

# LFS-Aware Surface Reconstruction from Unoriented 3D Point Clouds

Rao Fu, Kai Hormann and Pierre Alliez

**Abstract**—We present a novel approach for generating isotropic surface triangle meshes directly from unoriented 3D point clouds, with the mesh density adapting to the estimated local feature size (LFS). Popular reconstruction pipelines first reconstruct a dense mesh from the input point cloud and then apply remeshing to obtain an isotropic mesh. The sequential pipeline makes it hard to find a lower-density mesh while preserving more details. Instead, our approach reconstructs both an implicit function and an LFS-aware mesh sizing function directly from the input point cloud, which is then used to produce the final LFS-aware mesh without remeshing. We combine local curvature radius and shape diameter to estimate the LFS directly from the input point clouds. Additionally, we propose a new mesh solver to solve an implicit function whose zero level set delineates the surface without requiring normal orientation. The added value of our approach is generating isotropic meshes directly from 3D point clouds with an LFS-aware density, thus achieving a trade-off between geometric detail and mesh complexity. Our experiments also demonstrate the robustness of our method to noise, outliers, and missing data and can preserve sharp features for CAD point clouds. The source code will be publicly available <sup>1</sup>.

**Index Terms**—I.3.5 [Computing Methodologies]: Computer Graphics – Computational Geometry and Object Modeling;

## I. INTRODUCTION

THE increasing popularity of 3D scanners has facilitated the capture of large-scale 3D point cloud data that are required to generate accurate 3D models. This has fueled the interest in surface reconstruction that refers to recovering a continuous surface (here a triangle surface mesh) from a 3D point cloud. Surface mesh reconstruction is relevant for numerous applications in various fields, such as computer animation [1], computer-aided design [2], and compression [3]. Arbitrarily reconstructing a surface mesh is often insufficient, as additional properties are often sought. In this work, we focus on generating a mesh with well-shaped (isotropic) elements and variable mesh sizing in accordance with the local feature size (LFS). Such meshes exhibit a satisfactory balance between the quality of the elements, level of complexity, and geometric fidelity. They are relevant for downstream applications, such as simulation and visualization.

Surface reconstruction has been explored through various methods, including explicit interpolation with Delaunay triangulations or Voronoi diagrams, as well as implicit functions coupled with iso-surfacing techniques. However, when the goal is to obtain isotropic meshes with adaptive sizing that ensures

both isotropic triangle elements and adjustments to surface features like curvature or local feature size, the conventional pipeline involves first reconstructing a dense mesh from point clouds and subsequently applying remeshing to increase the mesh quality. This cascaded approach separates reconstruction and remeshing, making it hard to produce such a mesh that preserves more details while using fewer triangles. Furthermore, obtaining these adaptive isotropic meshes is complex, requiring a delicate balance between mesh complexity and reconstruction accuracy.

We tackle instead reconstruction and remeshing altogether by reconstructing an adaptive isotropic mesh directly from an unoriented 3D point cloud. Our method is motivated by the desire to apply LFS-aware meshing via Delaunay refinement [4] on an implicit function whose zero level set delineates the reconstructed surface. LFS captures important local topological information: curvature, thickness, and separation. We guide the size of triangle elements using the LFS to preserve fine details—obtaining lower reconstruction errors than remeshing as a postprocessing step.

We summarize our contributions as follows:

- 1) We introduce a new method to estimate the local feature size from a 3D point cloud, which avoids constructing the medial axis via the Voronoi diagram; the latter approach being sensitive to noise and sampling conditions. Instead, our approach estimates LFS via jet fitting and analyzing a Lipschitz distance function along random rays. Such an LFS estimation approach is reliable and robust to low noise and non-uniform sampling.
- 2) We introduce a novel approach for solving a signed implicit function in three main steps. First, we generate and discretize in tetrahedra a multi-domain composed of a thin envelope around the input points and the complement in a loose bounding sphere of the input. The envelope is derived from an unsigned implicit function with a level set of the estimated reach that casts off outliers. Second, we solve a signed implicit function from the signing guess of edges in a least-squares sense, which is capable of being resilient to outliers and filling large holes. Finally, we construct a robust implicit signed function from the solved implicit signed function that offers robustness to noise.
- 3) We conduct extensive experiments on synthetic and real-world 3D point clouds, including the AIM@Shape dataset [5] and the ABC dataset [6]. Our method generates output meshes that are valid (i.e., intersection-free), with adaptive sizing and smaller reconstruction errors than previous approaches.

Rao Fu is with INRIA, France, and also with Geometry Factory, France (e-mail: rao.fu@inria.com).

Kai Hormann is with Università della Svizzera italiana (USI), Switzerland (e-mail: kai.hormann@usi.ch).

Pierre Alliez is with INRIA, France (e-mail: pierre.alliez@inria.fr).

<sup>1</sup>Code: <https://github.com/bizerfr/LFS-Aware-Reconstruction>




Fig. 1: A step-by-step reconstruction example. (a) Our algorithm takes the unoriented point sets as input. (b) First, we estimate the local feature size (LFS) directly on the inputs. (c) Second, we construct a reach-aware multi-domain. We show the boundary of the multi-domain: an envelope domain embedded in a sphere domain. (d) Third, we use a new mesh solver to solve an implicit function defined on the multi-domain whose zero level set is the target surface. We present a clip view of the implicit function. (e) We extract the LFS-aware mesh from the solved implicit function.

## II. RELATED WORK

The proposed approach involves LFS estimation, surface reconstruction, and iso-surfacing and remeshing. We now briefly review each component.

### A. LFS Estimation

The local feature size (LFS) on a 3D shape is the distance from a query point to its closest point on the shape's medial axis. The reach of a shape refers to the minimum of the LFS [7]. There are two main approaches to estimating the LFS: indirect and direct. Indirect methods commonly involve computing the medial axis [8], defined as the locus of centers of spheres that touch the shape's boundary at two or more unique points. Amenta and Bern [9] proved that the Voronoi poles provide a good estimation of the medial axis if the point set used to generate the corresponding Voronoi diagram is noise-free and satisfies the  $\epsilon$ -sampling conditions. Several approaches have been proposed based on this idea [9]–[12], which estimate the medial axis and then estimate the LFS by finding the nearest valid Voronoi pole. The limitations of these methods are sensitivity to noise and non-uniform sampling, which significantly affect the accuracy of the estimated medial axis. Other approaches [13]–[15] attempt to gain robustness by directly computing a medial axis in an optimization scheme, free of the construction of the Voronoi diagram. However, some of them require oriented normals as a prior. Direct methods avoid estimating the medial axis explicitly [16], [17]. Instead, they estimate the shape diameter function from the input shape. While direct methods are less sensitive to noise, estimating the shape diameter function on a 3D point cloud is challenging, because finding the antipodal point given a ray cast from an input point is ill-posed. Furthermore, the shape diameter function only measures the thickness of the shape, while LFS also captures separation and curvature. Departing from previous approaches, our LFS estimation approach relies on both the shape diameter function and curvature estimation via jet fitting [18]. The shape diameter function is estimated by analyzing the unsigned distance function computed from the input point cloud.

### B. Surface Reconstruction

Many surface reconstruction approaches have been proposed over the years, and we review the learning-based and non-learning-based separately.

1) *Learning-based*: Learning-based methods have drawn lots of attention in recent years. The occupancy network and its convolutional variant [19], [20] predict an occupancy probability for the grid and then extract the mesh based on the prediction. Point2Surf [21] learns an implicit signed distance function from a local patch of the point cloud. The Shape-As-Points method [22] proposes a differentiable formulation of Poisson surface reconstruction. However, these methods are constrained by the dataset and lack scalability. To solve the scalability issue, the POCO method [23] uses convolutions and computes latent vectors at each point to deduce an occupancy for large-scale point clouds. To solve the generality issue, the neural kernel reconstruction methods [24], [25] solve for an implicit function using optimization with the data-driven kernel function and thus gain generalizability and scalability. There are also other approaches that use neural networks to learn implicit representations [26]–[30].

2) *Non-learning-based*: Non-learning-based methods can be categorized into Delaunay-based, implicit-based, primitive-based, and hybrid methods. Delaunay-based methods utilize the Delaunay triangulation or Voronoi diagram to generate a mesh that interpolates the input points. The popular crust and Power Crust approaches use the Voronoi diagram to compute the medial axis of a shape and then construct a mesh from the Delaunay triangulation of the crust or power diagram [9], [10]. The Tight Cocone algorithm [31] extends the Cocone algorithm [32] by introducing a tightness criterion that ensures that the mesh is close to the input point cloud. The advancing front algorithm [33] constructs a mesh by greedily advancing a front of Delaunay triangles. The restricted Delaunay method [34] interpolates input points by film-sticking and sculpting. While Delaunay-based methods provide provable guarantees under certain conditions, they are sensitive to noise and outliers. Implicit-based methods are devised to represent the reconstructed surface as an isolevel of an implicit function. The function is often represented by a discretized 3D domain such as an octree, or defined using smoothness priors like radial

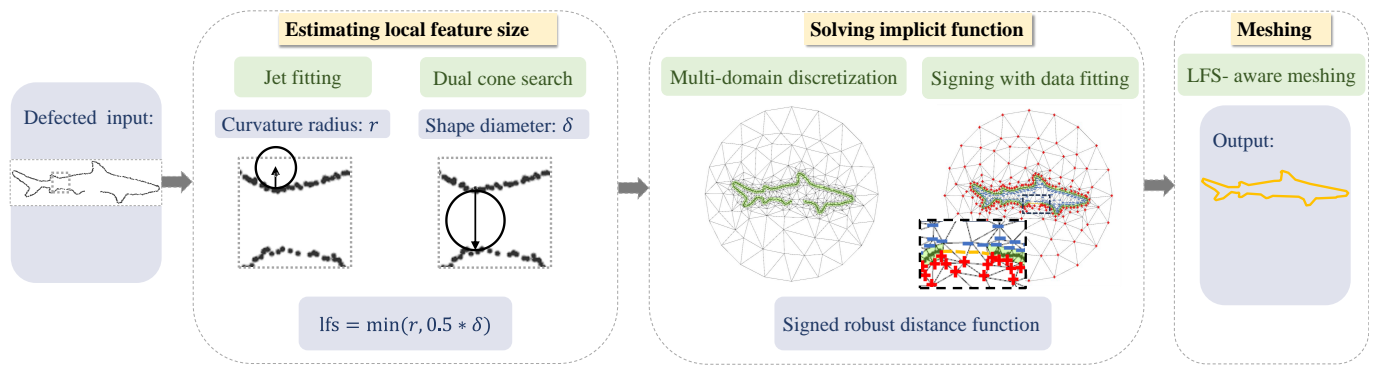


Fig. 2: Overview. The input is a defective point cloud, with or without normals. The algorithm first estimates LFS as the minimum of the local curvature radius and half of the shape diameter. An implicit function is solved on a multi-domain discretization obtained by Delaunay refinement, so as to fill large holes. Yellow dashed lines delineate the filled holes, “+” denotes the positive vertices, and “-” denotes the negative vertices. Finally, the output LFS-aware mesh is generated by Delaunay refinement.

