



Sharp Feature Consolidation from Raw 3D Point Clouds via Displacement Learning

Tong Zhao, Mulin Yu, Pierre Alliez, Florent Lafarge

► To cite this version:

Tong Zhao, Mulin Yu, Pierre Alliez, Florent Lafarge. Sharp Feature Consolidation from Raw 3D Point Clouds via Displacement Learning. Computer Aided Geometric Design, 2023, 13, pp.102204. 10.1016/j.cagd.2023.102204 . hal-03747150v2

HAL Id: hal-03747150

<https://inria.hal.science/hal-03747150v2>

Submitted on 14 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Sharp Feature Consolidation from Raw 3D Point Clouds via Displacement Learning

Tong Zhao^{a,1}, Mulin Yu^{a,1,*}, Pierre Alliez^a, Florent Lafarge^a

^a*Inria, Université Côte d’Azur, France*

Abstract

Detecting sharp features in raw 3D point clouds is an essential step for designing efficient priors in several 3D Vision applications. This paper presents a deep learning-based approach that learns to detect and consolidate sharp feature points on raw 3D point clouds. We devise a multi-task neural network architecture that identifies points near sharp features and predicts displacement vectors toward the local sharp features. The so-detected points are thus consolidated via relocation. Our approach is robust against noise by utilizing a dynamic labeling oracle during the training phase. The approach is also flexible and can be combined with several popular point-based network architectures. Our experiments demonstrate that our approach outperforms the previous work in terms of detection accuracy measured on the popular ABC dataset. We show the efficacy of the proposed approach by applying it to several 3D vision tasks.

Keywords: Sharp feature detection, Displacement learning, Multi-task neural network

1. Introduction

For scanned or Computer-Aided Design (CAD) 3D models, a sharp feature usually refers to creases or corners where the surface is not smooth. Recognizing such sharp features from raw point cloud is a preliminary step for several point cloud processing tasks such as surface reconstruction (Öztireli et al., 2009; Gal et al., 2007; Amenta et al., 1998; Guennebaud and Gross, 2007; Salman et al., 2010), extraction of feature graphs (Lee and Bo, 2016; Ni et al., 2016) or semantic segmentation (Gong et al., 2021; Loizou et al., 2020). With the development of point cloud acquisition techniques and the release of well-annotated 3D datasets (Koch et al., 2019), many data-driven approaches have been proposed in recent years. Most previous approaches reformulate sharp feature detection as a classification problem, which leads to several issues: (1) Feature points are not located exactly on sharp features in general, which leads to incomplete features; (2) To deal with this issue, several approaches introduce a distance threshold that defines the maximum distance between a feature point and the nearest sharp feature. Using such a distance threshold in common classification problems highly impacts the final prediction and it is delicate to choose (see Figure 1); (3) Noise has a large impact over the classification results; and (4) The network can be easily overfitted given patches or point clouds with ground truth labels.

Inspired by displacement-based point cloud denoising methods, we propose a deep learning framework for detecting and consolidating sharp features from a raw 3D point cloud, namely SFCNet. Our framework contains two main components: (1) A point-to-feature oracle that largely improves the data augmentation quality so as to prevent models from overfitting; (2) A multi-task neural network architecture that predicts altogether a binary label (sharp/smooth) for each input point and a displacement vector for each sharp feature point. Our approach outperforms popular unsupervised methods and supervised methods on the

*Corresponding author

Email addresses: tong.zhao@inria.fr (Tong Zhao), mulin.yu@inria.fr (Mulin Yu), pierre.alliez@inria.fr (Pierre Alliez), florent.lafarge@inria.fr (Florent Lafarge)

¹These authors contributed equally.




Figure 1: Illustrating the impact of distance thresholds and displacement vectors on sharp feature detection. (a) Noisy point cloud around a sharp feature. (b) Given a small distance threshold, existing classification-based methods detect few accurate sharp feature points, which do not capture well the entire sharp feature. (c) Given a large distance threshold, the detected sharp feature points become vague. (d) Our method can detect the sharp feature points within a large distance threshold, and the learned displacement vectors relocate the points closer to the sharp feature. (e) Resulting accurate detection.

ABC dataset (Koch et al., 2019). We also evaluate our method on noisy point clouds and real scanned models. Quantitative and qualitative results are provided in Section 4.

In summary, our insights and technical contributions are:

- We propose a novel end-to-end multi-task neural network architecture that performs accurate sharp feature detection and consolidation;
- Our neural network has a flexible structure that can be combined with many popular backbones and adapted to many datasets;
- We propose a point-to-feature oracle that largely improves the data augmentation quality during training to prevent models from overfitting;
- Our approach has a good resilience to noise.

2. Related work

In this section, we first review the popular neural network structures for extracting geometric features from raw point clouds. We next review related work on sharp feature detection and consolidation from raw point clouds, including unsupervised methods and supervised methods. We then review the pioneering work on feature-preserving mesh denoising, which can also be utilized for sharp feature detection and consolidation from raw point clouds, when combined with an extra surface reconstruction step. We also briefly review point cloud denoising methods, focusing on displacement-based methods that motivate our displacement learning approach.

2.1. Learning-based geometric feature extraction

While convolutional neural networks yield great results for many image-related tasks, point cloud processing remains a challenge due to the unstructured and irregular nature of point clouds. The PointNet approach proposed by Qi et al. (2017a) quickly became a popular learning approach for point cloud processing, and was applied to many tasks. It extracts point features from point coordinates by applying input and feature transformations and then aggregates global point features by max pooling. Qi et al. (2017b) then proposed PointNet++, which improves over previous work by considering local geometries with sampling and grouping layers. Guerrero et al. (2018) proposed another variant of PointNet, referred to as PCPNet. Instead of taking the whole point cloud as input, PCPNet deals with patches to learn specific local point features. More recently, Qian et al. (2021) proposed ASSANet, a novel separable set abstraction module that further improves PointNet++.

Another family of methods is spatial-based. Wang et al. (2019) introduced a graph neural network, referred to as DGCNN, which learns features on KNN neighbors using successive EdgeConv layers. The So-Net (Li et al., 2018) is devised to model the spatial distribution of the input points. The SPLATNet (Su et al., 2018) computes hierarchical and spatially-aware features of input points with sparse and efficient lattice filters. All the above networks can be utilized as a backbone of our framework, as long as they produce pointwise descriptors.

2.2. Sharp feature detection and consolidation

Unsupervised methods. Several unsupervised methods rely upon Moving Least Squares (MLS) surface reconstruction (Levin, 2004). Fleishman et al. (2005) proposed a Moving Least Squares method for reconstructing a piecewise smooth surface from point clouds. Sharp features are identified as points at the intersection of multiple iteratively fitted surfaces. Daniels et al. (2007) improved over this method by using an adaptive threshold. They define a projection operator which takes into account edge points, followed by a smoothing filter. However, the performance of these methods highly depends on the quality of surface reconstruction results, which may fail in noisy cases. Moreover, most of these methods are computationally intensive.

