

Denoising a Point Cloud for Surface Reconstruction

Siu-Wing Cheng Man-Kit Lau
Department of Computer Science and Engineering
HKUST, Hong Kong
{scheng,lmkaa}@ust.hk

Abstract

Surface reconstruction from an unorganized point cloud is an important problem due to its widespread applications. White noise, possibly clustered outliers, and noisy perturbation may be generated when a point cloud is sampled from a surface. Most existing methods handle limited amount of noise. We develop a method to denoise a point cloud so that the users can run their surface reconstruction codes or perform other analyses afterwards. Our experiments demonstrate that our method is computationally efficient and it has significantly better noise handling ability than several existing surface reconstruction codes.

1 Introduction

Surface reconstruction from a point cloud is an important problem due to its widespread applications. Algorithms for reconstructing from noise-free data have been developed. Some of them [2, 3, 4, 5] produce provably good approximations when the data satisfies certain sampling conditions. However, it is a challenge to handle white noise, outliers, and noisy perturbation that may be generated when sampling a point cloud from an unknown surface. Outliers may even be structured and clustered (e.g. sample points in the planar background when an object is scanned).

A popular strategy for denoising a point cloud is to define a smooth surface function using the point cloud (e.g. [12, 15, 17, 19, 21, 23, 24, 25]). If a denoised point set is desired, the input data points can be projected iteratively to the zero set of the surface function, or a meshing algorithm can be applied. The efficiency depends heavily on whether the surface function can be evaluated very quickly because many evaluations are needed. Most of these surface functions require accurate oriented surface normals at the data points. It has been reported that reconstruction algorithms may be brittle if they depend on accurate surface normal estimation [22].

In this paper, we take a different approach to denoise the point cloud, assuming that the data points are sampled densely from a smooth surface with or without boundaries. Our goal is to produce a denoised point set on which the users can run their surface reconstruction codes or perform other analyses. In our experiments, we run Robust Cocone [11] on the point sets denoised by our method. Our denoising code can be downloaded from the webpage for this project [1].

In our experiments, the number of white noise points and possibly clustered outliers can be more than 100% of the number of data points, and the noisy perturbation can be as large as 2% of the bounding box diameter of the point cloud. The amount of outliers and noise pose a serious challenge to several existing surface reconstruction codes as indicated in our experiments.

Our denoising method consists of three stages. First, we use a new octree decomposition to cluster the data points. The largest cluster contains the data points around the unknown surface; therefore, by extracting the largest cluster, we can remove white noise and outliers. If there are $k > 1$ surfaces, we work with the k largest clusters instead. Second, we perform a simple statistical analysis on the octree boxes to remove data points with relatively large noisy perturbation. Third, we use the remaining points to guide the construction of a sparser point set, and we develop a meshless Laplace smoothing procedure to denoise this new point set. Unlike previous approaches, we do not estimate surface normals, or perform surface fitting, or determine the inside/outside of the unknown surface, which helps to make our method simple and fast.

Our octree decomposition is based on a previous work with a collaborator [7] for surface reconstruction from clean data. We extend it so that the octree-induced clustering facilitates the removal of outliers and white noise.

We find it desirable to divide the removal of noisy perturbation of the data points into two steps, the first step for removing relatively large perturbation and the second step for removing any smaller perturbation left. Otherwise, if we apply the meshless Laplacian smoothing right after removing the outliers and white noise, the “weighted averages” will drift away from the unknown surface.

There is a novel feature in our meshless Laplacian smoothing step. Instead of smoothing the input data points, we construct a sparser point set and smooth these points. The reason is that it is computationally inefficient to shift an input data point based on all other input data points in its neighborhood. We use the input data points to guide the construction of a sparser point set, but it is undesirable to apply Laplacian smoothing yet because if a point is surrounded by other points, the unevenness in the local sampling will cause the point to drift in the direction of higher local density during smoothing. Poor output may be produced. Instead, we divide the neighborhood of a point in this sparser point set into groups, and we only pick one representative point from each group. This gives an even sparser point set to which smoothing is applied.

1.1 Previous Work

Previous work perform the noise and outlier filtering together with the surface reconstruction. Therefore, we briefly survey the related surface reconstruction work.

1.1.1 Point Cloud Denoising

There are several variants of *moving least squares* (MLS) for approximating the unknown surface under noise and outliers. Shen et al. [25] defined the *implicit moving least-squares* (IMLS) surface for a soup of polygons with their oriented normals as constraints. Kolluri [17] adapted IMLS to a point cloud in which every data point is tagged with an oriented surface normal. He proved that the approximation is good if a uniform sampling condition is satisfied. Dey and Sun [12] proposed the *adaptive moving least-squares* (AMLS) surface definition which allows the sampling to be non-uniform and sensitive to the local feature size. AMLS can handle a small amount of noise in our experiments.

Schall, Belyaev and Seidel [24] defined a likelihood function for data points associated with oriented surface normals. Outliers are removed by thresholding on the function values. The remaining data points are projected to local functional maxima to produce a denoised point set.

Xie, Wang, Hua, Qin and Kaufman [27] defined a surface function by blending quadric surfaces defined over different regions. It seems hard to determine the extent of these regions for a very noisy point set. Xie, McDonnell, and Qin [26] used an octree to classify the space into regions that are inside/outside the unknown surface, fit a surface to each octree cell, and blend these surfaces. Non-uniform sampling and sharp features are allowed. Their experiments show that a small amount of noise and outliers can be handled.

Nagai, Ohtake and Suzuki [23] presented the PoissonPU algorithm for constructing an implicit surface. Surface normals are required at the data points. They apply anisotropic Laplacian smoothing to the surface normals to deal with normal estimation error, while preserving features. A small amount of noise and outliers can be handled as shown in their experiments.

Mullen, De Goes, Desbrun, Cohen-Steiner and Alliez [22] proposed to impose signs on an unsigned distance function defined for a point cloud. The inside/outside test is performed by tracing rays and checking the parity of intersections with an ϵ band around the zero set of the distance function. Structured and clustered outliers may be an issue (e.g. sample points in the planar background in the Dragon data set) as they may fool the inside/outside decision. The large amount of outlier clusters in our experiments pose some difficulties.

