

Learning Manifolds in High-D Point Embedding for Anisotropic Surface Approximation from Unstructured Point Clouds

Hongbo Li¹, Haikuan Zhu¹, Xiaohu Guo², Wenping Wang³, Jing Hua¹, and
Zichun Zhong¹

¹ Department of Computer Science, Wayne State University, Detroit, USA

² Department of Computer Science, University of Texas at Dallas, USA

³ Department of Computer Science & Engineering, Texas A&M University, USA
hm9026@wayne.edu, hkzhu@wayne.edu, xguo@utdallas.edu, wenping@tamu.edu,
jinghua@wayne.edu, zichunzhong@wayne.edu

Abstract. Dense 3D sensors in various real-world fields produce point clouds that are geometrically redundant for real-time processing. In this paper, we propose an efficient and scalable learning-based anisotropic surface approximation framework, HD-PEA, that operates directly on unstructured point clouds, integrating anisotropic optimization into reconstruction to produce compact, geometry-aligned surface representations with higher fidelity, fewer elements, and improved numerical stability compared to isotropic and adaptive meshes. Firstly, we develop a novel learning-based high-dimensional (high-d) Euclidean point embedding method to map the input point clouds into a high-d manifold embedding space. For handling large-scale point clouds without retraining and fine-tuning, a patch-based meta-embedding scheme is designed during the inference stage. Then, we develop a new tangent subspace estimation for the high-d embedding manifold approximation and anisotropic manifold reconstruction in high-d space. The main contribution of this work is to propose a scalable deep learning framework and a variety of datasets for constructing a high-d Euclidean point embedding space aimed to 3D anisotropic surface mesh approximation and Riemannian curvature tensor estimation from point clouds. We extensively evaluate our method against state-of-the-art surface reconstruction approaches using several datasets, such as Thingi10K dataset, AIM@SHAPE and Stanford 3D Scanning Repository, ScanNet dataset, and further demonstrate its generalization and usability on diverse unseen shapes and applications from these datasets.

Keywords: High-d Euclidean embedding · Unstructured point clouds · Anisotropic surface mesh

1 Introduction

In robotics, autonomous driving, and AR / VR applications, dense 3D sensors produce massive point clouds that are often too computationally expensive for

real-time processing. Anisotropic surface meshes are therefore essential in geometric modeling, computer vision, and computer graphics, as they enable more accurate and efficient shape representations while significantly reducing computational cost. Unlike isotropic meshes, where the elements are as regular and uniform as possible, anisotropic surface mesh representation uses elongated triangles aligned with specific directions to optimize surface approximation, such as following the curvature tensor’s eigenvalues and eigenvectors [32, 58], leading to higher geometric accuracy with fewer elements and to improve downstream geometric processing efficiency. Different from existing anisotropic remeshing methods with the given input meshes [4, 7, 22, 26, 41, 71, 72], in this paper, we address a more challenging cross-modality problem in approximating the anisotropic surface mesh directly from unstructured 3D point clouds. The fundamental motivation of this work is to present a new manifold-guided surface reconstruction framework to the synergic integration of anisotropic meshing and surface reconstruction for effective 3D shape representation.

With the advancement of neural representations, learning implicit functions, such as signed distance functions (SDFs), from point clouds has become a prevalent approach to surface reconstruction. To achieve high reconstruction accuracy, they often rely on high-resolution Marching Cubes (MC) results [44], which generate densely sampled meshes with a large number of vertices and faces. Recent neural meshing methods, such as NMC [15] and NDC [14], enhance Marching Cubes and Dual Contouring [36] by learning vertex placements. VoroMesh [46] refines vertex positions using Voronoi diagrams, while PoNQ [47] applies quadric error metrics for optimization. However, most of these approaches still depend on voxel-based representations, and they often fail to capture smooth surfaces and geometric features accurately due to the resolution constraint. More recently, LMR [68] proposes an adaptive meshing method that uses implicit surfaces and surface curvature to generate resolution-adaptive and lightweight meshes. However, it does not address anisotropic mesh representation from point clouds, so that there is still a gap between their adaptive meshing and the L_2 optimal approximation to a smooth surface, i.e., the anisotropy of triangles conforms to the eigenvalues and eigenvectors of the curvature tensors [32, 58].

This work addresses the above challenges in anisotropic surface mesh approximation from unstructured point clouds. The *main contributions* are:

- Design a new computational framework to reconstruct high-quality anisotropic surface meshes from unstructured point clouds;
- Develop a scalable neural high-d Euclidean point embedding that maps large-scale point clouds onto a high-d manifold while preserving intrinsic Riemannian metrics, overcoming the mesh-dependent limitation of prior anisotropy losses, which prevents their direct application to point clouds;
- Develop a new tangent subspace approximation for the high-d embedding manifold to construct uniformly sparse samplings and reconstruct anisotropic meshes in the high-d space;
- Construct and release a variety of datasets for anisotropic surface mesh reconstruction and curvature tensor estimation along with the input point clouds.




Fig. 1: The overview pipeline of the proposed neural high-d point embedding for anisotropic surface approximation (HD-PEA). Our method includes three main components: neural high-d Euclidean point embedding, high-d embedding manifold approximation by uniformly sparse sampling, and anisotropic manifold mesh reconstruction.

The overview pipeline of the proposed neural high-d point embedding for anisotropic surface approximation (HD-PEA) is shown in Fig. 1.

2 Related Work

Neural Surface Meshing and Reconstruction. In recent years, the deep learning-based approaches for meshing of shapes become popular in computer vision and geometric modeling. For instances, Point2Mesh [31], Pixel2Mesh [64], Pixel2Mesh++ [65] reconstruct meshes by deforming a template mesh to fit the input point cloud, but this typically restricts the output to the template’s fixed topology. To address this, PointTriNet [57] directly predicts mesh connectivity by iteratively proposing and classifying triangle candidates. Although the method is local and differentiable, it often produces meshes with holes. NeuralMeshing [63] proposes differentiable meshing algorithm for extracting surface meshes from neural implicit representations, but it cannot guarantee watertightness and feature preservation. NMC [15] and NDC [14] improve the meshing templates of Marching Cubes (MC) [44] and Dual Contouring (DC) [36] by incorporating learned vertex predictions. VoroMesh [46] further refines vertex placement using Voronoi diagrams, and PoNQ [47] leverages a quadric error metric for optimization. Despite these advancements, most existing neural meshing techniques still rely on voxel-based representations, which constrain their accuracy and expressiveness, especially on low-resolution grids. To address this issue, most recently LMR [68] proposes an adaptive Delaunay meshing that uses implicit surfaces and surface curvatures, but this method is not generalizable and efficient, as it requires optimizing each input point cloud individually. In summary, the existing methods do not address anisotropic mesh representation and how to extend them into an arbitrary dimensional space is unexplored.

High-Dimensional Geometric Embeddings. Several approaches have been proposed to address specific classes of surface meshing and processing problems by embedding them in higher-dimensional spaces [5, 11, 18, 19, 38, 40], compared with parameterization-based methods [12, 51] often struggling with complicated topologies due to their reliance on low-dimensional canonical domains. Their core idea is to transform an anisotropic meshing problem defined on a 3D surface into an equivalent isotropic meshing problem in 6D or a higher dimension. In particular, SIFHDE² [72] introduces an optimization-based approach to compute

the high-d embedding space. NASM [41] uses dot products between edge vectors on mesh faces to model anisotropy to learn the high-d embedding. However, since point clouds lack explicit tangent spaces and vector representations, such dot product-based losses cannot be directly applied. Moreover, the above methods are designed for mesh-based representations and cannot be directly applied to point clouds. Fundamentally, extending the Nash Embedding Theorem [49, 50] to point clouds remains challenging, as existing formulations assume continuous manifolds and therefore cannot operate on discrete, unstructured point sets.

