

IterativePFN: True Iterative Point Cloud Filtering

Dasith de Silva Edirimuni¹, Xuequan Lu^{1*}, Zhiwen Shao^{2†}, Gang Li¹, Antonio Robles-Kelly^{1,4}, Ying He³

¹School of Information Technology, Deakin University

²School of Computer Science and Technology, China University of Mining and Technology

³School of Computer Science and Engineering, Nanyang Technological University

⁴Defense Science and Technology Group, Australia

{dtdesilva, xuequan.lu, gang.li, antonio.robles-kelly}@deakin.edu.au,
zhiwen_shao@cumt.edu.cn, yhe@ntu.edu.sg

Abstract

The quality of point clouds is often limited by noise introduced during their capture process. Consequently, a fundamental 3D vision task is the removal of noise, known as point cloud filtering or denoising. State-of-the-art learning based methods focus on training neural networks to infer filtered displacements and directly shift noisy points onto the underlying clean surfaces. In high noise conditions, they iterate the filtering process. However, this iterative filtering is only done at test time and is less effective at ensuring points converge quickly onto the clean surfaces. We propose **IterativePFN** (iterative point cloud filtering network), which consists of multiple **IterationModules** that model the true iterative filtering process internally, within a single network. We train our IterativePFN network using a novel loss function that utilizes an adaptive ground truth target at each iteration to capture the relationship between intermediate filtering results during training. This ensures that the filtered results converge faster to the clean surfaces. Our method is able to obtain better performance compared to state-of-the-art methods. The source code can be found at: <https://github.com/ddsediri/IterativePFN>.

1. Introduction

Point clouds are a natural representation of 3D geometric information and have a multitude of applications in the field of 3D Computer Vision. These applications range from robotics and autonomous driving to urban planning [14, 19, 35, 38]. They are captured using 3D sensors and comprise of unordered points lacking connectivity information. Furthermore, the capturing of point cloud data is error-prone as sensor quality and environmental factors may

*Corresponding author: X. Lu, supported by PJ03906.PG00507.F002.

†Z. Shao is supported by the NSFC (No. 62106268).

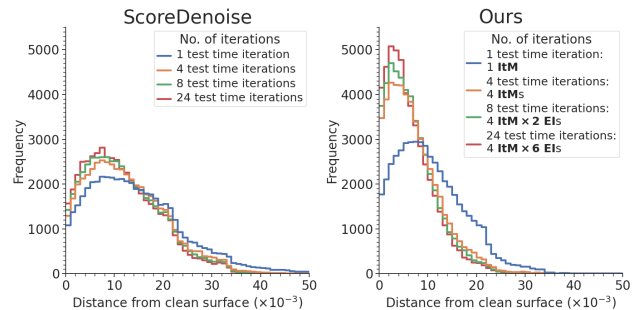


Figure 1. Histograms of filtered point distances from clean surface after 1, 4, 8 and 24 test time iterations for ScoreDenoise [22] on the Casting shape with 50K points and 2.5% Gaussian noise. We compare it with our proposed IterativePFN where 1 IterationModule (ItM) corresponds to 1 internal iteration and 4 ItMs equal 1 external iteration (EI). There are 4 ItMs in the proposed network. Note 1 ItM is analogous to 1 test time iteration of ScoreDenoise. Our filtering results converge closer to the surface.

introduce noisy artifacts. The process of removing noise is a fundamental research problem which motivates the field of point cloud filtering, also known as denoising. Filtering facilitates other tasks such as normal estimation and, by extension, 3D rendering and surface reconstruction.

Conventional point cloud filtering methods such as MLS based methods [2, 12], bilateral filtering mechanisms [9] and edge recovery algorithms [18, 34] rely on local information of point sets, i.e., point normals, to filter point clouds. However, such methods are limited by the accuracy of normals. Alternatives include the Locally Optimal Projection (LOP) family of methods [13, 17, 25], which downsample and regularize point clouds but incur the loss of important geometric details. More recently, deep learning based filtering methods have been proposed to alleviate the disadvantages and limitations of conventional methods [22–24, 28, 41].

Early **deep learning** based filtering methods, such as PointProNets [31], require pre-processed 2D height maps to filter point clouds. However, the advent of PointNet, PointNet++ and DGCNN, made direct point set convolution a possibility [26, 27, 37]. Feature encoders based on these architectures were exploited by recent methods to produce richer latent representations of point set inputs and filter noise more effectively [20, 22, 28, 41]. These methods can be broadly characterized as 1) **resampling**, 2) **probability** and 3) **displacement** based methods. Resampling based methods such as DMRDenoise [20] suffer from the loss of geometric details as the method relies on identifying down-sampled underlying clean surfaces and upsampling along them. ScoreDenoise, which models the gradient-log of the noise-convolved probability to find a point at a given position, iteratively performs Langevin sampling-inspired gradient ascent [22] to filter points. However, filtered points are slow to converge to the clean surface after many test time iterations of filtering, as illustrated in Fig. 1. By contrast, for an IterativePFN network with 4 IterationModules, where 1 iterationModule (ItM) represents 1 internal iteration of filtering and is analogous to 1 test time iteration of ScoreDenoise, we see that a higher number of filtered points converge closer to the clean surface within the same number of test time iterations.

Among displacement based methods, PointCleanNet (PCN) [28] shows sensitivity to high noise while Pointfilter [41] utilizes a bilateral filtering inspired weighted loss function that causes closely separated surfaces to collapse into a single surface during filtering. Moreover, gradient ascent and displacement based methods filter point clouds iteratively during test times and do not consider true iterative filtering during training. Although RePCDNet [7] offers a recurrent neural network inspired alternative to capture this information during training, at each iteration RePCDNet attempts to directly shift points onto the clean surface without considering, that at different iterations, their existing residual noise, in decreasing order w.r.t. iteration number. Furthermore, it uses a single network to filter points, increasing the burden on the network to correctly distinguish between noise scales and requires multiple test time iterations which lead to low efficiency. Based on these major limitations, we propose:

- a novel neural network architecture of stacked encoder-decoder modules, dubbed **IterationModule**, to model the true iterative filtering process internally (see Fig. 2). Each IterationModule represents an iteration of filtering and the output of the τ -th IterationModule becomes the input for the $\tau + 1$ -th IterationModule. Thereby, the $\tau + 1$ -th IterationModule represents the filtering iteration $t = \tau + 1$. This allows the network to develop an understanding of the filtering relationship across iterations.

