

THE ABC DATASET IS A BIG DATA SET OF CAD MODELS FOR GEOMETRIC DEEP LEARNING

Nabnit Panigrahi¹, Jibananda Jena²

^{1, 2}gandhi Institute for Education & Technology, Bhubaneswar

Abstract

We present the ABC-Dataset, a database of one million Computer-Aided Design (CAD) models for the study of geometric deep learning techniques and applications. For the purposes of differential quantities, patch segmentation, geometric feature recognition, and form reconstruction, each model is a set of explicitly parametrized curves and surfaces. In order to generate data in various for-mats and resolutions, sampling the parametric descriptions of surfaces and curves makes it possible to generate fair comparisons for a variety of geometric learning methods. We carry out a large-scale benchmark for the estimation of surface normals using our dataset as a use case. We compare current data-driven methods and assess their effectiveness against both the ground truth and conventional normal estimate methods.

1. Introduction

The combination of large data collections and deep learning algorithms is transforming many areas of computer science. Large data collections are an essential part of this transformation. Creating these collections for many types of data (image, video, and audio) has been boosted by the ubiquity of acquisition devices and mass sharing of these types of data on social media. In all these cases, the data representation is based on regular discretization in space and time providing structured and uniform input for deep learning algorithms.

The situation is different for three-dimensional geometric models. Acquiring or creating high-quality models of this type is still difficult, despite growing availability of 3D sensors, and improvements in 3D design tools. Inherent ir-



Figure 1: Random CAD models from the ABC-Dataset

regularity of the surface data, and still-significant level of skill needed for creating high-quality 3D shapes contributes to the limited availability of geometric datasets. Irregularity of geometric data is reflected in commonly used geometric formats, which differ in fundamental ways from formats for images, video, and audio. Existing datasets often lack reliable ground truth annotations. We further discuss currently available geometric datasets in Section 2.

Common shape analysis and geometry processing tasks that can benefit from *geometric deep learning* include estimation of differential surface properties (Section 5.1), feature detection, and shape reconstruction. Ground truth for some of these tasks is hard to generate, as marking features by hand is a laborious task and differential properties can only be approximated for sampled surfaces.

In this work, we make the following contributions:

Dataset. We introduce a dataset for geometric deep learning consisting of over 1 million individual (and high quality) geometric models, each defined by parametric surfaces and associated with accurate ground truth information on the decomposition into patches, sharp feature annotations, and analytic differential properties. We gather the models through a publicly available interface hosted by Onshape [5]. Similar to vector graphics for images, this representation allows resampling the surface data at arbitrary resolutions, with or without connectivity information (i.e. into a point cloud or a mesh).

Benchmark. We demonstrate the use of the dataset by building a benchmark for the estimation of surface normals. This benchmark is targeting methods that compute normals

(1) locally on small regions of a surface, and (2) globally over the entire surface simultaneously. We have chosen this problem for two reasons: most existing geometric deep learning methods were tested on this task, and very precise ground truth can be readily obtained for shapes in our data set. We run the benchmark on 7 existing deep learning algorithms, studying how they scale as the dataset size increases, and comparing them with 5 traditional geometry processing algorithms to establish an objective baseline for future algorithms. Results are presented in Section 5.1.

Processing Pipeline. We develop an open-source geometry processing pipeline that processes the CAD models to directly feed deep learning algorithms. Details are provided in Section 3. We will continually update the dataset and benchmark as more models are added to the public collection of models by Onshape. Our contribution adds a new resource for the development of geometric deep learning, targeting applications focusing on human-created, mechanical shapes. It will allow researchers to compare against existing techniques on a large and realistic dataset of man-made objects.

Dataset	#Models	CAD Files	Curv atures	Patch Annotations	Semantic Categories	Categories
ABC	1,000,000+	X	X	X	–	–
ShapeNet* [22]	3,000,000+	–	–	–	X	X
ShapeNetCore	51,300	–	–	–	X	X
ShapeNetSem	12,000	–	–	–	X	X
ModelNet [58]	151,128	–	–	–	X	X
Thing10K [62]	10,000	–	–	–	–	X
PrincetonSB[9]	6670	–	–	–	–	X
NIST [4]	≤ 30	X	X	X	–	–

Table 1: Overview of existing datasets and their capabilities. *: The full ShapeNet dataset is not yet publicly available, only the subsets ShapeNetCore and ShapeNetSem.

2. Related Work

We review existing datasets for data-driven processing of geometrical data, and then review both data-driven and analytical approaches to estimate differential qualities on smooth surfaces.

3D Deep Learning Datasets. The community has seen a growth in the availability of 3D models and datasets. Segmentation and classification algorithms have benefited greatly from the most prominent ones [22, 58, 62, 9]. Further datasets are available for large scenes [51, 23], mesh registration [14] and 2D/3D alignment [12]. The dataset proposed in this paper has the unique property of containing the analytic representation of surfaces and curves, which is ideal for a quantitative evaluation of data-driven methods. Table 1 gives an overview of the most comparable datasets and their characteristics and capabilities.

