

# Scan2Mesh: From Unstructured Range Scans to 3D Meshes

Angela Dai

Matthias Nießner

Technical University of Munich



Figure 1: 3D scans of objects suffer from sensor occlusions as well as noisy, oversmoothed reconstruction quality in very dense, triangle-heavy meshes, due to both sensor noise and resolution as well as reconstruction artifacts. We propose a novel approach, leveraging graph neural networks, which takes a partial scan of an object and generates a complete, lightweight 3D mesh of the object. Our approach is the first to propose a generative deep-learning based model for directly creating a 3D mesh as an indexed face set.

## Abstract

We introduce *Scan2Mesh*, a novel data-driven generative approach which transforms an unstructured and potentially incomplete range scan into a structured 3D mesh representation. The main contribution of this work is a generative neural network architecture whose input is a range scan of a 3D object and whose output is an indexed face set conditioned on the input scan. In order to generate a 3D mesh as a set of vertices and face indices, the generative model builds on a series of proxy losses for vertices, edges, and faces. At each stage, we realize a one-to-one discrete mapping between the predicted and ground truth data points with a combination of convolutional- and graph neural network architectures. This enables our algorithm to predict a compact mesh representation similar to those created through manual artist effort using 3D modeling software. Our generated mesh results thus produce sharper, cleaner meshes with a fundamentally different structure from those generated through implicit functions, a first step in bridging the gap towards artist-created CAD models.

## 1. Introduction

3D meshes are one of the most popular representations used to create and design 3D surfaces, from across content creation for movies and video games to architectural and mechanical design modeling. These mesh or CAD models are handcrafted by artists, often inspired by or mimicking real-world objects and scenes through expensive, significantly tedious manual effort. Our aim is to develop a generative model for such 3D mesh representations; that is, a mesh model described as an indexed face set: a set of vertices as 3D positions, and a set of faces which index into the vertices to describe the 3D surface of the model. In this way, we can begin to learn to generate 3D models similar to the handcrafted content creation process.

The nature of these 3D meshes, structured but irregular (e.g., irregular vertex locations, varying face sizes), make them very difficult to generate. In particular, with the burgeoning direction of generative deep learning approaches for 3D model creation and completion [5, 10, 30, 6], the irregularity of mesh structures provides a significant challenge, as these approaches are largely designed for regular grids. Thus, work in the direction of generating 3D models predominantly relies on the use of implicit functions stored

in regular volumetric grids, for instance the popular truncated signed distance field representation [3]. Here, a mesh representation can be extracted at the isosurface of the implicit function through Marching Cubes [21]; however, this uniformly-sampled, unwieldy triangle soup output remains fundamentally different from 3D meshes in video games or other artist-created mesh/CAD content.

Rather than generate 3D mesh models extracted from regular volumetric grids, we instead take inspiration from 3D models that have been hand-modeled, that is, compact CAD-like mesh representations. Thus, we propose a novel approach, Scan2Mesh, which constructs a generative formulation for producing a mesh as a lightweight indexed face set, and demonstrate our approach to generate complete 3D mesh models conditioned on noisy, partial range scans. Our approach is the first, to the best of our knowledge, to leverage deep learning to fully generate an explicit 3D mesh structure. From an input partial scan, we employ a graph neural network based approach to jointly predict mesh vertex positions as well as edge connectivity; this joint optimization enables reliable vertex generation for a final mesh structure. From these vertex and edge predictions, interpreting them as a graph, we construct the corresponding dual graph, with potentially valid mesh faces as dual graph vertices, from which we then predict mesh faces. These tightly coupled predictions of mesh vertices along with edge and face structure enable effective transformation of incomplete, noisy object scans to complete, compact 3D mesh models. Our generated meshes are cleaner and sharper, while maintaining fundamentally different structure from those generated through implicit functions; we believe this is a first step to bridging the gap towards artist-created CAD models.

To sum up, our contributions are as follows:

- A graph neural network formulation to generate meshes directly as indexed face sets.
- Demonstration of our generative model to the task of shape completion, where we achieve significantly cleaner and more CAD-like results than state-of-the-art approaches.

## 2. Related Work

Recent advances in machine learning, coupled with the increasing availability of 3D shape and scene databases [2, 30, 4], has spurred development of deep learning approaches on 3D data. 3D ShapeNets [35] and VoxNet [22] were among the first approaches to propose 3D convolutional neural networks, both leveraging occupancy-based representations encoded in regular volumetric grids in order to perform object recognition tasks. Various other approaches have since been developed upon 3D CNN-based

architectures, targeting applications such as object classification [24], object detection [29], 3D keypoint matching [36], and scan completion [5, 10, 30, 6].

Such approaches have largely been developed upon regular volumetric grid representations, a natural 3D analogue to image pixels. Earlier 3D CNN approaches leveraged occupancy-based volumetric representations [35, 22, 24], simply encoding whether each voxel is occupied, empty (or optionally unknown). Inspiration has also been taken from work in 3D scanning and reconstruction, where implicit volumetric representations, in particular truncated signed distance fields, are very popular. Such representations encode both finer-grained information about the surface as well as the empty space, and have recently been effectively leveraged for both discriminative and generative tasks [5, 30, 6]. For generative tasks, Liao et al. [18] proposed a learned marching cubes mesh extraction from a volumetric grid for further output refinement. Hierarchical strategies have also been developed to alleviate the cubic cost of such dense volumetric representations [26, 33], and have been shown to generate higher-resolution output grids [25, 31, 11, 10, 34]. However, the 3D surfaces extracted from these regular volumetric grids maintain fundamentally different structure from that of handcrafted CAD models.

Point-based representations have recently been popularized with the introduction of the PointNet architecture [23], which demonstrated 3D classification and segmentation on a more efficient 3D representation than dense volumetric grids. Generative approaches have also been developed upon point cloud representations [7], but 3D point cloud outputs lack the structured surface information of meshes.

