3D Understanding of Deformable Linear Objects: Datasets and Transferability Benchmark

Bare Luka Zagar[1*], Mingyu Liu[1*], Tim Hertel[1*], Ekim Yurtsever[2], Alois Knoll[1]

[1]Technical University of Munich, 85748 Garching b. München, Germany

[2] The Ohio State University, Columbus, OH 43212, USA

bare.luka.zagar@tum.de

[*]Authors contributed equally.

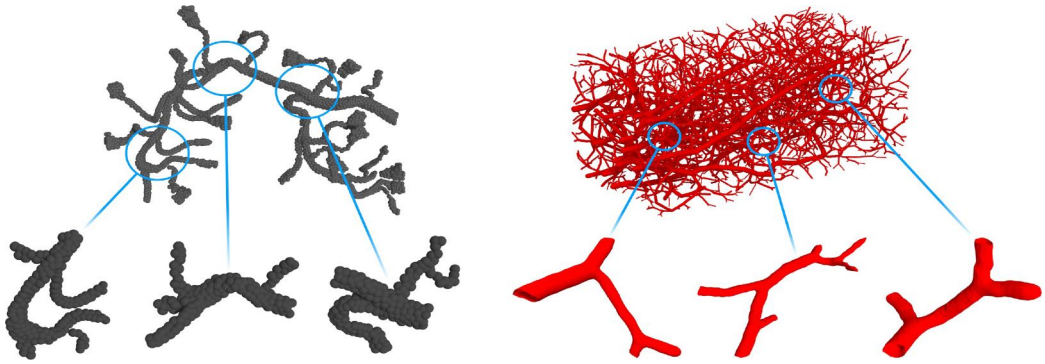

Figure 1. Deformable linear objects can manifest in many variations, shapes, and configurations. For example, automotive wiring harnesses (**left**) and blood vascular structures (**right**) are completely different types of objects, but when compared at a micro-level, it is clear that they share similar spatial topological properties.

Abstract

Deformable linear objects are commonly found in our daily lives. Understanding them visually can be challenging even for humans, as the same object can become entangled and look completely different. Examples of deformable linear objects include blood vessels and wiring harnesses, which are crucial for the proper functioning of systems like the human body and vehicles. Recently, some studies have focused on 2D image segmentation of wires. However, there are no point cloud datasets available for studying 3D deformable linear objects, which are more complex and challenging. To address this gap, we introduce two point cloud datasets, PointWire and PointVessel, generated using our proposed semi-automatic pipeline. We evaluated state-of-the-art methods on these large-scale 3D deformable linear object benchmarks. Additionally, we analyzed the generalization capabilities of these methods through transferability experiments on the PointWire and PointVessel datasets.

https://github.com/heti2000/cdlo-datasets.

1. Introduction

Deformable linear objects (DLOs), such as blood vessels and wiring harnesses, are crucial components of their corresponding higher-level systems. Blood vessels allow the vascular system to carry blood throughout the human body, while wiring harnesses enable the electrical and electronic architecture to transport current through the vehicle body. Despite their inherent differences, when observed through the prism of DLOs, these objects are based on the same underlying topological structure. Moreover, DLOs can come in very complex entangled configurations, as shown in Fig. 1. Therefore, understanding the topology and disentanglement of these structures is challenging.

The automation of wiring harness manufacturing remains significantly underdeveloped, with automation processes being nearly nonexistent [25, 27]. Research on the perception and analysis of wiring harnesses is notably lim-

ited. The work proposed in [28] concentrates on extracting the profiles of simple wiring harnesses, while [14] classifies different parts of complex automotive wiring harnesses. A sophisticated 2D instance segmentation approach of wires was introduced in [2]. Despite these advancements, a major challenge persists. Currently, there is no publicly available 3D dataset specifically designed for wiring harnesses with intricate 3D geometries. Furthermore, there is a lack of studies focusing on the extraction of topological elements, such as bifurcations and endpoints, which are crucial for achieving disentanglement in these complex structures.

On the other hand, deep learning methods have boosted the medical imaging field and are gaining more and more attention [37]. A significant amount of research in the medical imaging domain focuses on the segmentation of vascular systems, such as brain vasculature [15, 35] and retinal vessel segmentation [13]. Recently, the IntrA [47] dataset sparked research on point cloud segmentation of intracranial aneurysms [21,48]. However, there still does not exist a publicly available 3D point cloud dataset for studying blood vessels' spatial topology configurations.

To this end, we propose two deformable linear object datasets, PointWire and PointVessel. The PointWire dataset is based on 40 real-scanned complex automotive wiring harnesses, which we extended with our semi-automatic dataset generator to a total of 12000 samples. The high-precision raw point clouds, including a wiring harness, are first processed by removing the background and noise. Then, a skeleton of the wiring harness points is retrieved and manually refined. The refined skeleton is used for automating the process of rigging the wiring harness point cloud with Blender [8]. It enables us to extend the dataset size and introduce a wide variety of wiring harness spatial configurations. The PointVessel synthetic dataset is derived from [34], which includes 136 blood vessel volumes. Firstly, we convert the volumes to mesh data and separate each volume into 96 cropped-out smaller vascular structures. Then, the surface of these crops is uniformly subsampled to retrieve the final point cloud data.

There are multiple benefits of the introduced PointWire and PointVessel datasets: 1) more research regarding the automation of wiring harnesses will be encouraged, 2) the different data domains but still very similar data structure and topology will further boost the research in the medical imaging field of vascular systems by providing a different perspective on similar underlying issues, and 3) due to a general lack of data, the manufacturing and medical imaging fields can benefit from jointly using the given datasets in terms of transfer learning.

The key contributions of our work are the following:

- We are the first to introduce point cloud datasets - PointWire and PointVessel - to understand 3D DLOs by means of disentanglement and segmentation of topological structures.
- We propose a semi-automatic pipeline for processing, augmenting, and annotating real-scanned wiring harnesses and generating datasets.
- We provide large-scale benchmarks on segmentation approaches and conduct comprehensive transferability analysis to understand the typology of complex DLOs by combining segmentation and disentanglement.

