bottom are the orientations of the Pelvis model generated by 1, 4, 6, 18 and 26 view points, respectively. The outliers are highlighted in the red boxes.

3.3. Initial Orienting by Visibility Voting

With the confidences C, we estimate the orientation N^{vv} by flipping n_i^{in} if $c_i n_i^{in} \cdot v_i < 0$. The orientations of points with $c_i = 0$, even with very low value, may not be reliable. We apply a global smoothing procedure guided by the visibility confidence to adjust those unreliable orientations. Since we contract the input point set using equation (1), we can reuse the Laplacian matrix to speed up the process. The smoothed orientation satisfies the linear equation

$$\begin{bmatrix} L \\ W(C) \end{bmatrix} N^{sv} = \begin{bmatrix} 0 \\ W(C)N^{vv} \end{bmatrix}, \tag{4}$$

where $W(C)$ is a weighting function that maps the confidence c_i to a non-negative weight. $W(c_i) = c_i^2$ works well in all our experiments.

3.4. Final Orienting by Consistency Voting

As demonstrated in Fig. 2 (d), local inconsistent orientations may still exist after the smoothing operation via visibility voting, especially when the input points are noisy, extremely sparse or too complicated relative to the pre-defined viewpoints. Hence we correct these minority orientations by making them consistent with the orientation of the majority via the same Laplacian smoothing procedure again, with updated attraction weights $W(C^{maj})$:

$$\begin{bmatrix} L \\ W(C^{maj}) \end{bmatrix} N^{sm} = \begin{bmatrix} 0 \\ W(C^{maj})Flip(N^{in}, N^{sv}) \end{bmatrix}, \tag{5}$$

where C^{maj} are majority-determined confidences.

The majority can be identified by clustering the points P according to the current orientation $Flip(N^{in}, N^{sv})$. C^{maj} of the minority are set to zero and those of the majority are set to the maximum visibility confidence. Finally, we flip the smoothed normals N^{sm} again to generate the consistent orientation $N^{out} = Flip(N^{in}, N^{sm})$. Even the default number of the viewpoints we employed is small. It works fairly well in practice and the computed orientation N^{out} is consistent after two times of smoothing procedures in our experiments.

4. Experimental Results

To demonstrate the performance of our method, this section presents more experimental results on a wide range of raw point sets with noise, non-uniformities, sparsity, open boundaries and thin sharp features. We also demonstrate that consistent orientation of point clouds facilitates surface reconstruction algorithms. For comparison, we chose three state-of-the-art orienters: Cons [9], ORT [14] and BOT [4]. In all the experiments, our method uses the default parameters depicted previously. Those of the Cons orienter are chosen in a trial-and-error way - we try to maximize the sampling rates and apply the normal propagation and orientation-aware PCA iteration multiple times to estimate the normals more accurately. All results of BOT are provided by its authors.

Non-uniformity and sparsity: Non-uniformity and sparsity (Figures 1 and 5) always challenge conventional normal orienters. As stated in [14, 4], the Cons method is not suitable to handle non-uniform sparsity. The down-sampling strategy, especially in the highly sparse region, exacerbates the case. The BOT orienter proceeds more globally and produces more faithful orientations than the Cons orienter. However, it has issues with holes formed by serious non-uniform sparsity (see the mouth of the horse model in Fig. 5), and may generate local incorrect orientation near thin surface structures (see the feet of the horse model in Fig. 5), due to the visibility checking strategy they employed. Both ORT and our approach can handle these defects easily.