Several approaches for inferring the mesh structure of an object from an input image have recently been introduced, leveraging very strong priors on possible mesh structure in order to create the output meshes. AtlasNet [9] learns to generate a 2D atlas embedding of the 3D mesh of an object. Another approach is to learn to deform template meshes (e.g., an ellipsoid) to create an output 3D mesh model of an object [19, 32, 12]. Such approaches generate 3D mesh surfaces as output, but are constrained to a limited set of mesh structures, whereas we aim to generate the explicit mesh structure from scratch.

In contrast to previous approaches, we take inspiration from handcrafted CAD models and develop an approach to generate the full mesh graph structure, from vertices to edges to faces. To this end, we leverage developments in machine learning approaches on graphs, in particular graph neural networks [28], to formulate a method to generate 3D mesh vertices, edges, and faces.

## 3. Method Overview

Our method generates a 3D mesh as a set of vertices (3D positions) and a set of  $k$ -sided faces which index into the

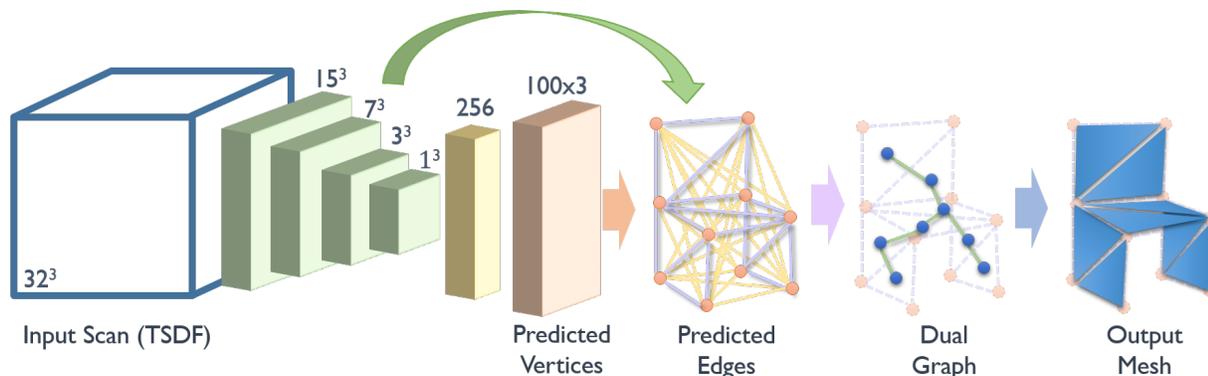


Figure 2: Our Scan2Mesh approach takes as input an partial scan of an object as a TSDF, and proposes a new graph neural network formulation to predict the mesh graph structure of vertices, edges, and faces. First, the input TSDF is used to jointly predict mesh vertices and edges as a graph, then this graph is transformed into its dual in order to predict the final mesh output faces (which need not contain all intermediate predicted edges). We maintain losses on each of the mesh vertex, edge, and face predictions to produce a final output mesh graph structure.

vertices to describe the mesh surface, conditional on an input partial range scan. Note that our approach is agnostic to the input data and representation as our focus lies in the formulation of a generative approach to explicitly generate mesh structure; in this paper, we use the task of shape completion to exemplify our approach. For shape completion, we aim to generate a complete mesh model from an input partial scan of an object. Here, the input scan is captured as depth image or set of depth images, which are then fused into a  $32^3$  volumetric grid as a truncated signed distance field through volumetric fusion [3]. Training is performed with supervised input-target pairs, with input scans generated by virtually scanning objects from the ShapeNet dataset [2].

We propose a new graph neural network in order to predict the vertices, edges, and then faces of the mesh graph structure. First, features from the input TSDF scan are computed through a series of 3D convolutions; from this feature space, we predict a set of 3D vertex locations. These vertex locations form the nodes of the mesh graph. We construct a fully connected graph on these mesh vertices, and employ graph neural network to predict which mesh edges belong to the mesh graph structure. Note that the vertices and edges are predicted jointly in order to learn reliable vertex generation for a final mesh structure.

From the graph of intermediate predicted vertices and edges, we then construct the dual graph in order to predict the final face structure of the mesh. The nodes of the dual graph characterize potential faces (i.e., each node represents a potential face, which is a set of  $k$  predicted edges that connect to form a valid  $k$ -sided face), and we employ another graph neural network to predict the final mesh faces. We maintain losses on the vertices, edges, and faces during this mesh generation process in order to learn to generate CAD-like mesh models.

## 4. Scan2Mesh Network Architecture

Our Scan2Mesh network architecture is visualized in Figure 2. It is composed of two main components: first, a 3D-convolutional and graph neural network architecture to jointly predict vertex locations and edge connectivity; and second, a graph neural network to predict the final mesh face structure. For the task of shape completion, we take as input a range scan represented as a truncated signed distance field (TSDF) in a  $32^3$  volumetric grid. We represent the TSDF as a 5-channel volumetric grid, in which the first two channels store the truncated unsigned distance field values and known/unknown space according to the camera trajectory of the scan, and the last three channels store the  $(x, y, z)$  coordinates of the volumetric grid in the coordinate system of the mesh vertex positions, so that the TSDF volume is spatially “aligned” with the mesh – in the same spirit as the CoordConv operator proposed by [20]. The TSDF data generation of the partially-scanned input is detailed in Sec. 5.

