

# Nesti-Net: Normal Estimation for Unstructured 3D Point Clouds using Convolutional Neural Networks

Yizhak Ben-Shabat  
 Mechanical Engineering  
 Technion IIT  
 Haifa, Israel  
 sitzikbs@gmail.com

Michael Lindenbaum  
 Computer Science  
 Technion IIT  
 Haifa, Israel  
 mic@cs.technion.ac.il

Anath Fischer  
 Mechanical Engineering  
 Technion IIT  
 Haifa, Israel  
 meranath@technion.ac.il

## Abstract

*In this paper, we propose a normal estimation method for unstructured 3D point clouds. This method, called Nesti-Net, builds on a new local point cloud representation which consists of multi-scale point statistics (MuPS), estimated on a local coarse Gaussian grid. This representation is a suitable input to a CNN architecture. The normals are estimated using a mixture-of-experts (MoE) architecture, which relies on a data-driven approach for selecting the optimal scale around each point and encourages sub-network specialization. Interesting insights into the network’s resource distribution are provided. The scale prediction significantly improves robustness to different noise levels, point density variations and different levels of detail. We achieve state-of-the-art results on a benchmark synthetic dataset and present qualitative results on real scanned scenes.*

## 1. Introduction

Commodity 3D sensors are rapidly becoming an integral part of autonomous systems. These sensors, *e.g.* RGB-D cameras or LiDAR, provide a 3D point cloud representing the geometry of the scanned objects and surroundings. This raw representation is challenging to process since it lacks connectivity information or structure, and is often incomplete, noisy and contains point density variations. In particular, processing it by means of the highly effective convolutional neural networks (CNNs) is problematic because CNNs require structured, grid-like data as input.

When available, additional local geometric information, such as the surface normals at each point, induces a partial local structure and improves performance of different tasks such as over-segmentation [3], classifica-

tion [21] and surface reconstruction [11].

Estimating the normals from a raw, point-only, cloud, is a challenging task due to difficulties associated with sampling density, noise, outliers, and detail level. The common approach is to specify a local neighborhood around a point and to fit a local basic geometric surface (*e.g.* a plane) to the points in this neighborhood. Then the normal is estimated from the fitted geometric entity. The chosen size (or scale) of the neighborhood introduces an unavoidable tradeoff between robustness to noise and accuracy of fine details. A large-scale neighborhood over-smoothes sharp corners and small details but is otherwise robust to noise. A small neighborhood, on the other hand, may reproduce the normals more accurately around small details but is more sensitive to noise. Thus it seems that an adaptive, data-driven scale may improve estimation performance.

We propose a normal estimation method for unstructured 3D point clouds. It features a mixture-of-experts network for scale prediction, which significantly increases its robustness to different noise levels, outliers, and varying levels of detail. In addition, this method overcomes the challenge of feeding point clouds into CNNs by extending the recently proposed 3D modified Fischer Vector (3DmFV) representation [4] to encode local geometry on a coarse multi-scale grid. It outperforms state-of-the-art methods for normal vector estimation.

The main contributions of this paper are:

- A new normal estimation method for unstructured 3D point clouds based on mixture of experts and scale prediction.
- A local point representation which can be used as input to a CNN.

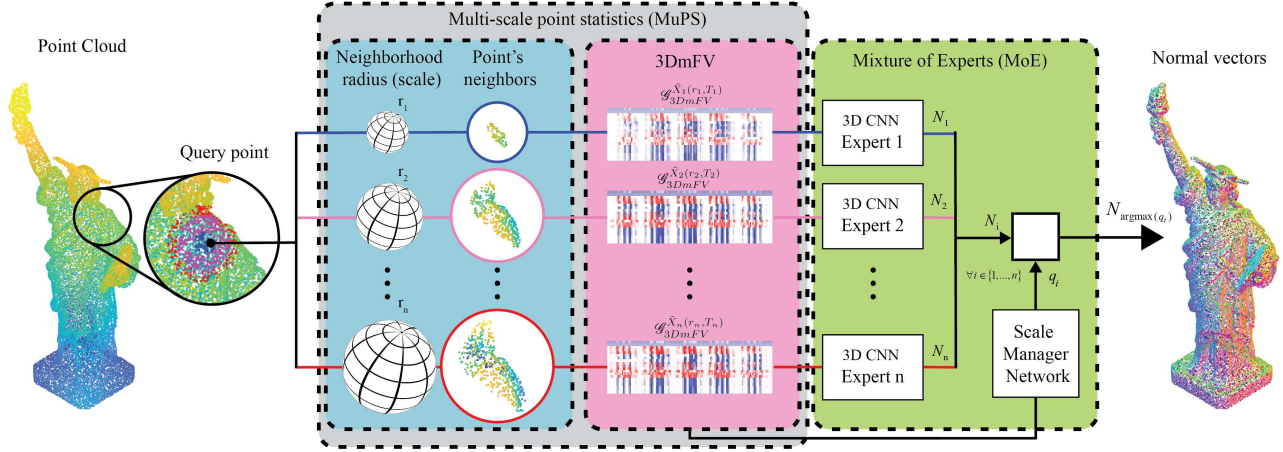


Figure 1. Nesti-Net pipeline for normal estimation. For each point in a given point cloud, we compute a multi-scale point statistics representation (MuPS). Then, a scale manager network is used to determine the optimal scale and uses the corresponding expert sub-network to estimate the normal.

