

# SpaceMesh: A Continuous Representation for Learning Manifold Surface Meshes

TIANCHANG SHEN, NVIDIA, University of Toronto, Vector Institute, Canada

ZHAOSHUO LI, NVIDIA, USA

MARC LAW, NVIDIA, Canada

MATAN ATZMON, NVIDIA, Canada

SANJA FIDLER, NVIDIA, University of Toronto, Vector Institute, Canada

JAMES LUCAS, NVIDIA, UK

JUN GAO, NVIDIA, University of Toronto, Vector Institute, Canada

NICHOLAS SHARP, NVIDIA, USA




Fig. 1. We present a mesh representation that enables learning to directly generate polygonal meshes as the output of a neural network. *Left*: 3D meshes produced from a generative model trained on a dataset with desirable mesh connectivity. *Right*: our model can be applied to challenging tasks such as mesh repair, and produces manifold meshes suitable for downstream processing like computing geodesic distance.

Meshes are ubiquitous in visual computing and simulation, yet most existing machine learning techniques represent meshes only indirectly, *e.g.* as the level set of a scalar field or deformation of a template, or as a disordered triangle soup lacking local structure. This work presents a scheme to directly generate manifold, polygonal meshes of complex connectivity as the output of a neural network. Our key innovation is to define a continuous latent connectivity space at each mesh vertex, which implies the discrete mesh. In

Authors' addresses: Tianchang Shen, NVIDIA, University of Toronto, Vector Institute, Canada, frshen@nvidia.com; Zhaoshuo Li, NVIDIA, USA, maxzhaoshuol@nvidia.com; Marc Law, NVIDIA, Canada, marcl@nvidia.com; Matan Atzmon, NVIDIA, Canada, matzmon@nvidia.com; Sanja Fidler, NVIDIA, University of Toronto, Vector Institute, Canada, sfidler@nvidia.com; James Lucas, NVIDIA, UK, jalucas@nvidia.com; Jun Gao, NVIDIA, University of Toronto, Vector Institute, Canada, jung@nvidia.com; Nicholas Sharp, NVIDIA, USA, nsharp@nvidia.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SA Conference Papers '24, December 3–6, 2024, Tokyo, Japan

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1131-2/24/12

<https://doi.org/10.1145/3680528.3687634>

particular, our vertex embeddings generate cyclic neighbor relationships in a halfedge mesh representation, which gives a guarantee of edge-manifoldness and the ability to represent general polygonal meshes. This representation is well-suited to machine learning and stochastic optimization, without restriction on connectivity or topology. We first explore the basic properties of this representation, then use it to fit distributions of meshes from large datasets. The resulting models generate diverse meshes with tessellation structure learned from the dataset population, with concise details and high-quality mesh elements. In applications, this approach not only yields high-quality outputs from generative models, but also enables directly learning challenging geometry processing tasks such as mesh repair.

CCS Concepts: • **Computing methodologies** → **Shape representations; Reconstruction.**

Additional Key Words and Phrases: Mesh Generation, 3D Machine Learning, Graph Representations

## ACM Reference Format:

Tianchang Shen, Zhaoshuo Li, Marc Law, Matan Atzmon, Sanja Fidler, James Lucas, Jun Gao, and Nicholas Sharp. 2024. SpaceMesh: A Continuous Representation for Learning Manifold Surface Meshes. In *SIGGRAPH Asia 2024 Conference Papers (SA Conference Papers '24)*, December 3–6, 2024, Tokyo, Japan. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3680528.3687634>

## 1 INTRODUCTION

Polygonal meshes play an essential role in computer graphics, favored for their simplicity, flexibility, and efficiency. They can represent surfaces of arbitrary topology with non-uniform polygons, and support a wide range of downstream processing and simulation. Additionally, meshes are ideal for rasterization and texture mapping, making them efficient for rendering. However, the benefits of meshes rely heavily on their quality. For example, meshes with non-manifold connectivity or too many elements may break operations that leverage local structure, or make processing prohibitively expensive. Consequently, developing automatic algorithms and tools for generating high-quality meshes is an ongoing research focus.

It is no surprise that recent advancements in deep learning have led to growing interest in learning-based mesh creation. Generating meshes as output, however, is a notoriously challenging task for machine learning algorithms, as meshes have a complex combination of continuous and discrete structure. Not only do mesh vertices and edges form a graph, but mesh faces add additional interconnected structure, and furthermore those faces ought to be arranged locally for manifold connectivity. Existing approaches range from implicit function isosurfacing [Gao et al. 2022; Mescheder et al. 2019; Shen et al. 2021, 2023], which offers easy optimization and a guarantee of validity at the expense of restricting to a limited family of meshes, to directly generating faces as an array of vertex triplets [Alliegro et al. 2023; Nash et al. 2020; Siddiqui et al. 2023], a discrete-first perspective which cannot be certain to respect the constraints of local structure. This work seeks a solution that offers the best of all worlds: the ease and utility that comes from working in a continuous parameterization, a guarantee to produce meshes with manifold structure by construction, and the generality to represent the full range of possible meshes.

We present SpaceMesh, a representation for meshes built on continuous embeddings well-suited for learning and optimization, which guarantees manifold output and supports complex polygonal connectivity. Our approach derives from the halfedge data structure [Weiler 1986], which inherently represents manifold, oriented polygonal meshes—the heart of our contribution is a continuous parameterization for halfedge mesh connectivity.

The main idea is to represent mesh connectivity by first constructing a set of edges and halfedges, and then constructing the so-called *next* relationship among those halfedges to implicitly define the faces of the mesh. We introduce a parameterization of edge adjacency and *next* relationships with low-dimensional, per-vertex embeddings. These embeddings, by construction, always produce a manifold halfedge mesh without additional constraints. Moreover, the per-vertex embedding is straightforward to predict as a neural network output and demonstrates fast convergence during optimization. The continuous property of our representation facilitates new architectures for mesh generation, and enables applications like mesh repair with learning.

We validate our representation against alternatives for representing graph adjacency and meshes, and demonstrate superior significantly faster convergence, which is fundamentally important for learning tasks. Combined with a generative model for vertices, we

showcase our representation in learning different surface discretization for meshing. Additionally, our representation enables mesh repair via deep learning, simultaneously predicting both vertices and topology.

## 2 RELATED WORK

Initial deep learning-based mesh generation techniques focused on vertex prediction while maintaining fixed connectivity, which are challenging to adapt for complex 3D objects [Chen et al. 2019; Groueix et al. 2018; Hanocka et al. 2020; Litany et al. 2018; Liu et al. 2021; Ranjan et al. 2018; Tanwar et al. 2020; Wang et al. 2018; Zhang et al. 2020, 2021]. Although local topology modifications are possible through subdivision [Liu et al. 2020a; Wang et al. 2018] or remeshing [Palfinger 2022], these methods still struggle to represent general, complex 3D objects. Recent methods utilize intermediary representations that are converted into meshes using techniques like Poisson reconstruction on point clouds [Kazhdan et al. 2006; Peng et al. 2021] or isosurfacing on implicit fields [Chen et al. 2022; Gao et al. 2022; Lin et al. 2023; Shen et al. 2021, 2023]. However, these conversion processes lack precise control over mesh connectivity.