### 4.1. Joint Vertex and Edge Prediction

The TSDF input then is used to predict a set of  $n$  mesh vertex locations through a series of 3D convolutions (kernel sizes 4, 3, 3, 3, all but the last followed by a  $1 \times 1 \times 1$  convolution). The resulting feature space,  $f(\text{TSDF})$  is used to predict an  $n \times 3$  tensor of  $n$  vertex position through a series of two fully-connected layers. We also denote the intermediary feature space after two sets of 3D convolutions as  $f_2(\text{TSDF})$ , which is used to capture spatial features of the input scan to inform the edge prediction.

We then construct a fully-connected graph with  $n$  nodes corresponding to the  $n$  vertex positions. The initial node features are characterized by the 3-dimensional vertex positions, in addition to the closest feature vector in  $f_2(\text{TSDF})$

by looking up the vertex positions into the  $f_2(\text{TSDf})$  grid. We propose a graph neural network on this graph, which remains agnostic to the vertex ordering. For a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  comprising vertices  $v \in \mathcal{V}$  and edges  $e = (v, v') \in \mathcal{E}$ , messages are passed from nodes to edges, and edges to nodes as follows, similar to [8, 16]:

$$v \rightarrow e : \mathbf{h}'_{i,j} = f_e([\mathbf{h}_i, \mathbf{h}_j])$$

$$e \rightarrow v : \mathbf{h}'_i = f_v\left(\sum_{\{e_{i,j}\}} \mathbf{h}_{i,j}\right)$$

where  $\mathbf{h}_i$  represents the feature of vertex  $v_i$ ,  $\mathbf{h}_{i,j}$  represents the feature of edge  $e_{i,j}$ , and  $[\cdot, \cdot]$  denotes concatenation. Thus, an edge  $e_{i,j}$  receives updates through the concatenated features of the vertices  $v_i, v_j$  it is defined by, and a vertex  $v_i$  receives updates through the sum of the features of the edges  $e_{i,j}$  incident on  $v_i$ .  $f_v$  and  $f_e$  are MLPs operating on nodes and edges, respectively. For full network architecture details regarding layer definitions, please see the supplemental material.

The vertices are optimized with an  $\ell_1$  loss, where we first compute a one-to-one mapping between the predicted vertices and ground truth vertices using the Hungarian algorithm [17]. This one-to-one mapping during training is essential for predicting reliable mesh structure; a greedy approach (e.g., Chamfer loss on vertices) results in collapse of smaller structures as shown in Figure 3.

The output predicted vertices along with the input scan features  $f_2(\text{TSDf})$  are then used to predict edge connectivity on the graph of the mesh with vertices as nodes. Each node is initially associated with two features, the 3-dimensional vertex positions and the closest feature vector in  $f_2(\text{TSDf})$  according to the respective vertex positions. These features are processed independently through small MLPs, then concatenated to the form vertex features which are then processed through graph message passing. For each edge in the fully-connected graph, we predict whether it belongs to the mesh graph structure using a (two-dimensional) cross entropy loss. The vertex positions and edges are optimized for jointly in order to reliably predict vertices belonging to a mesh structure.

## 4.2. Mesh Face Prediction

We predict the final mesh faces from these intermediate predicted vertices and edges by transforming the graph of predicted mesh vertices and edges into its dual graph. This dual graph comprises the set of valid potential faces as the nodes of the graph, with a (dual graph) edge between two nodes if the two potential faces share an edge. The nodes are represented by an 8-dimensional feature vector comprising the centroid, normal, surface area, and radius of its respective potential face. We then employ a graph neural network formulated similarly as that for the vertex and

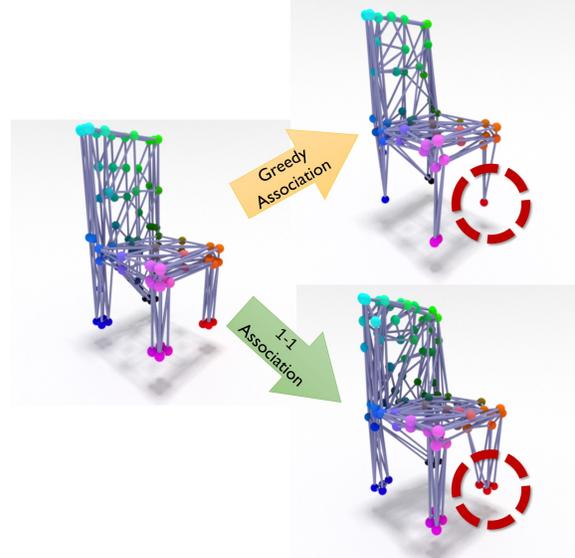


Figure 3: During training, we map the predicted graph with a one-to-one mapping on the vertices with the ground truth (top-left) using the Hungarian algorithm for bi-bipartite matching [17]. This enables prediction of both large structures as well as small structures, which might collapse with a greedy association, as seen in the chair legs (top, right).

edge prediction, this time predicting which faces belong to the final mesh structure. Note that final mesh face predictions need not contain all intermediary predicted edges. We first train the face prediction using a cross entropy loss on the nodes of the dual graph, and then use a chamfer loss between points sampled from the predicted mesh and the target mesh in order to better encourage all structurally important faces to be predicted.

## 4.3. Training

To train our model, we use the training data generated from the ShapeNet dataset [2] as described in Sec. 5.

We use the ADAM optimizer [15] with a learning rate of 0.0005 and batch size of 8 for all training. We train on eight classes of the ShapeNet dataset, following the train/test split of [5]. We additionally follow their training data augmentation, augmenting each train object by generating two virtual scanning trajectories for each object, resulting in 48,166 train samples and 10,074 validation samples.

