

PD-Flow: A Point Cloud Denoising Framework with Normalizing Flows

Aihua Mao¹, Zihui Du¹, Yu-Hui Wen², Jun Xuan¹, and Yong-Jin Liu²

¹ South China University of Technology, Guangzhou, China
ahmao@scut.edu.cn, {csusami, 202020143921}@mail.scut.edu.cn

² Tsinghua University, Beijing, China
{wenyh1616, liuyongjin}@tsinghua.edu.cn

Abstract. Point cloud denoising aims to restore clean point clouds from raw observations corrupted by noise and outliers while preserving the fine-grained details. We present a novel deep learning-based denoising model, that incorporates normalizing flows and noise disentanglement techniques to achieve high denoising accuracy. Unlike existing works that extract features of point clouds for point-wise correction, we formulate the denoising process from the perspective of distribution learning and feature disentanglement. By considering noisy point clouds as a joint distribution of clean points and noise, the denoised results can be derived from disentangling the noise counterpart from latent point representation, and the mapping between Euclidean and latent spaces is modeled by normalizing flows. We evaluate our method on synthesized 3D models and real-world datasets with various noise settings. Qualitative and quantitative results show that our method outperforms previous state-of-the-art deep learning-based approaches. The source code is available at <https://github.com/unknownue/pdflow>.

Keywords: point cloud, denoising, normalizing flows

1 Introduction

As one of the most widely used representations for 3D objects, point clouds have attracted considerable attention in many fields, including geometric analysis, robotic object detection, and autonomous driving. The rapid development of 3D scanning devices, such as depth cameras and LiDAR sensors, has made point cloud data increasingly popular. However, raw point clouds produced from these devices are inevitably contaminated by noise and outliers, due to inherent environment noise (e.g., lighting and background) and hardware limitation. Hence, point cloud denoising, which is a technique to restore high-quality and well-distributed points, is crucial for downstream tasks.

Despite decades of research, point cloud denoising remains a challenging problem, because of the intrinsic complexity of the topological relationship and

Y.H Wen and Y.J Liu are the corresponding authors.

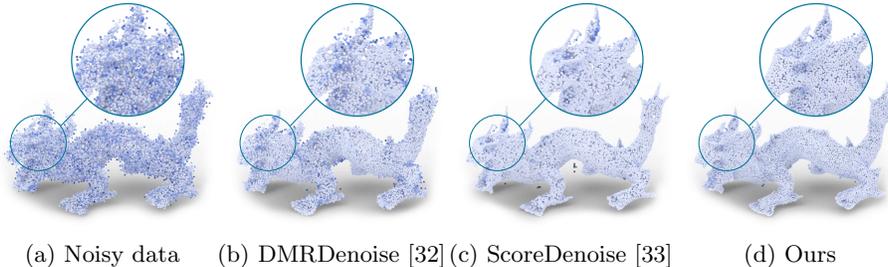


Fig. 1. The denoising results produced by (b) DMRDenoise [32], (c) ScoreDenoise [33] and (d) our method from noisy input (a). Deeper color indicates higher error. Our method preserves notably more fine details with less noise and outperforms others especially in uniformity.

connectivity among points. Traditional denoising methods [2,28,4,8,34,30] perform well in some circumstances. However, they generally rely on prior knowledge on point sets or some assumptions on noise distributions, and they may compromise the denoising quality for unseen noise (e.g., distortion, non-uniformity).

Recent promising deep learning approaches [12,41,17,32,33] bring new insight to point cloud denoising in a data-driven manner and exhibit superior performance over traditional methods. These works can be classified into two categories. The first class treats existing points as approximating the underlying surface by regressing points [12], predicting displacements [41,51], or progressive movement [33]. Nonetheless, the point features are extracted from the local receptive field independently. Therefore, consistent surface properties may not be preserved between neighborhood points, resulting in artifacts, such as outliers and scatter. The second class treats downsampling noisy data as a coarse point set and resampling/upsampling points from the learned manifold with a target resolution [17,32]. However, the downsampling scheme inevitably discards geometric details, leading to distorted distribution.

In this paper, we consider the noisy point clouds as samples of the joint distribution of 3D shape and corrupted noise. Based on this setup, it is intuitive to capture the characteristics of noise and underlying surface in the form of distribution. Thus, we can formulate the point cloud denoising problem as disentangling the clean section from its latent representation. We can also interpret this idea from the perspective of signal processing [36], where clean points and noise are analogous to low- and high-frequency part of signals, respectively. We can filter out the high-frequency contents and recover the smooth signal via the low-frequency counterpart that encodes the major information of raw signal.

Our denoising technique mainly consists of three phases: 1) learning the distribution of noisy point clouds by encoding the points into a latent representation, 2) filtering out the noise section from the latent representation, and 3) decoding/restoring noise-free points from the clean latent code. To realize this process, we require a generative model that can simultaneously learn the latent distribution and restore clean points. In this paper, we propose to exploit

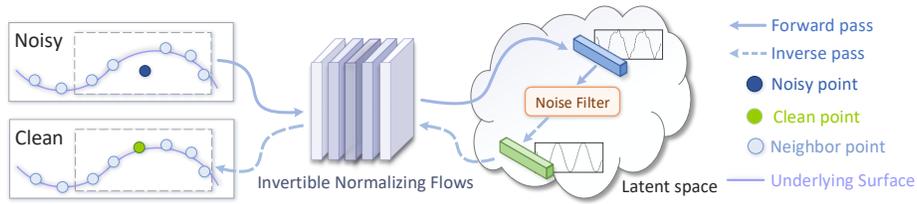


Fig. 2. Schematic illustration of the proposed method. We disentangle the noise factor from the latent representation of noisy point clouds and leverage NFs to model the transformation between Euclidean and latent spaces.

normalizing flows (NFs) in an invertible generative framework, to model the distribution mapping of point clouds. The whole process is illustrated in Fig. 2.