Giraudot, Cohen-Steiner, and Alliez [13] developed a noise-adaptive distance function that can handle any smooth submanifold of known dimension, variable noise, and outliers. They report that some data sets are challenging for reconstruction algorithms that rely on accurate surface normals. The clustered outliers in our experiments seem to pose a challenge to their code.

Guggeri, Scateni and Pajarola [14] proposed a depth carving algorithm to reconstruct the unknown surface. Their experiments show that some white noise can be handled. However, boundaries pose a problem and no result is given on handling noisy perturbation.

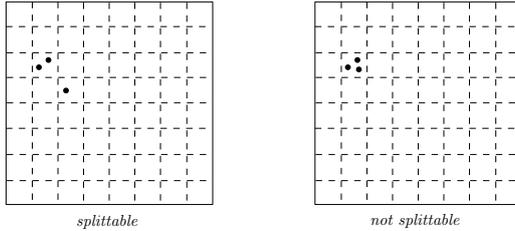


Figure 1: 2D case example of splittability.

1.1.2 Triangulation Algorithms

There are several triangulation algorithms [5, 6, 8, 9, 10, 11, 16, 20] for surface reconstruction and some are able to handle a small amount of noise [6, 11, 16].

Lorensen and Cline [20] proposed the seminal marching cube algorithm which is a fast way to obtain a triangular mesh of the zero set of an implicit function. However, it is sensitive to noise and outliers.

Bernardini, Mittleman, Rushmeier, Silva and Taubin [6] proposed a ball-pivoting algorithm for surface reconstruction of a point cloud by a region growing method. Oriented surface normals at the data points are required.

Kazhdan, Bolitho and Hoppe [16] expressed the surface reconstruction as a Poisson problem. A 3D indicator function that best-fits the point cloud is computed. Then, an appropriate isosurface is extracted to approximate the unknown surface. Oriented surface normals at the data points are required.

Dey and Goswami [11] developed the Robust Cocone algorithm to reconstruct a closed triangular mesh from data points with a small amount of noise. The mesh consists of an appropriate set of Delaunay triangles induced by the data points. It is a provably good approximation if the point cloud satisfies certain sampling conditions.

Kolluri, Shewchuk and O'Brien [18] proposed the Eigencrust algorithm. After computing the Delaunay tetrahedralization, spectral graph partitioning is used to determine whether a Delaunay tetrahedron lies inside the unknown surface. The reconstruction consists of the triangles shared by inside and outside tetrahedra. Their experimental results show that the reconstruction can be corrupted when there are many outliers or the noise level is high.

2 Octree Decomposition

Consider an axes-aligned minimum bounding cube B of the input point cloud. We present a new octree decomposition to subdivide B into smaller cubes.

Every cube that we refer to is aligned with the coordinate axes. The *size* of a cube x is its side length, and we denote it by ℓ_x . Every octree node corresponds to a cube, and therefore, we refer to an octree node as a cube. Two octree nodes are *neighbors* if their interiors are disjoint and their boundaries are in contact. Two nodes are *neighbors at the same level* if they are neighbors and have the same size. If a cube is split, it is partitioned by three planes orthogonal to the coordinate axes into eight smaller cubes of equal size.

A leaf node is *splittable* if it satisfies the following condition, which was proposed by the authors and a collaborator [7] for surface reconstruction from clean data.

Let x be a leaf node. Let S denote the partition of x into $8 \times 8 \times 8$ disjoint cubes with size $\ell_x/8$. The node x is *splittable* if at least two cubes in S contain some input data points.

Fig. 1 illustrates the idea in 2D. The left figure shows an octree node x containing 3 data points. After dividing x into 8^2 square of equal sizes, two of these squares contain some data points. Therefore, x is splittable. The right figure is a non-splittable node that contains 3 data points.

For every octree node x , we say that x is *empty* if there is no data point inside x . We call an octree *balanced* if neighboring leaf nodes differ in size by at most a factor 2.

The initial octree consists of just the root corresponding to the bounding cube B . We maintain a queue Q that contains all the current, non-empty leaf nodes whose splittability have not been checked yet. Initially, B is the only item in Q . In a generic step, we dequeue a leaf node x from Q and calls a procedure $\text{SPLITBALANCE}(x)$ until Q becomes empty. We maintain the invariant that the current octree is balanced after each call of SPLITBALANCE .

In $\text{SPLITBALANCE}(x)$, we check whether x is splittable, and if so, we call a procedure $\text{SPLIT}(x)$. In $\text{SPLIT}(x)$, we first partition x into eight equal-sized cubes z_i , $i \in [1, 8]$, which become the children of x in the octree, and then we call a procedure $\text{BALANCE}(z_i)$ for all $i \in [1, 8]$, and finally we enqueue the leaf nodes z_i , $i \in [1, 8]$, into Q . In $\text{BALANCE}(z_i)$, we check whether the octree is unbalanced, which can happen if and only if for some neighboring leaf node y of z_i , $\ell_y = 2\ell_x$ and $\ell_y = 4\ell_{z_i}$. That is, x and y were just “in balance” and z_i and y are not “in balance” after the splitting of x . For every such leaf node y , we call $\text{SPLIT}(y)$ recursively to restore the balance.

Remark 1: The non-empty leaf nodes in the final octree have the appropriate sizes in the sense that if we select one data point from each non-empty leaf node, then the selected data points form a locally uniform sample of the unknown surface,¹ i.e., the local density of the selected sample points varies smoothly over the unknown surface.

Let T denote the resulting octree. Although this octree definition cannot handle outliers and noise, it serves as a good starting point. The key observation is that if the point cloud is clean and it is a fairly uniform sample of the unknown surface, then most non-empty leaf nodes should have the same size. Even in the presence of outliers and noise, the mean octree cell size ℓ_{avg} is informative.