We train the vertex-edge prediction for 5 epochs ( $\approx 15$  hours). While we found it sufficient to train the joint vertex-edge prediction through finding a one-to-one mapping between the predicted vertices and ground truth mesh vertices (the edges following as vertex indices), we found that for training face prediction with cross entropy loss, the one-to-one mapping sometimes resulted in distorted target faces, and it was more reliable to train the model on dual graphs computed from the ground truth meshes. Thus we first train

Method	Mesh Distance	Mesh Normal Similarity
Poisson Surface Reconstruction [13, 14]	0.0136	0.60
Point Pred + Poisson [13, 14]	0.0089	0.67
ShapeRecon [27]	0.0075	0.60
3D ShapeNets [35]	0.0027	0.68
3D-EPN [5]	0.0023	0.76
<b>Ours</b>	<b>0.0016</b>	<b>0.83</b>

Table 1: Quantitative shape completion results for different methods on synthetic scans of ShapeNet objects. We measure the distance between the predicted meshes and the ground truth mesh as the average point distance between points uniformly sampled from the respective meshes, as well as the normal similarity to the ground truth mesh. *Point Pred + Poisson* refers to using our architecture to only predict 1024 “vertices,” followed by Poisson Surface Reconstruction [13, 14].

Input	Average		Chairs		Tables		Airplanes		Dressers		Lamps		Boats		Sofas		Cars	
	Dist	NSim																
Points	0.0019	0.80	0.0016	0.79	0.0022	<b>0.82</b>	<b>0.0008</b>	<b>0.93</b>	0.0015	0.76	0.0045	0.72	0.0013	0.82	0.0021	0.76	0.0009	0.78
TSDF	<b>0.0016</b>	<b>0.83</b>	<b>0.0015</b>	<b>0.82</b>	<b>0.0021</b>	<b>0.82</b>	0.0010	<b>0.93</b>	<b>0.0014</b>	<b>0.79</b>	<b>0.0029</b>	<b>0.80</b>	<b>0.0011</b>	<b>0.85</b>	<b>0.0016</b>	<b>0.80</b>	<b>0.0008</b>	<b>0.83</b>

Table 2: Evaluating the effect of different input scan representations. We compare point cloud inputs with TSDF inputs, measuring the distance between the predicted meshes and the ground truth mesh as the chamfer distance between points uniformly sampled from the respective meshes, as well as the normal similarity to the ground truth mesh. The regularity of the TSDF and encoding of known and unknown space result in improved mesh prediction results.

the face prediction network for 1 epoch ( $\approx 6$  hours) using a cross entropy loss and ground truth dual graph data, and then train on dual graphs from predicted vertices and edges using a chamfer loss between the predicted and target meshes (for 1 epoch,  $\approx 18$  hours).

## 5. Data Generation

For training data generation, we use the ShapeNet model database [2], and train on a subset of 8 classes (see Sec. 6). We follow the training data generation process of [5], generating training input-target pairs by virtually scanning the ShapeNet objects along the camera trajectories given by their ShapeNet virtual scans dataset. We use two trajectories for each object for training. The virtually captured depth map(s) along these trajectories are then fused into a  $32^3$  grid through volumetric fusion [3] to obtain input TSDFs. We use a truncation of 3 voxels for all experiments. An object is mapped from its world space into a  $32^3$  grid by scaling the largest bounding box extent to  $32 - 3 * 2$  (for 3 voxels of padding on each side).

For ground truth meshes, we use triangle meshes simplified from ShapeNet models. In order to both reduce the complexity of the graph sizes as well as unify some of the

irregularity of the ShapeNet meshes, we simplify all target meshes to 100 vertices each using the V-HCAD library [1], which approximately maintains the convex hull of the original mesh.

## 6. Results and Evaluation

In this section, we provide an evaluation of our proposed method with a comparison to existing approaches on the task of scan completion of 3D shapes. We evaluate on the ShapeNet [2] dataset, using the train/test split provided by 3D-EPN [5] comprising 8 classes: chairs, tables, sofas, dressers, lamps, boats, cars, and airplanes. We test on the 1200 object test set proposed by 3D-EPN of single depth image range scans (150 objects per class), where input scans are aligned with the ground truth complete meshes, which lie in the unit cube. We compare our mesh results to meshes produced by state-of-the-art approaches; in the case that an approach generates an implicit function, we extract an output mesh using Matlab’s *isosurface* function. To measure the mesh quality, we employ two metrics: (1) we measure the mesh completeness using a chamfer distance between uniformly sampled points from the predicted mesh and the ground truth mesh, and (2) we measure the nor-

	Average		Chairs		Tables		Airplanes		Dressers		Lamps		Boats		Sofas		Cars	
	Dist	NSim																
Greedy	0.0022	0.74	0.0020	0.73	0.0027	0.75	0.0012	0.85	0.0020	0.70	0.0047	0.70	0.0016	0.72	0.0021	0.71	0.0011	0.74
1-to-1	<b>0.0016</b>	<b>0.83</b>	<b>0.0015</b>	<b>0.82</b>	<b>0.0021</b>	<b>0.82</b>	<b>0.0010</b>	<b>0.93</b>	<b>0.0014</b>	<b>0.79</b>	<b>0.0029</b>	<b>0.80</b>	<b>0.0011</b>	<b>0.85</b>	<b>0.0016</b>	<b>0.80</b>	<b>0.0008</b>	<b>0.83</b>

Table 3: Evaluating greedy vs 1-to-1 association of predictions and ground truth during training. We measure the distance between the predicted meshes and the ground truth mesh as the chamfer distance between points uniformly sampled from the respective meshes, as well as the normal similarity to the ground truth mesh. Here, a 1-to-1 discrete mapping encourages higher quality vertex, edge, and face predictions.

