

PARSENET: A Parametric Surface Fitting Network for 3D Point Clouds

Gopal Sharma¹ , Difan Liu¹ , Subhansu Maji¹ ,
Evangelos Kalogerakis¹ , Siddhartha Chaudhuri^{2,3}, and Radomír Měch²
¹{gopalsharma, dliu, smaji, kalo}@cs.umass.edu, ²{sidch, rmech}@adobe.com

¹University of Massachusetts Amherst ²Adobe Research ³IIT Bombay

Abstract. We propose a novel, end-to-end trainable, deep network called PARSENET that decomposes a 3D point cloud into parametric surface patches, including B-spline patches as well as basic geometric primitives. PARSENET is trained on a large-scale dataset of man-made 3D shapes and captures high-level semantic priors for shape decomposition. It handles a much richer class of primitives than prior work, and allows us to represent surfaces with higher fidelity. It also produces repeatable and robust parametrizations of a surface compared to purely geometric approaches. We present extensive experiments to validate our approach against analytical and learning-based alternatives. Our source code is publicly available at: <https://hippogriff.github.io/parsenet>.

1 Introduction

3D point clouds can be rapidly acquired using 3D sensors or photogrammetric techniques. However, they are rarely used in this form in design and graphics applications. Observations from the computer-aided design and modeling literature [3, 4, 14, 16] suggest that designers often model shapes by constructing several non-overlapping patches placed seamlessly. The advantage of using several patches over a single continuous patch is that a much more diverse variety of geometric features and surface topologies can be created. The decomposition also allows easier interaction and editing. The goal of this work is to automate the time-consuming process of converting a 3D point cloud into a piecewise parametric surface representation as seen in Figure 1.

An important question is how surface patches should be represented. Patch representations in CAD and graphics are based on well-accepted geometric properties: (a) *continuity* in their tangents, normals, and curvature, making patches appear smooth, (b) *editability*, such that they can easily be modified based on a few intuitive degrees of freedom (DoFs), *e.g.*, control points or axes, and (c) *flexibility*, so that a wide variety of surface geometries can be captured. Towards this goal, we propose PARSENET, a **parametric surface fitting network** architecture which produces a compact, editable representation of a point cloud as an assembly of geometric primitives, including open or closed B-spline patches.

PARSENET models a richer class of surfaces than prior work which *only* handles basic geometric primitives such as planes, cuboids and cylinders [11,




Fig. 1: PARSENET decomposes point clouds (top row) into collections of assembled parametric surface patches including B-spline patches (bottom row). On the right, a shape is edited using the inferred parametrization.

[12, 17, 24]. While such primitives are continuous and editable representations, they lack the richness and flexibility of spline patches which are widely used in shape design. PARSENET includes a novel neural network (SPLINENET) to estimate an open or closed B-spline model of a point cloud patch. It is part of a *fitting module* (Section 3.2) which can also fit other geometric primitive types. The fitting module receives input from a *decomposition module*, which partitions a point cloud into segments, after which the fitting module estimates shape parameters of a predicted primitive type for each segment (Section 3.1). The entire pipeline, shown in Figure 2, is fully differentiable and trained end-to-end (Section 4). An optional geometric postprocessing step further refines the output.

Compared to purely analytical approaches, PARSENET produces decompositions that are more consistent with high-level semantic priors, and are more robust to point density and noise. To train and test PARSENET, we leverage a recent dataset of man-made parts [8]. Extensive evaluations show that PARSENET outperforms baselines (RANSAC and SPFN [11]) by 14.93% and 13.13% respectively for segmenting a point cloud into patches, and by 50%, and 47.64% relative error respectively for parametrizing each patch for surface reconstruction (Section 5).

To summarize, our contributions are:

- The first proposed end-to-end differentiable approach for representing a raw 3D point cloud as an assembly of parametric primitives *including* spline patches.
- Novel decomposition and primitive fitting modules, including SPLINENET, a fully-differentiable network to fit a cubic B-spline patch to a set of points.
- Evaluation of our framework vs prior analytical and learning-based methods.

2 Related Work

Our work builds upon related research on parametric surface representations and methods for primitive fitting. We briefly review relevant work in these areas. Of course, we also leverage extensive prior work on neural networks for general shape processing: see recent surveys on the subject [1].

Parametric surfaces. A parametric surface is a (typically diffeomorphic) mapping from a (typically compact) subset of \mathbb{R}^2 to \mathbb{R}^3 . While most of the geometric

primitives used in computer graphics (spheres, cuboids, meshes etc) can be represented parametrically, the term most commonly refers to curved surfaces used in engineering CAD modelers, represented as spline patches [3]. There are a variety of formulations – *e.g.* Bézier patches, B-spline patches, NURBS patches – with slightly different characteristics, but they all construct surfaces as weighted combinations of control parameters, typically the positions of a sparse grid of points which serve as editing handles. More specifically, a B-spline patch is a smoothly curved, bounded, parametric surface, whose shape is defined by a sparse grid of control points $\mathbf{C} = \{\mathbf{c}_{p,q}\}$. The surface point with parameters $(u, v) \in [u_{\min}, u_{\max}] \times [v_{\min}, v_{\max}]$ and *basis functions* [3] $b_p(u), b_q(v)$ is given by:

$$\mathbf{s}(u, v) = \sum_{p=1}^P \sum_{q=1}^Q b_p(u) b_q(v) \mathbf{c}_{p,q} \quad (1)$$

Please refer to supplementary material for more details on B-spline patches.

