



MSECNet: Accurate and Robust Normal Estimation for 3D Point Clouds by Multi-Scale Edge Conditioning

Haoyi Xiu
National Institute of Advanced
Industrial Science and Technology
(AIST)
Tokyo, Japan
hiroki-shuu@aist.go.jp

Xin Liu*
National Institute of Advanced
Industrial Science and Technology
(AIST)
Tokyo, Japan
xin.liu@aist.go.jp

Weimin Wang*[†]
National Institute of Advanced
Industrial Science and Technology
(AIST)
Tokyo, Japan
weimin.wang@aist.go.jp

Kyoung-Sook Kim
National Institute of Advanced
Industrial Science and Technology
(AIST)
Tokyo, Japan
ks.kim@aist.go.jp

Masashi Matsuoka
Tokyo Institute of Technology
Tokyo, Japan
matsuoka.m.ab@m.titech.ac.jp

ABSTRACT

Estimating surface normals from 3D point clouds is critical for various applications, including surface reconstruction and rendering. While existing methods for normal estimation perform well in regions where normals change slowly, they tend to fail where normals vary rapidly. To address this issue, we propose a novel approach called MSECNet, which improves estimation in normal varying regions by treating normal variation modeling as an edge detection problem. MSECNet consists of a backbone network and a multi-scale edge conditioning (MSEC) stream. The MSEC stream achieves robust edge detection through multi-scale feature fusion and adaptive edge detection. The detected edges are then combined with the output of the backbone network using the edge conditioning module to produce edge-aware representations. Extensive experiments show that MSECNet outperforms existing methods on both synthetic (PCPNet) and real-world (SceneNN) datasets while running significantly faster. We also conduct various analyses to investigate the contribution of each component in the MSEC stream. Finally, we demonstrate the effectiveness of our approach in surface reconstruction.

CCS CONCEPTS

• **Computing methodologies** → **Shape representations; Scene understanding.**

*Corresponding authors

[†]Also is affiliated with Dalian University of Technology

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '23, October 29–November 3, 2023, Ottawa, ON, Canada

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0108-5/23/10...\$15.00

<https://doi.org/10.1145/3581783.3613762>

KEYWORDS

normal estimation, edge detection, point cloud, multi-scale fusion

ACM Reference Format:

Haoyi Xiu, Xin Liu, Weimin Wang, Kyoung-Sook Kim, and Masashi Matsuoka. 2023. MSECNet: Accurate and Robust Normal Estimation for 3D Point Clouds by Multi-Scale Edge Conditioning. In *Proceedings of the 31st ACM International Conference on Multimedia (MM '23)*, October 29–November 3, 2023, Ottawa, ON, Canada. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3581783.3613762>

1 INTRODUCTION

3D point clouds can be used for many applications such as surface reconstruction [21], rendering [13], and virtual reality [12]. An essential step for many of them is to estimate normals at each point, which is to infer the local surface orientation from point samples. Therefore, the quality of normals estimated has a significant impact on the performance of those applications.

A classical approach for normal estimation is to fit a geometric primitive such as a plane [19] or surface [9] to the local neighborhood/patch of each point. However, a key challenge of these methods is that various parameters (e.g., the patch size) must be set manually without knowing the noise level and feature distribution, and their performance significantly depends on the choice of those parameters. Recently, data-driven approaches based on deep learning have been proposed to address this issue. They take a patch as input and estimate the normal of the patch center using deep neural networks (DNNs). There are two main types of approaches. The first one directly regresses the normal vector from features extracted by DNNs [15]. The other one uses DNNs to predict per-point weights for a weighted least square fitting and obtains the normal vector by solving it [24].

A recurrent issue of normal estimation is that the estimation in regions where normals change rapidly is not reliable. The problem is made even more challenging in the presence of data corruptions such as noise and density variations. While some studies attempt to address this issue with hand-crafted estimators [48], metric learning [45], or graph convolutions [25], they often require manual




Figure 1: MSECNet performs accurate and robust normal estimation while running significantly faster than recent works. The MSEC stream improves estimation at sharp edges with multi-scale fusion, edge detection, and edge conditioning.

preprocessing or still rely on surface fitting at the end, which may smooth out details and requires careful selection of the fitting order.

To address this issue, in this work, we propose Multi-Scale Edge Conditioning Network (MSECNet) for accurate and robust surface normal estimation¹. The overview of our approach is presented in Fig. 1. MSECNet achieves accurate estimation at normal varying regions by the explicit modeling of normal variations. We treat the modeling of normal variations as the modeling of geometric edges. An edge, in essence, describes a variation in certain properties. A smooth change represents a weak edge while a drastic change (e.g., a corner) indicates a strong edge. Therefore, explicit modeling of geometric edges can lead to better modeling of normal variations, which in turn leads to more accurate estimation.

To model geometric edges, we propose the MSEC stream which consists of multi-scale fusion, edge detection, and edge conditioning (see Fig. 2). The fusion step performs multi-scale feature fusion by reusing the multi-scale features of a backbone network. The fusion is necessary for successful edge detection as edges are multi-scale structures and the detection needs to be robust against various data corruption. Unlike previous works (e.g., [51] and [26]) whose fusion aggregates information from larger scales by pooling, we distribute information from larger scales to the smallest one by interpolation. This is because we aim to detect edges that are spatially localized and aggregation destroys such information. The fused features are then used for adaptive edge detection in which the adaptability allows for the enhancement of true edges. Lastly, the edge conditioning step merges the detected edges with the representation learned by the backbone to produce edge-aware features.