Noise: Our approach is robust against noise, as shown in Figures 1 and 6, and the third column of Fig. 11. In contrast, the ORT orienter may produce some isolated incorrect orientations using the default parameters, which may result in small artifacts in the reconstructed surface (see Fig. 10). In Fig. 6, noise is added to each point by randomly displacing the position by a fixed distance (0.5%, 1% and 3% of the length of the main diagonal of the bounding box of the point set). The ORT orienter handles noise by increasing the default spherical covering radius. 0.01, 0.05 *and* 0.1 are carefully chosen to deal with the increasing noise level in (d), (e) and (f) of Fig. 6 respectively. However, the ORT orienter still may produce inconsistent orientation due to the local property of the approach. Note that a too-large default spherical covering radius may make the generated coarse mesh oversimplified, which causesincorrect orientation. As illustrated in Fig. 6, a faithful surface can be recovered using the orientation of our method if the noise level is 0.5%. When the noise level increases to 1%, orientations of thin feature regions will be somewhat polluted, leading to small reconstruction artifacts (see the ears of the Bunny model in (h) of Fig. 6). A 3% rate of noise seriously damages the input shape, and causes the normals estimated by PCA to vary drastically. Although the recovered surface is not faithful, the generated orientation is still consistent. Furthermore, the third column of Fig. 11 also demonstrates that our method can tackle point sets with a lot of structured noise.

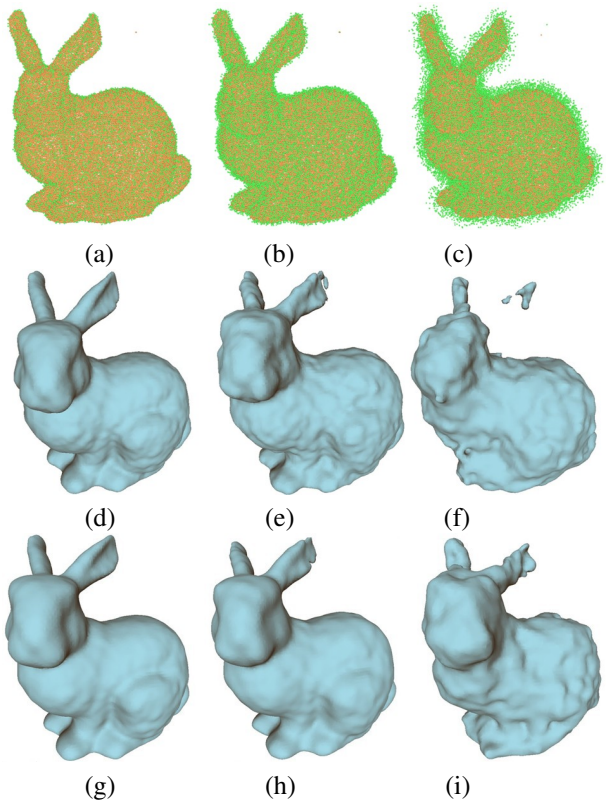

Figure 6: Our approach is robust to noise. The top row shows three noisy Bunny models with 0.5%, 1% and 3% random noise respectively (The original Bunny model is rendered in orange points and the noisy model in green points). The middle and bottom rows are the corresponding Poisson reconstruction [12] via orientation by the ORT orienter and our approach.

Holes and open boundaries: Fig. 7 demonstrates that both the ORT and our methods handle holes and open boundaries gracefully. The man model with 128K points in Fig. 7 is captured with only two column laser scanners. Hence the points are roughly two pieces (front and back), and contain sparsity, holes, or even large boundaries. As mentioned in subsection 3.2, only 14 standard viewpoints are leveraged to generate faithful orientations of the two models, without taking the boundaries' position and facing information into consideration. The ORT method produces few isolated errors near the boundary and sharp features for the Laurana model with 50K points. In general, it is very robust for various defective point sets.

Thin sharp features and nearby surface sheets: Thin features and nearby surface sheets are always difficult to deal with for normal estimation and orientation. All the state-of-the-art orienters tackle such cases to some degree; minor locally inconsistent orientations are often inevitable even with delicate parameter tuning. Due to the strategy of neighbor points selec-