basis functions (RBF), moving least squares (MLS) or kernel regression. Some implicit-based approaches require oriented normals as input, while others do not. More specifically, MLS-based methods [35], [36] rely on oriented normals, Poisson surface reconstruction approaches [37], [38] diffuse the input oriented normals on an octree before solving for a Poisson equation. If normals are unoriented, normal-orienting methods [35], [39], [40] shall be applied. Some approaches deduce an implicit signed function from the unoriented point set [41], [42], and a spectral approach solves a generalized eigenvalue problem to obtain a signed function from unoriented normals and covariance matrices [43]. VIPSS [44] solves for a signed function by Duchon’s energy. PGR [45] re-formulates the Poisson reconstruction by incorporating the Gauss formula, thus leading to a dense system that lacks scalability, while iPSR [46] runs Poisson reconstruction in an iterative manner and a recent work [47] orients point cloud normals by incorporating isovalue constraints to Poisson reconstruction. Some hybrid methods combine Delaunay triangulations and implicit functions. The AMLS approach [48] projects input points into an LFS-aware implicit surface and then uses the tight cone to interpolate the input points. A recent progressive approach interleaves refinement of a triangulated 3D domain with implicit mesh-based solvers [49]. RFEPS [50] preserves sharp features by using the restricted power diagram and a Poisson-reconstructed mesh. Implicit-based methods are often more robust to noise than Delaunay-based methods but may require a dense mesh. They often require a fine discretization of the 3D domain to capture the surface details, resulting in a complex uniform output mesh, even in flat areas. We refer to the surveys [51]–[53] for a detailed overview.

### C. Iso-surfacing and Remeshing

Iso-surfacing refers to extracting a sublevel from a given scalar field. Marching cubes [54] can extract a triangular mesh on a regular voxel grid or an octree. The Marching tetrahedra method [55] simplifies the marching cubes algorithm by replacing voxels with tetrahedral cells. In order to preserve the sharp features, dual contouring [56] with its variants

[57]–[59] optimizes vertex locations within the confines of individual cells using the quadratic error function. Besides, neural versions of marching cubes and dual contouring [60], [61] are also proposed to improve the performance. In order to obtain isotropic triangular meshes, GradNorm and its variant [62], [63] first tile space with fixed or adaptive tetrahedra and then extract a triangle mesh. Remeshing is used to improve mesh quality in terms of vertex sampling, regularity, and the shape of its elements [64]. The remeshing process involves computing a new mesh approximating the original mesh while meeting specific quality requirements. It proceeds either in some parameter space or directly on the input mesh. Only a few remeshing approaches [65]–[67] generate output meshes that are adapted to the local curvature or estimated LFS with an estimation process applied to the mesh. In contrast to these methods, our method estimates LFS directly on the input 3D point cloud. LFS is then used for isocontouring the implicit function via Delaunay refinement, resulting in a lower mesh complexity for a given level of detail.

## III. METHOD

The input of our algorithm is a 3D point cloud  $\mathbf{P} = \{\mathbf{p}_i \in \mathbb{R}^3\}_{i=1}^N$  with or without normals  $\mathbf{N} = \{\mathbf{n}_i \in \mathbb{R}^3\}_{i=1}^N$ , where  $N$  is the number of points. The output is an isotropic triangle surface mesh where the size of the triangle elements is controlled by the estimated LFS. Fig. 2 depicts an overview of our algorithm.

### A. LFS Estimation

Mathematically, LFS is defined as the distance from a query point on the surface to its closest point on the medial axis. Estimating the medial axis, e.g., from the Voronoi diagram [12] or finding the maximal empty ball [13], is difficult as it is sensitive to noise and sampling density. Instead, we estimate LFS directly from the input point set and bypass the construction of the medial axis. LFS captures a shape’s local curvature, thickness, and separation (see Fig. 3). Given a

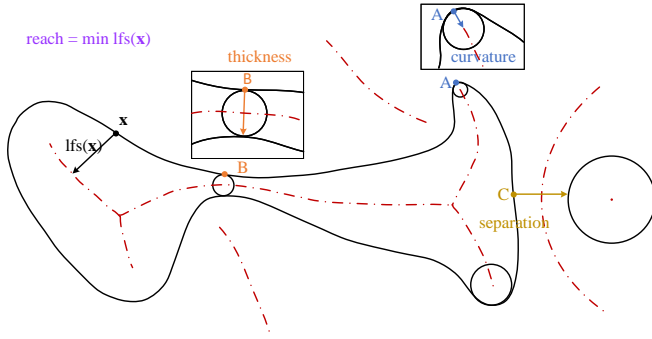


Fig. 3: LFS captures the local curvature, thickness, and separation. The reach is the minimum of the LFS for the whole shape. The red dashed line depicts the medial axis.

query point  $\mathbf{x}$  on a shape embedded in 3D space, we define LFS as

$$\text{lfs}(\mathbf{x}) = \min(r(\mathbf{x}), 0.5 * \delta(\mathbf{x})),$$

where  $r(\mathbf{x})$  denotes the estimated local curvature radius and  $\delta(\mathbf{x})$  denotes the estimated shape diameter function [16], [17]. The curvature radius captures the local curvature, while the shape diameter function captures the local thickness and separation of a shape. In addition, our surface reconstruction approach relies on the reach (i.e., the minimum LFS) to construct a reach-aware multi-domain discretization.

1) *Local curvature radius*: The local curvature radius is estimated by first fitting a differential jet locally [18]. Jet fitting constructs a local Monge coordinate system, whose basis  $(\mathbf{d}_1, \mathbf{d}_2, \mathbf{n})$  is defined by three orthogonal directions: maximum principal curvature, minimum principal curvature, and normal. The jet fitting approach fits a Monge form as

$$z = \frac{1}{2}(k_1x^2 + k_2y^2) + \text{hot},$$

where  $k_1$  and  $k_2$  denote the local principal curvatures,  $x$ ,  $y$  and  $z$  denote the coordinates, and *hot* denotes a higher-order polynomial term. The curvature radius for a query point  $\mathbf{x}$  is then defined as

$$r(\mathbf{x}) = 1/|h(\mathbf{x})|,$$

where  $|h(\mathbf{x})|$  is the curvature estimated via jet fitting with the maximum absolute value.

When the input 3D point cloud has no normal attributes, we directly derive the normals from the analytical jet surface. Note that the normals estimated via jet fitting are globally unoriented. We refer to [18] for more details.

2) *Shape Diameter*: The shape diameter function maps a query point  $\mathbf{x}$  on the surface of a 3D solid shape to a scalar defined as the distance from  $\mathbf{x}$  to its antipodal surface point with respect to the local normal direction. Previous approaches [16], [17] only consider the local thickness of the shape as they only search along the inward normal direction. Thus, the local separation is ignored. Furthermore, the original shape diameter function can only be estimated on a surface mesh where the antipodal points are well-defined given rays cast from a query point. Finding the antipodal points for an input 3D point cloud is hampered by the fact that the intersection point of a ray with a point cloud is ill-defined.

We address the two aforementioned issues by using a dual cone search algorithm that estimates an extended version of the shape diameter function  $\delta(\mathbf{x})$ , defined as

$$\delta(\mathbf{x}) = \min(\tau(\mathbf{x}), \sigma(\mathbf{x})),$$

where  $\tau(\mathbf{x})$  denotes the thickness defined as the distance from  $\mathbf{x}$  to its antipodal point along the inward normal direction, and  $\sigma(\mathbf{x})$  denotes the separation defined as the distance from  $\mathbf{x}$  to its antipode along the outward normal direction. It is worth noting that in our implementation, we conduct the search along the normal direction and its inverse direction, subsequently selecting the minimum of these two values. The final shape diameter value does not need any normal orientation.

To find the antipodal points, we analyze the unsigned distance function to the input point cloud  $\mathbf{P}$  along a ray  $\mathbf{r}$ . The unsigned distance function is defined as

$$d_u(\mathbf{x}) = \min_{\mathbf{p} \in \mathbf{P}} \{|\mathbf{x} - \mathbf{p}|\},$$

where  $\mathbf{x} \in \mathbb{R}^3$  denotes a query point. If  $d_u(\mathbf{x})$  is smaller than a small distance value referred to as  $\epsilon$ , we treat  $\mathbf{x}$  as the antipodal point. We estimate  $\epsilon$  from the unsigned robust distance function

$$\hat{d}_u(\mathbf{x}) = \sqrt{\frac{1}{k} \sum_{\mathbf{p} \in N_k(\mathbf{x})} \|\mathbf{x} - \mathbf{p}\|^2},$$

where  $N_k(\mathbf{x})$  denotes the  $k$  nearest neighbors to a query point  $\mathbf{x} \in \mathbb{R}^3$ . The parameter  $k$  is used to trade accuracy for robustness to noise and outliers [41]. Then  $\epsilon$  is estimated as the smallest  $\hat{d}_u(\mathbf{x})$  computed at all input points, i.e.,  $\epsilon = \min\{\hat{d}_u(\mathbf{p}) | \mathbf{p} \in \mathbf{P}\}$ . Equipped with this, the brute-force approach is to sample points densely on  $\mathbf{r}$  to find the antipodal points whose function value  $d_u(\mathbf{x})$  is below  $\epsilon$ . However, as  $d_u(\mathbf{x})$  is 1-Lipschitz continuous, we utilize a Lipschitz-guided recursive dichotomic search to avoid dense sampling via using the Lipschitz continuity as a prior (see Appendix). In addition, to capture both thickness and separation, we cast random searching rays inside two opposite cones with  $\mathbf{x}$  as the apex with  $\mathbf{n}$  and  $-\mathbf{n}$  as their axes, because it is challenging for a single ray to find an antipodal point directly. More specifically, we cast  $N_c$  random rays within each cone for a query point  $\mathbf{x}$  to collect  $k_c$  antipodal points and calculate the robust distance function from  $\mathbf{x}$  to the collected  $k_c$  antipodal points. Such a dual cone search may not find any antipodal points when the local point sampling is too sparse. In this case, we set the shape diameter as the diameter of the loose bounding sphere of the input point cloud. Fig. 4 illustrates the dual cone search for estimating the shape diameter.