- a novel loss function that formulates the nearest neighbor loss at each iteration as the L_2 norm minimization between the filtered displacements, inferred by the τ -th IterationModule, and the nearest point within a target point cloud at $t = \tau$, of a lower noise scale σ_τ compared to the noise scale $\sigma_{\tau-1}$ of the target at $t = \tau - 1$. This promotes a gradual filtering process that encourages convergence to the clean surface.
- a generalized patch-stitching method that designs Gaussian weights when determining best filtered points within overlapping patches. Patch stitching improves efficiency as it facilitates filtering multiple points simultaneously.

We conduct comprehensive experiments, in comparison with state-of-the-art methods, which demonstrate our method’s advantages on both synthetic and real world data.

2. Related Work

Traditional methods. Moving Least Squares (MLS) introduced by Levin [15] was extended to tackle point cloud filtering by the work of Alexa et al. [2] with the intent of minimizing the approximation error between filtered surfaces and their ground truth. Thereafter, Adamson and Alexa proposed Implicit Moving Least Squares (IMLS) [1] and Guennebaud and Gross introduced Anisotropic Point Set Surfaces (APSS) [12], where both methods attempted to enhance the filtering performance on point clouds with sharp geometric features. However, these MLS techniques rely on parameters of the local surface, such as point normals. In order to alleviate this, Lipman et al. [17] proposed the Locally Optimal Projection (LOP) method. In addition to being independent of local surface parameters, LOP based methods downsample and regularize the input point clouds. This work spawned a new class of point cloud filtering methods, including Continuous Locally Optimal Projection (CLOP) by Preiner et al. [25] and Weighted Locally Optimal Projection (WLOP) by Huang et al. [13]. In general, both classes of methods suffer from the same disadvantage: they are unable to preserve sharp geometric features of the underlying clean surfaces and are sensitive to noise.

Moreover, Cazals and Pouget proposed fitting n -order polynomial surfaces to point sets [6] to calculate quantities of interest such as point normals and curvature and has seen learning based applications in [4, 43]. Moreover, by projecting noisy points onto the fitted surface, their filtered counterparts can be obtained. Digne introduced a similarity based filtering method that uses local descriptors for each point and, subsequently, exploits the similarity between descriptors to determine filtered displacements at each point [8]. More recently, a generalization of the mesh bilateral filtering mechanism [11] for point clouds, was proposed by Digne and de Franchis that filters points based

on an anisotropic weighting of point neighborhoods. This weighting considers both point positions and their normals. Other methods such as the L_1 and L_0 minimization methods of Avron et al. and Sun, Schaefer, and Wang [3, 34] and the Low Rank Matrix Approximation of Lu et al. [18] estimate normals that are used in filtering algorithms to denoise point clouds. The filtering process has also been reformulated as a sparse optimization problem which can be solved by Augmented Lagrangian Multipliers, in the work of Remil et al. [30].

Deep learning based methods. Among learning based methods, PointProNets was proposed by Roveri et al. [31] and utilized a traditional CNN architecture to filter noisy 2D height maps which were re-projected into 3D space. By contrast, the Deep Feature Preserving (DFP) network of Lu et al. used a traditional CNN with 2D height maps to estimate point normals that can be used in conjunction with the position update algorithm of [18]. With the advent of PointNet and PointNet++ [26, 27], direct convolution of point sets became possible. Furthermore, Wang et al. [37] proposed a graph convolutional architecture which consumed nearest-neighbor graphs generated from point sets to produce rich feature representations of points. The following methods use PointNet or DGCNN inspired backbones for processing point cloud inputs. Yu et al. developed EC-Net to perform denoising by upsampling in an edge-aware manner and, thereby, consolidate points along a surface [40]. A similar approach is adopted by Luo and Hu in their DM-RDenoise mechanism [20] which sequentially identifies an underlying downsampled, less noisy, manifold and subsequently upsamples points along it. PCN, by Rakotosaona et al. [28], filters central points within point cloud patches. PCN infers these filtered displacements by utilizing a loss which minimizes the L_2 norm between predicted displacements and points within the ground truth patch. Furthermore, they add a repulsion term to ensure a regular distribution of points. Zhang et al. proposed Pointfilter [41] that develops this displacement based line of inquiry. They use ground truth normals during training to calculate a bilateral loss. Pistilli et al. proposed GPDNet, a graph convolutional architecture to infer filtered displacements [24]. Recently, Chen et al. proposed RePCD-Net, a recurrent architecture that iteratively filters point clouds during training and considers the training losses at each iteration. However, this method must still be applied iteratively during test time [7].

Langevin sampling of noisy data to iteratively recover less noisy data, at test-time, was initially proposed in the 2D generative modelling field [33] and has been applied to point cloud generation and reconstruction [5, 21]. This relies on learning the unnormalized gradient-log probability distribution of the data. Luo and Hu extended this to point cloud filtering. Their method, ScoreDenoise, models the gradient-log of the underlying noise convolved probability

distribution for point cloud patches [22]. Mao et al. proposed PDFlow to uncover latent representations of noise-free data at higher dimensions by utilizing normalizing flows and, thereafter, obtain the filtered displacements [23].

3. Research Problem and Motivation

As discussed in Sec. 1, previous point cloud filtering methods can be characterized as resampling, probability and displacement based methods, the second of which includes score-based methods. We are motivated by the interplay between filtering objectives of displacement and score-based methods and how a score-based method can inform the construction of a truly iterative displacement based method, that is iterative in its training objective. PCN [28] proposed the idea that the filtering objective should aim to regress a noisy point \mathbf{x}_i back onto the underlying clean surface using a displacement \mathbf{d}_i . This displacement is the output of the regression network, which takes a noisy point cloud patch \mathcal{X} , centered at \mathbf{x}_i , as its input. The filtering objective is expressed as $\tilde{\mathbf{x}}_i = \mathbf{x}_i + \mathbf{d}_i$ where $\tilde{\mathbf{x}}_i$ is the filtered point. During testing, the output at time $t = \tau$ is taken as input for the next iteration at time $t = \tau + 1$. This leads to the test-time iterative filtering objective:

$$\tilde{\mathbf{x}}_i^{(t)} = \tilde{\mathbf{x}}_i^{(t-1)} + \mathbf{d}_i^{(t)}, t = 1, \dots, T \quad (1)$$

