

Point2Skeleton: Learning Skeletal Representations from Point Clouds

Cheng Lin¹ Changjian Li^{2*} Yuan Liu¹ Nenglu Chen¹ Yi-King Choi¹ Wenping Wang^{1,3*}
¹The University of Hong Kong ²University College London ³Texas A&M University

Abstract

We introduce *Point2Skeleton*, an unsupervised method to learn skeletal representations from point clouds. Existing skeletonization methods are limited to tubular shapes and the stringent requirement of watertight input, while our method aims to produce more generalized skeletal representations for complex structures and handle point clouds. Our key idea is to use the insights of the medial axis transform (MAT) to capture the intrinsic geometric and topological natures of the original input points. We first predict a set of skeletal points by learning a geometric transformation, and then analyze the connectivity of the skeletal points to form skeletal mesh structures. Extensive evaluations and comparisons show our method has superior performance and robustness. The learned skeletal representation will benefit several unsupervised tasks for point clouds, such as surface reconstruction and segmentation.

1. Introduction

Generating skeleton-based representations to capture the underlying shape structures is a classic problem in computer vision and computer graphics. Skeletonization has been shown to benefit various tasks including shape recognition [4, 40], 3D reconstruction [38, 43], segmentation [24], shape matching [34, 36], pose estimation [28, 32], action recognition [21, 33] and animation [5]. Extracting skeletons of 3D shapes using hand-crafted rules [3, 26, 37] has been researched for decades. With the recent advances in 3D vision with deep learning, predicting curve skeletons for 3D shapes using networks [44] is beginning to be studied.

In fact, the existing methods only target a specific category of shapes that can be abstracted appropriately by curve segments, i.e., shapes composed of tubular parts. Also, generating an internal representation, such as the skeleton, for a 3D shape heavily relies on watertight surface meshes to explicitly give inside/outside classification labels and precisely compute certain geometric functions [12, 14]. These

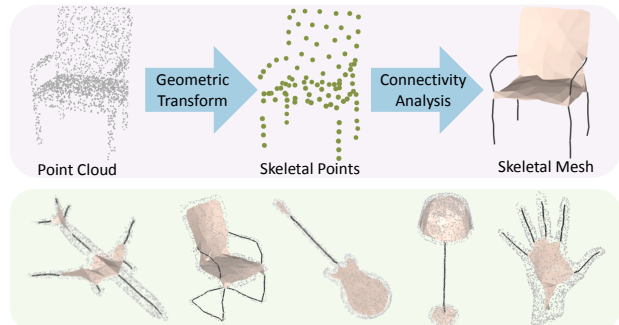


Figure 1. We introduce an unsupervised method to learn skeletal meshes from point clouds. The skeletal meshes contain both 1D curve segments and 2D surface sheets which can represent underlying structures of various shapes.

restrictions critically limit the applicability of the existing skeletonization methods. Hence, it is imperative to develop an effective method for computing a generalized skeletal representation for an arbitrary 3D shape. Such a skeletonization method should be able to handle general input beyond the closed surface, such as point clouds with missing parts, in order to significantly extend the utility of skeleton in various computer vision and graphics tasks.

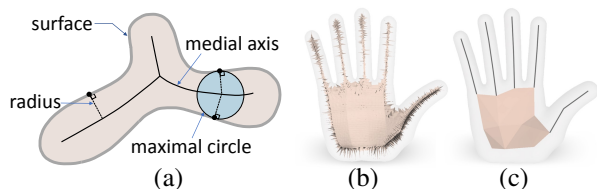


Figure 2. Illustration of the medial axis transformation (MAT) and the skeletal mesh. (a) The MAT of a 2D shape; (b) the original MAT of a 3D shape; (c) the skeletal mesh.

We observe that the medial axis transform (MAT), one of the best known examples of skeletal representation, has a rigorous mathematical definition for arbitrary shapes, unlike the curve skeleton which is only empirically understood for tubular objects. Given a 3D shape, the MAT [8] is defined as the set of points in the interior with more than one closest point on the boundary surface; the MAT encodes the shape to a lower-dimensional representation with the associated radius function. The examples are shown in Fig. 2. Despite its simple definition, the MAT is difficult to use

* corresponding authors

in practice for the following two reasons: (1) the computation of the MAT is expensive, since it requires the input 3D shape to be defined by a closed boundary surface and a substantial amount of time for geometric processing; (2) the MAT is notoriously sensitive to surface noise, i.e., small perturbations to the shape surface lead to numerous insignificant branches [2, 15] (Fig. 2 (b)); such a noisy MAT does not clearly reflect the structures of the given shape.

These difficulties motivate us to resort to the formulation of the MAT and the representational power of the deep neural network to learn a generalized skeletal representation. We name this representation *skeletal mesh*, which is an important extension to the curve skeleton. As shown in Fig 2, the skeletal mesh follows a similar but not identical definition to the MAT; it circumvents the drawbacks of the MAT and has its own merits: (1) it is structurally meaningful and topologically informative, where the tubular parts can be properly abstracted by a few curves while the planar or bulky parts by interior surfaces; (2) the skeletal mesh is simpler and more compact than the standard MAT; it focuses on the fundamental geometry of a shape and can cope with point clouds, making it robust to surface noise and partially missing data. Therefore, given such a representation, the expressive capacity for the geometry and topology of a 3D shape is significantly enhanced.

In this paper, we propose *Point2Skeleton*, an unsupervised method for learning skeletal meshes from 3D point clouds. Our method consists of two main steps as shown in Fig 1. The first step is to predict the skeletal points by learning a geometric transformation. The second step is to connect the skeletal points to form a mesh structure; we adopt a graph structure, and analyze the edge connectivity by jointly leveraging the properties of skeletal mesh and the correlations learned by a graph auto-encoder (GAE). Our main contributions are:

- To our best knowledge, *Point2Skeleton* is the first unsupervised learning method for generalized point cloud skeletonization.
- We present novel unsupervised formulations for geometric learning of 3D point clouds, i.e., learning intrinsic geometric transformations and predicting connectivity for mesh generation.
- We introduce a new representation, called *skeletal mesh*, which gives new insights into some unsupervised tasks for point clouds, such as surface reconstruction and segmentation.