Finally, the LFS value is taken as the minimum of the local curvature radius and half of the shape diameter. Then, we apply smoothing on the raw estimated LFS values.

## B. Implicit Function

Our objective is to solve for an implicit signed function, that is piecewise-linear and defined on a 3D triangulation, whose zero level set corresponds to the reconstructed surface. Instead of solving for the implicit function using a single global solve step, our rationale is to compute its components sequentially.

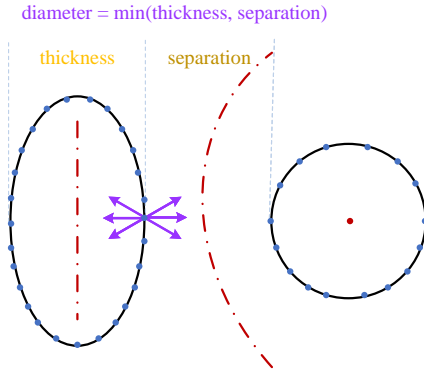


Fig. 4: Dual cone search. The red dashed line depicts the medial axis. The shape diameter is the minimum of the thickness and separation.

First, we construct a 3D multi-domain that discretizes an unsigned distance function to the input points (multi-domain). Next, we solve a least squares problem to obtain a signed implicit function from sign guesses estimated at the edges of the discretization (signing with data fitting). Finally, we utilize the signs of the solved implicit function to sign the unsigned robust distance function (signed robust distance function). We detail each step below.

1) *Multi-domain*: We construct a multi-domain discretization to represent the piecewise-linear implicit function. The multi-domain is the union of two 3D sub-domains: a thin envelope enclosing the input  $\mathbf{P}$  and a loose bounding sphere containing  $\mathbf{P}$  minus the above envelope. The thin envelope is made reach-aware to separate thin structures and preserve fine details, while the loose bounding sphere is added and discretized to help fill large holes with a smoothly graded 3D triangulation.

The thin envelope domain is derived from an unsigned implicit function  $I_u(\mathbf{x})$  and made read-aware by constraining  $I_u(\mathbf{x})$  to be smaller than the estimated reach  $I_R$ , i.e., the minimum of LFS. The definition is

$$I_E(\mathbf{x}) = \{\mathbf{x} : I_u(\mathbf{x}) \leq I_R\},$$

and  $I_u(\mathbf{x})$  is defined as

$$I_u(\mathbf{x}) = \frac{\sum_{\mathbf{p} \in N_k(\mathbf{x})} |(\mathbf{x} - \mathbf{p})^T \cdot \mathbf{n}| W(\mathbf{x})}{\sum_{\mathbf{p} \in N_k(\mathbf{x})} W(\mathbf{x})},$$

where  $W(\mathbf{x}) = e^{-\|\mathbf{x} - \mathbf{p}\|^2/h^2}$  denotes a Gaussian weighted function. Note that  $I_u(\mathbf{x})$  uses the absolute value of  $(\mathbf{x} - \mathbf{p})^T \cdot \mathbf{n}$ . Thus, normal orientation is not required here.

The other domain is the complement between the envelope and a loose bounding sphere. The loose bounding sphere is  $I_S(\mathbf{x}) = \{\mathbf{x} : \|\mathbf{x} - \mathbf{c}\| \leq r\}$ .  $\mathbf{c}$  is the centroid of  $\mathbf{P}$ , and  $r = \max_{\mathbf{p}, \mathbf{q}} \|\mathbf{p} - \mathbf{q}\|$ , where  $\mathbf{p} \in \mathbf{P}$ ,  $\mathbf{q} \in \mathbf{P}$  and  $\mathbf{p} \neq \mathbf{q}$ .

We utilize the Delaunay refinement paradigm [4] to generate the multi-domain discretization. In order to generate well-shaped cells, we specify a constant cell radius-edge ratio in all sub-domains. We also specify a constant facet shape criterion for the boundary of the loose bounding sphere. We refer to [4] for more details.

2) *Signing with Data Fitting*: Our multi-domain is a reach-aware envelope  $I_E(\mathbf{x})$  embedded in a 3D Delaunay triangulation  $\mathbf{Tr}$  bounded by a sphere  $I_S(\mathbf{x})$ . The sublevel of the unsigned function inside  $I_E(\mathbf{x})$  contains most of the inferred surface, but it remains to compute a signed function in order to fill holes and delineate the inferred surface as its zero level set. We proceed by solving for a signed function defined at the vertices of the multi-domain. We first estimate the sign guesses at the edges of the multi-domain. We then incorporate a data-fitting term to solve an implicit signed function value for each vertex in the least squares sense.

The sign guesses are estimated as follows. Let  $\text{sign}(\mathbf{e}_{mn})$  denote the binary sign guess for an edge  $\mathbf{e}_{mn} \in \mathbf{Tr}$ . If  $\mathbf{e}_{mn}$  crosses the inferred surface, then  $\text{sign}(\mathbf{e}_{mn})$  is set to  $-1$  to reflect that it connects two vertices with opposite signs and set to  $+1$  otherwise. An edge  $\mathbf{e}_{mn}$  has three possible locations: outside, inside, and on the boundary of the envelope. If  $\mathbf{e}_{mn}$  is outside or on the boundary of the envelope,  $\text{sign}(\mathbf{e}_{mn})$  is directly set to  $+1$  as  $\mathbf{e}_{mn}$  does not cross the inferred surface. If  $\mathbf{e}_{mn}$  is inside the envelope, we check whether  $\mathbf{e}_{mn}$  connects two vertices with opposite signs. Note that the edges of slivers (flat cells) inside the envelope may not cross the inferred surface. We still perform the Lipschitz-guided recursive dichotomic search (see Appendix) on  $\mathbf{e}_{mn}$  inside the envelope to check if the unsigned distance function  $I_u(\mathbf{x})$  is below a value  $\epsilon$ , which is set as half of the surface reach  $I_R$ . If  $I_u(\mathbf{x})$  is below  $\epsilon$ , we set  $\mathbf{e}_{mn}$  to  $-1$ , and  $+1$  otherwise. Equipped with sign guesses at edges, we then solve for a signed implicit function  $d_s(\cdot)$  at all vertices  $\mathbf{V}$  of  $\mathbf{Tr}$  by minimizing the objective function

$$E = \sum_{\mathbf{e}_{mn}} [d_s(\mathbf{v}_m) - \text{sign}(\mathbf{e}_{mn}) \cdot d_s(\mathbf{v}_n)]^2 + \lambda \sum_t \sum_{\mathbf{p}_t} \left[ \sum_{i=0}^3 \alpha_i(\mathbf{p}_t) \cdot d_s(\mathbf{v}_{t_i}) \right]^2, \quad (1)$$

where  $d_s(\cdot)$  denotes the signed implicit function solved for at the vertices  $\mathbf{V}$ ,  $\alpha_i(\mathbf{p}_t)$  is the  $i$ -th barycentric coordinate for the input point  $\mathbf{p}_t$  located in the tetrahedron  $t$  with vertices  $\mathbf{v}_{t_0}, \dots, \mathbf{v}_{t_3}$ .

The first term of the above objective function optimizes the sign consistency on the multi-domain. The second data fitting term favors the fact that the input points lie near the zero level set of the interpolated signed implicit function. The parameter  $\lambda$  provides a means to balance between the two terms. For all shown experiments,  $\lambda$  is set to 1.0.

To avoid the trivial solution, we enforce the sum of the signed implicit values to be a constant, i.e.,  $\sum_{\mathbf{v}} d_s(\mathbf{v}) = |\mathbf{V}|$ , where  $|\mathbf{V}|$  denotes the number of vertices. Fig. 5 illustrates the signing process.