Extensive experiments demonstrate that MSECNet achieves state-of-the-art performance on the synthetic PCPNet dataset [15]. To investigate its generalizability, we apply pretrained MSECNet to the real-world SceneNN dataset [20]. The results indicate that MSECNet outperforms all baselines by a large margin. Furthermore, we

verify that MSECNet runs $\sim 61\times$ faster than the previous works during inference. Moreover, we perform various analyses to show how each step of the MSEC stream helps. Finally, we demonstrate the utility of MSECNet in surface reconstruction.

Our main contributions are summarized as follows:

- We propose MSECNet, a new normal estimation approach that models normal variations by explicit edge modeling. It achieves state-of-the-art performance on both synthetic and real-world datasets while showing strong robustness against different noise levels and density variations. Also, it runs significantly faster than cutting-edge networks.
- We propose the MSEC stream that improves normal estimation at normal varying regions by the combination of multi-scale fusion, adaptive edge detection, and edge conditioning. We verify that the MSEC stream significantly advances the performance of the backbone network.

2 RELATED WORK

2.1 Traditional Normal Estimation

The classical methods for normal estimation are based on Principle Component Analysis (PCA) [2, 10, 19]. They perform a local co-variance analysis on a local patch and use the direction of minimal variance as the normal. On the other hand, methods for surface reconstruction such as moving least squares [1, 23], osculating jets [9], and spherical fitting [14] are also used for estimating normals. While their behavior is theoretically well-understood, they require careful selection of parameters such as the neighborhood size [33] and often need special treatment when data are corrupted.

2.2 Normal Estimation with DNNs

One of the earliest approaches first projects a Hough accumulator into an image and then the normal is estimated by a convolutional neural network (CNN) [6]. Similarly, PointProNets [40] and RefineNet [48] convert a local point patch to a 2D height map that can be fed into CNN. However, they tend to lose fine geometric details due to the conversion. To address this issue, some methods [15, 41] leverage PointNet [36], which directly takes point clouds as input, and aggregates points into a single vector from which the normal is regressed. This type of approach is further improved in accuracy and efficiency by introducing multi-scale structures [15], a point-voxel architecture [16], and local plane constraint [49].

On the other hand, some studies leverage intrinsic problem structures to improve data efficiency. For instance, Lenssen et al. [24] use a graph neural network to produce per-point weights and then use them for a weighted least-square plane fitting. Similarly, DeepFit [4] adopts PointNet to predict weights for jet fitting. Both of them try to improve the robustness of estimation and data efficiency with the fitting techniques. However, the fitting tends to smooth the surfaces. Besides, the polynomial order of the jet fitting is difficult to determine because the complexity of the surfaces varies greatly. AdaFit [51] addresses this issue by predicting an additional offset to adjust point distributions. Zhang et al. [47] learns the optimal set of points from which the normal vector is regressed. Zhou et al. [50] simplifies the surface fitting by focusing on critical points found by top-K selection. HSurf-Net [26] proposes a hypersurface fitting that frees the user from choosing the polynomial order. NeAF [27],

¹Code and pretrained models is available in <https://github.com/martianxiu/MSECNet>.

on the other hand, learns the angle field by predicting the offset between the random input vector and the ground truth normal. However, they tend to be slow in inference because they estimate only a single vector from a patch.

A recurrent issue of existing approaches is that they produce unreliable estimations at sharp edges. Hence, a line of research focuses on addressing this problem. Some approaches aim to accurately predict normal vectors in areas where the normal varies by separately modeling sharp and non-sharp patches [30, 45], but these methods require manual preprocessing. Other approaches use graph convolutions to learn more about local geometry [24, 25], but these techniques have limited access to multiple scales and may be less robust to data corruption. Additionally, many of these methods still rely on surface fitting, which can smooth out details. MSECNet differs from these methods in that it does not require any preprocessing or surface fitting. Furthermore, it can directly access features at all scales, allowing for robust and accurate estimation.

2.3 Edge Detection

Computational edge detection is an extremely well-studied topic. Pioneering works such as Sobel detector [22], zero-crossing [31], and Canny method [7] detect edges using gradients. However, these methods usually cannot cope with the complexity of real-world applications. To address this issue, learning-based edge detection methods (e.g., [3, 11, 32]) are proposed. Nevertheless, their performance depends on manual processing and expert knowledge. Recently, edge detectors based on DNNs have shown impressive performance. For instance, HED [46] performs edge detection using a cascaded architecture in which one layer produces a scale-specific edge prediction. Edge maps of different scales are then combined to produce accurate edge predictions. Building on top of HED, RCF [28] makes full use of multi-scale features by using all convolutional layers for prediction. BDCN [17] further enhances the prediction with scale-specific supervision. PiDiNet [42] integrates traditional operators to convolutions for efficient detection. RIND-Net [34] models different types of edges with dedicated decoders. EDTER [35] exploits the full image context with transformers.

A key observation is that multi-scale features are essential for successful edge detection because edges are multi-scale structures [31]. Regarding multi-scale fusion strategies for normal estimation, some studies use multiple patches with different patch sizes as input and build separate networks for each scale [5, 15]. However, creating separate networks for each scale is costly. Another approach is to aggregate large-scale information through max-pooling and then copy the resulting vector for all points [25, 26, 51]. Our approach differs from theirs in that we distribute information instead of aggregating it, since we aim to detect spatially localized edges. Also, their fusion is limited to adjacent scales, while ours has direct access to all scales, leading to a more accurate and robust representation.