Point Cloud Networks. Neural networks for point clouds are particularly popular, as they make minimal assumptions on input data. One of the earliest examples is PointNet [46] and its extension PointNet++ [47], which ensure that the network output is invariant with respect to point permutations. PCPNet [29] is a variation of PointNet tailored for estimating local shape properties: it extracts local patches from the point cloud, and estimates local shape properties at the central points of these patches. PointCNN [38] explores the idea of learning a transformation from the initial point cloud, which provides the weights for input points and associated features, and produces a permutation of the points into a latent order.

Point Convolutional Neural Networks by Extension Operators [11] is a fundamentally different way to process point clouds through mapping point cloud functions to volumetric functions and vice versa through extension and re-

striction operators. A similar volumetric approach has been proposed in Pointwise Convolutional Neural Networks [32] for learning pointwise features from point clouds.

PointNet-based techniques, however, do not attempt to use local neighborhood information explicitly. Dynamic Graph CNNs [55] uses an operation called EdgeConv, which exploits local geometric structures of the point set by generating a local neighborhood graph based on proximity in the feature space and applying convolution-like operation on the edges of this graph. In a similar fashion, FeaStNet [54] proposes a convolution operator for general graphs, which applies filter kernels in a data-driven manner to the local irregular neighborhoods.

Networks on Graphs and Manifolds. Neural networks for graphs have been introduced in [48], and extended in [39, 52]. A wide class of convolutional neural networks with spectral filters on graphs was introduced in [20] and developed further in [30, 24, 36]. Graph convolutional neural networks were applied to non-rigid shape analysis in [15, 60]. Surface Networks [37] further proposes the use of the differential geometry operators to extend GNNs to exploit properties of the surface. For a number of problems where quantities of interest are localized, spatial filters are more suitable than spectral filters. A special CNN model for meshes which uses intrinsic patch representations was presented in [41], and further generalized in [16, 42]. In contrast to these intrinsic models, Euclidean models [59, 57] need to learn the underlying invariance of the surface embedding, hence they have higher sample complexity. More recently, [40] presented a surface CNN based on the canonical representation of planar flat-torus. We refer to [19] for an extensive overview of geometric deep learning methods.

Analytic Normal Estimation Approaches. The simplest methods are based on fitting tangent planes. For a point set, these methods estimate normals in the points as directions of smallest co-variance (i.e., the tangent plane is the total least squares fit to the points in the neighborhood). For a triangle mesh, normals of the triangles adjacent to a vertex can be used to compute a vertex normal as a weighted average. We consider uniform, area, and angle weighting [34].

One can also fit higher order surfaces to the discrete surface data. These methods first estimate a tangent plane, and then fit a polynomial over the tangent plane that interpolates the point for which we want to estimate the normal (a so-called osculating jet [21]). A more accurate normal can then be recomputed from the tangents of the polynomial surface approximation at the point. The difference for triangle meshes and point sets is only in collecting the samples in the neighborhood. In addition, one can use robust statistics to weight the points in the neighborhood [35].

For the weighted triangle normals we use the implementation of libigl [33], for computation of co-variance normals

and osculating jets we use the functions in CGAL [45, 10]. For the robust estimation the authors have provided source code. There are many other techniques for estimating surface normals (e.g. [43]), however we only focus on a selected subset in the current work.

One important generalization is the detection of sharp edges and corners, where more than one normal can be assigned to a point [44, 17]. However, as we only train the machine learning methods to report a single normal per point, we leave an analysis of these extensions for future work.

3. Dataset

We identify six crucial properties that are desirable for an "ideal" dataset for geometric deep learning: (1) large size: since deep networks require large amounts of data, we want to have enough models to find statistically significant patterns; (2) ground truth signals: in this way, we can quantitatively evaluate the performance of learning on different tasks; (3) parametric representation: so that we can resample the signal at the desired resolution without introducing errors; (4) expandable: it should be easy to make the collection grow over time, to keep the dataset challenging as progress in learning algorithms is made; (5) variation: containing a good sampling of diverse shapes in different categories; (6) balanced: each type of objects should have a sufficient number of samples.

Since existing datasets are composed of acquired or synthesized point clouds or meshes, they do not satisfy property 2 or 3. We thus propose a new dataset of CAD models, which is complementary: it satisfies properties 1-4; it is restricted to a specific type of models (property 6), but has a considerable variation inside this class. While restriction to CAD models can be viewed as a downside, it strikes a good balance between having a sufficient number of similar samples and diversity of represented shapes, in addition to having high-quality ground truth for a number of quantities.

Acquisition. Onshape has a massive online collection of CAD models, which are freely usable for research purposes. By collecting them over a period of 4 months we obtained a collection of over 1 million models (see Figure 2).

Ground Truth Signals and Vector Representation. The data is encoded in a vectorial format that enables to resample it at arbitrary resolution and to compute analytically a large set of signals of interest (Section 4), which can be used as a ground truth.