## 2. Related-work

### 2.1. Deep learning for unstructured 3D point clouds

The point cloud representation is challenging for deep learning methods because it is both unstructured and point-wise unordered. In addition, the number of points in the point cloud is usually not constant. Several methods were proposed to overcome these challenges. Voxel-based methods embed the point cloud into a voxel grid but suffer from several accuracy-complexity tradeoffs [16]. The PointNet approach [20, 21] applies a symmetric, order-insensitive, function on a high-dimensional representation of individual points. The Kd-Network [14] imposes a kd-tree structure on the points and uses it to learn shared weights for nodes in the tree. The recently proposed 3DmFV [4] represents the points by their deviation from a Gaussian Mixture Model (GMM) whose Gaussians are uniformly positioned on a coarse grid.

In this paper, we propose a point-wise and multi-scale variation of 3DmFV. Instead of generating a structured representation for the entire point cloud, we represent each point and its neighbors within several scales.

### 2.2. Normal estimation

A classic method for normal estimation uses Principal Component Analysis (PCA) [12]. It first specifies the neighbors within some scale, and then uses PCA regression to estimate a tangent plane. Variants fitting local spherical surfaces [10] or jets [7] (truncated Taylor expansion) were also proposed. To be robust to noise, these methods usually choose a large-scale neighborhood, leading them to smooth sharp features and

to fail in estimating normals near edges. Computing the optimal neighborhood size can decrease the estimation error [18] but requires the (usually unknown) noise standard deviation value and a costly iterative process to estimate the local curvature and additional density parameters.

Other approaches rely on using Voronoi cells of point clouds [2, 17, 9]. These methods are characterized by robustness to sharp features but are sensitive to noise. To overcome this challenge, Alliez et al. [1] proposed PCA-Voronoi approach to create cell sets which group adjacent cells to provide some control over smoothness. While many of these methods hold theoretical guarantees on approximation and robustness, in practice they rely on a preprocessing stage in the presence of strong or unstructured noise in addition to a fine-tuned set of parameters.

A few deep learning approaches have been proposed to estimate normal vectors from unstructured point clouds. Boulch et al. proposed to transform local point cloud patches into a 2D Hough space accumulator by randomly selecting point triplets and voting for that plane’s normal. Then, the normal is estimated from the accumulator by designing explicit criteria [5] for bin selection or, more recently, by training a 2D CNN [6] to estimate it continuously as a regression problem. This method does not fully utilize the 3D information since it loses information during the transformation stage. We refer to this method as HoughCNN in the evaluation section. A more recent method, PCPNnet [11], uses a PointNet [20] architecture on local point neighborhoods of multiple scales. It achieves good normal estimation performance and has been extended to estimating other surface properties. However, it processes

the multi-scale point clouds jointly and does not select an optimal scale. This type of architecture tends to encourage averaging during training rather than specialization [13].

In this paper we propose a method that approximates the local normal vector using a point-wise, multi-scale 3DmFV representation which serves as an input to a deep 3D CNN architecture. In addition, we learn the neighborhood size that minimizes the normal estimation error using a mixture of experts (MoE) [13], which encourages specialization.

### 2.3. Representing point clouds using 3DmFV

The 3DmFV representation for point clouds [4] achieved good results for point cloud classification using a 3D CNN. See Section 3.1 for details. In this paper we propose the Multi-scale Point Statistics (MuPS) representation, which extends the 3DmFV and computes a point-wise multi-scale 3DmFV.

## 3. Approach

The proposed method is outlined in Figure 1. It receives a 3D point cloud as input and consists of two main stages. In the first stage, we compute a multiscale point representation, denoted MuPS. In the second stage we feed it into a mixture-of-experts (MoE) CNN architecture and estimate the normal at each point as output. The stages are detailed below.

### 3.1. MuPS - Multi-scale point statistics

MuPS is a local multi-scale representation which computes point statistics on a coarse Gaussian grid. It builds on the well-known Fisher Vector (FV) [22], and the recently proposed 3DmFV representation [4]. Therefore, we first outline the FV and the 3DmFV, and then continue to the MuPS representation and its attractive properties.