Anisotropic Triangular Meshing. Delaunay triangulation has been extended to anisotropic settings using two main approaches: refinement-based methods [7, 8, 21] and variational methods [13]. [4] employs a progressive vertex insertion strategy for anisotropic Delaunay refinement. [55] approximates Riemannian Voronoi diagrams on surfaces using geodesic distance computations. [26] proposes Locally Convex Triangulation (LCT), which aligns convex functions with Riemannian metrics for anisotropic meshing. [10] lifts points to convex functions to minimize reconstruction error, incorporating anisotropy, but their method is limited to anisotropies representable via convex functions. [17] introduces a method for high-quality triangular meshing with prescribed anisotropy using metric-adapted embeddings based on cone singularities. Another group of anisotropic triangular meshing methods are based on computing the dual mesh from the anisotropic centroidal Voronoi tessellation (ACVT), such as [22, 62, 71]. All of these methods are designed for mesh-based representations and computations. Recently, [70] develops a unified particle-based formulation for resamplings with specific patterns from original point clouds. However, it is not effective and scalable for a large-scale number of models, since it is a traditional model-based method and it also needs a curvature tensor for each model as input.

3 Method

3.1 Neural High-D Euclidean Point Embedding

The Nash embedding theorem states that every smooth Riemannian manifold can be isometrically embedded into a high-d Euclidean space [49, 50]. However, applying Nash Embedding Theorem to point clouds is challenging, as existing methods (e.g., SIFHDE² [72], NASM [41]) require continuous manifolds and cannot operate on unstructured points. We address this by introducing a barycentric-invariant isometric propagation that generates high-d embeddings from point clouds, by enforcing barycentric consistency, our method strictly satisfies the continuity and differentiability requirements of the embedding theorem, bridging Nash embedding to unstructured point clouds for the first time rather than continuous manifolds or mesh domains required by prior high-d embedding formulations. Motivated by this foundation, we propose a method to learn a high-d point embedding that preserves the intrinsic Riemannian geometry, given only its point cloud and normal vectors. This high-d point embedding is designed to approximate an isometric mapping of the underlying manifold which is the input for subsequent high-d embedding manifold approximation.

Data Generation for High-D Point Embedding. Existing mesh-based approaches [17, 41, 72] in geometry processing that consider surface anisotropy primarily operate on graph-like structures, where edges and connectivity are explicitly defined. These methods rely on well-structured meshes to model anisotropic behavior. In contrast, there is no existing method capable of generating ground truth high-d embeddings for unstructured point sets, particularly in a way that faithfully encodes manifold’s local geometry (e.g., curvatures, tangent space, principal directions).

To address this gap, our data generation for point cloud embedding leverages the mapping from the original 3D Riemannian space to the target high-d Euclidean space. It can be described by an affine transformation [52, 60, 72], where the transformation matrix inherently captures local manifold curvature and directional information. Since barycentric coordinates are preserved under affine transformations, we first sample points on the original 3D triangle surfaces and compute the barycentric coordinates of each sampled point with respect to its triangle. Using these same barycentric weights, we interpolate the corresponding high-d vertex embeddings to obtain the high-d coordinates of each sampled point. This process constructs a dense and high-quality ground truth point cloud in the target high-d manifold space, where it is both geometrically consistent and sensitive to local anisotropy. It provides a complete ground truth mapping from the 3D surface to its high-d embedding for supervised learning. Details are provided as follows.

Formally, let the original 3D triangle mesh denotes as $\mathcal{M} = (\mathcal{V}, \mathcal{F})$, where $\mathcal{V} = \{\mathbf{v}_i \in \mathbb{R}^3\}$ are the vertex positions and $\mathcal{F}$ is the set of triangle faces. The corresponding high-d embedded mesh is denoted as $\bar{\mathcal{M}} = (\bar{\mathcal{V}}, \bar{\mathcal{F}})$, where each vertex $\bar{\mathbf{v}}_i \in \mathbb{R}^d$ is the high-d embedding of $\mathbf{v}_i$ and $d = 8$ (as suggested in [72]), $\bar{\mathcal{F}} = \bar{\mathcal{F}}$. As shown in Fig. 2, for any point $\mathbf{p}$ sampled in a triangle face $\mathbf{f} = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$, we compute its barycentric coordinates (w_1, w_2, w_3) as:




Fig. 2: The procedure of data generation for high-d point embedding. The triangles represent the underlying mesh used in the invariant weight transfer.

$$\mathbf{p} = w_1 \mathbf{v}_1 + w_2 \mathbf{v}_2 + w_3 \mathbf{v}_3, \quad (1)$$

$$\text{where } w_1 + w_2 + w_3 = 1, w_i \geq 0. \quad (2)$$

The corresponding high-d embedding of $\mathbf{p}$, denoted $\bar{\mathbf{p}} \in \mathbb{R}^d$, is obtained by applying the invariant weight transfer (i.e., same barycentric weights) to the high-d vertices:

$$\bar{\mathbf{p}} = w_1 \bar{\mathbf{v}}_1 + w_2 \bar{\mathbf{v}}_2 + w_3 \bar{\mathbf{v}}_3. \quad (3)$$

By repeating this procedure across all triangles in $\mathcal{M}$, we obtain a dense set of point pairs $(\mathbf{p}, \bar{\mathbf{p}})$, where $\mathbf{p} \in \mathbb{R}^3$ and $\bar{\mathbf{p}} \in \mathbb{R}^d$. This provides a complete ground

truth mapping from the 3D surface to its high-d Euclidean embedding, which can be used for our follow-up high-d point embedding learning task.

While the generated embeddings may partially inherit biases from the source meshes, our framework uses the mesh only during this data generation stage to maintain consistency with an underlying smooth Riemannian manifold. Unlike two-stage pipelines that suffer from intermediate reconstruction artifacts (discussed in Section 4.2), our method learns the anisotropic metric field directly from raw point clouds in the high-d space.

High-D Point Embedding Network and Loss. Formally, let $\mathbf{P} \in \mathbb{R}^{N \times 3}$ denote a point cloud consisting of N points in 3D space, and let $\mathbf{n} \in \mathbb{R}^{N \times 3}$ represent the corresponding unit normal vectors. The input to our network is the concatenated representation $[\mathbf{P}, \mathbf{n}] \in \mathbb{R}^6$. A point-based transformer backbone $\mathcal{F}$, parameterized by $\theta_{\mathcal{F}}$, predicts the high-d coordinate extension $\tilde{\mathbf{P}} \in \mathbb{R}^{N \times (d-3)}$. The final high-d embedding is obtained by concatenating the original 3D coordinates with the predicted extension: $\bar{\mathbf{P}} = [\mathbf{P}, \mathcal{F}([\mathbf{P}, \mathbf{n}]; \theta_{\mathcal{F}})]$. We apply Point Transformer V3 [67] as our backbone architecture. To accommodate the required output dimensionality, we append a single linear layer to the end of the network.