Compared with other popular deep learning models, such as generative adversarial network (GAN) and variational autoencoder, NFs provide several advantages: (i) NFs are capable of transforming complex distributions into disentangled code space, which is a desired property for point cloud denoising task, (ii) an NF is an invertible and lossless propagation process, which ensures one-to-one mapping between point clouds and their latent representations, and (iii) NFs realize the encoding and decoding process in a unified framework and share weights between forward and inverse propagations.

In summary, the main contributions of this work include:

- We propose a simple yet intuitive framework for point cloud denoising, called PD-Flow, which learns the distribution of noisy point sets and performs denoising via noise disentanglement.
- We propose to augment vanilla flows to improve the flexibility and expressiveness of the latent representation of points. We investigate various noise filtering strategies to disentangle noise from latent points.
- To validate the effectiveness of our method, extensive evaluations are conducted on synthetic and real-world datasets. Qualitative experiments show that our method outperforms the state-of-the-art works on diverse metrics.

2 Related Works

2.1 Denoising Methods

Traditional denoising methods. Conventional methods for point cloud denoising can be coarsely classified into three categories: 1) Statistical-based filtering methods generally apply statistical analysis theories, such as kernel density estimation [42], sparse reconstruction principle [4,45,34], principal component analysis [35], Bayesian statistics [22] and curvature extraction [24]. 2) Projection-based filtering methods first construct a smooth surface (e.g., Moving Least Squares surface [2,13,3]) from a set of noisy points. Then, denoising is implemented by projecting points onto surfaces. According to projection strategies,

this class of methods can be further divided into, e.g., locally optimal projection [28,19,20], jet fitting [5], and bilateral filtering [9]. 3) Neighborhood-based filtering methods measure the correlation and similarity between a point and its neighbor points. Nonlocal-based methods [8,21,46,52] generally detect self-similarity among nonlocal patches and consolidate them into coherent noise-free point clouds. Graph-based denoising methods [14,18,43,50] naturally represent point cloud geometry with a graph. All above methods generally require user interaction or geometric priors (e.g., normals) and still lack the ability of filtering various noise levels.

Deep-learning-based denoising methods. In recent years, several deep learning based methods [41,17,32,12,33,37] have been proposed for point cloud denoising. PointCleanNet [41] first removes outliers and then predicts inverse displacement for each point [16]. It is the first learning-based method that directly inputs noisy data without the acquisition of normals nor noise/device specifications. Hermosilla et al. proposed Total Denoising (TotalDn) [17] to regress points from the distribution of unstructured *total* noise. This allows TotalDn to approximate the underlying surface without the supervision of clean points. Pistilli et al. proposed GPDNet [37], which is a graph convolutional network, to improve denoising robustness under high noise levels. In the denoising pipeline of DMRDenoise [32], noisy input is first downsampled by a differentiable pooling layer, and then the denoised points are resampled from estimated manifolds. However, using the downsampling schema [41,32] to remove outliers may cause unnecessary detail loss. Recently, Luo and Hu developed a score-based denoising algorithm (ScoreDenoise [33]), which utilizes the gradient ascent technique and iteratively moves points to the underlying surface via estimated scores.

Our method differs from the above methods in several aspects. First, we formulate the denoising process as disentangling noise from the factorized representation of noisy input. Second, instead of applying separate modules to extract high-level features and reconstruct coordinates, we unify the point encoding/decoding process with a bijective network design.

2.2 Normalizing Flows for Point Cloud Analysis

NFs define a probability distribution transformation for data, allowing exact density evaluation and efficient sampling. In recent years, NFs have become a promising method for generative modeling and have been adopted into various applications [39,27,31,1]. Representative models include discrete normalizing flows (DNF) [10,11,25] and continuous normalizing flows (CNF) [7,15].

As the first NF-based algorithm for point cloud generation, PointFlow [48] employs CNF to learn a two-level distribution hierarchy of given shapes. PointFlow is a flexible scheme for modeling point distribution. However, the expensive equation solvers and training instability issues still remain to be open problems. Sharing the similar idea, Discrete PointFlow [26] proposes to use discrete flow layers as an alternative to continuous flows to reduce computation overhead. Pumarola et al. [40] introduced C-Flow, which is a parallel conditional scheme in the DNF-based architecture, to bridge data between images and point clouds

domains. Postels et al. [38] recently presented mixtures of NFs to improve the representational ability of flows and show superior performance to a single NF model [26]. These works mainly focus on improving generative ability and are evaluated on toy datasets. However, there are few works concentrating on flow-based real-world point cloud applications.

In this paper, we take advantage of the invertible capacity of NFs, which enable exact latent variable inference and efficient clean point synthesis. To the best of our knowledge, no prior work has been proposed for the point cloud denoising task by developing a new framework with NFs.

3 Method

3.1 Overview

Given an input point set $\tilde{\mathcal{P}} = \{\tilde{p}_i = p_i + o_i\} \in \mathbb{R}^{N \times D_p}$ corrupted by noise $\mathcal{O} = \{o_i\}$, we aim to predict a clean point set $\mathcal{P} = \{\hat{p}_i\} \in \mathbb{R}^{N \times D_p}$, where N is the number of points, D_p is the point coordinate dimension, and \hat{p}_i is the prediction of clean point p_i . In our study, we consider the coordinate dimension with $D_p = 3$ and make no assumptions about the noise distribution of \mathcal{O} .