Challenge. However, the data representation is not suitable for most learning methods, and the conversion of CAD models to discrete formats is a difficult task. This paper presents a robust pipeline to process, and use CAD models as an ideal data source for geometry processing algorithms.




Figure 2: Random examples from the dataset. Most models are mechanical parts with sharp edges and well defined surfaces.

3.1. CAD Models and Boundary Representations

In the following, we will use the term CAD model to refer to a geometric shape defined by a boundary representation (B-Rep). The boundary representation describes models by their topology (faces, edges, and vertices) as well as their geometry (surfaces, curves, and points). The topology description is organized in a hierarchical way: *Solids* are bound by a set of oriented faces which are called shells; *Faces* are bound by wires which are ordered lists of edges (the order defines the face orientation); *Edges* are the oriented connections between 2 vertices; *Vertices* are the basic entities, corresponding to points in space.

Each of these entities has a geometric description, which allows us to embed the model in 3D space. In our dataset, each surface can represent a plane, cone, cylinder, sphere, torus, surface of revolution or extrusion or NURBS patch [26]. Similarly, curves can be lines, circles, ellipses, parabolas, hyperbolas or NURBS curves [26]. Each vertex has coordinates in 3D space. An additional complexity is added by trimmed patches (see [26] for a full description).

3.2. Processing Pipeline

Assembling a dataset of such a large size is a time-consuming task where many design decisions have to be made beforehand: as an example, it requires around 2 CPU years to extract triangle meshes for 1 million CAD models.

To encourage active community participation and easy adoption, we use accessible, well-supported open-source tools, instead of relying on commercial CAD software. It allows the community to use and expand our dataset in the future. Our pipeline is designed to run in parallel on large computing clusters.

The Onshape public collection is not curated. It contains all the public models created by their users, without any additional filtering. Despite the fact that all models have been manually created, there is a small percentage of imperfect models with broken boundaries, self-intersecting faces or edges, as well as duplicate vertices. In addition to that, there are many duplicate models, and especially models that are just made of single primitives such as a plane,

box or cylinder, probably created by novice users that were learning how to use Onshape.

Given the massive size of the dataset, we developed a set of geometric and topological criteria to filter low quality of defective models, which we describe in the supplementary material, and we leave a semantic, crowd-sourced filtering and annotation as a direction for future work.

Step 1: B-Rep Loading and Translation. The STEP files [8] we obtain from Onshape contain the boundary representation of the CAD model, which we load and query using the open-source software Open Cascade [6]. The translation process generates for each surface patch and curve segment an explicit parameterization, which can be sampled at arbitrary resolution.

Step 2: Meshing/Discretization. The parameterizations of the patches are then sampled and triangulated using the open-source software Gmsh [28]. We offer an option here to select between uniform (respecting a maximal edge length) and curvature adaptive sampling.

Step 3: Topology Tree Traversal/Data Extraction. The sampling of B-Rep allows to track correspondences between the continuous and discrete representations. As they can be differentiated at arbitrary locations, it is possible to calculate ground truth differential quantities for all samples of the discrete model. Another advantage is the explicit topology representation in B-Rep models, which can be transferred to the mesh in the form of labels. These labels define for each triangle of the discrete mesh to which surface patch it belongs. The same applies also for the curves, we can label edges in the discrete mesh as sharp feature edges. While CAD kernels provide this information, it is difficult to extract it in a format suitable for learning tasks. Our pipeline exports this information in yaml files with a simple structure (see supplementary material).

Step 4: Post-processing. We provide tools to filter our meshes depending on quality or number of patches, to compute mesh statistics, and to resample the generated surfaces to match a desired number of vertices [27].




Figure 3: Each model in our dataset is composed of multiple patches and feature curves. The two images show the distribution of types of patches (left) and curves (right) over the current dataset ($\approx 1\text{M}$ models).




Figure 4: Histograms over the number of patches and curves per CAD model. This shows that there are many simpler models which consist of less than 30 patches/100 curves as well as more complex ones. Both histograms are truncated at the right side.

3.3. Analysis and Statistics

We show an overview of the models in the dataset in Figures 1 and 2. In Figure 3 and 4 we show the distribution of surface and edge types and the histogram of patch and edge numbers, to give an impression of the complexity and variety of the dataset. Updated statistics about the growing dataset are available on our dataset website [1].

4. Supported Applications

We briefly overview a set of applications that may benefit from our dataset, that can be used as either training data or as a benchmark.

Patch Decomposition. Each object in our collection is naturally divided into surface regions, separated by feature lines. The decomposition is defined by the author when a shape is constructed, and is likely to be semantically meaningful. It is also constrained by a strong geometric criteria: each region should be representable by a (possibly trimmed) NURBS patch.

Surface Vectorization. The B-rep of a CAD models is the counterpart of a vector representation for images, that can be resampled at any desired resolution. The conversion of a surface triangle mesh into a B-rep is an interesting and challenging research direction, for which data driven methods are still at their infancy [49, 50, 25].

Estimation of Differential Quantities. Our models have ground truth normals and curvature values making them an ideal, objective benchmark for evaluating algorithms to predict these quantities on point clouds or triangle meshes of artificial origin.