2. Related Work

Skeletonization and medial axis transform. Skeleton-based shape representation has been extensively researched in computer vision and computer graphics. A widely-used form is the curve skeleton due to its simplicity. Traditional methods use hand-crafted rules to capture geometric proper-

ties for generating curve skeletons [3, 9, 17, 26, 37]. Recently, attempts have been made to use deep neural networks to predict curve skeletons [44]. The curve skeletons are only empirically understood for tubular geometries, so they can only be applied to a limited class of shapes.

Another form of skeletal representation is known as the medial axis transform (MAT) [8], which is a principled formulation that can encode arbitrary shapes. There are numerous methods for computing the MAT of 3D shapes, such as [1] [2] and [15]. However, the MAT is notoriously sensitive to surface noise, which usually leads to numerous insignificant spikes. Therefore, some methods use simplification techniques to produce a clean and structurally simple MAT [23, 35, 45]. These methods all need expensive geometric processing and require watertight input surfaces. There are also some rule-based methods for generating meso-skeletons [37, 43] to approximate the MAT. We use deep learning techniques to efficiently predict the clean skeletons based on the insights of the MAT, and also show better performance than the other related methods.

Learning transformations. Learning spatial transforms in the image domain is an important problem. Representative works on this problem include spatial transformer networks (STN) [19] and view synthesis by appearance flow [48], to name a few. However, learning geometric transforms in the 3D domain is still under-explored. Berkiten et al. [6] propose to learn a combination of geometric features to derive the transform of surface details. More relevantly, Yin et al. [47] introduce a deep neural framework to learn geometric transforms between two domains of 3D points. Yang et al. [46] explicitly predict the medial axis from an input point cloud. These two methods are able to learn geometric transformations, but the learning is driven by the pre-computed ground-truth points rather than capturing the intrinsic geometric natures. Also, the learned point-to-point transformation neither contains structured surfaces nor gives topological connectivity.

Inferring mesh structures. There has been growing interest in inferring mesh structures of 3D shapes using deep neural networks. Due to the difficulty of directly generating such compact structures, many approaches generate target meshes by deforming the pre-defined mesh templates [25, 41]. These methods lack the flexibility of handling complex geometries and typologies. Recently, learning implicit representations becomes popular. The common form is to first learn implicit volumetric functions to represent a 3D shape, such as signed-distance functions [29], indicator functions [27] and structured implicit functions [14], and then extract the surface. Some methods aim to directly predict mesh structures from scratch. Scan2Mesh [13] learns to explicitly generate vertices, edges and faces using graph neural networks. PointTriNet [31] explicitly learns the triangulation of a given point set to form the mesh.

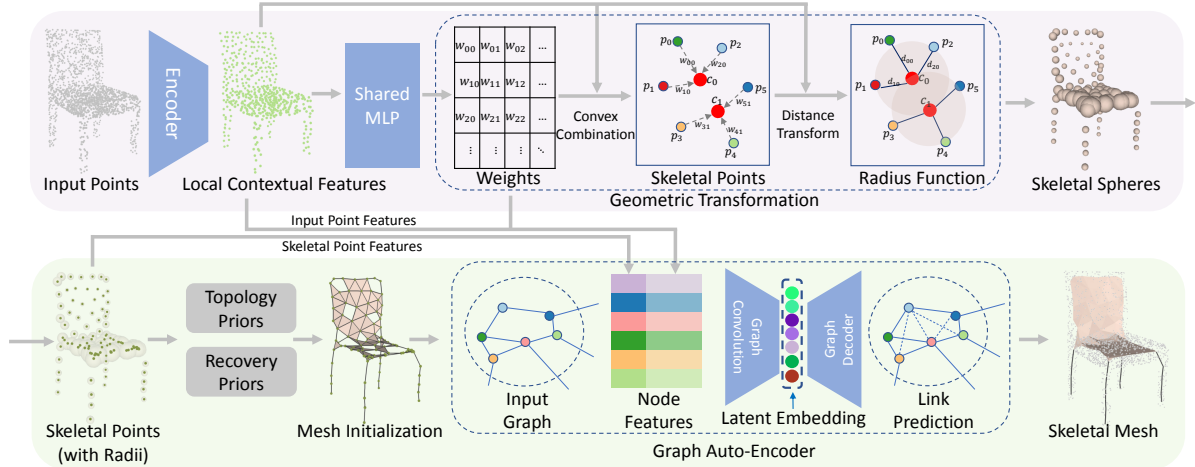


Figure 3. Overview of our Point2Skeleton pipeline. Given a point cloud as input, first, we learn a geometric transformation via convex combinations to predict the skeletal points together with their radii. Second, we connect the skeletal points to form a mesh structure; we initialize a graph structure using two simple priors and formulate a link prediction problem using a graph auto-encoder to obtain a complete skeletal mesh.

Most of these methods rely on ground-truth surfaces for reliable supervision; none of them can be used to predict skeleton representations that contain curves. In contrast, our method is directly trained on point clouds, aiming to effectively generate skeletal meshes in an unsupervised manner.

3. Skeletal Mesh

The skeletal mesh of a given 3D shape is a discrete 2D non-manifold defined by a collection of skeletal points, edges and faces that form the underlying structure of the 3D shape, as shown in Fig. 4 (a). A *skeletal sphere* is required to be maximally inscribed in the 3D shape, denoted as $s = (c, r(c)) \in \mathbb{R}^4$, where $c \in \mathbb{R}^3$ is the center of the sphere (i.e., the *skeletal point*), and $r(c) : \mathbb{R}^3 \rightarrow \mathbb{R}$ is the associated *radius*. We use $e_{ij} = (c_i, c_j)$ to represent an edge of the skeletal mesh connecting two skeletal points c_i and c_j , and use $f_{ijk} = (c_i, c_j, c_k)$ to represent a triangle face.

The skeletal mesh is a discrete representation while the MAT is continuously defined. The discrete nature of the skeletal mesh allows us to develop a learning-based approach to robustly compute the skeletal representation, such that it inherits the favorable properties of the MAT for shape representation while not suffering from the instability to boundary noise [2, 15]. Several good properties of the skeletal mesh are elaborated as follows, which makes it a useful representation for shape analysis (see Sec. 6).