Huang et al. (2013) contributed an edge-aware point cloud resampling technique built upon a normal estimation step followed by a robust projection operator. The outcome highly depends on the quality of the normal estimation.

Another line of methods detect sharp features via local geometric descriptors. Weber et al. (2010) rely upon point-sampled geometries. They first eliminate planar points by performing a flatness test. Clustering over the Gauss map is then performed to find sharp feature points. M  rigot et al. (2010) utilize convolved covariance matrices of Voronoi cells to characterize point properties and then filter out smooth points by a fixed threshold. Bazazian et al. (2015) accelerated it by computing tensors directly from the nearest neighbors of points. Chen et al. (2021) proposed a two-phase algorithm in which they first localize candidate feature points by estimating a joint metric of plane fitting residual and local surface variation, then extract consolidated feature points by filtering the normal field and optimizing the coordinates via an anisotropic contracting scheme. Guo et al. (2022) recently proposed another two-phase algorithm. An improved α -shape technique is devised to generate two graphs with potential feature points and neighborhood description, followed by a subgraph-based local binary pattern operator to remove non-feature points. The above methods are computationally efficient but defining relevant thresholds requires a trial-and-error process.

Supervised methods. With the development of neural networks, more and more supervised methods are proposed for point cloud processing. Yu et al. (2018a) proposed EC-net, in which they extend PU-Net (Yu et al., 2018b), a network designed to upsample and perform an edge-aware point cloud consolidation. Despite its success, its results are impacted by the way the patches are created. Wang et al. (2020) contributed an end-to-end neural network, called PIE-Net, that is trained for parameter inference of feature edges over a 3D point cloud, where the output consists of one or more parametric curves. However, it is limited in terms of scalability. Loizou et al. (2020) devised a convolutional neural network based on DGCNN (Wang et al., 2019), to detect the boundaries of parts in 3D point clouds. Himeur et al. (2022) introduced PCEDNet, a lightweight neural network that takes as input a series of multi-scale differential geometric descriptors of points and predicts three-class labels for a sharp feature and its surroundings. In contrast, our framework learns displacement while detecting sharp features, which offers higher genericity and robustness to noise. Matveev et al. (2022) introduced recently DEF, a learning-based framework which estimates a distance-to-feature field for sampled 3D shapes. A set of depth images are generated on local patches and passed to an image-based neural network. The output distance prediction of an individual point is blended from the overlapping patches. Compared to point cloud-based networks, DEF needs a more sophisticated data annotation procedure and the predictions are influenced by feature visibility.

2.3. Feature-preserving mesh denoising

Many filtering-based approaches have been proposed in the last two decades since its successful applications in image-related tasks. Fleishman et al. (2003) first proposed a bilateral mesh denoising approach by




Figure 2: Overview of the proposed learning architecture. It comprises four modules: *point-to-feature oracle* (purple dash), *point descriptor extractor* (red dash), *sharp feature detector* (green dash) and *displacement predictor* (blue dash). *Point-to-feature oracle* module is only used for training. It introduces noise to the input, computes the ground truth classification labels and the ground truth displacement vectors. A backbone is used to extract high dimensional representations $\mathbf{f}$ (size: $n \times D_2$) in the *point descriptor extractor* module. **MLP** stands for multi-layer perceptron. *Cla* and *Dis* represent binary sharp feature labels and displacement vectors, respectively. $\oplus$ adds the corresponding predicted displacement vectors to detected sharp feature points to make them lie closer to the ground truth sharp features. *BCE* and *Masked MSE* are two losses calculated base on *Cla*, *Dis* and the ground truth generated by the oracle.

filtering mesh vertices in the normal direction, iteratively. Jones et al. (2003) introduced another single-pass bilateral filtering approach that can deal with arbitrary triangle soups. Wang (2006a) further improved the previous approach by integrating the filter with remeshing operators to iteratively recover sharp features. Zheng et al. (2010) introduced a novel bilateral filter on the facet normal field, taking into consideration both spatial difference and signal difference into bilateral weighting. Zhang et al. (2015) devised a two-stage approach in which they apply a joint bilateral filter on face normals guided by an estimated normal field before updating the vertex position. Instead of filtering vertices, Attene et al. (2003) proposed the Edge sharpener that first filters chamfer triangles and then subdivides them by inserting new vertices located on sharp features. Wang (2006b) further improved this incremental approach by introducing a new sharp-fold detector and a skeleton-guided vertex relocation step.

Another family of approaches are optimization-based. He and Schaefer (2013) introduced a mesh denoising approach via L_0 minimization. The vertex locations are optimized by minimizing the L_0 norm of the signal gradients. Wang et al. (2014) devised a two-stage approach in which they compute at first a smoothly regularized mesh and then recover sharp features from the residual between the regularized mesh and the original mesh via l_1 analysis. With recent trend of data-driven techniques, Zhao et al. (2019) proposed NormalNet, a learning-based normal filtering approach for mesh denoising. They followed the same iterative scheme of Zhang et al. (2015) and replaced the estimated guidance field by a deep neural network.

While sharp feature detection and consolidation can be performed by combining surface reconstruction approaches and feature-preserving mesh denoising approaches, the former step may introduce extra errors into the pipeline and the latter step may depend on the quality, the density, or the manifoldness of the reconstructed mesh.

2.4. Displacement-based point cloud denoising

As the displacement property indicates the mismatches between the true model and the sampling, it can be utilized to reduce or even eliminate the intrinsic noise in point clouds. Rakotosaona et al. (2020) proposed PointcleanNet, which pioneered the idea of displacement learning for point cloud denoising. They designed a two-step neural network based on PCPNet (Guerrero et al., 2018) that first eliminates outliers before estimating displacement vectors for all inliers. Instead, our method uses a shared backbone for both classification and regression, which significantly reduces the number of parameters. Pistilli et al. (2020) introduced a fully convolutional network predicting displacement vectors. The success of the above methods shows the potential of displacement learning for sharp feature detection.

3. Our method

During training session, our method takes as input a ground truth mesh $\mathcal{M}$, a set of ground truth sharp features $\mathcal{F}$ (with parametric representations) and an initial point cloud $\mathbf{P} \in \mathbb{R}^{n \times 3}$ sampled near $\mathcal{M}$, where n is the number of points and 3 represents the $\{x, y, z\}$ coordinates of the points. During inference, our method only needs a 3D point cloud $\mathbf{P}$. It outputs a set of binary classification labels $Cl_a \in \{0, 1\}^n$ and a set of displacement vectors $Dis \in \mathbb{R}^{n \times 3}$, simultaneously.