3 MULTI-SCALE EDGE CONDITIONING

Let $\mathbf{P} \in \mathbb{R}^{N \times 3}$ be a patch consisting of N points and $\mathbf{p}_i = (x_i, y_i, z_i) \in \mathbb{R}^3$ denote the coordinates of point i . A point i can be associated with a feature $\mathbf{f}_i \in \mathbb{R}^d$. Our task is to estimate the unoriented normals $\mathbf{n} \in \mathbb{R}^{N \times 3}$ of this patch from $\mathbf{P}$. The overall architecture is

illustrated in Fig. 2. Given a patch of point clouds, a backbone extracts multi-scale features for each point. Then, features from each stage are fed into the MSEC stream. The motivation for introducing the MSEC stream is explained in Sec. 3.1. The first step of the stream is feature fusion (Sec. 3.2) which merges features from different scales. The resulting multi-scale features are subsequently used for adaptive edge detection (Sec. 3.3). Then, the detected edges are used to condition the output of the backbone to produce edge-aware representation (Sec. 3.4). Finally, normals are regressed from the resulting features. Other details are elaborated in Sec. 3.5.

Unlike previous work that estimates only one normal for a patch, we estimate normals for all points of the patch. This allows for a significant speedup in inference, which is described in Sec. 3.5.

3.1 Motivation

Successful edge modeling must rely on multi-scale analysis for two reasons. First, edges are naturally multi-scale [31], and thus combining multiple scales improves edge detection [38]. Second, single-scale edge detection is hardly effective in the presence of data corruption such as noise and density variations. Combining information from large and small scales can help to estimate normals accurately while being robust against data corruption.

To obtain features of different scales, it is natural to use hierarchical DNNs (e.g., PointNet++ [37]), which perform efficient multi-scale learning through the use of local aggregation and subsampling. One might think that the output of such networks could be used for edge detection. However, from an edge detection perspective, the problem with these networks is that each layer only has direct access to adjacent scales. Direct access to non-adjacent scales is hampered by intermediate processing such as transformations and subsampling, which can severely compromise the rich information contained in the original features. Therefore, we construct a new stream so that the edge detector can directly access all scales.

3.2 Multi-Scale Fusion

Effectively combining features of different scales is critical for edge detection. In this work, a feature from one stage of the backbone network represents a single-scale feature that characterizes a particular scale based on receptive field size and density.

A common multi-scale fusion strategy for normal estimation aggregates information from larger scales using max-pooling to extract more information about the underlying geometry [25, 26, 51]. However, since edges are spatial structures and aggregation destroys spatial relationships, such fusion methods are not suitable for our purpose. Therefore, instead of aggregating from larger scales, we distribute information from larger scales to preserve spatial structures. As shown in Fig. 2, an output from a stage undergoes a 3-neighbor interpolation before entering the MSEC stream. The interpolation ensures that the spatial density of the feature is the same as the input, since we want to detect edges from the raw point cloud. Unlike the popular gradual interpolation strategy (e.g., PointNet++ [37]), our strategy allows a jump across the hierarchy. This jump results in “blurred” point clouds that efficiently describe multi-scale structures. As a result, each point has direct access to rich features learned at different scales.




Figure 2: The architecture of MSECNet. The MSEC stream consists of multi-scale fusion, edge detection, and edge conditioning steps. LBR stands for a sequential application of a Linear transformation, a batch normalization, and a ReLU.

Let $\mathbf{f}_i^s$ denote the feature of point i whose scale is $s \in \{1, 2, \dots, S\}$. Given S features from different scales, the initial multi-scale feature is produced by merging all features of different scales and embedding them into a single vector:

$$\mathbf{f}_i^{\text{fused}} = \phi \left(\left[\mathbf{f}_i^1, \mathbf{f}_i^2, \dots, \mathbf{f}_i^S \right] \right), \quad (1)$$

where $\phi : \mathbb{R}^{d_1+d_2+\dots+d_S} \rightarrow \mathbb{R}^{d_{\text{fused}}}$ denotes an LBR and $[\cdot]$ denotes the concatenation. ϕ is shared among points and nonlinearly maps features of different scales to the fixed d_{fused} -dimensional space. Hence, the resulting features have the accuracy of small scales and robustness of large scales. However, a point-wise fusion is not sufficient for edge detection due to various factors: 1) small-scale features are too sensitive to noise; 2) large-scale features have poor localization; and 3) interpolation may produce biased multi-scale features due to biases in sampling density. A key observation is that the above issues can be addressed by spatial filtering. Therefore, we introduce space transformation that spatially filters the fused feature. Let $\mathcal{N}(i)$ denote the local neighborhood of point i . The space transformation is defined as:

$$\mathbf{f}_i^{\text{sp}} = \max_{j \in \mathcal{N}(i)} \alpha \left(\left[\mathbf{p}_j - \mathbf{p}_i, \mathbf{f}_j^{\text{fused}} \right] \right), \quad (2)$$

where $\alpha : \mathbb{R}^{3+d_{\text{fused}}} \rightarrow \mathbb{R}^{d_{\text{fused}}}$ is an MLP consisting of point-wise linear transformations, batch normalizations, and ReLUs. $\max$ denotes a max-pooling operation within the neighborhood of point i . The relative positions are concatenated to $\mathbf{f}_i^{\text{fused}}$ to better describe the 3D geometry of points. Consequently, the space transformation strives to address the above issues by sharpening important details on the one hand and smoothing harmful signals on the other. Furthermore, we introduce the channel transformation that performs intra-channel filtering:

$$\mathbf{f}_i^{\text{spch}} = \beta \left(\mathbf{f}_i^{\text{sp}} \right), \quad (3)$$

where $\mathbf{f}_i^{\text{spch}}$ indicates features that undergo spatial and channel transformations and $\beta : \mathbb{R}^{d_{\text{fused}}} \rightarrow \mathbb{R}^{d_{\text{fused}}}$ represents an MLP, similar to α . With more refined and noise-free features as input, the channel transformation aims to produce more discriminative

features by nonlinear mapping, in the same sense as a position-wise feed-forward network in the transformer [44].