We can re-write the minimization of the objective function (1) in matrix form as

$$\begin{aligned} \min_{\mathbf{x}} \quad & \|\mathbf{S}\mathbf{x}\|^2 + \lambda \|\mathbf{B}\mathbf{x}\|^2 \\ \text{s.t.} \quad & \mathbf{C}\mathbf{x} = |\mathbf{V}|, \end{aligned}$$

where  $\mathbf{S}$  denotes the sign guess matrix. Each row corresponds to an edge of  $\mathbf{Tr}$  and contains two elements, either set to  $+1$  or  $-1$ .  $\mathbf{B}$  are the barycentric coordinates of each input point, and  $\mathbf{x}$  are the values for the signed implicit function  $d_s(\cdot)$  that




Fig. 5: Solving for a signed implicit function from the sign guesses for edges. (a) The sign guess for an edge is determined by detecting the crossings with sublevel sets of the unsigned distance function to the points. An edge is labeled  $-1$  if it crosses the point clouds, and  $+1$  otherwise. (b) A signed implicit function is then solved from the sign guesses of the edges. Red crosses denote the positive vertices, and blue minus signs denote the negative vertices. The signed implicit function is piecewise linear.

we need to solve for at each vertex.  $\mathbf{C}$  is a one-dimensional constraint matrix of size  $1 \times |\mathbf{V}|$  and filled with 1. Both  $\mathbf{S}$  and  $\mathbf{B}$  are sparse. By applying the Karush–Kuhn–Tucker (KKT) conditions, we have

$$\min_{\mathbf{x}, \mathbf{z}} (\mathbf{S}\mathbf{x})^T(\mathbf{S}\mathbf{x}) + \lambda(\mathbf{B}\mathbf{x})^T(\mathbf{B}\mathbf{x}) + \mathbf{z}^T(\mathbf{C}\mathbf{x} - |\mathbf{V}|)$$

Finally, we solve for the signed implicit function values  $\mathbf{x}$  and the Lagrange multiplier  $\mathbf{z}$  in the least squares sense,

$$\begin{bmatrix} 2\mathbf{S}^T\mathbf{S} + 2\lambda\mathbf{B}^T\mathbf{B} & \mathbf{C}^T \\ \mathbf{C} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{z} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ |\mathbf{V}| \end{bmatrix}. \quad (2)$$

3) *Signed Robust Distance Function*: The previous step solves for a signed implicit function defined on the 3D multi-domain. To obtain a robust implicit function that is more resilient to noise, we use the binary sign value from the signed implicit function  $d_s(\cdot)$  to sign the unsigned robust distance function  $\hat{d}_u(\cdot)$  to obtain a new implicit function  $\hat{d}_s(\cdot)$  whose zero level set is the target surface

$$\hat{d}_s(\cdot) = \text{sign}(d_s(\cdot)) \cdot \hat{d}_u(\cdot), \quad (3)$$

where  $d_s(\cdot)$  denotes the solved signed implicit function values after solving (2).

### C. LFS-Aware Meshing

A simple iso-surfacing method such as marching tetrahedra [55] would not yield the final reconstructed surface mesh with the desired properties (sizing and well-shaped tetrahedra). We utilize instead a Delaunay refinement meshing [4] to extract the final size-varying isotropic mesh. The Delaunay refinement meshing takes an implicit function with a mesh sizing function as input, optionally with sharp feature curves. It outputs an isotropic triangular mesh contouring the zero level set of the input implicit function with the mesh sizing function controlling the refinement termination. If sharp feature curves are provided, the final mesh can preserve sharp features. In our case, we feed both the signed robust distance function  $\hat{d}_s(\cdot)$  together with mesh sizing function  $\text{size}(\cdot)$ , both of which are derived from the input point cloud directly, into the Delaunay refinement. Furthermore, we use NerVE [68] to detect the sharp feature curves for CAD shapes with sharp features.

We define the mesh sizing function, which is LFS-aware, as

$$\text{size}(\mathbf{x}) = \frac{\text{lfs}(\mathbf{x}) - I_R}{\text{lfs}_{\max} - I_R} \cdot (\text{size}_{\max} - \text{size}_{\min}) + \text{size}_{\min}, \quad (4)$$

where  $\text{lfs}(\mathbf{x})$  denotes the local feature size,  $I_R$  denotes the surface reach,  $\text{size}_{\max}$  denotes the user-specified maximal facet size, and  $\text{size}_{\min}$  denotes the minimal facet size set as a user-specified ratio to the thickness of the envelope. Finally, we apply smoothing to the facet sizing function.

## IV. EXPERIMENTS

The CGAL library [69] was used to implement the multi-domain discretization and Delaunay refinement steps. We use the conjugate gradient solver from the Eigen library [70] to solve for the implicit function. Experiments are conducted on a Dell laptop with a 2.60GHz Intel i7-10750H CPU and 32GB memory.

### A. LFS estimation

1) *Validation on canonical primitives*: We first validate the proposed LFS estimation algorithm on 3D point clouds sampled from canonical primitives. We generate canonical primitives for which we know the ground-truth LFS analytically. We then measure the error made by our estimation algorithm (without smoothing) and compare it with NormFet [12], which computes the medial axis via constructing the Voronoi diagram. More specifically, we sample three primitives (a sphere, a cone, and an ellipsoid) with non-uniform sampling and a small amount of added noise. We measure both the mean absolute error and max absolute error, as recorded in Table I. Our experimental

Capsule	Sphere (#648)		Cone (#2,610)		Ellipsoid (#16,374)	
	mean(↓)	max(↓)	mean(↓)	max(↓)	mean(↓)	max(↓)
NormFet	4.079E-1	8.191E-1	9.413E-1	2.233	7.835E-2	1.917E-1
Ours	<b>5.511E-3</b>	<b>3.190E-2</b>	<b>2.135E-1</b>	<b>7.751E-1</b>	<b>5.471E-2</b>	<b>1.786E-1</b>

TABLE I: LFS comparisons with NormFet on canonical primitives. “mean” denotes the mean absolute error. “max” denotes the maximal absolute error.

results show that our LFS estimation algorithm outperforms NormFet in terms of the mean and maximal absolute errors. The experiments highlight the robustness of our algorithm to non-uniform sampling and small amounts of noise. NormFet estimates LFS by first constructing the medial axis from the Voronoi diagram, but a good estimation of the medial axis requires a point cloud that is dense and noise-free. Recall that we define LFS as the minimum of the curvature radius and the shape diameter halved. We can thus depict the origin of LFS, as shown in Fig. 6 (second row). Green points depict the points with a curvature radius smaller than half of the shape diameter, and orange points depict the other points. As expected, the sphere, whose curvature radius equates to half of the shape diameter, yields a mixture of green and orange points. On the ellipsoid, the two tips correspond to a curvature-based LFS, while the rest are diameter-based. Fig. 6 offers a comprehensive visual depiction.

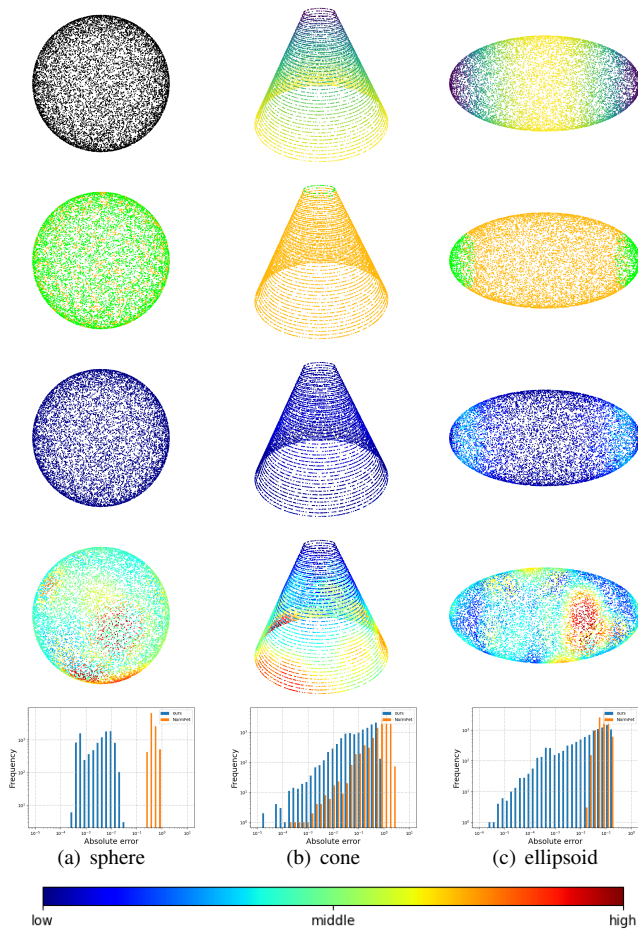


Fig. 6: Visual analysis of LFS estimation on canonical primitives (sphere, cone, and ellipsoid) sampled non-uniformly, with a little noise added. First row: ground truth LFS using the viridis color ramp. Second row: LFS type, where green denotes curvature-based and orange denotes diameter-based. Third row: estimated LFS error depicted for our algorithm using the jet color ramp. Fourth row: estimated LFS error depicted for the NormFet algorithm. Fifth row: distribution of absolute estimation errors for all input points.

2) *Free-form point clouds*: We now evaluate our LFS estimation approach for point clouds sampled on free-form shapes that lack ground-truth data. Smoothing is performed on the raw estimated LFS. Refer to Fig. 7 for visual results.

### B. Implicit Function

The implicit function is defined on the reach-aware 3D multi-domain and represented as a piecewise linear function on a 3D Delaunay triangulation. The multi-domain is used to ensure that the triangulation is dense inside the envelope where the input point set is located and sparser outside the envelope and inside the loose bounding sphere. The reach-aware property ensures that the piecewise linear function closely approximates the input point cloud. Meshing the entire bounding sphere increases the numerical stability of our solver. Fig. 8 depicts the implicit function defined on the reach-aware multi-domain.

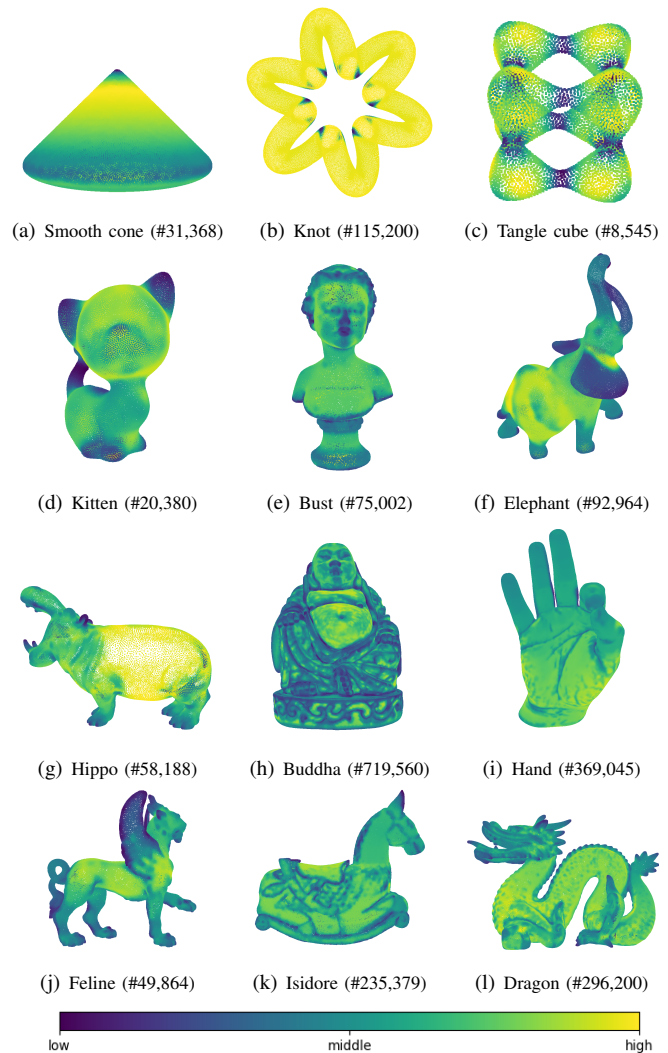


Fig. 7: LFS estimation for 3D point clouds sampled on free-form shapes. The viridis color ramp is used to map the smoothed LFS values for each point cloud.