Figure 2 provides an overview of our architecture with four modules: *point-to-feature oracle*, *point descriptor extractor*, *sharp feature detector* and *displacement predictor*.

3.1. Point-to-feature oracle

Most data-driven methods use pre-defined labels for point clouds during the training session. This operation would probably lead to overfitting, or designing complicated loss functions to compute gradients, which makes the network hard to adapt to other datasets or architectures. Instead, we propose a point-to-feature oracle. It updates dynamically the ground truth classification labels and the ground truth displacement vectors according to the current noisy point cloud and the ground truth mesh before forwarding inputs to the network. Our network can deal with, but is not limited to, four types of sharp creases: segments, circles, ellipsoids and B-splines.

Figure 3 shows the principle of our oracle. It takes two user-define parameters, a noise level n and a distance threshold $d_{\max}$, whose influences are well studied in Section 4. Given as input a sampled point cloud $\mathbf{P}$ from a CAD model and its associated ground truth sharp features in analytical expression, we first add a uniform random noise with the user-defined noise level (1% length of the bounding box diagonal by default) on $\mathbf{P}$. We then use the oracle to find, for each point p , its nearest ground truth sharp feature, and project p onto this feature. p is labeled as a sharp feature point if the distance between p and its projected point is smaller than the user-defined distance threshold, which is set to 0.03 times the length of the bounding box diagonal in our experiments.

For segments and (open or closed) circles, we compute the point’s projection onto sharp features by algebra computation. To simplify the computation for (open or closed) ellipsoids and B-splines, we generate a dense point sample on the sharp features and find the nearest one to the query point p as its projection.




Figure 3: Principle of the point-to-feature oracle. (a) Our oracle takes a point cloud and a set of ground truth sharp features as input. (b) We compute the point-to-feature distances and the nearest projections between the point cloud and each primitive, respectively. Yellow denotes $d = 0$ and black denotes $d \geq d_{\max}$ in the color maps. (c) Each point is then mapped to the closest sharp feature. The ones with a point-to-feature distance smaller than $d_{\max}$ are labeled as sharp feature points. (d) The offset vectors are computed for all sharp feature points.

3.2. Point descriptor extractor

The so-called backbone is used to extract pointwise descriptions $\mathbf{f}$ for point clouds. We experimented with DGCNN (Wang et al., 2019) and PCPNet (Guerrero et al., 2018) backbones, but users can plug their own backbone as long as it produces pointwise descriptors.

DGCNN is a graph-based convolutional neural network that takes as input a point cloud of fixed cardinality $\mathbf{N}$. It yields a global descriptor of size 1,024 for the whole point cloud and a multi-scale local descriptor of size 448 for each point. By concatenating global and local descriptors, each point has a descriptor of size 1,472. Since $\mathbf{N}$ is fixed for DGCNN, the point cloud must have the same cardinality during inference, which is its main limitation.

PCPNet is a PointNet-based neural network that takes as input a center (query) point along with its k -nearest neighbors and produces a local descriptor of size 1,024. While the cardinality of point clouds is not restricted, the inference is substantially slower than that of the DGCNN backbone.

3.3. Network architecture

After generating the features, two multi-layer perceptrons take $\mathbf{f}$ as input and output Dis and Cla , respectively. Both $\mathbf{M}_1$ and $\mathbf{M}_2$ comprise three layers. Cla is a set of binary classification labels, indicating whether a point is near a sharp feature. Dis is a set of pointwise vectors that represent the relocation from sharp feature points to their closest sharp features. The displacement vectors for smooth points are ignored by applying masks to the loss function.

3.4. Loss function and training

Our loss function is composed of two parts, a Binary Cross Entropy (BCE) with logits loss for sharp feature classification and a masked Mean Squared Error (MSE) loss for displacement prediction. The two losses cooperate in order to yield both a more accurate classification and a more precise displacement regression.

BCE with logits loss. Given a point cloud with $\mathbf{n}$ points, the ground truth labels $y = \{y_1, \dots, y_{\mathbf{n}}\}$, the predictions $x = \{x_1, \dots, x_{\mathbf{n}}\}$ from $\mathbf{M}_1$ where $x_i \in [-\infty, +\infty]$, and the sigmoid function σ , the loss is computed as:

$$l_c(x, y) = -\frac{1}{\mathbf{n}} \sum_{i=1}^{\mathbf{n}} [y_i \cdot \log \sigma(x_i) + (1 - y_i) \cdot \log(1 - \sigma(x_i))] \quad (1)$$

By combining the Sigmoid layer with BCE with logits loss, the back-propagation is numerically more stable by taking advantage of the log-sum-exp formulation, since the computation converts to:

$$l_c(x_i, y_i) = \begin{cases} (1 - y_i) \cdot x_i + \log(1 + e^{-x_i}) & \text{if } x_i > 0 \\ -x_i \cdot y_i + \log(e^{x_i} + 1) & \text{otherwise} \end{cases} \quad (2)$$

Masked MSE loss. Given ground truth displacement vectors $u = \{u_1, \dots, u_{\mathbf{n}}\}$ and the predicted vectors $v = \{v_1, \dots, v_{\mathbf{n}}\}$ from $\mathbf{M}_2$, the masked MSE is:

$$l_d(y, u, v) = -\frac{1}{\mathbf{n}} \sum_{i=1}^{\mathbf{n}} \mathbb{1}_1(y_i) \cdot \|v_i - u_i\|^2 \quad (3)$$

The mask makes the network concentrate on regressing displacement vectors for points around sharp features.

Training. The total loss energy is defined as:

$$l(x, y, u, v) = w_c \cdot l_c(x, y) + w_d \cdot l_d(y, u, v) \quad (4)$$

where w_c and w_d denote user-defined coefficients to balance the two tasks. An appropriate balance between these two losses is recommended to obtain a good cooperation of two tasks. During training, we first set a large w_c and a small w_d ($w_c = 1$ and $w_d = 0.01$ for both backbones) during the first 10 epochs, because a good classification result is required for learning displacement vectors. We then set a smaller w_c and a larger w_d ($w_c = 0.01$ and $w_d = 100$ for DGCNN-backbone; $w_c = 1$ and $w_d = 20$ for PCPNet-backbone) until the network converges to ensure accurate displacement prediction.

4. Experiments

We implemented our approach using the *Geomdl* library (Bingol and Krishnamurthy, 2019) for B-spline distance computation and the PyTorch deep learning framework (Paszke et al., 2019). We first introduce the setup of our experiments. Then we study the robustness of our approach to the hyper-parameters, followed by an ablation study of each component of our approach to demonstrate their functionalities. We show both the quantitative and the qualitative results, as well as comparisons with unsupervised and supervised methods.

Dataset. The ABC dataset (Koch et al., 2019) is a collection of CAD models with parameterized curves and annotated surfaces. We selected 5,093 models which contain at least one sharp crease. We split them into three non-overlapping sets used for training (3,623 models), evaluating (471 models) and testing (999 models).