**FV and 3DmFV for 3D point clouds:** Given a set of  $T$  3D points  $X = \{\mathbf{p}_t \in \mathbb{R}^3, t = 1, \dots, T\}$  and a set of parameters for a  $K$  component GMM,  $\lambda = \{(w_k, \mu_k, \Sigma_k), k = 1, \dots, K\}$ , where  $w_k, \mu_k, \Sigma_k$  are the mixture weight, center, and covariance matrix of the  $k$ -th Gaussian. The likelihood of a single 3D point  $\mathbf{p}$  associated with the  $k$ -th Gaussian density is

$$u_k(\mathbf{p}) = \frac{1}{(2\pi)^{D/2} |\Sigma_k|^{1/2}} \exp \left\{ -\frac{1}{2} (\mathbf{p} - \mu_k)' \Sigma_k^{-1} (\mathbf{p} - \mu_k) \right\}. \quad (1)$$

Therefore, the likelihood of a single point associated

with the GMM density is:

$$u_\lambda(\mathbf{p}) = \sum_{k=1}^K w_k u_k(\mathbf{p}). \quad (2)$$

The 3DmFV uses a uniform GMM on a coarse  $m \times m \times m$  3D grid, where  $m$  is an adjustable parameter usually chosen to be from  $m = 3$  to 8. The weights are set to be  $w_k = \frac{1}{K}$ , the standard deviation is set to be  $\sigma_k = \frac{1}{m}$ , and the covariance matrix to be  $\Sigma_k = \sigma_k I$ . Although the parameters in GMMs are usually set using maximum likelihood estimation, here uniformity is crucial for shared weight filtering (convolutions).

The FV is expressed as the sum of normalized gradients for each point  $\mathbf{p}_t$ . The 3DmFV is specified similarly using additional symmetric functions, *i.e.*  $\min$  and  $\max$ . They are symmetric in the sense proposed in [20] and are therefore adequate for representing the structureless and orderless set of points. Adding these functions makes the representation more informative and the associated classification more accurate [4]:

$$\mathcal{G}_{FV_\lambda}^X = \sum_{t=1}^T L_\lambda \nabla_\lambda \log u_\lambda(\mathbf{p}_t), \quad (3)$$

$$\mathcal{G}_{3DmFV_\lambda}^X = \begin{bmatrix} \sum_{t=1}^T L_\lambda \nabla_\lambda \log u_\lambda(\mathbf{p}_t) \Big|_{\lambda=\alpha, \mu, \sigma} \\ \max_t (L_\lambda \nabla_\lambda \log u_\lambda(\mathbf{p}_t)) \Big|_{\lambda=\alpha, \mu, \sigma} \\ \min_t (L_\lambda \nabla_\lambda \log u_\lambda(\mathbf{p}_t)) \Big|_{\lambda=\mu, \sigma} \end{bmatrix} \quad (4)$$

where  $L_\lambda$  is the square root of the inverse Fisher Information Matrix, and the normalized gradients are:

$$\mathcal{G}_{\alpha_k}^X = \frac{1}{\sqrt{w_k}} \sum_{t=1}^T (\gamma_t(k) - w_k), \quad (5)$$

$$\mathcal{G}_{\mu_k}^X = \frac{1}{\sqrt{w_k}} \sum_{t=1}^T \gamma_t(k) \left( \frac{\mathbf{p}_t - \mu_k}{\sigma_k} \right), \quad (6)$$

$$\mathcal{G}_{\sigma_k}^X = \frac{1}{\sqrt{2w_k}} \sum_{t=1}^T \gamma_t(k) \left[ \frac{(\mathbf{p}_t - \mu_k)^2}{\sigma_k^2} - 1 \right]. \quad (7)$$

Here, we follow [15] and ensure that  $u_\lambda(x)$  is a valid distribution by changing the variable  $w_k$  to  $\alpha_k$  to simplify the gradient calculation using :

$$w_k = \frac{\exp(\alpha_k)}{\sum_{j=1}^K \exp(\alpha_j)}. \quad (8)$$

In addition,  $\gamma_t(k)$  expresses the soft assignment of point  $\mathbf{p}_t$  to Gaussian  $k$ , as obtained from the derivatives:

$$\gamma_t(k) = \frac{w_k u_k(\mathbf{p}_t)}{\sum_{j=1}^K w_j u_j(\mathbf{p}_t)}. \quad (9)$$

The FV and 3DmFV are normalized by the number of points in order to be sample-size independent [22]:

$$\mathcal{G}_{FV_\lambda}^X \leftarrow \frac{1}{T} \mathcal{G}_{FV_\lambda}^X, \mathcal{G}_{3DmFV_\lambda}^X \leftarrow \frac{1}{T} \mathcal{G}_{3DmFV_\lambda}^X. \quad (10)$$

Note that these are global representations which are applied to the entire set, *i.e.*, the entire point cloud.

