

Point2CAD: Reverse Engineering CAD Models from 3D Point Clouds

Yujia Liu
ETH Zürich

Anton Obukhov
ETH Zürich

Jan Dirk Wegner
University of Zürich

Konrad Schindler
ETH Zürich



Figure 1. **Point2CAD** is a pipeline for reconstructing complex CAD models, including freeform surfaces, from 3D point clouds. A raw point cloud is segmented into clusters corresponding to CAD faces, and each face is fitted either with a geometric primitive or with a parametric freeform surface, using a novel, implicit neural representation. Due to the analytic representation the surfaces can be extended and intersected, such that topology emerges, which is then used to clip the surface primitives.

Abstract

Computer-Aided Design (CAD) model reconstruction from point clouds is an important problem at the intersection of computer vision, graphics, and machine learning; it saves the designer significant time when iterating on in-the-wild objects. Recent advancements in this direction achieve relatively reliable semantic segmentation but still struggle to produce an adequate topology of the CAD model. In this work, we analyze the current state of the art for that ill-posed task and identify shortcomings of existing methods. We propose a hybrid analytic-neural reconstruction scheme that bridges the gap between segmented point clouds and structured CAD models and can be readily combined with different segmentation backbones. Moreover, to power the surface fitting stage, we propose a novel implicit neural representation of freeform surfaces, driving up the performance of our overall CAD reconstruction scheme. We extensively evaluate our method on the popular ABC benchmark of CAD models and set a new state-of-the-art for that dataset. Code is available at <https://github.com/YujiaLiu76/point2cad>.

1. Introduction

The task of reverse engineering CAD models from 3D point clouds has gained increasing attention in recent years due to the rapid development of 3D scanning technologies. Reverse engineering involves transforming a physical object into a digital model in an editing-friendly format, which can be used for analysis, visualization, and manufacturing.

Most approaches to reverse engineering CAD models follow a typical sequence of steps. First, data capture and pre-processing are carried out to obtain point cloud data sampled densely from the object’s surface. Subsequently, the point cloud is segmented to identify individual regions of interest. The regions are then classified to differentiate between different types of surfaces and features. Analytical representations are generated based on the identified regions, followed by the reconstruction of the model’s topology. Commercial software systems offer interactive tools and functions for each step, such as a dedicated segmentation or shape-fitting tool. Similarly, research that aims to automate the process has focused chiefly on individual steps, whereas little work covers the complete workflow of CAD model reconstruction. Furthermore, most existing work is either limited to case studies of a few models or is still quite

far from producing satisfactory results.

In particular, reconstruction with simple parametric surface primitives has been studied extensively. For such surfaces, both fitting and computing intersections are comparatively straight-forward, and with sufficiently clean point clouds, a boundary representation (B -rep) can be derived with classical geometric methods [1, 4, 44]. On the contrary, while there are several works about fitting freeform surfaces, often with various flavors of splines, they mostly stop short of assembling individual surfaces into complete B -reps [8, 39, 47]. In summary, geometrically complex free-form surfaces and topologically complex assemblies of simple surfaces have mostly been studied separately. The present work addresses this gap by proposing **Point2CAD**, a method that recovers complete CAD models, including free-form surfaces, edges, and corners. Refer to Sec. A1 of sup.mat. for notation.¹

Our proposed pipeline for CAD model reconstruction from point clouds comprises several steps. First, a pre-trained neural network segments the point cloud into clusters corresponding to distinct surfaces. Second, surfaces are fitted to the clusters, which involves both basic primitive fitting and a novel neural fitting scheme for freeform surfaces. Third, adjacent surfaces are intersected to recover edges, and adjacent edges are further intersected to recover corners, thus obtaining a full B -rep. Taken together, these steps form a comprehensive and versatile pipeline for reverse engineering point clouds into CAD models; see Fig. 1. By combining modern, learning-based segmentation backbones with classical geometric primitive fitting and with recent neural field methods, we get the best of both worlds and obtain a reconstruction pipeline that sets a new state of the art on the large and diverse ABC benchmark [19]. Our contributions can be summarized as follows:

- We propose Point2CAD, a comprehensive workflow for reconstructing complex free-form parametric CAD models from raw 3D point clouds;
- To power the surface fitting stage, we propose a novel neural representation of freeform surfaces;
- We demonstrate the superior quantitative and qualitative performance on the challenging ABC dataset.

2. Related work

Our work aims to recover CAD models from raw point clouds. Towards this goal, we draw inspiration from several domains of computer vision and 3D data processing.

3D point cloud segmentation. Many studies segment point clouds by assigning each point a class or instance label. Neural architectures to learn feature embeddings of point clouds include: unordered sets of neighbors [31, 32], graph

convolutions [22, 45], point convolutions [14, 24], point-voxel learning [26, 34], and transformers [10, 49]. Several authors use point-wise labeling as a mechanism to structure point clouds into geometric primitives [11]. *E.g.*, PrimitiveNet [15] segments point clouds into primitives via local embedding and adversarial learning. In [27] graph convolution is used to detect surface boundaries in 3D point clouds.