To capture the local geometric structure of the underlying manifold from the input point clouds, we utilize the pairwise distances between neighboring points. Different from the prior work such as [41], which models anisotropy using dot products between edge vectors defined on the given mesh faces, point clouds lack explicit tangent spaces and well-defined vector representations. These challenges make utilization of the dot product-based loss infeasible in our setting. Instead, we approximate the local geometry by pairwise distances, which implicitly encode the intrinsic Riemannian metric governing local lengths and angles.

In prior works [33, 43, 45], L_1 -based losses tend to bias predictions toward the median of the target distribution, making them more effective for learning from diverse and non-uniform geometric structures. Building on this insight, we adopt Mean Absolute Error (MAE) between predicted and ground truth pairwise distances. This formulation promotes the preservation of relative geometric structure across diverse shapes and supports more accurate high-d embedding prediction. Our neighborhood distance-based loss is:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \sum_{j \in \mathcal{N}(i)} | \|f(\mathbf{p}_i) - f(\mathbf{p}_j)\| - \|\bar{\mathbf{p}}_i - \bar{\mathbf{p}}_j\| |, \quad (4)$$

where $f(\mathbf{p}_i)$ denotes the predicted embedding coordinates of point $\mathbf{p}_i$, $\bar{\mathbf{p}}_i$ denotes the corresponding ground truth embedding coordinates. $\mathcal{N}(i)$ represents the set of neighboring points of $\mathbf{p}_i$. In our experiment, $\mathcal{N}(i)$ is set to be 40 and the embedding dimension is 8 as suggested in [72]. This formulation encourages local geometric consistency between the prediction and the ground truth without relying on explicit edge connectivity. We demonstrate the effectiveness of our loss function in the ablation study.

Patch-Based Meta-Embedding Inference. Since the embedding network is trained on point clouds with a fixed number of points, applying it directly to larger point clouds can lead to degraded embeddings. To address this, we in-

introduce patch-based meta-embedding inference (PMEI), i.e., an inference-time strategy that divides large point clouds into overlapping patches, processes each independently, and aligns the individual patch-based embeddings into a consistent global meta-embedding. It can efficiently handle arbitrarily large-scale point clouds without retraining while preserving embedding quality and consistency.

For a point cloud $\mathbf{P} \in \mathbb{R}^{N \times 3}$, we first construct a set of patch centers $\{\mathbf{c}_1, \dots, \mathbf{c}_l\}$ where $l \ll N$ using farthest point sampling (FPS). For each $\mathbf{c}_i \in \{\mathbf{c}_1, \dots, \mathbf{c}_l\}$, we use its k nearest neighbors to form a local patch $\mathbf{q}_i \in \mathbb{R}^{k \times 3}$, and the k we choose guarantees the overlapping of the patches. Each patch $\mathbf{q}_i$ is then processed by the network to obtain its high-d Euclidean point embedding $\bar{\mathbf{q}}_i \in \mathbb{R}^d$. Specifically, for two overlapping patches $\bar{\mathbf{q}}_i$ and $\bar{\mathbf{q}}_j$ with shared domain $\mathcal{O}_{ij}$, the ideal consistency condition $\bar{\mathbf{q}}_j|_{\mathcal{O}_{ij}} = \bar{\mathbf{q}}_i|_{\mathcal{O}_{ij}}$ is typically not satisfied. This inconsistency arises because standard attention mechanisms are not equivariant to affine transformations unless explicitly designed as [27]. Let $\{\mathbf{p}_i^k, \mathbf{p}_j^k\}_{k=1}^m$ denotes the set of m corresponding point pairs between the overlapping regions $\bar{\mathbf{q}}_i|_{\mathcal{O}_{ij}}$ and $\bar{\mathbf{q}}_j|_{\mathcal{O}_{ij}}$. We formulate the patch-based embedding alignment as an optimization problem:

$$T_{ij}^* = \arg \min_{T_{ij}} \sum_{k=1}^m \|\mathbf{p}_i^k - T(\mathbf{p}_j^k)\|^2, \quad (5)$$

where $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is an affine transformation defined as $T(x) = \mathbf{A}x + \mathbf{b}$ with $\mathbf{A} \in \mathbb{R}^{d \times d}$ and $\mathbf{b} \in \mathbb{R}^d$.

The embedding alignment process is formulated as a minimum spanning tree (MST) problem over the set of patch centers $\{\mathbf{c}_i\}_{i=1}^l$. Each center $\mathbf{c}_i$ is treated as a node, and the edge weight between two centers $\mathbf{c}_i$ and $\mathbf{c}_j$ is given by their Euclidean distance $\|\mathbf{c}_i - \mathbf{c}_j\|$. Starting from an arbitrary root center, the algorithm iteratively aligns the next nearest patch using Eq. (5) and adds to the already aligned set, using a priority queue keyed by the inter-center distances-equivalent to Prim’s algorithm for constructing an MST. This process ensures that the newly selected patches potentially have the largest overlapping region with the already aligned set, promoting more stable and consistent transformations. The PME algorithm and illustration are provided in Section 8 of Appendix.

3.2 High-D Embedding Manifold Approximation

After obtaining the high-d point embedding, our goal is to obtain a piecewise linear approximation of the manifold $\mathcal{M}$ that follows a uniform sparse distribution in the high-d space. When this approximation is projected to original 3D manifold space, it naturally results in an anisotropic linear approximation. In this section, we present an approach for optimizing the vertex positions to achieve a uniform distribution in high-d point embedding space.

Embedding Tangent Space Approximation. To operate on a point-based manifold in a high-d space, it is essential to have a guiding structure that constrains operations to lie on the manifold itself. Different from normals are used in 3D space [6, 70] to define local orientation, we use the tangent subspace of the manifold to impose such constraints. The main reason is that if the surface

is not a hypersurface (e.g., a 2D surface embedded in $\mathbb{R}^8$ in our work), the normal space becomes a subspace of dimension d , and there is no unique normal vector; instead, a normal subspace exists [56, 69]. More generally, for an m -dimensional manifold $\mathcal{M}$ embedded in $\mathbb{R}^d$, the tangent space $\mathcal{T}_{\bar{\mathbf{p}}}\mathcal{M}$ at a point $\bar{\mathbf{p}} \in \mathcal{M}$ is a m -dimensional linear subspace of $\mathbb{R}^d$. It is spanned by linearly independent m tangent vectors, which form the basis of $\mathcal{T}_{\bar{\mathbf{p}}}\mathcal{M}$. In our case, where the manifold is a 2D surface embedded in a high-d space, we estimate the local tangent space at each point $\bar{\mathbf{p}} \in \mathbb{R}^d$ using Principal Component Analysis (PCA) [53]. To accurately capture the intrinsic low-dimensional structure, PCA is applied to the k -nearest neighbors $\mathcal{N}_{\bar{\mathbf{p}}} = \{\bar{\mathbf{p}}_1, \dots, \bar{\mathbf{p}}_k\}$.