Fitting geometric primitives. A variety of analytical (*i.e.* not learning-based) algorithms have been devised to approximate raw 3D data as a collection of geometric primitives: dominant themes include Hough transforms, RANSAC and clustering. The literature is too vast to cover here, we recommend the comprehensive survey of Kaiser *et al.* [7]. In the particular case of NURBS patch fitting, early approaches were based on user interaction or hand-tuned heuristics to extract patches from meshes or point clouds [2, 6, 10]. In the rest of this section, we briefly review recent methods that *learn* how to fit primitives to 3D data.

Several recent papers [20, 22, 24, 27] also try to approximate 3D shapes as unions of cuboids or ellipsoids. Paschalidou *et al.* [12, 13] extended this to superquadrics. Sharma *et al.* [17, 18] developed a neural parser that represents a test shape as a collection of basic primitives (spheres, cubes, cylinders) combined with boolean operations. Tian *et al.* [23] handled more expressive construction rules (*e.g.* loops) and a wider set of primitives. Because of the choice of simple primitives, such models are naturally limited in how well they align to complex input objects, and offer less flexible and intuitive parametrization for user edits.

More relevantly to our goal of modeling arbitrary curved surfaces, Gao *et al.* [5] parametrize 3D point clouds as extrusions or surfaces of revolution, generated by B-spline cross-sections detected by a 2D network. This method requires translational/rotational symmetry, and does not apply to general curved patches. Li *et al.* [11] proposed a supervised method to fit primitives to 3D point clouds, first predicting per-point segment labels, primitive types and normals, and then using a differential module to estimate primitive parameters. While we also chain segmentation with fitting in an end-to-end way, we differ from Li *et al.* in two important ways. First, our differentiable metric-learning segmentation produces improved results (Table 1). Second, a major goal (and technical challenge) for us is to significantly improve expressivity and generality by incorporating B-spline patches: we achieve this with a novel differentiable spline-fitting network. In a complementary direction, Yumer *et al.* [26] developed a neural network for fitting a single NURBS patch to an unstructured point cloud. While the goal is similar




Fig. 2: **Overview of PARSENET pipeline.** (1) The *decomposition module* (Section 3.1) takes a 3D point cloud (with optional normals) and decomposes it into segments labeled by primitive type. (2) The *fitting module* (Section 3.2) predicts parameters of a primitive that best approximates each segment. It includes a novel SPLINENET to fit B-spline patches. The two modules are jointly trained end-to-end. An optional postprocess module (Section 3.3) refines the output.

to our spline-fitting network, it is not combined with a decomposition module that jointly learns how to express a shape with *multiple* patches covering different regions. Further, their fitting module has several non-trainable steps which are not obviously differentiable, and hence cannot be used in our pipeline.

3 Method

The goal of our method is to reconstruct an input point cloud by predicting a set of parametric patches closely approximating its underlying surface. The first stage of our architecture is a *neural decomposition module* (Fig. 2) whose goal is to segment the input point cloud into regions, each labeled with a parametric patch type. Next, we incorporate a *fitting module* (Fig. 2) that predicts each patch’s shape parameters. Finally, an optional post-processing *geometric optimization* step refines the patches to better align their boundaries for a seamless surface.

The input to our pipeline is a set of points $\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^N$, represented either as 3D positions $\mathbf{p}_i = (x, y, z)$, or as 6D position + normal vectors $\mathbf{p}_i = (x, y, z, n_x, n_y, n_z)$. The output is a set of surface patches $\{\mathbf{s}_k\}$, reconstructing the input point cloud. The number of patches is automatically determined. Each patch is labeled with a type t_k , one of: sphere, plane, cone, cylinder, open/closed B-spline patch. The architecture also outputs a real-valued vector for each patch defining its geometric parameters, *e.g.* center and radius for spheres, or B-spline control points and knots.

3.1 Decomposition module

The first module (Fig. 2) decomposes the point cloud \mathbf{P} into a set of segments such that each segment can be reliably approximated by one of the abovementioned surface patch types. To this end, the module first embeds the input points into a representation space used to reveal such segments. As discussed in Section 4, the representations are learned using metric learning, such that points belonging to the same patch are embedded close to each other, forming a distinct cluster.