### 2.1 Generating Meshes

Much recent work has specifically studied approaches for generating surfaces meshes in learning-based pipelines.

*Volumetric 3D Reconstruction.* See Point2Surf [Erler et al. 2020], POCO [Boulch and Marlet 2022], NKSR [Huang et al. 2023] and BSP-Net [Chen et al. 2020], etc. These methods focus on reconstructing the geometric shape, rather than the mesh structure; output connectivity is always a marching-cubes mesh (or a union-of-planes in BSPNet). Our approach instead focuses on fitting particular discrete mesh connectivity structures from data. Figure 8 and 9 include a few representative methods from this family, although they generally target significantly different goals. A parallel class of methods leverages Voronoi/Delaunay-based formulations [Maruani et al. 2023, 2024]), but again these focus on fitting a surface’s geometric shape, rather than the particular mesh connectivity.

*Direct Mesh Learning.* See IER [Liu et al. 2020b], PointTriNet [Sharp and Ovsjanikov 2020], Delaunay Surface Elements (DSE) [Rakotasaona et al. 2021], DMesh [Son et al. 2024]. Like ours, these approaches aim to directly learn structured mesh connectivity. However, our approach offers a guarantee of manifoldness, and can encode general polygonal meshes. Additionally, we demonstrate the ability to encode concise artist/CAD-like tessellation via coupled learning of vertex positions and connectivity, rather than generating faces among a rough uniformly-sampled point set. Conversely, some of these methods scale to high resolution outputs, compared to our small-medium meshes. Figure 2 includes comparisons to DMesh [Son et al. 2024] as a representative method from this family, see also additional results from DSE in Figure 14 in the same setting.

*Sequence Modeling.* See Polygen [Nash et al. 2020], PolyDiff [Alliegro et al. 2023], MeshGPT [Siddiqui et al. 2023], and the concurrent MeshAnything [Chen et al. 2024]. These approaches use large-scale architectures to emit a mesh one face or vertex at a time. Unlike our method, they generally do not offer any guarantees of

connectivity or local structure, and all but Polygen produce triangle soup, connecting faces together only by generating vertices at coarsely-discretized categorical coordinates. However, by building on proven paradigms from language modeling, these models have been successfully trained at very large scale. Additionally, many of these approaches support only unconditional generation, and some are not publicly available. We include a gallery of qualitative comparisons in Figure 15.

## 2.2 Graph Learning

Our approach draws inspiration from graph learning representations, which have shown success for graphs including gene expression [Marbach et al. 2012], molecules [Kwon et al. 2020], stochastic processes [Backhoff-Veraguas et al. 2020], and social networks [Gehrke et al. 2003]. Based on the seminal work of Gromov [1987], Nickel and Kiela [2017] showed that hyperbolic embedding has fundamental properties which Euclidean embedding lacks (a relationship which has been well-studied in physics [Bombelli et al. 1987; Kronheimer and Penrose 1967; Meyer 1993]), to exploit the geometry of *spacetime* to represent graphs. In this paper, we leverage spacetime embeddings [Law and Stam 2020; Law and Lucas 2023] to put this perspective to work for generating meshes.

## 3 REPRESENTATION



We propose a continuous representation for the space of manifold polygonal meshes, which requires no constraints and is suitable for optimization and learning.

### 3.1 Background

*Manifold Surface Meshes.* A surface mesh  $\mathcal{M} = (\mathcal{V}, \mathcal{E}, \mathcal{F})$  consists of vertices  $\mathcal{V}$ , edges  $\mathcal{E}$ , and faces  $\mathcal{F}$ , where each vertex  $v \in \mathcal{V}$  has a position  $p_v \in \mathbb{R}^3$ . In a general polygonal mesh, each face is a cyclic ordering of 3 or more vertices. Each edge is an unordered pair of vertices which appear consecutively in one or more faces.

We are especially concerned with generating meshes which are not just a soup of faces, but which have coherent and consistent neighborhood connectivity. As such, we consider *manifold, oriented* meshes. Manifold connectivity is a topological property which does not depend on the vertex positions: *edge-manifoldness* means each edge has exactly two incident faces, while *vertex-manifoldness* means the faces incident on the vertex form a single edge-connected component homeomorphic to a disk. In an *oriented* mesh, all neighboring faces have a consistent outward orientation as defined by a counter-clockwise ordering of their vertices.

*Halfedge Meshes.* There are many possible data structures for mesh connectivity; we will leverage *halfedge meshes*, which by-construction encode manifold, oriented meshes with possibly polygonal faces, all using only a pair of references per element. As the name suggests, halfedge meshes are defined in terms of directed face-sides, called *halfedges* (see inset). Each halfedge stores two references: a *twin* halfedge, the oppositely-oriented halfedge along the same edge in a neighboring face, and a *next* halfedge, the subsequent halfedge within the same face.



The twin and next operators can be interpreted as a pair of permutations over the set of halfedges, this group-theoretic perspective is studied in combinatorics as a *rotation system*. A pair of permutations can be interpreted as a halfedge mesh as long as (a) neither operator maps any halfedge to itself, and (b) twin operator is an involution, i.e.  $\text{twin}(\text{twin}(h)) = h$ . The faces of the mesh are the *orbits* traversed by repeatedly following the next operator (see inset); we further require that these orbits all have a degree of at least three, to disallow two-sided faces. Our representation will construct a valid set of twin and next operators from a continuous embedding to define mesh connectivity.

### 3.2 Representing Edges

To begin, consider modeling a mesh simply as a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , later we will extend this model to capture manifold mesh structure via halfedge connectivity (Section 3.3). The vertex set  $\mathcal{V}$  can be viewed as a particular kind of point cloud, and point cloud generation is a well-studied problem ([Nichol et al. 2022; Zeng et al. 2022]). Likewise, continuous representations for generating undirected graph edges is a classic topic in graph representation learning [Law and Stam 2020; Nickel and Kiela 2017]. A basic approach is to associate an *adjacency embedding*  $x_v \in \mathbb{R}^k$  with each vertex, then define an edge between two vertices  $i, j$  if they are sufficiently close w.r.t. some distance function  $d$ :

$$\mathcal{E} := \{ \{i, j\} \text{ such that } d(x_i, x_j) < \tau \} \quad (1)$$

for some learned threshold  $\tau \in \mathbb{R}$ . Representing the vertices and edges of a mesh then amounts to two vectors for each vertex  $v$ : a 3D position  $p_v \in \mathbb{R}^3$  and an adjacency embedding  $x_v \in \mathbb{R}^k$ .

