

Bipartite Polar Classification for Surface Reconstruction

Yi-Ling Chen^{1,2} Tung-Ying Lee² Bing-Yu Chen³ Shang-Hong Lai²

¹Industrial Technology Research Institute, Taiwan ²National Tsing Hua University, Taiwan ³National Taiwan University, Taiwan




Figure 1: Left: Input unorganized point set \mathcal{P} . Middle: Voronoi voxelization of \mathcal{P} , where some Voronoi vertices are selected as poles (colored in navy blue). Right: Bipartite polar classification splits the poles into two disjoint sets lying in the opposite sides of \mathcal{P} (colored in purple and coral, respectively). Note that the Voronoi voxels were constructed by offsetting the underlying torus surface as the boundary constraint.

Abstract

In this paper, we propose bipartite polar classification to augment an input unorganized point set \mathcal{P} with two disjoint groups of points distributed around the ambient space of \mathcal{P} to assist the task of surface reconstruction. The goal of bipartite polar classification is to obtain a space partitioning of \mathcal{P} by assigning pairs of Voronoi poles into two mutually invisible sets lying in the opposite sides of \mathcal{P} through direct point set visibility examination. Based on the observation that a pair of Voronoi poles are mutually invisible, spatial classification is accomplished by carving away visible exterior poles with their counterparts simultaneously determined as interior ones. By examining the conflicts of mutual invisibility, holes or boundaries can also be effectively detected, resulting in a hole-aware space carving technique. With the classified poles, the task of surface reconstruction can be facilitated by more robust surface normal estimation with global consistent orientation and off-surface point specification for variational implicit surface reconstruction. We demonstrate the ability of the bipartite polar classification to achieve robust and efficient space carving on unorganized point clouds with holes and complex topology and show its application to surface reconstruction.

1 Introduction

In computer graphics, extensive efforts have been devoted to developing geometric modeling and surface reconstruction algorithms, which deal with point set data because of their wide availability. In addition to 3D positions of the surface points, some other prior knowledges, such as *oriented normal vectors*, play crucial roles in the success of many ex-

isting methods. In certain challenging scenarios, e.g. missing data or close-by surface sheets, directly interpolating the holes or smoothing the geometric details without considering the orientation information may not produce proper reconstruction in terms of topological correctness. Surface reconstruction from unorganized points [HDD*92] is thus an ill-posed and very challenging problem especially when

such side information is hard to infer due to the defects of the point clouds [HLZ*09].

Traditionally, the classical algorithm [HDD*92] and its variants start from estimating normal directions by local principal component analysis (local PCA) followed by a prioritized orientation propagation process exploiting the minimum spanning tree traversal to align the normal vectors such that they can have consistent in/out orientations. The oriented normal fields can then be utilized to reconstruct the implicit functions approximating the surface or serve as an input of other reconstruction algorithms [CBC*01, KBH06]. To handle some situations like close-by surface sheets or thin surface features, special priority measure and feature detection can be taken into account to improve the robustness of the orientation propagation [HLZ*09].

Despite the efforts taken by the traditional schemes, to correctly propagate the orientations remains a challenging problem, since normal estimation may be unreliable and only local considerations are taken to align the orientations between nearby points, which may actually drastically change. In [CCLN10], *binary orientation tree* (BOT) visually carves out the ambient space of a model from outside for surface reconstruction by exploiting the direct point set visibility [KTB07] and shows superior performance of resolving the point set orientation to traditional orientation propagation. However, it can only handle closed shapes due to the assumption that the interior region is everywhere occluded by the input point set. Recently, *cone carving* [SSZCO10] also takes advantage of global visibility property to carve out the exterior space of a shape to derive a more accurate distance-to-surface measure for surface reconstruction even with the presence of large holes or missing data. Its main drawback is the high computational complexity which degrades its feasibility in practical applications.

Given an input point set \mathcal{P} , we aim to establish a bipartite space partitioning by classifying a set of representative points distributed around \mathcal{P} into two disjoint subsets to facilitate the following surface reconstruction task. One key observation is that the points lying in the *opposite* sides of a shape (e.g. in/out of closed surfaces or front/back of open surfaces like frontal human faces) are mutually invisible from each other. Another intuition is that when observing a 3D model, once some points from its opposite sides are simultaneously visible, it implies that we are looking into a hole or across a boundary. Therefore, we take advantage of point set visibility as a means of classification and exploit the 3D Voronoi diagram constructed on \mathcal{P} to generate the representative points embedded in the ambient space of \mathcal{P} . Specifically, each input point $\mathbf{p} \in \mathcal{P}$ is associated with a pair of *poles*, \mathbf{p}^+ and \mathbf{p}^- , which are selected from the vertices of the corresponding Voronoi voxel of \mathbf{p} . Since the Voronoi poles (i.e., the set of \mathbf{p}^+ and \mathbf{p}^-) are highly likely to be opposite with respect to the shape described by \mathcal{P} if carefully selected, such mutually exclusive property of the Voronoi

pole pairs enables us to detect holes when examining their visibilities.