Embedding network. To learn these point-wise representations, we incorporate edge convolution layers (EdgeConv) from DGCNN [25]. Each EdgeConv layer performs a graph convolution to extract a representation of each point with an MLP on the input features of its neighborhood. The neighborhoods are dynamically defined via nearest neighbors in the input feature space. We stack 3 EdgeConv layers, each extracting a 256-D representation per point. A max-pooling layer is also used to extract a global 1024-D representation for the whole point cloud. The global representation is tiled and concatenated with the representations from all three EdgeConv layers to form intermediate point-wise (1024 + 256)-D representations $\mathbf{Q} = \{\mathbf{q}_i\}$ encoding both local and global shape information. We found that a global representation is useful for our task, since it captures the overall geometric shape structure, which is often correlated with the number and type of expected patches. This representation is then transformed through fully connected layers and ReLUs, and finally normalized to unit length to form the point-wise embedding $\mathbf{Y} = \{\mathbf{y}_i\}_{i=1}^N$ (128-D) lying on the unit hypersphere.

Clustering. A mean-shift clustering procedure is applied on the point-wise embedding to discover segments. The advantage of mean-shift clustering over other alternatives (*e.g.*, k-means or mixture models) is that it does not require the target number of clusters as input. Since different shapes may comprise different numbers of patches, we let mean-shift produce a cluster count tailored for each input. Like the pixel grouping of [9], we implement mean-shift iterations as differentiable recurrent functions, allowing back-propagation. Specifically, we initialize mean-shift by setting all points as seeds $\mathbf{z}_i^{(0)} = \mathbf{y}_i, \forall y_i \in R^{128}$. Then, each mean-shift iteration t updates each point’s embedding on the unit hypersphere:

$$\mathbf{z}_i^{(t+1)} = \sum_{j=1}^N \mathbf{y}_j g(\mathbf{z}_i^{(t)}, \mathbf{y}_j) / \left(\sum_{j=1}^N g(\mathbf{z}_i^{(t)}, \mathbf{y}_j) \right) \quad (2)$$

where the pairwise similarities $g(\mathbf{z}_i^{(t)}, \mathbf{y}_j)$ are based on a von Mises-Fisher kernel with bandwidth β : $g(\mathbf{z}_i, \mathbf{y}_j) = \exp(\mathbf{z}_i^T \mathbf{y}_j / \beta^2)$ (iteration index dropped for clarity). The embeddings are normalized to unit vectors after each iteration. The bandwidth for each input point cloud is set as the average distance of each point to its 150th neighboring point in the embedding space [19]. The mean-shift iterations are repeated until convergence (this occurs around 50 iterations in our datasets). We extract the cluster centers using non-maximum suppression: starting with the point with highest density, we remove all points within a distance β , then repeat.

Points are assigned to segments based on their nearest cluster center. The point memberships are stored in a matrix \mathbf{W} , where $\mathbf{W}[i, k] = 1$ means point i belongs to segment k , and 0 means otherwise. The memberships are passed to the fitting module to determine a parametric patch per segment. During training, we use soft memberships for differentiating this step (more details in Section 4.3).

Segment Classification. To classify each segment, we pass the per-point representation \mathbf{q}_i , encoding local and global geometry, through fully connected layers and ReLUs, followed by a softmax for a per-point probability $P(t_i = l)$, where l is a patch type (*i.e.*, sphere, plane, cone, cylinder, open/closed B-spline patch). The segment’s patch type is determined through majority voting over all its points.

3.2 Fitting module

The second module (Fig. 2) aims to fit a parametric patch to each predicted segment of the point cloud. To this end, depending on the segment type, the module estimates the shape parameters of the surface patch.

Basic primitives. Following Li *et al.* [11], we estimate the shape of basic primitives with least-squares fitting. This includes center and radius for spheres; normal and offset for planes; center, direction and radius for cylinders; and apex, direction and angle for cones. We also follow their approach to define primitive boundaries.

B-Splines. Analytically parametrizing a set of points as a spline patch in the presence of noise, sparsity and non-uniform sampling, can be error-prone. Instead, predicting control points directly with a neural network can provide robust results. We propose a neural network SPLINENET, that inputs points of a segment, and outputs a fixed size control-point grid. A stack of three EdgeConv layers produce point-wise representations concatenated with a global representation extracted from a max-pooling layer (as for decomposition, but weights are not shared). This equips each point i in a segment with a 1024-D representation ϕ_i . A segment’s representation is produced by max-pooling over its points, as identified through the membership matrix \mathbf{W} extracted previously:

$$\phi_k = \max_{i=1\dots N} (\mathbf{W}[i, k] \cdot \phi_i). \quad (3)$$

Finally, two fully-connected layers with ReLUs transform ϕ_k to an initial set of 20×20 control points \mathbf{C} unrolled into a 1200-D output vector. For a segment with a small number of points, we upsample the input segment (with nearest neighbor interpolation) to 1600 points. This significantly improved performance for such segments (Table 2). For closed B-spline patches, we wrap the first row/column of control points. Note that the network parameters to produce open and closed B-splines are not shared. Fig. 5 visualizes some predicted B-spline surfaces.