Sharp Feature Detection. Sharp features are explicitly encoded in the topological description of our models, and it is thus possible to obtain ground truth data for predicting sharp features on point clouds [56] and meshes.

Shape Reconstruction. Since the ground truth geometry is known for B-rep models, they can be used to simulate a scanning setup and quantitatively evaluate the reconstruction errors of both reconstruction [13, 53] and point cloud upsampling [61] techniques.

Image Based Learning Tasks. Together with the dataset, we provide a rendering module (based on Blender [2]) to generate image datasets. It supports rendering of models posed in physically static orientations on a flat plane, different lighting situations, materials, camera placements (half-dome, random) as well as different rendering modes (depth, color, contour). Note that all these images can be considered ground truth, since there is no geometric approximation error.

Robustness of Geometry Processing Algorithms. Besides data-driven tasks, the dataset can also be employed for evaluating the robustness of geometry processing algorithms. Even a simple task like normal estimation is problematic on such a large dataset. Most of the methods we evaluated in Section 5.1 fail (i.e. produce invalid normals with length zero or which contain NaNs) on at least one model: our dataset is ideal for studying and solving these challenging robustness issues.

5. Normal Estimation Benchmarks

We now introduce a set of large scale benchmarks to evaluate algorithms to compute surface normals, exploiting the availability of ground truth normals on the B-rep models. To the best of our knowledge, this is the first large scale study of this kind, and the insights that it will provide will be useful for the development of both data-driven and analytic methods.




Figure 5: Samples from the different categories in our normal estimation benchmark. From left to right: local patches of growing size and complexity (512, 1024, 2048 vertices), and full models at different densities (512, 1024, 2048 vertices).

Construction. To fairly compare a large variety of both data-driven and analytic methods, some targeting local estimation and some one-shot prediction of all normals of a 3D model, we build a series of datasets by randomly sampling points on our meshed B-reps, and growing patches of different sizes, ranging from 512 vertices to the entire model (Figure 5). For each patch size, we generate 4 benchmarks with an increasing number of patches, from 10k to 250k, to study how the data-driven algorithms behave when they are trained on a larger input set.

Split. All benchmark datasets are randomly split into training and test set with a distribution of 80% training data and 20% test data. The split will be provided to make the results reproducible.

5.1. Evaluation

Algorithms. We select 12 representative algorithms from the literature, and 5 of them are traditional ones, Robust Statistical Estimation on Discrete Surfaces (**RoSt**) [35] operating on point clouds (**PC**), and meshes (**M**), Osculating Jets (**Jets**) [21], also on point clouds and meshes, and Uniform weighting of adjacent face normals (**Uniform**) [34]. Notice that we replace K-ring neighborhoods with K-nearest neighbors for RoSt and Jets to support point cloud input. Also, 7 machine learning methods are selected, including PointNet++ (**PN++**) [47], Dynamic Graph CNN (**DGCNN**) [55], Pointwise CNN (**PwCNN**) [32], PointCNN (**PCNN**) [38], Laplacian Surface Network (**Laplace**) [37], PCP-Net (**PCPN**) [29] and Point Convolutional Neural Networks by Extension Operators (**ExtOp**) [11]. Of these methods, Laplace operates on triangle mesh input and the rest on point cloud input. Most of their output is one normal per vertex, except for PCPN and ExtOp the output is one normal for the center of the patch. We provide a detailed explanation of the (hyper-)parameters and modifications we did for each method in the supplementary material. For the statistics, we used only the valid normals reported by each method, and we filtered out all the degenerate ones.

Protocol. We compare the methods above on the benchmarks, using the following protocol: (1) for each method, we obtained the original implementation from the authors

(or a reimplementation in popular libraries); (2) we used the recommended or default values for all (hyper-)parameters of the learning approaches (if these were not provided, we fine-tuned them on a case by case basis); (3) if the implementation does not directly support normal estimation, we modified their code following the corresponding description in the original paper (4) we used the same loss function $1 - \mathbf{N}^T \mathbf{N}^*$, with \mathbf{N} as the estimated normal and \mathbf{N}^* as the ground truth normal, for all methods. Note that this loss function does not penalize normals that are inverted (flipped by 180°), which is an orthogonal problem usually fixed as in a postprocessing step [31].

Results. A listing of the statistical results for all methods is given in Table 2. Our experiments show that neural networks for normal estimation are stable across several runs; standard deviation of the losses is of the order of 10^{-3} . For the 10k dataset, most of the networks converge within 24 hours on a NVIDIA GeForce GTX 1080 Ti GPU. We capped the training time of all methods to 72 hours.

Comparison of Data-Driven Methods. We observe, as expected, that the error is reduced as we increase the number of samples in the training set, this is consistent for all methods on both patches and full models. However, the improvement is modest (Figures 6 and 7).

