This CVPR 2020 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Deep Iterative Surface Normal Estimation

Jan Eric Lenssen^{1,2,*} janeric.lenssen@udo.edu

Christian Osendorfer¹ christian@nnaisense.com Jonathan Masci¹ jonathan@nnaisense.com

¹ NNAISENSE ² TU Dortmund University

Abstract

This paper presents an end-to-end differentiable algorithm for robust and detail-preserving surface normal estimation on unstructured point-clouds. We utilize graph neural networks to iteratively parameterize an adaptive anisotropic kernel that produces point weights for weighted least-squares plane fitting in local neighborhoods. The approach retains the interpretability and efficiency of traditional sequential plane fitting while benefiting from adaptation to data set statistics through deep learning. This results in a state-of-the-art surface normal estimator that is robust to noise, outliers and point density variation, preserves sharp features through anisotropic kernels and equivariance through a local quaternion-based spatial transformer. Contrary to previous deep learning methods, the proposed approach does not require any hand-crafted features or preprocessing. It improves on the state-of-the-art results while being more than two orders of magnitude faster and more parameter efficient.

1. Introduction

Normal vectors are local surface descriptors that are used as an input for several computer vision tasks ranging from surface reconstruction [27] to registration [39] and segmentation [17]. For this reason, the task of surface normal estimation has been an important and well studied research topic for a long time, with several methods dating back up to 30 years [23]. Progress in the field, however, has been plateauing only until recently when a number of works has shown that improvements can be achieved with the use of data-driven deep learning techniques [5, 7, 19], as also shown in related fields like point cloud denoising [42] or finding correspondences on meshes and point clouds [10, 14, 32, 35]. Deep learning methods are known to often achieve better results compared to data-independent methods. However, they have downsides in terms of robustness



Figure 1: Simplified overview of the proposed method for deep iterative surface normal estimation. The figure shows the process for a subset of three points. (a) Surfaces are fitted by optimizing weighted least squares. (b) A graph neural network infers kernel parameters and local orientation from intermediate pair-wise point descriptors. (c) A trainable, adaptive kernel refines the weights for the next step of the least squares optimization.

to small input changes, adversarial attacks, interpretability, and sometimes also computational efficiency. Also, they do not make use of often well-known instrinsic problem structure, which leads to the necessity of having a large amount of training data and model parameters to learn that structure on their own.

It is well-known that surface normal estimation can be formulated as a least-squares optimization problem. A way to utilize this problem-specific knowledge with deep learning is to take an iteratively reweighting least squares (IRLS) scheme [22] for robust model fitting and modify it using deep data-dependent weighting, as it has been done recently (with or without iterations) for other tasks [25, 43, 44, 46].

^{*}Work performed during an internship at NNAISENSE.

It is a promising candidate to combine robustness, interpretability and efficiency with the data prior of deep neural networks (DNNs). From a deep learning perspective, the approach imposes a strong bias on the architecture, heavily constraining the space of solutions to those which are better suited for the given problem.

Contribution In this work, we present such a trainable reweighting procedure for input graphs with a large number of weighted least square problems and use it to design a fast and accurate algorithm for surface normal estimation on unstructured point clouds (c.f. Figure 1). The method consists of a light-weight graph neural network (GNN), which parameterizes a local quaternion transformer and a deep kernel function to iteratively re-weight graph edges in a largescale point neighborhood graph. We show that the resulting algorithm

- reaches state-of-the-art performance in surface normal estimation on unstructured point clouds,
- is more than two orders of magnitude faster and more parameter efficient than related deep learning approaches, and
- is robust to noise and point density variation, while being equivariant and able to preserve sharp features.

2. Related work

Traditional methods for surface normal estimation make use of plane fitting approaches like unweighted principal component analysis (PCA) [23] and singular value decomposition (SVD) (c.f. [28] for an overview). The performance of these approaches usually hinges upon the often cumbersome selection of data-specific hyper-parameters, such as point neighborhood sizes, and it is sensitive to noise, outliers and density variations. Because of this, several heuristics have been proposed to ease such selection, e.g. those for finding a neighborhood size for plane fitting [34]. Another limitation of plane fitting methods is that they tend to smoothen sharp details, in fact they can be seen as isotropic low-pass filters. In order to preserve sharp features methods that extract normal vectors from estimated Voronoi cells have been proposed [2, 33] and combined with PCA [1]. Alternative approaches include edgeaware sampling [24] or normal vector estimation in Hough space [6]. In addition, several methods arise from more complex surface reconstruction techniques, e.g. moving least squares (MLS) [30], spherical fitting [18], jet fitting [9] and multi-scale kernel methods [3].

Deep learning methods. Deep learning based approaches also found their way into surface normal estimation with

the recent success of deep learning in a wide range of domains. These approaches can be divided into two groups, depending on the actual type of input data they use. The first group aims at normal estimation from single images [4, 12, 15, 29, 31, 41, 45] and has received a lot of interest over the last few years due to the well understood properties of CNNs for grid-structured data.