### C. LFS-aware Meshing

The final output of our algorithm is an isotropic surface triangle mesh with a sizing function that adapts to the estimated LFS. The output mesh is generated by Delaunay refinement [4] on the zero level set of the signed robust distance function. An LFS-aware facet sizing function is designed to ensure that facets are larger where the LFS is larger and vice versa. Fig. 9 illustrates output LFS-aware meshes, where the facet sizing varies significantly, while preserving the details of the input 3D point cloud. The distribution of facet angles is reported to validate the meshes' isotropy.

We also conduct experiments to verify the impact of the user-defined parameter  $size_{max}$  on the output meshes.  $size_{max}$  controls the maximum sizing of the facets, while  $size_{min}$  is a fixed ratio to the thickness of the envelope. If  $size_{max}$  is set to equate to  $size_{min}$ , then the output mesh is uniform and isotropic. When increasing  $size_{max}$ , the variance in facet sizing increases. Fig. 10 illustrates the impact of  $size_{max}$  on the hippo 3D point cloud.

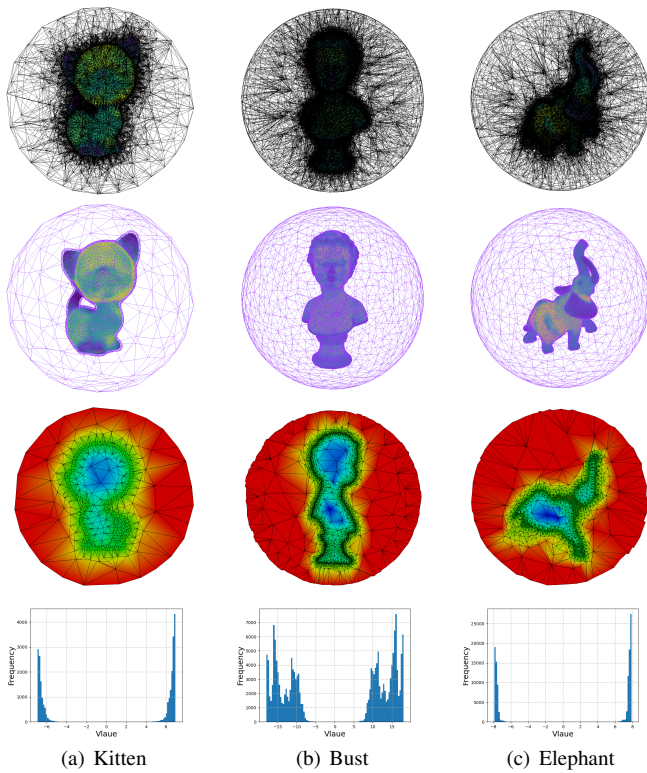


Fig. 8: Implicit functions for the kitten, bust, and elephant 3D point clouds. First row: global depiction of the multi-domain. Second row: boundary of the multi-domain. The reach-aware envelope encloses the input points embedded in a bounded sphere. Third row: cut view of the piecewise linear implicit function. The tetrahedra are well-shaped by Delaunay refinement. Fourth row: distribution of solved signed implicit values at the triangulation vertices. The two main peaks correspond to two sets of vertices located inside and outside the inferred surface. Again, we find the thin envelope layer sandwiching the inputs, and the tetrahedra are isotropic from the clip view.

#### D. Robustness

1) *Noise, outliers, and missing data:* We now evaluate the robustness of our surface reconstruction approach to noise, outliers, and missing data. We add Gaussian white noise at varying levels, ranging from 0.5% to 1.5% of the maximum edge length of the bounding box of the point cloud. We also introduce a moderate amount of random structured outliers at each noise level by generating one to three small clusters in the loose bounding sphere of the input, each containing five random points. Finally, we simulate missing data by creating two holes in the head and two holes in the body of the kitten point cloud. Our approach is robust to noise thanks to the property of the robust distance function. Our approach is also resilient to moderate amounts of random structured outliers thanks to the Delaunay refinement process [4]. The Delaunay refinement can eliminate the outliers outside the envelope. Thus, there is no impact on the implicit function solver, as all edges outside the envelope are considered to connect two vertices

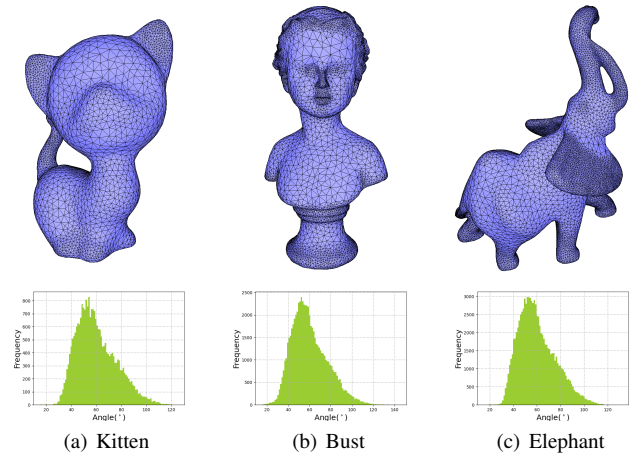


Fig. 9: Reconstructed LFS-aware meshes with distribution of facet angles. The facet sizing varies in accordance to the estimated LFS, and the output meshes are isotropic.

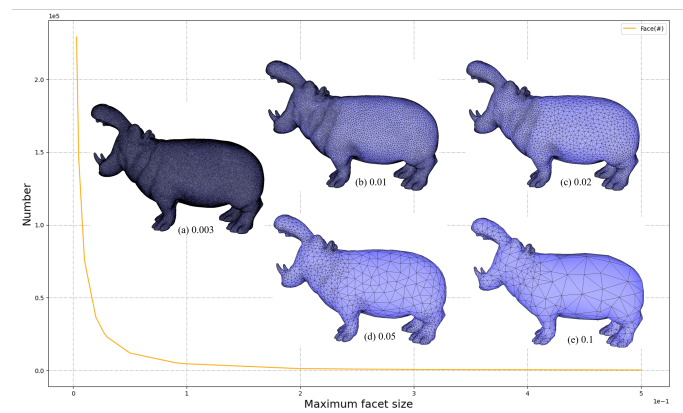


Fig. 10: Output of reconstructed meshes with different  $size_{max}$  parameters on the hippo 3D point cloud. The numbers under each hippo denote  $size_{max}$ , which determines the maximum sizing of the facets in the output mesh. The curve plots the number of output facets against  $size_{max}$ .

with similar signs. Fig. 11 illustrates the robustness of our approach on the kitten point cloud.

We now evaluate the resilience to random dense outliers of our implicit function solver by uniformly adding dense background outliers to the kitten point cloud. We generate 500 outliers from a random uniform distribution, bounded by a loose bounding box containing the kitten point cloud. The dense outliers affect the Delaunay refinement, resulting in a handful of bubbles in the multi-domain discretization. Nevertheless, thanks to the stability of our implicit function solver, the bubbles do not affect the signing with the data fitting step. The final mesh has no artifacts, as shown in Fig. 12.

2) *Large hole filling on real-scan data:* We now conduct experiments to verify the ability of our algorithm to fill large holes in real-world data acquired by laser scanners provided by the AIM@Shape dataset (see Fig. 13). Such data commonly feature non-uniform sampling and measurement noise, making hole filling challenging. On areas where the point set is dense,




Fig. 11: Evaluation of robustness on the kitten point cloud (5k points) with increasing levels of noise and number of outliers, as well as missing data. The noise level increases from 0.5% to 1.5%, and the numbers of outliers increases from 5 to 15. First row: estimated LFS. Second row: the boundary of the multi-domain ignores the outliers. Third and fourth rows: output LFS-aware meshes.

we observe a high density of edges with opposite sign guesses. Areas near holes are more ambiguous, with a mixture of edges with opposite or similar signs. Intuitively, the signing step is robust to large holes (with many missing data) thanks to the least-squares solver that propagates the sign on the triangulation vertices via consolidating sign guess hypotheses emitted for all triangulation edges.

3) *Sharp features*: Our algorithm also excels in preserving sharp features. To achieve this, we initially employ the NerVE method [68] to detect these sharp feature curves from the input point cloud, and then we convert the sharp feature curves into piecewise linear feature polylines. As we utilize the Delaunay refinement method [4] to extract the mesh from the solved implicit function, if feature polylines are provided, the Delaunay refinement method will generate protecting balls centered at these feature lines to avoid refinement on the sharp feature lines. Thus, the final mesh will preserve the sharp features. We verified the ability of sharp feature preservation on the

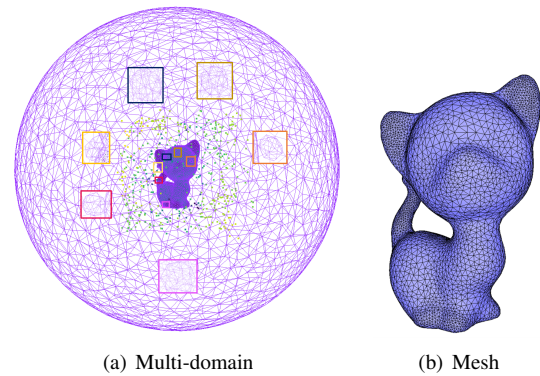


Fig. 12: Resilience to outliers evaluated on the kitten point cloud (5k points) with 500 random background outliers. A handful of bubbles are generated around some outliers inside the multi-domain discretization. A close-up of selected bubbles is delineated with a colored rectangle. Note that some bubbles are inside the point cloud, thus blocked by the front views. The bubbles do not affect the final result because the implicit function solver erases them.