Recoverability. The skeletal mesh can be considered as a complete shape descriptor, which means it can reconstruct the shape of the original domain. Fig. 4 (b)(c) shows how to reconstruct the original shape by the interpolation of all the skeletal spheres based on the mesh structure.

Abstraction. The skeletal mesh captures the fundamental

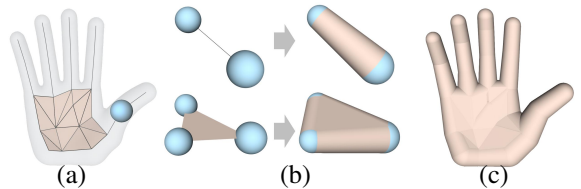


Figure 4. Properties of skeletal mesh. (a) A skeletal mesh structure of a shape; (b) interpolation of the skeleton spheres of an edge (top) and a triangle face (bottom); (c) the interpolation of all the spheres on the skeletal mesh reconstructs the original shape.

geometry of a 3D shape and extracts its global topology; the tubular parts are abstracted by simple 1D curve segments and the planar or bulky parts by 2D surface triangles.

Structure awareness. The 1D curve segments and 2D surface sheets as well as the non-manifold branches on the skeletal mesh give a structural differentiation of a shape.

Volume-based closure. The interpolation of the skeletal spheres gives solid cone-like or slab-like primitives (see Fig. 4 (b)); then a local geometry is represented by volumetric parts, which provides better integrity of shape context. The interpolation also forms a closed watertight surface.

4. Method

Our method is composed of two modules. As shown in Fig. 3, given a point cloud as input, the first module is to predict a set of skeletal points by learning a geometric transformation (Sec. 4.1). The second module is to connect the skeletal points to form a skeletal mesh, where we resort to a graph structure to analyze the edge connectivity (Sec. 4.2).

4.1. Skeletal Point Prediction

Given K input points $\{p_i\}$ with 3D coordinates, represented as $P \in \mathbb{R}^{K \times 3}$, our goal is to predict N skeletal spheres $\{s_i = (c_i, r(c_i))\}$, represented as $S \in \mathbb{R}^{N \times 4}$,

consisting of the coordinates of the centers $\mathbf{C} \in \mathbb{R}^{N \times 3}$ and their radii $\mathbf{R} \in \mathbb{R}^{N \times 1}$. To start with, we use PointNet++ [30] as the encoder to obtain the sampled input points $\mathbf{P}' \in \mathbb{R}^{K' \times 3}$ and their contextual features $\mathbf{F} \in \mathbb{R}^{K' \times D}$, where $K' < K$ and D is the dimension of the feature.

Convex combination. We observe a skeletal point can be considered as a local center of a set of surface points. Consequently, instead of directly predicting the coordinates of the skeletal points, we use the convex combination of input points to generate the skeletal points. To this end, we predict the weights $\mathbf{W} \in \mathbb{R}^{K' \times N}$ of the sampled input points \mathbf{P}' in the convex combination and all the skeletal points \mathbf{C} are derived by:

$$\mathbf{C} = \mathbf{W}^T \mathbf{P}' \quad \text{s.t. } j=1, \dots, N \quad \sum_{i=1}^{K'} \mathbf{W}(i, j) = 1. \quad (1)$$

Now we explain how to estimate the radius of each skeletal sphere. First, the closest distance from an input point p to all the skeleton points $\{c_i\}$ is defined as follows:

$$d(p, \{c_i\}) = \min_{c \in \{c_i\}} \|p - c\|_2. \quad (2)$$

The distances for all K' sampled input points p_i computed by Eq. 2 are summarized in a vector $\mathbf{D} \in \mathbb{R}^{K' \times 1}$. Now we express the radii of all the skeletal points as the linear combinations of their closest distances of all the sampled points, i.e., $\mathbf{R} = \mathbf{W}^T \mathbf{D}$. The rationale of using the same weights \mathbf{W} as Eq. 1 is based on the fact that the predicted weights for a skeletal point c are large only for the input points close to c , but diminish to zero for those far from c . The reader can refer to the supplementary material for a detailed argument.

The advantages of using the weights of convex combination to predict the skeletal points lie in: (1) the generated points tend to be located inside the local shape without using explicit inside/outside labels given by a watertight boundary surface; (2) the computed combinational weights \mathbf{W} can be considered as linear filters which are effective for denoising when computing the radius.

Similar to [11], we use a shared multi-layer-perceptron (MLP) followed by a softmax layer to produce the weights \mathbf{W} . By capturing the properties of the skeletal mesh, we are able to design a set of loss functions to train the network to obtain the combinational weights \mathbf{W} .

Based on the recoverability of the skeletal mesh, we first introduce two loss functions to measure the reconstruction error from two different perspectives.

Sampling loss. We sample points on the surface of each skeletal sphere and measure the Chamfer Distance (CD) between the sampled points $\{t_i\}$ and the input points $\{p_i\}$:

$$L_s = \sum_{p \in \{p_i\}} \min_{t \in \{t_i\}} \|p - t\|_2 + \sum_{t \in \{t_i\}} \min_{p \in \{p_i\}} \|t - p\|_2. \quad (3)$$

Here we adopt a uniform sampling strategy. For a skeletal sphere $(c, r(c))$, a surface point t can be obtained by $t = c + r(c)v$ given a unit direction vector v ; we sample 8 points for each skeletal sphere by giving 8 unit vectors $V = \{(\pm\eta, \pm\eta, \pm\eta), 3\eta^2 = 1\}$.

Point-to-sphere loss. In addition to sampling, we introduce a loss function that measures the reconstruction error by explicitly optimizing the coordinates of the skeletal points and their radii:

$$L_p = \sum_{p \in \{p_i\}} (\min_{c \in \{c_i\}} \|p - c\|_2 - r(c_p^{min})) + \sum_{c \in \{c_i\}} (\min_{p \in \{p_i\}} \|c - p\|_2 - r(c)), \quad (4)$$

where $\{c_i\}$ represents the predicted skeletal points, $\{p_i\}$ the input points, $r(c)$ the radius of a skeletal point c , and c_p^{min} the closest skeletal point to the input point p . The first term constrains each input point to be located on the surface of its closest skeletal sphere, while the second term encourages each skeletal sphere to touch its closest input point.