The second line of research directly uses unstructured point clouds and emerged only very recently, partially due to the advent of graph neural networks and geometric deep learning [8]. Boulch et al. [7] proposed to use a CNN on Hough transformed point clouds in order to find surface planes of the point cloud in Hough space. Based on the recently introduced point processing network, PointNet [40], Guerrero et al. [19] proposed a deep multi-scale architecture for surface normal estimation. Later, Ben-Shabat et al. [5] improved on those results using 3D point cloud fisher vectors as input features and a three-dimensional CNN architecture consisting of multiple expert networks.

3. Problem and background

Let S be a manifold in \mathbb{R}^3 , $\mathcal{P} = {\mathbf{p}_0, ..., \mathbf{p}_m}$ a finite set of sampled and possibly distorted points from that manifold and $\hat{\mathbf{N}} = {\hat{\mathbf{n}}_0, ..., \hat{\mathbf{n}}_m}$ the tangent plane normal vectors at sample points \mathbf{p}_i . Surface normal estimation for the point cloud \mathcal{P} can be described as the problem of estimating a set of normal vectors $\mathbf{N} = {\mathbf{n}_0, ..., \mathbf{n}_m}$ given \mathcal{P} , whose direction match those of the actual surface normals $\hat{\mathbf{n}}_i$ as close as possible. We consider the problem of unoriented normal estimation, determining the normal vectors up to a sign flip. Estimating the correct sign can be done in a post-processing step, depending on the task at hand, and is explicitly tackled by several works [37, 25, 47].

A standard approach to determine unoriented surface normals is fitting planes to the local neighborhood of every point \mathbf{p}_i [30]. Given a radius r or a neighborhood size k, we model the input as a nearest neighbor graph $G = (\mathcal{P}, \mathcal{E})$, where we have a directed edge $(i, j) \in \mathcal{E}$ if and only if $||\mathbf{p}_j - \mathbf{p}_j||_2 < r$ or if \mathbf{p}_j is one of the k nearest neighbors of \mathbf{p}_i , respectively. Let $\mathcal{N}(i)$ denote the local neighborhood of \mathbf{p}_i , with $k_i \equiv |\mathcal{N}(i)|$, containing all \mathbf{p}_j with $(i, j) \in \mathcal{E}$. Furthermore, let $\mathcal{P}(i) \in \mathbb{R}^{k_i \times 3}$ be the matrix of centered coordinates of the points from this neighborhood, that is

$$\mathcal{P}(i)_j = \mathbf{p}_j^\top - \frac{1}{k_i} \sum_{m \in \mathcal{N}(i)} \mathbf{p}_m^\top, \qquad \mathbf{p}_j \in \mathcal{N}(i).$$
(1)

Fitting a plane to this neighborhood is then described as finding the least squares solution of a homogeneous system of linear equations:

$$\mathbf{n}_{i}^{*} = \underset{\mathbf{n}:|\mathbf{n}|=1}{\arg\min} ||\mathcal{P}(i)\mathbf{n}||_{2}^{2} = \underset{\mathbf{n}:|\mathbf{n}|=1}{\arg\min} \sum_{j \in \mathcal{N}(i)} ||\mathcal{P}(i)_{j} \cdot \mathbf{n}||^{2}$$
(2)

The simple plane fitting of Eq. 2 is not robust and does not result in high-quality normal vectors: It produces accurate results only if there are no outliers in the data, which is never the case in practice. Additionally, this approach eliminates sharp details because it acts as a low-pass filter on the point cloud. Even when an isotropic radial kernel function $\theta(||\mathcal{P}(i)||)$ is used to weight points according to their distance to the local mean, fine details cannot be preserved.

Both problems can be resolved through integrating weighting functions into Eq. 2. Sharp features can be preserved with an anisotropic kernel that infers weights of point pairs based on their relative positions, i.e.:

$$\mathbf{n}_{i}^{*} = \operatorname*{arg\,min}_{\mathbf{n}:|\mathbf{n}|=1} \sum_{j \in \mathcal{N}(i)} \psi(\mathbf{p}_{j} - \mathbf{p}_{i}) \cdot ||\mathcal{P}(i)_{j} \cdot \mathbf{n}||^{2} \qquad (3)$$

where $\psi(\cdot)$ is an anisotropic kernel, considering the full Cartesian relationship between neighboring points, instead of only their distance. However, an anisotrop kernel is no longer rotation invariant, so that equivariance of output normals needs to be ensured additionally. Robustness to outliers can be achieved by another kernel that weights points according to an inlier score $s_{i,j}$. More specifically, Eq. 2 is changed to

$$\mathbf{n}_{i}^{*} = \operatorname*{arg\,min}_{\mathbf{n}:|\mathbf{n}|=1} \sum_{j \in \mathcal{N}(i)} s_{i,j} \cdot ||\mathcal{P}(i)_{j} \cdot \mathbf{n}||^{2}, \qquad (4)$$

where $s_{i,j}$ weights outliers with a low and inliers with a high score. However, in order to infer information about the outlier status of points an initial model estimation is necessary. A standard solution to this circular dependency is to formulate the problem as a sequence of weighted leastsquares problems [22, 43]. Given the residuals \mathbf{r}^{l} of the least squares solution from iteration l, the solution for iteration l + 1 is computed as

$$\mathbf{n}_{i}^{l+1} = \operatorname*{arg\,min}_{\mathbf{n}:|\mathbf{n}|=1} \sum_{j\in\mathcal{N}(i)} s(\mathbf{r}_{i,j}^{l}) \cdot ||\mathcal{P}(i)_{j}\cdot\mathbf{n}||^{2}.$$
 (5)

That is, the inlier score and the estimated model are refined in an alternating fashion.