**MuPS definition :** For each point  $\mathbf{p}$  in point set  $X$  we first extract  $n$  point subsets  $\tilde{X}_i(r_i)_{i=1,\dots,n} \subset X$  which contain  $T_i(\mathbf{p}, r_i)$  points and lie within a distance of radius  $r_i$  from  $\mathbf{p}$ . We refer to each of these subsets as a scale. Note that each scale may contain a different number of points. For scales with many points, we set a maximal point threshold, and sample a random subset of  $T_{max}$  points for that scale. Here,  $r_i$  and  $T_{max}$  are design parameters. Next, the scales (subsets) are independently translated and uniformly scaled so that they fit into a zero-centered unit sphere with  $\mathbf{p}$  mapped to the origin. Then, the 3DmFV representation is computed for each scale relative to a Gaussian grid positioned around the origin; see above. Concatenating the 3DmFVs of all scales yields the MuPS representation:

$$\mathcal{G}_{MuPS}^{\mathbf{p}} = \left( \mathcal{G}_{3DmFV}^{\tilde{X}_1(r_1)}, \dots, \mathcal{G}_{3DmFV}^{\tilde{X}_n(r_n)} \right). \quad (11)$$

*MuPS properties:* The MuPS representation overcomes the main challenges associated with feeding point clouds into CNNs. The symmetric functions make it independent of the number and order of points within each scale. In addition, the GMM gives it its grid structure, necessary for the use of CNNs. Furthermore, the multi-scale representation incorporates description of fine detail as well as robustness to noise.

### 3.2. The Nesti-Net architecture

The deep network architecture is outlined in Figure 1 (the green part). It is a mixture-of-experts architecture [13] which consists of two modules: a scale manager network module, and an experts module. The MoE architecture was chosen in order to overcome the averaging effect of typical networks when solving a regression problem.

*Scale manager network:* This module receives the MuPS representation as input and processes it using several 3D Inception inspired convolutions, and maxpool layers, followed by four fully connected layers, after which a softmax operator is applied. The architecture is specified in the top left part of Figure 2. The output is a vector of  $n$  scalars  $q_i$ , which can be intuitively interpreted as the probability of expert  $i$  to

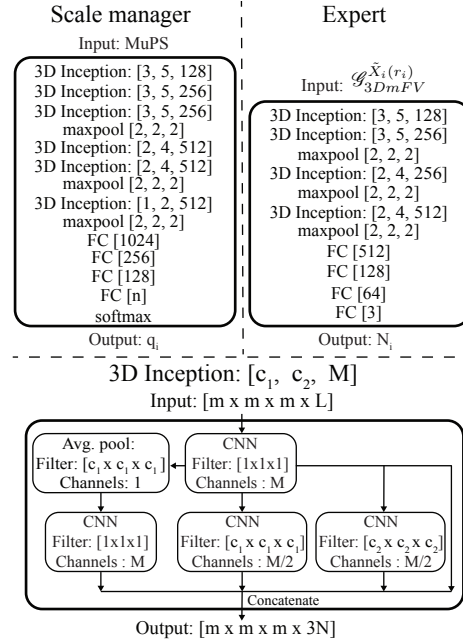


Figure 2. The mixture of experts and 3D Inception module architecture details. The scale manager and experts use several convolutional and maxpooling layers followed by fully connected layers.

estimate the normal correctly.

*Experts:* The normal is estimated using  $n$  separate "expert" networks. Each is a multi-layered 3D Inception inspired CNN followed by four fully connected layers. The MuPS representation is distributed to the experts. This distribution is a design choice. We obtained the best results when feeding each scale to two different experts in addition to one expert which receives the entire MuPS representation as input. Specifically, Nesti-Net uses 7 experts: experts 1-2 receive the smallest scale (1%), 3-4 the medium scale (3%), 5-6 the large scale (5%), and expert 7 receives all the scales. The last layer of each expert outputs a three-element vector  $N_i = (N_x, N_y, N_z)_i$ . The final predicted normal (for point  $\mathbf{p}$ ) is  $N_{argmax(q_i)}$ , *i.e.*, the normal associated with the expert expected to give the best results. The architecture is specified in the top right of Figure 2.