Sampling Resolution on Full Models. We also explore how data-driven methods behave when sampling resolution is growing. DGCNN, PCNN, and PwCNN clearly benefit from sampling resolution, while PN++ does not show clear improvements. This phenomenon is likely linked to the spatial subsampling mechanism that is used to ensure sublinear time in training, but prevents this method from leveraging the extra resolution. In case of Laplace surface network, it is difficult to understand the effect since it did not converge after 3 days of training on the highest resolution.

Comparison of Analytic Methods. Analytic methods are remarkably consistent across dataset sizes and improve as the mesh sampling resolution is increased. The methods based on surface connectivity heavily outperform those relying on K-nearest neighbour estimation, demonstrating that connectivity is a valuable information for this task.

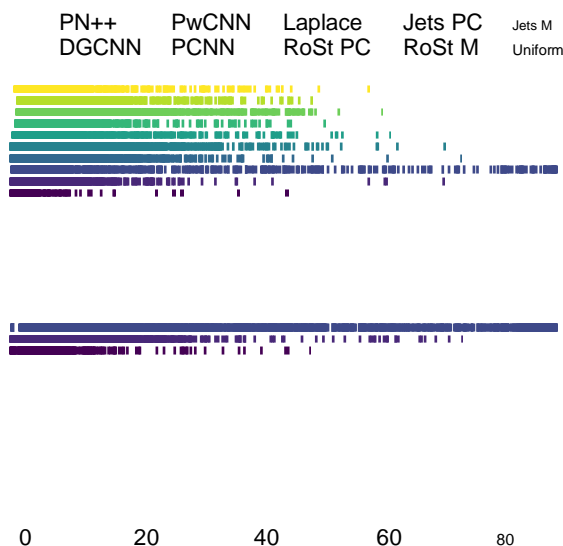


Figure 6: Plot of angle deviation error for the lower resolution patches (512 points) benchmark, using different sample size (top to bottom: 10k, 50k, and 100k).

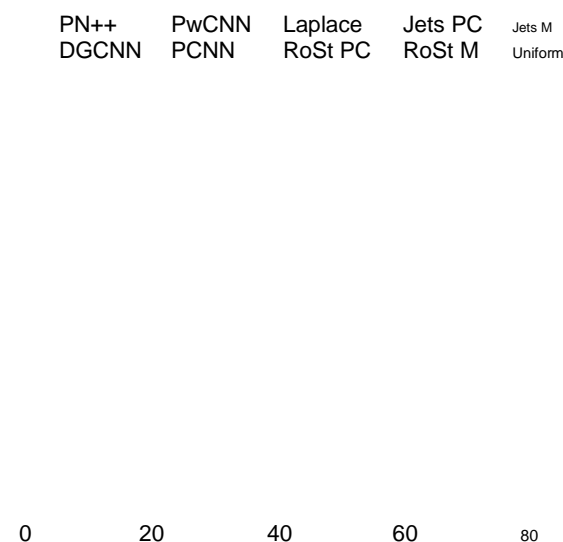


Figure 7: Plot of angle deviation error for the high-resolution (2048 points) full model benchmark, using different sample size (top to bottom: 10k, 50k, and 100k).

Data-Driven vs. Analytic Methods. Almost all data-driven methods perform well against analytic methods for point clouds, especially if the model resolution is low. However, if the analytic methods are allowed to use connectivity information, they outperform all learning methods by a large margin, even those also using connectivity information. To further support this conclusion, we run a similar experiment on a simpler, synthetic dataset composed of 10k and 50k random NURBS patches and observe similar results, which are available in the supplementary material.

6. Conclusion

We introduced a new large dataset of CAD models, and a set of tools to convert them into the representations used by deep learning algorithms. The dataset will continue to grow as more models are added to the Onshape public collection. Large scale learning benchmarks can be created using the ground truth signals that we extract from the CAD data, as we demonstrate for the estimation of differential surface quantities.

The result of our comparison will be of guidance to the development of new geometric deep learning methods. Our surprising conclusion is that, while deep learning methods which use only 3D points are superior to analytical methods, this is not the case when connectivity information is available. This suggests that existing graph architectures struggle at exploiting the connectivity efficiently and are considerably worse than the simplest analytical method (uniform), which simply averages the normals of neighbouring faces. It would be interesting to run a similar study by extending these algorithms to correctly identify and predict multiple normals on sharp features and compare them with specialized methods for this task [18].

Another surprising discovery is that even the uniform algorithm fails to produce valid normals on roughly 100 models in our dataset due to floating point errors. These kinds of problems are extremely challenging to identify, and we believe that the size and complexity of our dataset are an ideal stress test for robust geometry processing algorithms.

7. Distribution

The dataset and all information is available at: <https://deep-geometry.github.io/abc-dataset>

It is distributed under the MIT license and split into chunks of 10k models for each data type. The copyright owners of the models are the respective creators (listed in the meta information). The geometry processing pipeline is made available under the GPL license in form of a containerized software solution (Docker [3] and Singularity [7]) that can be run on every suitable machine.