4. Deep iterative surface normal estimation

In this section we present our method, which combines the described properties of robustness, anisotropy and equivariance with the deep learning property of adaptation to large data set statistics. In contrast to existing deep learning methods [5, 19], we do not directly regress normal vectors from point features but weights for a least-squares optimization step, utilizing the problem specific knowledge outlined above.

The core of the algorithm is a trainable kernel function $\psi: \mathbb{R}^3 \times \mathbb{R}^d \to \mathbb{R}$, which computes weights as

$$w_{i,j} = \psi(\mathbf{R}_i(\mathbf{p}_j - \mathbf{p}_i), \theta_i), \tag{6}$$

Alg

 $\mathbf{R} \leftarrow \text{QuatsToMats}(\mathbf{Q})$

 $\mathbf{U}, \boldsymbol{\Sigma} \leftarrow \text{ParallelEig}(\mathbf{C})$ $\mathbf{N}^l \leftarrow \text{Extract Solutions from } \mathbf{U}$

end for return \mathbf{N}^L

 $\mathbf{W} \leftarrow \text{ApplyKernel } \psi(\mathbf{R}, \mathcal{P}, \Theta, \mathcal{E})$

 $\mathbf{C} \leftarrow \text{WeightedCovMatrices}(\mathcal{P}, \mathbf{W}, \mathcal{E})$

where θ are kernel parameters and **R** is a rotation matrix. The kernel is shared by all local neighborhoods of the point graph while θ and \mathbf{R} are individual for each node. Because there is no apriori information about the structure of the input data, a reasonable approach is to model ψ as an MLP and to find kernel parameters through supervised learning from data. To this end, parameters θ and poses **R** for each neighborhood are jointly regressed by a graph neural network on the point neighborhood graph. Then, the kernel function ψ regresses anisotropic, equivariant weights $w_{i,i}$ for each edge in the graph, which are used to find the normal vectors using traditional weighted least-squares optimization

$$\mathbf{n}_{i} = \operatorname*{arg\,min}_{\mathbf{n}:|\mathbf{n}|=1} \sum_{j \in \mathcal{N}(i)} \operatorname{softmax}_{j \in \mathcal{N}(i)} (w_{i,j}) ||\mathcal{P}(i)_{j} \cdot \mathbf{n}||^{2}, \quad (7)$$

in parallel for all $i \in \mathcal{P}$. Similar to iterative re-weighting least squares (c.f. Eq. 5), we apply the method in an iterative fashion to achieve robustness and provide the residuals of the previous solution as input to the graph neural network.

The core algorithm is formulated as pseudo code in Algorithm 1. The initial weighting of the points in a neighborhood is chosen to be uniform, which results in unweighted least-squares plane fitting in the initial iteration. In the following, we present the graph neural network, the local quaternion rotation and our differentiable least square solver in more detail.

4.1. GNN for parameterization and rotation

For regressing parameters θ and rotations **R** for the whole point cloud, graph neural networks [13, 20] are a natural fit because the network must be invariant to the ordering of the points in a neighborhood and it must be able to allow weight sharing over neighborhoods with varying cardinality.

Our graph neural network architecture consists of a neighborhood aggregation procedure, which is applied three consecutive times. Given MLPs h and γ , the neighborhood aggregation scheme, similar to that of PointNet [40] and to general message passing graph neural network frameworks [20, 36], is given by message function

$$\mathbf{f}_{e}(i,j) = h\big(\mathbf{f}(i) \,|\, \mathbf{d}_{i,j} \,|\, \mathbf{prf}(i,j)\big),\tag{8}$$

and node update function

$$\mathbf{f}(i) = \gamma \Big(\frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \mathbf{f}_e(i, j) \Big), \tag{9}$$

with | denoting feature concatenation. Using this scheme, we alternate between computing new edge features $\mathbf{f}_e(i, j)$ and node features $\mathbf{f}(i)$. In addition to the Cartesian relation vector $\mathbf{d}_{i,j} = (\mathbf{p}_j - \mathbf{p}_i)$, pair-wise residual features, a modified version of *Point Pair Features (PPF)* [10, 11], are provided as input:

$$\mathbf{prf}(i,j) = (|\mathbf{n}_i \cdot \mathbf{d}_{i,j}|, |\mathbf{n}_j \cdot \mathbf{d}_{i,j}|, |\mathbf{n}_i \cdot \mathbf{n}_j|, ||\mathbf{d}_{i,j}||_2^2).$$
(10)

They are computed directly from the last set of least-squares solutions and contain the residuals as point-plane distances $|\mathbf{n}_i \cdot \mathbf{d}_{i,j}|$.