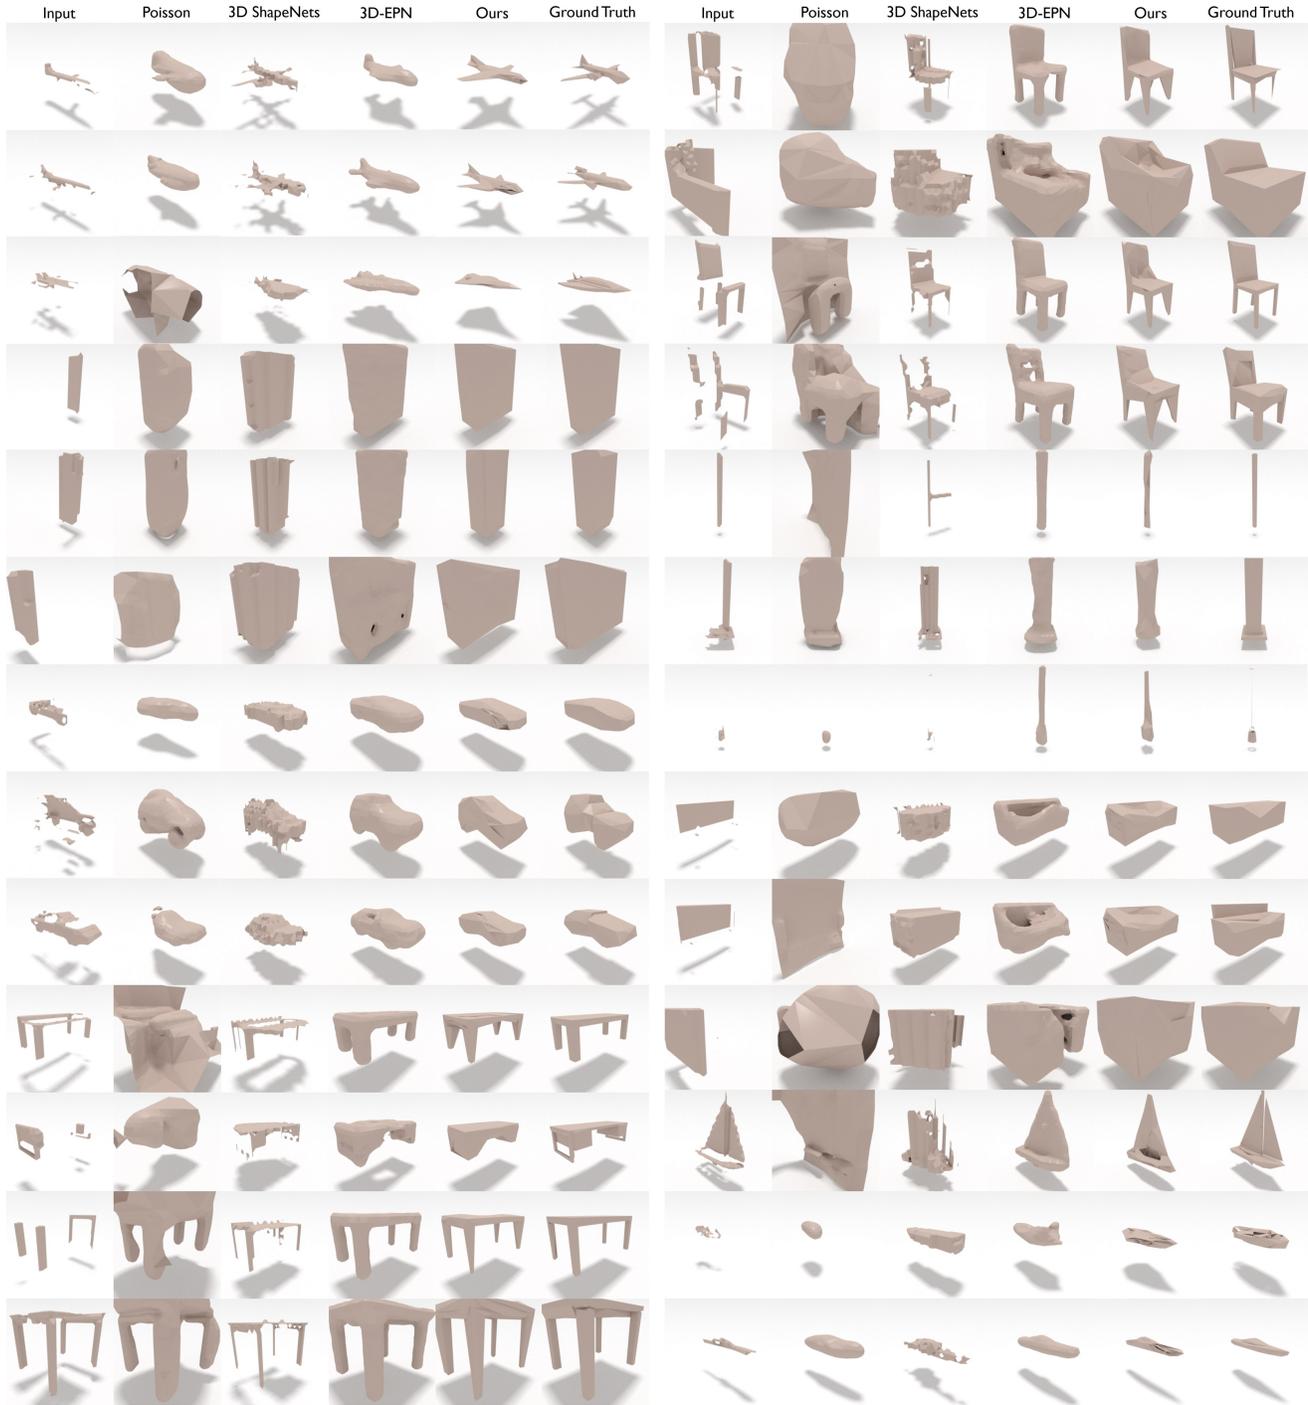