Primitive fitting [5, 17] is a common way to abstract 3D point clouds into CAD-like parametric elements. Classical solutions rely on robust statistical procedures like the Hough transform to recognize geometric primitives in point clouds despite the noise and missing data [35]. Several learning-based approaches have also been proposed to fit geometric primitives to point clouds. SPFN [23] is trained to detect varying numbers of primitives at different scales, supervised with ground truth surfaces and their membership in the primitives. ParSeNet [39] finds parametric surfaces in point clouds, including basic geometric primitives as well as B-spline surfaces, but does not put much effort into connecting them. HPNet [47] concentrates on partitioning the point cloud into segments using semantic and spectral features and edge information in the form of an adjacency matrix but does not fit actual primitives. Saporta et al. [37] proposed an unsupervised, recursive neural architecture that learns to fit geometric primitives to 3D points, while Vasu et al. [43] develop a hybrid representation that combines explicit geometric shapes with signed distance functions to represent and manipulate regular and freeform 3D shapes. Besides conventional geometric primitives, parametric freeform surfaces are an important building block of CAD systems. DeepSpline [8] utilizes a hierarchical RNN to reconstruct 2D spline curves, in conjunction with an unsupervised learning approach to recover surfaces of revolution and extrusion.

CAD representations. BRepNet [21] is a neural network architecture that operates directly on B -rep data structures. UV-Net [16] introduces another neural architecture that operates on B -rep data, without fitting to points. JOINABLE [46] starts from CAD primitives and learns to assemble them and form joints, using weak supervision available in standard CAD files. Another popular CAD representation is constructive solid geometry (CSG). CAPRI-Net [48] learns to fit a collection of quadric 3D primitives, and an associated CSG-tree, while CSGNet [38] learns to generate CSG programs for given 2D or 3D shapes, using a convolutional encoder and a recurrent decoder.

Generic CAD modeling. Reconstructing a complete CAD model [3, 30] is challenging, as the output should include both the geometric surface shapes and their connectivity. Traditional methods for reverse engineering often employ iterative RANSAC estimation to extract primitive shapes from point clouds, *e.g.*, [25]. Benko et al. [2] outline a hand-engineered, sequential procedure to reconstruct CAD

¹A note on terminology: many different terms have been used to denote the topological elements of CAD models, see list in the supplementary material. We use the terms *surface* (2D), *edge* (1D) and *corner* (0D).

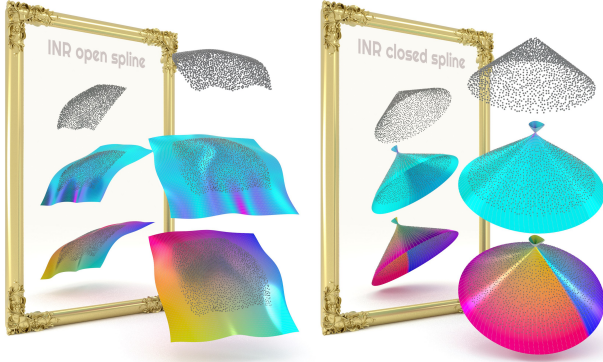


Figure 2. Visualization of open (**left**) and closed (**right**) surface fitting with the proposed INR. **Top**: point cloud sampled from a freeform surface; **middle**: INR fit, showing **high** or **low** principal curvature; **bottom**: same fit with color indicating uv coordinates. Our INR reproduces curvature where necessary and smoothly extrapolates without oscillations to facilitate subsequent steps.

models via segmentation, surface and topology fitting, and blending. Recent, learning-based approaches in the same spirit include Point2Cyl [42], a neural network that transforms a 3D point cloud into a set of extrusion cylinders by predicting point segmentation, base labels, normals, and extrusion parameters in closed form. ComplexGen [9] reconstructs CAD models by autoregressively detecting geometric primitives. The model consists of a sparse CNN encoder, three transformer decoders for geometric primitives and topology, and post-processing with global optimization to refine the rather rough transformer predictions.

Manifold learning refers to a family of methods that aim to discover a low-dimensional manifold underlying a higher-dimensional, possibly noisy, data set. Manifold learning methods can be understood as non-linear extensions of PCA [6] attempting to “flatten” curved manifolds. Most methods focus on visualization or fidelity to particular input points, by explicitly using those points to construct the mapping [28, 36, 41]. Autoencoders provide a natural way to not only project data non-linearly to a lower-dimensional latent space but also to decode back from the latent space to the original data space [20]. We extend this approach with the recent findings about implicit neural representations to design a fitting method for freeform surfaces.

Implicit Neural Representations (INR) are a framework to encode an arbitrary function, often observed in the form of sparse samples, into the weights of a neural network. Recent research about neural rendering has led to several useful and transferrable insights about implicit neural fields, for instance, the importance of positional encoding [29] and new (e.g., sinusoidal) activation functions [40]. These findings form the basis for our freeform surface fitting method.