As a result, we use a real-valued parameter $\alpha \geq 1$ and convert the octree T into another octree such that every leaf node has the same size ℓ_P , where $\ell_P \in (\frac{1}{2}\alpha\ell_{\text{avg}}, \alpha\ell_{\text{avg}}]$. (We set $\alpha = 2$ in almost all of our experiments.) This is done as follows. Take a leaf node x in the current octree. If $\ell_x > \alpha\ell_{\text{avg}}$, split x ; otherwise, if $\ell_x \leq \frac{1}{2}\alpha\ell_{\text{avg}}$, delete x . We repeat the above until $\frac{1}{2}\alpha\ell_{\text{avg}} < \ell_x \leq \alpha\ell_{\text{avg}}$ for every leaf node. We use T_P to denote the final octree obtained. Fig. 2 illustrates the octree construction in 2D. The fact that every leaf node has the same size ℓ_P is useful for our outlier detection to be discussed in Section 3.1.

3 Denoising

3.1 White Noise and Outlier Removal

Fig. 3 shows some noisy points sampled from a circle surrounded by white noise and outliers. Our outlier removal procedure is based on two ideas. First, the data points near the true surface are denser than the white noise. However, outliers may also form dense clusters which appear as dark, thick dots in Fig. 3. Our second idea is that the data points near the true surface occupy a bigger space than a coincidental cluster of outliers. We explain below how to turn these two ideas into efficient procedures.

We construct a graph G in which the vertices are the leaf nodes in T_P , and two vertices in G are connected by an edge if the corresponding leaf nodes are neighbors at the same level of T_P . We run a breadth-first search on G to obtain its connected components. Each connected component is a cluster of data points in T_P , and we expect the data points near the true surface to induce a large connected component. The white noise points are sparser, and therefore, a leaf node of T_P that contains white noise points may well be at distance greater than $\alpha\ell_{\text{avg}} \geq \ell_P$ from other leaf nodes containing white noise points. A cluster of outliers may induce a connected component in G . However, such a connected component has very few vertices when compared with the connected component induced by the data points near the true surface. The same can be said for connected components induced by white noise points that happen to be near each other. Consequently, if P is sampled from a single surface, we can simply extract the data points in the largest connected component in G . If P is sampled from $k \geq 1$ surfaces, then we extract the k largest connected components in G . The knowledge of k is a rather weak requirement because the user often knows how many surfaces are to be reconstructed. One can also try $k = 1$, check the output, and then increase k if necessary.

A larger α tolerates a higher non-uniformity in the density of the data points near the true surface, but fewer outliers and white noise points are removed. If α is smaller, more outliers and white noise points are removed, but the point density near the true surface should be more uniform.

¹This fact is proved in [7] for a slightly different octree decomposition in which the node splitting and node balancing do not interleave. That is, all splittable nodes are split first before the restoration of balance. Interleaving node splitting and balancing gives a better practical performance.