*Spacetime Distance.* We find that taking the adjacency features  $x$  as Euclidean vectors under pairwise Euclidean distance  $d^{\text{eu}}(x_i, x_j) = \|x_i - x_j\|_2$  is ineffective, with poor convergence in optimization and learning. There are many other possible choices of distance function for this embedding, but we find the recently proposed *spacetime* distance [Law and Lucas 2023] to be simple and highly effective. This distance function has deep interpretations in special relativity, defining pseudo-Riemannian structures. In our setting the *spacetime* distance  $d^{\text{st}}$  is computationally straightforward, splitting the components of  $x$  into a subvector  $x^s \in \mathbb{R}^{k^s}$  of space coordinates, and a subvector  $x^t \in \mathbb{R}^{k^t}$  of time coordinates:

$$d^{\text{st}}(x_i, x_j) = d^{\text{st}}([x_i^s, x_i^t], [x_j^s, x_j^t]) := \|x_i^s - x_j^s\|_2^2 - \|x_i^t - x_j^t\|_2^2, \quad (2)$$

where  $[\cdot, \cdot]$  denotes vector concatenation. Note that  $d^{\text{st}}$  is not a distance metric, and may be negative; this is of no concern, as we simply need to threshold it by some  $\tau \in \mathbb{R}$  to recover edges, treating  $\tau$  as an additional optimized parameter. In Figure 4 we show that this significantly accelerates convergence, see Section 4 for details.

*Loss Function.* At training time, we fit the adjacency embedding by supervising the distances under a cross entropy loss:


$$\sum_{i, j \in \mathcal{E}_{\text{gt}}} \log(\sigma(d(x_i, x_j) - \tau)) + \lambda \sum_{i, j \notin \mathcal{E}_{\text{gt}}} \log(\sigma(\tau - d(x_i, x_j))) \quad (3)$$

where  $\sigma$  is the logistic function (i.e. a sigmoid),  $\mathcal{E}_{\text{gt}}$  denotes the set of edges in the ground truth mesh, and  $\lambda > 0$  is a regularization parameter balancing positive and negative matches.

### 3.3 Representing Faces

To recover faces and manifold connectivity from a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , we further propose to parameterize halfedge connectivity for the mesh (Section 3.1). Given  $\mathcal{V}$  and  $\mathcal{E}$ , we construct the halfedge set by splitting each edge  $e_{ij}$  between vertices  $i, j$  into two oppositely-directed halfedges  $h_{ij}, h_{ji}$ . This pairing trivially implies the twin relationships as  $\text{twin}(h_{ij}) = h_{ji}$ ; we then only need to specify the next relationships to complete the halfedge mesh and define the face set.

*Neighborhood Orderings.* The next operator defines a cyclic permutation with a single orbit on the halfedges outgoing from each vertex. Thus the task of assigning the next operator (and implicitly, the potentially-polygonal faces of the mesh) comes down to learning this permutation for each vertex.



*Representing Neighborhood Orderings.* For each vertex, we define a triplet of continuous *permutation features*:  $y^{\text{root}}, y^{\text{prev}}, y^{\text{next}} \in \mathbb{R}^{k^p}$ . These are used to determine the local cyclic ordering of incident edges. Precisely, in the local neighborhood of each vertex  $i \in \mathcal{V}$  with degree  $D$ , for each pair of edges  $e_{ij}, e_{ik}$ , we combine the features of vertices  $i, j$  and  $k$  via a scalar-valued function  $F(y_i^{\text{root}}, y_j^{\text{prev}}, y_k^{\text{next}})$  (see Section 4.3). Gathering these pairwise entries yields a nonnegative matrix in the local neighborhood of each vertex:

$$\Phi^i \in \mathbb{R}^{D \times D}, \quad \Phi_{jk}^i := e^{F(y_i^{\text{root}}, y_j^{\text{prev}}, y_k^{\text{next}})}, \quad (4)$$

where each row corresponds to an incident edge. We then use Sinkhorn normalization [Sinkhorn 1964] to recover a doubly-stochastic matrix,  $\bar{\Phi}^i$ , representing a *softened* permutation matrix [Adams and Zemel 2011].

*Loss Function.* At training or optimization time, we simply supervise the matrices  $\bar{\Phi}$  directly with the ground truth permutation matrices using binary cross-entropy loss:

$$\sum_{\{i,j,k\} \in \mathcal{N}_{\text{gt}}} -\log(\bar{\Phi}_{jk}^i), \quad (5)$$

where  $\mathcal{N}_{\text{gt}}$  is the set of all next relationships in ground truth mesh such that  $\text{next}(h_{ij}) = h_{jk}$ . Note that we do not need to supervise the remaining entries of  $\bar{\Phi}^i$ , which is already Sinkhorn-normalized.

*Extracting Meshes.* At inference time to actually extract a mesh, for each vertex neighborhood we seek the lowest-cost matching under the pairwise cost matrix  $-\bar{\Phi}^i$ , among only those matchings which form a single orbit. To compute this matching, we first compute the optimal unconstrained lowest-cost matching [Jonker and Volgenant 1988]; often this matching already forms a single orbit, but when it does not we fall back on a greedy algorithm which starts at an arbitrary entry and repeatedly takes the next lowest-cost entry

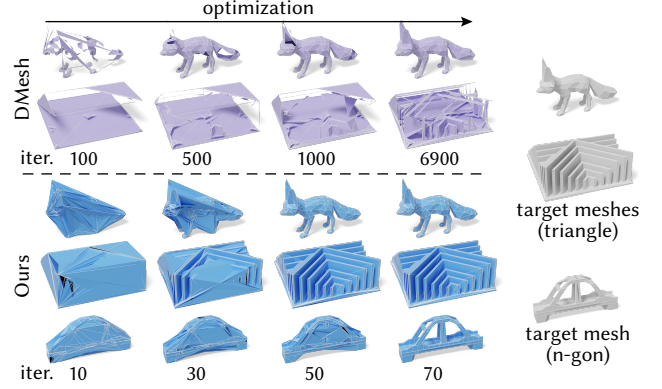


Fig. 2. Fitting the ground truth connectivity of a single mesh tessellated with triangles and n-gons.

without violating the single-orbit constraint. These neighborhood matchings then imply halfedge connectivity as

$$\text{next}(h_{ij}) := h_{ki} \quad \text{for } k = \text{match}_{\bar{\Phi}^i}(j). \quad (6)$$

This completes the halfedge mesh representation. Faces, potentially of any polygonal degree, can then be extracted as orbits of the next operator.

## 4 VALIDATION

In this section, we evaluate the basic properties of our method, by directly optimizing to fit both individual meshes and collections of meshes, as well as ablating design choices.