3. Method

Contrary to the recent trend towards generic, end-to-end deep learning pipelines, we found it advantageous to split the reconstruction process into steps and only use neural methods where necessary. Overall, our method consists of the following stages, *cf.* Fig. 1:

1. Partition the point cloud into clusters corresponding to the CAD model’s topological faces. We rely on existing (pretrained) neural network methods for that step.
2. Fit an analytical surface primitive to each cluster. Here we use a hybrid approach: First, we test a set of prevalent geometric surface primitives that admit efficient closed-form fitting. We propose a novel fitting scheme for freeform surfaces based on neural representations.
3. Identify the effective area of each (finite) parametric surface and clip it, leaving enough margin to ensure intersections of adjacent surfaces.
4. Perform pairwise surface intersections to obtain a set of topologically plausible object edges. Using these edges, remove parts not supported by input points.
5. Perform pairwise edge intersection to identify a set of topological corners. Clip edges based on proximity to the remaining surface regions using these corners.

As a result of applying Point2CAD, we obtain a CAD model in B -rep format, which includes analytical surfaces to represent the model’s faces, compatible edges and corners, and adjacency matrices that encode topology.

A cornerstone of the proposed pipeline is a novel, efficient method for fitting the implicit neural representation of a freeform, parametric 2D-manifold surface in 3D to an unordered set of points. A particular difficulty in this context is that the surface must be extrapolated beyond the supporting points to enable robust surface intersections, an ability that tends to be challenging for neural surface estimators. To that end, we design a novel freeform surface parameterization, drawing inspiration from recent advances in neural fields, positional encoding schemes, and efficient test-time optimization. See Sec. 3.2 for details.

While using neural methods, INR is a pure test-time optimization. We effectively side-step learning priors from training data except for the initial low-level segmentation. We argue that simple analytical procedures like surface fitting do not necessarily benefit from learned priors and that an over-reliance on data-driven learning may have hampered recent work on reverse engineering CAD models. Moreover, separately predicting – rather than constructing – the topology does not guarantee that it is consistent with the geometry, *e.g.*, edges could be instantiated such that they do not lie in a surface, see Fig. 5. We empirically support this claim in our experiments, where we outperform purely learning-based methods like [9], which are inherently plagued by even slight differences between the training and test distributions. When designing Point2CAD, we



Figure 3. Inspiration items of our INR surface fitting. UMAP [28] (1st column, middle) learns the underlying 2D manifold of 3D points along with the inverse mapping, but cannot capture its smooth analytic representation. Early experiments with ReLU activations confirmed its low-frequency bias (1st column, bottom) and piece-wise linearity. SiLU [12] (2nd column) suffers from low-frequency bias, too. We found that resolving it by adding positional encoding is challenging under our training protocol (3rd column). SIREN [40] (4th column) fits the data precisely, however, does not extrapolate well. A combination of the best parts of these ingredients led us to our INR surface fitting method, seen in Fig. 2 and discussed in Sec. 3.2.

aimed for a generalizable scheme that uses offline learning where necessary – in particular, for feature learning and segmentation, ill-posed tasks that are notoriously hard to engineer heuristically; but that does not replace well-understood analytical operations with less robust, data-driven feed-forward modeling. As a side effect, obviating parameter-heavy neural networks, where possible, also results in a computationally less demanding method.

3.1. Parametrization of standard primitives

The geometric primitives that our current implementation handles are planes, spheres, cylinders, and cones. The exact parameterizations are taken as in [39], see Sec. A2 of sup.mat.

We observed that utilizing the predicted surface type may harm the reconstruction since the prediction is not always correct. Therefore, we circumvent that step and instead rely on an exhaustive, data-driven comparison: we fit each primitive type (including INR) to point clusters and select the simplest model with the lowest reconstruction error.

3.2. Freeform surface parametrization with INR

In the context of the overall CAD reconstruction pipeline, this step should satisfy several requirements: (1) resilience to noise in the input point cloud, which is inevitable with real 3D sensors; (2) support for inverse mapping, to enable traversal of the latent space; (3) flexibility in interpolation mode, to ensure high data fidelity and avoid over-smoothing; (4) strong regularization in extrapolation mode, required to construct smooth, plausible margins around the observed data points for the subsequent surface intersection; (5) fast fitting with low computational cost.

Existing manifold learning techniques that support the inverse transform, such as UMAP [28], could not be readily used for this task, as demonstrated in Fig. 3 (1st column, middle). We thus developed a custom neural autoencoder [20] with a single hidden layer, which is sufficient for simple non-linear transforms [13].