2. Related Works

2.1. 3D Object Datasets

3D datasets enabled the tremendous advancement of deep learning in the point cloud domain. Many 3D object datasets, that employ synthetic CAD models, are widely used for common tasks, such as classification and segmentation. The ModelNet [42] and ShapeNet [5] datasets provide a large amount of object types. ABO [7] presented a more realistic and real-world relevant 3D object dataset. Additionally, the ABC [16] and DeepCAD [40] are introduced for mechanical objects. To address the lack of realistic propertiers inherent by CAD-based datasets, ScanObjectNN [36] and OmniObject3D [41] propose real-scanned 3D objects for the creation of their datasets. However, most of the aforementioned datasets contain only common objects. Recently, yang *et al.* [47] introduce the IntrA dataset to encourage the research on point cloud segmentation of aneurysms. However, it is relatively small and does not focus on the complexity of extracting the blood vessel topology. In contrast, our proposed PointWire and PointVessel datasets enable the research of segmenting the topological features of complex deformable linear objects, such as wiring harnesses and blood vessels.

2.2. Point Cloud Segmentation

In recent years, point cloud segmentation has been dramatically developed. PointNet [30] and PointNet++ [31] are prior works that extract point features directly from the input data, exhibiting promising performance. Other works [17, 19, 20, 44, 46] apply convolution neural network to learn point features. DGCNN [38] proposes the EdgeConv module to obtain edge features from the KNN-based local graphs. Furthermore, PointTransformer [50] and PCT [12] explore transformer-based networks, achieving comparable performance. GDANet [45] disentangles point clouds into the contour and flat for point cloud understanding. CurveNet [43] implements a curve grouping operator to aggregate point features. The most recent work RepSurf [32] exploits triangle-based and multi-surface representation for segmentation. DeltaConv [39] proposes a set of anisotropic convolution layers for point cloud surface representation.

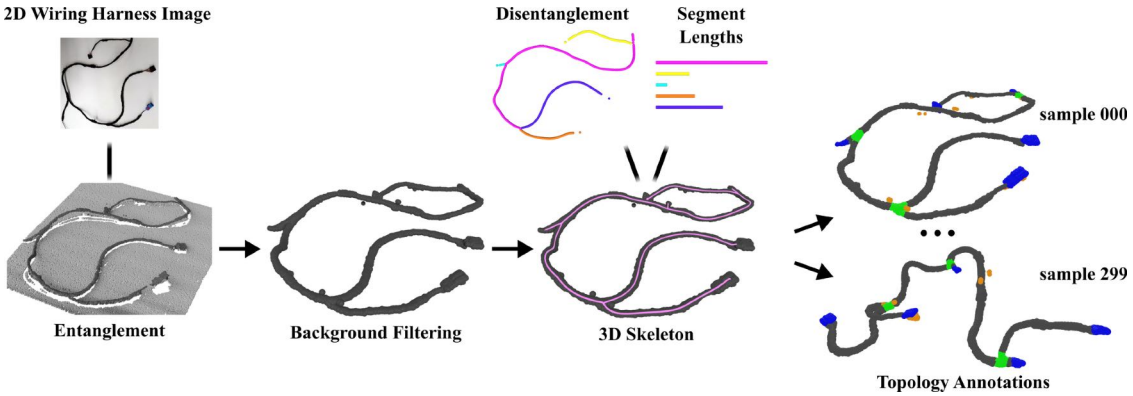

Figure 2. Wiring harness dataset generation pipeline: First, the raw input point cloud of an entangled wiring harness is taken with a Photoneo MotionCam-3D M. Then, we perform plane segmentation and statistical noise removal. Afterward, a skeleton is retrieved using the Laplace-based contraction method [1]. Additionally, the skeleton is transformed into a disentangled DLO graph representation. Further, the clean wire harness point cloud and the skeleton-based graph are combined and extended using the Blender toolkit. Finally, a diverse set of wire harness configurations is achieved.

Although the aforementioned methods have been extensively employed in the point cloud segmentation for common objects, their effectiveness and limitations have not been thoroughly validated on DLOs due to the lack of available datasets. To address this gap, we perform a comprehensive evaluation of several state-of-the-art methods using our large-scale benchmarks.

2.3. Deformable Linear Objects Perception

DLOs in the form of blood vessels are intensively studied in the medical imaging field [10, 24]. DeepVesselNet [34] introduces a synthetically generated blood vessel dataset and proposed a 3D U-Net architecture-like solution to retrieve the topological information of blood vessels. Recently, VesselGraph [29] proposes a large-scale graph dataset based on available blood vessel datasets. However, no available dataset focuses on the 3D spatial attributes that point cloud data could provide. Therefore, we propose to fill this gap with our PointVessel dataset, which enables the study of blood vessels in the rich 3D spatial space.

Cables, wires, ropes, and wiring harnesses represent a significant category of DLOs that remain underexplored within the field of perception research. Only recently a few researchers have focused on real-world 2D image segmentation of wires [3, 4, 6]. For instance, Lv *et al.* [22] propose a method for robustly estimating the 3D state of DLOs using point clouds. Additionally, Nguyen *et al.* [26] introduce a synthetic point cloud dataset for wiring harnesses. However, this dataset is not publicly available, limiting its potential for further research. Furthermore, the experimental analysis and the dataset's scale presented in [26] are relatively limited. To address these gaps, we propose the PointWire datasets, which comprise a large-scale, real-scanned point cloud dataset of wiring harnesses, aiming to advance the study and application of DLO understanding.

3. DLO Datasets

Since there are no available point cloud datasets for complex 3D deformable linear objects, such as wiring harnesses and blood vessels, we are the first to propose datasets of such kind. The PointWire and PointVessel datasets serve for topological DLOs key part segmentation and as a transferability benchmark to study the generalization capabilities of point cloud segmentation methods.

3.1. PointWire Dataset

The PointWire dataset generation pipeline is given in Fig. 2. This dataset is based on 40 real wiring harnesses used in the automotive sector. As shown in Fig. 4, we use the high-precision industrial-grade Photoneo MotionCam-3D M attached to a Kuka Iiwa LBR 14. To ensure a constant distance between the perception sensor and the wiring harness, we mount the 3D scanner on a robot arm and set the distance to the sensor's optimal distance based on [9].