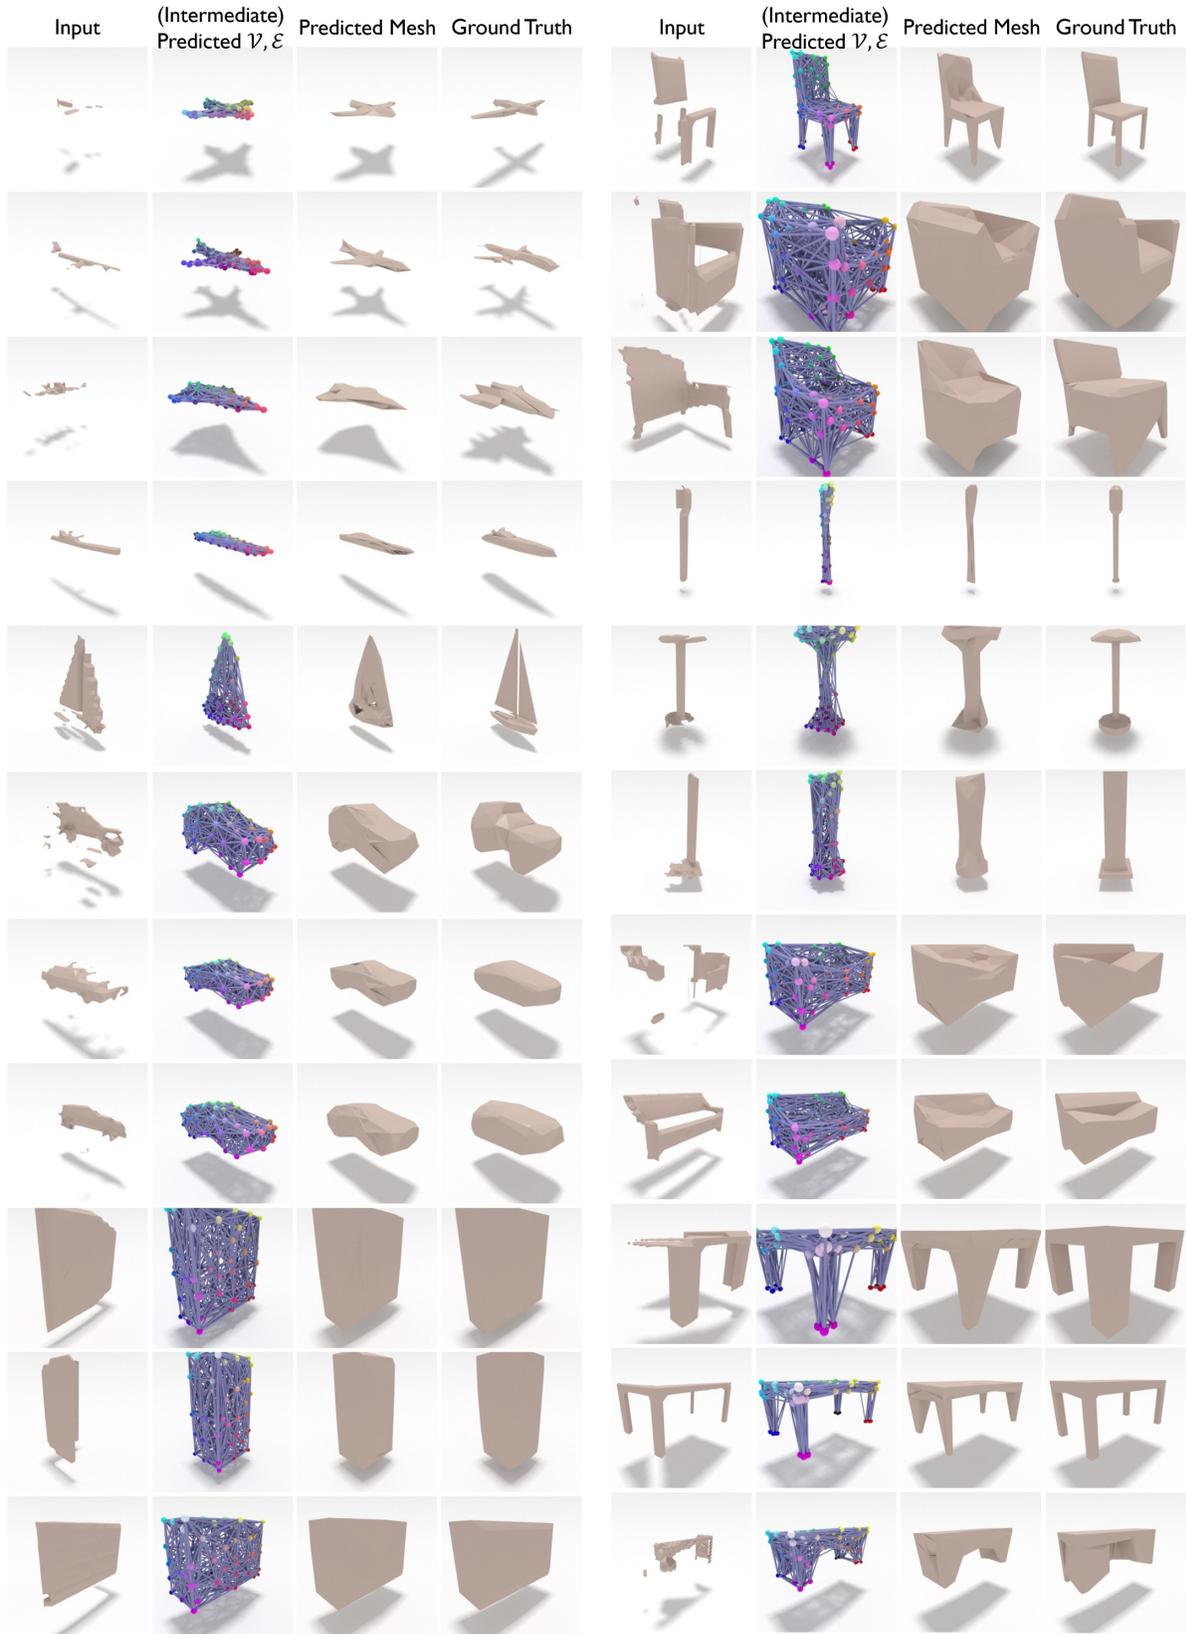