3.3 Post-processing module

SPLINENET produces an initial patch surface that approximates the points belonging to a segment. However, patches might not entirely cover the input

point cloud, and boundaries between patches are not necessarily well-aligned. Further, the resolution of the initial control point grid (20×20) can be further adjusted to match the desired surface resolution. As a post-processing step, we perform an optimization to produce B-spline surfaces that better cover the input point cloud, and refine the control points to achieve a prescribed fitting tolerance.

Optimization. We first create a grid of 40×40 points on the initial B-spline patch by uniformly sampling its UV parameter space. We tessellate them into quads. Then we perform a maximal matching between the quad vertices and the input points of the segment, using the Hungarian algorithm with L2 distance costs. We then perform an as-rigid-as-possible (ARAP) [21] deformation of the tessellated surface towards the matched input points. ARAP is an iterative, detail-preserving method to deform a mesh so that selected vertices (pivots) achieve targets position, while promoting locally rigid transformations in one-ring neighborhoods (instead of arbitrary ones causing shearing/stretching). We use the boundary vertices of the patch as pivots so that they move close to their matched input points. Thus, we promote coverage of input points by the B-spline patches. After the deformation, the control points are re-estimated with least-squares [14].

Refinement of B-spline control points. After the above optimization, we again perform a maximal matching between the quad vertices and the input points of the segment. As a result, the input segment points acquire 2D parameter values in the patch’s UV parameter space, which can be used to re-fit any other grid of control points [14]. In our case, we iteratively upsample the control point grid by a factor of 2 until a fitting tolerance, measured via Chamfer distance, is achieved. If the tolerance is satisfied by the initial control point grid, we can similarly downsample it iteratively. In our experiments, we set the fitting tolerance to 5×10^{-4} . In Fig. 5 we show the improvements from the post-processing step.

4 Training

To train the neural decomposition and fitting modules of our architecture, we use supervisory signals from a dataset of 3D shapes modeled through a combination of basic geometric primitives and B-splines. Below we describe the dataset, then we discuss the loss functions and the steps of our training procedure.

4.1 Dataset

The ABC dataset [8] provides a large source of 3D CAD models of mechanical objects whose file format stores surface patches and modeling operations that designers used to create them. Since our method is focused on predicting surface patches, and in particular B-spline patches, we selected models from this dataset that contain at least one B-spline surface patch. As a result, we ended up with a dataset of 32K models (24K, 4K, 4K train, test, validation sets respectively). We call this ABCPARTSDATASET. All shapes are centered in the origin and scaled so they lie inside unit cube. To train SPLINENET, we also extract 32K closed and




Fig. 3: **Standardization:** Examples of B-spline patches with a variable number of control points (shown in red), each standardized with 20×20 control points. Left: closed B-spline and Right: open B-spline. (Please zoom in.)

open B-spline surface patches each from ABC dataset and split them into 24K, 4K, 4K train, test, validation sets respectively. We call this SPLINEDATASET. We report the average number of different patch types in supplementary material.

Preprocessing. Based on the provided metadata in ABCPARTSDATASET, each shape can be rendered based on the collection of surface patches and primitives it contains (Figure 4). Since we assume that the inputs to our architecture are point clouds, we first sample each shape with 10K points randomly distributed on the shape surface. We also add noise in a uniform range $[-0.01, 0.01]$ along the normal direction. Normals are also perturbed with random noise in a uniform range of $[-3, 3]$ degrees from their original direction.

4.2 Loss functions

We now describe the different loss functions used to train our neural modules. The training procedure involving their combination is discussed in Section 4.3.

Embedding loss. To discover clusters of points that correspond well to surface patches, we use a metric learning approach. The point-wise representations \mathbf{Z} produced by our decomposition module after mean-shift clustering are learned such that point pairs originating from the same surface patch are embedded close to each other to favor a cluster formation. In contrast, point pairs originating from different surface patches are pushed away from each other. Given a triplet of points (a, b, c) , we use the triplet loss to learn the embeddings:

$$\ell_{emb}(a, b, c) = \max(0, \|\mathbf{z}_a - \mathbf{z}_b\|^2 - \|\mathbf{z}_a - \mathbf{z}_c\|^2 + \tau), \quad (4)$$

where τ the margin is set to 0.9. Given a triplet set \mathcal{T}_S sampled from each point set \mathcal{S} from our dataset \mathcal{D} , the embedding objective sums the loss over triplets:

$$L_{emb} = \sum_{\mathcal{S} \in \mathcal{D}} \frac{1}{|\mathcal{T}_S|} \sum_{(a,b,c) \in \mathcal{T}_S} \ell_{emb}(a, b, c). \quad (5)$$

Segment classification loss. To promote correct segment classifications according to our supported types, we use the cross entropy loss: $L_{class} = -\sum_{i \in \mathcal{S}} \log(p_i^t)$ where p_i^t is the probability of the i^{th} point of shape \mathcal{S} belonging to its ground truth type t , computed from our segment classification network.