These two complementary losses optimize the consistency between the skeleton and the input point cloud from different aspects, leading to more reliable predictions (see the evaluations in Sec. 5.2).

Radius regularizer. An inscribed sphere of a 3D shape is sensitive to surface noise; that is, a sphere can be stuck in the space among several noisy points thus resulting in a tiny radius. To focus on the fundamental geometry for better abstraction and avoid instability, therefore, we also include a radius regularization loss to encourage larger radii:

$$L_r = - \sum_{c \in \{c_i\}} r(c). \quad (5)$$

Instead of directly being predicted by the network, the radii are computed using a linear combination of the distances in Eq. 2 with the weights \mathbf{W} . Thus, the radius values are bounded to a certain range, which prevents them from growing too large and giving a very small value of L_r to destroy the optimization. The total loss function for predicting skeletal spheres is a weighted combination of these three terms:

$$L_{skel} = L_s + \lambda_1 L_p + \lambda_2 L_r. \quad (6)$$

4.2. Skeletal Mesh Generation

In this section, we aim to connect the predicted skeletal points to form skeletal meshes. Our mesh generation process is built on the edge connectivity of the skeletal points, which can encode both curve segments and face triangles. We thus resort to a graph representation for this problem.

Graph initialization. Based on the properties of the skeletal mesh, we first initialize a set of reliable links using two simple priors:

- **Topology prior:** A node has a link to its closest node (marked as *known existing links*) but no links to its k -farthest nodes (*known absent links*).
- **Recovery prior:** Given an input point p with its two closest skeletal points c_1, c_2 , there will be a link connecting (c_1, c_2) (*known existing links*).

The topology prior assumes the skeletal mesh is relatively regular. The recovery prior ensures each input point can be properly reconstructed by the skeletal mesh. Based on these reliable priors, we obtain an initialized graph with *known existing* and *known absent* links, and we mark the other links as *unknown*.

Link prediction by GAE. The initialized graph is usually incomplete. Therefore, we formulate a link prediction problem based on the graph auto-encoder (GAE) [22], which is an unsupervised formulation, to further predict the missing links by analyzing the correlations of the skeletal points in the latent space. As shown in Fig. 5, the GAE predicts additional links for the initialized graph, enabling our method to handle missing regions and noise of the input points.

The input to the encoder is the initialized graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ with $N = |\mathcal{N}|$ nodes representing N skeletal points and \mathcal{E} the edge connectivity. The graph \mathcal{G} is undirected and unweighted, represented by an adjacency matrix $\mathbf{A} \in \{0, 1\}^{N \times N}$. To effectively leverage the geometric correlations between the skeleton and the input surface points, the node features are jointly characterized by the coordinates \mathbf{C} of the skeleton points, their radii \mathbf{R} and the contextual features \mathbf{F} of the input, which is denoted as $[\mathbf{C}, \mathbf{R}, \mathbf{W}^T \mathbf{F}]$, where $[\cdot, \cdot, \cdot]$ represents the concatenation on the feature dimension and $\mathbf{W}^T \mathbf{F}$ combines the contextual features to the skeletal points using the predicted weights \mathbf{W} .

The encoder is a series of graph convolutional network (GCN) layers with residual blocks [16] between the consecutive layers. We use an inner product decoder to produce the reconstructed adjacency matrix $\hat{\mathbf{A}}$. Please refer to the supplementary material to find more details.

Learning. We define a reconstruction loss that captures the similarity of the reconstructed $\hat{\mathbf{A}}$ and the initialized \mathbf{A} :

$$L_{link} = \text{mean}(\mathcal{M} \odot (-\xi \mathbf{A} \log(\sigma(\hat{\mathbf{A}})) - (1 - \mathbf{A}) \log(\sigma(1 - \hat{\mathbf{A}})))) \quad (7)$$

This is referred to as a Masked Balanced Cross-Entropy (MBCE) loss [39]. Here ξ represents the balanced weight, which is the ratio of the amount of *absent* links to the *existing* links; $\mathcal{M} \in \mathbb{R}^{N \times N}$ is a binary mask indicating whether a link between two nodes is *known* or *unknown* (as marked in the graph initialization priors), by which we only consider the *known* links in the back-propagation; \odot is the element-wise product and $\sigma(\cdot)$ represents the Sigmoid function.

The capability of the GAE to predict missing links is attributed to the encoded meaningful latent embeddings that can reconstruct the graph. The decoder establishes the links by measuring the correlations of the latent features; thus the

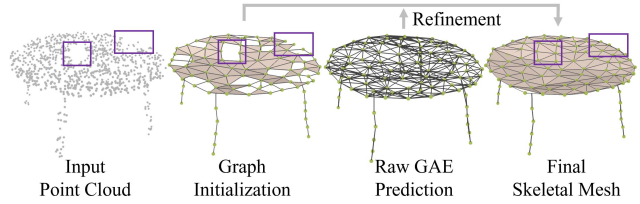


Figure 5. Skeletal mesh generation. Given the initialized graph, the GAE predicts potential links by measuring the correlations of the skeletal points in the latent embedding, based on which we refine the initialized graph to derive the final skeletal mesh.

nodes with unknown connectivity will form links if they exhibit strong correlations in the latent space.

Mesh generation. Once the training is finished, we do not directly use the raw output of GAE as the final skeletal mesh, since it will lead to many redundant triangles and unclear structures. Instead, as shown in Fig. 5, we use the GAE prediction to refine the initial graph, which mainly contains two steps. The first is *hole filling*. We extract all the triangle faces on the initial graph, and then detect the small polygonal loops that do not form face triangles on the remaining initial graph. A polygonal loop hole will be filled if the links predicted by the GAE show the hole does not exist. To explicitly consider the geometric recoverability for more reliable mesh generation, we also fill a hole if there are enough points on its corresponding area of the input points. The second is *boundary refinement*. We extract the skeletal points on the boundary and connect two points if the GAE prediction shows they are connected. With the predicted mesh, we can re-compute the closest distances based on Eq. 2 and then obtain a more accurate radius estimation.