First, a dense 3D scan of the wiring harness is obtained with the Photoneo MotionCam-3D M. An entangled wiring harness's captured point cloud scan contains more than 3 million sub-millimeter precise points. Secondly, the raw input point cloud is processed by removing the plane as well as the noise, using RANSAC and statistical outlier removal provided by Open3D [51]. The result is a clean point cloud of the wiring harness. To generate the skeleton from the cleaned point cloud, we applied the Laplacian-based contraction method [1], as illustrated in Fig. 2. After this automatic generation, we manually refined the skeleton by adjusting and relocating key points to ensure an accurate representation of bifurcation and endpoint nodes. Specifically,

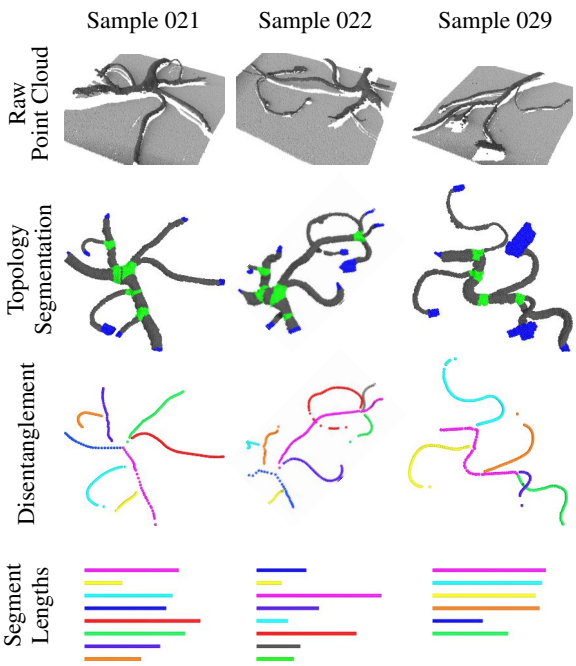

Figure 3. PointWire dataset description: The **first** row shows entangled wiring harness point clouds. The **second** row depicts the topology segmentation. The **third** row shows the disentangled wiring harness skeleton. The **fourth** row shows the disentangled segment lengths of the wiring harness.

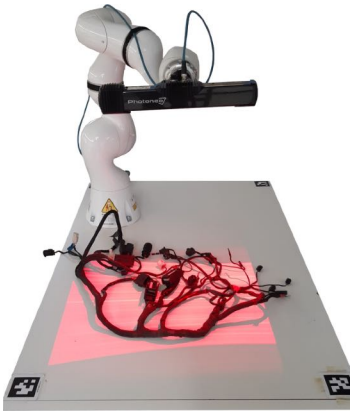

Figure 4. Experimental setup for the collection of 3D wiring harness scans. A Photoneo Motion-Cam 3D M is mounted on a Kuka Iiwa LBR 14 robot arm to ensure consistent measurement distances for all scans.

skeleton points were purged or relocated to eliminate ambiguities. For bifurcations, we identified nodes with a degree greater than two, while nodes with a degree of one were classified as endpoints.

The resulting skeleton graph was then imported into Blender [8] for rigging, where we assigned ground truth labels to the structural elements. We ensured annotation accuracy by assigning points based on the local thickness of

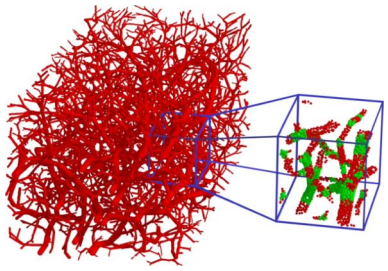

Figure 5. Blood vessel point cloud data generation. The entire synthetic blood vessel mesh is cropped into smaller subsets and uniformly sampled to 2048 points. Red points indicate the vessels, while the green points indicate the bifurcation.

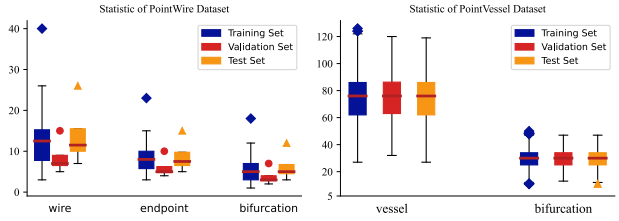

Figure 6. Statistics of the class distributions of the PointWire and PointVessel datasets. The PointWire dataset has a diverse class distribution, while PointVessel has a higher average of class instances per sample and maintains an equal distribution of classes.

neighboring segments. Furthermore, segment lengths were calculated using B-spline interpolation, providing more precision in the final annotated data. This manual adjustment process helped guarantee high-quality annotations by ensuring the correct labeling of nodes and accurate representation of the topological structure. Finally, by rigging and animating the cleaned point cloud, we extend each wiring harness point cloud scan to 300 samples. Each rigged real scan is manually moved in a physically feasible configuration. The in-between frames were interpolated by Blender as depicted in Fig. 2. Additional samples of entangled wiring harnesses are shown in Fig. 3. Furthermore, it demonstrates the challenging perception task setting that DLOs represent due to potentially complex entanglement.

3.2. PointVessel Dataset

The presented point cloud dataset of blood vessels, PointVessel, is derived from synthetic medical blood vessel data introduced in [34]. The original DeepVesselNet [34] dataset is available as NIfTI files, which is a common data type in the medical imaging field.

There are 136 different blood vessel volumes with the size of $325 \times 304 \times 600$ voxels. The dataset processing pipeline is shown in Fig. 5. The first step is to convert the NIfTI blood vessel volumes using the vtk toolkit [33] into a 3D mesh. As each blood vessel volume contains a large

Dataset	Class	Training set	Validation set	Test set	Total
PointWire	**w**	124.5K	10.2K	16.8K	151.5K
	e	78.9K	7.2K	10.5K	96.6K
	b	55.2K	4.5K	7.5K	67.2K
	samples	9.6k	1.2K	1.2K	12K
PointVessel	**v**	711.9K	129.2K	128.2K	969.3K
	b	283.3K	51.4K	51K	385.7K
	samples	9.6K	1.7K	1.7K	13K

Table 1. PointWire and PointVessel statistic. **w**-wire, **e**-endpoint, **b**-bifurcation, **v**-vessel, samples are the number of frames.

number of vessels and bifurcations, we divide each volume into 96 same-sized overlaying crops, similarly to [18]. The cropping strategy is depicted in Fig. 5 with the blue box inside the blood vessel volume. Afterward, the surface mesh of each cropped volume is uniformly subsampled to a point cloud of 2048 points, shown with the enlargened blue box. The ground truth bifurcation points and the vessel radii from [34] are used to assign the corresponding point annotations. More details and figures regarding the PointVessel dataset are given in the supplementary.