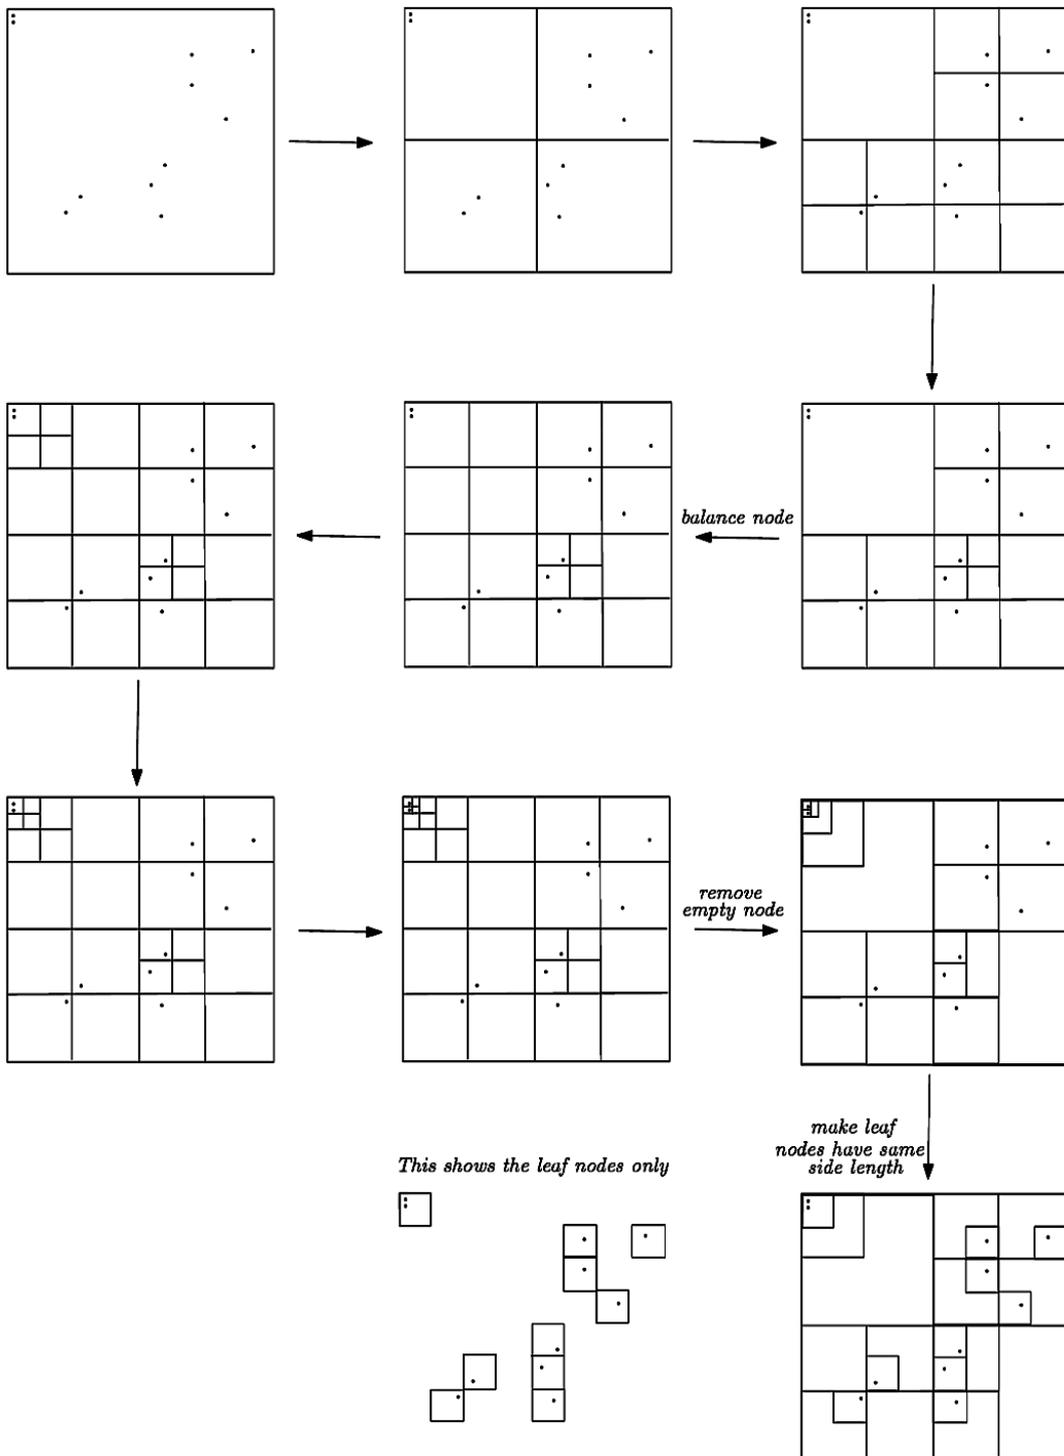


Figure 2: 2D case example of octree construction.

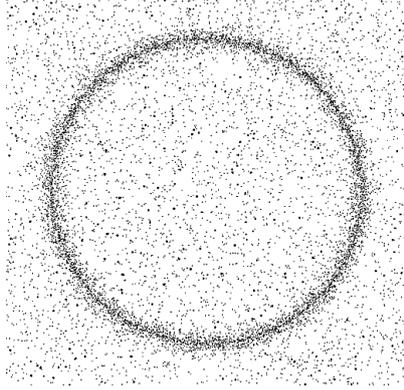


Figure 3: Noisy sample of a circle.

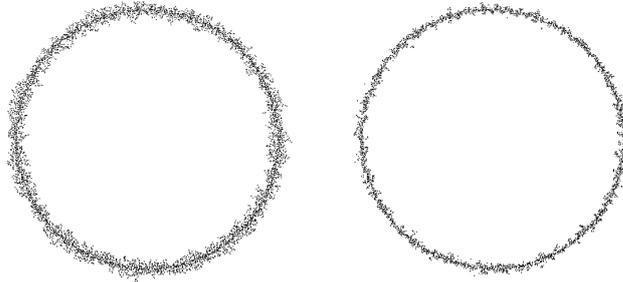


Figure 4: Left: after removing white noise and outliers. Right: after removing the very noisy points.

3.2 Filtering Very Noisy Points

Let P' denote the subset of P that we extracted from the largest k connected components of G . When the noisy perturbation is large, some points in P' may still be quite noisy and relatively far from the true surface. We eventually use a variant of meshless Laplacian smoothing to handle the noisy perturbation; however, according to our experience, the smoothing works less satisfactorily when some points in P' suffer from large noisy perturbation. We discuss in this subsection how to remove very noisy points from P' .

The left image in Fig. 4 shows the resulting point set after removing the white noise and outliers in Fig. 3. We want to remove the peripheral points (very noisy points). These peripheral points have a relatively sparser neighborhood (i.e., fewer points in the neighborhood). We introduce an efficient pruning method based on this observation.

We prune T_P to $T_{P'}$ by deleting the points not in P' and removing the nodes that become empty afterwards. Define the *neighborhood size* of a leaf node x in $T_{P'}$ be the number of data points in the cube that is centered at the center of x with side length $5\ell_P$. Let n_{avg} denote the average and n_{sd} denote the standard deviation of the neighborhood sizes among all the leaf nodes in $T_{P'}$. Let $\beta \geq 1$ be a parameter. (We set $\beta = 2$ in our experiments.) If $\beta n_{\text{sd}} > n_{\text{avg}}$, there is still a large variation in neighborhood sizes due to very noisy points, and therefore, we remove the points in leaf nodes that have neighborhood sizes at or below the 1 percentile. Afterwards, we update n_{avg} , n_{sd} , and the neighborhood sizes of the nodes affected. We repeat the above removal of points until $\beta n_{\text{sd}} \leq n_{\text{avg}}$. In Fig. 4, the right image shows the result of pruning the left image.

3.3 Meshless Laplacian Smoothing

Let P'' be the remaining point set after removing from P the white noise, outliers, and the very noisy points as described previously. Laplacian smoothing is a classic algorithm for smoothing a polygonal mesh. It works by moving a point p_i iteratively along the direction which is a weighted average of the vectors from p_i to its 1-ring neighbors. We use a variant of this scheme. In the following, we first introduce our scheme by showing how it works on P'' . Eventually, we construct another point set to replace P'' and run our scheme on this point set instead.

For every point $p_i \in P''$, let $N(p_i)$ denote a set of neighbors of p_i that will be defined later. For each $p_i \in P''$ and for each $p_j \in N(p_i)$, define

$$\begin{aligned} d(p_i) &= \max_{p_j \in N(p_i)} \|p_i - p_j\|, \\ w(p_i, p_j) &= \exp(-\|p_i - p_j\|^2/d(p_i)^2). \end{aligned}$$

Let λ be a parameter that controls the step size of moving p_i in each iteration. Let γ be a parameter so that a point p_i is moved only if p_i will move over a distance greater than the average distance between p_i and $N(p_i)$ divided by γ . (We set $\lambda = 0.25$ and $\gamma = 40$ in our experiments.) The pseudocode of the smoothing process is given below.