5. Experimental Results

Dataset. We collect 7088 shapes from 8 categories of ShapeNet [10], and use a virtual scanner [42] to generate the point clouds (2000 points are randomly sampled for each shape). We use 5/6 of the data from each category for training and the other 1/6 for testing. The network is trained on all the shape categories jointly.

Implementation details. We use a PointNet++ [30] encoder of 4 set abstraction levels to obtain the contextual features; then the features are processed by 5 shared MLPs to predict $N=100$ skeletal points for each shape. For the GAE, we use a 12 layers GCN with residual blocks. More details and evaluations are given in the supplementary material.

For the training process, the skeletal point prediction network is first pre-trained only by the CD loss between the skeletal points and the input points for 20 epochs. This is to make the skeletal points evenly distributed around the shape, leading to a better initialization. Then, the network is trained using L_{skel} (with $\lambda_1 = 0.3$ and $\lambda_2 = 0.4$) for 30 epochs. After that, we freeze the skeletal point prediction network and train the GAE for 30 epochs. The learning




Figure 6. Qualitative comparison with the competitive point cloud skeletonization methods, i.e., L_1 -medial skeleton [17] and deep point consolidation (DPC) [43].

rates for the two networks are $1e^{-3}$ and $5e^{-4}$ respectively and the batch size is 4.

5.1. Results and Comparisons

Evaluation metrics. We evaluate our method and the competitive skeletonization methods from two aspects. The first is the reconstruction quality for the original shape. We measure the difference between the shapes reconstructed from the skeletons and the ground truth shapes using the Chamfer distance (CD) and Hausdorff distance (HD), which are denoted as *CD-Recon* and *HD-Recon*, respectively.

Note that only the reconstruction accuracy alone is not enough to completely reflect the quality of the skeletonization. For example, a poorly designed algorithm that directly outputs the input points will yield small reconstruction error, but do not constitute a reasonable skeleton. Unfortunately, there are no existing metrics to evaluate if a skeletonization is reasonable, since the skeleton is a form of shape abstraction involving higher-level perceptions. Nevertheless, we observe a manually simplified MAT, where the unstable spikes are removed, not only shares some common characteristics with different forms of skeletons, but also has good geometric accuracy. A simplified MAT contains both curve-like and surface-like geometries; thus it can be used to evaluate the methods that are based on curves or beyond the curves.

We repair the ground truth meshes of the ShapeNet, convert them to watertight meshes [18] and compute the strictly defined MATs. We manually simplify the MATs using the handcrafted methods [23] by removing the spikes to make them visually simple and clean (see the supplementary material for more details). Again, we use the CD and HD to measure the difference between the output skeletal representations and the simplified MAT, denoted as *CD-MAT* and *HD-MAT*. All these distances are computed by randomly sampled points from the respective geometries.

Comparisons. Since learning-based skeletonization is barely studied, we evaluate our method with comparisons to closely relevant approaches that are not learning-based,

i.e., L_1 -medial skeleton [17] and deep point consolidation (DPC) [43]. We use the code released by the authors for comparison. Similar to ours, these methods are designed for the skeletonization of an arbitrary 3D shape given as a point cloud. Fig. 6 shows the qualitative comparisons with these methods. The L_1 -medial skeleton [17] can produce structured representations that only contain 1D curves, thus resulting in large errors when used to abstract non-tubular shapes. The DPC [43] method can generate both surface-like and curve-like skeletons, but the representations are unstructured points that lack topological constraints, leading to inconsistency for thin structures. In contrast, our method can generate more compact and structurally meaningful skeletal representations for various geometries.

As shown in Table 1, we quantitatively compare our method with these two methods using the four metrics aforementioned. The results show that our method not only more accurately encodes the information from the original input, but also produces more reasonable skeletonization results that are geometrically meaningful.

5.2. Discussions

We conduct a series of ablation studies to verify the various settings in our framework, and also some experiments to further explore the properties of the proposed method.

Skeletal point prediction. We evaluate the effect of the geometric constraints in the loss function (Eq. 6) for skeletal point prediction. We alternatively remove each constraint and analyze the quality of the predicted skeletal points. The qualitative and quantitative results are shown in Fig. 7 and Table 2, respectively. Under the effect of the radius regularizer that encourages larger radii, only the point-to-sphere loss struggles to constrain the skeletal spheres inside the shape without using sampling loss. The removal of the point-to-sphere loss results in inaccurate radii and uneven distribution of spheres, while the removal of the radius regularizer leads to spheres that are scattered instead of being maximally inscribed. Using the full configuration better captures the inner structures and achieves higher accuracy.

	CD-Recon			HD-Recon			CD-MAT			HD-MAT		
	L1	DPC	Ours	L1	DPC	Ours	L1	DPC	Ours	L1	DPC	Ours
Airplane	0.0378	0.0348	0.0363	0.2216	0.1436	0.1266	0.0793	0.1307	0.0611	0.2384	0.2580	0.1721
Chair	0.1126	0.0769	0.0441	0.4810	0.2478	0.1618	0.1885	0.2286	0.0974	0.4991	0.3707	0.2151
Table	0.1041	0.0853	0.0424	0.3453	0.2584	0.1745	0.1541	0.2683	0.0876	0.3583	0.3690	0.2085
Lamp	0.1542	0.0712	0.0335	0.3956	0.1850	0.1382	0.1870	0.1751	0.0884	0.4089	0.2627	0.2003
Guitar	0.0655	0.0212	0.0179	0.2180	0.0589	0.0625	0.0817	0.0672	0.0536	0.2262	0.1226	0.1216
Earphone	0.0437	0.0573	0.0399	0.1908	0.2059	0.1125	0.0607	0.2216	0.1638	0.1732	0.3403	0.2130
Mug	0.2864	0.1280	0.0417	0.9142	0.3510	0.1419	0.5316	0.4600	0.1179	0.9057	0.4308	0.2158
Rifle	0.0260	0.0215	0.0213	0.1078	0.0702	0.0767	0.0494	0.0427	0.0356	0.1234	0.1050	0.0957
Average	0.1038	0.0668	0.0372	0.3593	0.2049	0.1424	0.1665	0.2026	0.0828	0.3667	0.3047	0.1898

Table 1. Quantitative comparison with the competitive point cloud skeletonization methods.