Metrics. We selected two lines of quantitative metrics to evaluate our approach: (1) Common metrics used for binary classification and (2) Distance metrics used for regression. Similarly to previous approaches (Himeur et al., 2022; Loizou et al., 2020; Wang et al., 2020), we adopt the following metrics for classification: Precision, Recall, F1-score, Accuracy and Intersection over Union score (IoU). Besides, we estimate two metrics based on the Euclidean distance. More specifically, we compute the distance between detected sharp feature points and continuous ground truth sharp features, by taking advantage of the annotation of the ABC dataset and our point-to-feature oracle, which evaluates the correctness of sharp feature points. We refer to this distance with *PtF* distance. We then estimate the Chamfer distance between the detected sharp feature points and a set of uniformly-sampled points on the ground truth sharp features, which evaluates the completeness of detected features. The definitions of above error metrics are provided in the appendix.

Setups. We experimented with two backbones to validate our approach: PCPNet (Guerrero et al., 2018), an architecture based on PointNet; and DGCNN (Wang et al., 2019), an architecture based on graph neural networks. PCPNet takes a point cloud with arbitrary cardinality, while DGCNN takes a point cloud with a fixed cardinality (set by default to 10k in our experiments). We refer to our model with PCPNet backbone as SFCNet-P and the one with DGCNN backbone as SFCNet-D. The point cloud is normalized inside the bounding box from $[-1, -1, -1]$ to $[1, 1, 1]$ and centered at the origin. By default, we add a random uniform noise with a noise level set to 0.01 during training. To compare different methods, we add a similar noise level in the test set. The default backbone is DGCNN and the default sharp feature threshold is 0.03. Figure 4 shows several results produced by our default model on the test set.

4.1. Robustness study

We evaluated the sensitivity of our approach to the sharp distance threshold, defined as the maximum point-to-feature distance for sharp feature points, and to the noise level of the point clouds.

Distance threshold $d_{\max}$. A large sharp distance threshold helps the framework detect more sharp feature points, but displacement learning becomes more challenging, as shown in Figure 5. We trained four SFCNet-D networks using the same parameters, except for the distance threshold. Table 1 records the measured metrics for $d_{\max} \in \{0.02, 0.03, 0.05, 0.1\}$. We observe that the model with $d_{\max} = 0.1$ achieves good scores for most of the classification criteria while the distance criteria are much worse than $d_{\max} = 0.03$. Figure 5 shows the visualization results on one model. The right one detects more than 5k sharp feature points while we can see from the figure that some of the detected sharp feature points can not be well displaced to the sharp features.




Figure 4: 3D models from the ABC dataset. Our method displaces the detected sharp feature points closer to the ground truth. Top row: ground truth. Middle row: detected raw sharp feature points. Bottom row: displaced sharp feature points. Sharp feature points are depicted in red, and smooth points are depicted in blue. a , d and $\#s$ denote accuracy, point-to-feature distance and number of sharp feature points, respectively.

		Classification(%)					Distance($\times 10^{-2}$)	
		Recall $\uparrow$	Precision $\uparrow$	IoU $\uparrow$	F1 $\uparrow$	Accuracy $\uparrow$	Chamfer $\downarrow$	PtF $\downarrow$
Threshold	(0.01, 0.02)	76.20	66.31	54.74	68.29	93.05	5.43	2.77
	(0.01, 0.03)	79.65	82.68	67.65	78.98	92.65	4.55	2.65
	(0.01, 0.05)	80.58	84.53	69.88	80.77	90.87	5.16	3.26
	(0.01, 0.1)	86.15	89.51	78.14	86.69	89.96	5.84	4.17
Noise	(0, 0.03)	87.27	70.91	63.96	76.37	91.83	4.81	3.02
	(0.01, 0.03)	79.65	82.68	67.65	78.98	92.65	4.55	2.65
	(0.03, 0.03)	71.34	65.50	51.26	65.63	90.48	6.87	3.73
	(0.06, 0.03)	45.32	67.22	36.52	50.71	93.63	11.42	3.81
	(0.06, 0.1)	84.3	83.67	71.63	81.89	85.88	7.82	5.31
Components	D1: DGCNN+CLA	76.43	79.03	60.26	72.45	87.58	7.99	4.40
	D2: DGCNN+ORACLE+CLA	80.62	80.54	67.40	78.38	92.34	6.41	3.61
	D3: DGCNN+CLA+DIS	87.32	74.91	66.44	77.68	89.11	5.72	4.01
	D4: SFCNet-D	79.65	82.68	67.65	78.98	92.65	4.55	2.65
	P1: PCPNet+CLA	96.01	70.82	68.88	79.79	92.15	7.09	5.36
	P2: PCPNet+ORACLE+CLA	88.96	87.20	79.14	86.92	95.54	5.97	3.69
	P3: PCPNet+CLA+DIS	9.24	60.10	8.01	14.14	82.20	51.6	6.04
	P4: SFCNet-P	88.47	88.58	79.86	87.38	95.74	4.31	2.63

Table 1: Quantitative results for robustness study and ablation study. Quantitative results with different distance thresholds, noise levels and components are recorded. Multiple pairs of ($noise$, $threshold$) are used to study the user-defined parameters. We also compared our SFCNet-P model and SFCNet-D model with different combinations of the proposed modules. D and P are abbreviations for DGCNN and PCPNet. 1 to 4 indicate the experiment numbers.




Figure 5: Visual comparison for different distance thresholds. From left to right: distance thresholds set to 0.02, 0.03, 0.05 and 0.1. The one with $d_{\max} = 0.1$ misclassified some smooth points as sharp feature points, while the ones with $d_{\max} \in \{0.03, 0.05\}$ keep a good trade-off between the classification accuracy and the distance criterion.

Noise level n . This parameter has a strong relationship with $d_{\max}$ thus we cannot fix $d_{\max}$ for all experiments. We choose to train a series of SFCNet-D networks with different pairs of $(n, d_{\max})$ and then compare the distance metrics. Table 1 records the quantitative results, which indicates that our method is robust to noise when the distance threshold is larger than the noise level. When the distance threshold is smaller than the noise level, the sharp feature points can not be accurately detected, as shown in the eighth row of Table 1 and in the fourth model of Figure 6 where few sharp feature points are detected.




Figure 6: Visual comparison with different pairs of sharp distance thresholds and noise levels. From left to right, the pairs of sharp distance thresholds and noise levels have the same order as in Table 1. An incompatible choice of $d_{\max}$ and n (the 4th model) leads to a bad result.

4.2. Ablation study

We evaluate the role of the different components by removing one or several components from our approach and comparing the quantitative results on the test set.