### 4.1 Encoding a Given Mesh

The most basic task for a mesh representation is to directly fit it to encode a particular mesh. Though straightforward in principle, this optimization could fail if a representation is unable to represent all possible meshes, or if local minima and slow convergence make fitting ineffective in practice. We consider three different challenging meshes with thin parts, anisotropic faces, and varying geometric details. For each single shape, we optimize to encode its connectivity with our per-vertex embeddings  $(x_i, y_i^{\text{root}}, y_i^{\text{prev}}, y_i^{\text{next}})$  using the loss functions from Equation 3 and 5. In Figure 2 we show the result of this optimization with our approach, as well as with the recent DMesh [Son et al. 2024], which proposes a Delaunay-based mesh representation. Our method not only converges much faster to the correct connectivity, but also is applicable to polygonal meshes, making it more suitable for general mesh generation tasks. Further experimental details are provided in the Supplement.




Fig. 3. Meshes encoded by our autoencoder.




Fig. 4. Comparison of convergence speed with different distance functions.




Fig. 5. Comparison of convergence speed with different permutation feature reduction functions.

## 4.2 Fitting Mesh Collections

As a next basic test of the ability of our method to encode collections of shapes in a learning setting, we train a simple auto-decoder architecture on a subset of 200 shapes from the Thingi10k dataset, a challenging set of real-world models originally for 3D printing [Zhou and Jacobson 2016]. To be clear, we do not aim to demonstrate downstream learning tasks with this experiment, we simply validate that our representation can simultaneously represent a variety of complex shapes, even when the embeddings are parameterized by a neural network, see Section 5 for large-scale learning and applications. In particular, here we allocate a latent code for each mesh, and optimize those latent codes as well as the parameters of a simple transformer model [Vaswani et al. 2017] that decodes each latent code into the mesh, in the form of per-vertex positions and connectivity embeddings of our representation. See the Supplement for further experimental details. As shown in Figure 3, our model faithfully overfits the shape collection. Quantitatively, the encoded meshes achieve a mean L2 loss of 0.00062, an F1 score of 0.99 for adjacency prediction, and an accuracy of 0.98 for permutation predictions. This is positive evidence that the representation is able to simultaneously represent many complex shapes, even with significant geometric complexity and the nonconvexity of the neural parameterization.

## 4.3 Ablating Design Choices

*Spacetime Distance.* We find spacetime distance to be a dramatically more effective representation than Euclidean or other metrics to define adjacency embeddings (Section 3.2), in the sense that it can be optimized much more easily. To demonstrate this, we fit the edges of the bridge mesh appearing on the bottom right of Figure 4 using each of three formulations for  $d(x_i, x_j)$ : (1) the spacetime distance introduced in Section 3.2, (2) the squared Euclidean distance  $\|x_i - x_j\|_2^2$ , and (3) the negative dot product  $-x_i^T x_j$ . Figure 4 shows the speed of convergence—spacetime distance converges much faster compared to other distance formulations, which we observed consistently across all experiments.

*Permutation Feature Reduction.* We also investigate several choices for the permutation feature reduction function  $F$  (Equation 7), including elementwise maximum, addition, or concatenation. Figure 5

shows the results. We find elementwise multiplication followed by summation of all elements to be most effective. Precisely, we use

$$F(y_i^{\text{root}}, y_j^{\text{prev}}, y_k^{\text{next}}) := \text{trace}(\text{diag}(y_j^{\text{prev}})\text{diag}(y_i^{\text{root}})\text{diag}(y_k^{\text{next}})), \quad (7)$$

where  $\text{diag}$  denotes constructing a diagonal matrix from a vector.

## 5 APPLICATION: LEARNING TO MESH

Equipped with a continuous representation for manifold polygonal meshes, we can then begin large-scale learning atop the representation. In this section, we integrate SpaceMesh with a 3D generative model to generate meshes conditioned on geometry provided as a point cloud. This conditioned model can then be directly applied to mesh repair without fine-tuning (Section 5.5).

### 5.1 Model Architecture

Our model architecture (Figure 6), consists of three modules: a point cloud encoding network for processing geometry information, a vertex diffusion model to generate 3D locations for vertices, and a connectivity prediction network to predict per-vertex embeddings.

*Point Cloud Encoder.* We encode the point cloud using PVCNN [Liu et al. 2019] to generate the feature volumes at multiple spatial resolutions. These feature volumes, as geometry context, guide the subsequent mesh generation. Note that this input point cloud is *not* the resulting mesh vertex set, it is conditioning information indicating the geometry that we are trying to generate a mesh of.

*Vertex Position Generation Network.* We re-purpose Point-E [Nichol et al. 2022], a diffusion transformer network that was originally designed for point cloud generation, to generate sparse mesh vertices conditioned on the geometry context from the encoder. Specifically, we first initialize the vertex position by sampling from a Gaussian distribution, and iteratively denoise the vertex location through the diffusion transformer. At each denoising step, we feed the input to the transformer by concatenating the vertices’ positions with features that are tri-linearly interpolated with the multi-resolution feature volumes from the encoder to capture the geometry information. If needed, we handle varying vertex counts by padding to a predefined maximum size, and additionally diffusing a binary mask at each vertex to indicate which vertices are artificial padding.

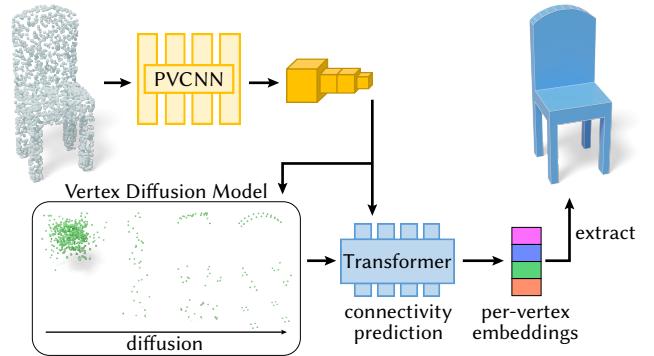


Fig. 6. Network architecture for learning to generate meshes.




Fig. 7. Conditioned on the same geometry, our model can generate different styles of meshes depending to the distribution it was trained on. Each row denotes a style of mesh, for which we construct a dataset of meshed primitive surfaces and fit our model. Because our model is generative, it matches the distribution but does not exactly replicate vertex positions or connectivity.

*Vertex Connectivity Prediction Network.* We leverage a transformer architecture [Vaswani et al. 2017] to predict the per-vertex connectivity embeddings given vertex positions. Similar to the vertex position generation network, we concatenate vertex position with the interpolated feature from the encoder for each vertex, and predict the adjacency embeddings  $x$  and permutation embeddings  $y^{\text{root}}, y^{\text{prev}}, y^{\text{next}}$ . We remove the positional embedding from the original transformer and predict the embeddings for all the vertices simultaneously by using the self-attention across the vertices.

*Training Details.* We train all the neural networks together. To train the vertex position generation network, we adopt the  $\epsilon$ -prediction from the diffusion model [Ho et al. 2020; Nichol et al. 2022]. To train the connectivity generation model, we combine the losses Equation 3 and 5, supervising on meshes from the dataset. Further details are provided in the Supplement.