Fig. 13: Reconstructions on real-world data with large holes. First row: estimated LFS. Second row: clipped view of the signed robust distance function. Third and fourth rows: LFS-aware meshes.




Fig. 14: Experiments on CAD models with sharp features. The top row shows the input point clouds with detected sharp features using NerVE [68]. The middle row is the results of RFEPS [50]. The bottom row shows the results of our algorithm.

ABC dataset [6] and compared our algorithm with RFEPS [50]. Compared to RFEPS, our method can split sharp features that are close to each other and maintain the mesh’s isotropy. Please refer to Fig. 14.

### E. Ablation Study

1) *LFS vs. curvature*: We now conduct experiments to verify the added value of using LFS instead of curvature. A first experiment is conducted on two adjacent capsules with a smaller separation distance than the curvature radius, see Fig. 15. We use three settings: (1) the estimated minimum LFS is used to construct the reach-aware envelope, and LFS is used for sizing the facets; (2) we also construct a reach-aware envelope but use only the curvature for sizing the facets; (3) we use only the minimum curvature radius to construct the envelope and the curvature to size the facets. Setting (1) yields the best results, where the curvature and topology are captured, while settings (2) and (3) fail to reconstruct the topology as the two capsules merge. More specifically, using only the minimum curvature considers neither the thickness nor the separation, thus leading to a thicker envelope and insufficient resolution for the implicit function. In addition, using only the curvature for sizing the output mesh is insufficient even when the implicit function is solved on the reach-aware domain, as the Delaunay refinement process fails to separate nearby surface sheets (see Fig. 15(f)). The middle of the two capsules is connected as the size of the facets derived from the curvature is larger than the separation distance.

### F. Comparisons

We now conduct more experiments to compare our algorithm with other baseline methods. Comparisons with RFEPS [50] on CAD models with sharp features are already presented in Fig. 14. We compare against AMLS [48], screened Poisson




Fig. 15: LFS vs. curvature. (a) The estimated LFS. (b) The estimated curvature. (c) The reach-aware multi-domain. (d) The minimum-curvature-aware multi-domain. (e) The LFS-aware mesh extracted from the implicit function solved on (c). (f) The curvature-aware mesh extracted from the implicit function solved on (c). (g) The curvature-aware mesh extracted from the implicit function solved on (d).

surface reconstruction (sPSR) [38], iPSR [46], and NKSR [25]. AMLS falls into the class of explicit reconstruction methods that also use LFS as a prior. In contrast, sPSR and iPSR fall into the category of implicit surface reconstruction, and NKSR is a learning-based method. In addition, to obtain isotropic meshes with LFS-aware sizing, we perform LFS-based remeshing using the Geogram library [71]. This involves sampling points on the reconstructed raw mesh and constructing the Voronoi diagram to estimate the LFS, which is then used to size the facets. To make a fair comparison, we ensure that all algorithms output roughly the same mesh complexity in terms of the number of faces. We then evaluate the reconstruction fidelity using the point-to-mesh Chamfer distance and the Hausdorff distance [28], [49]. Fig. 16 provides a visual comparison and Table II reports the evaluation metrics. In Fig. 16, our method may appear to produce more triangles for each model visually compared to the baselines. But this is not the case. The baselines generate extremely dense and tiny triangles in detailed regions, such as the creases of the hand model, which gives the impression of fewer triangles overall in Fig. 16. From the experiment, our approach achieves the smallest reconstruction errors for both the Chamfer and Hausdorff distance, and a visual inspection shows that it preserves more details, such as the back of the Buddha and the tail of the Isidore horse point cloud. Our approach better captures the topology by leveraging the LFS estimation, while AMLS and NKSR fail to reconstruct from the Metatron and PI point clouds. Additionally, the output of our method is water-tight, while NKSR can leave holes. We mark those holes in the red rectangles in Fig. 16.

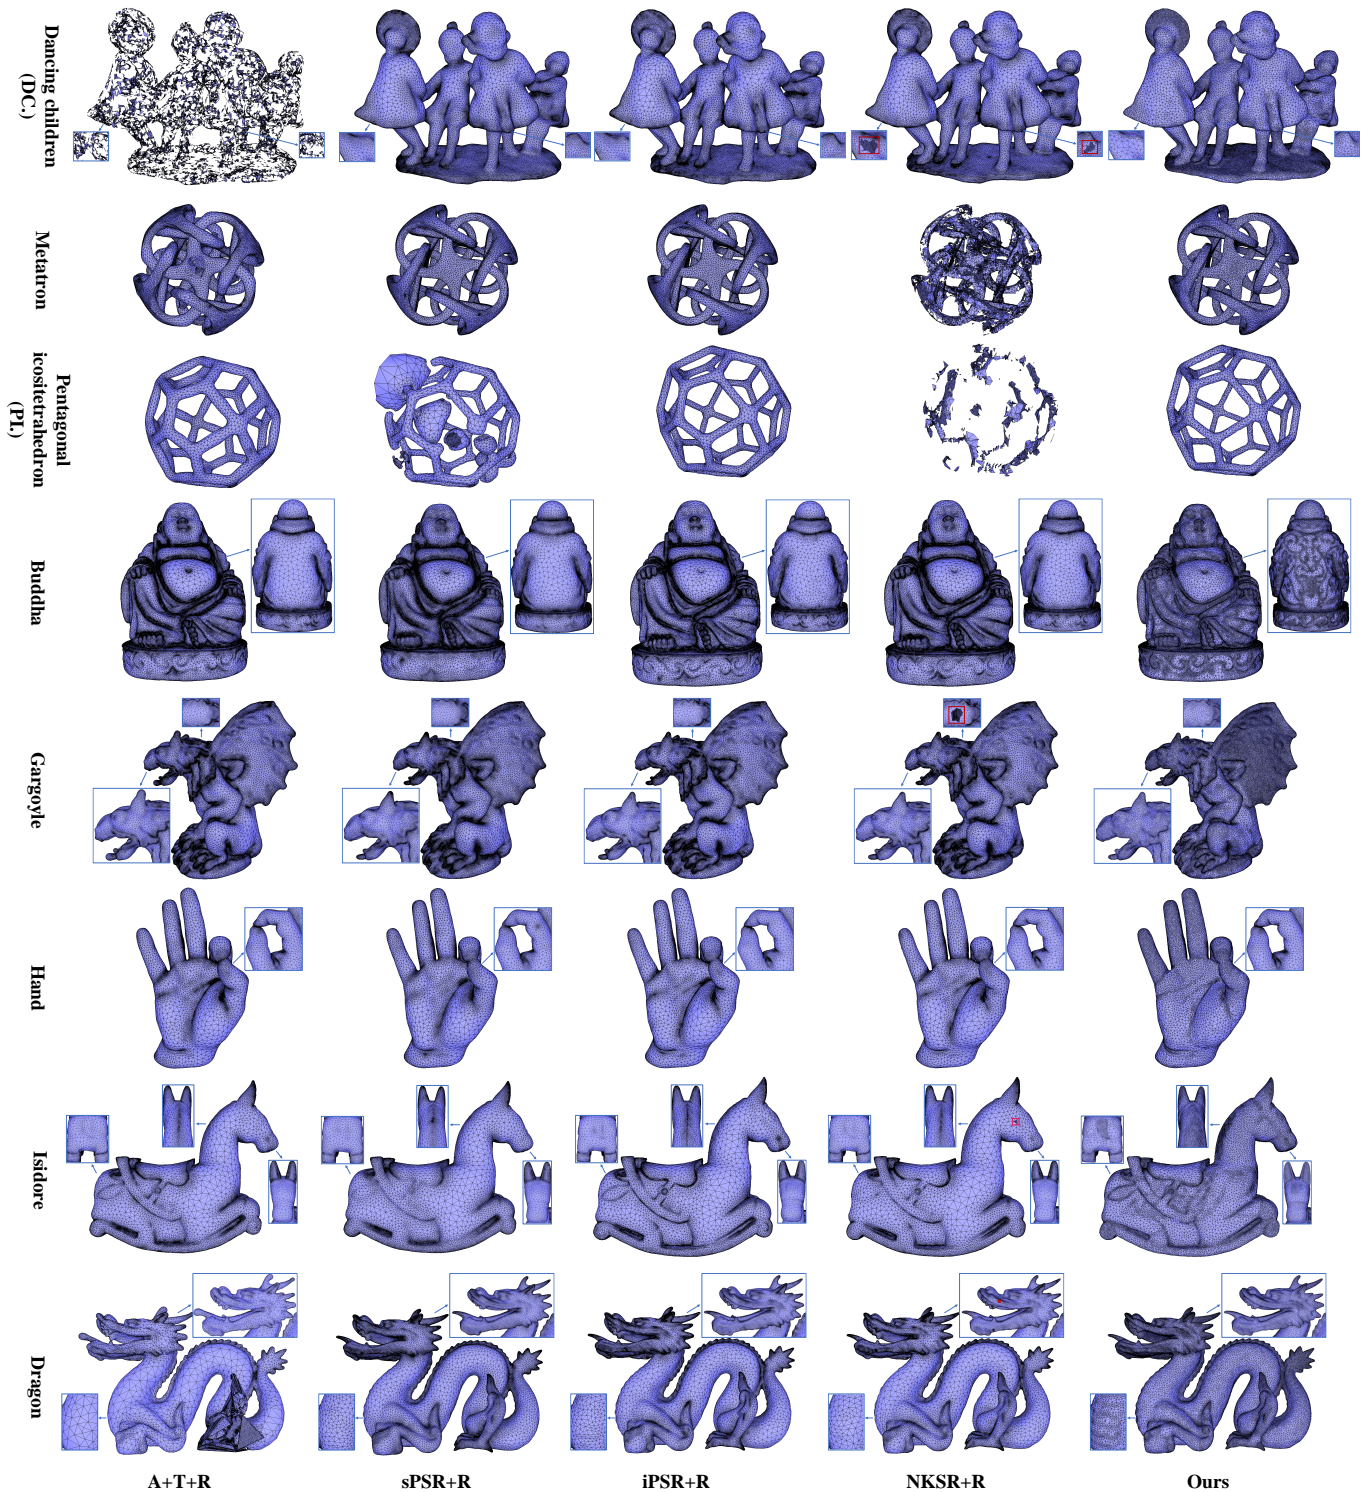


Fig. 16: Visual comparison with different algorithms: (A+T+R) AMLS reconstruction, followed by Tight Cocone and remeshing; (sPSR+R) screened Poisson surface reconstruction, followed by remeshing; (iPSR+R) iterative Poisson surface reconstruction, followed by remeshing; (NKSR+R) NKSR, followed by remeshing.