We found that for our purposes, the choice of activation function is critical: Besides causing a non-smooth, piece-wise linear approximation, ReLU activations [7] exhibit a noticeable low-frequency bias [33], as can be seen in Fig. 3 (1st column, bottom). High-frequency undulations are over-smoothed, such that entire point cloud regions noticeably deviate from the fit on either side. Replacing ReLU with SiLU [12] enhances surface smoothness, but does not overcome the low-frequency bias (2nd column). Adding positional encoding [29] in conjunction with SiLU resolves the issue, as one would expect from the literature, but provokes new artifacts, due to the fixed set of spatial frequencies (3rd column). Sinusoidal activations [40] (4th column) achieve excellent data fitting, but impose little regularization when not supported by data, and thus extrapolate poorly.

A combination of activations with desired properties within the same layer overcomes these individual problems: equipping some neurons with SiLU and the remaining ones with sinusoidal activations produces geometrically faithful fits that extrapolate smoothly, see Fig. 2. Remarkably, this behavior emerges without any dedicated mechanisms to steer the collaboration between the two types of neurons.

Fig. 4 depicts a block diagram of the proposed autoencoder. Each surface is autoencoded independently, by feeding batches of its 3D points (X, Y, Z) through a 1-layer MLP encoder and a corresponding decoder, with a 2D bottleneck

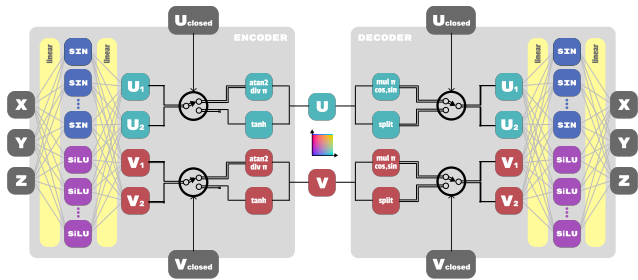


Figure 4. Block scheme of our proposed INR for freeform surface fitting. We encode groups of 3D surface points into a latent 2D uv space. We use a mixture of activations to achieve high fidelity of fitting and smooth extrapolation. Both open and closed surface fitting is supported via preconfigured routing of activations in the latent space. See Sec. 3.2 for more details.

corresponding to manifold coordinates (u, v) . The 2D latent space is regularized to a unit square $[-1, 1] \times [-1, 1]$. Both closed, and open surfaces (independently in the u and v dimensions) are supported via pre-configured routing.

For each surface, the weights of a template INR are initialized randomly and optimized with standard mini-batch (or full-batch) descent. We run Adam [18] for 1000 steps, with 50 steps warm-up of the learning rate, followed by a linear decay that reaches zero at the last step. The complete optimization takes only a few seconds on a single GPU, and multiple surfaces can be fitted in parallel.

For extrapolation and latent space traversal, we encode all cluster points into the latent uv space and store the bounding box parameters along with the autoencoder. We reset the corresponding axis range for surfaces with closed dimensions to $[-1, 1]$. To sample the extended surface with the margin, we extend the bounding box by 10% in both dimensions and compute 3D surface points using the decoder.

3.3. Topology reconstruction

Following the instantiation of all individual surfaces, it becomes necessary to establish their boundaries, which amounts to finding the intersection curves between adjacent surfaces. Since this step is crucial for subsequent stages of CAD reconstruction, it is important to correctly identify and recover the intersection curves. This requirement again calls for deterministic geometric methods, as independent detection can hardly guarantee a set of complete edges consistent with the surfaces.

While it is theoretically possible to work out the surface intersections analytically, multiple freeform surfaces lead to extremely complicated, and potentially unstable calculations. Hence, we opt for a generic solver and convert all surfaces to triangle meshes for this step to be accurately intersected with mature, numerically stable computational tools. For each surface, the (infinite) geometric primitive is

trimmed to form a margin of width ϵ around the supporting points to ensure enough space for the subsequent steps and then fed to a standard meshing algorithm by tessellation and then triangulation. We identify all intersecting surfaces for each meshed surface, compute the pairwise mesh intersections to obtain poly-line edges, and re-mesh the surfaces along those edges. Similarly, we intersect adjacent edge poly-lines to obtain corner points and, in turn, use those corners to trim the edges (which still extend beyond the surfaces since they were generated from unclipped surfaces). See Sec. A3 of sup.mat. for a detailed algorithm.

Such a discretization step temporarily side-steps the B -rep format. However, it is possible to compute the true analytical intersections starting from the discrete approximations. This arguably simpler task is left for future work.

4. Experiments

The rather flexible constraints on the input to our method permit employing it in various setups. Such may include:

- Evaluation of the ground truth point cloud clustering or segmentation as a way to quantify the contribution of point cloud sampling sparsity and sample noise on the reconstruction quality (aliased “Point2CAD GT”);
- Usage of geometry fitting and topology extraction on top of any pretrained point cloud clustering or segmentation methods, such as ParSeNet [39] or HPNet [47] (aliased “Point2CAD <method>”).

Under both setups, we first individually evaluate the surface reconstruction quality of the most challenging freeform surfaces, which is one of the main contributions of our work. We then analyze the geometry and topology of the entire reconstructed CAD dataset. This enables us to provide an overall assessment of the effectiveness of our approach in reconstructing surfaces and CAD models from point clouds.