Point-to-feature oracle. To show the importance of the dynamic labeling oracle, we consider four pairs of contrast experiments: D1 versus D2, D3 versus D4, P1 versus P2 and P3 versus P4 as shown in Table 1. The results show that the oracle can efficiently improve both the classification and the regression performance. Moreover, the oracle helps the network converge faster and better. In our experiments, networks without oracle are more likely to diverge (see P3) than the ones with oracle.

Displacement learning. The following four pairs of contrast experiments show the improvement by combining displacement learning with classical binary classification learning: D1 versus D3, D2 versus D4, P1 versus P3 and P2 versus P4 as shown in Table 1. Displacement learning boosts not only the distance metrics, but




Figure 7: Visual comparisons with supervised and unsupervised methods on selected models from the test set. According to the closeups, our methods can detect more meaningful sharp feature points that are closed to the sharp features.

also the classification metrics in most cases. However, the comparison between P1 and P3 clearly shows that when the oracle is disabled during the training phase, the displacement learning introduces instability into the framework, since both MLPs share the same backbone. By combining our oracle with displacement learning, our network achieves the best scores in terms of most metrics.

Backbone. We conducted two groups of experiments: D1-D4 for DGCNN backbone and P1-P4 for PCPNet backbone. For both backbones, our proposed approach performs best within the group. Generally speaking, PCPNet backbone has a better performance than DGCNN backbone, while the price is a much longer testing time for a given point cloud.

Sharp feature type. We consider four types of sharp features in our experiments: segments, circles, ellipsoids and B-splines. In order to evaluate the capacity of handling all these types, we first computed the distribution of sharp feature types on the testing set and then compared it with the distributions of detected sharp feature points by both SFCNet-D and SFCNet-P. We observe a similarity among the distributions, while our models find slightly less sharp feature points than the ground truth. We then compute respectively the accuracy of detected sharp feature points and the average PtF distance of consolidation sharp feature points for each type. Our models achieve a good and balanced results on the four types, while the performance of B-splines is slightly worse than the other types due to its high complexity. Note that our models can not directly predict the type of feature points, the type of each predicted sharp feature point is obtained by finding the closest ground truth feature.




Figure 8: The statistics of consolidated sharp feature points by our approach on the test set. Ground truth: computed by our oracle on noisy point clouds. SFCNet-D: D4. SFCNet-P: P4. (a) The distributions of sharp feature points belonging to each sharp feature type. (b) The accuracy of detected sharp feature points for each sharp feature type. (c) The PtF distance of consolidated sharp feature points for each sharp feature type.

4.3. Comparisons

Point-based approaches. We compare our approach with PB-DGCNN (Loizou et al., 2020) (re-trained from scratch on our dataset for 50 epochs, since there is no released pre-trained model), ECNet (Yu et al., 2018a) (using the released pre-trained model), PIE-Net (Wang et al., 2020) (using the released pre-trained model), together with two unsupervised methods: Feature edge estimation via Voronoi Covariance Measure (VCM) provided by the CGAL Library (Mérigot et al., 2010) and Covariance Analysis (CA) (Bazazian et al., 2015). Table 2 records the error metrics on our test set. We observe that both of our networks have smaller Chamfer distances and smaller PtF distances compared to other methods. SFCNet-P performs best on the various metrics. Figure 7 offers a visual comparison of all methods. VCM and CA are sensitive to user-defined parameters: we find it hard to define through trial-and-error parameters that fit all point clouds in our test set. Compared to PB-DGCNN, our SFCNet-D achieves similar scores for classification metrics while performing better in terms of distance, owing to the displacement learning. EC-Net seems to have more missing parts than our models and both distance criteria are worse than the ours. The results show that PIE-Net is not noise-resilient and the detected sharp feature points may distribute over the whole point cloud. While the paper mentions that it also predicts 3D offset vectors to relocate points onto edges, such a prediction step is commented out in the code and no pre-trained model contains such a relocation operation.

	Cla (%)					Dis ($\times 10^{-2}$)	
	Rec	Pre	IoU	F1	Acc	Cha	PtF
CA	67.18	63.06	43.50	57.58	78.98	11.80	8.23
VCM	81.29	38.10	34.09	48.99	72.29	12.01	7.91
PB-DGCNN	85.81	77.90	68.29	79.55	92.20	6.17	3.81
EC-Net	-	-	-	-	-	6.90	4.25
PIE-Net	48.87	45.11	20.32	32.47	62.86	16.65	14.07
SFCNet-D	79.65	82.68	67.65	78.98	92.65	4.55	2.65
SFCNet-P	88.47	88.58	79.86	87.38	95.74	4.31	2.63

Table 2: Quantitative comparison with unsupervised and deep learning sharp feature points detection methods on the ABC dataset. Note that we cannot compute classification metrics of EC-Net since the produced point clouds are upsampled. CA: [Bazazian et al. \(2015\)](#), VCM: [Mérigot et al. \(2010\)](#), PB-DGCNN: [Loizou et al. \(2020\)](#), EC-Net: [Yu et al. \(2018a\)](#), PIE-Net: [Wang et al. \(2020\)](#).

We evaluate the average inference time on our test set. VCM ([Mérigot et al., 2010](#)), implemented in C++, takes around 1.877s for 1 model with 10k points. All the other methods are implemented in Python, and the approximate test times are: 0.134s for CA ([Bazazian et al., 2015](#)), 1.26s for PB-DGCNN ([Loizou et al., 2020](#)), 2.4s for EC-Net ([Yu et al., 2018a](#)), 0.177s for PIE-Net ([Wang et al., 2020](#)), 0.174s for SFCNet-D and 5s for SFCNet-P.

Mesh-based approaches. Feature-preserving mesh denoising approaches can be utilized for sharp feature detection and consolidation from point clouds, combining with an extra surface reconstruction step. We compared our approach with Bilateral Normal Filtering ([Zheng et al., 2010](#)), Mesh Denoising via L0 Optimization ([He and Schaefer, 2013](#)), Guided Mesh Normal Filtering ([Zhang et al., 2015](#)), Edge Sharpener ([Attene et al., 2003](#)) and Non-iterative Feature-preserving Mesh Denoising ([Jones et al., 2003](#)).

Typical mesh-based approaches work on orientable meshes with noisy vertex positions, as shown in the *Fandisk* model in Figure 9. We sampled the noisy mesh to obtain a noisy point cloud. As a consequence, the resulting point cloud contains structural piecewise linear noises which are rarely present in raw point clouds and our approach cannot achieve a good performance on such inputs.