In summary, the proposed method, which we call *bipartite polar classification* (BPC), aims to identify the inside/outside regions of a shape. Specifically, the proposed method assigns the Voronoi poles to two disjoint sets by iteratively carving out the visible poles lying outside the surface with the opposite poles simultaneously determined. It exploits the *hidden point removal* (HPR) operator [KTB07] for visibility examination and is simple to implement and efficient to compute. In addition, the proposed method is hole-aware and capable of dealing with incomplete raw point sets without additional information. Compared with existing algorithms, the proposed method is more robust to estimate oriented surface normals with global consistency.

2 Related Work

Visibility information has been widely exploited by various space carving techniques. One of the classical approaches is the volume carving method [CL96], which eliminates the empty voxels along the line-of-sight from the sampled points to scanner. Another class of approaches is the image-based *visual hull*, which utilizes silhouette information from a number of reference images to carve out the outside regions not belonging to an object [Lau94]. Instead of volumetric representations, BOT [CCLN10] accomplishes space carving on raw input point clouds by taking “*snapshots*” from various viewpoints to tag and remove certain exterior auxiliary points depending on their visibilities. However, such a simple carving process cannot prevent from penetrating into holes which reveal the inside space of the point set. In [SSZCO10], cone carving creates *visibility cones* apexed at each point and extended outward to carve away the outside space, whose silhouette is traced by a view-dependent point rendering process. Since splatting of all other points is required, it is not an efficient algorithm even with the acceleration of GPUs. The HPR operator [KTB07] is directly related to our method. It determines the visible points among a point set from a given viewpoint through a *spherical flipping* transformation and convex hull computation. Recently, Mehra et al. [MTSM10] extended the HPR operator to deal with point sets containing concavities, non-uniformly spaced samples, and possibly corrupted with noise. Some other related works [NT00, ZZBW08] include the visibility-driven algorithms aiming to identify the interior and exterior parts of polygonal models.

The pioneering work of Hoppe et al. [HDD*92] computes a signed distance function approximating an unorganized point cloud by using estimated tangent planes, which is followed by iso-surface extraction. Many implicit surface modeling methods have been developed, such as variational implicit surfaces of radial basis functions (RBFs) [CBC*01, TO02], multi-level partition-of-unity [OBA*03] and Poisson reconstruction [KBH06]. All of these methods require prescribed surface normals to facilitate the re-




Figure 2: (a) Hole detection with the poles and main axes of the conflict Voronoi voxels shown in blue spheres and lines, respectively. (b) Spatial classification without γ -sphere carving. The green spheres are the unclassified poles, which cannot be carved away by visual carving.

construction task. Some other approaches achieve 3D reconstruction by first computing unsigned distance approximation to the input data followed by stochastic signing of the function [MdGD*10]. To obtain precise geometric and orientation information of a 3D model [NRDR05] is a process subject to the data acquisition conditions, which may be affected by the presence of material artifacts, shadowing or inaccurate registration. To deal with incomplete scan data caused by physical inaccessibility or poor visibility, an interactive technique is introduced to recover the topology of complex undersampled regions with user input [SLS*07]. To reliably estimate the surface normals, particularly with globally consistent orientation, is considerably difficult and has been intensively investigated in previous researches [HDD*92, ACSTD07, GG07, HLZ*09].

Voronoi diagrams and Delaunay triangulations have been widely applied in surface reconstruction [EM94, ABK98, DG03, CG06], which provide provable theoretical guarantees under proper sampling conditions. However, the quality of Voronoi-based approaches typically degrade due to under- or non-uniform sampling. Although, the idea of using poles for surface reconstruction is not new [ABK98, ACK01], the proposed method differs from the previous ones in identifying sufficient poles of high confidence by checking their visibilities, which is a weak assumption that is true under most situations. The contribution of this work is thus a novel algorithm to classify the initially unorganized poles into two sets with explicit spatial identification to facilitate surface reconstruction.