PCN’s training objective is motivated by the need to regress the noisy point back to the clean patch \mathcal{Y} , while ensuring it is centered within it. Thus, the PCN training objective is achieved by the following loss:

$$L_i^{PCN} = \alpha \min_{\mathbf{x}_j \in \mathcal{Y}} \|\mathbf{d}_i - \delta \mathbf{x}_j\|_2^2 + (1 - \alpha) \max_{\mathbf{x}_j \in \mathcal{Y}} \|\mathbf{d}_i - \delta \mathbf{x}_j\|_2^2, \quad (2)$$

where $\delta \mathbf{x}_j = \mathbf{x}_j - \mathbf{x}_i$. The score-based method, ScoreDenoise [22], has a similar filtering objective given by:

$$\tilde{\mathbf{x}}_i^{(t)} = \tilde{\mathbf{x}}_i^{(t-1)} + \alpha^{(t)} \mathcal{E}_i(\tilde{\mathbf{x}}_i^{(t-1)}), t = 1, \dots, T \quad (3)$$

where $\mathcal{E}_i(\mathbf{x}) = (1/K) \sum_{\mathbf{x}_j \in kNN(\mathbf{x}_i)} \mathcal{S}_j(\mathbf{x})$ is an ensemble average of scores for the k -neighbors of the given point. These scores are predicted by the ScoreDenoise network and correspond to the gradient-log of the noise-convolved probability $\nabla_{\mathbf{x}} \log[(p*n)(\mathbf{x})]$. Its training objective is:

$$L_i^{SD} = \mathbb{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{x}_i)} \left[\|s(\mathbf{x}) - \mathcal{S}_i(\mathbf{x})\|_2^2 \right], \quad (4)$$

where $s(\mathbf{x}) = NN(\mathbf{x}, \mathcal{Y}) - \mathbf{x}$, and $NN(\mathbf{x}, \mathcal{Y})$ is the nearest point to \mathbf{x} in the clean patch.

Based on the above analysis, we note the training objective of [22] considers a concentration of neighbors \mathbf{x} near \mathbf{x}_i while PCN’s training objective only attempts to infer the displacement for \mathbf{x}_i . Despite the differences of these two

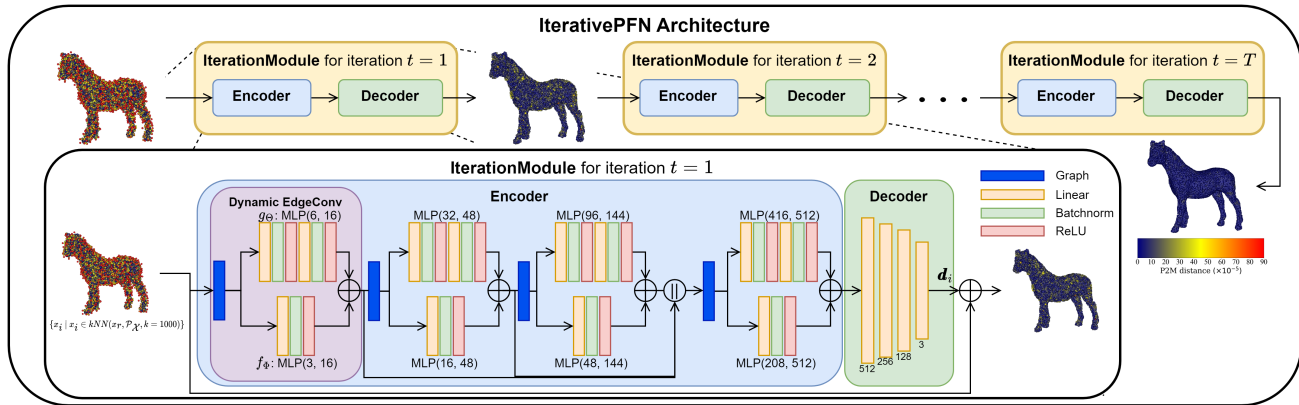


Figure 2. Overview of our IterativePFN method. Unlike existing methods which need to be iteratively applied during test-times, we explicitly model T iterations of filtering in our network using T IterationModules.




Figure 3. Histograms of distance to clean surface for each filtered point cloud. We look at the initial Casting shape at 50K point resolution and Gaussian noise of 2.5%.

methods, both groups of methods perform reasonably well at removing noise. This is in part due to the fact that during the test phase they emulate the Langevin equation of a reverse Markov process that iteratively removes noise. However, we also observe several drawbacks: **Firstly**, both displacement and score-based methods are iterative only at test time. During training, they see noisy patches with underlying Gaussian noise distributions and attempt to infer the filtered displacements [28] or scores [22] to shift points directly onto the clean patch. Consequently, the PCN position update, Eq. (1), and gradient ascent equation, Eq. (3), both neglect the fact that filtered points have an additive noise contribution $\sigma_\tau \xi$ where $\xi \sim \mathcal{N}(0, I) \wedge \sigma_\tau \in \mathbb{R}$ which is due to the stochastic nature of the Langevin equation.

Secondly, ScoreDenoise [22] relies on a decaying step size, $\alpha^{(t)}$, to ensure the convergence of Eq. (3) to a steady state. However, this can indeed be very slow to converge [5]. These observations motivate us to propose a learning based filtering method that effectively incorporates iterative filtering into the training objective such that the network is able to successfully recognize the underlying noise distribution at each iteration and, thereby, impose a more robust filtering objective wherein filtered results converge faster to the clean surface. Fig. 3 depicts histograms of the number of filtered points at a given distance from the clean surface.

We use the default test time iteration numbers for PCN and ScoreDenoise, which are 4 and 30, respectively. For our method, we use 4 IterationModules, i.e., 4 *internal iterations* or 1 *external iteration*. For PCN, the mean distance from the clean surface is 0.014 (in units w.r.t. the bounding sphere’s radius) while for ScoreDenoise it is 0.012 which are both much higher than the mean value of our method: 0.007. For point clouds filtered using PCN and ScoreDenoise, larger numbers of filtered points lie further away from the underlying clean surfaces after applying multiple iterations of filtering at test time but our method filters points closer to the surface after a single external iteration.

4. True Iterative Point Cloud Filtering

In this section we look in detail at our proposed **IterativePFN** network that consists of a stack of encoder-decoder pairs. Each pair, dubbed **IterationModule**, models a single internal filtering iteration and corresponds to a test time iteration of a method such as ScoreDenoise or PCN. At each internal iteration, the corresponding IterationModule takes a patch of noisy points and constructs a directed graph. The patch, and its associated graph, are then consumed by Dynamic EdgeConv layers to generate latent representations of vertices for regressing filtered displacements for each patch point. To model the effect of T iterations, we use T IterationModules. An external iteration represents the effect of filtering a point cloud using all IterationModules in the network. During inference, we use farthest point sampling on the noisy point cloud to generate input patches. However, this implies overlapping regions between different patches. Therefore, we propose a generalized method of patch stitching to recover the best filtered points.