Successive transformations inevitably increase the difficulty of optimization. Transformations may also result in the loss of useful information contained in the original features. Therefore, we adopt a residual structure [18] to ease the optimization while reducing information loss. Specifically, the output of the fusion is:

$$\mathbf{f}_i^{\text{ms}} = \mathbf{f}_i^{\text{fused}} + \mathbf{f}_i^{\text{spch}}, \quad (4)$$

where $\mathbf{f}_i^{\text{ms}}$ denotes the resulting multi-scale features.

3.3 Adaptive Edge Detection

Given a multi-scale feature, the detection step transforms the features so that the output describes the “edginess” of the point. For our purposes, the desired detector should satisfy two requirements. First, the operator should be computationally efficient, since it must be computed for many training iterations. Second, the detector should be permutation invariant because the point clouds can come in any order. To this end, we advocate the use of the discrete Laplacian (or the umbrella operator [43]). Unlike sophisticated edge detectors (e.g., [11]), the Laplacian is a linear operator and can be computed by simply taking the differences between neighbors. In addition, the Laplacian is permutation invariant because it involves only pairwise differences and a symmetric function (i.e., averaging).

However, there are situations where the Laplacian may not be sufficient or even problematic. For instance, Laplacian is known to be sensitive to noise and may amplify noisy signals, leading to less reliable edge detection. On the other hand, there might be weak edges that could be smoothed out easily, which is undesirable since these edges reflect small changes in surface normals that should be accurately modeled. To address these concerns, we propose introducing adaptability into the Laplacian to sharpen useful edges while suppressing noisy signals. Specifically, edge detection using the adaptive Laplacian can be expressed as follows:

$$\mathbf{e}_i = \varphi \left(\frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \theta \left(\mathbf{f}_j^{\text{ms}} - \mathbf{f}_i^{\text{ms}} \right) \right), \quad (5)$$

where $\mathbf{e}_i$ denotes the detected edges and $\theta, \varphi : \mathbb{R}^{d_{fused}} \rightarrow \mathbb{R}^{d_{fused}}$ denote point-wise MLPs. θ filters the pairwise difference in a local neighborhood to emphasize significant differences and smooth harmful ones. φ is shared among all points to capture a more complex relationship between detected edges, making the representation more discriminative. Consequently, $\mathbf{e}_i$ can effectively describe the “edgeness” of points. Note that Eq. (5) degenerates to the raw Laplacian when θ and φ are removed.

3.4 Edge Conditioning

While detected edges are accurate and robust to data corruption, the edges have limited information about surface orientation as the detector is based on the isotropic Laplacian. Consider a scenario where two perpendicular planes intersect. Our edge detector would ideally produce zero vectors for every point within the planes. However, this would result in the detector being unable to distinguish between points on the different planes, which in turn would prevent the differentiation of their normal vectors. To address this issue, we must supplement the detected edges with global information that contains surface orientation. This can be achieved by combining the detected edges and the output of the backbone. As the backbone network essentially learns the representation of the surface, it naturally contains information about surface orientations (not necessarily the globally correct one because our task is to estimate unoriented normals). Specifically, we propose an edge conditioning operation that conditions edges on the backbone output and generates an edge-aware surface representation:

$$\mathbf{f}_i^{cond} = \mathbf{f}_i^b + \gamma \left(\left[\mathbf{f}_i^b, \mathbf{e}_i \right] \right), \quad (6)$$

where $\mathbf{f}_i^{cond}$ denotes the edge-aware surface representation, $\mathbf{f}_i^b \in \mathbb{R}^{d_{out}}$ is the backbone output, and $\gamma : \mathbb{R}^{d_{fused}+d_{out}} \rightarrow \mathbb{R}^{d_{out}}$ denotes an LBR. It conditions edges to global information with nonlinear mapping to construct a feature that can be used for accurate normal estimation. Similar to Eq. (4), we adopt a residual structure to facilitate optimization while reducing information loss.

3.5 MSECNet

Backbone. As shown in Fig. 2, MSECNet consists of a backbone and an MSEC stream. We adopt the popular PointNet++ [37] except that each encoder layer is adapted into a residual structure [18] to save computation. The network performs subsamplings three times in the encoder, and each time the number of points is halved while the feature dimension is doubled. Then, three upsampling layers are used to recover the original point cloud with nearest-neighbor interpolation. Every interpolation layer adopts a U-Net [39]-like skip connection to help feature reconstruction. The feature dimension of the first stage is set to 64. Unlike many prior studies that generate a single normal for a patch center, we estimate the normals for all points in the patch. In the case of a single-output network, it is necessary for the network to iterate through all points to compute normal vectors for the entire point cloud. In contrast, MSECNet only needs to run until the patches cover the whole cloud, making the inference dramatically faster (Sec. 4.5).




Figure 3: Examples of point clouds in the PCPNet dataset.

Loss functions. We adopt two standard loss functions for normal estimation. The first one is the regression loss [15, 47]:

$$\mathcal{L}^{reg} = \frac{1}{N} \sum_{i=1}^N \min(\|\hat{\mathbf{n}}_i - \mathbf{n}_i\|^2, \|\hat{\mathbf{n}}_i + \mathbf{n}_i\|^2), \quad (7)$$

where $\hat{\mathbf{n}}_i$ denotes an estimated normal. The regression loss is used to penalize the Euclidean distance between estimated and ground truth normals. Additionally, we adopt a sine loss [4, 26] to penalize the angular distance:

$$\mathcal{L}^{sin} = \frac{1}{N} \sum_{i=1}^N |\hat{\mathbf{n}}_i \times \mathbf{n}_i|, \quad (8)$$

where $\times$ denotes the cross product. Consequently, the total loss is

$$\mathcal{L} = \mathcal{L}^{reg} + \mathcal{L}^{sin}. \quad (9)$$

We do not include other popular loss functions like the consistency loss [4, 25, 51] or geometry-guiding loss [26, 47] because they are not applicable or involves manual thresholding.