Conversely, mesh-based approaches are not dealing well with noise in raw point clouds and errors originating from the reconstruction step. The *Mechanical part* model belongs to our test set. The raw point cloud contains random noise. To reconstruct an oriented mesh, we estimate and orient normals, then reconstruct the surface using the popular Screened Poisson Reconstruction ([Kazhdan and Hoppe, 2013](#)). Since smoothness priors are introduced during reconstruction, most mesh-based approaches fail to detect connected and clean sharp features, and outliers are inevitable. In addition, not all point clouds are orientable, as shown in the *Art piece* model. We utilize a recent learning-based surface reconstruction approach, referred to as DSE meshing ([Rakotosaona et al., 2021](#)). DSE combines 3D Delaunay triangulations with learned local parametrizations to yield quality meshes, even if it may generate non-manifold edges. The advantage of this approach is that it does not require any normal information. Most mesh-based approaches fail on this model. The Edge Sharpener contributed by ([Attene et al., 2003](#)) successfully orients the noisy mesh while it introduces artifacts around the sharp feature in the intersection of two nearby surface patches with opposite orientation. Since our approach is normal-free, it yields satisfactory outputs on the whole model.

4.4. Results on real scans

We evaluate our SFCNet-P model on real scanned 3D point clouds taken from the online Visionair repository ([Kopecki et al., 2011](#)). These point clouds have never been seen during the training phase. Compared to the training dataset, we highlight three major differences in real scanned point clouds: (1) The cardinality of the point clouds (from 460k to 1M) is substantially larger than that of our training dataset (10k); (2) The intrinsic noise in scanned point clouds originates from the laser, instead of being simulated randomly; (3) The models scanned by real-world scanners are often open, while all models of our training dataset are closed. Despite the above differences, our model predicts accurate sharp feature points from the real scanned point clouds with displacement learning, see Figure 10.




Figure 9: Qualitative comparison with mesh-based approaches. The first two rows: *Fandisk* model. The third and fourth rows: *Mechanical part* model from ABC Dataset. The last two rows: *Art piece* model from Xu et al. (2022) (Bilateral Normal fails on this model). Input: point cloud for SFCNet-P (odd rows), mesh for the other approaches (even rows). The sharp features are extracted from mesh by comparing the maximum angle between the normal vectors of adjacent triangles using The CGAL Project (2023).




Figure 10: Scanned point clouds of Visionair repository. The first and the third image show the detected sharp feature points without displacement. The second and the forth ones show the predicted sharp feature points with displacement. $\#i$ and $\#s$ denote the number of input points and the number of detected sharp feature points, respectively.

Due to the lack of ground truth labels, it is hard to obtain quantitative results on real-world data. We compared our proposed approach with other approaches qualitatively on a scanned point cloud.




Figure 11: Comparisons on a scanned point cloud (718k points). SFCNet-P(wo disp) denotes our SFCNet-P model without adding displacement vectors. The output of VCM is sensitive to the quality of estimated normals, since the ground truth normals are not available for this point cloud. Inputs for PB-DGCNN and PIE-Net are subsampled due to the restriction for their input.

4.5. Limitations

Our current approach presents several limitations. First, the detected sharp feature points may not distribute uniformly over the sharp features. We intend to devise a loss function that favors the uniform distribution of inliers along sharp creases. However, the difficulty lies in the fact that we have no knowledge of continuous sharp feature graphs. Second, we classify points into two classes: sharp or smooth, instead of classifying them into instances. In future work, we wish to extend our framework to perform instance segmentation of sharp features, in order to cluster sharp feature points into several creases, possibly meeting at corners, cusps or darts. To achieve this goal, the network must predict the number of sharp features, which is substantially more involved than binary classification.

5. Applications

We show two potential applications of our approach: 3D feature line extraction and 3D surface reconstruction. Nevertheless, our approach is not limited to these two applications. It can be further adapted to other applications such as instance segmentation or feature graph extraction.

5.1. Feature line extraction

After consolidating the sharp feature points, we can apply primitive detection methods to extract parametric representations of sharp feature creases. We adapt the RANSAC (Fischler and Bolles, 1981) algorithm to extract feature lines and the results are shown in Figure 12. The parameters are carefully tuned for all experiments. We observe that the extractions of consolidated sharp feature points are more robust and have a better quality than the one of the detected sharp feature points without displacement.




Figure 12: Feature line extraction results using RANSAC. From top to bottom: results of dense, uniformly sampled points on ground truth feature creases, results of the detected sharp feature points without displacement, results of the consolidated sharp feature points by our SFCNet-P model. $\#pr$ denotes the number of detected primitives. The ones without displacement detect many wrong primitives and there are more outliers (black points) left after the detection.

5.2. Surface reconstruction

Some surface reconstruction methods can benefit from consolidated point clouds. Again we utilize DSE meshing (Rakotosaona et al., 2021) to reconstruct meshes from unoriented point clouds. We show the quality improvement of our approach by comparing the surface reconstruction results from the input 3D point clouds with the consolidated 3D point clouds. For generating a consolidated 3D point cloud, we first combine the input point cloud with the displaced sharp feature points, then 10k points are sampled using farthest point sampling to meet the input size requirement of the DSE.

Figure 13 depicts the reconstruction results before and after consolidation using our SFCNet-P model, in which we selected three 3D point clouds without noise from the test set in order to better compare the reconstruction quality near sharp creases.

6. Conclusion

We contributed a novel deep learning based-framework devised to detect and consolidate sharp feature points from raw 3D point clouds. Compatible with existing backbones devised to extract features from 3D point clouds, our framework comprises two learnable modules: The first module learns to predict binary smooth/sharp labels for all points, and the second module learns to regress the displacement vectors used for relocating sharp feature points onto sharp features. Our framework is capable of identifying and consolidating sharp features, altogether. Our approach is robust to intrinsic noise thanks to the point-to-feature oracle that performs data augmentation during the training phase, and displacement learning that relocates sharp feature points onto their nearest sharp features. Our experiments demonstrate that our framework outperforms the state-of-the-art in terms of detection accuracy and distance criteria.




Figure 13: Reconstructed meshes using DSE meshing. First row: reconstruction from the consolidated 3D point clouds. Second row: reconstruction from the 3D input point clouds. c denotes the Chamfer distance between the reconstructed mesh and the ground truth mesh in the ABC dataset. DSE meshing is able to capture more details on sharp creases with consolidated point clouds, and the Chamfer distance shows that the meshes are closer to the ground truth compared to the original point clouds.

7. Acknowledgments

This work was supported by the French government, through the 3IA Côte d’Azur Investments in the Future project managed by the French National Research Agency (ANR) with the reference number ANR-19-P3IA-0002.