In this paper, we propose to utilize NFs to model the mapping of point distribution between Euclidean and latent spaces, thereby allowing us to formulate point cloud denoising as the problem of disentangling the noise factor from its latent representation. The overall denoising pipeline is shown in Fig. 3.

3.2 Flow-based Denoising Method

We consider the point cloud denoising problem from the perspective of distribution learning and disentanglement. We suppose the distribution of noisy point set $\tilde{\mathcal{P}}$ is the joint distribution of clean point set $\mathcal{P} = \{p_i\}$ and noise \mathcal{O} . Given a dataset of observation $\tilde{\mathcal{P}}$, we aim to learn a bijective mapping f_θ , which is parameterized by θ to approximate the data distribution:

$$\tilde{z} = f_\theta(\tilde{\mathcal{P}}) = f_\theta(\mathcal{P}, \mathcal{O}), \quad (1)$$

where $\tilde{z} \sim p_\theta(\tilde{z})$ is a random variable with known probability density. Note that $p_\theta(\tilde{z})$ follows a factorized distribution [10], such that $p_\theta(\tilde{z}) = \prod_i p_\theta(\tilde{z}_i)$ (i.e. the dimensions of \tilde{z} are independent of each other).

We further assume that f_θ can simultaneously learn to embed noise factor and intrinsic structure of point cloud into a disentangled latent code space (i.e. where noise is uniquely controlled by some dimensions). Based on this assumption, we approximate the clean latent representation z by

$$\hat{z} = \psi(\tilde{z}), \quad (2)$$

where $\psi : \mathbb{R}^D \rightarrow \mathbb{R}^D$ is a disentanglement function defined in latent space, and \hat{z} is an estimation of z . In this way, clean point samples $\hat{\mathcal{P}}$ can be derived by taking the inverse transformation

$$\hat{\mathcal{P}} = g_\theta(\hat{z}), \quad (3)$$

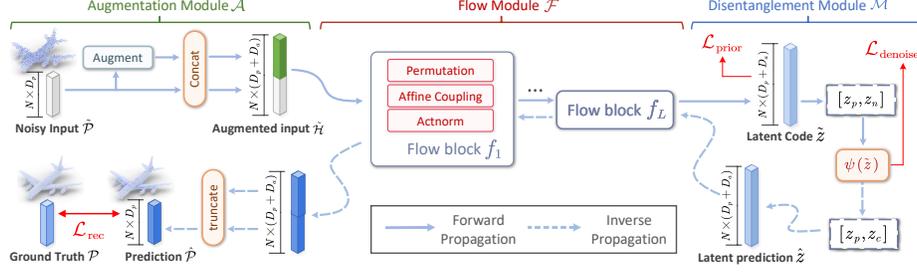


Fig. 3. The proposed flow-based denoising framework. Given a noisy point set $\tilde{\mathcal{P}} = \{\tilde{p}_i \in \mathbb{R}^{D_p}\}$, we first augment it with an additional D_a dimensional variable $\tilde{p}_i \in \mathbb{R}^{D_a}$ and obtain augmented point $h_i = [\tilde{p}_i, \tilde{p}_i] \in \mathbb{R}^{D_p + D_a}$. We transform the augmented points $\tilde{\mathcal{H}} = \{h_i\}$ to latent distribution $\tilde{z} \sim p_\vartheta(\tilde{z})$ by NFs. To estimate noise-free latent point \hat{z} , we filter out the noise factor from noisy \tilde{z} . We restore the noise-free point set $\hat{\mathcal{P}}$ from \hat{z} by the inverse propagation of \mathcal{F} , which utilizes the invertible capacity of NFs. Finally, the coordinates of the clean point set are derived from truncating the first D_p dimensions and discarding the augmented D_a dimensions.

where $g_\theta(\cdot) = f_\theta^{-1}(\cdot)$. The bijective mapping f_θ , which consists of a sequence of invertible transformations $f_\theta^1, \dots, f_\theta^L$, is referred to as *normalizing flows*. Denote by h^l the output of l -th flow transformation. Then h^{l+1} can be formulated as

$$h^{l+1} = f_\theta^{l+1}(h^l), \quad (4)$$

where $h^0 = \tilde{\mathcal{P}}$, $h^L = \tilde{z}$. Applying the change-of-variables formula and chain rule [10], the output probability density of $\tilde{\mathcal{P}}$ can be obtained as

$$\begin{aligned} \log p(\tilde{\mathcal{P}}; \theta) &= \log p_\vartheta(f_\theta(\tilde{\mathcal{P}})) + \log \left| \det \frac{\partial f_\theta}{\partial \tilde{\mathcal{P}}}(\tilde{\mathcal{P}}) \right| \\ &= \log p_\vartheta(f_\theta(\tilde{\mathcal{P}})) + \sum_{l=1}^L \log \left| \det \frac{\partial f_\theta^l}{\partial h^l}(h^l) \right|, \end{aligned} \quad (5)$$

where $\left| \det \frac{\partial f_\theta}{\partial \tilde{\mathcal{P}}}(\tilde{\mathcal{P}}) \right|$ is the log-absolute-determinant of the Jacobian of mapping f_θ , which measures the volume change [10] caused by f_θ . f_θ can be trained with the maximum likelihood principle using the gradient descent technique.

3.3 Augmentation Module

Dimensional bottleneck. To maintain the analytical invertibility, flow models impose more constraints on the network architecture than non-invertible models. One particular constraint is that the flow components f^1, \dots, f^L must output the same dimensionality D with the input data (where $D = 3$ for raw point clouds). The network bandwidth bottleneck sacrifices the model expressiveness, as shown in Fig. 4a.