4 EXPERIMENT

4.1 Datasets and Settings

We evaluate MSECNet with the PCPNet dataset [15] which consists of several synthetic point clouds. Each clean cloud has corrupted variants with different noise levels or density variations. Ground truth normals of corrupted variants are the same as clean ones. Some examples are shown in Fig. 3. Next, we evaluate the generalizability of MSECNet on the SceneNN dataset [20], a real-world dataset consisting of indoor scenes. We use the dataset prepared by [26], which contains clean and noisy scenes. All models are pre-trained on the PCPNet dataset and applied to the SceneNN dataset without fine-tuning. For the PCPNet dataset, we follow the settings in [15] including the train-test split and data augmentation. During training, we randomly sample 1000 patches for each sample. The patch size is set to $N = 700$, following previous works [26, 51]. The patch size (k) of kNN in the backbone is set to 16 and 9 for the MSEC stream. The AdamW [29] optimizer with an initial learning rate of 2×10^{-3} is used and the learning rate is decayed by cosine annealing schedule until it is reduced to 2×10^{-5} . The model is

Table 1: Normal angle RMSE results on the PCPNet and SceneNN datasets. The lower the better.

Method	Year	PCPNet Dataset							SceneNN Dataset		
		Noise σ				Density		Average	Clean	Noise	Average
		None	0.12%	0.6%	1.2%	Stripes	Gradient				
PCA [19]	1992	12.29	12.87	18.38	27.52	13.66	12.81	16.25	15.93	16.32	16.12
Jet [9]	2005	12.35	12.84	18.33	27.68	13.39	13.13	16.29	15.17	15.59	15.38
PCPNet [15]	2018	9.64	11.51	18.27	22.84	11.73	13.46	14.58	20.86	21.40	21.13
Zhou <i>et al.</i> [49]	2020	8.67	10.49	17.62	24.14	10.29	10.66	13.62	-	-	-
Nesti-Net [5]	2019	7.06	10.24	17.77	22.31	8.64	8.95	12.49	13.01	15.19	14.10
Lenßen <i>et al.</i> [24]	2020	6.72	9.95	17.18	21.96	7.73	7.51	11.84	10.24	13.00	11.62
DeepFit [4]	2020	6.51	9.21	16.73	23.12	7.92	7.31	11.80	10.33	13.07	11.70
MTRNet [8]	2021	6.43	9.69	17.08	22.23	8.39	6.89	11.78	-	-	-
Refine-Net [48]	2022	5.92	9.04	16.52	22.19	7.70	7.20	11.43	18.09	19.73	18.91
Zhang <i>et al.</i> [47]	2022	5.65	9.19	16.78	22.93	6.68	6.29	11.25	9.31	13.11	11.21
AdaFit [51]	2021	5.19	9.05	16.45	21.94	6.01	5.90	10.76	8.39	12.85	10.62
GraphFit [25]	2022	4.45	8.74	16.05	21.64	5.22	5.48	10.26	-	-	-
NeAF [27]	2023	4.20	9.25	16.35	21.74	4.89	4.88	10.22	-	-	-
HSurf-Net [26]	2022	4.17	8.78	16.25	21.61	4.98	4.86	10.11	7.55	12.23	9.89
Ours	-	3.84	8.74	16.10	21.05	4.34	4.51	9.76	6.94	11.66	9.30




Figure 4: RMSE visualization. Numbers indicate RMSEs. The “Point cloud” columns show point clouds colored by ground truth normals. (a) Examples from the PCPNet dataset. From top to bottom: clean, density corrupted, and noisy point clouds. (b) Examples from SceneNN dataset. The top one is clean and the others are noisy point clouds.

trained with a batch size of 128 and trained for 150 epochs on four NVIDIA Tesla V100 GPUs.

For the evaluation metric, we use the common angle Root Mean Squared Error (RMSE) between the estimated normal and the ground truth to evaluate performance.

The normal estimation results are reported in Sec. 4.2. Also, we perform various analyses including the ablation study (Sec. 4.3), visualization of edge conditioning (Sec. 4.4), efficiency analysis (Sec. 4.5), and the application to surface reconstruction (Sec. 4.6).

4.2 Normal Estimation

As shown in Tab. 1, MSENet achieves the best scores in terms of average performance, demonstrating its effectiveness. In addition, MSENet achieves the best scores in most categories. Thus, MSENet consistently performs well in a variety of situations. Compared to baselines, MSENet performs particularly well when data are significantly corrupted (see “1.2%” and “Stripes”). We speculate that our fusion strategy and adaptability of the detector lead to a better understanding of the underlying geometry, leading to improved robustness. MSENet also achieves the best score in the real-world scenario, which verifies its good generalizability.

As shown in Fig. 4 (a), the topmost example shows that explicit edge modeling of MSENet significantly reduces errors at sharp edges. When the density is significantly biased (the middle one), MSENet produces a robust prediction despite that many points are missing. We believe that the fusion of all scales leads to more reliable surface descriptions. In the bottom example, strong noise causes many false edges to appear, turning a smooth surface into a rough surface. MSENet produces smooth and accurate predictions, whereas existing methods produce salt-and-pepper errors. This also proves that MSENet learns the underlying geometry well.