4.1. Graph convolution and network architecture

Graph convolution based neural networks have been shown to generate rich latent representations of 3D point

based graphs [24, 37]. In particular, the work of Pistilli et al. [24] focuses on graph convolution for filtering point clouds. Apart from generating rich feature representations, graph convolution also provides an advantage on efficiency and lower inference times over point set convolution methods, that use PointNet inspired backbones, such as PCN and Pointfilter. These methods are designed to consume patches of points to filter a single central patch point, which leads to slower processing. In graph convolution based methods, all vertices within the graph are simultaneously processed, with filtered displacements being inferred for each single vertex (point) within the graph. For our IterativePFN network, we use a modified form of the Dynamic EdgeConv layers, proposed by Wang et al. [37], within each encoder, to generate rich feature representations of the input graphs. Formally, the vertex feature is updated as $\mathbf{h}_i^{l+1} = f_{\Phi}(\mathbf{h}_i^l) + \sum_{j:(i,j) \in \mathcal{E}} g_{\Theta}(\mathbf{h}_i^l \parallel \mathbf{h}_j^l - \mathbf{h}_i^l)$, where i is a vertex on the graph, (i, j) form an edge and (\parallel) represents concatenation. Here, $f_{\Phi} : \mathbb{R}^{F^l} \rightarrow \mathbb{R}^{F^{l+1}}$ and $g_{\Theta} : \mathbb{R}^{F^l} \times \mathbb{R}^{F^l} \rightarrow \mathbb{R}^{F^{l+1}}$ are parametrized by MLPs. F^l is the feature dimension at layer l . Thereafter, decoders consisting of 4 Fully Connected (FC) layers, consume these latent features to infer displacements at each iteration.

4.2. Graph construction

Given a clean point cloud $\mathcal{P}_{\mathcal{Y}} = \{\mathbf{y}_i \mid \mathbf{y}_i \in \mathbb{R}^3, i = 1, \dots, n\}$, perturbed points are formed by adding Gaussian noise with standard deviation σ_0 between 0.5% and 2% of the radius of the bounding sphere. Therefore, a noisy point cloud is given by $\mathcal{P}_{\mathcal{X}} = \mathcal{P}_{\mathcal{Y}} + \sigma_0 \xi, \xi \sim \mathcal{N}(0, I)$. Subsequently, given a reference point $\mathbf{x}_r \in \mathcal{P}_{\mathcal{X}}$, we obtain patches \mathcal{X} such that $\mathcal{X} = \{\mathbf{x}_i \mid \mathbf{x}_i \in \mathcal{P}_{\mathcal{X}} \wedge \mathbf{x}_i \in kNN(\mathbf{x}_r, \mathcal{P}_{\mathcal{X}}, k = 1000)\}$ where $kNN(\mathbf{x}, \mathcal{P}_{\mathcal{X}}, k = m)$ refers to the m nearest neighbors of the point \mathbf{x} in the point set $\mathcal{P}_{\mathcal{X}}$. Specifically, we select the 1000 nearest neighbors of \mathbf{x}_r . The corresponding clean patch \mathcal{Y} w.r.t. \mathcal{X} is given by $\mathcal{Y} = \{\mathbf{y}_i \mid \mathbf{y}_i \in \mathcal{P}_{\mathcal{Y}} \wedge \mathbf{y}_i \in kNN(\mathbf{x}_r, \mathcal{P}_{\mathcal{Y}}, k = 1200)\}$. Boundary points of \mathcal{X} may originate from points on the clean surface outside the 1000 nearest neighbors of \mathbf{x}_r in \mathcal{Y} , that are perturbed onto \mathcal{X} due to noise. To account for this, we follow ScoreDenoise’s strategy by setting the number of points in \mathcal{Y} to be slightly greater than \mathcal{X} . Thereafter, patches are centered at the reference point, i.e., $\mathcal{X} = \mathcal{X} - \mathbf{x}_r$ and $\mathcal{Y} = \mathcal{Y} - \mathbf{x}_r$. Patch point indices are then used to construct a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ of vertices and edges where $\mathcal{V} = \{i \mid \mathbf{x}_i \in \mathcal{X}\}$ and $\mathcal{E} = \{(i, j) \mid \mathbf{x}_i, \mathbf{x}_j \in \mathcal{X} \wedge \mathbf{x}_j \in kNN(\mathbf{x}_i, \mathcal{X}, k = 32)\}$. For each vertex, edges are formed with its 32 nearest neighbors.

4.3. Displacement based iterative filtering during training

While previous displacement based methods, such as PCN and Pointfilter, attempt to infer the filtered displacement




Figure 4. An adaptive ground truth target is used during intermediate filtering iterations to encourage convergence to the surface.

required to directly shift a noisy point onto the clean patch, our method uses multiple IterationModules to iteratively reduce noise. Thus we propose a novel training objective aimed at gradually reducing point cloud noise as illustrated in Fig. 4. Our loss function is given by,

$$L_i^{(\tau)}(\mathcal{Y}^{(\tau)}) = \left\| \mathbf{d}_i^{(\tau)} - \delta \mathbf{x}_i^{(\tau)}(\mathcal{Y}^{(\tau)}) \right\|_2^2, \quad (5)$$

where $\mathbf{d}_i^{(\tau)}$, the filtered displacement, is the output of IterationModule $t = \tau$, $\delta \mathbf{x}_i^{(\tau)}(\mathcal{Y}^{(\tau)}) = NN(\mathbf{x}_i^{(\tau-1)}, \mathcal{Y}^{(\tau)}) - \mathbf{x}_i^{(\tau-1)}$ and $NN(\mathbf{x}_i^{(\tau-1)}, \mathcal{Y}^{(\tau)})$ is the nearest neighbor of $\mathbf{x}_i^{(\tau-1)}$ in $\mathcal{Y}^{(\tau)}$. Moreover, $\mathcal{Y}^{(\tau)} = \mathcal{Y} + \sigma_{\tau} \xi \wedge \xi \sim \mathcal{N}(0, I)$ is the **adaptive ground truth** patch at iteration $t = \tau$. This target patch contains less noise compared to $\mathcal{Y}^{(\tau-1)}$. We ensure $\sigma_0 > \sigma_1 > \dots > \sigma_T$ where we set $\sigma_T = 0$. That is, $\mathcal{Y}^{(T)} = \mathcal{Y}$. For simplicity, we set $\sigma_{\tau+1} = \sigma_{\tau} / \delta$, where $\delta = 16/T$ is the decay hyperparameter controlling the noise scale of the target patch, $12 \geq T \geq 2$ and $T - 2 \geq \tau \geq 0$. More precisely, we distinguish our training objective from PCN and Pointfilter by minimizing the L_2 loss between the filtered displacements and the nearest neighbor in a less noisy version of the same, underlying, clean surface.