After applying the message passing scheme, the output node feature matrix $\mathbf{F} \in \mathbb{R}^{N \times (d+4)}$ is interpreted as a tuple $(\Theta \in \mathbb{R}^{N \times d}, \mathbf{Q} \in \mathbb{R}^{N \times 4})$, containing kernel and rotation parameters for all nodes. We use the row-normalized \mathbf{Q} as unit quaternions to efficiently parameterize the rotation group SO(3). We found that using a rotation matrix instead of an arbitrary 3×3 matrix (as in the *Spatial Transformer Network* [26]) heavily improves training stability, as also observed by Guerrero et al. [19]. By applying a custom, differentiable map from quaternion space to the space of rotation matrices we efficiently compute the local rotation matrices \mathbf{R} for all nodes in parallel.

All in all, the graph neural network is permutation invariant, can be efficiently applied in parallel on varying neighborhood sizes, and is a local operator. Locality is an advantage which allows the algorithm to be applied on partial point clouds and scans, without relying on global features or semantics.

4.2. Parallel differentiable least-squares

In every iteration of the presented algorithm, the plane fitting problem of Eq. 7 needs to be solved. A standard approach is to utilize the Singular Value Decomposition of the weighted matrix $\operatorname{diag}(\sqrt{\mathbf{w}_i^l})\mathcal{P}(i)$: Let $\mathbf{U}\Sigma\mathbf{V}^{\mathbf{T}}$ be its decomposition, then the column vector of \mathbf{V} corresponding to the smallest singular value is the optimal solution for the given least squares problem [21, 43]. However, *n* SVDs (for potentially varying matrix sizes) need to be solved in our scenario, one for every neighborhood, which makes this approach prohibitive. A much more efficient approach in this case is to consider the eigendecomposition of the weighted 3×3 covariance matrix $\mathbf{C}(i) = \mathcal{P}(i)^{\top} \operatorname{diag}(\mathbf{w}_i^l)\mathcal{P}(i)$ which has the columns of \mathbf{V} as its eigenvectors [21]. The solution for Eq. 7 is then the eigenvector associated with the smallest eigenvalue. The computational complexity for the eigendecomposition of this 3×3 matrix is O(1) and hence for one overall iteration O(n).

Our algorithm is trained end-to-end by minimizing the distance between ground truth normals and the least squares solution, requiring backpropagation through the eigendecomposition. We follow the work of Giles [16]: Given partial derivatives $\partial L/\partial U$ and $\partial L/\partial \Sigma$ for eigenvectors and eigenvalues, respectively, we compute the partial derivatives for a real symmetric 3×3 covariance matrix C as

$$\frac{\partial L}{\partial \mathbf{C}} = \mathbf{U} \Big(\Big(\frac{\partial L}{\partial \boldsymbol{\Sigma}} \Big)_{diag} + \mathbf{F} \circ \mathbf{U}^{\top} \frac{\partial L}{\partial \mathbf{U}} \Big) \mathbf{U}^{\top}, \quad (11)$$

where $\mathbf{F}_{i,j} = (\lambda_j - \lambda_i)^{-1}$ contains inverse eigenvalue differences. We implemented forward and backward steps for eigendecomposition of a large number of symmetric 3×3 matrices, where we parallelize over graph nodes, leading to an O(1) implementation (using O(n) processors) of parallel least squares solvers.

Handling numerical instability. Backpropagation through the eigendecomposition can lead to numerical instabilities due to at least two reasons: 1) Low-rank input matrices with two or more zero eigenvalues. 2) Exploding gradients when two eigenvalues are very close to each other and values of \mathbf{F} go to infinity. We apply two tricks to avoid these problems. First, a small amount of noise is added to the diagonal elements of all covariance matrices, making them full-rank. Second, gradients are clipped after the backward step on very large values, to tackle the cases of nearly equal eigenvalues that lead to exploding gradients.

4.3. Training

Training is performed by minimizing the Euclidean distance between estimated normals N and ground truth normals \hat{N} , averaged over all normal vectors in the training set:

$$L(\hat{\mathbf{N}}, \mathbf{N}) = \frac{1}{n} \sum_{i=1}^{n} \min(||\hat{\mathbf{n}}_{i} - \mathbf{n}_{i}||_{2}, ||\hat{\mathbf{n}}_{i} + \mathbf{n}_{i}||_{2}), \quad (12)$$

where the minimum of the distances to the flipped or nonflipped ground truth vectors is used. While we also experi-

	Ours $(k = 64, L = 4)$	Nesti-Net [5]	PCPNet [19]	HoughCNN [7]	PCA	Jet [9]
No noise	6.72	6.99	9.68	10.23	12.29	12.23
Noise ($\sigma = 0.00125$)	9.95	10.11	11.46	11.62	12.87	12.84
Noise ($\sigma = 0.006$)	17.18	17.63	18.26	22.66	18.38	18.33
Noise ($\sigma = 0.012$)	21.96	22.28	22.8	33.39	27.5	27.68
Varying Density (Stripes)	7.73	8.47	11.74	12.47	13.66	13.39
Varying Density (Gradients)	7.51	9.00	13.42	11.02	12.81	13.13
Average	11.84	12.41	14.56	16.9	16.25	16.29

Table 1: Results for unoriented normal estimation. Shown are normal estimation errors in angle RMSE. For PCA and Jet, optimal neighborhood size for average error is chosen. For our approach, we display results for a balanced neighborhood size k = 64, which improves on the state of the art for all noise levels. Results for different k are shown in Table 2.

mented with different angular losses, we found that the Euclidean distance loss still provides the best result and the most stable training. A loss is computed after each least squares step and the network is trained iteratively by performing a gradient descent step after each iteration of the algorithm. This fights vanishing gradients that occur due to the normalization of vectors in quaternion and eigenvector computations. The weights of our network are shared over iterations, allowing generalization to further iterations.