3.3. Statistics

We show the statistic of our proposed two datasets, PointWire and PointVessel, in Tab. 1. In general, there are 12,000 and 13,056 samples in PointWire and PointVessel, respectively. For both PointWire and PointVessel, each sample includes 2048 points. The PointWire dataset consists of the wire, bifurcation, and endpoint classes, while the PointVessel dataset has two classes: vessel and bifurcation. On average, PointWire and PointVessel have 26 and 104 instances per sample, respectively. Additionally, we split each dataset into training, validation and test set with the ratio 80 %, 10 % and 10 %. The distribution of class instances for both datasets is given in Fig. 6. Overall, the distribution of the PointWire dataset varies across the training, validation, and test sets. In contrast, the PointVessel dataset has a more balanced distribution and a relatively higher number of class instances per sample in each set.

4. Experiments

This section provides the baseline benchmark experiments for our proposed two datasets, PointWire and PointVessel, on the topology segmentation task. Additionally, we conduct transferability experiments of the select state-of-the-art point cloud segmentation methods on the PointWire and PointVessel datasets. The intention is to analyze whether the methods are capable of learning general topological concepts of two different DLO object types, i.e., wiring harnesses and blood vascular systems.

Method	mIoU↑	mAcc↑	OA↑	w IoU↑	e IoU↑	b IoU↑
PointNet++ [31]	49.08	69.53	73.09	68.49	54.86	23.89
DGCNN [38]	54.84	80.56	73.86	67.39	66.90	30.25
PCT [12]	36.68	59.27	62.01	56.59	30.64	22.8
CurveNet [43]	53.92	68.96	79.91	76.92	53.31	31.53
DeltaConv [39]	60.47	81.06	80.34	75.64	70.49	35.29
RepSurf [32]	**73.6**	**85.6**	**85.22**	**81.52**	**72.72**	**40.14**

Table 2. PointWire (all classes) benchmark. We evaluate different segmentation methods on the test set. RepSurf [32] outperforms other approaches by significant gaps.

4.1. Experiment Metrics

Point cloud segmentation evaluation metrics. We evaluate different methods on our proposed datasets in terms of mean Intersection over Union (mIoU, %), mean accuracy (mACC, %), overall accuracy (OA, %), Length-Mean Average Error (L-MAE), Segment-Mean Average Error (S-MAE), and IoU for each class: wire (w IoU), endpoint (e IoU), and bifurcation (b IoU).

Disentanglement evaluation metrics. We evaluate the disentanglement by using the mean absolute error to calculate the total length error (L-MAE) and the segment-wise error (S-MAE). L-MAE aggregates the lengths of all segments of a skeleton. For S-MAE we do a bipartite matching of the bifurcation and endpoints. We then compare the individual segment lengths of the ground truth to the length of the path between the matched graph nodes.

4.2. Implementation Details

We evaluate PointNet++ [31], DGCNN [38], PCT [12], CurveNet [43], DeltaConv [39], and RepSurf [32] on our large-scale benchmarks to get the baseline results for topological segmentation. Furthermore, we show the transferability performance of each network over the two datasets. For the baseline method of the disentanglement benchmark, we generate the nodes with semantic laplacian-based contraction [23] and connect them with minimum spanning tree. All the experiments are conducted using one NVIDIA GeForce RTX 3090 (24Gb). For all network training on the two datasets, we set the batch size as 8 and modified the learning rate according to the Linear Scaling Rule [11].

4.3. PointWire Benchmark

PointWire (all classes). The quantitative evaluation of the chosen point cloud segmentation methods on the PointWire all classes are given in Tab. 2. RepSurf [32] and DeltaConv [39] outperform them by a large margin the other methods. The reason for that lies in the point cloud surface learning strategies used by RepSurf and DeltaConv. By utilizing local 2D manifold feature aggregation, these methods are able to generalize well on the challenging PointWire

dataset. Comparing the qualitative results of the evaluation on the PointWire test set, given in the first row of Fig. 7, we can observe that both RepSurf and DeltaConv correctly segment the bifurcations and endpoints with the least amount of false positives. PointNet++ [31] and DGCNN [38] heavily under-segmented the bifurcations, while PCT [12] and CurveNet [43] have relatively more false positives located at wire segments. However, based on the results from Tab. 2, it is clear that the PointWire dataset is rather challenging, especially for the bifurcation class since the highest IoU score is $40.14\ \%$, achieved by RepSurf [32].

PointWire (bifurcation). We show the quantitative results of PointWire bifurcation in Tab. 3. The reason for conducting the PointWire$_{bifurcation}$ experiment is to have an identical setting, w.r.t. the same number of classes, as in the PointVessel benchmark. Similar to the results on PointWire all classes benchmark, the surface-based method DeltaConv performs best on the PointWire bifurcation experiment. Compared with the PointWire$_{all}$, CurveNet and PointNet++ achieve better performance on the PointWire$_{bifurcation}$ benchmark in terms of mIoU metrics. Moreover, by comparing the qualitative results shown in the second row in Fig. 7 DeltaConv and PointNet++ have the visually best results. CurveNet fails to identify bifurcation segments and collapses into a mode of segmenting almost the entire wiring harness as the wire class. The other methods [12, 32, 38] have the tendency to under-segment the bifurcation class, which results in less overall accuracy (OA).