## 5.2 Basic Validation on a Synthetic Dataset

Our model learns to fit *distributions* of meshes; the tessellation pattern and element shapes of generated meshes will mimic the training population. We first demonstrate this behavior with a simple synthetic dataset, constructed by generating shapes as a union of randomly arranged cubes, tetrahedra, and spheres. For each shape, we extract a 3D iso-surface using Dual Marching Cubes [Nielson 2003], and mesh it according to several strategies: (1) isotropic remeshing [Hoppe et al. 1993] with Meshlab [Cignoni et al. 2008] (2) planar decimation from Blender [Community 2018] to create N-gon mesh. (3) QEM for surface simplification [Garland and Heckbert 1997] from Meshlab, and (4) InstantMesh [Jakob et al. 2015] to create a quad-dominant mesh with the official implementation

In Figure 7, we show how training on each of these datasets causes our model to generate different styles of meshes as outputs. The four models, when each given the same point cloud as input specifying the desired geometry, produce respectively (1) isotropic triangle meshes, (2) minimal planar-decimated meshes, (3) QEM-simplified meshes, and (4) quad-dominant meshes.

Table 1. Accuracy and quality statistics for mesh reconstruction.

Method	CD ( $10^{-3}$ ) ↓	F1↑	ECD( $10^{-2}$ ) ↓	EF1↑	#V	#F	IN↓
PSR	46.35	0.44	56.81	0.03	2406	4736	63.31
PSR*	46.72	0.42	51.86	0.03	494	968	61.61
OccNet	11.31	0.47	33.08	0.08	7344	14688	48.53
Pixel2Mesh	6.37	0.48	29.52	0.09	2466	4928	52.03
Ours	1.39	0.66	3.21	0.42	512	1818	34.54

## 5.3 Learning Meshes from the ABC Dataset

To evaluate learning at scale on a realistic dataset, train our model on ABC dataset [Koch et al. 2019a], which consists of watertight triangle meshes of CAD shapes with isotropic triangle distribution. The meshes in the ABC dataset exhibit considerable diversity, featuring both sharp and smooth curved geometric features. We employed a benchmark [Koch et al. 2019b] subset of 10,000 shapes, all with 512 vertices, randomly split into 80% for training and 20% for testing. To obtain the input conditioning point cloud, we uniformly sampled 2048 points from the mesh surface.

*Baselines.* We compare our model against both classic and learning-based point cloud reconstruction methods. As a representative classic approach, we compare to Poisson Surface Reconstruction (PSR) [Kazhdan et al. 2006] as implemented in Open3D [Zhou et al. 2018], with meshes extracted via marching cubes [Lorensen and Cline 1998]. We also consider isotropic remeshing [Hoppe et al. 1993] on the output of Poisson reconstruction to obtain a more compact mesh tessellation, which is denoted PSR\*. For representative learning-based approaches, we choose Pixel2Mesh [Wang et al. 2018], which deforms a template sphere to generate a mesh, and OccNet [Mescheder et al. 2019], which predicts an implicit field and extracts the mesh using Marching Cubes [Lorensen and Cline 1998] afterwards. For a fair comparison among deep learning based methods, we adopt the same point cloud encoder as our approach.

*Metrics.* Our primary goal is to evaluate the ability to capture the desired distribution of surface discretization, as measured by intrinsic mesh statistics such as edge lengths and corner angles for each polygon. Furthermore, although our method is not directly designed to minimize reconstruction error, we additionally evaluate our method against baselines on how well the generated meshes align with ground truth geometry. To this end, we follow the methodology from NDC [Chen et al. 2022] and compute Chamfer Distance (CD), F-Score (F1), Edge Chamfer Distance (ECD), Edge F-Score (EF1), and the percentage of Inaccurate Normals (IN $> 10^\circ$ ) with respect to the ground truth mesh. A detailed description of these metrics is provided in the Supplement.

*Results.* As shown in Figure 9 and Table 1, both qualitative and quantitative results demonstrate that our method outperforms baselines under the target metrics, particularly in recovering sharp features. The vertices and edges align accurately with sharp features, highlighting the advantage of directly generating meshes as the output representation. As shown in Figure 8, the distribution of element shapes from our generated meshes aligns much better with the ground truth than the baselines, demonstrating the ability of our model to predict connectivity which aligns with the target training

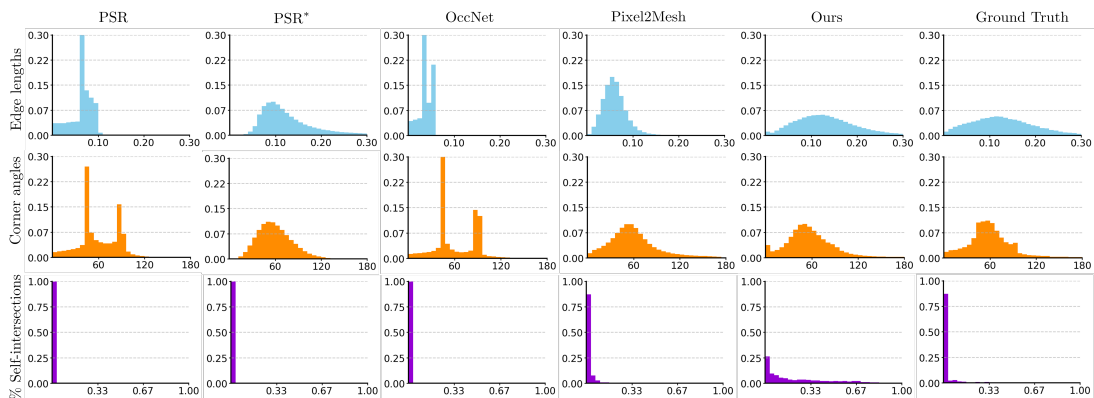


Fig. 8. Quantitative comparison of the intrinsic quality of reconstructed meshes. Our method produces meshes with distributions of edge lengths and corner angles more closely aligned with the ground truth. This is because our model learns the surface discretization from the data, unlike other methods that primarily focus on reconstructing geometry. We additionally report the percentage of faces with self-intersections in each mesh.

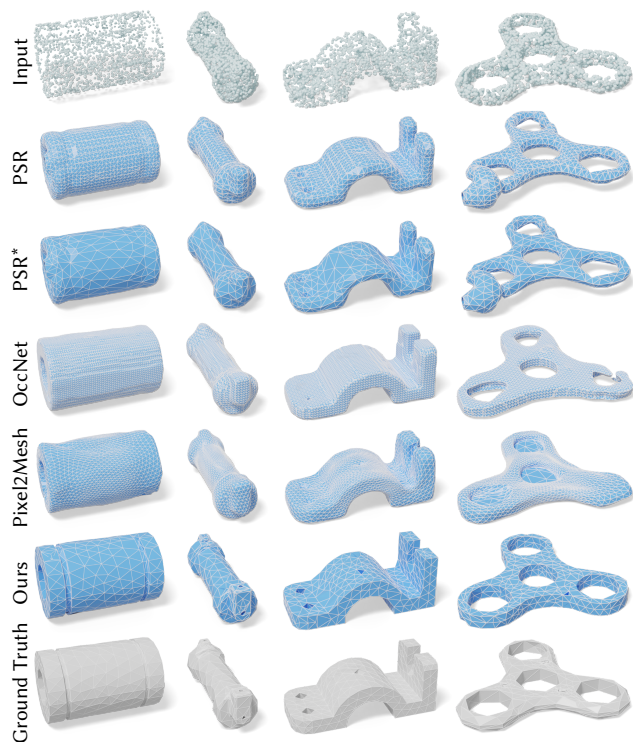


Fig. 9. Generated meshes for the ABC Dataset.

population. Note that although our representation guarantees manifold connectivity, there may still be geometric self-intersections between faces. We report the fraction of faces in each mesh with self-intersections in Figure 8, and provide further discussion in Section 6.