5. Experiments

Experiments were conducted to compare the proposed Differentiable Iterative Surface Normal Estimation with state-of-the-art methods both quantitatively, measuring normal estimation accuracy, and qualitatively, on a Poisson reconstruction and on a transfer learning task. Section 5.1 introduces the dataset used to train our model whereas Section 5.2 details the architecture and the protocol followed in our experiments. Then, qualitative (Section 5.3) and quantitative (Section 5.5) results are presented and an analysis of complexity and execution time (Section 5.4) is given.

5.1. PCPNet dataset

Our method is trained and validated quantitatively on the PCPNet dataset as provided by Guerrero et al. [19]. It consists of a mixture of high-resolution scans, point clouds sampled from handmade mesh surfaces and differentiable surfaces. Each point cloud consists of 100k points. We reproduce the experimental setup of [5, 19], training on the provided split containing 32 point clouds under different levels of noise. The test set consists of six categories, containing four sets with different levels of noise (no noise, $\sigma = 0.00125, \sigma = 0.0065$ and $\sigma = 0.012$) and two sets with different sampling density (striped pattern and gradient pattern). We evaluate unoriented normal estimation, same as the related approaches. The Root Mean Squared Error (RMSE) on the provided 5k points subset is used as performance metric following the protocol of related work, where the RMSE is first computed for each test point cloud before the results are averaged over all point clouds in one category. Model selection is performed using the provided validation set.

5.2. Experimental setup and architecture

The presented graph neural network was implemented using the *Pytorch Geometric* library [13]. The neural networks h, γ and ψ each consist of two linear layers, with ReLU non-linearity. A detailed description of the architecture is presented in the supplemental materials. During training, output weights from the kernel are randomly set to zero with probability of 0.25.

It should be noted that despite inheriting the neighborhood size parameter from traditional PCA, it is possible for a network trained on a specific k to be applied for other kas well. This is because all networks can be shared across an arbitrary number of points and the softmax function normalizes weights for neighborhoods of varying sizes. We observed that generalization across different k only leads to a very small increase in average error. However, to fairly evaluate our method for different k, a network is trained for each $k \in \{32, 48, 64, 96, 128\}$. Trained consists of 300 epochs using the RMSProp optimization method. All reported test results are given after 4 re-weighting iterations of our algorithm. Iterating longer does not show significant improvements. Quantitative results over iterations, results for extrapolation over iterations and generalization between different k are presented in the supplemental materials. For further realization details, we refer to our implementation, which we will make available online prior to publication.

5.3. Quantitative Evaluation

RMSE results for the PCPNet test set of our approach (with k = 64) and related works are shown in Table 1. We improve on the state of the art on all noise levels and varying densities. While the improvement is only small, it should be noted that we reach it while being orders of magnitude faster and more parameter efficient (c.f. Section 5.4, which is of importance for many applications in resource constraint environments. For the non deep learn-

		C	Durs $L =$	4				PCA		
Neighborhood size k	32	48	64	96	128	32	48	64	96	128
No noise	6.09	6.63	6.72	6.82	7.35	9.10	9.94	10.68	11.93	12.54
Noise ($\sigma = 0.00125$)	10.22	9.63	9.95	10.45	9.64	11.22	11.56	12.08	12.71	12.97
Noise ($\sigma = 0.006$)	18.17	17.36	17.18	17.03	16.90	28.41	23.00	20.68	18.81	18.12
Noise ($\sigma = 0.012$)	25.17	22.40	21.96	21.80	22.13	45.35	38.48	33.67	28.81	26.67
Varying Density (Stripes)	7.22	7.63	7.73	7.87	8.67	10.48	11.40	12.07	13.18	14.07
Varying Density (Gradients)	6.84	7.19	7.51	7.69	8.49	9.96	10.74	11.35	12.36	13.21
Average	12.28	11.81	11.84	11.94	12.20	19.09	17.52	16.75	16.30	16.26

Table 2: Comparison of unoriented normal estimation RMSE between the proposed method and PCA for different neighborhood sizes k. It can be seen that our method consistently provides lower errors while being significantly more robust to changes of that parameter, compared to PCA.



Figure 2: Comparison for varying angle error threshold. For error thresholds on the x-axis, the y-axis shows the percentage of normals which have an error lower than that threshold. Our method and PCA use neighborhood size k = 64. For low noise settings and varying density, our method succeeds in recovering sharp features, as shown by the higher accuracies for low angle thresholds.

ing approaches, PCA and Jet, results for medium neighborhood sizes are displayed. In addition, results for different k are provided in Table 2 and compared to errors obtained by PCA with the same respective neighborhood size. Our method performs stronger than the PCA baseline in all scenarios. As expected, varying k leads to a behavior similar to that of PCA, with large k's performing better on more noisy data. However, it can be observed that our approach is more robust to changes of k: Even for small neighborhood sizes, high noise is handled significantly better than by PCA and large neighborhoods still produce satisfactory results for low noise data. It should be noted that for all evaluated k we improve on the state of the art w.r.t. average error. An evaluation for smaller k down to k = 2 is provided in the supplemental materials.