## V. CONCLUSION

This paper introduces a novel approach for estimating the local feature size (LFS) and leveraging it for LFS-aware surface reconstruction from unorganized 3D point cloud data. The primary output of our method is a valid and intersection-free

isotropic triangle surface mesh with facet sizing that captures the estimated LFS. One of the key contributions of our approach is the incorporation of LFS into the surface reconstruction process. By doing so, we generate a mesh with, ideally, “just-enough” complexity, seeking a balance between capturing fine

Model	Chamfer(L)					Hausdorff(L)					Face(#)				
	A+T+R	sPSR+R	iPSR+R	NKSR+R	Ours	A+T+R	sPSR+R	iPSR+R	NKSR+R	Ours	A+T+R	sPSR+R	iPSR+R	NKSR+R	Ours
DC.(#724k)	4.88E-1	1.37E-2	1.19E-2	1.07E-2	<b>6.10E-3</b>	3.99	1.72E-1	1.65E-1	1.48	<b>1.63E-1</b>	95k	92k	92k	92k	92k
Metatron(#69k)	4.73E-1	6.35E-2	5.72E-2	6.32E-1	<b>4.76E-2</b>	4.01	6.11E-1	5.49E-1	5.98	<b>5.40E-1</b>	61k	61k	61k	64k	61k
PL(#9k)	3.03E-1	1.79E-1	1.48E-1	1.53	<b>9.69E-2</b>	9.31E-1	1.92	7.20E-1	7.24	<b>4.02E-1</b>	15k	15k	15k	15k	15k
Buddha(#719k)	4.60E-4	3.50E-4	3.03E-4	3.13E-4	<b>2.49E-4</b>	3.81E-3	4.30E-3	3.78E-3	3.71E-3	<b>2.73E-3</b>	128k	128k	128k	128k	128k
Gargoyle(#963k)	6.12E-2	1.34E-2	1.03E-2	1.31E-2	<b>4.82E-3</b>	6.49E-1	1.89E-1	1.23E-1	2.17	<b>1.21E-1</b>	180k	180k	180k	180k	180k
Hand(#369k)	1.12E-3	2.09E-4	2.14E-4	2.28E-4	<b>2.06E-4</b>	7.08E-3	5.31E-3	5.60E-3	6.48E-3	<b>5.28E-3</b>	35k	35k	35k	35k	35k
Isidore(#235k)	4.12E-4	3.78E-4	3.62E-4	3.99E-4	<b>3.15E-4</b>	6.93E-3	7.09E-3	6.95E-3	7.18E-3	<b>6.91E-3</b>	72k	72k	72k	72k	72k
Dragon(#296k)	1.02E-3	3.05E-4	1.85E-4	3.10E-4	<b>1.60E-4</b>	3.29E-2	1.44E-2	1.37E-2	9.98E-3	<b>6.99E-3</b>	164k	158k	158k	158k	158k

TABLE II: Comparison with different algorithms: (A+T+R) AMLS reconstruction, followed by Tight Cocone and remeshing; (sPSR+R) screened Poisson surface reconstruction, followed by remeshing; (iPSR+R) iterative Poisson surface reconstruction, followed by remeshing; (NKSR+R) NKSR, followed by remeshing.

details in the data, while maintaining robustness to various data defects such as noise, outliers, and missing data. This results in a more accurate and visually pleasing reconstructed surface. Our experiments demonstrate the robustness of our LFS estimation method and show that it can handle moderate noise and non-uniform sampling effectively. Additionally, our surface reconstruction algorithm outperforms some existing baselines in terms of preserving details and accurately representing thin topological features.

#### ACKNOWLEDGEMENTS

This research is part of a project that has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 860843. We thank Prof. Tamal K. Dey for providing the official implementation of AMLS.

#### REFERENCES

[1] A. Agudo and F. Moreno-Noguer, "Shape basis interpretation for monocular deformable 3-d reconstruction," *IEEE Transactions on Multimedia*, vol. 21, no. 4, pp. 821–834, 2018.

[2] C. Lv, W. Lin, and B. Zhao, "Voxel structure-based mesh reconstruction from a 3d point cloud," *IEEE Transactions on Multimedia*, vol. 24, pp. 1815–1829, 2021.

[3] J.-K. Ahn, Y. J. Koh, and C.-S. Kim, "Efficient fine-granular scalable coding of 3d mesh sequences," *IEEE Transactions on Multimedia*, vol. 15, no. 3, pp. 485–497, 2012.

[4] C. Jamin, P. Alliez, M. Yvinec, and J.-D. Boissonnat, "CGALmesh: a generic framework for Delaunay mesh generation," *ACM Transactions on Mathematical Software (TOMS)*, vol. 41, no. 4, pp. 1–24, 2015.

[5] B. Falcidieno, "Aim@shape project presentation," in *Proceedings of Shape Modeling International*. IEEE Computer Society, 2004, pp. 329–329.

[6] S. Koch, A. Matveev, Z. Jiang, F. Williams, A. Artemov, E. Burnaev, M. Alexa, D. Zorin, and D. Panozzo, "ABC: A big CAD model dataset for geometric deep learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2019.

[7] E. Aamari, J. Kim, F. Chazal, B. Michel, A. Rinaldo, and L. Wasserman, "Estimating the reach of a manifold," *Electronic Journal of Statistics*, vol. 13, no. 1, pp. 1359–1399, 2019.

[8] A. Tagliasacchi, T. Delame, M. Spagnuolo, N. Amenta, and A. Telea, "3D skeletons: A state-of-the-art report," *Computer Graphics Forum*, vol. 35, no. 2, pp. 573–597, 2016.

[9] N. Amenta and M. Bern, "Surface reconstruction by Voronoi filtering," in *Proceedings of the 14th annual symposium on Computational geometry*, 1998, pp. 39–48.

[10] N. Amenta, S. Choi, and R. K. Kolluri, "The power crust, unions of balls, and the medial axis transform," *Computational Geometry*, vol. 19, no. 2-3, pp. 127–153, 2001.

[11] J.-D. Boissonnat and F. Cazals, "Natural neighbor coordinates of points on a surface," *Computational Geometry*, vol. 19, no. 2-3, pp. 155–173, 2001.

[12] T. K. Dey and J. Sun, "Normal and feature approximations from noisy point clouds," in *FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science*. Springer, 2006, pp. 21–32.

[13] J. Ma, S. W. Bae, and S. Choi, "3d medial axis point approximation using nearest neighbors and the normal field," *The Visual Computer*, vol. 28, pp. 7 – 19, 2011. [Online]. Available: <https://api.semanticscholar.org/CorpusID:10652045>

[14] D. Rebain, B. Angles, J. Valentin, N. Vining, J. Peethambaran, S. Izadi, and A. Tagliasacchi, "LSMAT Least squares medial axis transform," *Computer Graphics Forum*, vol. 38, no. 6, pp. 5–18, 2019.

[15] Y. Lee, J. Baek, Y. M. Kim, and F. C. Park, "IMAT: The iterative medial axis transform," *Computer Graphics Forum*, vol. 40, no. 6, pp. 162–181, 2021.

[16] L. Shapira, A. Shamir, and D. Cohen-Or, "Consistent mesh partitioning and skeletonisation using the shape diameter function," *The Visual Computer*, vol. 24, no. 4, pp. 249–259, 2008.

[17] X. Rolland-Neviere, G. Doërr, and P. Alliez, "Robust diameter-based thickness estimation of 3D objects," *Graphical Models*, vol. 75, no. 6, pp. 279–296, 2013.

[18] F. Cazals and M. Pouget, "Estimating differential quantities using polynomial fitting of osculating jets," in *Symposium on Geometry Processing*, ser. SGP '03, L. Kobbelt, P. Schroeder, and H. Hoppe, Eds. The Eurographics Association, 2003.

[19] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, "Occupancy networks: Learning 3D reconstruction in function space," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4460–4470.

[20] S. Peng, M. Niemeyer, L. Mescheder, M. Pollefeys, and A. Geiger, "Convolutional occupancy networks," in *Computer Vision—ECCV 2020*. Springer, 2020, pp. 523–540.

[21] P. Erler, P. Guerrero, S. Ohrhallinger, N. J. Mitra, and M. Wimmer, "Points2Surf: Learning implicit surfaces from point clouds," in *Computer Vision – ECCV 2020*. Springer, 2020, pp. 108–124.

[22] S. Peng, C. M. Jiang, Y. Liao, M. Niemeyer, M. Pollefeys, and A. Geiger, "Shape as points: A differentiable poisson solver," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

[23] A. Boulch and R. Marlet, "POCO: Point convolution for surface reconstruction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6302–6314.

[24] J. Huang, H. Chen, and S. Hu, "A neural Galerkin solver for accurate surface reconstruction," *ACM Transactions on Graphics (TOG)*, vol. 41, no. 6, pp. 1–16, 2022.

[25] J. Huang, Z. Gojcic, M. Atzmon, O. Litany, S. Fidler, and F. Williams, "Neural kernel surface reconstruction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 4369–4379.

[26] Z. Wang, S. Zhou, J. J. Park, D. Paschalidou, S. You, G. Wetzstein, L. Guibas, and A. Kadambi, "ALTO: Alternating latent topologies for implicit 3D reconstruction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 259–270.