#### 5.4 Learning Meshes from the ShapeNet Dataset

Following recent work on mesh generation [Alliegro et al. 2023; Gao et al. 2022; Nash et al. 2020; Siddiqui et al. 2023], we further evaluate our model on ShapeNet dataset [Chang et al. 2015].




Fig. 10. By leveraging a diffusion model, we can generate different meshes from the same input condition. Notice how the chair legs are modeled with different topologies, which all conform to the input condition.

*Dataset Details.* As in prior work [Nash et al. 2020; Siddiqui et al. 2023], we note that the raw meshes from ShapeNet consist largely of non-manifold meshes with duplicated faces and T-junctions at intersections, and thus we preprocess all shapes by removing duplicated faces and applying planar decimation with varying thresholds to simplify them into minimal polygonal meshes. After this preprocessing, the majority of the shape are still non-manifold, making them unsuitable for our goal of generating manifold meshes with clean connectivity. We thus remove all non-manifold shapes, resulting in a total of 20,255 shapes. We adhere to an 80-20 train-test split and randomly sample 2,048 surface points as geometry conditioning input. Additionally, we apply random scaling augmentation during training. Unlike previous autoregressive methods, our approach does not require quantization of vertices.

*Baselines.* Many relevant baselines [Alliegro et al. 2023; Nash et al. 2020; Siddiqui et al. 2023] do not have either training or inference code available, and regardless there are many differences in experimental protocols and target task. As such, we instead focus on qualitative comparisons to give intuition about the differences between these methods, primarily in regard to mesh quality.

*Results.* Figure 15 shows a variety of results generated by our method, as well as a sampling of published results from baselines. Our method generates sharp and compact polygonal meshes that match with the input condition and are guaranteed to be manifold. We also note a promising *diversity* in our outputs on this dataset: because our model uses a probabilistic diffusion model to generate




Fig. 11. Our trained conditioned meshing model can be repurposed for mesh repair. For a mesh with good geometry but poor tessellation in certain regions (highlighted in red), the user can mark those regions and pass the mesh to our model to re-predict both vertices and connectivity, effectively repairing the mesh (highlighted in green).




Fig. 12. A visual comparison of mesh repair methods. Note that our method additionally takes surface points sampled from the whole mesh as input, unlike other methods which use only the partial mesh.

vertices, we are able to produce distinct meshes conditioned on the same point cloud input by repeatedly sampling the model (Figure 10).

### 5.5 Mesh Repair

Lastly, we demonstrate the application of our model to the downstream geometry processing task of mesh repair. As illustrated in Figure 11, we envision a workflow where a user identifies a region of a mesh with poor tessellation such as self-intersections, skinny triangles, or non-manifold structures, and wishes to re-triangulate that region in a way that seamlessly blends with the surrounding mesh. We show that we can repurpose our model for this task without retraining, by viewing it as *mesh inpainting*, in the same sense that image models are used to inpaint undesired regions of images according to some conditioning while matching the surrounding context. We inpaint the mesh by sampling a point cloud from the desired geometry and applying our generative model, projecting during diffusion to ensure the fixed region of the input mesh is retained—see the Supplement for an in-depth explanation. Note that MeshGPT [Siddiqui et al. 2023] also demonstrated completion of a partial mesh; however, it was limited to bottom-up completion due to auto-regressive inference with sorted vertices.

**Results.** We visualize the results in Figure 12. Our approach generates high-quality patches to fill the removed regions in the partial meshes while preserving the geometry and connectivity of the input. For comparison, we also include the most similar results of which we are aware: a classic mesh repair framework, MeshFix [Attene 2010], and a recent learning-based method, SeMIGCN [Hattori et al. 2024]. However, note that this is not exactly an apples-to-apples comparison, our method additionally takes the surface point cloud




Fig. 13. Failure cases from our learning to mesh model. Although they still have manifold connectivity, large erroneous faces and excessive self-intersections yield a tangled mesh with poor geometric accuracy.

of the complete shape as input, with a focus on re-generating surface discretization while preserving geometry. MeshFix is designed only for hole filling and cannot generate a repaired mesh conditioned on the geometry. In contrast, SeMIGCN re-meshes the shape for running GCN, resulting in an overly dense mesh that might not be desirable. We compare quantitatively with 100 randomly sampled examples from the ABC dataset validation shapes. SpaceMesh achieved a Chamfer Distance (CD) of  $0.77 (10^{-3})$  and a 0.76 F1 score. The baselines, SeMIGCN and MeshFix, achieve a CD of  $39.50 (10^{-3})$  and  $31.59 (10^{-3})$ , and an F1 score of 0.57 and 0.72, respectively.

## 6 DISCUSSION

**Scalability and Runtime.** Our approach represents discrete connectivity via a fixed-size continuous embedding per-vertex. Concrete results about the size of such an embedding needed to represent all possible discrete structures remain an open problem in graph theory [Nickel et al. 2014; Nickel and Kiela 2017]. In practice, we find low-dimensional embeddings  $k < 10$  to be sufficient to represent every mesh in our experiments. Encoding a 10,000-vertex mesh via direct optimization, as shown in Figure 2, converges in 600 iterations (approximately 2 minutes) with  $k^P = 6$ .

For learning, the bottleneck is memory usage in transformer blocks. We demonstrate generations up to 2,000 vertices in the auto-decoder setting; this is modest compared to high-resolution meshes, but it already captures many CAD and artist-created assets, and exceeds other recent mesh generation works (e.g., around 200 vertices in MeshGPT [Siddiqui et al. 2023]). Our generative model takes less than 2 seconds to generate a single mesh, which is notably faster than recent auto-regressive models like MeshGPT, which require 30-90 seconds. All inference and optimization times are measured on an NVIDIA A6000 GPU.

**Limitations.** Although our representation guarantees manifold connectivity, it may contain other errors such as self-intersections, spurious high-degree polygons, or significantly non-planar faces. The frequency of such errors depends on how the representation is generated or optimized: often they have little effect on the approximated surface (Figure 9), but in other cases they may significantly degrade the generated geometry, as shown in Figure 13. Note that such artifacts are not always erroneous—meshes designed by artists often intentionally include self-intersections; if desired, we could potentially mitigate self-intersections by penalizing them with regularizers during training.