While the RMSE error metric is well suited for a general comparison, it is not a good proxy to estimate the ability of recovering sharp features since it does not take into account the error distribution over angles. Therefore, as an additional metric, Figure 2 presents the percentage of angle errors falling below different angle thresholds. The results confirm that our approach is better at preserving details and sharp edges, especially for low noise point clouds and varying density, where it outperforms other approaches. For higher noise, results similar to Nesti-Net are achieved.

5.4. Efficiency

Our model is small, consisting of only 7981 trainable parameters, shared over iterations and spatial locations. On a single Nvidia Titan Xp, a point cloud with 100k points is processed in 5.67 seconds (0.0567 ms per point). A large part of this is the kd-tree used to compute the nearest neighbor graph, which takes 2.1 seconds of the 5.67 seconds. It is run on the CPU and could be further sped-up by utilizing GPUs.

	Ours	Nesti-Net [5]	PCPNet [19]
Num. parameters Exec. time, 100k p.	7981 3.57 s	179M 1350 s 378 ×	22M 470 s

Table 3: Comparison of efficiency between the approaches using deep learning. We list number of model parameters as well as average execution times for estimating normals on a point cloud with 100k points.

In Table 3 we compare our approach against the related deep learning approaches Nesti-Net and PCPNet. Our approach is orders of magnitude $(378 \times \text{ and } 131 \times)$ faster than the related approaches. The comparison was made



Figure 3: Qualitative comparison between our method (k = 64, L = 4) and related work. We show diverse examples from the test set, sampled from different categories, noise levels and density variations. The color encodes the angle error of estimated normals in degrees. Best viewed in the digital version of the paper.



Figure 4: Local behaviour of our method over several iterations for a sharp edge (a) and a noisy surface (b). The partial point clouds where sampled from the PCPNet test dataset. The colors in the first rows show the weights from the kernel network for one normal in the neighborhood while the colors in the second row show the angle error of all neighborhood normals.

as fair as possible by excluding nearest neighbor queries (note that this favors the other approaches since they need larger neighborhoods) and the original implementations. The speedup of our method can be contributed to the much smaller network size and the parallel design of the GNN and least-squares optimization steps.

5.5. Qualitative Evaluation

This section visually presents surface normal errors for various elements of the PCPNet test set in Figure 3 and compares them against results from the PCA baseline and related deep learning approaches. It can be seen that the biggest improvements are obtained for low noise scenarios and varying density, where our method is able to preserve sharp features of objects better than the other methods. In general it can be observed that our approach tends to provide sharp, discriminative normals for points on edges instead of smooth averages. In rare cases, this can lead to a false assignment of points to planes, as we can see in the example in column 8. It can be observed that, in contrast to Nesti-Net, our approach behaves equivariant to input rotation as is seen clearly on the diagonal edge of the box example in column 3. Sharp edges are kept also in uncommon rotations, which we can attribute to our local rotational transformer. Results for more examples are displayed in the supplemental material.

Interpretability. In order to interprete the results of our method, Figure 4 shows a detailed view of local neighborhoods over several iterations of our algorithm. An example for a sharp edge is shown in Figure 4a and a high noise surface in Figure 4b. Both sets of points were sampled from the real test data. For the sharp edge, the algorithm initially fits a plane with uniform weights, leading to smoothed normals. Over the iterations, high weights concentrate on the more plausible plane, leading to recovering of the sharp edge. In the noisy example, we can see that outliers are iteratively receiving lower weights, leading to more stable estimates.



Figure 5: Selected results after applying Poisson surface reconstruction using the estimated normal vectors. In most cases, differences between the methods are very small. Examples 2 and 3 show reconstructions from point clouds with varying density, which show the largest differences.

Surface reconstruction. To further evaluate the quality of the produced normals when used as input to other computer vision pipelines, Figure 5 shows the results for Poisson surface reconstruction. Since the methods in this comparison all perform unoriented normal estimation (Guerrero et al. [19] evaluates both, unoriented and oriented, where we chose the unoriented version for a fair comparison), we determine the signs of the output normals from all four methods using the ground truth normals. Most of the reconstructions show only small differences, with our approach and Nesti-Net retaining slightly more details than the others. Significant differences can be observed for point clouds with varying density, displayed in rows 2 and 3. Here, our approach successfully retains the original structure of the object while still providing sharp edges.