Figure 4: Qualitative scan completion results on virtual scans of ShapeNet [2] objects, in comparison to Poisson Surface Reconstruction [13, 14], as well as the volumetric generative approaches of 3D ShapeNets [35] and 3D-EPN [5]. We show results on a variety of object classes, and produce both sharp and complete mesh structure in contrast to the volumetrically regular triangulation and noisy or oversmoothed results from approaches using implicit representations on a volumetric grid.

mal deviation from the ground truth mesh to characterize mesh sharpness and cleanliness. The normal deviation metric is computed bi-directionally: for meshes  $M_a, M_b$ , we

sample points from each of their surfaces and compute the normal deviation  $N(M_b, M_a)$  from  $M_b$  to  $M_a$  as the average of the cosine of the normal angle difference for the



	Average		Chairs		Tables		Airplanes		Dressers		Lamps		Boats		Sofas		Cars	
	Dist	NSim																
Direct(GT)	0.0042	0.66	0.0035	0.62	0.0042	0.66	0.0078	0.69	0.0030	0.63	0.0053	0.67	0.0022	0.71	0.0058	0.59	0.0014	0.69
Direct(Surf)	0.0031	0.69	0.0031	0.64	0.0028	0.69	0.0025	0.81	0.0028	0.62	0.0077	0.66	0.0016	0.72	0.0033	0.62	0.0010	0.73
Dual Pred	<b>0.0016</b>	<b>0.90</b>	<b>0.0015</b>	<b>0.90</b>	<b>0.0021</b>	<b>0.89</b>	<b>0.0010</b>	<b>0.93</b>	<b>0.0014</b>	<b>0.88</b>	<b>0.0029</b>	<b>0.86</b>	<b>0.0011</b>	<b>0.88</b>	<b>0.0016</b>	<b>0.91</b>	<b>0.0008</b>	<b>0.91</b>

Table 4: Evaluating direct prediction of faces using a mesh graph with mesh vertices as nodes in comparison to using the dual graph with potential faces as nodes. We measure the distance between the predicted meshes and the ground truth mesh as the chamfer distance between points uniformly sampled from the respective meshes, as well as the normal similarity to the ground truth mesh. The dual graph significantly reduces the large combinatorics of the possible faces, providing much more reliable mesh prediction results.

closest sampled point from  $M_b$  to each that of  $M_a$ , and take  $0.5(N(M_a, M_b) + N(M_b, M_a))$  as the global normal deviation (taking the best normal deviation from a search window of 0.03, to disambiguate small positional misalignments).

**Comparison to state of the art.** We evaluate against state-of-the-art shape completion approaches in Table 1 and Figure 4. Additionally, we evaluate various design choices in Tables 2, 3, and 4. Here, we see that our approach generates sharper, cleaner meshes than previous volumetric-based approaches while producing accurate completeness in global shape structure.

**What is the impact of the input scan representation?** We evaluate our approach using a point cloud representation of the input range scan (uniformly sampled from the range image inputs) in comparison to a TSDF in Table 2. To process the point cloud input, we replace the volumetric convolutions of the encoder with a PointNet-based architecture [23]. Both representations produce good mesh results, but we find that regularity and encoding of known and unknown space in the TSDF produces better completion and mesh quality.

**Do we need a 1-to-1 mapping between prediction and target during training?** In Table 3, we evaluate using a greedy mapping between predicted vertices and target vertices during vertex-edge training. Using a greedy mapping degrades the quality of vertex predictions with respect to the final mesh structure (e.g., we want a cluster of vertices at the end of a chair leg instead of one vertex), and results in worse mesh reconstruction quality (see Figure 3 for an example visualization).

**Why use the dual graph for face prediction?** We evaluate our face prediction approach, which leverages the dual graph of the mesh vertex-edge graph, in Table 4. Here, we compare against directly predicting mesh faces using the same formulation as the joint vertex-edge prediction, where instead of predicting edges as which two vertices are connected, we predict faces as which sets of three vertices are connected, resulting in  $\mathcal{O}(n^3)$  possible faces from which the mesh faces must be predicted (we refer to the supplemental for more detail regarding directly predicting faces). Given the large combinatorics here, where the number of ground truth mesh faces is approximately 0.2% of

the number of total possible faces (for  $n = 100$ ), we evaluate two possibilities for training the direct face prediction: *Direct(GT)* uses only the ground truth mesh faces as target faces, and *Direct(Surf)* which uses all possible faces close to the ground truth mesh surface as target faces. Both approaches nonetheless suffer from the heavy combinatorics, whereas our approach of predicting faces by using the dual graph of the mesh vertex-edge graph produces significantly better mesh structure and completeness.

## 6.1. Limitations

We propose one of the first approaches to explicitly generate a 3D mesh as an indexed face set, and believe that this is a stepping stone towards future work in constructing CAD-like 3D models akin to those currently handcrafted. For instance, our use of fully-connected graphs limits the size of our models; adapting the graphs and message passing to enable learning on significantly larger mesh graphs would open up generation of higher resolution or larger scale models. Additionally, we do not explicitly enforce mesh regularity or surface continuity (which are also not given in the ShapeNet models); adding hard constraints into the optimization to guarantee these attributes would open up many more applications for these models.