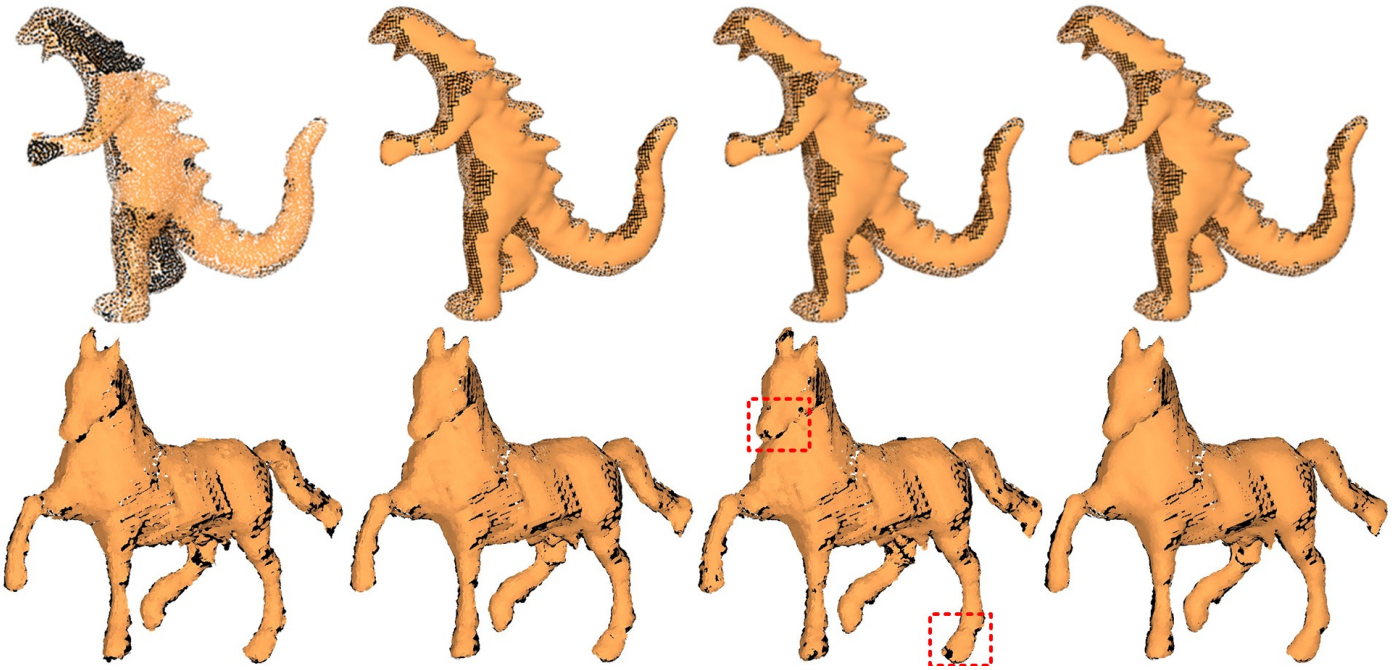

Figure 5: Orientations of Dinosaur (36K points) and Horse (18K points) models with non-uniformities and sparsity. From the left to right are results of the Cons, ORT, BOT and our approaches.

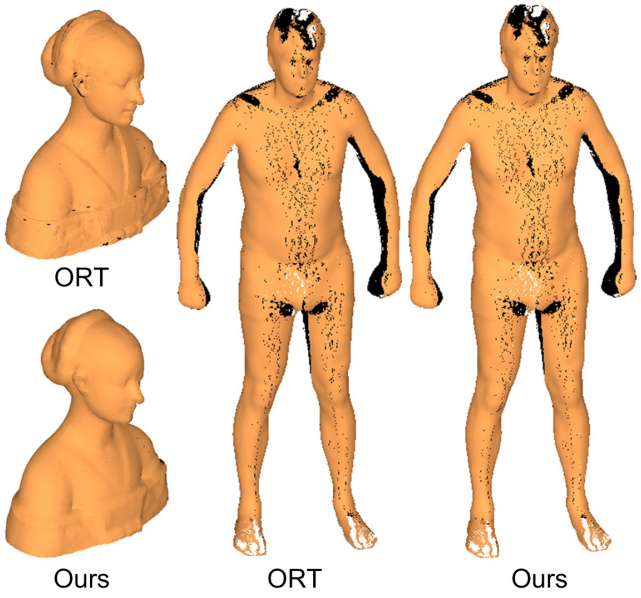

Figure 7: Orientations of point sets with open boundaries.

tion and the global nature of our scheme, we can orient models with such characteristics and achieve slightly better results, as illustrated in Figures 8 and 9.