Control point regression loss. This loss function is used to train SPLINET. As discussed in Section 3.2, SPLINET produces 20×20 control points per B-spline patch. We include a supervisory signal for this control point grid prediction. One issue is that B-spline patches have a variable number of control points in our dataset. Hence we reparametrize each patch by first sampling $M = 3600$ points and estimating a new 20×20 reparametrization using least-squares fitting [10, 14], as seen in the Figure 3. In our experiments, we found that this standardization produces no practical loss in surface reconstructions in our dataset. Finally, our reconstruction loss should be invariant to flips or swaps of control points grid in u and v directions. Hence we define a loss that is invariant to such permutations:

$$L_{cp} = \sum_{\mathcal{S} \in \mathcal{D}} \frac{1}{|\mathcal{S}^{(b)}|} \sum_{\mathbf{s}_k \in \mathcal{S}^{(b)}} \frac{1}{|\mathbf{C}_k|} \min_{\pi \in \Pi} \|\mathbf{C}_k - \pi(\hat{\mathbf{C}}_k)\|^2 \quad (6)$$

where $\mathcal{S}^{(b)}$ is the set of B-spline patches from shape \mathcal{S} , \mathbf{C}_k is the predicted control point grid for patch \mathbf{s}_k ($|\mathbf{C}_k| = 400$ control points), $\pi(\hat{\mathbf{C}}_k)$ is permutations of the ground-truth control points from the set Π of 8 permutations for open and 160 permutations for closed B-spline.

Laplacian loss. This loss is also specific to B-Splines using SPLINET. For each ground-truth B-spline patch, we uniformly sample ground truth surface, and measure the surface Laplacian capturing its second-order derivatives. We also uniformly sample the predicted patches and measure their Laplacians. We then establish Hungarian matching between sampled points in the ground-truth and predicted patches, and compare the Laplacians of the ground-truth points $\hat{\mathbf{r}}_m$ and corresponding predicted ones \mathbf{r}_n to improve the agreement between their derivatives as follows:

$$L_{lap} = \sum_{\mathcal{S} \in \mathcal{D}} \frac{1}{|\mathcal{S}^{(b)}| \cdot M} \sum_{\mathbf{s}_k \in \mathcal{S}^{(b)}} \sum_{\mathbf{r}_n \in \mathbf{s}_k} \|\mathcal{L}(\mathbf{r}_n) - \mathcal{L}(\hat{\mathbf{r}}_m)\|^2 \quad (7)$$

where $\mathcal{L}(\cdot)$ is the Laplace operator on patch points, and $M = 1600$ point samples.

Patch distance loss. This loss is applied to both basic primitive and B-splines patches. Inspired by [11], the loss measures average distances between predicted primitive patch s_k and uniformly sampled points from the ground truth patch as:

$$L_{dist} = \sum_{\mathcal{S} \in \mathcal{D}} \frac{1}{K_{\mathcal{S}}} \sum_{k=1}^{K_{\mathcal{S}}} \frac{1}{M_{\hat{\mathbf{s}}_k}} \sum_{n \in \hat{\mathbf{s}}_k} D^2(\mathbf{r}_n, \mathbf{s}_k), \quad (8)$$

where $K_{\mathcal{S}}$ is the number of predicted patches for shape \mathcal{S} , $M_{\hat{\mathbf{s}}_k}$ is number of sampled points \mathbf{r}_n from ground patch $\hat{\mathbf{s}}_k$, $D^2(\mathbf{r}_n, \mathbf{s}_k)$ is the squared distance from \mathbf{r}_n to the predicted primitive patch surface \mathbf{s}_k . These distances can be computed analytically for basic primitives [11]. For B-splines, we use an approximation based on Chamfer distance between sample points.

4.3 Training procedure

One possibility for training is to start it from scratch using a combination of all losses. Based on our experiments, we found that breaking the training procedure into the following steps leads to faster convergence and to better minima:

- We first pre-train the networks of the decomposition module using ABCPARTS-DATASET with the sum of embedding and classification losses: $L_{emb} + L_{class}$. Both losses are necessary for point cloud decomposition and classification.
- We then pre-train the SPLINENET using SPLINEDATASET for control point prediction exclusively on B-spline patches using $L_{cp} + L_{lap} + L_{dist}$. We note that we experimented training the B-spline patch prediction only with the patch distance loss L_{dist} but had worse performance. Using both the L_{cp} and L_{lap} loss yielded better predictions as shown in Table 2.
- We then jointly train the decomposition and fitting module end-to-end with all the losses. To allow backpropagation from the primitives and B-splines fitting to the embedding network, the mean shift clustering is implemented as a recurrent module (Equation 2). For efficiency, we use 5 mean-shift iterations during training. It is also important to note that during training, Equation 3 uses *soft* point-to-segment memberships, which enables backpropagation from the fitting module to the decomposition module and improves reconstructions. The soft memberships are computed based on the point embeddings $\{\mathbf{z}_i\}$ (after the mean-shift iterations) and cluster center embedding $\{\mathbf{z}_k\}$ as follows:

$$\mathbf{W}[i, k] = \frac{\exp(\mathbf{z}_k^T \mathbf{z}_i / \beta^2)}{\sum_{k'} \exp(\mathbf{z}_{k'}^T \mathbf{z}_i / \beta^2)} \quad (9)$$

Please see supplementary material for more implementation details.