Let $\{\bar{\mathbf{p}}_i\}_{i=1}^k \subset \mathbb{R}^d$ denote the k -nearest neighbors of a center point $\bar{\mathbf{p}} \in \mathbb{R}^d$. The local centroid of these neighbors is given by $\hat{\mathbf{p}} = \frac{1}{k} \sum_{i=1}^k \bar{\mathbf{p}}_i$. The local covariance matrix at $\bar{\mathbf{p}}$ is defined as $\mathbf{C}_{\bar{\mathbf{p}}} = \frac{1}{k} \sum_{i=1}^k (\bar{\mathbf{p}}_i - \hat{\mathbf{p}})(\bar{\mathbf{p}}_i - \hat{\mathbf{p}})^T$. The tangent basis at $\bar{\mathbf{p}}$ is then obtained by performing eigen decomposition of $\mathbf{C}_{\bar{\mathbf{p}}}$. For the target dimension m of tangent subspace, the top m eigenvectors $\bar{\mathbf{v}}$ of $\mathbf{C}_{\bar{\mathbf{p}}}$ corresponding to its largest m eigenvalues, form an orthonormal basis: $\bar{\mathbf{V}} = [\bar{\mathbf{v}}_1, \bar{\mathbf{v}}_2, \dots, \bar{\mathbf{v}}_d] \in \mathbb{R}^{d \times m}$ for the estimated tangent space $\mathcal{T}_{\bar{\mathbf{p}}}\mathcal{M} \subset \mathbb{R}^{d \times m}$. In our setting, we use $d = 8$ and $m = 2$. Fig. 3 illustrates the tangent space estimation in the high-d space.




Fig. 3: The tangent space estimation in high-d space.

Sampling Uniformity on Manifold. Given a set of sparse seed points lying on the manifold of the target domain, our goal is to achieve an isotropic distribution over this manifold. Particle-based methods have been proven to be effective in low-dimensional settings (e.g., 2D or 3D) [66, 70, 71] and have also been extended to high-d manifolds when an explicit mesh representation is available [72]. However, their application to manifolds represented purely by point clouds in high-d spaces remains unexplored. To address this gap, we leverage the introduced tangent space approximation to constrain particle-based optimization directly to the local geometry of point-based high-d manifolds. More details about particle formulation and computations are discussed in Section 1 and Section 2 of Appendix.

It is important to note that both energy and forces are defined in terms of Euclidean distances between particles. However, the desired optimization domain is the target embedding manifold $\mathcal{M}$, rather than the entire ambient space $\mathbb{R}^d$. To constrain the optimization within the manifold, we use the tangent basis $\bar{\mathbf{V}}$ to project each particle back onto a local linear approximation of the embedding manifold. Specifically, for a given particle $\bar{\mathbf{x}}$, we first identify its nearest neighbors $\bar{\mathbf{x}}_p \in \mathcal{M}$ in the original 3D space, and perform the projection onto the tangent space $\mathcal{T}_{\bar{\mathbf{x}}_p}\mathcal{M}$ at $\bar{\mathbf{x}}_p$ as: $\tilde{\mathbf{x}}_{\mathcal{T}_{\bar{\mathbf{x}}_p}\mathcal{M}} = \bar{\mathbf{x}}_p + \bar{\mathbf{V}}\bar{\mathbf{V}}^T(\bar{\mathbf{x}} - \bar{\mathbf{x}}_p)$, where $\bar{\mathbf{V}} \in \mathbb{R}^{d \times m}$ is the local orthonormal basis spanning the estimated tangent space $\mathcal{T}_{\bar{\mathbf{x}}_p}\mathcal{M}$. This projection ensures that particle's coordinate updates remain locally constrained to the manifold geometry in high-d space.

By minimizing the total energy subject to the constraint that all particles remain on the manifold $\mathcal{M} \subset \mathbb{R}^m$, the system evolves toward a state of equilibrium where the forces acting on all particles are balanced. At equilibrium, this process yields a uniformly distributed and isotropic sampling over the target manifold.

3.3 Anisotropic Manifold Reconstruction

The next step is to compute a piecewise linear approximation of this manifold $\mathcal{M}$ by using the output seeds / particles $\bar{\mathbf{x}}$ from the previous optimization, which is a set of well-distributed sparse seeds in $\mathbb{R}^d$ that lies on a manifold $\mathcal{M}$ of intrinsic dimension m , where $m < d$. In our case, we assume that $\mathcal{M}$ is a smooth manifold with intrinsic dimension $m = 2$. Inspired by [6], we adopt the concept of Tangential Delaunay Complex [2, 3, 24] to generalize the restricted Voronoi cells (RVCs) from 3D to a higher dimension, constructing our high-d tangential complex. Our method primarily operates within the m -dimensional tangent spaces at sampling points. Therefore, the computational complexity depends on the manifold dimension m rather than high dimension d in the ambient space.

Manifold Tangent Space Approximation. To reconstruct a piecewise linear approximation of the manifold in a higher dimension, we begin by forming a m -dimensional unit disk with a radius r and uniformly discretizing into n disks, where each radius is defined as $r = \alpha \cdot l_{\text{diag}}$, where l_{diag} denotes the diagonal length of the shape’s bounding box in high-d ambient space, and α is a user-defined scaling coefficient. Then, we estimate the tangent basis $\bar{\mathbf{V}}$ at each point $\bar{\mathbf{x}}_i \in \mathbb{R}^d$, using its nearest neighbors identified in the same space, and define an affine transformation $\mathbf{A}_{\bar{\mathbf{x}}} : \mathbb{R}^m \rightarrow \mathbb{R}^d$ that maps the intrinsic disk $\mathbf{D}^m \subset \mathbb{R}^m$ into the ambient space with dimension d . This mapping is: $\mathbf{D}^d(\bar{\mathbf{x}}) = \bar{\mathbf{V}}\mathbf{D}^m + \bar{\mathbf{x}}$, where $\mathbf{D}^d \in \mathbb{R}^d$ is an image of the disk in the ambient space centered at $\bar{\mathbf{x}}$. It forms a local linear approximation of the manifold around $\bar{\mathbf{x}}$.

High-D RVC and Mesh Extraction. Once the tangent space is constructed at each seed, the restricted Voronoi cells (RVCs) forming the tangential complex are computed by clipping the embedded disk $\mathbf{D}^d$ with the bisectors between the seeds and their neighbors. An example is illustrated in Fig. 4. Since our target dimension is higher than 3, we apply re-entrant clipping method [61]. Details are provided in Section 3 of Appendix.




Fig. 4: Illustration of a high-d RVC.

As in [6, 20], the result of the previous clipping step is a triangle soup. We adopt a manifold extraction method following [6]. Finally, we apply the post-processing step provided by Geogram [39] to fill any remaining holes from the previous step if necessary. Details are provided in Section 4 of Appendix. When the extracted mesh approximation is projected (via truncation) to original 3D manifold space, it naturally results in an anisotropic mesh approximation.

4 Experiment

We evaluate the effectiveness of our method in approximating surfaces by reconstructing high-fidelity lightweight anisotropic meshes and its applications from

point clouds across a wide variety of shapes. Due to the page limit, the implementation details are provided in Section 5 of Appendix and some additional results are provided in Section 7 of Appendix. Detailed implementations of other comparison methods in the following section can be referred from their original papers. *The source code and data will be publicly released after acceptance.*

4.1 Datasets and Metrics

To thoroughly evaluate our method, we employ multiple datasets and applications that target different aspects of generalization and robustness.

Training Dataset. We apply FPS to sample 40K points from each of the 240 selected models in the Thingi10K dataset [73]. To mitigate biases arising from geometric variations in the training set, we apply data augmentation strategy followed by NASM [41]. After applying this data augmentation strategy, the training set comprises 2,400 models of point clouds. The sampled point clouds are then normalized and estimated oriented normals using the Winding Number Normal Consistency (WNNC) method [42]. The ground-truth high-d point embeddings are generated following Section 3.1.