4.1. Dataset

To evaluate the performance of Point2CAD, we conduct experiments on the ABC dataset [19], a large-scale collection of CAD models ($\sim 1.000.000$ models) widely used in geometric deep learning. We use the same split as ParseNet, where each model contains at least one freeform surface, to facilitate a fair comparison with existing methods and to better demonstrate the feasibility of our approach.

4.2. Evaluation metrics

- **Residual error** measures the discrepancy between the reconstructed surface and the corresponding ground truth surface obtained by Hungarian matching. Following the established protocol [39, 47], we densely sample points on the ground truth surface and evaluate their correspondence with the reconstructed surface. Specifically, we compute the distance between each sampled point and its projected point on the reconstructed surface, which

Table 1. Comparison of surface fitting. ComplexGen uses a transformer decoder to convert a surface latent code to a 20x20 control-points grid. ParseNet employs a neural network for the same kind of output. Our method utilizes the proposed INR surface format.

	Open surfaces		Closed surfaces	
	Res-err ↓	P-cover ↑	Res-err ↓	P-cover ↑
ComplexGen	0.021	0.938	0.023	0.900
ParseNet	0.006	0.930	0.008	0.902
Ours	0.002	0.999	0.003	0.998

Table 2. Ablation of INR design performance with and w/o noise added to the input points. Our method demonstrates favorable properties as measured by residual error and P-coverage.

Noise	Activation	Open surfaces		Closed surfaces	
		Res-err ↓	P-cover ↑	Res-err ↓	P-cover ↑
$\sigma = 0.05$	ReLU	0.016	0.519	0.018	0.509
	SiLU	0.024	0.522	0.030	0.543
	SiLU+PosEnc	0.014	0.561	0.015	0.516
	SIREN	0.012	0.524	0.014	0.521
	Ours	0.011	0.585	0.013	0.581
No noise	ReLU	0.006	0.986	0.008	0.981
	SiLU	0.020	0.905	0.023	0.887
	SiLU+PosEnc	0.004	0.992	0.007	0.982
	SIREN	0.003	0.996	0.006	0.990
	Ours	0.002	0.999	0.003	0.998

Table 3. Ablation of INR under surface margin extrapolation by a fraction c of model diameter. Point2CAD uses $c = 10\%$.

Res-err ↓	$c = 0\%$	$c = 10\%$	$c = 20\%$
SiLU	0.0582	0.0565	0.0714
SIREN	0.0195	0.0376	0.1205
Ours	0.0052	0.0156	0.0474

Table 4. Geometric evaluation of reconstructed CADs. Segmentation denotes the method used for point cloud clustering. ‘‘GT’’ stands for oracle ground truth segmentation, which is also an upper bound of the performance of our method.

	Segmentation	Res-err ↓	P-cover ↑	Chamfer ↓
ComplexGen	N/A	0.020	0.950	0.042
Point2CAD	ParseNet	0.018	0.942	0.017
Point2CAD	HPNet	0.020	0.937	0.018
Point2CAD	GT	0.011	0.968	0.016

obtain the results of all three methods on distinct point cloud clusters. While this is straightforward with ParseNet and our method, ComplexGen does not operate on individual clusters but instead requires the entire point cloud at once and outputs all B -rep elements at once, too. Thus, to obtain ComplexGen predictions for a particular cluster, we run it with the entire point cloud input, find the nearest predicted surface to the queried cluster out of all the predictions, and

perform the surface-to-surface evaluation on that pair.

We repeat this procedure for a subset of freeform surfaces of the ABC dataset and average the geometry metrics; see Tab. 1 for quantitative evaluation. As can be seen, our INR fitting method generates freeform surfaces that represent the underlying points more faithfully than prior parameterizations. Qualitative results of Point2CAD freeform surface fitting confirm the qualitative study, which can be seen in Fig. 2, and the side-by-side comparison is in Fig. 5.

Additionally, we studied the behavior of different INR activations under noise-induced corruption of the input point clouds (Tab. 2) and surface margin extrapolation (Tab. 3), a crucial property of our pipeline. For the latter, given a point cloud patch of a freeform surface, we select a subset of points from the center of the patch for fitting, leaving a margin of $c\%$ for extrapolation evaluation. Our INR design consistently performs well across various settings.

4.4. Reconstructed CAD evaluation

We conduct an evaluation of Point2CAD in terms of the reconstruction quality for surfaces, edges, and corners. We compare our results to those obtained with ComplexGen [9], which is also capable of generating all those geometric elements and their topological relationships. However, it should be noted that the models generated by ComplexGen are often invalid, in the sense that the predicted geometry and topology are inconsistent. We are not aware of an easy way to rectify this issue. A possible solution could be only to utilize the surfaces generated by ComplexGen and apply our method to refit and intersect them.