5 Experiments

Our experiments compare our approach to alternatives in three parts: (a) evaluation of the quality of segmentation and segment classification (Section 5.1), (b) evaluation of B-spline patch fitting, since it is a major contribution of our work (Section 5.2), and (c) evaluation of overall reconstruction quality (Section 5.3). We include evaluation metrics and results for each of the three parts next.

5.1 Segmentation and labeling evaluation

Evaluation metrics. We use the following metrics for evaluating the point cloud segmentation and segment labeling based on the test set of ABCPARTSDATASET:

- **Segmentation mean IOU** (“seg mIOU”): this metric measures the similarity of the predicted segments with ground truth segments. Given the ground-truth point-to-segment memberships $\tilde{\mathbf{W}}$ for an input point cloud, and the predicted ones \mathbf{W} , we measure: $\frac{1}{K} \sum_{k=1}^K IOU(\tilde{\mathbf{W}}[:, k], h(\mathbf{W}[:, k]))$ where h represents a membership conversion into a one-hot vector, and K is the number of ground-truth segments.

Method	Input	seg iou	label iou	res (all)	res (geom)	res (spline)	P cover
NN	p	54.10	61.10	-	-	-	-
RANSAC	p+n	67.21	-	0.0220	0.0220	-	83.40
SPFN	p	47.38	68.92	0.0238	0.0270	0.0100	86.66
SPFN	p+n	69.01	79.94	0.0212	0.0240	0.0136	88.40
PARSENET	p	71.32	79.61	0.0150	0.0160	0.0090	87.00
PARSENET	p+n	81.20	87.50	0.0120	0.0123	0.0077	92.00
PARSENET + e2e	p+n	82.14	88.60	0.0118	0.0120	0.0076	92.30
PARSENET + e2e + opt	p+n	82.14	88.60	0.0111	0.0120	0.0068	92.97

Table 1: **Primitive fitting on ABCPARTSDATASET.** We compare PARSENET with nearest neighbor (NN), RANSAC [15], and SPFN [11]. We show results with points (p) and points and normals (p+n) as input. The last two rows shows our method with end-to-end training and post-process optimization. We report ‘seg iou’ and ‘label iou’ metric for segmentation task. We report the residual error (res) on all, geometric and spline primitives, and the coverage metric for fitting.

- **Segment labeling IOU** (“label mIOU”): this metric measures the classification accuracy of primitive type prediction averaged over segments: $\frac{1}{K} \sum_{k=1}^K \mathcal{I}[t_k = \hat{t}_k]$ where t_k and \hat{t}_k is the predicted and ground truth primitive type respectively for k^{th} segment and \mathcal{I} is an indicator function.

We use Hungarian matching to find correspondences between predicted segments and ground-truth segments.

Comparisons. We first compare our method with a nearest neighbor (NN) baseline: for each test shape, we find its most similar shape from the training set using Chamfer distance. Then for each point on the test shape, we transfer the labels and primitive type from its closest point in \mathcal{R}^3 on the retrieved shape.

We also compare against efficient RANSAC algorithm [15]. The algorithm only handles basic primitives (cylinder, cone, plane, sphere, and torus), and offers poor reconstruction of B-splines patches in our dataset. Efficient RANSAC requires per point normals, which we provide as the ground-truth normals. We run RANSAC 3 times and report the performance with best coverage.

We then compare against the supervised primitive fitting (SPFN) approach [11]. Their approach produces per point segment membership, and their network is trained to maximize relaxed IOU between predicted membership and ground truth membership, whereas our approach uses learned point embeddings and clustering with mean-shift clustering to extract segments. We train SPFN network using their provided code on our training set using their proposed losses. We note that we include B-splines patches in their supported types. We train their network in two input settings: (a) the network takes only point positions as input, (b) it takes point and normals as input. We train our PARSENET on our training set in the same two settings using our loss functions.

The performance of the above methods are shown in Table 1. The lack of B-spline fitting hampers the performance of RANSAC. The SPFN method with




Fig. 4: Given the input point clouds with normals of the first row, we show surfaces produced by SPFN [11] (second row), PARSENET without post-processing optimization (third row), and full PARSENET including optimization (fourth row). The last row shows the ground-truth surfaces from our ABCPARTSDATASET.

points and normals as input performs better compared to using only points as input. Finally, PARSENET with only points as input performs better than all other alternatives. We observe further gains when including point normals in the input. Training PARSENET end-to-end gives 13.13% and 8.66% improvement in segmentation mIOU and label mIOU respectively over SPFN with points and normals as input. The better performance is also reflected in Figure 4, where our method reconstructs patches that correspond to more reasonable segmentations compared to other methods. In the supplementary material we evaluate methods on the TraceParts dataset [11], which contains only basic primitives (cylinder, cone, plane, sphere, torus). We outperform prior work also in this dataset.