*Loss function:* We train the network to minimize the difference between a predicted normal  $N_i$  and a ground truth normal  $N_{GT}$ . This difference is quantified by the metric  $D_N = \sin \theta$ , where the angle  $\theta$  is the difference between the vectors, and  $D_N$  is calculated as the magnitude of the cross product between these two vectors; see Eq. 12. In addition, to encourage specialization of each expert network, we follow [13] and minimize the

loss:

$$L = \sum_{i=1}^n q_i \cdot D_N = \sum_{i=1}^n q_i \frac{\|N_i \times N_{GT}\|}{\|N_i\| \cdot \|N_{GT}\|}. \quad (12)$$

Using this loss, each expert is rewarded for specializing in a specific input type. Note that during training, all  $n$  normal vectors are predicted and used to compute the loss and derivatives. However, at test time, we compute only one normal, which is associated with the maximal  $q_i$ .

## 4. Evaluation

### 4.1. Datasets

For training and testing we used the PCPNet shape dataset [11]. The trainset consists of 8 shapes: four CAD objects (fandisk, boxunion, flower, cup) and four high quality scans of figurines (bunny, armadillo, dragon and turtle). All shapes are given as triangle meshes and densely sampled with 100k points. The data is augmented by introducing Gaussian noise for each point’s spacial location with a standard deviation of 0.012, 0.006, 0.00125 w.r.t the bounding box. This yields a set with 3.2M points and 3.2M corresponding training examples. The test set consists of 22 shapes, including figurines, CAD objects, and analytic shapes. For evaluation we use the same 5000 point subset per shape as in [11].

For qualitative testing on scanned data, we used the NYU Depth V2 dataset [19] and the recent ScanNet dataset [8], which include RGB-D images of indoor scenes.

### 4.2. Training details

All variations of our method were trained using 32,768 (1024 samples  $\times$  32 shapes) random subsets of the 3.2M training samples at each epoch. For each point, we extract 512 neighboring points enclosed within a sphere of radius  $r$ . For neighborhoods with more than 512 points, we perform random sampling, and for those with fewer points we use the maximum number of points available. For the MuPS representation we chose to use an  $m = 8$  Gaussian grid. We used Tensorflow on a single NVIDIA Titan Xp GPU.

### 4.3. Normal estimation performance

We use the RMS normal estimation error metric for comparing the proposed NestiNet to other deep learning based [11, 6] and geometric methods [12, 7]. We also analyze robustness for two types of data corruptions (augmentations):

- Gaussian noise - perturbing the points with three levels of noise specified by  $\sigma$ , given as a percentage of the bounding box.
- Density variation - selecting a subset of the points based on two sampling regimes: gradient, simulating effects of distance from the sensor, and stripes, simulating occlusions.

For the geometric methods, we show results for three different scales: small, medium and large, which correspond to 18, 112, 450 nearest neighbors. For the deep learning based methods we show the results for the single-scale (ss) and multi-scale (ms) versions. Additional evaluation results using other metrics are available in the supplemental material.

Table 1 shows the unoriented normal estimation results for the methods detailed above. It can be seen that our method outperforms all other methods across all noise levels and most density variations. It also shows that both PCA and Jet perform well for specific noise-scale pairs. In addition, for PCPNet and HoughCNN, using a multi-scale approach only mildly improves performance.

Figure 3 illustrates Nesti-Net’s results on three point clouds. For visualization, the normal vector is mapped to the RGB cube. It shows that for complex shapes (pillar, liberty) with high noise levels, the general direction of the normal vector is predicted correctly, but, the fine details and exact normal vector are not obtained. For a basic shape (Boxysmooth) the added noise does not affect the results substantially. Most notably, Nesti-Net shows robustness to point density corruptions. The angular error in each point is visualized in Figure 4 for the different methods using a heat map. For PCA and Jet we display the best result out of the three scales (small, medium, and large, specified above), and for PCPNet the best out of its single-scale and multi-scale options. For all methods, it can be seen that more errors occur near edges, corners and small regions with a lot of detail and high curvature. Nesti-Net suffers the least from this effect due to its scale manager, which allows it to adapt to the different local geometry types. Figure 5 shows the performance of the scale manager network. A color is assigned to each expert and the chosen expert color is visualized over the point cloud. This provides some insight regarding each expert’s specialization. For example, the figure shows that experts 1, 2 (small scale) specialize in points in regions with high curvatures (near corners). Experts 3 and 4 (medium scale) specialize in the complex cases where multiple surfaces are close to each other, or in the presence of noise. As for the large-scale experts, expert 5 specializes in planar surfaces with normal vec-