4.4. Generalized patch stitching



Figure 5. Patches of nearest neighbor points constructed using farthest point sampling along the surface.

During inference, we use farthest point sampling to obtain R reference points $\{\mathbf{x}_r\}_{r=1}^R$ and construct input patches using these points. All points within a patch are filtered simultaneously, as opposed to methods such as PCN and Pointfilter where a given patch is used to filter only a single central point. As shown in Fig. 5, patches may have overlapping regions with repeated points. Therefore, we must find the patches that optimally filter these repeated points.

Zhou et al., proposed patch stitching as a mechanism to obtain the best filtered points within such regions [42]. However, this method relies on network inferred weights, obtained using a self-attention module, to infer the best patch that yields the best filtered point. However, this is computationally expensive as it relies on a self-attention module and is not easy to generalize to methods that do not use such a module. Hence, we are motivated to design a more general patch stitching method that can be used in conjunction with any graph based neural network architecture. We observe that filtering results of points near the boundary of patches are less favorable than those closer to the reference points, \mathbf{x}_r , which are located at the center. This is due to the fact that points close to the patch boundary have asymmetric neighborhoods and, as such, receive biased signals during graph convolution that affect their filtered output. Therefore, an intuitive strategy is to weight input points based on their proximity to the reference point \mathbf{x}_r according to a Gaussian distribution, $w_i = \frac{\exp(-\|\mathbf{x}_i - \mathbf{x}_r\|_2^2 / r_s^2)}{\sum_i \exp(-\|\mathbf{x}_i - \mathbf{x}_r\|_2^2 / r_s^2)}$, where r_s is the support radius, which we empirically set to $r_s = r/3$ with r being the patch radius.

4.5. Iterative filtering and patch stitching

Now, the loss at $t = \tau$ is given by $L^{(\tau)} = \sum_i w_i L_i^{(\tau)}$, where the individual loss contribution at each point is given by $L_i^{(\tau)}$, according to Eq. (5). Finally, we sum loss contributions across iterations to obtain the final loss,

$$\mathcal{L}_a = \sum_{\tau=1}^T L^{(\tau)}, \quad (6)$$

which is used to train our network. It enjoys the distinction of being a truly iterative train time filtering solution that, at test times, can consume noisy patches and filter points without needing to resort to multiple external iterations.

5. Experimental Results

5.1. Dataset and implementation

We follow ScoreDenoise [22] and PDFlow [23] and utilize the PUNet [39] dataset to train our network. We retrain all other methods including PCN [28], GPDNet [24], DMRDenoise [20] and Pointfilter [41]. An implementation of RePCDNet [7] was not available for comparison. Poisson disk sampling is used to sample 40 training meshes, at resolutions of 10K, 30K and 50K points, which yields 120 point clouds for the training set. Subsequently, Gaussian noise with standard deviation ranging from 0.5% to 2% of the bounding sphere’s radius is added to each point cloud. For testing, we utilize 20 test meshes sampled at resolutions of 10K and 50K, similar to [22, 23]. These 40 point clouds are perturbed by Gaussian noise with standard deviations of 1%, 2% and 2.5% of the bounding sphere’s radius.

For comparisons on real world scans, we look at test results on the following datasets: The Paris-Rue-Madame database consisting of scans acquired by the Mobile Laser Scanning system L3D2 [32], which capture the street of Rue Madame in the 6th district of Paris. It contains real world noisy artifacts, a consequence of the limitations of scanning technology, and provides an excellent basis for comparing performance on real-world data. As ground truth point clouds and meshes are unavailable, we present qualitative results. Next, we consider the Kinect v1 and Kinect v2 datasets of [36] consisting of 71 and 72 real-world scans acquired using Microsoft Kinect v1 and Kinect v2 cameras. We compare performance of all methods on the Chamfer distance (CD) [10] and the Point2Mesh distance (P2M) [16] evaluation metrics. Both metrics are calculated using their latest implementations in PyTorch3D [29].

Implementation. Our IterativePFN network is trained on NVIDIA A100 GPUs using PyTorch 1.11.0 with CUDA 11.3. We train the network for 100 epochs, with the Adam optimizer and a learning rate of 1×10^{-4} . All methods are tested on a NVIDIA GeForce RTX 3090 GPU to ensure fair comparison of test times.

5.2. Results on synthetic data

The comparison of methods on synthetic data is given by Table 1 and Fig. 6. The baseline displacement based methods, PCN and GPDNet, show sub-optimal performance at both low and high resolutions. Specifically, PCN shows sensitivity to high noise and filtered point clouds suffer from shrinkage. The resampling based DMRDenoise and score-matching based ScoreDenoise both show sensitivity to noise. The normalizing flow based PDFlow performs well at low resolution but has trouble at high resolution and high noise. Our method shows an advantage across all resolutions and noise scales. Pointfilter is also able to generalize well across resolutions and noise scales as it is based on a weighted bilateral filtering inspired loss function that uses ground truth normal information. However, it still compares less favorably to our method. Furthermore, for complex shapes such as the Casting shape of Fig. 6, Pointfilter’s weighted loss causes the collapse of closely neighboring surfaces onto a single surface.

5.3. Results on scanned data

Fig. 7 demonstrates filtering results on two scenes of the RueMadame database. As it contains only noisy scanned data, we only consider visual results. As shown by the results on scene 1, only our method consistently filters surfaces such as the sign post and vehicle hood. PDFlow leaves behind many noisy artifacts while ScoreDenoise performs only a little better. Likewise, Pointfilter leaves noisy artifacts while taking a large amount of time to finish filtering the scene. This is due to their filtering strategy that takes one

Method	10K points						50K points					
	1% noise		2% noise		2.5% noise		1% noise		2% noise		2.5% noise	
	CD	P2M	CD	P2M	CD	P2M	CD	P2M	CD	P2M	CD	P2M
Noisy	36.9	16.03	79.39	47.72	105.02	70.03	18.69	12.82	50.48	41.36	72.49	62.03
PCN	36.86	15.99	79.26	47.59	104.86	69.87	11.03	6.46	19.78	13.7	32.03	24.86
GPDNet	23.1	7.14	42.84	18.55	58.37	30.66	10.49	6.35	32.88	25.03	50.85	41.34
DMRDenoise	47.12	21.96	50.85	25.23	52.77	26.69	12.05	7.62	14.43	9.7	16.96	11.9
PDFlow	21.26	6.74	32.46	13.24	36.27	17.02	6.51	4.16	12.7	9.21	18.74	14.26
ScoreDenoise	25.22	7.54	36.83	13.8	42.32	19.04	7.16	4.0	12.89	8.33	14.45	9.58
Pointfilter	24.61	7.3	35.34	11.55	40.99	15.05	7.58	4.32	9.07	5.07	10.99	6.29
Ours	20.56	5.01	30.43	8.45	33.52	10.45	6.05	3.02	8.03	4.36	10.15	5.88

Table 1. Filtering results on the PUNet dataset. CD and P2M distances are multiplied by 10^5 .