	CD-Recon	HD-Recon	CD-MAT	HD-MAT
w/o sampling(L_s)	0.3032	0.4886	0.7319	0.7848
w/o point2sphere(L_p)	0.0535	0.1619	0.1102	0.2236
w/o radius(L_r)	0.0634	0.2206	0.1276	0.2517
full configuration	0.0525	0.1592	0.1060	0.2079

Table 2. Quantitative ablation study using different configurations of the geometric constraints for skeletal point prediction.

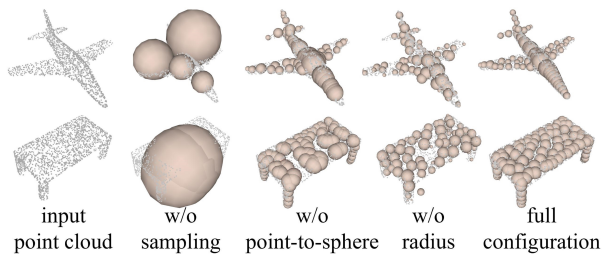


Figure 7. Qualitative ablation study on different configurations of the geometric constraints for skeletal point prediction.

Mesh generation. Given a set of discrete points, there are some frequently used methods to generate meshes in the same manner as our method, i.e., connecting the vertices to form mesh structures. The alternative methods tested include Ball pivoting [7], Delaunay triangulation (deleting the triangle faces with overlong edges using a threshold) and K Nearest Neighbor (KNN) (connecting the K nearest neighbors of each point and extracting the formed triangles). The qualitative and quantitative results are shown in Fig. 8 and Table 3 (the first three rows). Our strategy for connecting points into a skeletal mesh is significantly better than the alternative methods, thanks to the joint use of the properties of skeletal mesh and the point correlations in latent embeddings; the other methods cannot preserve the typologies and thus are not suitable for generating skeletal meshes.

As an ablation study, in Table 3, we also report the errors without using the graph initialization (only skeletal spheres) and without using the GAE (only graph initialization), respectively. The results demonstrate the improvements brought about by the learned latent embeddings over the raw output of skeletal spheres and the initialization.

Noise and sparsity. To study the effect of the input quality on our algorithm, we evaluate our method using the input of different point numbers and noise levels, i.e., 2000 points with no noise, 1000 points with 0.5% noise and 500

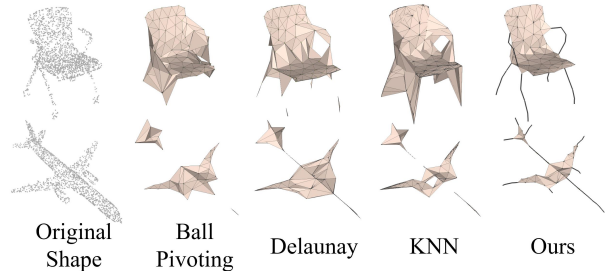


Figure 8. Qualitative results of alternative mesh generation methods by connecting vertices to form mesh structures.

	CD-Recon	HD-Recon	CD-MAT	HD-MAT
Delaunay	0.0388	0.3313	0.1196	0.3386
Ball Pivoting	0.0443	0.3903	0.1095	0.3965
KNN	0.0394	0.2811	0.0946	0.3243
w/o graph	0.0525	0.1592	0.1060	0.2079
w/o GAE	0.0383	0.1462	0.0842	0.1886
w/ GAE	0.0372	0.1424	0.0828	0.1898

Table 3. Quantitative comparisons with different methods for mesh generation by connecting vertices (first 3 rows); ablation study for the mesh generation (last 3 rows).

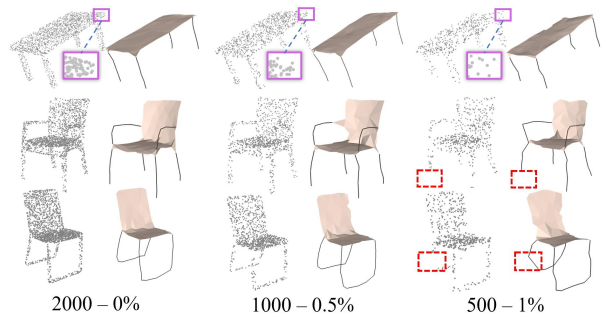


Figure 9. Qualitative evaluation on the effect of the quality of input point cloud, i.e., density (point number) and noise.

points with 1% noise. The qualitative results are shown in Fig. 9 and more results can be found in the supplementary material, where we can find the results are largely similar.

However, since we use the convex combination of the input points to generate the skeletal points, if the input points are too sparse and some points on the convex hull are missing, our method struggles to recover the complete geometry (see the red boxes in Fig. 9).

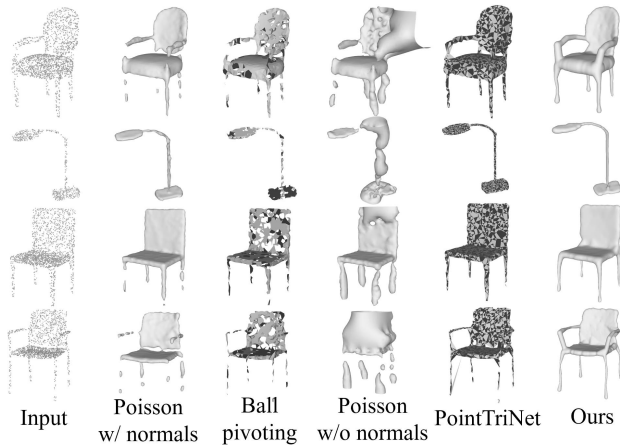


Figure 10. Unsupervised surface reconstruction from point clouds. Our method can produce watertight surfaces without the need of surface normals and can capture the details of thin structures.

6. More Applications

Some good properties (see Sec. 3) of the skeletal mesh, i.e., recoverability, abstraction, structure-awareness and volume-based closure, provide new insights into several unsupervised tasks on point clouds. In this section, we demonstrate the benefits to four such applications.