Aug.	Our Nesti-Net	PCA [12]			Jet [7]			PCPNet [11]		HoughCNN [6]	
scale	ms	small	med	large	small	med	large	ss	ms	ss	ms
None	<b>6.99</b>	8.31	12.29	16.77	7.60	12.35	17.35	9.68	9.62	10.23	10.02
<b>Noise</b>											
$\sigma = 0.00125$	<b>10.11</b>	12.00	12.87	16.87	12.36	12.84	17.42	11.46	11.37	11.62	11.51
$\sigma = 0.006$	<b>17.63</b>	40.36	18.38	18.94	41.39	18.33	18.85	18.26	18.87	22.66	23.36
$\sigma = 0.012$	<b>22.28</b>	52.63	27.5	23.5	53.21	27.68	23.41	22.8	23.28	33.39	36.7
<b>Density</b>											
Gradient	9.00	9.14	12.81	17.26	<b>8.49</b>	13.13	17.8	13.42	11.7	11.02	10.67
Stripes	<b>8.47</b>	9.42	13.66	19.87	8.61	13.39	19.29	11.74	11.16	12.47	11.95
<b>average</b>	<b>12.41</b>	21.97	16.25	18.87	21.95	16.29	19.02	14.56	14.34	16.9	17.37

Table 1. Comparison of the RMS angle error for unoriented normal vector estimation of our Nesti-Net method to classic geometric methods (PCA [12], Jet [7]) with three scales, and deep learning methods (PCPNet [11], HoughCNN [6])

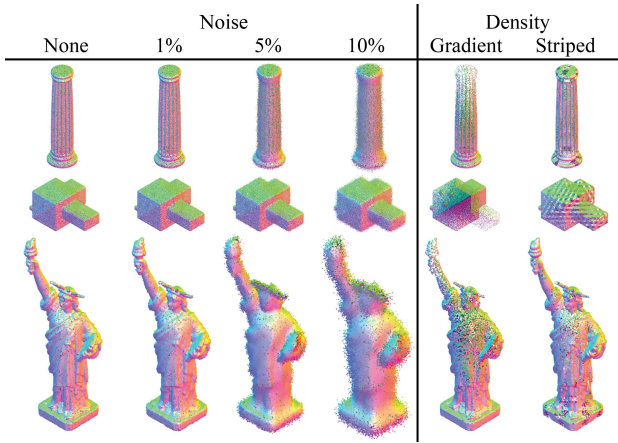


Figure 3. Nesti-Net normal prediction results for different noise levels (columns 1-4), and density distortions (columns 5-6). The colors of the points are normal vectors mapped to RGB. This figure is best viewed digitally on a large screen.

tors, which have a large component in the  $x$  direction, whereas expert 6 specializes in planar surfaces, which have a large component in the  $y$  direction. Expert 5 also specializes in very noisy planar surfaces with a large component in the  $z$  direction. Expert 7 (combined scales) plays multiple roles; it handles points on planar surfaces which have a large component in the  $z$  direction, complex geometry, and low to medium noise. Figure 6 shows the number of points assigned to each expert for all points in the test set, and the average error per expert. It shows an inverse relation between the number of points assigned to an expert and its average error: the more points assigned to the expert, the lower the error. This is consistent with the definition of the cost function. Timing performance and visualization of additional results are provided in the supplemental material.

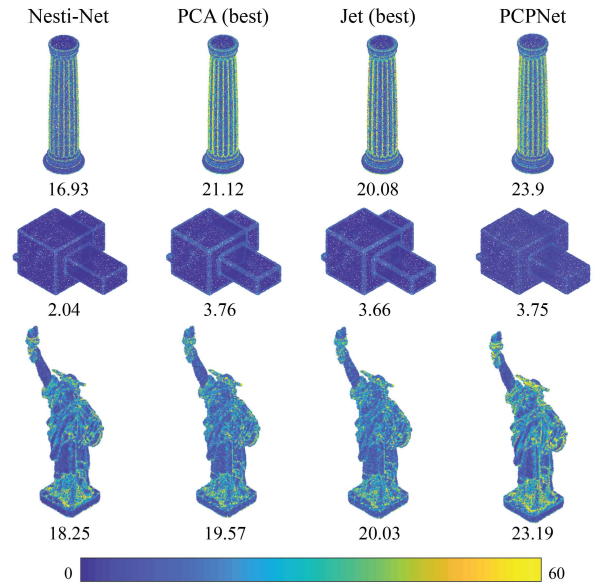


Figure 4. Normal estimation error results for Nesti-Net compared to other methods for three types of point clouds with low noise level ( $\sigma = 1\%$ ). The colors of the points correspond to angular difference, mapped to a heatmap ranging from 0-60 degrees; see bottom color bar. The number under each point cloud is its RMS error.

#### 4.4. Scale selection performance

We analyze the influence of scale selection on the normal estimation performance. We create several ablations of our method.