5.2 B-Spline fitting evaluation

Evaluation metrics. We evaluate the quality of our predicted B-spline patches by computing the Chamfer distance between densely sampled points on the ground-truth B-spline patches and densely sampled points on predicted patches. Points are uniformly sampled based on the 2D parameter space of the patches. We use 2K samples. We use the test set of our SPLINEDATASET for evaluation.

Loss				Open splines		Closed splines	
cp	dist	lap	opt	w/ ups	w/o ups	w/ ups	w/o ups
✓				2.04	2.00	5.04	3.93
✓	✓			1.96	2.00	4.9	3.60
✓	✓	✓		1.68	1.59	3.74	3.29
✓	✓	✓	✓	0.92	0.87	0.63	0.81

Table 2: **Ablation study for B-spline fitting.** The error is measured using Chamfer Distance (CD is scaled by 100). The acronyms “cp”: control-points regression loss, “dist” means patch distance loss, and “lap” means Laplacian loss. We also include the effect of post-processing optimization “opt”. We report performance with and without upsampling (“ups”) for open and closed B-splines.




Fig. 5: **Qualitative evaluation of B-spline fitting.** From top to bottom: input point cloud, reconstructed surface by SPLINET, reconstructed surface by SPLINET with post-processing optimization, reconstruction by SPLINET with control point grid adjustment and finally ground truth surface. Effect of post process optimization is highlighted in red boxes.

Ablation study. We evaluate the training of SPLINET using various loss functions while giving 700 points per patch as input, in Table 2. All losses contribute to improvements in performance. Table 2 shows that upsampling is effective for closed splines. Figure 5 shows the effect of optimization to improve the alignment of patches and the adjustment of resolution in the control point grid. See supplementary material for more experiments on SPLINET’s robustness.

5.3 Reconstruction evaluation

Evaluation metrics. Given a point cloud $\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^N$, ground-truth patches $\{\cup_{k=1}^K \hat{\mathbf{s}}_k\}$ and predicted patches $\{\cup_{k=1}^K \mathbf{s}_k\}$ for a test shape in ABCPARTS-DATASET, we evaluate the patch-based surface reconstruction using the following:

- **Residual error** (“res”) measures the average distance of input points from the predicted primitives following [11]: $L_{dist} = \sum_{k=1}^K \frac{1}{M_k} \sum_{n \in \hat{\mathbf{s}}_k} D(\mathbf{r}_n, \mathbf{s}_k)$ where K is the number of segments, M_k is number of sampled points \mathbf{r}_n from ground patch $\hat{\mathbf{s}}_k$, $D(\mathbf{r}_n, \mathbf{s}_k)$ is the distance of \mathbf{r}_n from predicted primitive patch \mathbf{s}_k .
- **P-coverage** (“P-cover”) measures the coverage of predicted surface by the input surface also following [11]: $\frac{1}{N} \sum_{i=1}^N \mathbf{I}[\min_{k=1}^K D(\mathbf{p}_i, \mathbf{s}_k) < \epsilon]$ ($\epsilon = 0.01$).

We note that we use the matched segments after applying Hungarian matching algorithm, as in Section 5.1, to compute these metrics.

Comparisons. We report the performance of RANSAC for geometric primitive fitting tasks. Note that RANSAC produces a set of geometric primitives, along with their primitive type and parameters, which we use to compute the above metrics. Here we compare with the SPFN network [11] trained on our dataset using their proposed loss functions. We augment their per point primitive type prediction to also include open/closed B-spline type. Then for classified segments as B-splines, we use our SPLINENET to fit B-splines. For segments classified as geometric primitives, we use their geometric primitive fitting algorithm.

Results. Table 1 reports the performance of our method, SPFN and RANSAC. The residual error and P-coverage follows the trend of segmentation metrics. Interestingly, our method outperforms SPFN even for geometric primitive predictions (even without considering B-splines and our adaptation). Using points and normals, along with joint end-to-end training, and post-processing optimization offers the best performance for our method by giving 47.64% and 50% reduction in relative error in comparison to SPFN and RANSAC respectively.

6 Conclusion

We presented a method to reconstruct point clouds by predicting geometric primitives and surface patches common in CAD design. Our method effectively marries 3D deep learning with CAD modeling practices. Our architecture predictions are editable and interpretable. Modelers can refine our results based on standard CAD modeling operations. In terms of limitations, our method often makes mistakes for small parts, mainly because clustering merges them with bigger patches. In high-curvature areas, due to sparse sampling, PARSENET may produce more segments than ground-truth. Producing seamless boundaries is still a challenge due to noise and sparsity in our point sets. Generating training point clouds simulating realistic scan noise is another important future direction.

Acknowledgements. This research is funded in part by NSF (#1617333, #1749833) and Adobe. Our experiments were performed in the UMass GPU cluster funded by the MassTech Collaborative. We thank Matheus Gadelha for helpful discussions.