4.3 Ablation Study

(a) Number of scales. The effectiveness of the MSEC stream is investigated by varying the number of scales involved, where $S = 1$ represents the smallest scale and each increment of S adds an adjacent larger one. As shown in Tab. 2 (a), the model achieves mediocre performance when the MSEC stream is not used. However, adding just one scale significantly improves performance, demonstrating the effectiveness of explicit edge modeling. Also, consistent improvement for noisy samples as S increases from 1 to 4 suggests that integrating more scales provides better robustness to noise.

(b) Size of d_{fused} . The amount of information retained from different scales is determined by the size of d_{fused} . Tab. 2 (b) shows that increasing d_{fused} from 256 to 1024 results in a reduction in the average error, indicating that a larger d_{fused} can preserve more useful information. The performance of “None” and density variations also improves as d_{fused} increases. Therefore, a larger d_{fused} can provide a more comprehensive descriptor and enhance the network’s robustness against sampling biases.

(c) Interpolation strategy. We investigate the effectiveness of our direct interpolation strategy by comparing it to the standard gradual interpolation, i.e., the feature propagation layer [37]. Note that no transformation is applied for a fair comparison. As shown in Tab. 2 (c), our direct interpolation results in lower errors consistently across all corruption types.

(d) Space and channel transformations. As shown in Tab. 2 (d), an increase in average error indicates that a channel transformation cannot make full use of multi-scale features. The significant decrease in estimation under density variations indicates that sampling biases cannot be well resolved. As expected, the situation improves when the space transformation is used instead of the channel one. The combination of space and channel transformations achieves the best performance, which fully verifies their necessities.

(e) Effect of the adaptability. To validate the effect of the adaptability in the adaptive Laplacian, we remove θ and φ from Eq. (5).




Figure 5: Effect of the MSEC stream on RMSEs.




Figure 6: Visualization of learned features. The sample patches are the same ones as Fig. 5. “Edge” indicates detected edges. “Before” shows the output of the backbone (f_i^b) and “After” shows the features after edge conditioning.

As shown in Tab. 2 (e), the adaptability is effective for estimation under low to medium levels of noise and “Striped” corruption. Thus, the adaptability is most effective in situations where surface geometry is significantly distorted. Also, no improvement is seen under the strongest noise level. For complex point clouds, it is assumed that the noise almost completely destroys the raw local geometry, making edge detection challenging.

(f) Effect of the edge conditioning. As shown in Tab. 2 (f), the importance of global information is confirmed by the 0.35 drop in performance when f_i^b is removed from Eq. (6). To verify the effect of edge conditioning, we only remove f_i^b from γ , which leads to a performance drop of 0.06. This result shows that conditioning is useful, but the degradation is not as severe as when f_i^b is entirely removed thanks to the residual connection.

4.4 Visualization of Edge Conditioning

As shown in Fig. 5, high RMSEs are observed mainly at the edges, where errors are corrected by the MSEC stream. Thus, the effect of the MSEC stream is most pronounced in the accurate estimation at the edges. Further, using the same sample as in Fig. 5, we visualize the detected edges and the effect of edge conditioning using l^2 -norm.

Table 2: Results of ablations. The lower the better. The Δ column shows the difference of “Average” compared with “Default”.

Ablation		Noise σ				Density		Average	Δ
		None	0.12%	0.6%	1.2%	Stripes	Gradient		
Default	$S=4$, $d_{fused} = 1024$, Direct interpolation, Space + Channel, w/ adaptive Lap., w/ edge cond.	3.84	8.74	16.10	21.05	4.34	4.51	9.76	-
(a)	w/o MSEC stream	5.10	8.97	16.22	21.04	5.92	5.73	10.50	-0.74
	$S=1$	3.75	9.04	16.34	21.12	4.43	4.54	9.86	-0.10
	$S=2$	3.73	8.87	16.25	21.14	4.41	4.42	9.80	-0.04
	$S=3$	3.80	8.83	16.24	21.06	4.44	4.44	9.80	-0.04
(b)	$d_{fused}=256$	3.97	8.81	16.10	20.98	4.64	4.67	9.86	-0.10
	$d_{fused}=512$	3.85	8.75	16.07	21.04	4.46	4.58	9.79	-0.03
(c)	Gradual interpolation	3.91	8.92	16.15	21.12	4.59	4.58	9.88	-0.12
(d)	Channel transformation	4.03	8.82	16.12	21.09	4.71	4.75	9.92	-0.16
	Space transformation	3.89	8.83	16.16	21.04	4.52	4.55	9.83	-0.07
(e)	w/o adaptability	3.86	8.85	16.17	21.05	4.44	4.52	9.82	-0.06
(f)	edge only	4.37	8.89	16.22	21.09	5.10	5.01	10.11	-0.35
	w/o conditioning	3.87	8.84	16.11	21.07	4.48	4.53	9.82	-0.06

Table 3: Number of parameters and inference Time. The average inference time for point clouds with 100k points is reported.

	Ours ($d_{fused} = 256$)	Ours ($d_{fused} = 512$)	Ours ($d_{fused} = 1024$)	HSurf-Net [26]	AdaFit [51]	DeepFit [4]
Param. (M)	1.92	3.83	10.40	2.16	4.87	3.53
Time (s)	0.45	0.45	0.93	56.70	26.54	45.12

As can be seen from Fig. 6, the edges are correctly detected, as strong responses are obtained in the edge regions. After conditioning, the feature distribution changes significantly, confirming that edge-conditioned representation is obtained. Combining Fig. 5 and 6, we confirm that the MSEC stream contributes to the improvement.