3 Bipartite Polar Classification

Given a point set \mathcal{P} , which possibly contains noise, holes and non-uniformities, our goal is to partition the space into the two opposite sides of the shape described by \mathcal{P} . To achieve this, we formulate it as a space carving problem by grouping a set of representative points \mathcal{P}' into two disjoint sets \mathcal{P}^+ and \mathcal{P}^- (i.e., $\mathcal{P}^+ \cup \mathcal{P}^- = \mathcal{P}'$ and $\mathcal{P}^+ \cap \mathcal{P}^- = \emptyset$),

which are mutually invisible to obtain a bipartite classification. Such a classification problem is difficult in selecting sufficient and discriminative points to identify the spatial relationship and also to fit to the application in question. For surface reconstruction, we associate each point $\mathbf{p}_i \in \mathcal{P}$ with a pair of poles \mathbf{p}_i^+ and \mathbf{p}_i^- , which are supposed to be opposite and thus invisible to each other. Specifically, the HPR operator [KTB07] is utilized to carve out the visible exterior poles[†] from the space and assign them to \mathcal{P}^+ with their opposite poles assigned to \mathcal{P}^- simultaneously. Note that initially \mathbf{p}_i^+ and \mathbf{p}_i^- do not necessarily correspond to either \mathcal{P}^+ or \mathcal{P}^- , and are then classified by visual carving. Hole detection is first performed by checking the conflicts of mutual invisibility between the pole pairs from some preselected viewpoints before the iterative carving process. The awareness of holes prevents us from carving the poles away from both sides of \mathcal{P} to obtain incorrect spatial classification. To summarize, the proposed method, *bipartite polar classification* (BPC), consists of the following main components: *pole selection*, *hole detection* and *iterative polar classification*.

3.1 Pole Selection

For spatial classification, it would be most desirable to create a number of representative points which can be potentially divided into two mutually invisible sets. To achieve this, we choose to construct a 3D Voronoi diagram \mathcal{V} of \mathcal{P} with each voxel $\mathcal{V}_i \in \mathcal{V}$ occupied by only a single point $\mathbf{p}_i \in \mathcal{P}$ based on the observation exploited by some previous methods [ABK98, ACK01]. Since the Voronoi voxels tend to be long and skinny under a good sampling quality, the poles can thus be simply selected from the two vertices of \mathcal{V}_i farthest from \mathbf{p}_i . However, the Voronoi voxels would become irregularly shaped with the influence of noise. Similar to [ACSTD07], we perform a covariance analysis on \mathcal{V}_i to obtain a principal axis \mathbf{n}_i and its anisotropy $\sigma_i \in [0, 1]$, which is defined as $\sigma_i = 1 - \frac{\lambda_{min}}{\lambda_{max}}$, where λ_{min} and λ_{max} correspond to the smallest and largest eigenvalues, respectively. A high σ_i implies an elongated \mathcal{V}_i extending perpendicularly to the shape, and \mathbf{p}_i^+ and \mathbf{p}_i^- can thus be selected to be the two farthest vertices of \mathcal{V}_i along \mathbf{n}_i . In contrast, a low σ_i indicates that the vertices of \mathcal{V}_i are isotropically distributed. Therefore, no poles are included for spatial classification. We empirically adopted a threshold of 0.9 for pole selection. A voxel \mathcal{V}_i and its pole pair \mathbf{p}_i^+ and \mathbf{p}_i^- are referred to as *reliable* if $\sigma_i > 0.9$, otherwise they are *vulnerable*. For noisy data sets, we optionally apply the *weighted locally optimal projection* (WLOP) operator [HLZ*09] to preprocess \mathcal{P} before the Voronoi voxelization.

[†] Without loss of generality, we refer to the poles observed by the HPR operator as the exterior ones, since they are usually visible when viewing from outside.