4.4. PointVessel Benchmark

The quantitative results on the PointVessel benchmark are reported in Tab. 3. Overall, the selected methods achieve better mAcc and mIoU compared to the PointWire benchmark. The reasons for that are twofold: 1) Class distributions in PointVessel are similar, and 2) PointVessel is synthetic with no noise and a complete surface of the underlying structures. Notably, PointNet++ [31] along with CurveNet [43] achieve the best results. The superior performance of CurveNet and PointNet++ can likely be attributed to their feature aggregation strategies, which are particularly well-suited for clean, synthetic data. In particular, the curve feature aggregation approach introduced by CurveNet [43] achieves the best results on smooth and continuous point cloud surfaces.

Examining the quantitative and qualitative results for the bifurcation class, as presented in Tab. 3 and the first row of Fig. 8, it is evident that surface learning-based methods generally perform well. In contrast, PCT [12] achieves the worst results. This suboptimal performance can primarily be attributed to the inherent limitations of transformers, particularly the need for large amounts of data to function effectively. Another interesting observation is that all the models have difficulties in correctly segmenting bifurcations in larger blood vessels. However, they perform successfully when segmenting bifurcations located at blood vessels with a smaller diameter.

4.5. Transferability Benchmark

The transferability benchmarks are designed to demonstrate the capability of existing point cloud segmentation methods to learn the general topological concepts of DLOs. These experiments provide a comprehensive evaluation of point learning techniques by assessing not only their domain-specific learning capabilities but also their potential to generalize high-level concepts across similar domains.

Therefore, we train six frameworks [12,31,32,38,39,43] on the PointWire and the PointVessel datasets separately, and then cross-evaluated on the other's dataset test set. For example, we train the networks on the PointVessel train set and evaluated them on the PointWire test set. The experimental results are shown in Tab. 3. Moreover, to show the performance of models pretrained on non-DLO data, we directly evaluate GDCNN [38] and PAConv [44], which are trained on the widely used ShapeNetPart [49] dataset on our proposed two datasets without fine-tuning (see Tab. 4).

PointVessel $\rightarrow$ PointWire. The quantitative results of the transferability experiments by training first on the PointVessel train set and then evaluating on the PointWire test set are reported in Tab. 3. PointNet++ [31] outperforms the other state-of-the-art methods by achieving a mIoU of $47.98\ \%$. This demonstrates quite good generalization capabilities of PointNet++ when trained on complete synthetic point cloud data. A potential reason is its ability to effectively capture and process local geometric features and topological relationships. Furthermore, a general observation is that training on complete synthetic PointVessel data leads to unsatisfied generalization on the more challenging incomplete, and noisy PointWire dataset for most of the other state-of-the-art methods. This is the case, especially for the bifurcation class. Most of the baseline methods achieve single-digit IoU results. The qualitative results shown in the third row of Fig. 7 further confirm the aforementioned observations.

PointWire $\rightarrow$ PointVessel. We present the transferability experiments from PointWire to the Point Vessel test set in Table 3. Compared to the *PointVessel $\rightarrow$ PointWire* benchmark, most of the models achieve relatively good performances, especially on the bifurcation class. In terms of mIoU, DGCNN outperforms the second best performing method DeltaConv when trained on all PointWire classes by almost $6\ \%$, and is on par when trained on the PointWire$_{bifurcation}$ train set. This indicates that point cloud graph representations learn well general topological

Benchmark	PointWire$_{bifurcation}$ PointVessel → PointWire$_{bifurcation}$ -					PointVessel PointWire$_{all}$ → PointVessel PointWire$_{bifurcation}$ → PointVessel					PointWire$_{bifurcation}$ PointVessel -	
Method	**mIoU↑**	**mAcc↑**	**OA↑**	**w IoU↑**	**b IoU↑**	**mIoU↑**	**mAcc↑**	**OA↑**	**v IoU↑**	**b IoU↑**	**L-MAE↓**	**S-MAE↓**
PointNet++ [31]	56.74	74.22	84.36	83.22	30.27	84.79	94.32	92.34	88.64	80.93	0.89	0.48
	47.98	75.20	71.71	68.94	27.02	31.27	42.48	52.21	52.81	9.73	0.98	0.27
	-	-	-	-	-	43.83	57.91	70.64	68.43	19.23	-	-
DGCNN [38]	49.28	78.71	72.87	69.48	29.08	77.44	89.15	88.21	83.19	71.70	**0.79**	0.48
	30.33	48.11	55.4	50.85	9.81	52.02*	66.70	72.21	66.70	37.35	0.99	0.27
	-	-	-	-	-	45.94	66.08	63.67	52.36	39.51	-	-
PCT [12]	45.78	60.86	75.23	73.84	17.72	57.38	76.70	73.62	63.46	51.31	0.90	0.49
	35.31	51.31	65.71	60.19	10.42	40.34	55.02	62.9	58.51	22.17	1.01	**0.26**
	-	-	-	-	-	38.14	54.95	57.36	49.46	26.82	-	-
CurveNet [43]	57.3	67.86	87.63	87.02	27.57	84.99	93.02	92.63	89.38	80.60	0.87	0.49
	45.5	53.54	81.86	81.51	9.49	41.41	55.81	69.52	67.78	15.05	0.99	0.27
	-	-	-	-	-	40.49	55.13	69.62	68.21	12.77	-	-
DeltaConv [39]	**59.22**	81.09	84.47	83.03	35.41	84.64	92.75	92.45	89.15	80.13	0.86	0.46
	37.72	63.96	62.17	59.17	16.27	46.03	59.88	71.16	68.31	23.75	0.98	0.27
	-	-	-	-	-	**46.14**	59.97	71.4	68.62	23.66	-	-
RepSurf [32]	52.94	83.81	77.97	74.78	31.10	79.1	90.24	89.19	84.48	73.72	0.87	0.49
	37.04	50.72	79.89	71.21	2.87	44.27	58.30	69.68	66.91	21.63	0.99	0.27
	-	-	-	-	-	42.05	62.68	59.74	47.63	36.47	-	-

Table 3. PointWire bifurcation benchmark, PointVessel benchmark, and transferability benchmark. **PointWire**$_{bifurcation}$: experiments using only the wire and bifurcation classes. → indicates that the model is pretrained on the dataset on the left side of the arrow and tested on the right side dataset. _, ∗ and **bold** denote the best results for each benchmark.