Our implementation does not handle open surfaces, this could be addressed by predicting a flag for boundary edges much like we predict a mask for padded vertices. Also, like other diffusion-based generative models, our large-scale learning experiments may produce nonsensical outputs for difficult or out-of-distribution input.

*Future Work.* Looking forward, we see many possibilities to build upon our representation for directly generating meshes in learning pipelines. In the short term, this could mean generating connectivity embeddings as well as vertex positions from a diffusion model, and in the longer term, one might even fit SpaceMesh generators in an unsupervised fashion using energy functions to remove the reliance on mesh datasets for supervision entirely.

## ACKNOWLEDGMENTS

The authors are grateful to Yawar Siddiqui for providing the results of MeshGPT and Polygen, as well as the anonymous reviewers for their valuable comments and feedback.

## REFERENCES

- Ryan Prescott Adams and Richard S Zemel. 2011. Ranking via sinkhorn propagation. *arXiv preprint arXiv:1106.1925* (2011).
- Antonio Alliegro, Yawar Siddiqui, Tatiana Tommasi, and Matthias Nießner. 2023. PolyDiff: Generating 3D Polygonal Meshes with Diffusion Models. *arXiv preprint arXiv:2312.11417* (2023).
- Marco Attene. 2010. A lightweight approach to repairing digitized polygon meshes. *The visual computer* 26 (2010), 1393–1406.
- Julio Backhoff-Veraguas, Daniel Bartl, Mathias Beiglböck, and Manu Eder. 2020. All adapted topologies are equal. *Probability Theory and Related Fields* 178 (2020), 1125–1172.
- Luca Bombelli, Joohan Lee, David Meyer, and Rafael D Sorkin. 1987. Space-time as a causal set. *Physical review letters* 59, 5 (1987), 521.
- Alexandre Boulch and Renaud Marlet. 2022. Poco: Point convolution for surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6302–6314.
- Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. 2015. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012* (2015).
- Wenzheng Chen, Huan Ling, Jun Gao, Edward Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. 2019. Learning to predict 3d objects with an interpolation-based differentiable renderer. *Advances in neural information processing systems* 32 (2019).
- Yiwen Chen, Tong He, Di Huang, Weicai Ye, Sijin Chen, Jiayang Tang, Xin Chen, Zhongang Cai, Lei Yang, Gang Yu, et al. 2024. MeshAnything: Artist-Created Mesh




Fig. 14. Fitting the ground truth connectivity of a single mesh with DSE [Rakotosaona et al. 2021]. The experiment setting is described in Section 4.1. Here DSE is overfit to encode a single shape, but even then its representation struggles when vertices are sparse and geometry is highly nonconvex.




Fig. 15. Visual comparison of generated meshes from models trained on ShapeNet. The top part showcases results from three unconditioned mesh generation methods: GET3D, PolyGen, and MeshGPT. The bottom part shows meshes generated by our model, which takes point clouds as input.