3.2 Hole Detection

The aforementioned polarized Voronoi voxelization provides us a number of reliable poles, which are eligible for the following spatial classification. It is important to be conscious of the presence of holes in \mathcal{P} for a space carving method like BPC to succeed. In geometric models, holes are the locus where the spatial identification conflicts occur. With the mutual opposite property between the Voronoi pole pairs \mathbf{p}_i^+ and \mathbf{p}_i^- , such conflicts can be easily detected by examining their visibilities from a certain viewpoint. If both a pair of Voronoi poles are observed at the same time, we are aware that a hole is visible under current viewing direction, which should be avoided. Hence, we perform a *prescanning* process for hole detection before space carving. From a set of preselected viewpoints, the HPR operator [KTB07] is applied to all reliable pole pairs \mathcal{P}' as well as \mathcal{P} to detect the conflicts. Note that holes or boundaries are typically not visible under all viewing directions. Hence, it is important to record the detected Voronoi pairs of conflict and *freeze* them during the following space carving. The frozen poles are never removed and function as a *barrier* across which line-of-sight cannot pass. Undoubtedly, increasing the viewpoints in prescanning will certainly find more such protective Voronoi pole pairs. Empirically, we found that it is sufficient to prescan a model along the opposite directions of x -, y - and z -axes if it is located at the origin. Figure 2 (a) shows the Voronoi poles of conflicts and the detected hole of the DISTCAP data set.

3.3 Iterative Polar Classification

After hole detection, we are ready to classify $\mathcal{P}' = \{\mathbf{p}_i^+, \mathbf{p}_i^- \mid \forall \mathbf{p}_i \in \mathcal{P}, \sigma_i > 0.9\}$ into \mathcal{P}^+ and \mathcal{P}^- . The following iterative procedure composes the proposed polar classification method:

1. Initially, a number of viewing directions on a unit sphere are uniformly sampled.
2. Randomly choose an unused viewing direction and place a virtual camera at c sufficiently far away along this direction. Then, apply the HPR operator to observe $\mathcal{P} \cup \mathcal{P}'$ from c .
3. Remove the observed \mathbf{p}_i^+ or \mathbf{p}_i^- from \mathcal{P}' and assign every visible pole to \mathcal{P}^+ and its counterpart to \mathcal{P}^- , but skip the poles if conflicts occur (i.e., both \mathbf{p}_i^+ and \mathbf{p}_i^- are observed).
4. For every carved pole pair, recursively perform γ -sphere carving.
5. Repeat from Step 2 if \mathcal{P}' is not empty (excluding the frozen poles) and there are still unused viewing directions.

In Step 1, we uniformly sample 300 viewing directions on a unit sphere which sufficiently exposes \mathcal{P} and \mathcal{P}' when performing the HPR operation. The actual viewpoints c are positioned by extending the viewing directions from the origin with a distance of $20 \sim 30$ (relative to a unit bounding




Figure 3: Noise effect on Voronoi voxelization and the results of BPC. (a) A clean KNOT data set consisting of 10,000 points, and there are 9,529 reliable pole pairs. (c) The KNOT data set is perturbed by 1% random noise, and the number of reliable pole pairs hence decreases to 4,059. (b) and (d) are the corresponding reconstructed RBF surfaces of (a) and (c), respectively.

sphere of \mathcal{P}). For efficiency consideration, we do not perform visibility examination from all candidate viewpoints and empirically eliminate the nearest $10 \sim 15$ viewpoints to the current c from the candidate set since the fields of vision of such viewpoints do not vary much from c . Typically, the first few iterations identify most exterior poles. Some exterior poles occluded by \mathcal{P} or other exterior poles in \mathcal{P}' depending on the viewpoints are gradually observed as the iterative carving algorithm proceeds. Recall that the protective poles detected in the prescanning stage indicate the existence of holes or boundaries. Therefore, they are never removed in Step 3 even though conflict does not occur under the current viewing direction. The co-existence of pairs of protective \mathbf{p}_i^+ and \mathbf{p}_i^- is important for hole detection as soon as a hole becomes visible from a new viewpoint.

In Step 4, γ -sphere carving stands for the process of placing a ball of radius γ centered at a classified pole and carving away any other unclassified poles lying within the γ -sphere with the same polar classification, which is conceptually similar to the α -balls used to construct the α -shapes [EM94]. Such a supplementary carving operation possesses significant impact on BPC. We conservatively set γ to a small value due to two reasons. First, the Voronoi poles tend to cluster, which implies distant points in \mathcal{P} may have neighboring poles. This phenomenon is best understood by con-




Figure 4: Surface reconstruction of a sparse and non-uniform point set HAND. (a) The BPC result. (b) The normal estimation and orientation propagation by [HLZ*09]. (c) and (d) The reconstructed RBF surface by the normal fields obtained from (a) and (b), respectively.

sidering the TORUS (Figure 1) or KNOT (Figure 3) model, which is symmetrically shaped. The interior poles form a curve with each pole extremely close to its neighbors. By carving away the nearby interior poles, the corresponding exterior poles can also be explicitly classified, which may be hard to be reached by visual carving. In practice, a small γ suffices to globally deliver local classification results over the entire model. Second, a small γ is more appropriate since \mathcal{P} possibly contains holes. Otherwise, a γ -sphere of adaptive radius which contains no $\mathbf{p} \in \mathcal{P}$ may be used. Although the protective poles indicate the existence of holes, they are usually not dense enough to block the γ -spheres from growing across the shape boundaries.