As seen in Tab. 4, Point2CAD outperforms ComplexGen in most metrics. Furthermore, Point2CAD achieves excellent results when fed with ground truth segmentations, indicating that it will likely further improve as better point cloud segmentation engines become available. We also observed that our P-coverage results were slightly inferior to ComplexGen’s. This can be attributed to the fact that, in the final step of our reconstruction, we trim the margins of the individual surfaces to establish a clear and well-defined boundary. As a result, the coverage of our reconstructed CAD models is somewhat reduced compared to ComplexGen, which does not perform this step and typically generates more surfaces than the ground truth has. The presence of numerous spurious surfaces naturally inflates the P-coverage. This limitation of the metric is confirmed by the fact that we obtain higher P-coverage if we evaluate Point2CAD without clipping the surfaces at the edges (Point2CAD+ParseNet: 0.956, Point2CAD+HPNet: 0.953, Point2CAD+GT: 0.971).

We also employ several standard evaluation metrics, including precision, recall, and F-score, to measure the accuracy and completeness of the surfaces after the overall reconstruction of CAD under different distance thresh-

Table 5. Evaluation on CAD Reconstruction of *Surfaces*, *Edges* and *Corners* from the aspects of accuracy and completeness. The surface evaluation assesses the reconstruction performance in geometric terms, while these metrics on edges and corners reflect the performance of reconstruction in terms of topological properties. Given that the distance calculation is based on sampling, we adopt various threshold values for evaluating surfaces, edges and corners. Specifically, the threshold for a single-point corner is expected to be smaller than that of edges, while the threshold for edges should be smaller than that of surfaces, which aligns with the intuitive expectations in CAD evaluation.

Surfaces		$\theta_{\text{surface}} = 0.08$			$\theta_{\text{surface}} = 0.06$			$\theta_{\text{surface}} = 0.03$		
Method	Segmentation	precision \uparrow	recall \uparrow	F-score \uparrow	precision \uparrow	recall \uparrow	F-score \uparrow	precision \uparrow	recall \uparrow	F-score \uparrow
ComplexGen	N/A	0.732	0.732	0.731	0.633	0.641	0.637	0.370	0.388	0.379
Point2CAD	ParseNet	0.817	0.750	0.782	0.764	0.693	0.727	0.624	0.562	0.591
Point2CAD	HPNet	0.836	0.745	0.788	0.796	0.695	0.742	0.697	0.593	0.641
Point2CAD	GT	0.976	0.920	0.947	0.962	0.898	0.929	0.873	0.801	0.836
Edges		$\theta_{\text{edge}} = 0.05$			$\theta_{\text{edge}} = 0.03$			$\theta_{\text{edge}} = 0.02$		
Method	Segmentation	precision \uparrow	recall \uparrow	F-score \uparrow	precision \uparrow	recall \uparrow	F-score \uparrow	precision \uparrow	recall \uparrow	F-score \uparrow
ComplexGen	N/A	0.620	0.576	0.597	0.421	0.397	0.409	0.291	0.279	0.285
Point2CAD	ParseNet	0.636	0.596	0.615	0.523	0.486	0.504	0.440	0.414	0.427
Point2CAD	HPNet	0.637	0.586	0.611	0.526	0.473	0.498	0.449	0.397	0.421
Point2CAD	GT	0.863	0.774	0.816	0.766	0.673	0.717	0.686	0.597	0.639
Corners		$\theta_{\text{corner}} = 0.03$			$\theta_{\text{corner}} = 0.02$			$\theta_{\text{corner}} = 0.01$		
Method	Segmentation	precision \uparrow	recall \uparrow	F-score \uparrow	precision \uparrow	recall \uparrow	F-score \uparrow	precision \uparrow	recall \uparrow	F-score \uparrow
ComplexGen	N/A	0.667	0.633	0.650	0.483	0.454	0.468	0.217	0.203	0.210
Point2CAD	ParseNet	0.661	0.641	0.651	0.553	0.529	0.541	0.398	0.392	0.395
Point2CAD	HPNet	0.646	0.613	0.629	0.549	0.521	0.535	0.389	0.386	0.388
Point2CAD	GT	0.780	0.696	0.736	0.704	0.662	0.661	0.581	0.515	0.546

olds θ_{surface} , see Tab. 5 (top). The results suggest that Point2CAD with HPNet segmentation backbone achieves the highest overall performance, while reconstruction based on the ground truth point cloud segmentation provides the best results, as expected.

We utilize the precision, recall, and F-score of the reconstructed curves and corners as metrics to evaluate the topological reconstruction performance; results are shown in Tab. 5 (middle, bottom).

5. Conclusion

We present Point2CAD, a comprehensive and versatile approach for reverse engineering CAD models from point clouds, which can yield various types of surfaces and topologically correct reconstructions. The proposed method segments point clouds with a pre-trained segmentation backbone but then employs learning-free optimization methods to fit geometric primitives, including freeform surfaces that are optimized with a novel neural surface estimation scheme. The recovered surfaces are analytically intersected to obtain the edges and corners of the model. Empirically, our proposed method qualitatively and quantitatively outperforms prior art on the popular ABC dataset of CAD models. In our view, these results support our hypothesis that the main role of machine learning in the context of

CAD reconstruction is actually segmentation, whereas subsequent steps are adequately solved by classical geometric operations and may not always require learning.