4.5 Number of Parameters and Inference Time


The results are presented in Tab. 3. We use an NVIDIA Tesla V100 with a batch size of 128. As shown, our default model with $d_{fused} = 1024$ is significantly faster than other methods. Additionally, we can significantly reduce the model size by decreasing d_{fused} without compromising performance too much (see Tab. 2 (b)). Therefore, we conclude that MSECNet can be a practical choice for normal estimation in situations where resources are limited.

4.6 Application to Surface Reconstruction

We use surface reconstruction to demonstrate the utility of MSECNet. Poisson reconstruction [21] is used to reconstruct surfaces. Fig. 7 shows that the reconstructed surfaces using our normals are more precise than those produced by other methods.

5 CONCLUSION

We propose MSECNet, a novel method for improving normal estimation at normal varying regions. The key is that the modeling of normal variations is converted to the modeling of geometric edges. We design the MSEC stream that performs multi-scale edge conditioning. MSECNet achieves state-of-the-art performance on

**Figure 7: Result of surface reconstruction.**

both synthetic and real-world datasets while running significantly faster than existing methods. We also confirm that the MSEC stream better models the geometry in the presence of data corruptions. Therefore, a promising direction for future research is to extend MSECNet to tasks such as point cloud filtering and generation.

ACKNOWLEDGMENTS

This paper is based on the results of a project, JPNP20006, commissioned by the New Energy and Industrial Technology Development Organization (NEDO). We would also like to acknowledge the support from JSPS Grant-in-Aid for Scientific Research (21K12042).