Acknowledgements

We are grateful to Onshape for providing the CAD models and support. This work was supported in part through the NYU IT High Performance Computing resources, services, and staff expertise. Funding provided by NSF award MRI-1229185. We thank the Skoltech CDISE HPC Zhores cluster staff for computing cluster provision. This work was supported in part by NSF CAREER award 1652515, the NSF grants IIS-1320635, DMS-1436591, and 1835712, the Russian Science Foundation under Grant 19-41-04109, and gifts from Adobe Research, nTopology Inc, and NVIDIA.

Method/ Vertices		Full Models						Patches					
		Loss			Angle Deviation [°]			Loss			Angle Deviation [°]		
		10k	50k	100k	10k	50k	100k	10k	50k	100k	10k	50k	100k
PN++	512	0.168	0.155	0.142	7.83	7.32	6.43	0.034	0.032	0.025	1.70	1.42	1.10
	1024	0.180	0.163	0.160	7.49	6.06	6.17	0.056	0.052	0.048	2.11	1.57	1.51
	2048	0.171	0.156	0.149	6.75	5.47	5.08	0.081	0.071	0.063	2.48	1.77	1.44
DGCN	512	0.177	0.167	0.144	9.61	8.32	7.13	0.054	0.049	0.025	3.20	3.00	1.12
	1024	0.126	0.104	0.099	5.91	4.59	4.34	0.048	0.036	0.024	2.98	2.15	1.13
	2048	0.090	0.070	0.068	4.54	2.80	2.77	0.045	0.035	0.023	2.66	1.95	0.98
PwCN	512	0.273	0.260	0.252	18.73	17.27	16.36	0.092	0.069	0.067	4.71	3.45	3.43
	1024	0.217	0.218	0.198	12.78	13.10	11.38	0.107	0.110	0.089	5.50	6.11	4.58
	2048	0.188	0.176*	0.168*	11.34	10.54*	10.05*	0.107	0.120*	0.094	5.98	6.36*	4.83
PCNN	512	0.146	0.153	0.139	6.47	6.96	6.15	0.037	0.043	0.028	1.84	1.84	1.42
	1024	0.104	0.099	0.103	3.56	3.46	3.69	0.025	0.030	0.025	0.94	1.37	0.92
	2048	0.065	0.070	0.067	2.05	2.44	2.22	0.023	0.025	0.023*	0.88	1.01	0.83*
Laplac	512	0.282	0.203	0.133	20.01	11.94	8.47	0.041	0.047	0.022	1.93	3.13	1.12
	1024	0.211	0.138	0.146*	34.24	9.43	9.85*	0.030	0.027	0.029*	1.65	1.36	1.46*
	2048	0.197	0.148*	0.158*	9.99	9.95*	10.57*	0.031	0.040	0.040*	1.60	1.67	1.81*
PCPNNet	512	–	–	–	–	–	–	0.098†	0.081†	–	9.95†	9.28†	–
	1024	–	–	–	–	–	–	0.123†	0.097†	–	13.89†	9.55†	–
	2048	–	–	–	–	–	–	0.142†	0.200†	–	16.24†	16.45†	–
ExtOp	512	–	–	–	–	–	–	0.074†	0.073†	–	2.42†	2.05†	–
	1024	–	–	–	–	–	–	0.095†	0.096†	–	3.32†	2.50†	–
	2048	–	–	–	–	–	–	0.091†	–	–	3.00†	–	–
PC	512	0.298	0.300	–	21.32	21.36	–	0.083	0.082	–	0.82	0.79	–
	1024	0.220	0.223	–	14.47	14.63	–	0.078	0.077	–	0.74	0.72	–
	2048	0.164	0.166	–	9.96	10.18	–	0.073	0.072	–	0.59	0.62	–
RoSt	512	0.260	0.261	–	17.84	17.97	–	0.050	0.050	–	0.05	0.05	–
	1024	0.183	0.186	–	12.19	12.39	–	0.048	0.048	–	0.05	0.05	–
	2048	0.129	0.132	–	8.41	8.63	–	0.045	0.044	–	0.04	0.04	–
Jets	512	0.082	0.084	0.084	2.15	2.17	2.18	0.108	0.103	0.102	0.06	0.06	0.06
	1024	0.053	0.055	0.056	0.25	0.29	0.29	0.107	0.105	0.105	0.06	0.06	0.06
	2048	0.047	0.048	0.050	0.08	0.08	0.08	0.112	0.108	0.107	0.06	0.06	0.06
M	512	0.175	0.176	0.175	7.26	7.29	7.29	0.036	0.036	0.036	0.00	0.00	0.00
	1024	0.118	0.118	0.117	0.10	0.10	0.11	0.033	0.033	0.033	0.00	0.00	0.00
	2048	0.078	0.079	0.079	0.01	0.02	0.02	0.029	0.031	0.031	0.00	0.00	0.00
Jets _{in}	512	0.024	0.025	0.024	0.26	0.29	0.28	0.007	0.007	0.007	0.00	0.00	0.00
	1024	0.013	0.013	0.013	0.00	0.00	0.00	0.005	0.005	0.005	0.00	0.00	0.00
	2048	0.009	0.010	0.009	0.00	0.00	0.00	0.004	0.005	0.004	0.00	0.00	0.00