Limitations While our method offers a solution to CAD model reconstruction, it is subject to certain limitations. Notably, our method is sensitive to the quality of the point cloud segmentation. A number of outliers in the prediction of points belonging to the same surface can adversely impact the subsequent surface fitting and topology reconstruction. Despite the fact that our method is not end-to-end-trainable, the main guideline in our design is to divide the pipeline into isolated, accountable steps and only employ neural methods where needed. While this runs against a puristic deep learning spirit, our modular pipeline outperforms the end-to-end learned solutions. We see this as an important message, although we do not rule out that it could serve as a foundation for future end-to-end works.

Future work One promising direction would be to exploit the proposed INR fitting scheme in a feedback loop of any other algorithm to drive point cloud partitioning into clusters, thus lifting the dependence of the reconstruction pipeline on the quality of point cloud segmentation. Differently, bringing back learned priors and the end-to-end property should make it possible to further improve the results.

References

- [1] Shashank Alai. A review of 3d design parameterization using reverse engineering. *International Journal of Emerging Technology and Advanced Engineering*, 3(10):171–179, 2013. 2
- [2] Pál Benkő, Ralph R Martin, and Tamás Várady. Algorithms for reverse engineering boundary representation models. *Computer-Aided Design*, 33(11):839–851, 2001. 2
- [3] Francesco Buonamici, Monica Carfagni, Rocco Furferi, Lapo Governi, Alessandro Lapini, and Yary Volpe. Reverse engineering modeling methods and tools: a survey. *Computer-Aided Design and Applications*, 15(3):443–464, 2018. 2
- [4] Alexandre Durupt, Sebastien Remy, Guillaume Ducellier, and Benoît Eynard. From a 3d point cloud to an engineering CAD model: a knowledge-product-based approach for reverse engineering. *Virtual and Physical Prototyping*, 3(2): 51–59, 2008. 2
- [5] Gerald Farin. *Curves and surfaces for CAGD: a practical guide*. Morgan Kaufmann, 2002. 2
- [6] Karl Pearson F.R.S. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901. 3
- [7] Kunihiro Fukushima. Cognitron: A self-organizing multilayered neural network. *Biological cybernetics*, 20(3-4): 121–136, 1975. 4
- [8] Jun Gao, Chengcheng Tang, Vignesh Ganapathi-Subramanian, Jiahui Huang, Hao Su, and Leonidas J Guibas. DeepSpline: Data-driven reconstruction of parametric curves and surfaces. *arXiv preprint arXiv:1901.03781*, 2019. 2
- [9] Haoxiang Guo, Shilin Liu, Hao Pan, Yang Liu, Xin Tong, and Baining Guo. ComplexGen: CAD reconstruction by B-rep chain complex generation. *ACM TOG*, 41(4):1–18, 2022. 3, 6, 7
- [10] Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R Martin, and Shi-Min Hu. PCT: Point cloud transformer. *Computational Visual Media*, 7:187–199, 2021. 2
- [11] Chen He and Chunmeng Wang. A survey on segmentation of 3d models. *Wireless Personal Communications*, 102:3835–3842, 2018. 2
- [12] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 4
- [13] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989. 4
- [14] Binh-Son Hua, Minh-Khoi Tran, and Sai-Kit Yeung. Point-wise convolutional neural networks. In *CVPR*, 2018. 2
- [15] Jingwei Huang, Yanfeng Zhang, and Mingwei Sun. PrimitiveNet: Primitive instance segmentation with local primitive embedding under adversarial metric. In *ICCV*, 2021. 2
- [16] Pradeep Kumar Jayaraman, Aditya Sanghi, Joseph G Lambourne, Karl DD Willis, Thomas Davies, Hooman Shayani, and Nigel Morris. UV-Net: Learning from boundary representations. In *CVPR*, 2021. 2
- [17] Adrien Kaiser, Jose Alonso Ybanez Zepeda, and Tamy Boubekeur. A survey of simple geometric primitives detection methods for captured 3d data, 2019. 2
- [18] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. 5
- [19] Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. ABC: A big CAD model dataset for geometric deep learning. In *CVPR*, 2019. 2, 5
- [20] Mark A Kramer. Nonlinear principal component analysis using autoassociative neural networks. *AIChE journal*, 37(2):233–243, 1991. 3, 4
- [21] Joseph G Lambourne, Karl DD Willis, Pradeep Kumar Jayaraman, Aditya Sanghi, Peter Meltzer, and Hooman Shayani. BRepNet: A topological message passing system for solid models. In *CVPR*, 2021. 2
- [22] Guohao Li, Matthias Muller, Ali Thabet, and Bernard Ghanem. DeepGCNs: Can GCNs go as deep as CNNs? In *ICCV*, 2019. 2
- [23] Lingxiao Li, Minhyuk Sung, Anastasia Dubrovina, Li Yi, and Leonidas J Guibas. Supervised fitting of geometric primitives to 3d point clouds. In *CVPR*, 2019. 2
- [24] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. PointCNN: Convolution on x-transformed points. *NeurIPS*, 2018. 2
- [25] Jin Liu. An adaptive process of reverse engineering from point clouds to CAD models. *International Journal of Computer Integrated Manufacturing*, 33(9):840–858, 2020. 2
- [26] Zhijian Liu, Haotian Tang, Yujun Lin, and Song Han. Point-voxel CNN for efficient 3d deep learning. *NeurIPS*, 2019. 2
- [27] Marios Loizou, Melinos Averkiou, and Evangelos Kalogerakis. Learning part boundaries from 3d point clouds, 2020. 2
- [28] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction, 2018. 3, 4
- [29] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM*, 65(1):99–106, 2021. 3, 4
- [30] Mojtaba Mozaffar, Shuheng Liao, Xiaoyu Xie, Sourav Saha, Chanwook Park, Jian Cao, Wing Kam Liu, and Zhengtao Gan. Mechanistic artificial intelligence (mechanistic-AI) for modeling, design, and control of advanced manufacturing processes: Current state and perspectives. *Journal of Materials Processing Technology*, page 117485, 2021. 2
- [31] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 2
- [32] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 2017. 2
- [33] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and