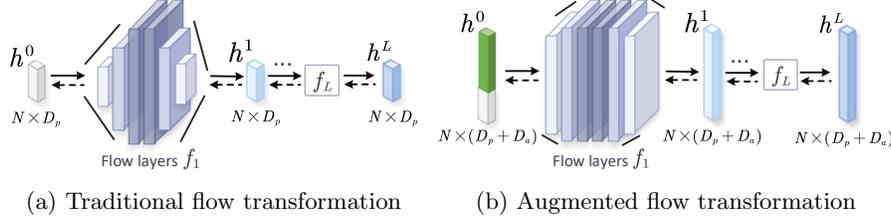


Fig. 4. Illustration of dimension augmentation. (a) The dimensionality of raw data limits the bandwidth of the network in vanilla flows. (b) Augmented flows can share considerable information among flow blocks, thereby improving model expressiveness.

Previous works [11,25] generally use a squeezing operator to alleviate this limitation by exchanging spatial dimensions for feature channels. However, the squeezing operator is mainly designed for image manipulation. It is non-trivial to adopt squeezing to point cloud due to the unordered nature in point sets.

Dimension augmentation. Inspired by VFlow [6], we resolve the bottleneck by increasing the dimensionality of input data. To be specific, for each input point $\tilde{p}_i \in \mathbb{R}^{D_p}$ in $\tilde{\mathcal{P}}$, we augment it with a random variable $\bar{p}_i \in \mathbb{R}^{D_a}$. This process is modeled by an augmentation module \mathcal{A} :

$$\bar{p}_i = \mathcal{A}(\tilde{p}_i, \mathcal{N}(\tilde{p}_i)), \bar{\mathcal{P}} = \{\bar{p}_i\}, \quad (6)$$

where $\mathcal{N}(\tilde{p}_i)$ denotes the k -nearest neighbors of \tilde{p}_i , and $\bar{\mathcal{P}}$ represents the set of augmented dimensions. We feed the augmented point set $\tilde{\mathcal{H}} = \{h_i = [\tilde{p}_i, \bar{p}_i]\}$ as input of flow module \mathcal{F} , and the underlying NFs become $\tilde{z} = f_\theta(\tilde{\mathcal{H}})$, where $\tilde{z} \in \mathbb{R}^{D_p+D_a}$, as shown in Fig. 4b.

Variational augmentation. To model the distribution of the augmented data space, VFlow [6] resorts to optimizing the evidence lower bound observation (ELBO) on the log-likelihood of augmented data as an alternative of Eq. 5:

$$\log p(\tilde{\mathcal{P}}; \theta) \geq \mathbb{E}_{q(\bar{\mathcal{P}}|\tilde{\mathcal{P}}; \phi)} \left[\log p(\tilde{\mathcal{P}}, \bar{\mathcal{P}}; \theta) - \log q(\bar{\mathcal{P}} | \tilde{\mathcal{P}}; \phi) \right], \quad (7)$$

where $q(\bar{\mathcal{P}} | \tilde{\mathcal{P}}; \phi)$ indicates the distribution of augmented data, which is modeled by the augmentation module \mathcal{A} , θ and ϕ denote the parameters of \mathcal{F} and \mathcal{A} , respectively. We briefly explain Eq. (7) in supplementary material.

3.4 Flow Module

The flow module \mathcal{F} transforms augmented points $\tilde{\mathcal{H}} = \{h_i\} \in \mathbb{R}^{N \times (D_p+D_a)}$ from the Euclidean space to the latent space, and vice versa.

The architecture of flow module \mathcal{F} comprises L blocks, where each block consists of a couple of flow components, as shown in Fig. 3. Each component is designed to satisfy the efficient invertibility and tractable Jacobian, including affine coupling layer [11], actnorm [25], and permutation layer [25]. The descriptions of each flow component are detailed in supplementary material.

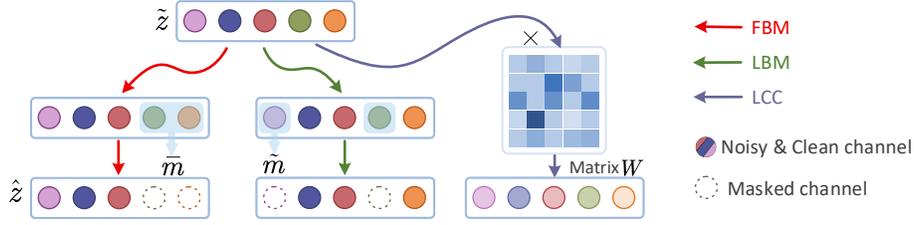


Fig. 5. Illustration of different filtering strategies.

3.5 Disentanglement Module

Let z_p and z_n be the clean point and noisy parts of the latent point \tilde{z} , i.e. $\tilde{z} = [z_p, z_n]$. We aim to disentangle noise z_n from \tilde{z} by a smooth operator $\psi : \mathbb{R}^D \rightarrow \mathbb{R}^D$

$$\hat{z} = [z_p, z_c] = \psi(\tilde{z}), \quad (8)$$

where z_c is the denoised feature and $D = D_p + D_a$. \hat{z} denotes the prediction of noise-free point representation, which is fed as the input of inverse propagation of flow module \mathcal{F} . However, how to implement $\psi(\cdot)$ is non-trivial. In this paper, we investigate three types of noise filtering strategies (Fig. 5) to formulate $\psi(\cdot)$ as follows:

Fix Binary Mask (FBM). Similar to a previous work [29], we explicitly divides the channels of latent code into two groups, i.e. clean and noisy channels. FBM simply sets noisy channels to 0 by

$$\psi(\tilde{z}) = \bar{m} \odot \tilde{z}, \quad (9)$$

$$\mathcal{L}_{\text{denoise}} = \mathcal{L}_{\text{FBM}} = 0, \quad (10)$$

where $\bar{m} \in \{0, 1\}^D$ is a fixed binary mask specified by the user, \odot denotes element-wise product and \mathcal{L}_{FBM} is the corresponding loss function of FBM.