Testing Datasets. We conduct experiments on several challenging datasets, including Thingi10k dataset [73], Myles et al.’s dataset [48] (i.e., 3D shapes from AIM@Shape and Stanford shape repository), TetWeave dataset [1], and ScanNet dataset [16]. Thingi10k dataset serves as a baseline for evaluating our method. Others are unseen datasets, which are only used in the testing stage. All evaluations are conducted without any fine-tuning or retraining. As for point clouds acquisition, we utilize randomly sampling from surface meshes, synthetic scanning [29, 35], noisy and irregular sampling, real scanning, etc. More details are provided in Section 6 of Appendix.

Metrics. Besides the qualitative evaluation and comparison, we provide quantitative evaluation. We adopt standard metrics for *reconstructed surface accuracy*, including Chamfer distance (CD), F-Score (F1), normal consistency (NC), and Hausdorff distance (HD). To measure the *anisotropic mesh quality* G , for each triangle $\triangle_{abc}$ in the final mesh, we use its approximated metric $\mathbf{Q}(\triangle_{abc}) = (\mathbf{Q}(\mathbf{x}_a) + \mathbf{Q}(\mathbf{x}_b) + \mathbf{Q}(\mathbf{x}_c))/3$, where $\mathbf{Q}(\cdot) = \sqrt{\mathbf{M}(\cdot)}$, to affine-transform it from the original anisotropic space into the Euclidean space. $\mathbf{M}(\cdot)$ is the curvature metric tensor. After that, we employ the isotropic triangular criteria [25], to evaluate the quality of generated anisotropic mesh as [41, 71]. The quality of a triangle is measured by $G_{\triangle_{abc}} = 2\sqrt{3}\frac{S}{ph}$, where S is the triangle area, p is its half-perimeter, and h is the length of its longest edge. G is the average qualities of all triangles. We report the *average runtime* of all the methods.

4.2 Evaluations

Baseline Sampling. We first evaluate our HD-PEA method on the testing set with 80 point cloud models selected from Thingi10K as shown in Fig. 5, and compare the results to several state-of-the-art in surface mesh reconstruction, including Poisson [37], NDC [14], POCO [9], PoNQ [47], and LMR [68] in Table 1. We report both quantitative and qualitative results to demonstrate that

HD-PEA achieves the best CD, F1, NC, and HD values. It justifies that incorporating anisotropic surface meshing strategy not only better captures geometric detail but also leads to more accurate surface approximations with fewer mesh elements. Compared with other deep learning-based methods, our approach requires only a few seconds for the entire computation.



Fig. 5: Our anisotropic surface mesh approximation from input point clouds. The models are selected from Thingi10K dataset [73]. Left to right: input point clouds, reconstructed anisotropic meshes with zoom-in views, and estimated curvature tensors with stretching ratios for input point clouds in colors.




Fig. 6: Our anisotropic surface mesh approximation results comparison with the state-of-the-art in surface mesh reconstruction. The models are selected from Myles et al.’s dataset [48]. NDC, POCO, and Poisson use grid size of 128^3 .

For Marching Cubes-like methods, such as Poisson, NDC, and POCO, often exhibit undersampling artifacts such as zigzag patterns, due to their inflexible sampling directions in Fig. 6. Although PoNQ employs QEM [28] for surface approximation, it still relies on Delaunay triangulation, which typically produces more vertices than desired. LMR (one of the most recent adaptive mesh reconstruction methods) conforms to surface curvatures by using adaptive meshes; however, it does not adequately control stretching ratios for anisotropic triangles, leading to noticeable number of degenerate triangles. Another limitation lies in the reliance of LMR on an optimization-based process that is computationally expensive, thereby limiting its practicality. The detailed runtime is shown in Table 1.

Synthetic Scanning. To further justify capability of our method, this experiment evaluates our method on synthetic scanned point clouds without any fine-tuning or retraining. We evaluate our HD-PEA method on 75 synthetic scanned point clouds generated from Myles et al.’s dataset using Blensor [29] as in [35]. For each model, we simulate realistic scanning on point clouds containing 40K points. To provide a more comprehensive evaluation under these realistic scanning conditions, we expand our baselines from the previous subsection to include three additional state-of-the-art methods: NKSR [34], SAP [54], and SIREN [59]. Fig. 6 and Table 2 show that our method consistently achieves the best qualitative and quantitative performance with lightweight mesh elements.

Table 1: Quantitative comparison with our HD-PEA and state-of-the-art surface reconstructions on 80 models selected from Thingi10K dataset [73]. The best results are highlighted in bold, and the second best results are underlined. Note: CD ($\times 10^{-5}$), HD ($\times 10^{-2}$), and Time (seconds). All metrics are averaged over 80 models.

Method	# V_{out}	# f_{out}	CD $\downarrow$	F1 $\uparrow$	NC $\uparrow$	HD $\downarrow$	Time
Poisson	38,346	76,684	39.128	0.912	0.970	1.821	19.271
NDC(64^3)	4,134	8,327	0.798	0.986	0.978	0.921	4.971
NDC(128^3)	16,638	33,470	<u>0.676</u>	<u>0.991</u>	0.987	<u>0.716</u>	5.532
POCO(64^3)	8,237	16,497	29.817	0.853	0.951	6.521	4.367
POCO(128^3)	33,590	67,196	33.107	0.888	0.960	6.862	15.588
PoNQ	15,760	31,531	11.729	0.967	<u>0.982</u>	1.813	31.047
LMR	7,848	15,721	2.936	0.978	0.978	0.778	2,710.6
Ours	5,180	10,358	0.496	0.996	0.987	0.678	14.280

Table 2: Quantitative comparison with our HD-PEA and state-of-the-art surface reconstructions on 75 synthetic scanned models selected from Myles et al.’s dataset [48]. The best results are highlighted in bold, and the second best results are underlined. Note: CD ($\times 10^{-5}$), HD ($\times 10^{-2}$), and Time (seconds). All metrics are averaged over 75 models.

Method	# V_{out}	# f_{out}	CD $\downarrow$	F1 $\uparrow$	NC $\uparrow$	HD $\downarrow$	G $\uparrow$	Time
Poisson	77,296	154,594	26.346	0.922	0.966	2.147	0.528	39.298
NDC(64^3)	5,182	10,575	1.308	0.939	0.956	1.230	0.589	5.276
NDC(128^3)	21,011	42,721	<u>0.749</u>	<u>0.982</u>	0.971	<u>0.928</u>	0.588	5.928
+NASM	6,964	14,210	0.792	0.980	<u>0.975</u>	1.265	0.576	24.701
POCO(64^3)	14,246	28,600	94.402	0.608	0.859	11.665	0.537	6.329
POCO(128^3)	58,871	117,837	90.462	0.626	0.875	11.601	0.530	22.869
PoNQ	24,685	49,368	7.610	0.963	0.973	1.511	0.596	35.834
NKSR	12,023	24,040	1.461	0.954	0.966	1.745	0.525	0.705
SAP	31,692	63,320	2.081	0.976	0.972	1.447	0.526	6.575
SIREN	696,845	1,393,631	58.96	0.843	0.953	3.593	0.522	4658.97
LMR	7,590	15,124	10.048	0.964	0.972	2.471	<u>0.629</u>	3,153.8
Ours	7,028	13,989	0.649	0.989	0.979	0.798	0.702	14.594

Although our method and NASM (one of the most recent anisotropic meshing methods) are inspired by the Nash Embedding Theorem, target fundamentally different problems. HD-PEA reconstructs anisotropic surfaces directly from raw point clouds, while NASM performs remeshing on an existing mesh. HD-PEA operates on raw point clouds without any intermediate mesh. In practice (see Fig. 7), we use NDC(128^3) to generate the intermediate mesh for NASM, as NDC performed is the second-best method (besides ours) in our previous experiments. However, such two-stage pipelines (e.g., NDC+NASM) are sensitive to the quality of the intermediate mesh. As shown in Fig. 6, NDC produces severe aliasing artifacts, which subsequently cause NASM to overfit. By learning the anisotropic metric field directly in high-d embedding space from raw points, our HD-PEA avoids this artifact-prone stage, resulting in more stable and faithful reconstructions with consistently better surface accuracy (CD, F1, NC, HD) and anisotropic mesh quality metric G.