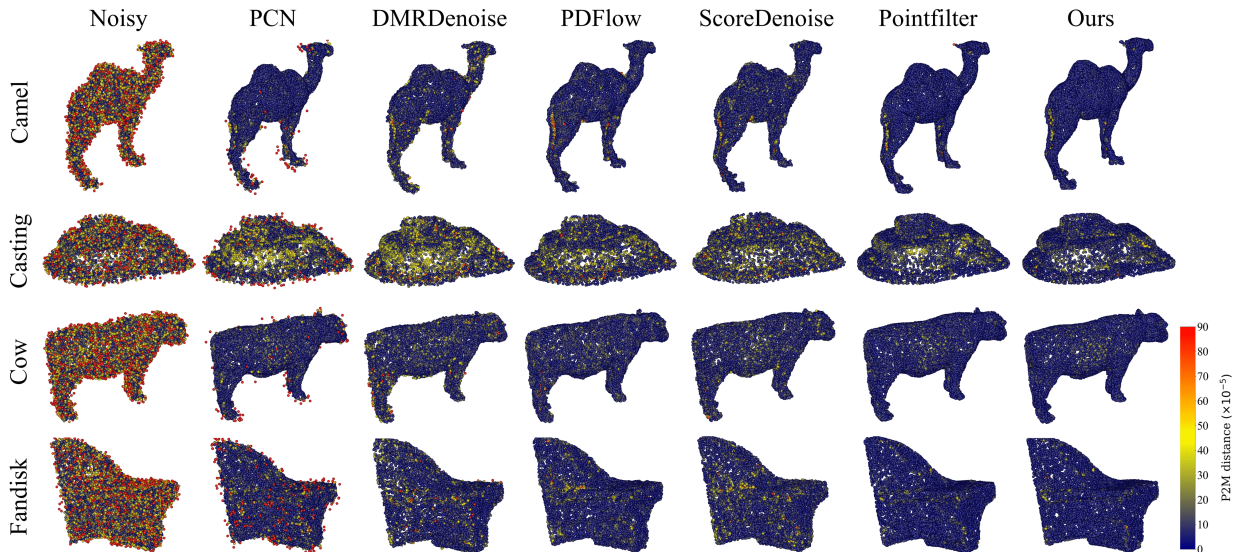


Figure 6. Visual results of of point-wise P2M distance for 50K resolution shapes with Gaussian noise of 2% of the bounding sphere radius.

patch per central point (OPPP). Therefore, patches need to be constructed for each point, which is inefficient. Please refer to the supplementary document for details on each method’s runtime. Table 2 presents results on the Kinect v1 and v2 datasets. Our method, with a smaller patch size of 100, with edges between 32 k -nearest neighbors, achieves the best results on the CD metric and second best results on the P2M metric. Additional results on synthetic and scanned data is provided in the supplementary.

5.4. Ablation study

In this section we look at the importance of the following elements to our method: 1) The number of IterationModules modelling the iteration number in our IterativePFN network, 2) the impact of true iterative filtering and 3) the impact of a fixed ground truth target, given by Eq. (7), versus an adaptive ground truth target, Eq. (5) and Eq. (6). The fixed ground truth target-based loss is given by,

$$\mathcal{L}_b = \sum_{\tau=1}^T \left[\sum_i w_i \left(\left\| \mathbf{d}_i^{(\tau)} - \delta \mathbf{x}_i^{(\tau)}(\mathcal{Y}) \right\|_2^2 \right) \right], \quad (7)$$

where $\delta \mathbf{x}_i^{(\tau)}(\mathcal{Y}) = NN(\mathbf{x}_i^{(\tau-1)}, \mathcal{Y}) - \mathbf{x}_i^{(\tau-1)}$. Table 3 demonstrates the impact of iteration number on filtering. Although 8 iterations, i.e., 8 IterationModules, provide the best results for 1% and 2% noise, it does not generalize well, compared to 4 iterations, on the unseen noise of 2.5%. This is due to the higher number of IterationModules causing the network to specialize on the training noise scales of 0.5% to 2%. Furthermore, 8 IterationModules lead to larger network size and memory usage. Runtime memory consumption for 1, 2, 4, 8 and 12 iterations are 3.9GB, 7.6GB, 15GB, 29.5GB and 44.1GB, respectively. Hence, we select 4 iterations as the optimal number. Moreover, to confirm the efficacy of true iterative filtering, we train a deep network of 1 ItM (DPFN), with 6 Dynamic EdgeConv layers,

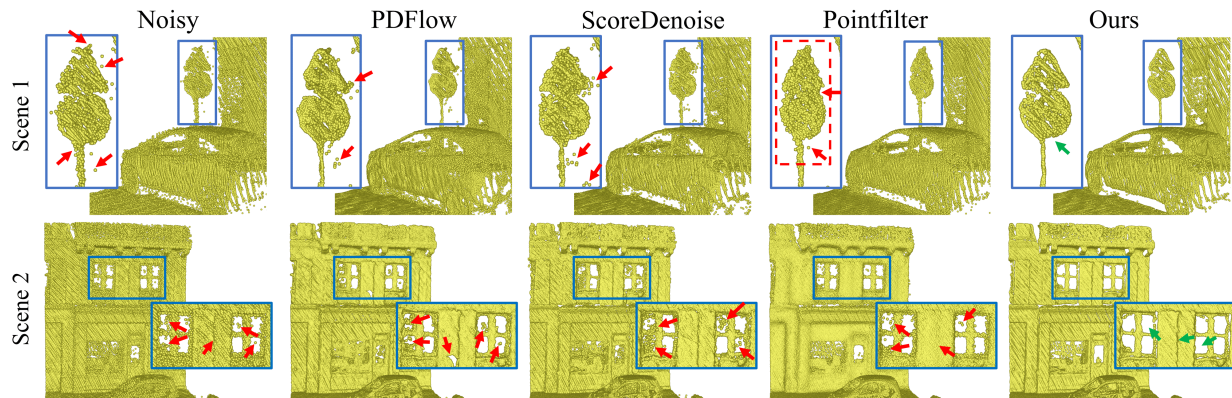


Figure 7. Results on two scenes of the real-world RueMadame dataset. Green and red arrows are used to indicate accurately and inaccurately filtered regions, respectively. PDFlow and ScoreDenoise leave behind large outliers while Pointfilter distorts underlying shapes.