4 Applications

4.1 Surface Normal Estimation

For a reliable voxel \mathcal{V}_i , its principal axis \mathbf{n}_i and exterior pole assigned to \mathcal{P}^+ explicitly specify a normal vector of its corresponding point \mathbf{p}_i directed outward. As for the vulnerable voxels, an unsigned normal direction \mathbf{n}'_i may be estimated by local PCA and then orientated with the k -nearest classified poles of \mathbf{p}_i by maximizing

$$\sum_{j=1}^k \text{sign}(\mathbf{p}'_j) \cdot (\mathbf{p}'_j - \mathbf{p}_i) \cdot \mathbf{n}'_i, \quad (1)$$

where $\text{sign}(\mathbf{p}'_j)$ is a binary function that returns 1 or -1 if a pole \mathbf{p}'_j is assigned to \mathcal{P}^+ or \mathcal{P}^- , respectively. Similar to cone carving [SSZCO10] and BOT [CCLN10], BPC takes the global visibility into account to obtain a spatial partition compliant with the model structure and is thus more robust than traditional propagation-based approaches [HDD*92, HLZ*09], which take only local conditions into consideration for orientation determination.

4.2 Variational Implicit Surface Reconstruction

The classified poles \mathcal{P}^+ and \mathcal{P}^- are semantically similar to the *off-surface constraints* typically required in solving the

variational implicit surfaces [TO02,CBC*01], which specify a number of locations to assist the implicit surface to separate the inside/outside space. Voronoi voxelization and BPC can thus simplify the task of reconstructing the variational implicit surfaces from unorganized points. Specifically, a subset of \mathcal{P} may be augmented by off-surface points $\hat{\mathbf{p}}_i$ explicitly specified by their corresponding classified poles. Figure 3 illustrates two examples of reconstructing the variational implicit surfaces from clean and noisy data sets, respectively. Note that despite the influence of noise, which affects the Voronoi voxelization and decreases the size of \mathcal{P}' , \mathcal{P}' still suffice to deliver the topology and orientation information of \mathcal{P} , resulting in topologically correct surface reconstruction. Nevertheless, it will still be desirable to *consolidate* \mathcal{P} [HLZ*09] to enrich the pole set especially when the noise level is high.

5 Results and Discussions

With many well-established libraries of computational geometry algorithms, the implementation of BPC is straightforward. We adopted Qhull [BDH96] for the HPR operation [KTB07] and voro++ [Ryc07] for the Voronoi voxelization, respectively. Initially, the input point set \mathcal{P} is first normalized into the interval of $[-1, 1]$ and translated to be centered at the origin. voro++ enables us to impose a bounding sphere of radius 5 on \mathcal{P} as boundary constraint and the Voronoi voxelization of \mathcal{P} is performed to extract \mathcal{P}' . $\mathcal{P} \cup \mathcal{P}'$ are then collected together for hole detection (Section 3.2) and iterative polar classification (Section 3.3) with the HPR operator. Note that the presence of \mathcal{P} is essential for separating \mathcal{P}^+ and \mathcal{P}^- during the iterative carving process.

Table 1 summarizes the CPU runtimes of processing the models shown in this paper. Generally, BPC takes only a few seconds to process a point set of moderate size. Despite the overhead of the Voronoi voxelization and WLOP, BPC is still quite efficient to process the point clouds with similar data size when comparing to cone carving [SSZCO10]. When dealing with a large \mathcal{P} , a subset of \mathcal{P} will be sufficient to obtain an appropriate spatial classification, which is demon-

Table 1: Computation times for point consolidation (WLOP), Voronoi voxelization (VD) and bipartite polar classification (BPC) on several data sets. Point# and the numbers in parentheses under BPC stand for the number of points and the iteration times of the HPR operation. All results were obtained on a desktop PC with an Intel i7 2.8GHz CPU with 2GB RAM, and the computation times are represented in seconds.