Challenging Scanning. In this section, we evaluate the performance of our method on degraded point clouds to demonstrate its robustness and practical applicability. Raw point clouds acquired from 3D scanners or multi-view reconstructions often suffer from measurement noise, outliers, and uneven sampling densities. To simulate these challenging conditions, we build upon the synthetic scanned dataset introduced in the previous section and apply some perturbations. Specifically, we add Gaussian noise with standard deviations of 0.001, 0.003, and 0.005 to simulate three levels of point-wise perturbations. Following the standard practice, we incorporate a lightweight preprocessing using the Moving Least Squares (MLS) method [23]. MLS performs local surface fitting to smooth point clouds, effectively reducing high-frequency noise while preserving fine geometric details. This denoising step is applied prior to inference and adds minimal computational overhead.

To simulate irregular sampling density, we resample scanned point clouds using striped and gradient density variations, such that certain regions of the




Fig. 7: Comparison between ours and NDC+NASM, where HD-PEA is better on CD↓ and G↑.




Fig. 8: Our anisotropic surface mesh approximation results on a Lion model from input point clouds with different noise levels and irregular sampling patterns.



Fig. 9: Our reconstruction of a ScanNet scene (some holes remain due to missing regions in the scanned input). Portions of the walls and roof are removed to reveal the interior geometry.

surface are more sparsely sampled, followed by PCP-Net [30]. Fig. 8 and Table 3 demonstrate the good robustness of our method.

Complicated Geometry and Large-Scale Scene. We further evaluate our method on point clouds with complex geometry details and scene-level datasets containing a significantly large number of points. In this setting, we employ the proposed patch-based meta-embedding inference (PMEI) to demonstrate the scalability and adaptability to large-scale data. This evaluation covers both model-level shapes with fine geometric details (from Myles et al.’s dataset) and scene-level point clouds (ScanNet dataset [16]) with hundreds of thousands of input points. Figure 9 demonstrates that our method has strong performance on large-scale point clouds on scene levels, even without fine-tuning or retraining. We also include both qualitative and quantitative comparisons between our method and NDC. Our method achieves a lower CD than NDC (0.786 vs. 0.814), while using only less than one-third of the mesh resolution. More results can be found in Section 7 of Appendix.

Curvature Tensor Estimation. One downstream application of our neural high-d point embedding is to estimate the curvature tensors from the input point clouds. For each point $\mathbf{p}$ in the original 3D space, we identify its six nearest neighbors $\mathbf{p}_j$ and construct a matrix $\mathbf{A} \in \mathbb{R}^{6 \times 3}$, where each row corresponds the

Table 3: Quantitative evaluation on our HD-PEA with different levels of noise on samplings as well as irregular sampling schemes on 75 synthetic scanned point clouds from Myles et al.’s dataset. Note: CD ($\times 10^{-5}$), HD ($\times 10^{-2}$), and Time (seconds). All metrics are averaged over 75 models.

Method	# V_{out}	# f_{out}	CD $\downarrow$	F1 $\uparrow$	NC $\uparrow$	HD $\downarrow$	Time
real scan+0.1%	8,684	17,255	0.822	0.978	0.974	0.893	14.809
real scan+0.3%	9,341	18,521	1.000	0.966	0.959	1.125	14.973
real scan+0.5%	9,601	19,043	1.417	0.933	0.947	1.358	14.997
gradient	7,526	14,907	0.730	0.982	0.976	0.970	14.630
striped	7,186	14,296	0.688	0.985	0.978	0.894	14.597

Table 4: Ablation study on different loss functions for neural high-d point embedding on 80 models selected from Thingi10K dataset. The best results are highlighted in bold. Note: CD ($\times 10^{-5}$) and HD ($\times 10^{-2}$). All metrics are averaged over 80 models.

Method	# V_{out}	# f_{out}	CD $\downarrow$	F1 $\uparrow$	NC $\uparrow$	HD $\downarrow$
MSE	12,632	48,368	30.864	0.953	0.949	2.148
MAE	12,116	48,421	29.737	0.958	0.952	2.044
NDHDE w/ MSE	8,688	48,991	30.072	0.984	0.979	1.065
NDHDE w/ MAE (Ours)	5,180	10,358	0.496	0.996	0.987	0.678

vector $\mathbf{p} - \mathbf{p}_j$. We also get the $\bar{\mathbf{p}}_j$ which are the corresponding predicted high-d embeddings of these 3D neighbors and construct a matrix $\mathbf{B} \in \mathbb{R}^{6 \times 8}$ where each row represents the vector $\bar{\mathbf{p}} - \bar{\mathbf{p}}_j$. We then solve the linear system of $\mathbf{A}\mathbf{x} = \mathbf{B}$ in least squares sense, yielding a matrix $\mathbf{x} \in \mathbb{R}^{3 \times 8}$. Applying singular value decomposition (SVD) to $\mathbf{x}$, we extract the three principal directions from its left singular vectors as the direction of the principal curvatures. The anisotropy stretching ratio is defined as s_1/s_2 where s_1 and s_2 denote the second and third singular values, respectively. Fig. 5 illustrates the estimated curvature tensors. We visualize each tensor using ellipsoid-style glyphs, where the axes represent the principal curvature directions. The color encodes the anisotropy, defined by the stretching ratios of the tensors. Another application of our HD-PEA method is to analyze its ability to preserve surface curvature characteristics, which directly influence the rendering quality of the surface roughness or smoothness as shown in Section 7 of Appendix.

4.3 Ablation Study

To demonstrate the effectiveness of our proposed neighborhood distance-based high-d embedding loss, i.e., NDHDE w/ MAE (in Section 3.1), we perform an ablation study comparing it against two commonly-used losses, i.e., Mean Squared Error (MSE) and Mean Absolute Error (MAE), both of which compute point-wise differences between predicted and ground truth embeddings. We also evaluate an MSE-based NDHDE and the ablation study is performed on Thingi10k dataset. Table 4 shows that our proposed NDHDE w/ MAE loss is most effective on the anisotropic surface mesh reconstruction w.r.t. all surface accuracy metrics. Furthermore, we also show the necessity of applying the proposed PMEI (in Section 3.1) on large-scale complicated data in Fig. 10. Fig. 11 shows that applying PCA to 3D neighbors for embedding tangent space approximation and particle optimization constraints (in Section 3.2) preserves topology better than using high-d neighbors.

4.4 Failure Case

Figure 12 illustrates a failure case on a challenging geometric shape with incomplete input point clouds from practical scanning. The clipped view removes

front- and back-facing points to better highlight the severely missing interior samples. Such incomplete sampling makes neural high-d point embedding and manifold approximation particularly challenging for thin structures (e.g., the slender cylinder), resulting in invalid or missing mesh elements in these regions.