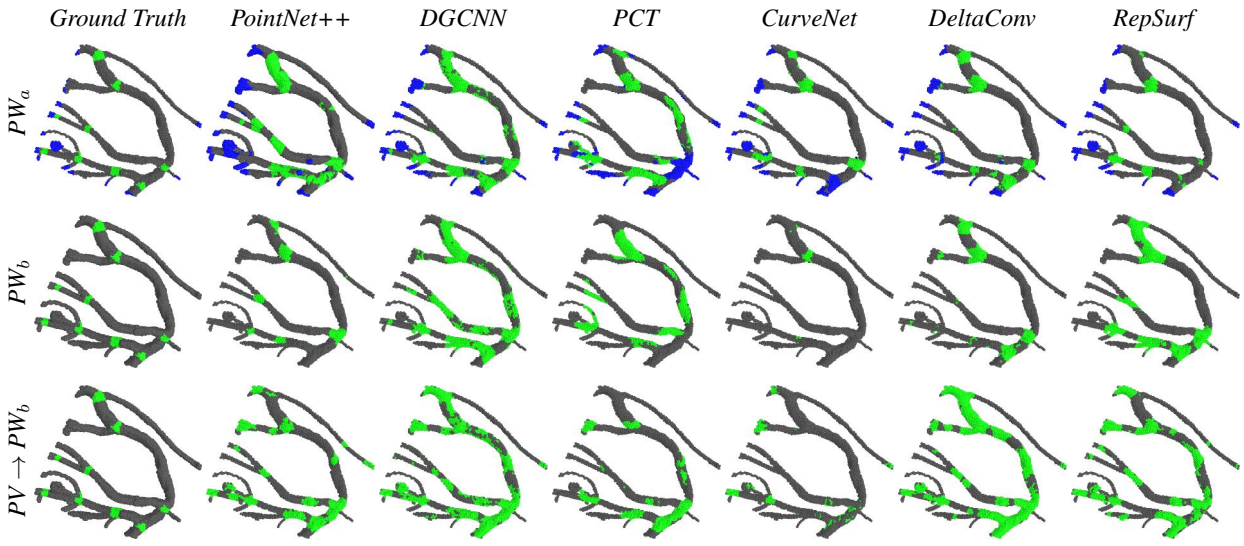

Figure 7. Qualitative results on PointWire all classes (PW$_a$), PointWire bifurcation (PW$_b$) and the transferability benchmark on the PointWire bifurcation (PV → PW$_b$). For PW$_a$ and PW$_b$, gray points are the wires, endpoints are blue, and green ones are bifurcations.

DLO segmentation concepts. Moreover, the surface representation learning-based methods [32, 39] are also reasonably good at capturing general DLO topology. The qualitative results are given in the second and third rows of Fig. 8.

ShapeNetPart → PointWire and PointVessel. The results of this evaluation are presented in Tab. 4. Overall, the performance of DGCNN is significantly lower compared to when it is pretrained on DLO datasets. This discrepancy highlights the challenges associated with understanding DLOs and underscores the importance of our proposed datasets for advancing research in this area.

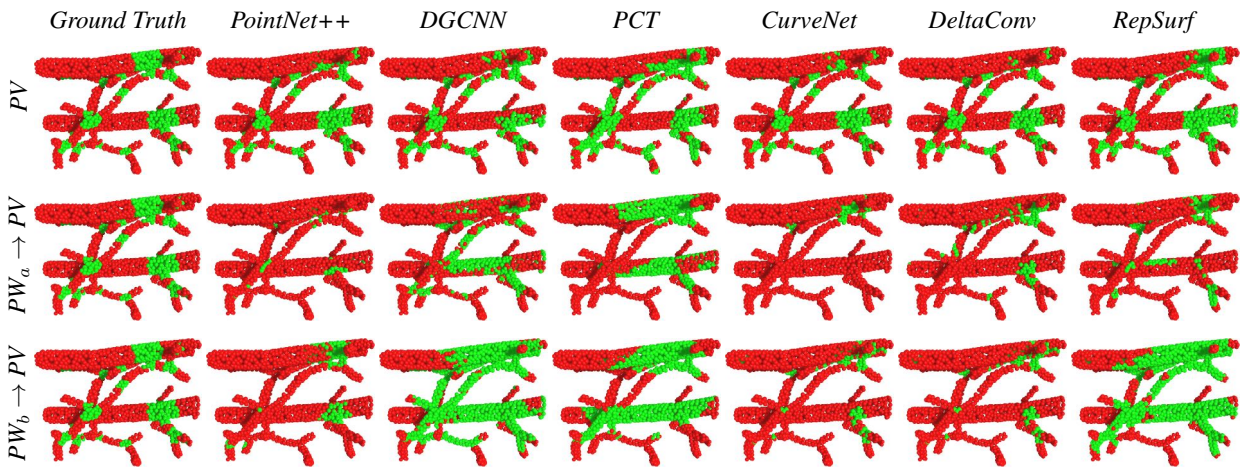

Figure 8. Qualitative results on PointVessel and transferability benchmark on PointVessel. PointWire all classes to PointVessel ($PW_a \rightarrow PV$), PointWire bifurcation to PointVessel ($PW_b \rightarrow PV$). We show vessels and bifurcations in red and green.

Method	$PointWire_b$			PointVessel		
	mIoU	**mAcc**	**OA**	**mIoU**	**mAcc**	**OA**
DGCNN [38]	6.13	31.87	8.19	16.43	50.05	32.68
PAConv [44]	10.64	11.6	15.85	33.73	49.98	67.46

Table 4. Evaluation results of pretrained models from ShapeNet-Part [49] on $PointWire_{bifurcation}$ and PointVessel.

4.6. Disentanglement Benchmark

The results of the disentanglement benchmark are shown in the last two columns of Tab. 3. DGCNN yields the best results in terms of L-MAE on the $PointWire_{bifurcation}$ dataset, likely due to its heavy under-segmentation in thick DLO parts. In contrast, DeltaConv shows superior performance in terms of S-MAE, which can be attributed to its high-quality segmentation results. On the PointVessel disentanglement benchmark, PointNet++ and DeltaConv achieve the best performance in terms of L-MAE, which is consistent with their promising segmentation results.