SMOOTH(P'' , λ , γ)

1. For each point $p_i \in P''$, compute

$$m(p_i) = \frac{1}{|N(p_i)|} \sum_{p_j \in N(p_i)} \|p_i - p_j\|.$$

2. For each point $p_i \in P''$, compute

$$p'_i = p_i + \lambda \cdot \frac{\sum_{j \in N(p_i)} w(p_i, p_j)(p_j - p_i)}{\sum_{j \in N(p_i)} w(p_i, p_j)}.$$

3. For each point $p_i \in P''$, if $\|p_i - p'_i\| > m(p_i)/\gamma$, then set $p_i = p'_i$.

SMOOTH(P'' , λ , γ) is called repeatedly until no point in P'' moves anymore or until the number of calls reaches a predefined threshold.

In the absence of a mesh, we cannot define $N(p_i)$ to be the 1-ring neighbor of p_i . The obvious alternative is to define $N(p_i)$ to be the points in P'' within some distance from p_i . However, there are two problems with this approach. First, there can be quite a lot of points within a predefined distance from p_i . Working with all these points for every p_i slows down the computation, and according to our experience, the output does not appear to be better. Second, the point distribution in P'' can be uneven and this uneven distribution may make p_i drift in the direction of higher local density. We observe that this may produce poorer output.

We address these two problems by building an octree for P'' and use the octree to define a smaller point set Q that is representative of P'' . The set Q has fewer points than P , so it improves the computational efficiency. We still have to define $N(q)$ for every $q \in Q$ to address the issue of non-uniform local density. Eventually, we forget about P'' and smooth Q instead, i.e., we call SMOOTH(Q , λ , γ) repeatedly until no point in Q moves or until the number of repetitions reaches a predefined threshold. In the following, we describe the construction of Q and $N(q)$ for every $q \in Q$.

First, construct an octree $T_{P''}$ for P'' as described in Section 2. Second, for each leaf node x of $T_{P''}$, compute the mean point q_x of the subset of P'' in x . The set Q is the resulting collection of mean points. By Remark 1 in Section 2, we expect Q to be a noisy locally uniform sample.

Consider a point $q_x \in Q$. Recall that ℓ_x denotes the size of the leaf node x in $T_{P''}$. Let B be the ball centered at q_x with radius $4\ell_x$. Let C be the axes-aligned cube centered at q_x with side length ℓ_x . Partition the boundary of C into 24 equal squares of side length $\ell_x/2$. For any two points $q_y, q_z \in Q \cap B$, we put q_y and q_z into the same group if the rays from q_x through q_y and q_z , respectively, intersect the same square in the partition of the boundary of C . As a result, $Q \cap B$ are divided into at most 24 groups. For each group, we pick the point in the group that is closest to q_x , and include this point in $N(q_x)$. Therefore, $N(q_x)$ consists of one point from each non-empty group. Selecting one point per group addresses the problem of non-uniform local density. Fig. 5 illustrates this grouping method in 2D.

Running SMOOTH(Q , λ , γ) repeatedly gives our final denoised point set. Fig. 6 shows the effect of our meshless Laplacian smoothing on the right image in Fig. 4. Only the upper right quarter of the circle is shown.

4 Postprocessing

Let P_f denote the final denoised point set. In our experiments, we run Robust Cocone [11] on P_f because no normal information is required.

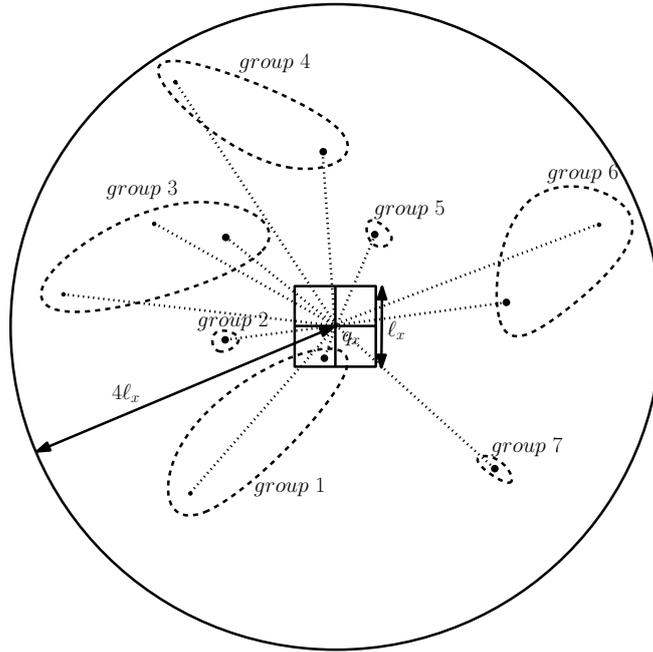


Figure 5: A 2D example of the grouping process. The boundary of the square centered at q_x is divided into 8 segments. The dotted lines are parts of rays from q_x to points in $Q \cap B$. The points enclosed by the same dashed curve are in the same group. The thicker dots are the points in $N(q_x)$.

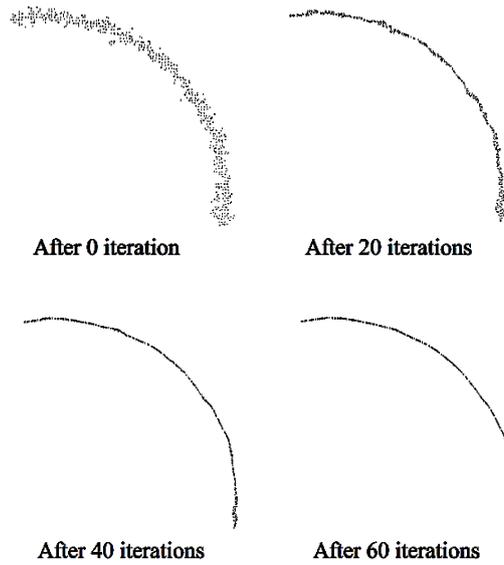


Figure 6: A 2D example of the meshless Laplacian smoothing effect.