[27] Y. Ben-Shabat, C. Hewa Koneputugodage, and S. Gould, "DiGS: Divergence guided shape implicit neural representation for unoriented point clouds," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 19 323–19 332.

[28] Z. Wang, P. Wang, P. Wang, Q. Dong, J. Gao, S. Chen, S. Xin, C. Tu, and W. Wang, "Neural-IMLS: Self-supervised implicit moving least-squares network for surface reconstruction," *IEEE Transactions on Visualization and Computer Graphics*, 2023.

[29] S.-L. Liu, H.-X. Guo, H. Pan, P.-S. Wang, X. Tong, and Y. Liu, "Deep implicit moving least-squares functions for 3D reconstruction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 1788–1797.

[30] J. Ye, Y. Chen, N. Wang, and X. Wang, "GIFS: Neural implicit function for general shape representation," in *Proceedings of the IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition*, 2022, pp. 12 829–12 839.
- [31] T. K. Dey and S. Goswami, "Tight Cocone: A water-tight surface reconstructor," in *Proceedings of the 8th ACM symposium on Solid modeling and applications*, 2003, pp. 127–134.
- [32] N. Amenta, S. Choi, T. K. Dey, and N. Leekha, "A simple algorithm for homeomorphic surface reconstruction," in *Proceedings of the 16th annual symposium on Computational geometry*, 2000, pp. 213–222.
- [33] D. Cohen-Steiner and F. Da, "A greedy Delaunay-based surface reconstruction algorithm," *The Visual Computer*, vol. 20, pp. 4–16, 2004.
- [34] P. Wang, Z. Wang, S. Xin, X. Gao, W. Wang, and C. Tu, "Restricted Delaunay triangulation for explicit surface reconstruction," *ACM Transactions on Graphics (TOG)*, vol. 41, no. 5, pp. 1–20, 2022.
- [35] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle, "Surface reconstruction from unorganized points," in *Proceedings of the 19th annual conference on computer graphics and interactive techniques*, 1992, pp. 71–78.
- [36] R. Kolluri, "Provably good moving least squares," *ACM Transactions on Algorithms (TALG)*, vol. 4, no. 2, pp. 1–25, 2008.
- [37] M. Kazhdan, M. Bolitho, and H. Hoppe, "Poisson surface reconstruction," in *Symposium on Geometry Processing*, ser. SGP '06, A. Sheffer and K. Polthier, Eds. The Eurographics Association, 2006.
- [38] M. Kazhdan and H. Hoppe, "Screened Poisson surface reconstruction," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 3, pp. 1–13, 2013.
- [39] G. Metzger, R. Hanocka, D. Zorin, R. Giryes, D. Panozzo, and D. Cohen-Or, "Orienting point clouds with dipole propagation," *ACM Transactions on Graphics (TOG)*, vol. 40, no. 4, pp. 1–14, 2021.
- [40] R. Xu, Z. Dou, N. Wang, S. Xin, S. Chen, M. Jiang, X. Guo, W. Wang, and C. Tu, "Globally consistent normal orientation for point clouds by regularizing the winding-number field," *ACM Transactions on Graphics (TOG)*, vol. 42, no. 4, pp. 1–15, 2023.
- [41] 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*, vol. 29, no. 5, pp. 1733–1741, 2010.
- [42] S. Giraudot, D. Cohen-Steiner, and P. Alliez, "Noise-adaptive shape reconstruction from raw point sets," *Computer Graphics Forum*, vol. 32, no. 5, pp. 229–238, 2013.
- [43] P. Alliez, D. Cohen-Steiner, Y. Tong, and M. Desbrun, "Voronoi-based Variational Reconstruction of Unoriented Point Sets," in *Symposium on Geometry Processing*, ser. SGP '07, A. Belyaev and M. Garland, Eds. The Eurographics Association, 2007.
- [44] Z. Huang, N. Carr, and T. Ju, "Variational implicit point set surfaces," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 4, pp. 1–13, 2019.
- [45] S. Lin, D. Xiao, Z. Shi, and B. Wang, "Surface reconstruction from point clouds without normals by parametrizing the Gauss formula," *ACM Transactions on Graphics (TOG)*, vol. 42, no. 2, pp. 1–19, 2022.
- [46] F. Hou, C. Wang, W. Wang, H. Qin, C. Qian, and Y. He, "Iterative Poisson surface reconstruction (iPSR) for unoriented points," *ACM Transactions on Graphics (TOG)*, vol. 41, no. 4, pp. 1–13, 2022.
- [47] D. Xiao, Z. Shi, S. Li, B. Deng, and B. Wang, "Point normal orientation and surface reconstruction by incorporating isovalue constraints to Poisson equation," *Computer Aided Geometric Design*, vol. 103, p. 102195, 2023.
- [48] T. K. Dey and J. Sun, "An adaptive MLS surface for reconstruction with guarantees," in *Symposium on Geometry Processing*, ser. SGP '05, M. Desbrun and H. Pottmann, Eds. The Eurographics Association, 2005.
- [49] T. Zhao, P. Alliez, T. Boubekeur, L. Busé, and J.-M. Thiery, "Progressive discrete domains for implicit surface reconstruction," *Computer Graphics Forum*, vol. 40, no. 5, pp. 143–156, 2021.
- [50] R. Xu, Z. Wang, Z. Dou, C. Zong, S. Xin, M. Jiang, T. Ju, and C. Tu, "RFEPS: Reconstructing feature-line equipped polygonal surface," *ACM Transactions on Graphics (TOG)*, vol. 41, no. 6, pp. 1–15, 2022.
- [51] M. Berger, A. Tagliasacchi, L. M. Seversky, P. Alliez, G. Guennebaud, J. A. Levine, A. Sharf, and C. T. Silva, "A survey of surface reconstruction from point clouds," *Computer Graphics Forum*, vol. 36, no. 1, pp. 301–329, 2017.
- [52] Z. Huang, Y. Wen, Z. Wang, J. Ren, and K. Jia, "Surface reconstruction from point clouds: A survey and a benchmark," *arXiv preprint arXiv:2205.02413*, 2022.
- [53] R. Sulzer, L. Landrieu, R. Marlet, and B. Vallet, "A survey and benchmark of automatic surface reconstruction from point clouds," 2023, working paper or preprint. [Online]. Available: <https://hal.science/hal-03968453>
- [54] W. E. Lorensen and H. E. Cline, "Marching Cubes: A high resolution 3D surface construction algorithm," in *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, vol. 21, no. 4. ACM New York, NY, USA, 1987, pp. 163–169.
- [55] A. Doi and A. Koide, "An efficient method of triangulating equivalent surfaces by using tetrahedral cells," *IEICE TRANSACTIONS on Information and Systems*, vol. 74, no. 1, pp. 214–224, 1991.
- [56] T. Ju, F. Losasso, S. Schaefer, and J. Warren, "Dual contouring of Hermite data," in *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, 2002, pp. 339–346.
- [57] G. M. Nielson, "Dual marching cubes," in *Proceedings of the IEEE visualization 2004 conference*. IEEE, 2004, pp. 489–496.
- [58] S. Schaefer and J. Warren, "Dual marching cubes: Primal contouring of dual grids," in *Proceedings of the 12th Pacific Conference on Computer Graphics and Applications*. IEEE, 2004, pp. 70–76.
- [59] T. Shen, J. Munkberg, J. Hasselgren, K. Yin, Z. Wang, W. Chen, Z. Gojic, S. Fidler, N. Sharp, and J. Gao, "Flexible isosurface extraction for gradient-based mesh optimization," *ACM Transactions on Graphics (TOG)*, vol. 42, no. 4, pp. 1–16, 2023.
- [60] Z. Chen and H. Zhang, "Neural marching cubes," *ACM Transactions on Graphics (TOG)*, vol. 40, no. 6, pp. 1–15, 2021.
- [61] Z. Chen, A. Tagliasacchi, T. Funkhouser, and H. Zhang, "Neural dual contouring," *ACM Transactions on Graphics (TOG)*, vol. 41, no. 4, pp. 1–13, 2022.
- [62] J. Hass and M. Trnkova, "Approximating isosurfaces by guaranteed-quality triangular meshes," *Computer Graphics Forum*, vol. 39, no. 5, pp. 29–40, 2020.
- [63] C. Zong, J. Zhao, P. Wang, S. Chen, S. Xin, Y. Zhou, C. Tu, and W. Wang, "A region-growing gradnormal algorithm for geometrically and topologically accurate mesh extraction," *Computer-Aided Design*, p. 103559, 2023.
- [64] P. Alliez, G. Ucelli, C. Gotsman, and M. Attene, "Recent advances in remeshing of surfaces," *Shape Analysis and Structuring*, pp. 53–82, 2008.
- [65] V. Surazhsky and C. Gotsman, "Explicit Surface Remeshing," in *Symposium on Geometry Processing*, ser. SGP '03, L. Kobbelt, P. Schroeder, and H. Hoppe, Eds. The Eurographics Association, 2003.
- [66] K. Su, N. Lei, W. Chen, L. Cui, H. Si, S. Chen, and X. Gu, "Curvature adaptive surface remeshing by sampling normal cycle," *Computer-Aided Design*, vol. 111, pp. 1–12, 2019.
- [67] V. Nivoliere, B. Lévy, and C. Geuzaine, "Anisotropic and feature sensitive triangular remeshing using normal lifting," *Journal of Computational and Applied Mathematics*, vol. 289, pp. 225–240, 2015.
- [68] X. Zhu, D. Du, W. Chen, Z. Zhao, Y. Nie, and X. Han, "NerVE: Neural volumetric edges for parametric curve extraction from point cloud," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 13 601–13 610.
- [69] CGAL, "The computational geometry algorithms library," <https://www.cgal.org/>, 2009.
- [70] G. Guennebaud, B. Jacob *et al.*, "Eigen v3," <http://eigen.tuxfamily.org>, 2010.
- [71] B. Lévy and A. Filbois, "Geogram: a library for geometric algorithms," 2015.