- Generation with Autoregressive Transformers. *arXiv preprint arXiv:2406.10163* (2024).
- Zhiqin Chen, Andrea Tagliasacchi, Thomas Funkhouser, and Hao Zhang. 2022. Neural Dual Contouring. *ACM Transactions on Graphics (Special Issue of SIGGRAPH)* 41, 4 (2022).
- Zhiqin Chen, Andrea Tagliasacchi, and Hao Zhang. 2020. BSP-Net: Generating Compact Meshes via Binary Space Partitioning. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2020).
- Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, Guido Ranzuglia, et al. 2008. Meshlab: an open-source mesh processing tool. In *Eurographics Italian chapter conference*, Vol. 2008. Salerno, Italy, 129–136.
- Blender Online Community. 2018. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam. <http://www.blender.org>
- Philipp Erler, Paul Guerrero, Stefan Ohrhallinger, Niloy J Mitra, and Michael Wimmer. 2020. Points2surf learning implicit surfaces from point clouds. In *European Conference on Computer Vision*. Springer, 108–124.
- Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojic, and Sanja Fidler. 2022. GET3D: A Generative Model of High Quality 3D Textured Shapes Learned from Images. In *Advances In Neural Information Processing Systems*.
- Michael Garland and Paul S Heckbert. 1997. Surface simplification using quadric error metrics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 209–216.
- Johannes Gehrke, Paul Ginsparg, and Jon Kleinberg. 2003. Overview of the 2003 KDD Cup. *Acm Sigkdd Explorations Newsletter* 5, 2 (2003), 149–151.
- Mikhael Gromov. 1987. Hyperbolic groups. In *Essays in group theory*. Springer, 75–263.
- Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. 2018. AtlasNet: A Papier-Mâché Approach to Learning 3D Surface Generation. *CVPR* (2018).
- Rana Hanocka, Gal Metzer, Raja Giryes, and Daniel Cohen-Or. 2020. Point2Mesh: A Self-Prior for Deformable Meshes. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 1–12.
- Shota Hattori, Tatsuya Yatagawa, Yutaka Ohtake, and Hiromasa Suzuki. 2024. Learning Self-Prior for Mesh Inpainting Using Self-Supervised Graph Convolutional Networks. *IEEE Transactions on Visualization and Computer Graphics* (2024).
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems* 33 (2020), 6840–6851.
- Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. 1993. Mesh optimization. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*. 19–26.
- Jiahui Huang, Zan Gojic, Matan Atzmon, Or Litany, Sanja Fidler, and Francis Williams. 2023. Neural kernel surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4369–4379.
- Wenzel Jakob, Marco Tarini, Daniele Panozzo, and Olga Sorkine-Hornung. 2015. Instant Field-Aligned Meshes. *ACM Transactions on Graphics (Proceedings of SIGGRAPH ASIA)* 34, 6 (Nov. 2015). <https://doi.org/10.1145/2816795.2818078>
- Roy Jonker and Ton Volgenant. 1988. A shortest augmenting path algorithm for dense and sparse linear assignment problems. In *DGOR/NSOR: Papers of the 16th Annual Meeting of DGOR in Cooperation with NSOR/Vorträge der 16. Jahrestagung der DGOR zusammen mit der NSOR*. Springer, 622–622.
- Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. 2006. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, Vol. 7.
- Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. 2019a. ABC: A Big CAD Model Dataset For Geometric Deep Learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. 2019b. Abc dataset normal estimation benchmark. (2019).
- Erwin H Kronheimer and Roger Penrose. 1967. On the structure of causal spaces. In *Mathematical Proceedings of the Cambridge Philosophical Society*, Vol. 63. Cambridge University Press, 481–501.
- Youngchun Kwon, Dongseon Lee, Youn-Suk Choi, Kyoham Shin, and Seokho Kang. 2020. Compressed graph representation for scalable molecular graph generation. *Journal of Cheminformatics* 12, 1 (2020), 1–8.
- Marc Law and Jos Stam. 2020. Ultrahyperbolic representation learning. *Advances in neural information processing systems* 33 (2020), 1668–1678.
- Marc T. Law and James Lucas. 2023. Spacetime Representation Learning. In *The Eleventh International Conference on Learning Representations*. [https://openreview.net/forum?id=qV\\_M\\_rhYajc](https://openreview.net/forum?id=qV_M_rhYajc)
- Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiao-hui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. 2023. Magic3D: High-Resolution Text-to-3D Content Creation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Or Litany, Alex Bronstein, Michael Bronstein, and Ameesh Makadia. 2018. Deformable Shape Completion with Graph Convolutional Autoencoders. *CVPR* (2018).
- Hsueh-Ti Derek Liu, Vladimir G. Kim, Siddhartha Chaudhuri, Noam Aigerman, and Alec Jacobson. 2020a. Neural Subdivision. *ACM Trans. Graph.* 39, 4 (2020).
- Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. 2021. Mesh Graphormer. In *ICCV*.
- Minghua Liu, Xiaoshuai Zhang, and Hao Su. 2020b. Meshing point clouds with predicted intrinsic-extrinsic ratio guidance. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VIII* 16. Springer, 68–84.
- Zhijian Liu, Haotian Tang, Yujun Lin, and Song Han. 2019. Point-voxel cnn for efficient 3d deep learning. *Advances in neural information processing systems* 32 (2019).
- William E Lorensen and Harvey E Cline. 1998. Marching cubes: A high resolution 3D surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*. 347–353.
- Daniel Marbach, James C Costello, Robert Küffner, Nicole M Vega, Robert J Prill, Diogo M Camacho, Kyle R Allison, Manolis Kellis, James J Collins, and Gustavo Stolovitzky. 2012. Wisdom of crowds for robust gene network inference. *Nature methods* 9, 8 (2012), 796–804.
- Nissim Maruani, Roman Klokov, Maks Ovsjanikov, Pierre Alliez, and Mathieu Desbrun. 2023. Voromesh: Learning watertight surface meshes with voronoi diagrams. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 14565–14574.
- Nissim Maruani, Maks Ovsjanikov, Pierre Alliez, and Mathieu Desbrun. 2024. PoNQ: a Neural QEM-based Mesh Representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3647–3657.
- Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. 2019. Occupancy Networks: Learning 3D Reconstruction in Function Space. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- David A Meyer. 1993. Spherical containment and the Minkowski dimension of partial orders. *Order* 10 (1993), 227–237.
- Charlie Nash, Yaroslav Ganin, SM Ali Eslami, and Peter Battaglia. 2020. Polygen: An autoregressive generative model of 3d meshes. In *International conference on machine learning*. PMLR, 7220–7229.
- Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. 2022. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751* (2022).
- Maximilian Nickel, Xueyan Jiang, and Volker Tresp. 2014. Reducing the rank in relational factorization models by including observable patterns. *Advances in Neural Information Processing Systems* 27 (2014).
- Maximilian Nickel and Douwe Kiela. 2017. Poincaré embeddings for learning hierarchical representations. *Advances in neural information processing systems* 30 (2017).
- Gregory M. Nielson. 2003. On marching cubes. *IEEE Transactions on visualization and computer graphics* 9, 3 (2003), 283–297.
- Werner Palfinger. 2022. Continuous remeshing for inverse rendering. *Computer Animation and Virtual Worlds* 33, 5 (2022), e2101.
- Songyou Peng, Chiyu "Max" Jiang, Yiyi Liao, Michael Niemeyer, Marc Pollefeys, and Andreas Geiger. 2021. Shape As Points: A Differentiable Poisson Solver. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Marie-Julie Rakotosoana, Paul Guerrero, Noam Aigerman, Niloy J Mitra, and Maks Ovsjanikov. 2021. Learning delaunay surface elements for mesh reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 22–31.
- Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J. Black. 2018. Generating 3D Faces using Convolutional Mesh Autoencoders. In *ECCV*.
- Nicholas Sharp and Maks Ovsjanikov. 2020. Pointtrinet: Learned triangulation of 3d point sets. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII* 16. Springer, 762–778.
- Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. 2021. Deep Marching Tetrahedra: a Hybrid Representation for High-Resolution 3D Shape Synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Tianchang Shen, Jacob Munkberg, Jon Hasselgren, Kangxue Yin, Zian Wang, Wenzheng Chen, Zan Gojic, Sanja Fidler, Nicholas Sharp, and Jun Gao. 2023. Flexible Isosurface Extraction for Gradient-Based Mesh Optimization. *ACM Trans. Graph.* 42, 4, Article 37 (jul 2023), 16 pages. <https://doi.org/10.1145/3592430>
- Yawar Siddiqui, Antonio Alliegro, Alexey Artemov, Tatiana Tommasi, Daniele Sirigatti, Vladislav Rosov, Angela Dai, and Matthias Nießner. 2023. Meshgpt: Generating triangle meshes with decoder-only transformers. *arXiv preprint arXiv:2311.15475* (2023).
- Richard Sinkhorn. 1964. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The annals of mathematical statistics* 35, 2 (1964), 876–879.
- Sanghyun Son, Matheus Gadelha, Yang Zhou, Zexiang Xu, Ming C Lin, and Yi Zhou. 2024. DMesh: A Differentiable Representation for General Meshes. *arXiv preprint arXiv:2404.13445* (2024).

- Siddharth Tanwar, Rajiv Ratn Shah, and Roger Zimmermann. 2020. Variational Autoencoders for Deformable 3D Mesh Generation. In *ECCV*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. 2018. Pix2Mesh: 3D Mesh Model Generation via Image Guided Deformation. In *ECCV*.
- Kevin J Weiler. 1986. *Topological structures for geometric modeling (Boundary representation, manifold, radial edge structure)*. Rensselaer Polytechnic Institute.
- Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. 2022. LION: Latent Point Diffusion Models for 3D Shape Generation. In *Advances in Neural Information Processing Systems*.
- Yuxuan Zhang, Wenzheng Chen, Huan Ling, Jun Gao, Yinan Zhang, Antonio Torralba, and Sanja Fidler. 2020. Image gans meet differentiable rendering for inverse graphics and interpretable 3d neural rendering. *arXiv preprint arXiv:2010.09125* (2020).
- Yuxuan Zhang, Huan Ling, Jun Gao, Kangxue Yin, Jean-Francois Lafleche, Adela Barriuso, Antonio Torralba, and Sanja Fidler. 2021. Datasetgan: Efficient labeled data factory with minimal human effort. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10145–10155.
- Qingnan Zhou and Alec Jacobson. 2016. Thingi10k: A dataset of 10,000 3d-printing models. *arXiv preprint arXiv:1605.04797* (2016).
- Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. 2018. Open3D: A modern library for 3D data processing. *arXiv preprint arXiv:1801.09847* (2018).