References

1. Ahmed, E., Saint, A., Shabayek, A.E.R., Cherenkova, K., Das, R., Gusev, G., Aouada, D., Ottersten, B.E.: Deep learning advances on different 3D data representations: A survey. *Computing Research Repository (CoRR)* **abs/1808.01462** (2019)
2. Eck, M., Hoppe, H.: Automatic reconstruction of B-spline surfaces of arbitrary topological type. In: *SIGGRAPH* (1996)
3. Farin, G.: *Curves and Surfaces for CAD*. Morgan Kaufmann, 5th edn. (2002)
4. Foley, J.D., van Dam, A., Feiner, S.K., Hughes, J.F.: *Computer Graphics: Principles and Practice*. Addison-Wesley Longman Publishing Co., Inc., USA, 2nd edn. (1990)
5. Gao, J., Tang, C., Ganapathi-Subramanian, V., Huang, J., Su, H., Guibas, L.J.: DeepSpline: Data-driven reconstruction of parametric curves and surfaces. *Computing Research Repository (CoRR)* **abs/1901.03781** (2019)
6. Hoppe, H., DeRose, T., Duchamp, T., Halstead, M., Jin, H., McDonald, J., Schweitzer, J., Stuetzle, W.: Piecewise smooth surface reconstruction. In: *Proc. SIGGRAPH* (1994)
7. Kaiser, A., Ybanez Zepeda, J.A., Boubekeur, T.: A survey of simple geometric primitives detection methods for captured 3D data. *Computer Graphics Forum* **38**(1), 167–196 (2019)
8. Koch, S., Matveev, A., Jiang, Z., Williams, F., Artemov, A., Burnaev, E., Alexa, M., Zorin, D., Panozzo, D.: ABC: A big cad model dataset for geometric deep learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019)
9. Kong, S., Fowlkes, C.: Recurrent pixel embedding for instance grouping. In: *2018 Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)
10. Krishnamurthy, V., Levoy, M.: Fitting smooth surfaces to dense polygon meshes. In: *SIGGRAPH* (1996)
11. Li, L., Sung, M., Dubrovina, A., Yi, L., Guibas, L.J.: Supervised fitting of geometric primitives to 3d point clouds. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019)
12. Paschalidou, D., Ulusoy, A.O., Geiger, A.: Superquadrics revisited: Learning 3d shape parsing beyond cuboids. In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 10336–10345 (2019)
13. Paschalidou, D., Gool, L.V., Geiger, A.: Learning unsupervised hierarchical part decomposition of 3d objects from a single rgb image. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2020)
14. Piegl, L., Tiller, W.: *The NURBS Book*. Springer-Verlag, Berlin, Heidelberg, 2nd edn. (1997)
15. Schnabel, R., Wahl, R., Klein, R.: Efficient RANSAC for point-cloud shape detection. *Computer Graphics Forum* **26**, 214–226 (06 2007)
16. Schneider, P.J., Eberly, D.: *Geometric Tools for Computer Graphics*. Elsevier Science Inc., USA (2002)
17. Sharma, G., Goyal, R., Liu, D., Kalogerakis, E., Maji, S.: Csgnet: Neural shape parser for constructive solid geometry. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 5515–5523 (2018)
18. Sharma, G., Goyal, R., Liu, D., Kalogerakis, E., Maji, S.: Neural shape parsers for constructive solid geometry. *Computing Research Repository (CoRR)* **abs/1912.11393** (2019), <http://arxiv.org/abs/1912.11393>

19. Silverman, B.W.: Density Estimation for Statistics and Data Analysis. Chapman & Hall, London (1986)
20. Smirnov, D., Fisher, M., Kim, V.G., Zhang, R., Solomon, J.: Deep parametric shape predictions using distance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
21. Sorkine, O., Alexa, M.: As-rigid-as-possible surface modeling. In: Symposium on Geometry Processing. pp. 109–116 (2007)
22. Sun, C., Zou, Q., Tong, X., Liu, Y.: Learning adaptive hierarchical cuboid abstractions of 3D shape collections. *ACM Transactions on Graphics (SIGGRAPH Asia)* **38**(6) (2019)
23. Tian, Y., Luo, A., Sun, X., Ellis, K., Freeman, W.T., Tenenbaum, J.B., Wu, J.: Learning to infer and execute 3d shape programs. In: International Conference on Learning Representations (2019)
24. Tulsiani, S., Su, H., Guibas, L.J., Efros, A.A., Malik, J.: Learning shape abstractions by assembling volumetric primitives. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017)
25. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics* **38**(5) (Oct 2019)
26. Yumer, M.E., Kara, L.B.: Surface creation on unstructured point sets using neural networks. *Computer-Aided Design* **44**(7), 644 – 656 (2012)
27. Zou, C., Yumer, E., Yang, J., Ceylan, D., Hoiem, D.: 3d-prnn: Generating shape primitives with recurrent neural networks. In: The IEEE International Conference on Computer Vision (ICCV) (2017)