Robust Cocone fills all holes in the point cloud to produce a water-tight surface. So boundary holes are patched as well. We observe that the triangles that fill a boundary hole usually have larger circumradii because the Robust Cocone algorithm reconstructs a surface based on Delaunay triangulation. We recreate the boundary holes by removing triangles according to a simple statistical criterion. Let r_{avg} be the average circumradius of the triangles in the reconstruction. Let r_{sd} be the standard deviation of the circumradii of these triangles. Let ε be a parameter. (We set $\varepsilon = 10$ in our experiments.) We remove all triangles whose circumradii are greater than $r_{\text{avg}} + \varepsilon r_{\text{sd}}$. This step can be skipped if a water-tight surface is desired.

The resulting triangular mesh may not be sufficiently smooth as mentioned in [11]. Therefore, we smooth the mesh using the classic Laplacian smoothing based on the 1-ring neighbors in the mesh.

5 Experiments

We implemented our algorithm and did the experiments on a Linux platform on a PC with an AMD Opteron(tm) processor 844 (1792.513 MHz) and about 8GB free memory available to our experiments. We compare our code with the following reconstruction codes:

- Method 1: Robust Cocone [11]
- Method 2: AMLS [12]
- Method 3: Sign-the-Unclassified [22]
- Method 4: Noise-adaptive distance function [13]

The codes for methods 1, 2, 3, and 4 were obtained from the authors.

Methods 1 and 2 are not catered for outliers and white noise. Therefore, the input to methods 1 and 2 in our experiments are obtained after removing outliers and white noise as described in Section 3.1. Noisy perturbation may remain though. Methods 1, 2, and 4 were run on the same machine as our code. Method 3 was run on the Window 7 platform on a PC with an Intel Core2 duo CPU E8400 3.00GHz and 4GB RAM because the executable of Sign-the-Unclassified runs on Window platform. We are mostly concerned with comparing the noise handling capability of our code with the other codes. Although the execution times of our code are almost always the smallest in the experiments, it is not meaningful to compare execution times because the platforms may be different, the codes may not be optimized fully, and the codes may not be compiled with the best possible options for the respective computing environments. Thus, we only show our execution times but not the others.

Due to storage constraint, all images in Figures 7–16 are not included in this document. They can be found at <https://www.cse.ust.hk/faculty/scheng/self-paper-short-images.pdf>.

5.1 Data Sets

We used 7 data sets: Armadillo, Bunny, Chair, Dragon, Happy Buddha, Ramesses, and wFish. Bunny, Chair, Dragon, Ramesses, and wFish are raw point sets. In particular, the Dragon data set contains sample points from the background. Armadillo and Happy Buddha consist of vertices of reconstructed meshes.

To each point in each data set, we add a Gaussian noise with mean zero and standard deviation x in a direction chosen uniformly at random. There are four different settings of x : 0%, 0.5%, 1% and 2% of the diagonal length D of the minimum axes-aligned bounding box of the data set.

To each data set, 5000 white noise points are picked from a uniform distribution in the minimum axes-aligned bounding box of the data set. Each white noise point p has a probability 0.05 to become a cluster of outliers if there is no sample point inside the ball centered at p with radius $0.05D$. Let r_1 be an integer chosen uniformly at random from $[1, R]$, where R is some integer fixed *a priori*. Let r_2 be a real number chosen uniformly at random from $[0, 1]$. If p is chosen to be a cluster of outliers, then we pick r_1 outlier points uniformly at random inside the ball centered at p with radius $0.001r_2D$.

The first column in Fig. 7 shows two noisy input data sets with 5000 white noise points, outliers generated as described above with $R = 400$, and 1% Gaussian noise. The second column shows the point sets after removing the white noise and outliers. The third column shows the final denoised point sets produced by our denoising code.

To push the codes to the extreme, we also experimented with data sets contaminated with 60%, 80% and 100% white noise points followed by random outlier cluster generation as discussed above (no Gaussian noise is added). Fig. 8 shows two examples.

5.2 Parameters

As discussed in Section 3, our algorithm requires several parameters $\alpha, \beta, \lambda, \gamma, \epsilon$, and the threshold on the number of calls to $\text{SMOOTH}(Q, \lambda, \gamma)$. In our experiments, we set $\alpha = 2, \beta = 2, \lambda = 0.25, \gamma = 40$ and $\epsilon = 10$, but we change α and λ for Happy Buddha, the Tortile column, and Pierre’s clenched fist.

The larger α is, the larger the octree leaf nodes (Section 2). Larger leaf nodes are needed when the sampling is significantly non-uniform at some places. We increase α to 4 for Pierre’s clenched fist due to the significantly non-uniform sampling (Fig. 9(b)). When we work with data sets with 60%, 80% and 100% white noise points, the setting of $\alpha = 2$ still works for almost all data sets except that we need to set $\alpha = 2.5$ for Happy Buddha.

The larger λ is, the greater the step size in Laplacian smoothing, which gives a greater smoothing effect. The setting of $\lambda = 0.25$ works for the Gaussian noise levels of 0%, 0.5%, 1%, and 2% that we tried. If it is known that the noise level is small, say less than 0.25%, then λ can be reduced and the effect is that more features can be preserved. For the Tortile column, the setting of $\lambda = 0.1$ preserves more features than the setting of $\lambda = 0.25$ (Fig. 9(a)).

The threshold on how many times $\text{SMOOTH}(Q, \lambda, \gamma)$ is called is determined before smoothing. Translate and scale Q so that a minimum axes-aligned bounding cube of Q centered at the origin has side length 2. For each point $q \in Q$, let d_q denotes the average distance between q and $N(q)$ (Section 3.3). Let d_{avg} be the average of d_q over all $q \in Q$. Define $L = \lfloor (d_{\text{avg}}^2 \cdot |Q|) / 2 \rfloor$ to be the threshold desired, which works well in our experiments.

Sign-the-Unsigned has a maximum depth level parameter. A higher depth level means finer level of details in the output, possibly at the expense of a higher running time. We set the maximum depth level to be 10. The depth levels used in [22] are between 8 and 12.

5.3 Experimental Results and Comparison

Consider the experiments with 5000 white noise points, randomly generated clusters of outliers, and $x\%$ Gaussian noise, where $x \in \{0, 0.5, 1, 2\}$. Fig. 10–14 show some reconstructions by our code and methods 1–4. Tables 1 gives the statistics of these experiments. Recall that the input to Robust Cocone and AMLS has been subject to white noise and outlier removal using our code (noisy perturbation may remain). Noise-Adaptive seems to have difficulty in handling the white noise and clustered outliers in our experiments, and it produced corrupted output.