Table 2: Statistical results for all evaluated methods for the full model and patch benchmarks. The loss is calculated as $1 - \mathbf{N}^T \mathbf{N}^*$ and the angle deviation is calculated as the angle in degrees (\mathbf{N} , \mathbf{N}^*) between ground truth normal and estimated normal. For the loss we report the mean over all models, for the angle deviation we report the median of all models in the according datasets. Osculating Jets and Robust Statistical Estimation are evaluated both on point cloud inputs (PC suffix; comparable to the learning methods) as well as mesh inputs (M suffix). †: PCPNNet and ExtOp were not run on full models since they compute only 1 normal per patch (and their loss, differently from all other rows, is computed only on the vertex in the center of the patch). *: the training was not completed before the time limit is reached, and the partial result is used for inference.

References

- [1] ABC-Dataset. <https://deep-geometry.github.io/abc-dataset>. Accessed: 2019-03-20.
- [2] Blender. <https://www.blender.org/>. Accessed: 2018-11-14.
- [3] Docker. <https://www.docker.com/>. Accessed: 2018-11-11.
- [4] NIST CAD Models and STEP Files with PMI. <https://catalog.data.gov/dataset/nist-cad-models-and-step-files-with-pmi>. Accessed: 2018-11-11.
- [5] Onshape. <https://www.onshape.com/>. Accessed: 2018-11-14.
- [6] Open CASCADE Technology OCCT. <https://www.opencascade.com/>. Accessed: 2018-11-11.
- [7] Singularity. <https://singularity.lbl.gov/>. Accessed: 2018-11-11.
- [8] STEP File Format. https://en.wikipedia.org/wiki/ISO_10303-21. Accessed: 2018-11-11.
- [9] The princeton shape benchmark. In *Proceedings of the Shape Modeling International 2004*, SMI '04, pages 167–178, Washington, DC, USA, 2004. IEEE Computer Society.
- [10] P. Alliez, S. Giraudot, C. Jamin, F. Lafarge, Q. Merigot, J. Meyron, L. Saboret, N. Salman, and S. Wu. Point set processing. In *CGAL User and Reference Manual*. CGAL Editorial Board, 4.13 edition, 2018.
- [11] M. Atzmon, H. Maron, and Y. Lipman. Point convolutional neural networks by extension operators. *ACM Trans. Graph.*, 37(4):71:1–71:12, July 2018.
- [12] M. Aubry, D. Maturana, A. Efros, B. Russell, and J. Sivic. Seeing 3d chairs: exemplar part-based 2d-3d alignment using a large dataset of cad models. In *CVPR*, 2014.
- [13] M. Berger, J. A. Levine, L. G. Nonato, G. Taubin, and C. T. Silva. A benchmark for surface reconstruction. *ACM Transactions on Graphics (TOG)*, 32(2):20, 2013.
- [14] F. Bogo, J. Romero, G. Pons-Moll, and M. J. Black. Dynamic FAUST: Registering human bodies in motion. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [15] D. Boscaini, J. Masci, S. Melzi, M. M. Bronstein, U. Castellani, and P. Vandergheynst. Learning class-specific descriptors for deformable shapes using localized spectral convolutional networks. In *Computer Graphics Forum*, volume 34, pages 13–23. Wiley On-line Library, 2015.
- [16] D. Boscaini, J. Masci, E. Rodola, and M. Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 3189–3197, 2016.
- [17] A. Boulch and R. Marlet. Fast and robust normal estimation for point clouds with sharp features. *Comput. Graph. Forum*, 31(5):1765–1774, Aug. 2012.
- [18] A. Boulch and R. Marlet. Deep learning for robust normal estimation in unstructured point clouds. In *Proceedings of the Symposium on Geometry Processing*, SGP '16, pages 281–290, Goslar Germany, Germany, 2016. Eurographics Association.
- [19] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- [20] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.
- [21] F. Cazals and M. Pouget. Estimating differential quantities using polynomial fitting of osculating jets. *Computer Aided Geometric Design*, 22(2):121 – 146, 2005.
- [22] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015.
- [23] 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.
- [24] M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pages 3844–3852, 2016.
- [25] T. Du, J. P. Inala, Y. Pu, A. Spielberg, A. Schulz, D. Rus, A. Solar-Lezama, and W. Matusik. Inverse-csg: Automatic conversion of 3d models to csg trees. *ACM Trans. Graph.*, 37(6):213, 2018.
- [26] G. Farin. *Curves and Surfaces for CAGD: A Practical Guide*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 5th edition, 2002.
- [27] M. Garland and P. S. Heckbert. Surface simplification using quadric error metrics. In *Proceedings of the 24th*

Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '97, pages 209–216, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.