Learnable Binary Mask (LBM). We employ a soft masking to latent \tilde{z} by

$$\psi(\tilde{z}) = \tilde{m} \odot \tilde{z}, \quad (11)$$

$$\mathcal{L}_{\text{denoise}} = \mathcal{L}_{\text{LBM}} = |\tilde{m}(1 - \tilde{m})|, \quad (12)$$

where $\tilde{m} \in \mathbb{R}^D$ is a learnable parameter, $|\cdot|$ denotes L_1 norm, and \mathcal{L}_{LBM} is the corresponding loss of LBM that encourages \tilde{m} to approximate the binary mask.

Latent Code Consistency (LCC). We minimize the latent representation between clean points and noisy points by

$$\psi(\tilde{z}) = W\tilde{z}, \quad (13)$$

$$\mathcal{L}_{\text{denoise}} = \mathcal{L}_{\text{LCC}} = \sum_{i=1}^N \|W\tilde{z}^{(i)} - z^{(i)}\|, \quad (14)$$

where $W \in \mathbb{R}^{D \times D}$ is a learnable matrix to transform \tilde{z} , N is the number of points, $\|\cdot\|$ denotes L_2 norm, and $z = f_\theta(\mathcal{P})$ is the latent representation encoded from reference points \mathcal{P} by the forward propagation of flow (similar to Eq. (1)). \mathcal{L}_{LCC} is the corresponding loss of LCC that encourages the transformed \hat{z} to be consistent with noise-free representation z , which is analogous to perceptual loss [23] that evaluates difference on high-level features.

3.6 Joint Loss function

We present an objective function for training PD-Flow that combines the *reconstruction loss*, *prior loss* and *denoise loss* (Section 3.5) as follows:

Reconstruction loss quantifies the similarity between the generated points $\hat{\mathcal{P}} \in \mathbb{R}^{N \times D_p}$ and reference clean points $\mathcal{P} \in \mathbb{R}^{N \times D_p}$. In this paper, we use the Earth Mover’s Distance (EMD) metric as \mathcal{L}_{rec} by minimizing

$$\mathcal{L}_{\text{rec}} = \mathcal{L}_{\text{EMD}}(\hat{\mathcal{P}}, \mathcal{P}) = \min_{\varphi: \hat{\mathcal{P}} \rightarrow \mathcal{P}} \sum_{\hat{p} \in \hat{\mathcal{P}}} \|\hat{p} - \varphi(\hat{p})\|, \quad (15)$$

where $\varphi: \hat{\mathcal{P}} \rightarrow \mathcal{P}$ is a bijection and $\|\cdot\|$ denotes L_2 norm.

Prior loss optimizes the transformation capability of flow module \mathcal{F} by maximizing the likelihood of observation $\tilde{\mathcal{P}}$. We implement the prior loss by minimizing the negative ELBO in Eq. (7):

$$\mathcal{L}_{\text{prior}}(\tilde{\mathcal{P}}) = \mathcal{L}(\tilde{\mathcal{P}}; \theta, \phi) = - \left[\log p(\tilde{\mathcal{P}}, \bar{\mathcal{P}}; \theta) - \log q(\bar{\mathcal{P}} | \tilde{\mathcal{P}}; \phi) \right] \quad (16)$$

where $\bar{\mathcal{P}} = \mathcal{A}(\tilde{\mathcal{P}})$ are augmented dimensions (Section 3.3). Intuitively, $\mathcal{L}_{\text{prior}}$ encourages the input points $\tilde{\mathcal{P}}$ to reach high probability under the predefined prior $p_\theta(\tilde{z})$.

Total Loss. Combining the preceding formulas, our method can be trained in an end-to-end manner by minimizing

$$\mathcal{L}(\theta, \phi, \sigma) = \alpha \mathcal{L}_{\text{rec}} + \beta \mathcal{L}_{\text{prior}} + \gamma \mathcal{L}_{\text{denoise}}, \quad (17)$$

where θ , ϕ and σ denotes the network parameters of \mathcal{F} , \mathcal{A} and \mathcal{M} , respectively. And, α , β , γ are the hyper-parameters to balance the loss.

3.7 Discussion

Benefit of dimension augmentation. The dimension augmentation setting provides extra benefits to vanilla flows: (i) The augmented NFs are generalization of vanilla flows, where the extra dimensionality D_a can be freely adjusted by users, allowing it to model more complex function. (ii) The augmented dimensions afford more flexibility and expressiveness to intermediate point features (i.e. h^l in Section 3.2) between flow transformations, avoiding extracting high dimensional features from scratch. (iii) The augmented dimensions increase the

Table 1. Comparison of denoising algorithms on PUSet.