Fig. 10: Comparison w/o and w/ patch-based meta-embedding inference (PMEI) on a complicated geometric model.




Fig. 11: Comparison on the proposed embedding tangent space approximation (TSA) with 3D and high-d neighbors.




Fig. 12: Failure case. Left: input point cloud; Middle: our result from the surface view; Right: our result from the interior view.

5 Conclusion

In this paper, we introduce a new scalable anisotropic surface approximation method through the high-d Euclidean point embedding. To our knowledge, this is the first deep learning-based method capable of generating high-quality and high-fidelity anisotropic surface meshes from unstructured point clouds without requiring a pre-computed mesh or metric field. Compared with existing surface mesh reconstruction methods, our HD-PEA achieves significantly better performance in surface accuracy, mesh quality, lightweightness, and scalability. Moreover, it shows strong potential for applications in robust large-scale surface reconstruction, surface rendering, and curvature tensor estimation.

Acknowledgments. The authors would like to thank the anonymous reviewers for their valuable comments and suggestions. Hongbo Li, Haikuan Zhu, Jing Hua, and Zichun Zhong were partially supported by National Science Foundation (OAC-1910469 and OAC-2311245), National Institutes of Health (R61NS119434), and General Motors research grant.

References

1. Binninger, A., Wiersma, R., Herholz, P., Sorkine-Hornung, O.: TetWeave: Isosurface extraction using on-the-fly Delaunay tetrahedral grids for gradient-based mesh optimization. *ACM Transactions on Graphics* **44**(4), 1–19 (2025)
2. Boissonnat, J.D., Flötotto, J.: A coordinate system associated with points scattered on a surface. *Computer-Aided Design* **36**(2), 161–174 (2004)
3. Boissonnat, J.D., Ghosh, A.: Manifold reconstruction using tangential Delaunay complexes. In: *Proceedings of the twenty-sixth annual symposium on Computational geometry*. pp. 324–333 (2010)
4. Boissonnat, J.D., Shi, K.L., Tournais, J., Yvinec, M.: Anisotropic Delaunay meshes of surfaces. *ACM Transactions on Graphics* **34**(2), 1–11 (2015)
5. Boissonnat, J.D., Wormser, C., Yvinec, M.: Anisotropic diagrams: Labelle Shewchuk approach revisited. *Theoretical Computer Science* **408**(2-3), 163–173 (2008)
6. Boltcheva, D., Lévy, B.: Surface reconstruction by computing restricted Voronoi cells in parallel. *Comput. Aided Des.* **90**(C), 123–134 (2017)
7. Borouchaki, H., George, P.L., Hecht, F., Laug, P., Saltel, E.: Delaunay mesh generation governed by metric specifications. part i. algorithms. *Finite elements in analysis and design* **25**(1-2), 61–83 (1997)
8. Borouchaki, H., George, P.L., Mohammadi, B.: Delaunay mesh generation governed by metric specifications part ii. applications. *Finite Elements in Analysis and Design* **25**(1-2), 85–109 (1997)
9. Boulch, A., Marlet, R.: POCO: Point convolution for surface reconstruction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6302–6314 (2022)
10. Budninskiy, M., Liu, B., De Goes, F., Tong, Y., Alliez, P., Desbrun, M.: Optimal Voronoi tessellations with Hessian-based anisotropy. *ACM Transactions on Graphics* **35**(6), 1–12 (2016)
11. Canas, G.D., Gortler, S.J.: Surface remeshing in arbitrary codimensions. *The Visual Computer* **22**(9), 885–895 (2006)
12. Chen, C., Liu, Y.S., Han, Z.: NeuralTPS: Learning signed distance functions without priors from single sparse point clouds. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **47**(1), 565–582 (2024)
13. Chen, L., Sun, P., Xu, J.: Optimal anisotropic meshes for minimizing interpolation errors in l^p -norm. *Mathematics of Computation* **76**(257), 179–204 (2007)
14. Chen, Z., Tagliasacchi, A., Funkhouser, T., Zhang, H.: Neural dual contouring. *ACM Transactions on Graphics* **41**(4), 1–13 (2022)
15. Chen, Z., Zhang, H.: Neural marching cubes. *ACM Transactions on Graphics (TOG)* **40**(6), 1–15 (2021)
16. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2017)
17. Dai, Y., Su, J.P., Fu, X.M.: Anisotropic triangular meshing using metric-adapted embeddings. *Computer Aided Geometric Design* **111**, 102314 (2024)
18. Dassi, F., Mola, A., Si, H.: Curvature-adapted remeshing of CAD surfaces. *Procedia Engineering* **82**, 253–265 (2014)
19. Dassi, F., Si, H., Perotto, S., Streckenbach, T.: Anisotropic finite element mesh adaptation via higher dimensional embedding. *Procedia Engineering* **124**, 265–277 (2015)

20. Dey, T.K., Dyer, R., Wang, L.: Localized Cocone surface reconstruction. *Computers Graphics* **35**(3), 483–491 (2011)
21. Dobrzynski, C., Frey, P.: Anisotropic Delaunay mesh adaptation for unsteady simulations. In: *Proceedings of the 17th international Meshing Roundtable*. pp. 177–194 (2008)
22. Du, Q., Wang, D.: Anisotropic centroidal Voronoi tessellations and their applications. *SIAM Journal on Scientific Computing* **26**(3), 737–761 (2005)
23. Fleishman, S., Cohen-Or, D., Silva, C.T.: Robust moving least-squares fitting with sharp features. *ACM Transactions on Graphics* **24**(3), 544–552 (2005)
24. Freedman, D.: Efficient simplicial reconstructions of manifolds from their samples. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **24**(10), 1349–1357 (2002)
25. Frey, P.J., Borouchaki, H.: Surface mesh quality evaluation. *International journal for numerical methods in engineering* **45**(1), 101–118 (1999)
26. Fu, X.M., Liu, Y., Snyder, J., Guo, B.: Anisotropic simplicial meshing using local convex functions. *ACM Transactions on Graphics* **33**(6), 1–11 (2014)
27. Fuchs, F.B., Worrall, D.E., Fischer, V., Welling, M.: Se(3)-transformers: 3D rotation equivariant attention networks. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems* (2020)
28. Garland, M., Heckbert, P.S.: Surface simplification using quadric error metrics. In: *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. pp. 209–216 (1997)
29. Gschwandtner, M., Kwitt, R., Uhl, A., Pree, W.: Blensor: Blender sensor simulation toolbox. In: *Bebis, G., Boyle, R., Parvin, B., Koracin, D., Wang, S., Kyungnam, K., Benes, B., Moreland, K., Borst, C., DiVerdi, S., Yi-Jen, C., Ming, J. (eds.) Advances in Visual Computing*. pp. 199–208 (2011)
30. Guerrero, P., Kleiman, Y., Ovsjanikov, M., Mitra, N.J.: PCPNet: Learning local shape properties from raw point clouds. *Computer Graphics Forum* **37**(2), 75–85 (2018)
31. Hanocka, R., Metzger, G., Giryas, R., Cohen-Or, D.: Point2mesh: A self-prior for deformable meshes. *arXiv preprint arXiv:2005.11084* (2020)
32. Heckbert, P.S., Garland, M.: Optimal triangulation and quadric-based surface simplification. *Computational Geometry* **14**(1–3), 49–65 (1999)
33. Huang, H., Li, D., Zhang, H., Ascher, U., Cohen-Or, D.: Consolidation of unorganized point clouds for surface reconstruction. *ACM Transactions on Graphics* **28**(5), 1–7 (Dec 2009)
34. Huang, J., Gojcic, Z., Atzmon, M., Litany, O., Fidler, S., Williams, F.: Neural kernel surface reconstruction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4369–4379 (2023)
35. Huang, Z., Wen, Y., Wang, Z., Ren, J., Jia, K.: Surface reconstruction from point clouds: A survey and a benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2024)
36. Ju, T., Losasso, F., Schaefer, S., Warren, J.: Dual contouring of hermite data. In: *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. pp. 339–346 (2002)
37. Kazhdan, M., Bolitho, M., Hoppe, H.: Poisson surface reconstruction. In: *Proceedings of the fourth Eurographics symposium on Geometry processing*. vol. 7 (2006)
38. Kovacs, D., Myles, A., Zorin, D.: Anisotropic quadrangulation. In: *Proceedings of the 14th ACM Symposium on solid and physical modeling*. pp. 137–146 (2010)
39. Lévy, B.: Geogram. *GitHub Repository*. URL: <https://github.com/BrunoLevy/geogram> (2015)