Model	Point#	WLOP	VD	BPC
TORUS	4,800	-	8.36	0.73 (7)
HAND	7,609	-	12.27	2.86 (38)
KNOT	10,000	-	14.59	1.28 (7)
INUKSHUK	10,293	35.89	15.47	4.25 (36)
MANNEQUIN	10,722	11.95	18.86	6.97 (40)
FERTILITY	12,081	17.7	33.69	4.86 (42)
DANCER	12,428	26.61	27.52	9.19 (43)
CHILDREN	12,500	10.11	15.63	5.36 (39)
DANCER2	12,535	27.75	27.14	10.12 (41)
DISTCAP	12,745	-	51.97	5.37 (36)
HEPTOROID	14,334	23.31	15.25	4.63 (36)
BIMBA	15,002	-	43.13	11.63 (40)
DINOSAUR	18,494	31.28	99.67	19.33 (40)
HORSE	18,532	39.06	266.1	21.95 (43)

strated in Figure 5. In this example, BPC is performed on a uniformly sampled subset (15,002 points) of a dense point set \mathcal{P} (74,764 points). The orientated normal field of \mathcal{P} can thus be estimated by local PCA and the classified poles of the sampled subset, as described in Section 4.1, and the Poisson implicit surface [KBH06] can then be reconstructed. In this example, one can hardly see any visual differences from Figure 5 (b) when compared to Figure 5 (c) and the result of the same model by [ACSTD07]. It is considered to be necessary to extract a reasonably good sampling adhered to the underlying shape of \mathcal{P} for such a scenario to succeed, because a proper Voronoi voxelization is required to obtain sufficient reliable poles. The WLOP operator [HLZ*09], which pre-processed several originally denser point sets as indicated in Table 1, is potentially suitable for this purpose.

Being a similar space carving method exploiting direct point set visibility, BPC brings significant performance enhancement in dealing with complex models when comparing with BOT [CCLN10]. Figure 2 (b) demonstrates an example of space carving on a point cloud by purely using the HPR operation, which is taken by BOT. Since the ambient space of a topologically complex model, such as KNOT (Figure 3) or HEPTOROID (Figure 9 (d)), is self-occluded from various viewing directions, it is thus difficult to be handled by simply visual carving. In addition, the poles lying within the occluded regions actually form concave shapes of high curvature, which are hard to be observed by the HPR operation.

The property of being opposite with respect to a shape of pole pairs simplifies the task of space carving since their spatial classification can be simultaneously determined. In ad-




Figure 5: (a) Uniform sampling of the BIMBA point set (15K points) and the corresponding BPC result; (b) and (c) Reconstructed Poisson implicit surfaces by (a) and the original point set (75K points), respectively.




Figure 6: RBF implicit surfaces reconstructed from (a) the DINOSAUR and (b) HORSE data sets.

dition, since Voronoi poles are close to the medial axis of a shape, they capture the reflectional symmetries of a geometric model and approximate to a 2D manifold. The γ -sphere carving explained in Section 3.3 can carry local classification results around to carve out the regions which are actually not observed by the HPR operation. Concave regions hidden by topological details which cannot be observed under most viewing directions can thus be gracefully handled. As shown in Figure 9, BPC successfully deals with a variety of shapes of high topological complexity. It is worth noting that for every point set, the first six iterations of the HPR operators account for the prescanning process of the hole detection. As indicated in Table 1, one iteration of the HPR operation was sufficient for TORUS and KNOT to classify all poles since their interior poles form continuous curved medial axes and local classification results are propagated over the model to obtain the global solution by γ -sphere carving. Unlike traditional volume carving techniques, the requirement of sufficient “scans” to carve out all exterior regions is greatly alleviated by BPC.

In Figure 4, the performance of BPC and the traditional normal propagation scheme [HLZ*09] on a sparse and non-uniform point set with close-by surface sheets is compared. In Figure 4 (a), BPC successfully classified the Voronoi




Figure 7: Surface normal estimation with globally consistent orientation: (a) DINOSAUR by [HLZ*09], (b) DINOSAUR by BPC, (c) HORSE by [HLZ*09], (d) HORSE by BPC.

poles (with the interior ones shown in coral spheres), which are then utilized to estimate and orientate the normal vectors of every data point to assist RBF implicit surface reconstruction (Figure 4 (c)). In Figure 4 (b), the previous method [HLZ*09] failed to orientate the normal vectors of some data points (invisible due to back culling of point splatting), resulting in incorrectly reconstructed surface (Figure 4 (d)). Figure 7 shows two models whose 3D structures were recovered by multiview epipolar geometry [HL06] and contain holes, non-uniformities and noise. The previous method obviously mis-aligned some normal vectors, while the results by BPC do not have such problems. Figure 6 shows the corresponding RBF implicit surfaces of the DINOSAUR and HORSE models by BPC. Some more examples showing the ability of BPC to facilitate surface reconstruction with thin structures are demonstrated in Figure 8 (a) and (b).