Method	Kinect v1		Kinect v2	
	CD	P2M	CD	P2M
Noisy	14.49	9.32	22.63	13.39
PCN	13.73	8.75	22.48	13.29
GPDNet	14.83	8.69	23.09	11.78
DMRDenoise	22.78	12.89	-	-
PDFlow	13.34	8.69	19.87	11.69
ScoreDenoise	13.22	8.18	19.66	11.08
Pointfilter	13.77	7.91	18.85	10.29
Ours ($ \mathcal{X} = 1000$)	14.07	9.16	19.68	11.69
Ours ($ \mathcal{X} = 100$)	13.2	8.43	18.69	10.92

Table 2. Results on the Kinect v1 and Kinect v2 datasets. CD and P2M distances are multiplied by 10^5 .

and the same number of parameters (3.2M) as our IterativePFN with 4 ItMs. This variant performs sub-optimally, especially at high noise. This indicates the importance of the true iterative filtering approach. Finally, we consider the impact of keeping the ground truth target fixed during training, i.e., \mathcal{L}_b , where \mathcal{Y} corresponds to the clean patch. We see that the adaptive ground truth \mathcal{L}_a makes a noticeable difference, especially at the unseen noise scale of 2.5%. The supplementary contains additional ablation studies on 1) the impact of iteration number at higher noise, 2) the effect of patch stitching on filtering and 3) filtering results when using PointNet++ as the encoder backbone within each ItM.

6. Limitations and Future Work

The adaptive ground truth targets rely on adding noise, of a given distribution and noise scale (i.e., standard deviation), during training. However, this implies the noise distribution should be easy to replicate, such as Gaussian noise, with different scales in decreasing order. In future, it would be interesting to generalize this approach to utilize noisy data that simulates real world noise, with noise dis-

Ablation	10K points					
	1% noise		2% noise		2.5% noise	
	CD	P2M	CD	P2M	CD	P2M
\mathcal{L}_a & 1 it.	21.95	5.42	32.38	9.55	36.98	12.71
\mathcal{L}_a & 2 it.	21.13	5.14	30.82	8.67	34.33	11.0
\mathcal{L}_a & 4 it.	20.56	5.01	30.43	8.45	33.52	10.45
\mathcal{L}_a & 8 it.	19.78	4.9	30.12	8.3	33.88	10.78
\mathcal{L}_a & 12 it.	20.49	5.23	30.64	8.87	34.46	11.25
\mathcal{L}_a & DPFN	21.03	5.05	30.96	8.53	35.2	11.4
\mathcal{L}_b & 4 it.	20.64	5.04	30.59	8.54	34.17	10.87

Table 3. Ablation results for different iteration numbers and different loss functions. CD and P2M distances are multiplied by 10^5 .

tributions similar to that of real world scanners, such as Lidar scanner data. Additionally, we hope to apply our novel network architecture, of stacked encoder-decoder pairs, to other tasks, e.g., point cloud upsampling, that may benefit from a true iterative approach.

7. Conclusion

In this paper, we present IterativePFN, which consists of multiple IterationModules that explicitly model the iterative filtering process internally, unlike state-of-the-art learning based methods which only perform iterative filtering at test time. Furthermore, state-of-the-art methods attempt to learn filtered displacements that directly shift noisy points to the underlying clean surfaces and neglect the relationship between intermediate filtered points. Our IterativePFN network employs a re-imagined loss function that utilizes an adaptive ground truth target at each iteration to capture this relationship between intermediate filtered results during training. Our method ensures filtered results converge to the clean surface faster and, overall, performs better than state-of-the-art methods.