#Points	10K									50K								
	1%			2%			3%			1%			2%			3%		
Noise Level	CD	P2M	HD															
Method	10^{-4}	10^{-4}	10^{-3}	10^{-4}	10^{-4}	10^{-3}	10^{-4}	10^{-4}	10^{-3}	10^{-4}	10^{-4}	10^{-3}	10^{-4}	10^{-4}	10^{-3}	10^{-4}	10^{-4}	10^{-3}
Jet [5]	3.47	1.20	1.58	4.83	1.89	2.97	6.15	2.86	7.00	0.82	0.19	0.90	2.38	1.35	3.36	5.64	4.16	9.05
GPF [30]	3.28	1.17	1.52	4.18	1.54	3.45	5.37	2.75	8.14	0.76	0.23	1.42	2.04	0.94	4.25	3.82	2.87	10.4
MRPCA [34]	3.14	1.01	1.77	3.87	1.26	2.59	5.13	2.03	4.84	0.70	0.12	0.79	2.11	1.06	3.21	5.64	3.97	7.65
GLR [50]	2.79	0.92	1.16	3.66	1.14	2.88	4.84	2.08	6.80	0.71	0.18	0.93	1.61	0.85	4.90	3.74	2.67	11.7
PCNet [41]	3.57	1.15	1.54	7.54	3.92	6.25	13.0	8.92	13.9	0.95	0.27	2.21	1.56	0.62	9.84	2.32	1.32	8.42
GPDNet [37]	3.75	1.33	3.03	8.00	4.50	6.08	13.4	9.33	13.5	1.97	1.09	1.94	5.08	3.84	7.56	9.65	8.14	16.7
Pointfilter [51]	2.86	0.75	2.87	3.97	1.30	6.21	4.94	2.14	9.26	0.82	0.24	2.38	1.46	0.77	4.58	2.25	1.44	8.70
DMR [32]	4.54	1.70	6.72	5.04	2.13	7.02	5.87	2.86	8.60	1.17	0.46	2.26	1.58	0.81	4.29	2.45	1.54	7.32
Score [33]	2.52	0.46	4.30	3.68	1.08	5.78	4.69	1.94	10.5	0.71	0.15	2.30	1.28	0.57	4.95	1.92	1.05	9.30
Ours	2.12	0.38	1.36	3.25	1.02	3.71	4.45	2.05	5.31	0.65	0.16	1.71	1.18	0.60	2.58	1.94	1.26	5.21

degrees of freedom for noise filtering in the disentanglement phase, which is particularly helpful because raw point clouds contain only one dimension of $D_p = 3$. We investigate the influence of dimension augmentation in Section 4.3.

Although the augmented dimensions increase the network size of flow module \mathcal{F} , the overhead is only marginal. The computation overhead mainly depends on the hidden layer size D_h of the internal transformation unit of \mathcal{F} instead of the output dimensionality $D_p + D_a$.

Unified noise disentanglement pipeline. Considering the invertible property of NFs, raw points $\tilde{\mathcal{P}}$ and latent \tilde{z} share the identical information in different domains. We only manipulate the point features in the disentanglement module throughout the whole denoising pipeline, demonstrating the feature disentanglement capability of NFs.

Additionally, we do not explicitly introduce extra network modules to predict point-wise displacement [41] or upsample to a target resolution [32] for point generation. Utilizing the flow invertibility can share parameters between forward and inverse propagations, which help us to reduce the network size and avoid the use of a decoding module.

4 Experiments

4.1 Datasets

We evaluate our method on the following datasets: (i) **PUSet**. This dataset is a subset of PUNet [49] provided by [33], which contains 40 meshes for training and 20 meshes for evaluation. (ii) **DMRSet**. This dataset collects meshes from ModelNet40 [47] provided by [32], which contains 91 meshes for training and 60 meshes for evaluation. These point clouds are perturbed by Gaussian noise of various noise levels at resolutions ranging from 10K to 50K points.

We implement PD-Flow with the PyTorch framework. The training settings, datasets and network configurations are detailed in the supplementary material.

4.2 Comparisons with State-of-the-art Methods

Evaluation metrics. We use four evaluation metrics in quantitative comparison, including (i) Chamfer distance (CD), (ii) Point-to-mesh (P2M) distance,

Table 2. Comparison of uniformity on 10K points under various Gaussian noise levels. This metric is estimated in the local area of different radii p . Besides, we also show the corresponding CD loss ($\times 10^{-4}$) of the full point clouds.

Noise	Methods	CD	Uniformity for different p				
			0.4%	0.6%	0.8%	1.0%	1.2%
1%	MRPCA [34]	3.14	1.89	2.30	2.42	2.59	2.83
	DMR [32]	4.54	4.02	5.06	6.02	7.03	7.95
	Score [33]	2.52	1.10	1.38	1.69	2.05	2.45
	Ours	2.12	0.33	0.43	0.55	0.71	0.89
2%	MRPCA [34]	3.87	2.21	2.56	2.85	2.97	3.14
	DMR [32]	5.04	3.45	4.04	4.68	5.35	6.03
	Score [33]	3.68	1.95	2.39	2.91	3.44	4.04
	Ours	3.25	0.89	1.18	1.49	1.83	2.18
3%	MRPCA [34]	5.13	2.28	2.29	2.32	2.45	2.60
	DMR [32]	5.87	3.81	4.53	5.16	5.85	6.55
	Score [33]	4.69	4.19	5.42	6.43	7.46	8.37
	Ours	4.45	1.80	2.33	2.83	3.34	3.87

(iii) Hausdorff distance (HD), and (iv) Uniformity (Uni). The detailed description of each metric is available in the supplementary material.

Quantitative comparison. We compare our method with traditional methods (including Jet [5], MRPCA [34], GPF [30], GLR [50]) and state-of-the-art deep learning-based methods (including PointCleanNet (PCNet) [41], Pointfilter [51], DMRDenoise (DMR) [32], GPDNet [37], ScoreDenoise (Score) [33]).

We use LCC as the default noise filter. The benchmark is based on 10K and 50K points disturbed by isotropic Gaussian noise with the standard deviation of noise ranging from 1% to 3% of the shape’s bounding sphere radius.