Limitations. Due to the goal of bipartite space partitioning, BPC is not able to deal with the shapes which cannot explicitly separate the ambient space into two opposite sides. As a Voronoi-based method, BPC is affected by the quality of Voronoi voxelization. In Figure 8 (c), a proper Voronoi voxelization was not available in-between the extremely close-by surface sheets, so BPC failed to separate the inside/outside space due to the lack of reliable Voronoi poles, and hence resulting in geometrical errors in surface reconstruction. For similar reasons, it will also be difficult for BPC to carve out the exterior regions between opposite holes.

6 Conclusion and Future Work

To summarize, BPC utilizes Voronoi voxelization to generate a set of representative points, i.e., poles, suitable to discriminate the opposite sides of a shape. It is based on the assumption that all exterior poles are visible from outside and accomplishes spatial classification in a fashion of space carving on point clouds. As a visibility-driven space carving technique, some advantages of BPC include its hole-awareness and robustness to derive orientation information from incomplete raw point sets. The classified poles individually provide reliable surface normal estimation and

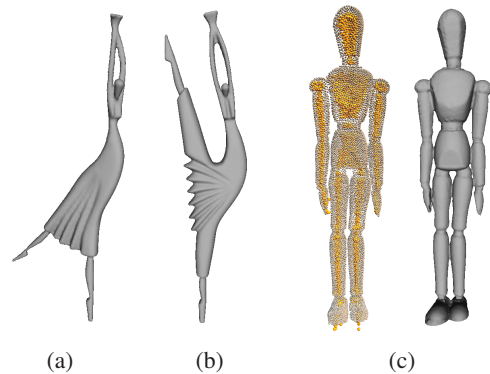


Figure 8: (a) and (b) The examples of BPC assisted surface reconstruction from the point sets with thin structures. (c) BPC failed to separate the close-by surface sheets due to the lack of reliable poles.

collectively convey topology information if properly post-processed. Despite the overhead of Voronoi voxelization, BPC is efficient to compute and simple to implement when compared to other visibility-driven method [SSZCO10]. For future work, we plan to investigate the feasibility to derive higher level topological representations of point clouds by BPC, e.g. skeletons, to assist surface reconstruction under difficult situations.

References

- [ABK98] AMENTA N., BERN M., KAMVYSSELIS M.: A new voronoi-based surface reconstruction algorithm. In *Proc. ACM SIGGRAPH* (1998), pp. 415–421. [3](#)
- [ACK01] AMENTA N., CHOI S., KOLLURI R. K.: The power crust. In *Proc. ACM SMA* (2001), pp. 249–266. [3](#)
- [ACSTD07] ALLIEZ P., COHEN-STEINER D., TONG Y., DESBRUN M.: Voronoi-based variational reconstruction for unoriented point sets. In *Proc. EG SGP* (2007), pp. 39–48. [3, 6](#)
- [BDH96] BARBER C. B., DOBKIN D. P., HUHDANPAA H.: The quickhull algorithm for convex hulls. *ACM TOMS* 22, 4 (1996), 469–483. [5](#)
- [CBC*01] CARR J. C., BEATSON R. K., CHERRIE J. B., MITCHELL T. J., FRIGHT W. R., MCCALLUM B. C., EVANS



Figure 9: BPC assisted surface reconstruction. Top: The input point sets and their interior poles \mathcal{P}^- colored in coral. Bottom: The corresponding reconstructed RBF implicit surfaces.