40. Lévy, B., Bonneel, N.: Variational anisotropic surface meshing with Voronoi parallel linear enumeration. In: Proceedings of the 21st international meshing roundtable. pp. 349–366 (2013)
41. Li, H., Zhu, H., Zhong, S., Wang, N., Lin, C., Guo, X., Xin, S., Wang, W., Hua, J., Zhong, Z.: NASM: Neural anisotropic surface meshing. In: SIGGRAPH Asia 2024 Conference Papers (2024)
42. Lin, S., Shi, Z., Liu, Y.: Fast and globally consistent normal orientation based on the winding number normal consistency. *ACM Transactions on Graphics* **43**(6) (2024)
43. Lipman, Y., Cohen-Or, D., Levin, D., Tal-Ezer, H.: Parameterization-free projection for geometry reconstruction. *ACM Transactions on Graphics* **26**(3), 22–es (Jul 2007)
44. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. In: Seminal graphics: pioneering efforts that shaped the field, pp. 347–353 (1998)
45. Lu, X., Wu, S., Chen, H., Yeung, S.K., Chen, W., Zwicker, M.: GPF: GMM-inspired feature-preserving point set filtering. *IEEE Transactions on Visualization and Computer Graphics* **24**(8), 2315–2326 (2018)
46. Maruani, N., Klovov, R., Ovsjanikov, M., Alliez, P., Desbrun, M.: VoroMesh: Learning watertight surface meshes with Voronoi diagrams. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14565–14574 (2023)
47. Maruani, N., Ovsjanikov, M., Alliez, P., Desbrun, M.: PoNQ: a neural QEM-based mesh representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3647–3657 (2024)
48. Myles, A., Pietroni, N., Zorin, D.: Robust field-aligned global parametrization. *ACM Transactions on Graphics* **33**(4) (2014)
49. Nash, J.: C^1 -isometric embeddings. *Annals of Mathematics* **60**(3), 383–396 (1954)
50. Nash, J.: The imbedding problem for Riemannian manifolds. *Annals of mathematics* **63**(1), 20–63 (1956)
51. Noda, T., Chen, C., Zhou, J., Zhang, W., Liu, Y.S., Han, Z.: Learning bijective surface parameterization for inferring signed distance functions from sparse point clouds with grid deformation. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 22139–22149 (2025)
52. Panozzo, D., Puppo, E., Tarini, M., Sorkine-Hornung, O.: Frame Fields: anisotropic and non-orthogonal cross fields. *ACM Transactions on Graphics* **33**(4), 134:1–134:11 (2014)
53. Pearson, K.: Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science* **2**(11), 559–572 (1901)
54. Peng, S., Jiang, C.M., Liao, Y., Niemeyer, M., Pollefeys, M., Geiger, A.: Shape as points: A differentiable poisson solver. In: Advances in Neural Information Processing Systems (NeurIPS) (2021)
55. Rouxel-Labbé, M., Wintraecken, M., Boissonnat, J.D.: Discretized Riemannian Delaunay triangulations. *Procedia engineering* **163**, 97–109 (2016)
56. Sakai, T.: Riemannian geometry, vol. 149. American Mathematical Soc. (1996)
57. Sharp, N., Ovsjanikov, M.: PointTriNet: Learned triangulation of 3D point sets. In: Proceedings of the European Conference on Computer Vision. pp. 762–778. Springer (2020)
58. Simpson, R.B.: Anisotropic mesh transformations and optimal error control. *Applied Numerical Mathematics* **14**(1–3), 183–198 (1994)

59. Sitzmann, V., Martel, J.N., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit neural representations with periodic activation functions. In: Proc. NeurIPS (2020)
60. Sumner, R.W., Popović, J.: Deformation transfer for triangle meshes. *ACM Transactions on Graphics* **23**(3), 399–405 (Aug 2004)
61. Sutherland, I.E., Hodgman, G.W.: Reentrant polygon clipping. *Communications of the ACM* **17**(1), 32–42 (1974)
62. Valette, S., Chassery, J.M., Prost, R.: Generic remeshing of 3D triangular meshes with metric-dependent discrete Voronoi diagrams. *IEEE Transactions on Visualization and Computer Graphics* **14**(2), 369–381 (2008)
63. Vetsch, M., Lombardi, S., Pollefeys, M., Oswald, M.R.: Neuralmeshing: Differentiable meshing of implicit neural representations. In: DAGM German Conference on Pattern Recognition. pp. 317–333. Springer (2022)
64. Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.G.: Pixel2Mesh: Generating 3D mesh models from single RGB images. In: Proceedings of the European Conference on Computer Vision. pp. 52–67 (2018)
65. Wen, C., Zhang, Y., Li, Z., Fu, Y.: Pixel2mesh++: Multi-view 3D mesh generation via deformation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1042–1051 (2019)
66. Witkin, A.P., Heckbert, P.S.: Using particles to sample and control implicit surfaces. In: Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques. p. 269–277 (1994)
67. Wu, X., Jiang, L., Wang, P.S., Liu, Z., Liu, X., Qiao, Y., Ouyang, W., He, T., Zhao, H.: Point transformer v3: Simpler faster stronger. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4840–4851 (2024)
68. Zhang, C., Wang, W., Li, X., Liao, X., Su, W., Tao, W.: High-fidelity lightweight mesh reconstruction from point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11739–11748 (2025)
69. Zhang, Z., Zha, H.: Principal manifolds and nonlinear dimensionality reduction via tangent space alignment. *SIAM journal on scientific computing* **26**(1), 313–338 (2004)
70. Zhong, S., Zhong, Z., Hua, J.: Surface reconstruction by parallel and unified particle-based resampling from point clouds. *Computer Aided Geometric Design* **71**, 43–62 (2019)
71. Zhong, Z., Guo, X., Wang, W., Lévy, B., Sun, F., Liu, Y., Mao, W.: Particle-based anisotropic surface meshing. *ACM Transactions on Graphics* **32**(4) (2013)
72. Zhong, Z., Wang, W., Lévy, B., Hua, J., Guo, X.: Computing a high-dimensional Euclidean embedding from an arbitrary smooth Riemannian metric. *ACM Transactions on Graphics* **37**(4) (2018)
73. Zhou, Q., Jacobson, A.: Thingi10K: A dataset of 10,000 3D-printing models. arXiv preprint arXiv:1605.04797 (2016)