## 7. Conclusion

We presented Scan2Mesh, a generative model for creating 3D mesh models as indexed face sets, inspired by 3D model representations used in handcrafted 3D models. We proposed a new graph neural network formulation to generate a mesh representation directly, and demonstrated our mesh generation on the task of shape completion, achieving cleaner and more CAD-like mesh models from noisy, partial range scans. We believe that this opens up myriad possibilities towards bridging the gap of 3D model generation towards the quality of artist-created CAD models.

## Acknowledgments

This work was supported by the ZD.B, a Google Research Grant, a TUM Foundation Fellowship, a TUM-IAS Rudolf Mößbauer Fellowship, and the ERC Starting Grant *Scan2CAD*.

## References

- [1] V-HCAD. <https://github.com/kmammou/v-hacd>. 5
- [2] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 2, 3, 4, 5, 6, 7
- [3] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312. ACM, 1996. 2, 3, 5
- [4] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2017. 2
- [5] A. Dai, C. R. Qi, and M. Nießner. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2017. 1, 2, 4, 5, 6
- [6] A. Dai, D. Ritchie, M. Bokeloh, S. Reed, J. Sturm, and M. Nießner. Scancomplete: Large-scale scene completion and semantic segmentation for 3d scans. In *Proc. Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2
- [7] H. Fan, H. Su, and L. J. Guibas. A point set generation network for 3d object reconstruction from a single image. In *CVPR*, volume 2, page 6, 2017. 2
- [8] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212*, 2017. 4
- [9] T. Groueix, M. Fisher, V. G. Kim, B. Russell, and M. Aubry. AtlasNet: A Papier-Mâché Approach to Learning 3D Surface Generation. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [10] X. Han, Z. Li, H. Huang, E. Kalogerakis, and Y. Yu. High Resolution Shape Completion Using Deep Neural Networks for Global Structure and Local Geometry Inference. In *IEEE International Conference on Computer Vision (ICCV)*, 2017. 1, 2
- [11] C. Häne, S. Tulsiani, and J. Malik. Hierarchical surface prediction for 3d object reconstruction. *arXiv preprint arXiv:1704.00710*, 2017. 2
- [12] A. Kanazawa, S. Tulsiani, A. A. Efros, and J. Malik. Learning category-specific mesh reconstruction from image collections. In *ECCV*, 2018. 2
- [13] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, 2006. 5, 6
- [14] M. Kazhdan and H. Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics (TOG)*, 32(3):29, 2013. 5, 6
- [15] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 4
- [16] T. Kipf, E. Fetaya, K.-C. Wang, M. Welling, and R. Zemel. Neural relational inference for interacting systems. *arXiv preprint arXiv:1802.04687*, 2018. 4
- [17] H. W. Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 4
- [18] Y. Liao, S. Donné, and A. Geiger. Deep marching cubes: Learning explicit surface representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2916–2925, 2018. 2
- [19] O. Litany, A. Bronstein, M. Bronstein, and A. Makadia. Deformable shape completion with graph convolutional autoencoders. *arXiv preprint arXiv:1712.00268*, 2017. 2
- [20] R. Liu, J. Lehman, P. Molino, F. P. Such, E. Frank, A. Sergeev, and J. Yosinski. An intriguing failing of convolutional neural networks and the coordconv solution. *arXiv preprint arXiv:1807.03247*, 2018. 3
- [21] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *ACM siggraph computer graphics*, volume 21, pages 163–169. ACM, 1987. 2
- [22] D. Maturana and S. Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 922–928. IEEE, 2015. 2
- [23] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 1(2):4, 2017. 2, 8
- [24] C. R. Qi, H. Su, M. Nießner, A. Dai, M. Yan, and L. Guibas. Volumetric and multi-view cnns for object classification on 3d data. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2016. 2
- [25] G. Riegler, A. O. Ulusoy, H. Bischof, and A. Geiger. Octnetfusion: Learning depth fusion from data. *arXiv preprint arXiv:1704.01047*, 2017. 2
- [26] G. Riegler, A. O. Ulusoy, and A. Geiger. Octnet: Learning deep 3d representations at high resolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [27] J. Rock, T. Gupta, J. Thorsen, J. Gwak, D. Shin, and D. Hoiem. Completing 3d object shape from one depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2484–2493, 2015. 5
- [28] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009. 2
- [29] S. Song and J. Xiao. Deep sliding shapes for amodal 3d object detection in rgb-d images. *arXiv preprint arXiv:1511.02300*, 2015. 2
- [30] S. Song, F. Yu, A. Zeng, A. X. Chang, M. Savva, and T. Funkhouser. Semantic scene completion from a single depth image. *Proceedings of 30th IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 1, 2
- [31] M. Tatarchenko, A. Dosovitskiy, and T. Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. *arXiv preprint arXiv:1703.09438*, 2017. 2
- [32] N. Wang, Y. Zhang, Z. Li, Y. Fu, W. Liu, and Y.-G. Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. *arXiv preprint arXiv:1804.01654*, 2018. 2

- [33] P.-S. Wang, Y. Liu, Y.-X. Guo, C.-Y. Sun, and X. Tong. O-cnn: Octree-based convolutional neural networks for 3d shape analysis. *ACM Transactions on Graphics (TOG)*, 36(4):72, 2017. [2](#)
- [34] P.-S. Wang, C.-Y. Sun, Y. Liu, and X. Tong. Adaptive o-cnn: a patch-based deep representation of 3d shapes. In *SIGGRAPH Asia 2018 Technical Papers*, page 217. ACM, 2018. [2](#)
- [35] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1912–1920, 2015. [2](#), [5](#), [6](#)
- [36] A. Zeng, S. Song, M. Nießner, M. Fisher, and J. Xiao. 3dmatch: Learning the matching of local 3d geometry in range scans. *arXiv preprint arXiv:1603.08182*, 2016. [2](#)