- T. R.: Reconstruction and representation of 3d objects with radial basis functions. In *Proc. ACM SIGGRAPH* (2001), pp. 67–76. [2](#), [5](#)
- [CCLN10] CHEN Y.-L., CHEN B.-Y., LAI S.-H., NISHITA T.: Binary orientation trees for volume and surface reconstruction from unoriented point clouds. *Computer Graphics Forum* 29, 7 (2010), 2011–2019. [2](#), [5](#), [6](#)
- [CG06] CAZALS F., GIESEN J.: Delaunay triangulation based surface reconstruction: Ideas and algorithms. In *Effective Computational Geometry for Curves and Surfaces*. Springer, 2006, pp. 231–273. [3](#)
- [CL96] CURLESS B., LEVOY M.: A volumetric method for building complex models from range images. In *Proc. ACM SIGGRAPH* (1996), pp. 303–312. [2](#)
- [DG03] DEY T. K., GOSWAMI S.: Tight cocone: a water-tight surface reconstructor. In *Proc. ACM SMA* (2003), pp. 127–134. [3](#)
- [EM94] EDELSBRUNNER H., MÜCKE E. P.: Three-dimensional alpha shapes. *ACM TOG* 13, 1 (1994), 43–72. [3](#), [4](#)
- [GG07] GUENNEBAUD G., GROSS M.: Algebraic point set surfaces. *ACM TOG* 26, 3 (2007), 23:1–23:10. [3](#)
- [HDD*92] HOPPE H., DE ROSE T., DUCHAMP T., McDONALD J., STUETZLE W.: Surface reconstruction from unorganized points. In *Proc. ACM SIGGRAPH* (1992), pp. 71–78. [1](#), [2](#), [3](#), [5](#)
- [HL06] HUANG P.-H., LAI S.-H.: Contour-based structure from reflection. In *Proc. IEEE CVPR* (2006), pp. 379–386. [7](#)
- [HLZ*09] HUANG H., LI D., ZHANG H., ASCHER U., COHEN-OR D.: Consolidation of unorganized point clouds for surface reconstruction. *ACM TOG* 28, 5 (2009), 176:1–176:7. [2](#), [3](#), [5](#), [6](#), [7](#)
- [KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson surface reconstruction. In *Proc. EG SGP* (2006), pp. 61–70. [2](#), [6](#)
- [KTB07] KATZ S., TAL A., BASRI R.: Direct visibility of point sets. *ACM TOG* 26, 3 (2007), 24:1–24:12. [2](#), [3](#), [4](#), [5](#)
- [Lau94] LAURENTINI A.: The visual hull concept for silhouette-based image understanding. *IEEE TPAMI* 16, 2 (1994), 150–162. [2](#)
- [MdGD*10] MULLEN P., DE GOES F., DESBRUN M., COHEN-STEINER D., ALLIEZ P.: Signing the unsigned: Robust surface reconstruction from raw pointsets. *Computer Graphics Forum* 29, 5 (2010), 1733–1741. [3](#)
- [MTSM10] MEHRA R., TRIPATHI P., SHEFFER A., MITRA N. J.: Visibility of noisy point cloud data. *Computers and Graphics* 34, 3 (2010), 219–230. [2](#)
- [NRDR05] NEHAB D., RUSINKIEWICZ S., DAVIS J., RAMAMOORTHY R.: Efficiently combining positions and normals for precise 3D geometry. *ACM TOG* 24, 3 (2005), 536–543. [3](#)
- [NT00] NOORUDDIN F. S., TURK G.: Interior/exterior classification of polygonal models. In *Proc. IEEE VIS* (2000), pp. 415–422. [2](#)
- [OBA*03] OHTAKE Y., BELYAEV A., ALEXA M., TURK G., SEIDEL H.-P.: Multi-level partition of unity implicits. *ACM TOG* 22, 3 (2003), 463–470. [2](#)
- [Ryc07] RYCROFT C. H.: *Multiscale Modeling in Granular Flow*. PhD thesis, Massachusetts Institute of Technology, 2007. [5](#)
- [SLS*07] SHARF A., LEWINER T., SHKLARSKI G., TOLEDO S., COHEN-OR D.: Interactive topology-aware surface reconstruction. *ACM TOG* 26, 3 (2007), 43:1–43:9. [3](#)
- [SSZCO10] SHALOM S., SHAMIR A., ZHANG H., COHEN-OR D.: Cone carving for surface reconstruction. *ACM TOG* 29, 6 (2010), 150:1–150:10. [2](#), [5](#), [7](#)
- [TO02] TURK G., O'BRIEN J. F.: Modelling with implicit surfaces that interpolate. *ACM TOG* 21, 4 (2002), 855–873. [2](#), [5](#)
- [ZZBW08] ZHOU K., ZHANG E., BITTNER J., WONKA P.: Visibility-driven mesh analysis and visualization through graph cuts. *IEEE TVCG* 14, 6 (2008), 1667–1674. [2](#)