Surface reconstruction. Generating watertight surfaces from point clouds for thin structures and preserving the original typologies, especially when there is no supervision from ground-truth, is a highly challenging problem. The skeletal mesh provides a suitable vehicle for solving this problem, since we can reconstruct the surface of an input point cloud by interpolating the skeletal spheres. We qualitatively compare to some unsupervised surface reconstruction methods, including Ball pivoting [7], Screened Poisson reconstruction [20] and PointTriNet [31]. The results are shown in Fig. 10. On the one hand, the reconstructions using skeletal meshes preserve the complex typologies and also capture the thin structures of the input. On the other hand, unlike the Poisson reconstruction, our method does not need to input any surface normal information, and is still able to produce high-quality watertight surfaces.

Structural decomposition. The skeletal mesh is structure-aware as introduced in Sec. 3; thus it naturally induces a structural decomposition of a shape without data annotation. Fig. 11 shows a set of segmentation results by detecting the dimensional changes and the non-manifold branches on the skeletal meshes.

Shape in-painting. The skeletal mesh encodes volume-based shape context which provides better integrity of part geometry, leading to robustness against partially missing data. Thus we can recover the complete geometry from an incomplete point cloud. Some results are given in Fig. 12.

Skeletonization with consistent correspondence. Skeletons play a vital role in pose recognition and animation, in which a challenge is to find the correspondences between

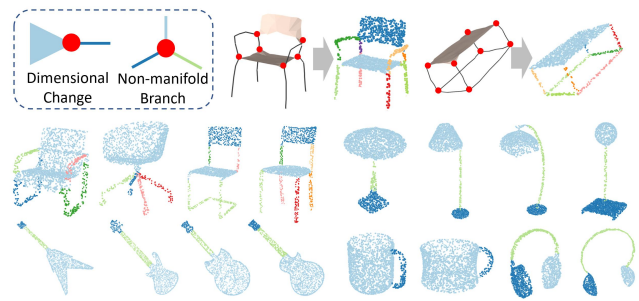


Figure 11. Unsupervised structural decomposition for point clouds by detecting dimensional changes and non-manifold branches on the skeletal mesh.

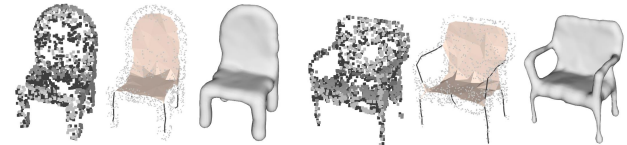


Figure 12. Skeletal mesh prediction for reconstructing complete surfaces from point clouds with missing regions.



Figure 13. Skeletal meshes with consistent correspondences. We color the skeletal points of each shape individually based on their indices in each prediction; some key skeletal points are highlighted to show the semantic consistency.

different poses of a human or an animal. Fig. 13 shows the skeletal meshes generated for a sequence of horse models in different poses. Without explicitly enforcing the correspondence, the results for different poses are semantically consistent due to the learning of convex combination (see [11] for the detailed explanation). This would be beneficial for tasks that require pose invariance in animation.

7. Conclusion

We propose *Point2Skeleton*, a novel unsupervised method for generating skeletal meshes from point clouds. We first predict skeletal points by learning a geometric transformation, and then analyze the connectivity of the skeletal points to form meshes. Experiments show that the skeletal mesh generated by our method effectively captures the underlying structures for general 3D shapes, even when represented as point clouds with noises or missing parts.

We believe our method for learning skeletons can benefit a variety of 3D applications, given the good properties of the skeletal mesh. In the future, how to combine the geometric and topological properties of the skeletal mesh and higher-level supervision from humans (e.g., semantics) for 3D learning tasks, would be a promising direction.

Acknowledgement. We thank the anonymous reviewers for their valuable feedback. This work was supported by the General Research Fund (GRF) of the Research Grant Council of Hong Kong (17210718).

References

- [1] Nina Amenta and Marshall Bern. Surface reconstruction by voronoi filtering. *Discrete & Computational Geometry*, 22(4):481–504, 1999. 2
- [2] Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust. In *Proceedings of the sixth ACM symposium on Solid modeling and applications*, pages 249–266, 2001. 2, 3
- [3] Oscar Kin-Chung Au, Chiew-Lan Tai, Hung-Kuo Chu, Daniel Cohen-Or, and Tong-Yee Lee. Skeleton extraction by mesh contraction. *ACM transactions on graphics (TOG)*, 27(3):1–10, 2008. 1, 2
- [4] Xiang Bai, Xinggong Wang, Longin Jan Latecki, Wenyu Liu, and Zhuowen Tu. Active skeleton for non-rigid object detection. In *2009 IEEE 12th International Conference on Computer Vision*, pages 575–582. IEEE, 2009. 1
- [5] Ilya Baran and Jovan Popović. Automatic rigging and animation of 3d characters. *ACM Transactions on graphics (TOG)*, 26(3):72–es, 2007. 1
- [6] Sema Berkiten, Maciej Halber, Justin Solomon, Chongyang Ma, Hao Li, and Szymon Rusinkiewicz. Learning detail transfer based on geometric features. In *Computer Graphics Forum*, volume 36, pages 361–373. Wiley Online Library, 2017. 2
- [7] Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Cláudio Silva, and Gabriel Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE transactions on visualization and computer graphics*, 5(4):349–359, 1999. 7, 8
- [8] Harry Blum. A transformation for extracting new descriptors of shape. *Models for Perception of Speech and Visual Forms, 1967*, pages 362–380, 1967. 1, 2
- [9] Junjie Cao, Andrea Tagliasacchi, Matt Olson, Hao Zhang, and Zhinxun Su. Point cloud skeletons via laplacian based contraction. In *2010 Shape Modeling International Conference*, pages 187–197. IEEE, 2010. 2
- [10] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 5
- [11] Nenglu Chen, Lingjie Liu, Zhiming Cui, Runnan Chen, Duygu Ceylan, Changhe Tu, and Wenping Wang. Unsupervised learning of intrinsic structural representation points. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9121–9130, 2020. 4, 8
- [12] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019. 1
- [13] Angela Dai and Matthias Nießner. Scan2mesh: From unstructured range scans to 3d meshes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5574–5583, 2019. 2
- [14] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser. Learning shape templates with structured implicit functions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7154–7164, 2019. 1, 2
- [15] Joachim Giesen, Balint Miklos, Mark Pauly, and Camille Wormser. The scale axis transform. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*, pages 106–115, 2009. 2, 3
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [17] Hui Huang, Shihao Wu, Daniel Cohen-Or, Minglun Gong, Hao Zhang, Guiqing Li, and Baoquan Chen. L1-medial skeleton of point cloud. *ACM Trans. Graph.*, 32(4):65–1, 2013. 2, 6
- [18] Jingwei Huang, Hao Su, and Leonidas Guibas. Robust watertight manifold surface generation method for shapenet models. *arXiv preprint arXiv:1802.01698*, 2018. 6
- [19] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *Advances in neural information processing systems*, pages 2017–2025, 2015. 2
- [20] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics (ToG)*, 32(3):1–13, 2013. 8
- [21] Qihong Ke, Mohammed Bennamoun, Senjian An, Ferdous Sohel, and Farid Boussaid. A new representation of skeleton sequences for 3d action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3288–3297, 2017. 1
- [22] Thomas N Kipf and Max Welling. Variational graph autoencoders. *NIPS Workshop on Bayesian Deep Learning*, 2016. 5
- [23] Pan Li, Bin Wang, Feng Sun, Xiaohu Guo, Caiming Zhang, and Wenping Wang. Q-mat: Computing medial axis transform by quadratic error minimization. *ACM Transactions on Graphics (TOG)*, 35(1):1–16, 2015. 2, 6
- [24] Cheng Lin, Lingjie Liu, Changjian Li, Leif Kobbelt, Bin Wang, Shiqing Xin, and Wenping Wang. Seg-mat: 3d shape segmentation using medial axis transform. *IEEE Transactions on Visualization and Computer Graphics*, 2020. 1
- [25] Or Litany, Alex Bronstein, Michael Bronstein, and Ameesh Makadia. Deformable shape completion with graph convolutional autoencoders. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1886–1895, 2018. 2
- [26] Wan-Chun Ma, Fu-Che Wu, and Ming Ouhyoung. Skeleton extraction of 3d objects with radial basis functions. In *2003 Shape Modeling International.*, pages 207–215. IEEE, 2003. 1, 2
- [27] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019. 2
- [28] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016. 1