- ss - A single scale version which directly feeds a 3DmFV representation into a CNN architecture (a single-scale MuPS).
- ms - A multi-scale version which feeds the MuPS representation into a CNN architecture.
- ms-sw - A multi scale version which first tries to

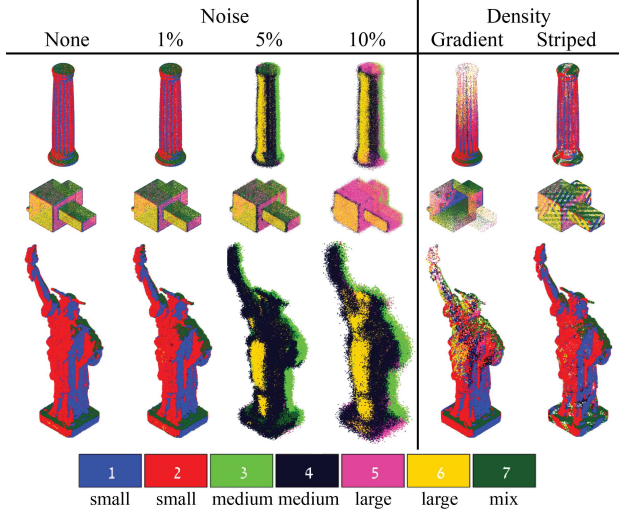


Figure 5. Nesti-Net predicted experts (scales). Each color represents the predicted expert for optimal normal estimation. Color coding is given at the bottom.




Figure 6. Nesti-Net expert (scale) prediction statistics. Number of points assigned to each expert (left), and average expert error (right).

estimate the noise level and then feeds the 3DmFV representation of the corresponding input scale into different sub-networks for a discrete number of noise levels (switching). Note that for this version, the noise level is provided during training.

- NestiNet - The method described in Section 3 which uses an MoE network to learn the scale.

Details of the architectures for the above methods are provided in the supplemental material.

Table 2 summarizes the results of the scale selection performance analysis. It shows that Nesti-Net’s scale selection performs better than all other variations. This is due to the trained scale-manager network within the MoE. The single-scale version performs well for specific noise-scale pairs but inferior performance for an inadequate scale selection. The multi-scale variations show improvement; however, selecting the correct scale yields improved performance over concatenating multiple scales. The main advantage of Nesti-Net over the switching variation is that the scale prediction is

Aug.	ss		ms	ms-sw	NestiNet	
scale	0.01	0.05	0.01 0.05	0.01 0.05	0.01 0.05	0.01 0.03 0.05
None	9.32	12.73	10.83	7.88	7.76	<b>6.99</b>
<b>Noise</b>						
$\sigma = 0.00125$	11.31	13.36	12.98	10.46	10.29	<b>10.11</b>
$\sigma = 0.006$	36.5	18.37	21.06	18.43	18.45	<b>17.63</b>
$\sigma = 0.012$	55.24	23.14	26.03	22.59	22.25	<b>22.28</b>
<b>Density</b>						
Gradient	16.61	14.65	12.81	11.89	9.44	<b>9.00</b>
Stripes	14.5	14.57	12.97	10.06	9.65	<b>8.47</b>
<b>average</b>	23.91	16.14	16.11	13.55	12.97	<b>12.41</b>

Table 2. Comparison of the RMS angle error for unoriented normal vector estimation of our method using single-scale (SS), multi-scale (MS), multi-scale with switching (MS-Sw) and multi-scale with mixture of experts (Nesti-Net)

unsupervised, *i.e.*, does not need the additional noise parameters as input during training.

#### 4.5. Results on scanned data

We show qualitative results on scanned point clouds from the ScanNet [8] and NYU Depth V2 [19] datasets in Figure 7. For visualization we project the normal vectors’ color back to the depth image plane. column (c) shows the results for Nesti-Net, trained on synthetic data with Gaussian noise. The estimated normals reveal the nonsmoothness of the scanned data associated with the correlated, non-Gaussian, noise signature associated with the scanning process. Essentially it shows normal estimation of the raw data, rather than the desired normal of the underlying surface. The raw point clouds suffer from “bumps” which get bigger as the distance from the sensor increases. Further improvement may be obtained by training Nesti-Net on data corrupted with scanner noise and with ground truth normals, but such data is currently not available and is difficult to manually label. Instead, we train Nesti-Net with normal vectors obtained from applying a Total Variation (TV) algorithm on the depth map, provided by Ladicky et al. [23] for the NYU depth V2 dataset. Note that TV substantially smooths fine detail and uses the depth image rather than unstructured point clouds. Column (d) in Figure 7 shows that after training on the TV data, the normal vector estimation of the underlying surface improves significantly. Column (b) shows the results of PCA with a medium scale for reference, for small radius, the result is significantly noisier and for large radius it over-smooths detail, see supplemental material. Note that Nesti-Net performs the estimation on the raw point cloud and does not use the depth image grid structure.