References

- Amenta, N., Bern, M., Kamvysselis, M., 1998. A new Voronoi-based surface reconstruction algorithm, in: Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, pp. 415–421.
- Attene, M., Falcidieno, B., Spagnuolo, M., Rossignac, J., 2003. Edge-Sharpener: Recovering sharp features in triangulations of non-adaptively re-meshed surfaces, in: Eurographics Symposium on Geometry Processing (SGP), Eurographics Association. pp. 62–69. URL: <https://doi.org/10.2312/SGP/SGP03/062-071>, doi:10.2312/SGP/SGP03/062-071.
- Bazazian, D., Casas, J.R., Ruiz-Hidalgo, J., 2015. Fast and robust edge extraction in unorganized point clouds, in: International Conference on Digital Image Computing: Techniques and Applications (DICTA), IEEE. pp. 1–8.
- Bingol, O.R., Krishnamurthy, A., 2019. NURBS-Python: An open-source object-oriented NURBS modeling framework in Python. SoftwareX 9, 85–94. doi:<https://doi.org/10.1016/j.softx.2018.12.005>.
- Chen, H., Huang, Y., Xie, Q., Liu, Y., Zhang, Y., Wei, M., Wang, J., 2021. Multiscale feature line extraction from raw point clouds based on local surface variation and anisotropic contraction. IEEE Transactions on Automation Science and Engineering (T-ASE) 19, 1003–1016.
- Daniels, J.I., Ha, L.K., Ochotta, T., Silva, C.T., 2007. Robust smooth feature extraction from point clouds, in: IEEE International Conference on Shape Modeling and Applications (SMI), IEEE. pp. 123–136.
- Fischler, M.A., Bolles, R.C., 1981. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. Communications of the ACM 24, 381–395.
- Fleishman, S., Cohen-Or, D., Silva, C.T., 2005. Robust moving least-squares fitting with sharp features. ACM Transactions on Graphics (TOG) 24, 544–552.
- Fleishman, S., Drori, I., Cohen-Or, D., 2003. Bilateral mesh denoising. ACM Transactions on Graphics (TOG) 22, 950–953. URL: <https://doi.org/10.1145/882262.882368>, doi:10.1145/882262.882368.
- Gal, R., Shamir, A., Hassner, T., Pauly, M., Cohen-Or, D., 2007. Surface reconstruction using local shape priors, in: Eurographics Symposium on Geometry Processing (SGP), pp. 253–262.
- Gong, J., Xu, J., Tan, X., Zhou, J., Qu, Y., Xie, Y., Ma, L., 2021. Boundary-aware geometric encoding for semantic segmentation of point clouds, in: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 1424–1432.

- Guennebaud, G., Gross, M.H., 2007. Algebraic point set surfaces. *ACM Transactions on Graphics (TOG)* 26, 23. URL: <https://doi.org/10.1145/1276377.1276406>, doi:10.1145/1276377.1276406.
- Guerrero, P., Kleiman, Y., Ovsjanikov, M., Mitra, N.J., 2018. PCPNet learning local shape properties from raw point clouds. *Computer Graphics Forum (CGF)* 37, 75–85.
- Guo, B., Zhang, Y., Gao, J., Li, C., Hu, Y., 2022. SGLBP: Subgraph-based local binary patterns for feature extraction on point clouds. *Computer Graphics Forum (CGF)* 41, 51–66. URL: <https://doi.org/10.1111/cgf.14500>, doi:10.1111/cgf.14500.
- He, L., Schaefer, S., 2013. Mesh denoising via L_0 minimization. *ACM Transactions on Graphics (TOG)* 32, 64:1–64:8. URL: <https://doi.org/10.1145/2461912.2461965>, doi:10.1145/2461912.2461965.
- Himeur, C., Lejembre, T., Pellegrini, T., Paulin, M., Barthe, L., Mellado, N., 2022. PCEDNet: A lightweight neural network for fast and interactive edge detection in 3D point clouds. *ACM Transactions on Graphics (TOG)* 41, 10:1–10:21. URL: <https://doi.org/10.1145/3481804>, doi:10.1145/3481804.
- Huang, H., Wu, S., Gong, M., Cohen-Or, D., Ascher, U., Zhang, H., 2013. Edge-aware point set resampling. *ACM Transactions on Graphics (TOG)* 32, 1–12.
- Jones, T.R., Durand, F., Desbrun, M., 2003. Non-iterative, feature-preserving mesh smoothing. *ACM Transactions on Graphics (TOG)* 22, 943–949. URL: <https://doi.org/10.1145/882262.882367>, doi:10.1145/882262.882367.
- Kazhdan, M., Hoppe, H., 2013. Screened Poisson surface reconstruction. *ACM Transactions on Graphics (ToG)* 32, 1–13.
- Koch, S., Matveev, A., Jiang, Z., Williams, F., Artemov, A., Burnaev, E., Alexa, M., Zorin, D., Panozzo, D., 2019. ABC: A big CAD model dataset for geometric deep learning, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9601–9611.
- Kopecki, A., Wössner, U., Mavrikios, D., Rentzos, L., Weidig, C., Roucoules, L., Ntouni, O.D., Reed, M., Dumont, G., Büdingers, D., Milecki, A., Baranyi, P., Noël, F., Masclet, C., Attene, M., Giannini, F., Spagnuolo, M., 2011. VISIONAIR: Vision advanced infrastructure for research. *SBC Journal on Interactive Systems* 2, 40–43. URL: <https://hal.inria.fr/hal-00644745>.
- Lee, K.W., Bo, P., 2016. Feature curve extraction from point clouds via developable strip intersection. *Journal of Computational Design and Engineering* 3, 102–111.
- Levin, D., 2004. Mesh-independent surface interpolation, in: *Geometric Modeling for Scientific Visualization*. Springer, pp. 37–49.
- Li, J., Chen, B.M., Lee, G.H., 2018. SO-Net: Self-organizing network for point cloud analysis, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9397–9406.
- Loizou, M., Averkiou, M., Kalogerakis, E., 2020. Learning part boundaries from 3D point clouds. *Computer Graphics Forum (CGF)* 39, 183–195.
- Matveev, A., Rakhimov, R., Artemov, A., Bobrovskikh, G., Egiastian, V., Bogomolov, E., Panozzo, D., Zorin, D., Burnaev, E., 2022. DEF: Deep estimation of sharp geometric features in 3D shapes. *ACM Transactions on Graphics (TOG)* 41, 1–22.
- Mérogot, Q., Ovsjanikov, M., Guibas, L.J., 2010. Voronoi-based curvature and feature estimation from point clouds. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 17, 743–756.
- Ni, H., Lin, X., Ning, X., Zhang, J., 2016. Edge detection and feature line tracing in 3D-point clouds by analyzing geometric properties of neighborhoods. *Remote Sensing* 8, 710.
- Öztireli, A.C., Guennebaud, G., Gross, M., 2009. Feature preserving point set surfaces based on non-linear kernel regression. *Computer Graphics Forum (CGF)* 28, 493–501.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al., 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems (NeurIPS)* 32.
- Pistilli, F., Fracastoro, G., Valsesia, D., Magli, E., 2020. Learning graph-convolutional representations for point cloud denoising, in: *European Conference on Computer Vision (ECCV)*, Springer. pp. 103–118.
- Qi, C.R., Su, H., Mo, K., Guibas, L.J., 2017a. PointNet: Deep learning on point sets for 3D classification and segmentation, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 652–660.
- Qi, C.R., Yi, L., Su, H., Guibas, L.J., 2017b. PointNet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in Neural Information Processing Systems (NeurIPS)* 30.
- Qian, G., Hammoud, H., Li, G., Thabet, A., Ghanem, B., 2021. ASSANet: An anisotropic separable set abstraction for efficient point cloud representation learning. *Advances in Neural Information Processing Systems (NeurIPS)* 34.
- Rakotosaona, M.J., Guerrero, P., Aigerman, N., Mitra, N.J., Ovsjanikov, M., 2021. Learning Delaunay surface elements for mesh reconstruction, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 22–31.
- Rakotosaona, M.J., La Barbera, V., Guerrero, P., Mitra, N.J., Ovsjanikov, M., 2020. PointCleanNet: Learning to denoise and remove outliers from dense point clouds. *Computer Graphics Forum (CGF)* 39, 185–203.
- Salman, N., Yvinec, M., Mérogot, Q., 2010. Feature preserving mesh generation from 3D point clouds. *Computer Graphics Forum (CGF)* 29, 1623–1632.
- Su, H., Jampani, V., Sun, D., Maji, S., Kalogerakis, E., Yang, M.H., Kautz, J., 2018. SPLATNet: Sparse lattice networks for point cloud processing, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2530–2539.
- The CGAL Project, 2023. CGAL user and reference manual. 5.5.2 ed., CGAL Editorial Board. URL: <https://doc.cgal.org/5.5.2/Manual/packages.html>.
- Wang, C.C.L., 2006a. Bilateral recovering of sharp edges on feature-insensitive sampled meshes. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 12, 629–639. URL: <https://doi.org/10.1109/TVCG.2006.60>, doi:10.1109/TVCG.2006.60.