References

- [1] A. Adamson and M. Alexa. Point-sampled cell complexes. *ACM Trans. Graph.*, 25:671–680, 2006. [2](#)
- [2] M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. Levin, and Cláudio T. Silva. Computing and rendering point set surfaces. *IEEE Trans. Vis. Comput. Graph.*, 9:3–15, 2003. [1](#), [2](#)
- [3] H. Avron, Andrei Sharf, C. Greif, and D. Cohen-Or. L1-sparse reconstruction of sharp point set surfaces. *ACM Trans. Graph.*, 29:135:1–135:12, 2010. [3](#)
- [4] Yizhak Ben-Shabat and Stephen Gould. Deepfit: 3d surface fitting via neural network weighted least squares. In *Computer Vision – ECCV 2020*, pages 20–34. Springer International Publishing, 2020. [2](#)
- [5] Ruojin Cai, Guandao Yang, Hadar Averbuch-Elor, Zekun Hao, Serge Belongie, Noah Snavely, and Bharath Hariharan. Learning gradient fields for shape generation. In *Computer Vision – ECCV 2020*, pages 364–381. Springer International Publishing, 2020. [3](#), [4](#)
- [6] Frederic Cazals and Marc Pouget. Estimating differential quantities using polynomial fitting of osculating jets. *Computer Aided Geometric Design*, 22(2):121–146, 2005. [2](#)
- [7] Honghua Chen, Zeyong Wei, Xianzhi Li, Yabin Xu, Mingqiang Wei, and Jun Wang. Repcd-net: Feature-aware recurrent point cloud denoising network. *International Journal of Computer Vision*, 130(3):615–629, 2022. [2](#), [3](#), [6](#)
- [8] Julie Digne. Similarity based filtering of point clouds. *2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 73–79, 2012. [2](#)
- [9] Julie Digne and C. D. Franchis. The bilateral filter for point clouds. *Image Process. Line*, 7:278–287, 2017. [1](#)
- [10] Haoqiang Fan, Hao Su, and Leonidas J. Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [6](#)
- [11] S. Fleishman, Iddo Drori, and D. Cohen-Or. Bilateral mesh denoising. *ACM SIGGRAPH 2003 Papers*, 2003. [2](#)
- [12] Gaël Guennebaud and M. Gross. Algebraic point set surfaces. In *SIGGRAPH 2007*, 2007. [1](#), [2](#)
- [13] Hui Huang, Dan Li, Hongxing Zhang, U. Ascher, and D. Cohen-Or. Consolidation of unorganized point clouds for surface reconstruction. *ACM SIGGRAPH Asia 2009 papers*, 2009. [1](#), [2](#)
- [14] Youngki Kim, Kiyoun Kwon, and Duhwan Mun. Mesh-offset-based method to generate a delta volume to support the maintenance of partially damaged parts through 3d printing. *Journal of Mechanical Science and Technology*, 35(7):3131–3143, 2021. [1](#)
- [15] D. Levin. The approximation power of moving least-squares. *Math. Comput.*, 67:1517–1531, 1998. [2](#)
- [16] Ruihui Li, Xianzhi Li, Pheng-Ann Heng, and Chi-Wing Fu. Point cloud upsampling via disentangled refinement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [6](#)
- [17] Y. Lipman, D. Cohen-Or, D. Levin, and H. Tal-Ezer. Parameterization-free projection for geometry reconstruction. *ACM SIGGRAPH 2007 papers*, 2007. [1](#), [2](#)
- [18] Xuequan Lu, S. Schaefer, Jun Luo, Lizhuang Ma, and Y. He. Low rank matrix approximation for 3d geometry filtering. *IEEE transactions on visualization and computer graphics*, PP, 2020. [1](#), [3](#)
- [19] Chenxu Luo, Xiaodong Yang, and A. Yuille. Self-supervised pillar motion learning for autonomous driving. In *CVPR*, 2021. [1](#)
- [20] Shitong Luo and Wei Hu. Differentiable manifold reconstruction for point cloud denoising. In *Proceedings of the 28th ACM International Conference on Multimedia*, page 1330–1338. Association for Computing Machinery, 2020. [2](#), [3](#), [6](#)
- [21] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2837–2845, 2021. [3](#)
- [22] Shitong Luo and Wei Hu. Score-based point cloud denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4583–4592, October 2021. [1](#), [2](#), [3](#), [4](#), [6](#)
- [23] Aihua Mao, Zihui Du, Yu-Hui Wen, Jun Xuan, and Yong-Jin Liu. Pd-flow: A point cloud denoising framework with normalizing flows. In *The European Conference on Computer Vision (ECCV)*, 2022. [1](#), [3](#), [6](#)
- [24] Francesca Pistilli, Giulia Fracastoro, Diego Valsesia, and Enrico Magli. Learning graph-convolutional representations for point cloud denoising. In *Computer Vision – ECCV 2020*, pages 103–118. Springer International Publishing, 2020. [1](#), [3](#), [5](#), [6](#)
- [25] R. Preiner, O. Mattausch, Murat Arıkan, R. Pajarola, and M. Wimmer. Continuous projection for fast I1 reconstruction. *ACM Transactions on Graphics (TOG)*, 33:1 – 13, 2014. [1](#), [2](#)
- [26] C. Qi, Hao Su, Kaichun Mo, and L. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 77–85, 2017. [2](#), [3](#)
- [27] 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*, volume 30. Curran Associates, Inc., 2017. [2](#), [3](#)
- [28] Marie-Julie Rakotosaona, Vittorio La Barbera, Paul Guerrero, N. Mitra, and M. Ovsjanikov. Pointcleannet: Learning to denoise and remove outliers from dense point clouds. *Computer Graphics Forum*, 39, 2020. [1](#), [2](#), [3](#), [4](#), [6](#)
- [29] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *ArXiv*, 2020. [6](#)
- [30] Oussama Remil, Qian Xie, Xingyu Xie, Kai Xu, and J. Wang. Data driven sparse priors of 3d shapes. *Computer Graphics Forum*, 36, 2017. [3](#)
- [31] Riccardo Roveri, A. Cengiz Öztireli, Ioana Pandele, and Markus Gross. Pointpronets: Consolidation of point clouds with convolutional neural networks. *Computer Graphics Forum*, 37(2):87–99, 2018. [2](#), [3](#)

- [32] Andrés Serna, Beatriz Marcotegui, François Goulette, and Jean-Emmanuel Deschaud. Paris-rue-madame database - a 3d mobile laser scanner dataset for benchmarking urban detection, segmentation and classification methods. In *ICPRAM*, 2014. 6
- [33] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. Curran Associates Inc., 2019. 3
- [34] Yujing Sun, S. Schaefer, and Wenping Wang. Denoising point sets via l0 minimization. *Comput. Aided Geom. Des.*, 35-36:2–15, 2015. 1, 3
- [35] Philipp R. W. Urech, M. Dissegna, C. Girot, and A. Grêt-Regamey. Point cloud modeling as a bridge between landscape design and planning. *Landscape and Urban Planning*, 203:103903, 2020. 1
- [36] Peng-Shuai Wang, Yang Liu, and Xin Tong. Mesh denoising via cascaded normal regression. *ACM Trans. Graph.*, 35(6):Article 232, 2016. 6
- [37] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (TOG)*, 2019. 2, 3, 5
- [38] Bekiroglu Yasemin, Björkman Mårten, Gandler Gabriela Zarzar, Exner Johannes, Ek Carl Henrik, and Kragic Danica. Visual and tactile 3d point cloud data from real robots for shape modeling and completion. *Data in Brief*, 30(105335-), 2020. 1
- [39] Lequan Yu, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng. Pu-net: Point cloud upsampling network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [40] L. Yu, X. Li, C. W. Fu, P. A. Heng, and D. Cohen-Or. *EC-Net: An edge-aware point set consolidation network*, volume 11211 LNCS of *Lecture Notes in Computer Science*. Springer Verlag, 2018. 3
- [41] Dongbo Zhang, Xuequan Lu, Hong Qin, and Y. He. Point-filter: Point cloud filtering via encoder-decoder modeling. *IEEE Transactions on Visualization and Computer Graphics*, 27:2015–2027, 2021. 1, 2, 3, 6
- [42] Jun Zhou, Wei Jin, Mingjie Wang, Xiuping Liu, Zhiyang Li, and Zhaobin Liu. Fast and accurate normal estimation for point clouds via patch stitching. *Computer-Aided Design*, 142, 2022. 6
- [43] Runsong Zhu, Yuan Liu, Zhen Dong, Yuan Wang, Tengping Jiang, Wenping Wang, and Bisheng Yang. Adafit: Rethinking learning-based normal estimation on point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6118–6127, October 2021. 2