- [28] C. Geuzaine and J. F. Remacle. Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 2009.
- [29] P. Guerrero, Y. Kleiman, M. Ovsjanikov, and N. J. Mitra. PCPNet: Learning local shape properties from raw point clouds. *Computer Graphics Forum*, 37(2):75–85, 2018.
- [30] M. Henaff, J. Bruna, and Y. LeCun. Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163*, 2015.
- [31] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. *Surface reconstruction from unorganized points*, volume 26. ACM, 1992.
- [32] B.-S. Hua, M.-K. Tran, and S.-K. Yeung. Point-wise convolutional neural networks, 2017. cite arxiv:1712.05245Comment: 10 pages, 6 figures, 10 tables. Paper accepted to CVPR 2018.
- [33] A. Jacobson, D. Panozzo, et al. libigl: A simple C++ geometry processing library, 2018. <http://libigl.github.io/libigl/>.
- [34] S. Jin, R. R. Lewis, and D. West. A comparison of algorithms for vertex normal computation. *Vis. Comput.*, 21(1-2):71–82, Feb. 2005.
- [35] E. Kalogerakis, P. Simari, D. Nowrouzezahrai, and K. Singh. Robust statistical estimation of curvature on discretized surfaces. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing, SGP '07*, pages 13–22, Aire-la-Ville, Switzerland, Switzerland, 2007. Eurographics Association.
- [36] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [37] I. Kostrikov, Z. Jiang, D. Panozzo, D. Zorin, and B. Joan. Surface networks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2018.
- [38] Y. Li, R. Bu, M. Sun, and B. Chen. Pointcnn. *CoRR*, abs/1801.07791, 2018.
- [39] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.
- [40] H. Maron, M. Galun, N. Aigerman, M. Trope, N. Dym, E. Yumer, V. G. Kim, and Y. Lipman. Convolutional neural networks on surfaces via seamless toric covers. *ACM Trans. Graph*, 36(4):71, 2017.
- [41] J. Masci, D. Boscaini, M. Bronstein, and P. Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 37–45, 2015.
- [42] F. Monti, D. Boscaini, J. Masci, E. Rodola, J. Svoboda, and M. M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proc. CVPR*, volume 1, page 3, 2017.
- [43] C. Mura, G. Wyss, and R. Pajarola. Robust normal estimation in unstructured 3d point clouds by selective normal space exploration. *Vis. Comput.*, 34(6-8):961–971, June 2018.
- [44] Y. Ohtake, A. Belyaev, M. Alexa, G. Turk, and H.-P. Seidel. Multi-level partition of unity implicits. *ACM Trans. Graph.*, 22(3):463–470, July 2003.
- [45] M. Pouget and F. Cazals. Estimation of local differential properties of point-sampled surfaces. In *CGAL User and Reference Manual*. CGAL Editorial Board, 4.13 edition, 2018.
- [46] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Point-net: Deep learning on point sets for 3d classification and segmentation. *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 1(2):4, 2017.
- [47] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5099–5108. Curran Associates, Inc., 2017.
- [48] 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.
- [49] V. Shapiro and D. L. Vossler. Construction and optimization of csg representations. *Comput. Aided Des.*, 23(1):4–20, Feb. 1991.
- [50] G. Sharma, R. Goyal, D. Liu, E. Kalogerakis, and S. Maji. Csgnet: Neural shape parser for constructive solid geometry. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [51] S. Song, S. P. Lichtenberg, and J. Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 567–576, 2015.
- [52] S. Sukhbaatar, R. Fergus, et al. Learning multiagent communication with backpropagation. In *Advances in Neural Information Processing Systems*, pages 2244–2252, 2016.

- [53] X. Sun, P. L. Rosin, R. R. Martin, and F. C. Langbein. Noise analysis and synthesis for 3d laser depth scanners. *Graphical Models*, 71(2):34 – 48, 2009. IEEE International Conference on Shape Modeling and Applications 2008.
- [54] N. Verma, E. Boyer, and J. J. Verbeek. Feastnet: Feature-steered graph convolutions for 3d shape analysis. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2598–2606, 2018.
- [55] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon. Dynamic graph cnn for learning on point clouds. *CoRR*, abs/1801.07829, 2018.
- [56] C. Weber, S. Hahmann, and H. Hagen. Sharp feature detection in point clouds. In *Proceedings of the 2010 Shape Modeling International Conference, SMI '10*, pages 175–186, Washington, DC, USA, 2010. IEEE Computer Society.
- [57] L. Wei, Q. Huang, D. Ceylan, E. Vouga, and H. Li. Dense human body correspondences using convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1544–1553, 2016.
- [58] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3d shapenets: A deep representation for volumetric shapes. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1912–1920, June 2015.
- [59] 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.
- [60] L. Yi, H. Su, X. Guo, and L. J. Guibas. Syncspeccnn: Synchronized spectral cnn for 3d shape segmentation. In *CVPR*, pages 6584–6592, 2017.
- [61] L. Yu, X. Li, C.-W. Fu, D. Cohen-Or, and P.-A. Heng. Pu-net: Point cloud upsampling network. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [62] Q. Zhou and A. Jacobson. Thingi10k: A dataset of 10,000 3d-printing models. *arXiv preprint arXiv:1605.04797*, 2016.