- Wang, C.C.L., 2006b. Incremental reconstruction of sharp edges on mesh surfaces. *Computer-Aided Design* 38, 689–702. URL: <https://doi.org/10.1016/j.cad.2006.02.009>, doi:10.1016/j.cad.2006.02.009.
- Wang, R., Yang, Z., Liu, L., Deng, J., Chen, F., 2014. Decoupling noise and features via weighted l1-analysis compressed sensing. *ACM Transactions on Graphics (TOG)* 33, 1–12.
- Wang, X., Xu, Y., Xu, K., Tagliasacchi, A., Zhou, B., Mahdavi-Amiri, A., Zhang, H., 2020. PIE-NET: Parametric inference of point cloud edges. *Advances in Neural Information Processing Systems (NeurIPS)* 33, 20167–20178.
- Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M., 2019. Dynamic graph CNN for learning on point clouds. *ACM Transactions on Graphics (TOG)* 38, 1–12.
- Weber, C., Hahmann, S., Hagen, H., 2010. Sharp feature detection in point clouds, in: *Shape Modeling International Conference (SMI)*, IEEE. pp. 175–186.
- Xu, R., Wang, Z., Dou, Z., Zong, C., Xin, S., Jiang, M., Ju, T., Tu, C., 2022. RFEPS: Reconstructing feature-line equipped polygonal surface. *ACM Transactions on Graphics (TOG)* 41, 1–15.
- Yu, L., Li, X., Fu, C.W., Cohen-Or, D., Heng, P.A., 2018a. EC-Net: An edge-aware point set consolidation network, in: *Proceedings of the European conference on computer vision (ECCV)*, pp. 386–402.
- Yu, L., Li, X., Fu, C.W., Cohen-Or, D., Heng, P.A., 2018b. PU-Net: Point cloud upsampling network, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2790–2799.
- Zhang, W., Deng, B., Zhang, J., Bouaziz, S., Liu, L., 2015. Guided mesh normal filtering. *Computer Graphics Forum (CGF)* 34, 23–34. URL: <https://doi.org/10.1111/cgf.12742>, doi:10.1111/cgf.12742.
- Zhao, W., Liu, X., Zhao, Y., Fan, X., Zhao, D., 2019. NormalNet: Learning based guided normal filtering for mesh denoising. CoRR abs/1903.04015. URL: <http://arxiv.org/abs/1903.04015>, arXiv:1903.04015.
- Zheng, Y., Fu, H., Au, O.K.C., Tai, C.L., 2010. Bilateral normal filtering for mesh denoising. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 17, 1521–1530.

Appendix A: Definition of the error metrics

According to the ground truth binary label $y \in \{0, 1\}$ and the predicted label $\hat{y} \in \{0, 1\}$, we divide the predictions into 4 classes:

- *TP*: True Positive where $y = \hat{y} = 1$,
- *TN*: True Negative where $y = \hat{y} = 0$,
- *FP*: False Positive where $y = 0, \hat{y} = 1$,
- *FN*: False Negative where $y = 1, \hat{y} = 0$.

Recall. indicates how many sharp feature points are correctly classified among all ground truth sharp feature points. It is defined as:

$$Recall = \frac{TP}{TP + FN}$$

Precision. indicates how many points are correctly classified as sharp feature points among all predicted sharp feature points. It is defined as:

$$Precision = \frac{TP}{FP + TP}$$

Accuracy. indicates how many points are correctly classified among all predictions. It is defined as:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

Intersection over Union score (IoU). quantifies the percent overlap between the groundtruth labels and our predictions.

$$IoU = \frac{TP}{TP + FP + FN}$$

F1-score. seeks a balance between *Precision* and *Recall*. It is computed as the harmonic mean of the two metrics:

$$F1\text{-score} = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}$$

Chamfer distance. We compute the Chamfer distance between a set of predicted sharp feature points $\hat{\mathcal{X}}$ and a dense, uniformly sampled point cloud $\mathcal{X}$ on the ground truth sharp creases. The Chamfer distance measures not only the fitting quality of the prediction but also its completeness. It is computed as:

$$CD(\mathcal{X}, \hat{\mathcal{X}}) = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} \min_{y \in \hat{\mathcal{X}}} \|x - y\| + \frac{1}{|\hat{\mathcal{X}}|} \sum_{y \in \hat{\mathcal{X}}} \min_{x \in \mathcal{X}} \|x - y\|$$

Point-to-feature (PtF) distance. Given the annotations of the ABC dataset, we extract the parametric representations L for all sharp creases. The *PtF* distance is then computed as:

$$PtF(\hat{\mathcal{X}}, L) = \frac{1}{|\hat{\mathcal{X}}|} \sum_{x \in \hat{\mathcal{X}}} \min_{l \in L} d(x, l)$$

where $d(x, l)$ is the Euclidean distance between a point x and its projection point onto the primitive l .