Transfer to NYU depth dataset. In order to show generality of our approach, our models trained on the PCPNet dataset are validated on the NYU depth v2 dataset [38], a common benchmark dataset in the field of estimating normals from single images. It contains 1449 aligned and preprocessed RGBD frames, which are transformed to a point cloud before applying our method. After performing unoriented estimation, the normals are flipped towards the camera position. Evaluation is done qualitatively, since the dataset does not contain ground truth normal vectors. Results for three different neighborhood sizes in comparison to PCA are shown in Figure 6. Our approach behaves as expected, as it is able to infer plausible normals for the given scenes. For all k, our approach is able to preserve sharp fea-



Figure 6: Examples for normal estimation on scanned data from the NYU depth v2 dataset. Colors encode the orientation of normals. Our model generalizes to this dataset while being able to retain more details and sharper edges than PCA. However, scanning artifacts are also kept and visible. Best viewed in the digital version of the paper.

tures while PCA produces very smooth results. However, this also leads to the sharp extraction of scanning artifacts, which can be seen on the walls of the scanned room.

6. Conclusion and future work

We presented a novel method for deep surface normal estimation on unstructured point clouds, consisting of parallel, differentiable least-squares optimization and deep reweighting. In each iterations, the weights are computed using a kernel function that is individually parameterized and rotated for each neighborhood by a task-specific graph neural network. The algorithm is much more efficient than previous deep learning methods, reaches state-of-the-art accuracy, and has favorable properties like equivariance and robustness to noise. For future work, investigating the possibility of utilizing deep data priors to parameterize leastsquares problems holds large potential. We suspect that introducing data-dependency to other traditional methods can lead to progress in other fields of research, by reducing common disadvantages of pure deep learning approaches. On the theoretical side, it is interesting to dive deeper into convergence properties of IRLS with deep re-weighting.

Acknowledgements

We thank Matthias Fey for his work on Pytorch Geometric. We also thank him and Prof. Dr. Heinrich Müller for helpful advice and discussions.

References

- P. Alliez, D. Cohen-Steiner, Y. Tong, and M. Desbrun. Voronoi-based variational reconstruction of unoriented point sets. In *Proceedings of the Fifth Eurographics Symposium* on Geometry Processing, SGP '07, pages 39–48, 2007.
- [2] Nina Amenta and Marshall Bern. Surface reconstruction by voronoi filtering. In *Proceedings of the Fourteenth Annual Symposium on Computational Geometry*, SCG '98, pages 39–48, 1998.
- [3] Samir Aroudj, Patrick Seemann, Fabian Langguth, Stefan Guthe, and Michael Goesele. Visibility-consistent thin surface reconstruction using multi-scale kernels. ACM Transaction on Graphics, 36(6):187:1–187:13, 2017.
- [4] Aayush Bansal, Bryan C. Russell, and Abhinav Gupta. Marr revisited: 2d-3d alignment via surface normal prediction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5965–5974, 2016.
- [5] Yizhak Ben-Shabat, Michael Lindenbaum, and Anath Fischer. Nesti-net: Normal estimation for unstructured 3d point clouds using convolutional neural networks. *CoRR*, abs/1812.00709, 2018.
- [6] Alexandre Boulch and Renaud Marlet. Fast and Robust Normal Estimation for Point Clouds with Sharp Features. *Computer Graphics Forum*, 2012.
- [7] Alexandre Boulch and Renaud Marlet. Deep Learning for Robust Normal Estimation in Unstructured Point Clouds. *Computer Graphics Forum*, 2016.
- [8] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18– 42, July 2017.
- [9] F. Cazals and M. Pouget. Estimating differential quantities using polynomial fitting of osculating jets. In *Proceedings* of the 2003 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing, pages 177–187, 2003.
- [10] Haowen Deng, Tolga Birdal, and Slobodan Ilic. Ppf-foldnet: Unsupervised learning of rotation invariant 3d local descriptors. In *European Conference of Computer Vision (ECCV)*, pages 620–638, 2018.
- [11] Haowen Deng, Tolga Birdal, and Slobodan Ilic. Ppfnet: Global context aware local features for robust 3d point matching. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 195–205, 2018.
- [12] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*, pages 2650–2658, 2015.
- [13] Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. *CoRR*, abs/1903.02428, 2019.
- [14] Matthias Fey, Jan Eric Lenssen, Frank Weichert, and Heinrich Müller. Splinecnn: Fast geometric deep learning with continuous b-spline kernels. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 869– 877, 2018.