Orientation and reconstruction: Sparsity, non-uniformities and large missing parts may cause the existing orienters to generate low-quality orientations, which causes failure of the state-of-the-art reconstruction algorithms, such as the RBF-based reconstruction [19] and Poisson reconstruction [12], as illustrated in [9, 14]. The ORT orienter handles these defects more robustly. However, as a local optimization approach, small isolated inconsistent orientations may still be generated as shown in Figures 1, 7 and 8.

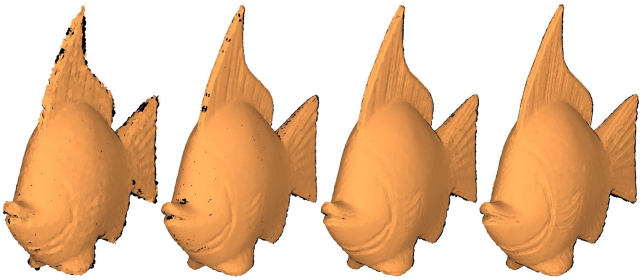

Figure 8: Orientations of the Fish model(24K points) with thin sharp features and nearby surface sheets. The columns from left to right are the results of the Cons, ORT, BOT and our approaches, respectively.

Occasionally, they may bring small artifacts to the reconstructed surface. Our global orienter leads to a more satisfactory result (see Fig. 10). Fig. 11 presents more examples to show that consistent orientation benefits surface reconstruction from defective raw point sets.

Occlusions: Thanks to the global smoothing operations, our method works well for the commonly used models (see Fig. 11 and Fig. 12), and is not sensitive to the occlusion issue due to the number and location of the viewpoints. Note that the genus-176 MegaCube model (Fig. 12) contains a large number of occlusions. By using the default 14 viewpoints, our method leads to a satisfacotry result with a very small number of inconsistent orientations. As a result, the reconstructed surface is faithful.

Implementation details: The proposed approach has been implemented in Matlab and is not optimized for efficiency. A typical example such as the Venus model (72.5K points) in Fig. 1,

Figure 11: More results. Top row: the orientation of various point clouds (Ness, Filigree, Lion, Statute and Decocube with respectively 14K, 88K, 23K, 120K and 107K points) by our approach; Bottom row: the reconstructed meshes of BRF method by using our orientation results.

Figure 9: Orientations of the Scissor model (105K points) with thin sharp features and nearby surface sheets. From the left column to right column are results of the Cons, ORT and our approaches, respectively.

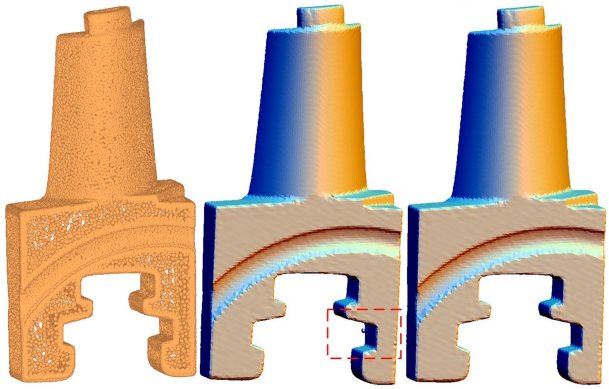

Figure 10: The locally inconsistent orientations via the ORT orienter may result in small artifacts (circled by the red dash line) when the RBF-based reconstruction [19] is employed (middle). The reconstructed surface (right) employing our global orienter does not suffer from this. The left is the point set of the Blade model (29K points).

took a total of 266 seconds. Of that time, the construction of the neighborhood took 194 seconds, the construction of the combinational Laplacian matrix took 58 seconds, solving the linear system three times took 4 seconds, the visibility checking (14 viewpoints) took 5 seconds, three times of normal flipping took 4.5 seconds, and the clustering by normals and updating of confidence took 15 seconds. The most time-consuming step is to determine neighbors, the construction of Laplacian matrix, and the point set clustering operation, which are non-vectorized Matlab code. A C++ implementation could increase the performance remarkably.