REFERENCES

- [1] Marc Alexa, Johannes Behr, Daniel Cohen-Or, Shachar Fleishman, David Levin, and Claudio T Silva. 2001. Point set surfaces. In *Proceedings Visualization, 2001. VIS'01. IEEE*, 21–29.
- [2] Pierre Alliez, David Cohen-Steiner, Yiyi Tong, and Mathieu Desbrun. 2007. Voronoi-based variational reconstruction of unoriented point sets. In *Symposium on Geometry processing*, Vol. 7. 39–48.
- [3] Pablo Arbelaez, Michael Maire, Charles Fowlkes, and Jitendra Malik. 2010. Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence* 33, 5 (2010), 898–916.
- [4] Yizhak Ben-Shabat and Stephen Gould. 2020. DeepFit: 3D surface fitting via neural network weighted least squares. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I* 16. Springer, 20–34.
- [5] Yizhak Ben-Shabat, Michael Lindenbaum, and Anath Fischer. 2019. Nesti-net: Normal estimation for unstructured 3d point clouds using convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10112–10120.
- [6] Alexandre Boulch and Renaud Marlet. 2016. Deep learning for robust normal estimation in unstructured point clouds. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 281–290.
- [7] John Canny. 1986. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence* 6 (1986), 679–698.
- [8] Junjie Cao, Hairui Zhu, Yunpeng Bai, Jun Zhou, Jinshan Pan, and Zhixun Su. 2021. Latent tangent space representation for normal estimation. *IEEE Transactions on Industrial Electronics* 69, 1 (2021), 921–929.
- [9] Frédéric Cazals and Marc Pouget. 2005. Estimating differential quantities using polynomial fitting of osculating jets. *Computer Aided Geometric Design* 22, 2 (2005), 121–146.
- [10] Tamal K Dey and Jian Sun. 2006. Normal and feature approximations from noisy point clouds. In *FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science: 26th International Conference, Kolkata, India, December 13–15, 2006. Proceedings* 26. Springer, 21–32.
- [11] Piotr Dollár and C Lawrence Zitnick. 2014. Fast edge detection using structured forests. *IEEE transactions on pattern analysis and machine intelligence* 37, 8 (2014), 1558–1570.
- [12] Daniel Garrido, Rui Rodrigues, A Augusto Sousa, Joao Jacob, and Daniel Castro Silva. 2021. Point cloud interaction and manipulation in virtual reality. In *2021 5th Int'l Conf. on Artificial Intelligence and Virtual Reality (AIVR)*. 15–20.
- [13] Henri Gouraud. 1971. Continuous shading of curved surfaces. *IEEE transactions on computers* 100, 6 (1971), 623–629.
- [14] Gaël Guennebaud and Markus Gross. 2007. Algebraic point set surfaces. In *ACM siggraph 2007 papers*. 23–es.
- [15] Paul Guerrero, Yanir Kleiman, Maks Ovsjanikov, and Niloy J Mitra. 2018. Pcpnet learning local shape properties from raw point clouds. In *Computer graphics forum*, Vol. 37. Wiley Online Library, 75–85.
- [16] Taisuke Hashimoto and Masaki Saito. 2019. Normal Estimation for Accurate 3D Mesh Reconstruction with Point Cloud Model Incorporating Spatial Structure.. In *CVPR workshops*, Vol. 1.
- [17] Jianzhong He, Shiliang Zhang, Ming Yang, Yanhu Shan, and Tiejun Huang. 2019. Bi-directional cascade network for perceptual edge detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3828–3837.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [19] Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. 1992. Surface reconstruction from unorganized points. In *Proceedings of the 19th annual conf. on computer graphics and interactive techniques*. 71–78.
- [20] Binh-Son Hua, Quang-Hieu Pham, Duc Thanh Nguyen, Minh-Khoi Tran, Lap-Fai Yu, and Sai-Kit Yeung. 2016. Scenenn: A scene meshes dataset with annotations. In *2016 fourth international conference on 3D vision (3DV)*. Ieee, 92–101.
- [21] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. 2006. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, Vol. 7. 0.
- [22] Josef Kittler. 1983. On the accuracy of the Sobel edge detector. *Image and Vision Computing* 1, 1 (1983), 37–42.
- [23] Peter Lancaster and Kes Salkauskas. 1981. Surfaces generated by moving least squares methods. *Mathematics of computation* 37, 155 (1981), 141–158.
- [24] Jan Eric Lenssen, Christian Osendorfer, and Jonathan Masci. 2020. Deep iterative surface normal estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 11247–11256.
- [25] Keqiang Li, Mingyang Zhao, Huaiyu Wu, Dong-Ming Yan, Zhen Shen, Fei-Yue Wang, and Gang Xiong. 2022. GraphFit: Learning Multi-scale Graph-Convolutional Representation for Point Cloud Normal Estimation. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*. Springer, 651–667.
- [26] Qing Li, Yu-Shen Liu, Jin-San Cheng, Cheng Wang, Yi Fang, and Zhizhong Han. 2022. HSurf-Net: Normal Estimation for 3D Point Clouds by Learning Hyper Surfaces. *Advances in Neural Information Processing Systems (NeurIPS)* (2022).
- [27] Shujuan Li, Junsheng Zhou, Baorui Ma, Yu-Shen Liu, and Zhizhong Han. 2023. NeAF: Learning Neural Angle Fields for Point Normal Estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [28] Yun Liu, Ming-Ming Cheng, Xiaowei Hu, Kai Wang, and Xiang Bai. 2017. Richer convolutional features for edge detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3000–3009.
- [29] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Bkg6RiCqY7>
- [30] Dening Lu, Xuequan Lu, Yangxing Sun, and Jun Wang. 2020. Deep feature-preserving normal estimation for point cloud filtering. *Computer-Aided Design* 125 (2020), 102860.
- [31] David Marr and Ellen Hildreth. 1980. Theory of edge detection. *Proceedings of the Royal Society of London. Series B. Biological Sciences* 207, 1167 (1980), 187–217.
- [32] David R Martin, Charles C Fowlkes, and Jitendra Malik. 2004. Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE transactions on pattern analysis and machine intelligence* 26, 5 (2004), 530–549.
- [33] Niloy J Mitra and An Nguyen. 2003. Estimating surface normals in noisy point cloud data. In *Proceedings of the nineteenth annual symposium on Computational geometry*. 322–328.
- [34] Mengyang Pu, Yaping Huang, Qingji Guan, and Haibin Ling. 2021. Rindnet: Edge detection for discontinuity in reflectance, illumination, normal and depth. In *Proceedings of the IEEE/CVF international conf. on computer vision*. 6879–6888.
- [35] Mengyang Pu, Yaping Huang, Yuming Liu, Qingji Guan, and Haibin Ling. 2022. Edge detection with transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1402–1412.
- [36] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. 2017. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 652–660.
- [37] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. 2017. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems* 30 (2017).
- [38] Xiaofeng Ren. 2008. Multi-scale improves boundary detection in natural images. In *Computer Vision—ECCV 2008: 10th European Conference on Computer Vision, Marseille, France, October 12–18, 2008, Proceedings, Part III* 10. Springer, 533–545.
- [39] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi (Eds.). Springer International Publishing, Cham, 234–241.
- [40] Riccardo Roveri, A Cengiz Öztireli, Ioana Pandele, and Markus Gross. 2018. Pointprons: Consolidation of point clouds with convolutional neural networks. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 87–99.
- [41] Rajat Sharma, Tobias Schwanndt, Christian Kunert, Steffen Urban, and Wolfgang Broll. 2021. Point cloud upsampling and normal estimation using deep learning for robust surface reconstruction. *arXiv preprint arXiv:2102.13391* (2021).
- [42] Zhuo Su, Wenzhe Liu, Zitong Yu, Dewen Hu, Qing Liao, Qi Tian, Matti Pietikäinen, and Li Liu. 2021. Pixel difference networks for efficient edge detection. In *Proceedings of the IEEE/CVF international conf. on computer vision*. 5117–5127.
- [43] Gabriel Taubin. 1995. A signal processing approach to fair surface design. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. 351–358.
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [45] Weijia Wang, Xuequan Lu, Dasith De Silva Edirimuni, Xiao Liu, and Antonio Robles-Kelly. 2022. Deep point cloud normal estimation via triplet learning. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 1–6.
- [46] Saining Xie and Zhuowen Tu. 2015. Holistically-nested edge detection. In *Proceedings of the IEEE international conference on computer vision*. 1395–1403.
- [47] Jie Zhang, Jun-Jie Cao, Hai-Rui Zhu, Dong-Ming Yan, and Xiu-Ping Liu. 2022. Geometry guided deep surface normal estimation. *Computer-Aided Design* 142 (2022), 103119.
- [48] Haoran Zhou, Honghua Chen, Yingkui Zhang, Mingqiang Wei, Haoran Xie, Jun Wang, Tong Lu, Jing Qin, and Xiao-Ping Zhang. 2022. Refine-net: Normal refinement neural network for noisy point clouds. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 1 (2022), 946–963.
- [49] Jun Zhou, Hua Huang, Bin Liu, and Xiuping Liu. 2020. Normal estimation for 3D point clouds via local plane constraint and multi-scale selection. *Computer-Aided Design* 129 (2020), 102916.
- [50] Jun Zhou, Wei Jin, Mingjie Wang, Xiuping Liu, Zhiyang Li, and Zhaobin Liu. 2023. Improvement of normal estimation for point clouds via simplifying surface fitting. *Computer-Aided Design* 161 (2023), 103533.
- [51] Runsong Zhu, Yuan Liu, Zhen Dong, Yuan Wang, Tengping Jiang, Wenping Wang, and Bisheng Yang. 2021. AdaFit: rethinking learning-based normal estimation on point clouds. In *Proceedings of the IEEE/CVF international conference on computer vision*. 6118–6127.