As shown in Table 1, traditional methods can achieve good performance in some cases depending on the manual tuning parameters, but they have difficulty in extending to all metrics. PCNet [41] and DMRDenoise [32] perform less satisfactory under 10K points, while GPDNet [37] fails to handle high noise levels. In most cases, our method outperforms Pointfilter [51] and ScoreDenoise [33], especially in CD and HD metrics. The quantitative comparison on DMRSet and more results of various noise types are provided in the supplementary material.

Furthermore, we present the quantitative results on uniformity metric under various noise levels in Table 2. Although MRPCA [34] achieves the best uniformity under 3% Gaussian noise, it fails to keep good generation accuracy on CD metric. Compared with other state-of-the-art methods [32,33], our method considerably promotes the uniformity of generated points.

Qualitative comparison. Fig. 6 visualizes the qualitative denoising results between ours and competitive works. We observe that our method achieves the most robust estimation under high noise corruption. In particular, our method can keep consistent density across different regions and avoid clustering phenomenon, resulting in better uniformity.

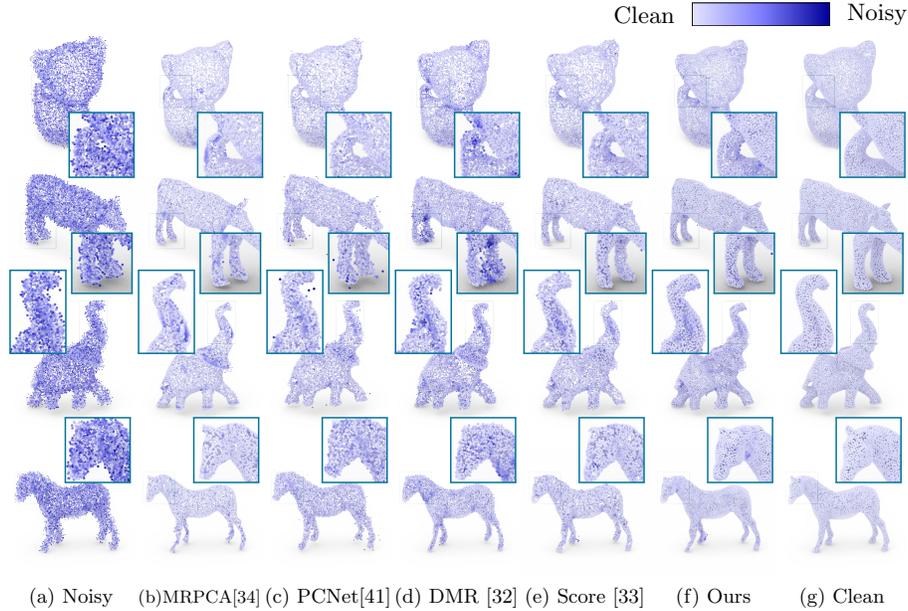


Fig. 6. Visual comparison of state-of-the-art denoising methods under 2% isotropic Gaussian noise. The color of each point indicates its reconstruction error measured by P2M distance. See supplementary material for visual results under more noise settings.

We also compare the denoising result under the *Paris-rue-Madame* [44] dataset, which contains real-world scene data captured by laser scanner. As shown in Fig. 7, our method improves the surface smoothness and preserves better details than DMRDenoise [32] and ScoreDenoise [33].

4.3 Ablation Study

We conduct ablation studies to demonstrate the contribution of the network design of PD-Flow. The evaluation is based on 10K points with 2% Gaussian noise in PUSet.

Flow architecture. The number of parameters mainly depends on the depth of flow module \mathcal{F} (e.g., the number of flow blocks L). As shown in Table 3, the fitting capacity improve as L increases. However, when L increases to 12, the relative performance boost becomes marginal, with the cost of a large number of network parameters and training instability. We find that $L = 8$ achieves the best balance between performance and training stability.

To verify the effectiveness of inverse propagation of \mathcal{F} , we replace the inverse pass with MLP layers, which are commonly used in other deep-learning-based methods [41,32]. As shown in Table 3, the inverse pass achieves better generation quality without introducing extra network parameters, which demonstrates the feasibility of flow invertibility. The similar result is also verified by the “forward

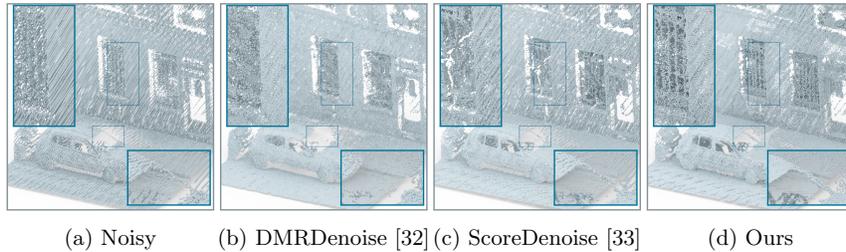


Fig. 7. Visual results of our denoiser on the real-world dataset *Paris-rue-Madame* [44].

Table 3. Ablation study of flow architectures.

Pass	#Flow block	#Params	CD	P2M
Forward + Inverse	4	299K	3.55	1.22
Forward + Inverse	8	470K	3.25	1.03
Forward + Inverse	12	647K	3.57	1.23
Forward + MLP	8	578K	4.65	2.14

+ mlp” curve in Fig. 8a, where using MLP as point generator leads to degraded performance.

Dimension augmentation. To investigate the impact of the number of augmentation channels D_a on model convergence, we show the training curve in Fig. 8a. The baseline model with $D_a = 0$ (i.e. vanilla flows) fails to converge to reasonable results. As the D_a increases, we observe faster convergence and better fitting capability. The similar trending can also be observed in quantitative evaluation under different D_a . This indicates that the dimension augmentation makes a key contribution to activate the denoising capability of NFs.