- [15] David F. Fouhey, Abhinav Gupta, and Martial Hebert. Datadriven 3D primitives for single image understanding. In *International Conference on Computer Vision (ICCV)*, 2013.
- [16] Mike B. Giles. Collected matrix derivative results for forward and reverse mode algorithmic differentiation. In Christian H. Bischof, H. Martin Bücker, Paul Hovland, Uwe Naumann, and Jean Utke, editors, *Advances in Automatic Differentiation*, pages 35–44. Springer Berlin Heidelberg, 2008.
- [17] Eleonora Grilli, Fabio Menna, and Fabio Remondino. A review of point clouds segmentation and classification algorithms. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-2/W3:339–344, 02 2017.
- [18] Gaël Guennebaud and Markus Gross. Algebraic point set surfaces. ACM Transactions on Graphics, 26(3), 2007.
- [19] Paul Guerrero, Yanir Kleiman, Maks Ovsjanikov, and Niloy J. Mitra. Pepnet learning local shape properties from raw point clouds. *Computer Graphics Forum*, 37(2):75–85, 2018.
- [20] William L. Hamilton, Zhitao Ying, and Jure Leskovec. Representation learning on graphs: Methods and applications. *IEEE Data Eng. Bull.*, 40:52–74, 2017.
- [21] Richard Hartley and Andrew Zisserman. Multiple View Geometry in Computer Vision. Cambridge University Press, New York, NY, USA, 2 edition, 2003.
- [22] Paul W. Holland and Roy E. Welsch. Robust regression using iteratively reweighted least-squares. *Communications in Statistics - Theory and Methods*, 6(9):813–827, 1977.
- [23] Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Surface reconstruction from unorganized points. In *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '92, pages 71–78, 1992.
- [24] Hui Huang, Shihao Wu, Minglun Gong, Daniel Cohen-Or, Uri Ascher, and Hao (Richard) Zhang. Edge-aware point set resampling. ACM Trans. Graph., 32(1):9:1–9:12, Feb. 2013.
- [25] Xiangru Huang, Zhenxiao Liang, Xiaowei Zhou, Yao Xie, Leonidas J. Guibas, and Qixing Huang. Learning transformation synchronization. *CoRR*, abs/1901.09458, 2019.
- [26] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. In Proceedings of the 28th International Conference on Neural Information Processing Systems (NeurIPS), pages 2017–2025, 2015.
- [27] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the Fourth Eurographics Symposium on Geometry Processing*, SGP '06, pages 61–70, 2006.
- [28] K. Klasing, D. Althoff, D. Wollherr, and M. Buss. Comparison of surface normal estimation methods for range sensing applications. In 2009 IEEE International Conference on Robotics and Automation, pages 3206–3211, May 2009.
- [29] Lubor Ladicky, Bernhard Zeisl, and Marc Pollefeys. Discriminatively trained dense surface normal estimation. In *ECCV*, 2014.
- [30] David Levin. The approximation power of moving leastsquares. *Math. Comput.*, 67(224):1517–1531, Oct. 1998.

- [31] Bo Li, Chunhua Shen, Yuchao Dai, A. van den Hengel, and Mingyi He. Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs. In 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 1119–1127, 2015.
- [32] Jonathan Masci, Davide Boscaini, Michael M. Bronstein, and Pierre Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. In *Proceedings of the* 2015 IEEE International Conference on Computer Vision Workshop (ICCVW), pages 832–840, 2015.
- [33] Quentin Merigot, Maks Ovsjanikov, and Leonidas J. Guibas. Voronoi-based curvature and feature estimation from point clouds. *IEEE Transactions on Visualization and Computer Graphics*, 17(6):743–756, 2011.
- [34] N. J. Mitra, A. Nguyen, and L. Guibas. Estimating surface normals in noisy point cloud data. In *special issue of International Journal of Computational Geometry and Applications*, volume 14, pages 261–276, 2004.
- [35] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodolà, Jan Svoboda, and Michael M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5425–5434, 2017.
- [36] Christopher Morris, Martin Ritzert, Matthias Fey, William L. Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In AAAI Conference on Artificial Intelligence (AAAI), 2019.
- [37] Patrick Mullen, Fernando de Goes, Mathieu Desbrun, David Cohen-Steiner, and Pierre Alliez. Signing the unsigned: Robust surface reconstruction from raw pointsets. *Computer Graphics Forum*, 29(5):1733–1741, 2010.
- [38] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *European Conference on Computer Vision* (ECCV), 2012.
- [39] François Pomerleau, Francis Colas, and Roland Siegwart. A review of point cloud registration algorithms for mobile robotics. *Found. Trends Robot*, 4(1):1–104, May 2015.
- [40] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In Advances in Neural Information Processing Systems 30, pages 5099–5108. 2017.
- [41] Xiaojuan Qi, Renjie Liao, Zhengzhe Liu, Raquel Urtasun, and Jiaya Jia. Geonet: Geometric neural network for joint depth and surface normal estimation. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 283–291, 2018.
- [42] Marie-Julie Rakotosaona, Vittorio La Barbera, Paul Guerrero, Niloy J. Mitra, and Maks Ovsjanikov. POINTCLEAN-NET: learning to denoise and remove outliers from dense point clouds. *CoRR*, abs/1901.01060, 2019.
- [43] Rene Ranftl and Vladlen Koltun. Deep fundamental matrix estimation. In *European Conference on Computer Vision* (ECCV), September 2018.

- [44] Supasorn Suwajanakorn, Noah Snavely, Jonathan J. Tompson, and Mohammad Norouzi. Discovery of latent 3d keypoints via end-to-end geometric reasoning. In Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems (NeurIPS), pages 2063–2074, 2018.
- [45] Xiaolong Wang, David F. Fouhey, and Abhinav Gupta. Designing deep networks for surface normal estimation. 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 539–547, 2015.
- [46] Yue Wang and Justin M. Solomon. Deep closest point: Learning representations for point cloud registration. In *International Conference on Computer Vision (ICCV)*, 2019.
- [47] Shihao Wu, Hui Huang, Minglun Gong, Matthias Zwicker, and Daniel Cohen-Or. Deep points consolidation. ACM Transaction on Graphics, 34(6):176:1–176:13, 2015.