In the experiments with 0% Gaussian noise, Robust Cocone may preserve more detailed features as seen in the Armadillo and Ramesses data sets (Figs. 10 and 11), but it may give corrupted output as seen in the Dragon data set. Our code, AMLS, and Sign-the-Unsigned also preserve a lot of detailed features. But AMLS and Sign-the-Unsigned may lose some detailed features as seen in the Ramesses data set. Sign-the-Unsigned had difficulty with the structured and clustered outliers in the raw Dragon data set (Fig. 13).

The Chair and Happy Buddha data sets are challenging. Refer to Fig. 12. In the experiments with 0% Gaussian noise in the Chair data set, AMLS produced a rough surface, and Sign-the-Unsigned produced wrong topology. In the experiments with 0.5%, 1% and 2% Gaussian noise in the Chair data set, Robust Cocone, AMLS, and Sign-the-Unsigned produced corrupted output. Our code still recovered the topology correctly and the surface quality degrades gradually as the Gaussian noise increases. Similar trends are observed in the experiments with the Happy Buddha data set (Fig. 14), except that our code produced wrong topology at 1% and 2% Gaussian noise. The Dragon data set is also difficult for our code (Fig. 13). The output is already corrupted at 0.5% Gaussian noise (in the sense that the neck is merged with the back).

The output of our code, Robust Cocone and AMLS show that it is insufficient to remove white noise and outliers alone. (The input to Robust Cocone and AMLS has been subject to removal of white noise and outliers by our code.) The remaining noisy perturbation is still so large that Robust Cocone and AMLS produced corrupted output.

Our code shows a better noise handling ability in the above experiments. Its execution times are reasonably small as seen in Table 1. Sign-the-Unsigned is also quite robust in the above experiments. Therefore, we push our code and Sign-the-Unsigned to the extreme by introducing a lot more white noise

Table 1: Each data set is contaminated with 5000 white noise points, randomly generated outlier clusters, and Gaussian noise. The input size includes the perturbed data points, white noise, and outliers. The column “Denoising Time” gives the execution time of our denoising step. The column “Total Time” is the total running time including the denoising step and the surface reconstruction step. Our code, Sign-the-Unsigned, and Noise-Adaptive were run on the same contaminated data sets, but Robust Cocone and AMLS were run on point sets that have been subject to white noise and outlier removal by our code. A cell is marked \times if the corresponding reconstruction is corrupted or unsuccessful. Otherwise, the cell is marked \checkmark .

	Gaussian Noise Level	Input size	Denoising Time	Total Time	Our Code	Robust Cocone	AMLS	Sign-the-Unsigned	Noise-Adaptive
Armadillo	0%	197640	14.36s	132.51s	\checkmark	\checkmark	\checkmark	\checkmark	\times
	0.5%	199518	29.38s	113.70s	\checkmark	\times	\times	\checkmark	\times
	1%	199117	44.48s	112.99s	\checkmark	\times	\times	\checkmark	\times
	2%	199002	45.71s	84.33s	\checkmark	\times	\times	\times	\times
Bunny	0%	385113	38.62s	346.24	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
	0.5%	392522	91.51s	240.89s	\checkmark	\checkmark	\times	\checkmark	\times
	1%	380930	125.29s	270.58s	\checkmark	\times	\times	\checkmark	\times
	2%	382294	145.87s	290.91s	\checkmark	\times	\times	\checkmark	\times
Ramesses	0%	824488	54.87s	1055.11s	\checkmark	\checkmark	\checkmark	\checkmark	\times
	0.5%	821297	143.45s	591.97s	\checkmark	\checkmark	\checkmark	\checkmark	\times
	1%	824296	161.52s	597.21s	\checkmark	\checkmark	\times	\checkmark	\times
	2%	825584	176.78s	389.87s	\checkmark	\checkmark	\times	\checkmark	\times
Chair	0%	246042	16.48s	242.10s	\checkmark	\checkmark	\checkmark	\times	\times
	0.5%	242940	28.74s	91.03s	\checkmark	\times	\times	\times	\times
	1%	252343	38.64s	105.36s	\checkmark	\times	\times	\times	\times
	2%	242693	45.00s	113.00s	\checkmark	\times	\times	\times	\times
Dragon	0%	2767387	351.39s	3512.06s	\checkmark	\checkmark	\checkmark	\times	\times
	0.5%	2767110	1018.29s	3537.03s	\times	\times	\times	\times	\times
	1%	2767461	1030.59s	3690.72s	\times	\times	\times	\times	\times
	2%	2768586	1401.62s	3026.28s	\times	\times	\times	\times	\times
Happy Buddha	0%	590779	36.72s	605.94s	\checkmark	\checkmark	\times	\times	\times
	0.5%	583277	89.37s	412.26s	\checkmark	\times	\times	\times	\times
	1%	588542	111.55s	247.29s	\times	\times	\times	\times	\times
	2%	587481	137.32s	298.02s	\times	\times	\times	\times	\times
wFish	0%	303478	14.87s	122.41s	\checkmark	\checkmark	\checkmark	\checkmark	\times
	0.5%	306185	43.28s	144.17s	\checkmark	\times	\checkmark	\times	\times
	1%	304348	62.88s	157.26s	\checkmark	\times	\times	\checkmark	\times
	2%	303475	75.95s	170.97s	\checkmark	\checkmark	\times	\checkmark	\times

in the another set of experiments. To each of the data sets Armadillo, Bunny, Chair, Dragon, Happy Buddha, Ramesses, and wFish, we introduce 60%, 80%, and 100% uniformly distributed white noise points. Then, we probabilistically turn white noise points into clusters of outliers as described in Section 5.1. No Gaussian noise is added. Sign-the-Unclassified could only produce useful output for Armadillo and Bunny in the experiments with 60% white noise points. Fig. 15 shows these two output surfaces. Our code handled these contaminated data sets quite well. Fig. 16 shows some of our reconstructions in comparison with those obtained in the first set of experiments when 5000 white noise points were introduced. The white noise and outliers may cause too many point deletions so that some areas become under-sampled. This will lead to a rough surface and/or holes. For example, there are dents in the base of the Happy Buddha, and there are some holes in Ramesses and wFish.