5. Discussion

Research Impact. Manufacturing processes including wiring harnesses are still carried out manually due to the significant complexities associated with both the perception and manipulation. The proposed datasets can encourage research on the understanding of such DLOs with intricate shapes, thereby facilitating advancements in automation and promoting more robotized environments. Moreover, this could relieve human workers from repetitive and effortful tasks as well as improve the supply chain where the wiring harness is a major bottleneck.

Effective Model Architectures. The experimental results in Tab. 2 and Tab. 3 suggest that incorporating topological relationship learning and adaptive surface-based representations is essential for accurately understanding 3D DLOs. For instance, models like RepSurf and CurveNet consistently exhibit compatible performance across both datasets and most tasks.

Limitations. Due to intellectual property constraints, significantly increasing the volume of data in the PointWire dataset is challenging. However, the proposed semi-automatic processing pipeline enables researchers to generate data according to their specific needs. On the other hand, although it lacks the properties of real-world medical data, this limitation can be mitigated through the application of data augmentation techniques.

6. Conclusion

In this study, we introduced two point cloud datasets for DLO understanding, i.e. the real-world PointWire dataset and the synthetic PointVessel dataset. PointWire is derived from real wiring harness scans, while PointVessel contains synthetic blood vessel point clouds. We proposed a semi-automatic pipeline to generate point cloud-based wiring harness data. Furthermore, we conducted transferability benchmarks to comprehensively analyze the generalization capabilities of six popular point cloud segmentation methods in terms of 3D DLO topology segmentation and disentanglement. The benchmarks indicate that DLO topology segmentation is challenging for existing methods. Therefore, we strongly believe that the joint topology segmentation and disentanglement tasks can have a significant impact on the industry, and further research is encouraged.

References

[1] Junjie Cao, Andrea Tagliasacchi, Matt Olson, Hao Zhang, and Zhixun Su. Point cloud skeletons via laplacian-based contraction. In *Proc. of IEEE Conf. on Shape Modeling and Applications*, 2015.

[2] Alessio Caporali, Kevin Galassi, Riccardo Zanella, and Gianluca Palli. Fastdlo: Fast deformable linear objects instance segmentation. *IEEE Robotics and Automation Letters*, 7(4):9075–9082, 2022.

[3] Alessio Caporali, Kevin Galassi, Bare Luka Žagar, Riccardo Zanella, Gianluca Palli, and Alois C Knoll. Rt-dlo: Real-time deformable linear objects instance segmentation. *IEEE Transactions on Industrial Informatics*, pages 1–10, 2023.

[4] Alessio Caporali, Riccardo Zanella, Daniele De Greogrio, and Gianluca Palli. Ariadne+: Deep learning–based augmented framework for the instance segmentation of wires. *IEEE Transactions on Industrial Informatics*, 18(12):8607–8617, 2022.

[5] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.

[6] Andrew Choi, Dezhong Tong, Brian Park, Demetri Terzopoulos, Jungseock Joo, and Mohammad Khalid Jawed. mbest: Realtime deformable linear object detection through minimal bending energy skeleton pixel traversals. *arXiv preprint arXiv:2302.09444*, 2023.

[7] Jasmine Collins, Shubham Goel, Kenan Deng, Achleshwar Luthra, Leon Xu, Erhan Gundogdu, Xi Zhang, Tomas F Yago Vicente, Thomas Dideriksen, Himanshu Arora, et al. Abo: Dataset and benchmarks for real-world 3d object understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21126–21136, 2022.

[8] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018.

[9] Konrad P Cop, Arne Peters, Bare L Žagar, Daniel Hettegger, and Alois C Knoll. New metrics for industrial depth sensors evaluation for precise robotic applications. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5350–5356. IEEE, 2021.

[10] Fan Fu, Jianyong Wei, Miao Zhang, Fan Yu, Yueting Xiao, Dongdong Rong, Yi Shan, Yan Li, Cheng Zhao, Fangzhou Liao, et al. Rapid vessel segmentation and reconstruction of head and neck angiograms using 3d convolutional neural network. *Nature communications*, 11(1):4829, 2020.

[11] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.

[12] 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.

[13] Ali Khandouzi, Ali Ariafar, Zahra Mashayekhpour, Milad Pazira, and Yasser Baleghi. Retinal vessel segmentation, a review of classic and deep methods. *Annals of Biomedical Engineering*, 50(10):1292–1314, 2022.

[14] Piotr Kicki, Michał Bednarek, Paweł Lembicz, Grzegorz Mierzwiak, Amadeusz Szymko, Marek Kraft, and Krzysztof Walas. Tell me, what do you see?—interpretable classification of wiring harness branches with deep neural networks. *Sensors*, 21(13):4327, 2021.

[15] Christoph Kirst, Sophie Skriabine, Alba Vieites-Prado, Thomas Topilko, Paul Bertin, Gaspard Gerschenfeld, Florine Verny, Piotr Topilko, Nicolas Michalski, Marc Tessier-Lavigne, et al. Mapping the fine-scale organization and plasticity of the brain vasculature. *Cell*, 180(4):780–795, 2020.

[16] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9601–9611, 2019.

[17] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. *Advances in neural information processing systems*, 31, 2018.

[18] Yuxin Li, Tong Ren, Junhuai Li, Huaijun Wang, Xiangning Li, and Anan Li. Vbnet: An end-to-end 3d neural network for vessel bifurcation point detection in mesoscopic brain images. *Computer Methods and Programs in Biomedicine*, 214:106567, 2022.

[19] Yiqun Lin, Zizheng Yan, Haibin Huang, Dong Du, Ligang Liu, Shuguang Cui, and Xiaoguang Han. Fpconv: Learning local flattening for point convolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4293–4302, 2020.

[20] Yongcheng Liu, Bin Fan, Shiming Xiang, and Chunhong Pan. Relation-shape convolutional neural network for point cloud analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8895–8904, 2019.

[21] Yifan Liu, Jie Liu, and Yixuan Yuan. Edge-oriented point-cloud transformer for 3d intracranial aneurysm segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part V*, pages 97–106. Springer, 2022.