Figure 7. Normal estimation results on scanned point clouds from the ScanNet [8] (top), and NYU Depth V2 dataset [19] (bottom). (a) RGB image, (b) PCA results using a medium scale, (c) Nesti-Net results trained on synthetic data (d) Nesti-net results trained on TV algorithm data.

## 5. Summary

In this work, we propose multi-scale point statistics, a new representation for 3D point clouds that encodes fine and coarse details while using a grid structure. The representation is effectively processed by a CNN architecture (Nesti-Net) for provide accurate normal estimation, which can be used for various applications, *e.g.* surface reconstruction. The mixture-of-experts design of the architecture enables the prediction of an optimal local scale and provides insights into the network’s resource distribution. The proposed representation and architecture achieve state-of-the-art results relative to all other methods and demonstrate robustness to noise and occlusion data corruptions.

## 6. Acknowledgment

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

## References

- [1] P. Alliez, D. Cohen-Steiner, Y. Tong, and M. Desbrun. Voronoi-based variational reconstruction of unoriented point sets. In *Symposium on Geometry Processing*, volume 7, pages 39–48, 2007.
- [2] N. Amenta and M. Bern. Surface reconstruction by voronoi filtering. *Discrete & Computational Geometry*, 22(4):481–504, 1999.
- [3] Y. Ben-Shabat, T. Avraham, M. Lindenbaum, and A. Fischer. Graph based over-segmentation methods for 3d point clouds. *Computer Vision and Image Understanding*, 2018.
- [4] Y. Ben-Shabat, M. Lindenbaum, and A. Fischer. 3dmfv: Three-dimensional point cloud classification in real-time using convolutional neural networks. *IEEE Robotics and Automation Letters*, 3(4):3145–3152, 2018.
- [5] A. Boulch and R. Marlet. Fast and robust normal estimation for point clouds with sharp features. In *Computer Graphics Forum*, volume 31, pages 1765–1774. Wiley Online Library, 2012.
- [6] A. Boulch and R. Marlet. Deep learning for robust normal estimation in unstructured point clouds. *Computer Graphics Forum*, 35(5):281–290, 2016.
- [7] F. Cazals and M. Pouget. Estimating differential quantities using polynomial fitting of osculating jets. *Computer Aided Geometric Design*, 22(2):121–146, 2005.
- [8] A. Dai, A. X. Chang, M. Savva, M. Halber, T. A. Funkhouser, and M. Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, page 10, 2017.
- [9] T. K. Dey and S. Goswami. Provable surface reconstruction from noisy samples. *Computational Geometry*, 35(1-2):124–141, 2006.
- [10] G. Guennebaud and M. Gross. Algebraic point set surfaces. *ACM Transactions on Graphics (TOG)*, 26(3):23, 2007.
- [11] P. Guerrero, Y. Kleiman, M. Ovsjanikov, and N. J. Mitra. Pcpnet learning local shape properties from raw point clouds. *Computer Graphics Forum*, 37(2):75–85, 2018.
- [12] H. Hoppe, T. DeRose, T. Duchampt, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Computer Graphics*, 26:2, 1992.
- [13] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton. Adaptive mixtures of local experts. *Neural Computation*, 3(1):79–87, 1991.
- [14] R. Klokov and V. Lempitsky. Escape from cells: Deep kd-networks for the recognition of 3d point cloud models. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 863–872, Oct 2017.
- [15] J. Krapac, J. Verbeek, and F. Jurie. Modeling spatial layout with fisher vectors for image categorization. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 1487–1494. IEEE, 2011.
- [16] D. Maturana and S. Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 922–928. IEEE, 2015.
- [17] Q. Mérigot, M. Ovsjanikov, and L. J. Guibas. Voronoi-based curvature and feature estimation from point clouds. *IEEE Transactions on Visualization and Computer Graphics*, 17(6):743–756, 2011.
- [18] N. J. Mitra and A. Nguyen. Estimating surface normals in noisy point cloud data. In *Proceedings of the Nineteenth Annual Symposium on Computational geometry*, pages 322–328. ACM, 2003.
- [19] P. K. Nathan Silberman, Derek Hoiem and R. Fergus. Indoor segmentation and support inference from RGBD images. In *European Conference on Computer Vision (ECCV)*, pages 746–760, 2012.



- [20] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [21] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017.
- [22] J. Sánchez, F. Perronnin, T. Mensink, and J. Verbeek. Image classification with the fisher vector: Theory and practice. *International Journal of Computer Vision*, 105(3):222–245, 2013.
- [23] B. Zeisl, M. Pollefeys, et al. Discriminatively trained dense surface normal estimation. In *European conference on computer vision*, pages 468–484. Springer, 2014.