Limitations: Nearby surface sheets are challenging to all the existing point sets orienters, including ours. In our method, we try to utilize a small number of k-nearest neighbors filtered by the tangential-direction-favored distance measure to avoid neighbors from different sheets. This strategy may fail and lead to incorrect orientation if the sampling is too sparse relative to the average local distance between the two sheets, i.e., the Pelvis model (25K points) in (b) of Fig. 13. As shown in (c) of Fig. 13, this can be avoided by dense sampling. In such cases, correct up-sampling of a point set with sparsely sampled nearby surface sheets may be a potential solution.

5. Conclusion

In this paper, we presented a simple and effective algorithm to consistently orient defective raw point sets. By making use of the Laplacian contraction and visibility voting, our approach tackles the problem without resorting to any complicated data

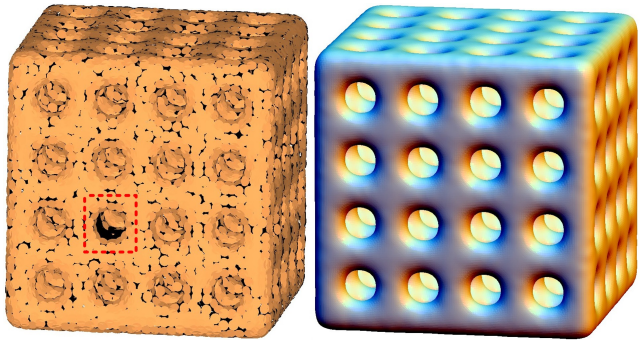

Figure 12: High genus model with large occlusions. The orientation of the genus-176 MegaCube model with 14K points and the reconstructed mesh of BRF method by using our orientation result.

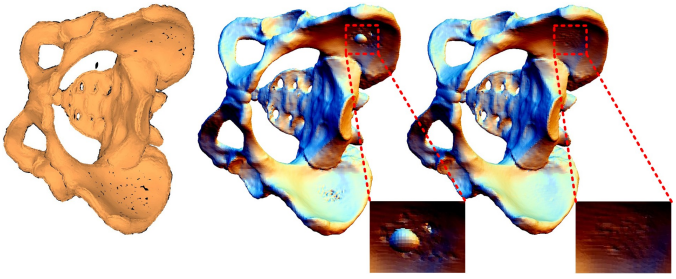

(a) (b) (c)

Figure 13: Nearby surface sheets with sparse sampling may give rise to locally inconsistent orientation. (a) The orientation of the Pelvis model (25K points) by our approach. (b) The corresponding RBF-based reconstruction has artifacts (see the close-up view). (c) With a densely sampled point clouds (51K), our method can avoid these artifacts.

structures and advanced solvers. As a global approach, our method is able to handle noise, non-uniformities, sparsity and thin sharp features. Thus it helps to improve the quality of reconstructed surfaces via conventional reconstruction schemes from raw and defect-ridden point sets. Our approach can also deal with point sets with holes and open boundaries, which pose a challenge for most orienters possessing the aid of voxelization. We demonstrated the efficacy of our method through numerous experiments.

Similar to other orientation approaches, our method cannot handle near-by surface sheets with sparse sampling. A possible future project would aim to adaptively determine the neighborhood by detecting these characteristics.

Acknowledgement

The authors would like to thank all the reviewers for their careful reviews and constructive comments, and the authors of [9, 14, 4] for sharing the executable programs and test models of their approaches. This work is supported by the Fundamental Research Funds for the Central Universities, the NSFC-Guangdong Joint Fund (No. U0935004), the NSFC (No. 60873181), the NRF2008IDM-IDM004-006 and AcRF 69/07.

References

[1] Alliez, P., Cohen-Steiner, D., Tong, Y., Desbrun, M., 2007. Voronoi-based variational reconstruction of unoriented point sets. In: SGP '07. pp. 39–48.

[2] Au, O. K.-C., Tai, C.-L., Chu, H.-K., Cohen-Or, D., Lee, T.-Y., 2008. Skeleton extraction by mesh contraction. TOG 27 (3).

[3] Cao, J., Tagliasacchi, A., Olson, M., Zhang, H., Su, Z., 2010. Point cloud skeletons via laplacian-based contraction. In: IEEE SMI '10. pp. 187–197.