- Aaron Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pages 5301–5310. PMLR, 2019. 4
- [34] Dario Rethage, Johanna Wald, Jurgen Sturm, Nassir Navab, and Federico Tombari. Fully-convolutional point networks for large-scale point clouds. In *ECCV*, 2018. 2
- [35] Chiara Romanengo, Andrea Raffo, Silvia Biasotti, and Bianca Falcidieno. Recognising geometric primitives in 3d point clouds of mechanical CAD objects. *Computer-Aided Design*, page 103479, 2023. 2
- [36] Sam T. Roweis and Lawrence K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000. 3
- [37] Tsahi Saporta and Andrei Sharf. Unsupervised recursive deep fitting of 3d primitives to points. *Computers & Graphics*, 102:289–299, 2022. 2
- [38] Gopal Sharma, Rishabh Goyal, Difan Liu, Evangelos Kalogerakis, and Subhansu Maji. Neural shape parsers for constructive solid geometry. *IEEE TPAMI*, 44(5):2628–2640, 2020. 2
- [39] Gopal Sharma, Difan Liu, Subhansu Maji, Evangelos Kalogerakis, Siddhartha Chaudhuri, and Radomír Měch. ParSeNet: A parametric surface fitting network for 3d point clouds. In *ECCV*, 2020. 2, 4, 5, 6
- [40] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems*, pages 7462–7473. Curran Associates, Inc., 2020. 3, 4
- [41] Joshua B. Tenenbaum, Vin de Silva, and John C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000. 3
- [42] Mikaela Angelina Uy, Yen-Yu Chang, Minhyuk Sung, Purvi Goel, Joseph G Lambourne, Tolga Birdal, and Leonidas J Guibas. Point2cyl: Reverse engineering 3d objects from point clouds to extrusion cylinders. In *CVPR*, 2022. 3
- [43] Subeesh Vasu, Nicolas Talabot, Artem Lukoianov, Pierre Baqué, Jonathan Donier, and Pascal Fua. HybridSDF: Combining deep implicit shapes and geometric primitives for 3d shape representation and manipulation. *arXiv preprint arXiv:2109.10767*, 2021. 2
- [44] Jun Wang, Dongxiao Gu, Zeyun Yu, Changbai Tan, and Laishui Zhou. A framework for 3d model reconstruction in reverse engineering. *Computers & Industrial Engineering*, 63(4):1189–1200, 2012. 2
- [45] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *ACM TOG*, 38(5): 1–12, 2019. 2
- [46] Karl DD Willis, Pradeep Kumar Jayaraman, Hang Chu, Yunsheng Tian, Yifei Li, Daniele Grandi, Aditya Sanghi, Linh Tran, Joseph G Lambourne, Armando Solar-Lezama, et al. JoinABLE: Learning bottom-up assembly of parametric CAD joints. In *CVPR*, 2022. 2
- [47] Siming Yan, Zhenpei Yang, Chongyang Ma, Haibin Huang, Etienne Vouga, and Qixing Huang. HPNet: Deep primitive segmentation using hybrid representations. In *ICCV*, 2021. 2, 5, 6
- [48] Fenggen Yu, Zhiqin Chen, Manyi Li, Aditya Sanghi, Hooman Shayani, Ali Mahdavi-Amiri, and Hao Zhang. CAPRI-net: learning compact CAD shapes with adaptive primitive assembly. In *CVPR*, 2022. 2
- [49] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *ICCV*, 2021. 2