- [29] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 2
- [30] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in neural information processing systems*, pages 5099–5108, 2017. 4, 5
- [31] Nicholas Sharp and Maks Ovsjanikov. Pointtrinet: Learned triangulation of 3d point sets. *arXiv preprint arXiv:2005.02138*, 2020. 2, 8
- [32] Jamie Shotton, Andrew Fitzgibbon, Mat Cook, Toby Sharp, Mark Finocchio, Richard Moore, Alex Kipman, and Andrew Blake. Real-time human pose recognition in parts from single depth images. In *CVPR 2011*, pages 1297–1304. Ieee, 2011. 1
- [33] Chenyang Si, Ya Jing, Wei Wang, Liang Wang, and Tieniu Tan. Skeleton-based action recognition with spatial reasoning and temporal stack learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 103–118, 2018. 1
- [34] Kaleem Siddiqi, Ali Shokoufandeh, Sven J Dickinson, and Steven W Zucker. Shock graphs and shape matching. *International Journal of Computer Vision*, 35(1):13–32, 1999. 1
- [35] Feng Sun, Yi-King Choi, Yizhou Yu, and Wenping Wang. Medial meshes—a compact and accurate representation of medial axis transform. *IEEE transactions on visualization and computer graphics*, 22(3):1278–1290, 2015. 2
- [36] Hari Sundar, Deborah Silver, Nikhil Gagvani, and Sven Dickinson. Skeleton based shape matching and retrieval. In *2003 Shape Modeling International.*, pages 130–139. IEEE, 2003. 1
- [37] Andrea Tagliasacchi, Ibraheem Alhashim, Matt Olson, and Hao Zhang. Mean curvature skeletons. In *Computer Graphics Forum*, volume 31, pages 1735–1744. Wiley Online Library, 2012. 1, 2
- [38] Jiapeng Tang, Xiaoguang Han, Junyi Pan, Kui Jia, and Xin Tong. A skeleton-bridged deep learning approach for generating meshes of complex topologies from single rgb images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4541–4550, 2019. 1
- [39] Phi Vu Tran. Learning to make predictions on graphs with autoencoders. In *2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 237–245. IEEE, 2018. 5
- [40] Nhon H Trinh and Benjamin B Kimia. Skeleton search: Category-specific object recognition and segmentation using a skeletal shape model. *International Journal of Computer Vision*, 94(2):215–240, 2011. 1
- [41] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 52–67, 2018. 2
- [42] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-cnn: Octree-based convolutional neural networks for 3d shape analysis. *ACM Transactions on Graphics (SIGGRAPH)*, 36(4), 2017. 5
- [43] Shihao Wu, Hui Huang, Minglun Gong, Matthias Zwicker, and Daniel Cohen-Or. Deep points consolidation. *ACM Transactions on Graphics (TOG)*, 34(6):1–13, 2015. 1, 2, 6
- [44] Zhan Xu, Yang Zhou, Evangelos Kalogerakis, and Karan Singh. Predicting animation skeletons for 3d articulated models via volumetric nets. In *2019 International Conference on 3D Vision (3DV)*, pages 298–307. IEEE, 2019. 1, 2
- [45] Yajie Yan, David Letscher, and Tao Ju. Voxel cores: Efficient, robust, and provably good approximation of 3d medial axes. *ACM Transactions on Graphics (TOG)*, 37(4):1–13, 2018. 2
- [46] Baorong Yang, Junfeng Yao, Bin Wang, Jianwei Hu, Yiling Pan, Tianxiang Pan, Wenping Wang, and Xiaohu Guo. P2mat-net: Learning medial axis transform from sparse point clouds. *Computer Aided Geometric Design*, page 101874, 2020. 2
- [47] Kangxue Yin, Hui Huang, Daniel Cohen-Or, and Hao Zhang. P2p-net: Bidirectional point displacement net for shape transform. *ACM Transactions on Graphics (TOG)*, 37(4):1–13, 2018. 2
- [48] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A Efros. View synthesis by appearance flow. In *European conference on computer vision*, pages 286–301. Springer, 2016. 2