[4] Chen, Y.-L., Chen, B.-Y., Lai, S.-H., Nishita, T., 2010. Binary orientation trees for volume and surface reconstruction from unoriented point clouds. CGF 29 (7).

[5] Dey, T. K., Goswami, S., 2003. Tight cocone: a water-tight surface reconstructor. In: ACM SPM '03. pp. 127–134.

[6] Guennebaud, G., Gross, M., 2007. Algebraic point set surfaces. TOG 26 (3).

[7] Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., Stuetzle, W., 1992. Surface reconstruction from unorganized points. In: SIGGRAPH '92. pp. 71–78.

[8] Hornung, A., Kobbelt, L., 2006. Robust reconstruction of watertight 3d models from non-uniformly sampled point clouds without normal information. In: SGP '06. pp. 41–50.

[9] Huang, H., Li, D., Zhang, H., Ascher, U., Cohen-Or, D., 2009. Consolidation of unorganized point clouds for surface reconstruction. TOG 28 (5).

[10] Jalba, A. C., Roerdink, J. B. T. M., 2009. Efficient surface reconstruction from noisy data using regularized membrane potentials. TIP 18 (5).

[11] Katz, S., Tal, A., Basri, R., 2007. Direct visibility of point sets. TOG 26 (3).

[12] Kazhdan, M., Bolitho, M., Hoppe, H., 2006. Poisson surface reconstruction. In: SGP '06. pp. 61–70.

[13] Li, B., Schnabel, R., Klein, R., Cheng, Z., Dang, G., Jin, S., 2010. Robust normal estimation for point clouds with sharp features. Computers & Graphics 34 (2).

[14] Liu, S., Wang, C. C. L., 2010. Orienting unorganized points for surface reconstruction. Computers & Graphics 34 (3).

[15] Ma, W.-C., Hawkins, T., Peers, P., Chabert, C.-F., Weiss, M., Debevec, P., 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In: 18th Eurographics Workshop on Rendering. pp. 183–194.

[16] Mehra, R., Tripathi, P., Sheffer, A., Mitra, N. J., 2010. Visibility of noisy point cloud data. Computers & Graphics 34 (3).

[17] Mello, V. c., Velho, L., Taubin, G., 2003. Estimating the in/out function of a surface represented by points. In: ACM SPM '03.

[18] Mullen, P., De Goes, F., Desbrun, M., Cohen-Steiner, D., Alliez, P., 2010. Signing the unsigned: Robust surface reconstruction from raw pointsets. CGF 29 (5).

[19] Ohtake, Y., Belyaev, A., Seidel, H.-P., 2003. A multi-scale approach to 3d scattered data interpolation with compactly supported basis functions. In: IEEE SMI '03. pp. 153–161.

[20] Ohtake, Y., Belyaev, A., Seidel, H.-P., 2005. An integrating approach to meshing scattered point data. In: ACM SPM '05.

[21] Sharf, A., Lewiner, T., Shklarski, G., Toledo, S., Cohen-Or, D., 2007. Interactive topology-aware surface reconstruction. TOG 26 (3).

[22] Sheung, H., Wang, C. C. L., 2009. Robust mesh reconstruction from unoriented noisy points. In: ACM SPM '09. pp. 13–24.

[23] Sorkine, O., Cohen-Or, D., 2004. Least-squares meshes. In: IEEE SMI '04. pp. 191–199.

[24] Sun, J., Smith, M., Smith, L., Farooq, A., 2007. Examining the uncertainty of the recovered surface normal in three light photometric stereo. Image and Vision Computing 25 (7), 1073–1079.

[25] Taubin, G., 1995. A signal processing approach to fair surface design. In: SIGGRAPH '95. pp. 351–358.

[26] Xie, H., McDonnell, K. T., Qin, H., 2004. Surface reconstruction of noisy and defective data sets. In: IEEE VIS '04.

[27] Xie, H., Wang, J., Hua, J., Qin, H., Kaufman, A., 2003. Piecewise c^1 continuous surface reconstruction of noisy point clouds via local implicit quadric regression. In: IEEE VIS '03.

[28] Zhao, H.-K., Osher, S., Fedkiw, R., 2001. Fast surface reconstruction using the level set method. In: IEEE VLSM '01.