6 Conclusion

We propose a fast and simple algorithm to denoise an unorganized point cloud for surface reconstruction. It can efficiently handle a large amount of white noise, outliers, and fairly large noisy perturbation. In particular, we can also filter outliers that are clustered together, which may happen to be sample points from the background during the scanning of an object in the foreground.

There are two open research problems. One is to determine the number of surfaces automatically. Another problem is handle sharp features and non-manifold features.

Acknowledgment

The authors would like to thank Jiongxin Jin for helpful discussion and the authors of [11, 12, 22, 13] for providing their codes.

References

- [1] https://ihome.ust.hk/~lmkaa/cgi-bin/DenoisingSurfRecon/DenoisingSurfRecon_projectPage.html.
- [2] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. In *Proceedings of the Annual Symposium on Computational geometry*, pages 39–48, 1998.
- [3] N. Amenta, M. Bern, and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. In *Proceedings of SIGGRAPH*, pages 415–421, 1998.
- [4] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proceedings of the Annual Symposium on Computational geometry*, pages 213–222, 2000.
- [5] N. Amenta, S. Choi, and R.K. Kolluri. The power crust, unions of balls, and the medial axis transform. *Computational Geometry: Theory and Applications*, 19:127–153, 2000.
- [6] F. Bernardini, J. Mittleman, H. Rushmeier, C. Silva, and G. Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 5:349–359, 1999.
- [7] S.-W. Cheng, J. Jin, and M.-K. Lau. A fast and simple surface reconstruction algorithm. In *Proceedings of the Annual Symposium on Computational Geometry*, pages 69–78, 2012.
- [8] T.K. Dey and J. Giesen. Detecting undersampling in surface reconstruction. In *Proceedings of the Annual Symposium on Computational geometry*, pages 257–263, 2001.
- [9] T.K. Dey, J. Giesen, and J. Hudson. Delaunay based shape reconstruction from large data. In *Proceedings of the IEEE Symposium on Parallel and Large-Data Visualization and Graphics*, pages 19–27, 2001.
- [10] T.K. Dey and S. Goswami. Tight cocone: a water-tight surface reconstructor. In *Proceedings of the ACM Symposium on Solid Modeling and Applications*, pages 127–134, 2003.
- [11] T.K. Dey and S. Goswami. Provable surface reconstruction from noisy samples. *Computational Geometry: Theory and Applications*, 35(1):124–141, 2006.
- [12] T.K. Dey and J. Sun. An adaptive MLS surface for reconstruction with guarantees. In *Proceedings of the Eurographics Symposium on Geometry Processing*, 2005.
- [13] S. Giraudot, D. Cohen-Steiner, and P. Alliez. Noise-adaptive shape reconstruction from raw point sets. In *Proceedings of the Eurographics Symposium on Geometry Processing*, volume 32, 2013.

- [14] F. Guggeri, R. Scateni, and R. Pajarola. Shape reconstruction from raw point clouds using depth carving. In *Proceedings of EUROGRAPHICS*, 2012.
- [15] H. Huang, D. Li, H. Zhang, U. Ascher, and D. Cohen-Or. Consolidation of unorganized point clouds for surface reconstruction. *ACM Transactions on Graphics*, 28(5): Article No. 176, 2009.
- [16] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Proceedings of the Eurographics Symposium on Geometry Processing*, pages 61–70, 2006.
- [17] R. Kolluri. Provably good moving least squares. *ACM Transactions on Algorithms*, 4(2): Article No. 18, 2008.
- [18] R. Kolluri, J.R. Shewchuk, and J.F. O’Brien. Spectral surface reconstruction from noisy point clouds. In *Proceedings of the Symposium on Geometry Processing*, pages 11–21, 2004.
- [19] D. Levin. Mesh-independent surface interpolation. In G. Brunett, B. Hamann, K. Mueller, and L. Linsen, editors, *Geometric Modeling for Scientific Visualization*. Springer-Verlag, 2003.
- [20] W.E. Lorensen and H.E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *Proceedings of SIGGRAPH*, volume 21, pages 163–169, 1987.
- [21] B. Mederos, L. Velho, and L.H. de Figueiredo. Robust smoothing of noisy point clouds. In *Proceedings of the SIAM Conference on Geometric Design and Computing*, 2003.
- [22] P. Mullen, F. de Goes, M. Desbrun, D. Cohen-Steiner, and P. Alliez. Signing the unsigned: Robust surface reconstruction from raw pointsets. *Computer Graphics forum*, 29(5):1733–1741, 2010.
- [23] Y. Nagai, Y. Ohtake, and H. Suzuki. Smoothing of partition of unity implicit surfaces for noise robust surface reconstruction. In *Proceedings of the Eurographics Symposium on Geometry Processing*, volume 28, 2009.
- [24] O. Schall, A. Belyaev, and H.-P. Seidel. Adaptive fourier-based surface reconstruction. In *Proceedings of the International Conference on Geometric Modeling and Processing*, pages 34–44, 2006.
- [25] C. Shen, J.F. O’Brien, and J.R. Shewchuk. Interpolating and approximating implicit surfaces from polygon soup. *ACM Transactions on Graphics*, 23(3):896–904, 2004.
- [26] H. Xie, K.T. McDonnell, and H. Qin. Surface reconstruction of noisy and defective data sets. In *Proceedings of IEEE Visualization*, pages 259–266, 2004.
- [27] H. Xie, J. Wang, J. Hua, and A. Kaufman. Piecewise C^1 continuous surface reconstruction of noisy point clouds via local implicit quadric regression. In *Proceedings of the IEEE Visualization*, pages 91–98, 2003.

Appendix: Images

Due to storage constraint, the images are not included in this document. Please find them at <https://www.cse.ust.hk/faculty/scheng/self-paper-short-images.pdf>.