[22] Kangchen Lv, Mingrui Yu, Yifan Pu, and Xiang Li. Learning to occlusion-robustly estimate 3-d states of deformable linear objects from single-frame point clouds. *arXiv preprint arXiv:2210.01433*, 2022.

[23] Lukas Meyer, Andreas Gilson, Oliver Scholz, and Marc Stamminger. Cherrypicker: Semantic skeletonization and topological reconstruction of cherry trees. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6243–6252, 2023.

[24] Muthu Rama Krishnan Mookiah, Stephen Hogg, Tom J MacGillivray, Vijayaraghavan Prathiba, Rajendra Pradeepa, Viswanathan Mohan, Ranjit Mohan Anjana, Alexander S

Doney, Colin NA Palmer, and Emanuele Trucco. A review of machine learning methods for retinal blood vessel segmentation and artery/vein classification. *Medical Image Analysis*, 68:101905, 2021.

[25] Gabriel E Navas-Reascos, David Romero, Johan Stahre, and Alberto Caballero-Ruiz. Wire harness assembly process supported by collaborative robots: Literature review and call for r&d. *Robotics*, 11(3):65, 2022.

[26] Huong Giang Nguyen, Resul Habiboglu, and Jörg Franke. Enabling deep learning using synthetic data: A case study for the automotive wiring harness manufacturing. *Procedia CIRP*, 107:1263–1268, 2022.

[27] Huong Giang Nguyen, Marlene Kuhn, and Jörg Franke. Manufacturing automation for automotive wiring harnesses. *Procedia CIRP*, 97:379–384, 2021.

[28] Thong Phi Nguyen and Jonghun Yoon. A novel vision-based method for 3d profile extraction of wire harness in robotized assembly process. *Journal of Manufacturing Systems*, 61:365–374, 2021.

[29] Johannes C Paetzold, Julian McGinnis, Suprosanna Shit, Ivan Ezhov, Paul Büschl, Chinmay Prabhakar, Anjany Sekuboyina, Mihail Todorov, Georgios Kaissis, Ali Ertürk, et al. Whole brain vessel graphs: A dataset and benchmark for graph learning and neuroscience. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.

[30] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.

[31] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017.

[32] Haoxi Ran, Jun Liu, and Chengjie Wang. Surface representation for point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18942–18952, 2022.

[33] Will Schroeder, Ken Martin, and Bill Lorensen. *The Visualization Toolkit–An Object-Oriented Approach To 3D Graphics*. Kitware, Inc., fourth edition, 2006.

[34] Giles Tetteh, Velizar Efremov, Nils D. Forkert, Matthias Schneider, Jan Kirschke, Bruno Weber, Claus Zimmer, Marie Piraud, and Björn H. Menze. Deepvesselnet: Vessel segmentation, centerline prediction, and bifurcation detection in 3-d angiographic volumes. *Frontiers in Neuroscience*, 14, 2020.

[35] Mihail Ivilinov Todorov, Johannes Christian Paetzold, Oliver Schoppe, Giles Tetteh, Suprosanna Shit, Velizar Efremov, Katalin Todorov-Völgyi, Marco Düring, Martin Dichgans, Marie Piraud, et al. Machine learning analysis of whole mouse brain vasculature. *Nature methods*, 17(4):442–449, 2020.

[36] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Duc Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *International Conference on Computer Vision (ICCV)*, 2019.

[37] Aly Al-Amyn Valliani, Daniel Ranti, and Eric Karl Oermann. Deep learning and neurology: a systematic review. *Neurology and therapy*, 8:351–365, 2019.

[38] 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 Transactions On Graphics (tog)*, 38(5):1–12, 2019.

[39] Ruben Wiersma, Ahmad Nasikun, Elmar Eisemann, and Klaus Hildebrandt. Deltaconv: anisotropic operators for geometric deep learning on point clouds. *ACM Transactions on Graphics (TOG)*, 41(4):1–10, 2022.

[40] Rundi Wu, Chang Xiao, and Changxi Zheng. Deepcad: A deep generative network for computer-aided design models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6772–6782, 2021.

[41] Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Jiawei Ren, Liang Pan, Wayne Wu, Lei Yang, Jiaqi Wang, Chen Qian, et al. Omniobject3d: Large-vocabulary 3d object dataset for realistic perception, reconstruction and generation. *arXiv preprint arXiv:2301.07525*, 2023.

[42] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong 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.

[43] Tiange Xiang, Chaoyi Zhang, Yang Song, Jianhui Yu, and Weidong Cai. Walk in the cloud: Learning curves for point clouds shape analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 915–924, 2021.

[44] Mutian Xu, Runyu Ding, Hengshuang Zhao, and Xiaojuan Qi. Paconv: Position adaptive convolution with dynamic kernel assembling on point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3173–3182, 2021.

[45] Mutian Xu, Junhao Zhang, Zhipeng Zhou, Mingye Xu, Xiaojuan Qi, and Yu Qiao. Learning geometry-disentangled representation for complementary understanding of 3d object point cloud. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 3056–3064, 2021.

[46] Yifan Xu, Tianqi Fan, Mingye Xu, Long Zeng, and Yu Qiao. Spidercnn: Deep learning on point sets with parameterized convolutional filters. In *Proceedings of the European conference on computer vision (ECCV)*, pages 87–102, 2018.

[47] Xi Yang, Ding Xia, Taichi Kin, and Takeo Igarashi. Intra: 3d intracranial aneurysm dataset for deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2656–2666, 2020.

[48] Xi Yang, Ding Xia, Taichi Kin, and Takeo Igarashi. A two-step surface-based 3d deep learning pipeline for segmentation of intracranial aneurysms. *Computational Visual Media*, 9(1):57–69, 2023.

[49] Li Yi, Vladimir G Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas. A scalable active framework for

region annotation in 3d shape collections. *ACM Transactions on Graphics (ToG)*, 35(6):1–12, 2016.

[50] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16259–16268, 2021.

[51] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Open3d: A modern library for 3d data processing. *arXiv preprint arXiv:1801.09847*, 2018.