Effect on noise filtering strategies. We compare the performances of various filtering strategies with $D_a = 32$. As shown in Table 4, all these strategies can achieve competitive performance, and the LCC filter achieves the best results. For LBM with $D_a = 16$ to $D_a = 64$, we observe that about 2/3 channels approximate zeros in \tilde{m} after training.

We further investigate the impact of the number of masked channels D_m on FBM filter in Fig. 8b, which shows the training curve under different D_m settings. For $D_m = 0$, the flow network produces the same result as input due to the invertible nature of NFs. For $D_m = 1$ to $D_m = 16$, our method can converge to reasonable performance and presents little differences in convergence. This indicates that our method can adaptively embed noisy and clean channels to disentangled position, even if D_m is set to a low value. For $D_a = 32$ and $D_m \geq 32$, the latent code space only contains 3 or less channels to embed intrinsic information of clean points. We can observe that the insufficient channels for clean point embedding obviously lead to degrade performance, demonstrating that the dimension of latent point has great impact on model expressiveness and denoising flexibility. In Section 3.5, we introduce $\mathcal{L}_{\text{denoise}}$ as a regularization term to improve noise disentanglement capability. Table 4 compares the effects

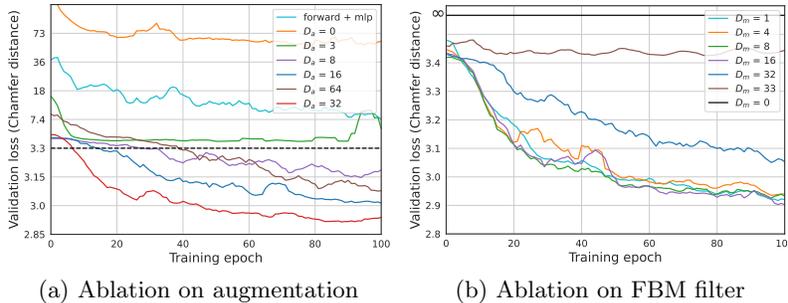


Fig. 8. (a) Loss curves of training PD-Flow with augmentation channels and LCC noise filter. For better visualization, we use linear mapping below dashed line and non-linear mapping above dashed line. (b) Loss curves of training PD-Flow with FBM noise filter. All these models are augmented with $D_a = 32$ and vary in mask channels D_m .

Table 4. Ablation study of different filtering strategies.

#Points, Noise	10K,1%		10K,2%		10K,3%	
Strategy	CD	P2M	CD	P2M	CD	P2M
FBM	2.22	0.44	3.40	1.14	4.56	2.16
LBM (w/o \mathcal{L}_{LBM})	2.54	0.62	3.65	1.38	5.06	2.32
LBM (w/ \mathcal{L}_{LBM})	2.37	0.53	3.50	1.23	4.80	2.23
LCC (w/o \mathcal{L}_{LCC})	2.32	0.45	3.48	1.19	4.62	2.20
LCC (w/ \mathcal{L}_{LCC})	2.13	0.38	3.26	1.03	4.50	2.09

of $\mathcal{L}_{\text{denoise}}$ term. The “LBM (w/o \mathcal{L}_{LBM})” term means that the experiment is trained without the \mathcal{L}_{LCC} loss (Eq. (12)), but still uses LBM as the filter strategy (Eq. (11)). As shown in Table 4, both \mathcal{L}_{FBM} and \mathcal{L}_{LCC} loss make positive contributions to model performance.

5 Conclusion

In this paper, we present PD-Flow, a point cloud denoising framework that combines NFs and distribution disentanglement techniques. It learns to transform noise perturbation and clean points into a disentangled latent code space by leveraging NFs, whereas denoising is formulated as channel masking. To alleviate the dimensional bandwidth bottleneck and improve the network expressiveness, we propose to extend additional channels to latent variables by an dimension augmentation module. Extensive experiments and ablation studies illustrate that our method outperforms existing state-of-the-art methods in terms of generation quality across various noise levels.

Acknowledgements: This work was supported by the NSF of Guangdong Province (2019A1515010833, 2022A1515011573), the Natural Science Foundation of China (61725204) and Tsinghua University Initiative Scientific Research Program, China Postdoctoral Science Foundation (2021M701891).

References

1. Abdal, R., Zhu, P., Mitra, N.J., Wonka, P.: Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. *ACM Transactions on Graphics (TOG)* **40**(3), 1–21 (2021)
2. Alexa, M., Behr, J., Cohen-Or, D., Fleishman, S., Levin, D., Silva, C.T.: Point set surfaces. In: *Proceedings of the Conference on Visualization '01*. p. 21–28. IEEE Computer Society, USA (2001)
3. Alexa, M., Behr, J., Cohen-Or, D., Fleishman, S., Levin, D., Silva, C.T.: Computing and rendering point set surfaces. *IEEE Transactions on visualization and computer graphics* **9**(1), 3–15 (2003)
4. Avron, H., Sharf, A., Greif, C., Cohen-Or, D.: ℓ_1 -sparse reconstruction of sharp point set surfaces. *ACM Transactions on Graphics (TOG)* **29**(5), 1–12 (2010)
5. Cazals, F., Pouget, M.: Estimating differential quantities using polynomial fitting of osculating jets. *Computer Aided Geometric Design* **22**(2), 121–146 (2005)
6. Chen, J., Lu, C., Chenli, B., Zhu, J., Tian, T.: Vflow: More expressive generative flows with variational data augmentation. In: *International Conference on Machine Learning*. pp. 1660–1669 (2020)
7. Chen, R.T., Rubanova, Y., Bettencourt, J., Duvenaud, D.: Neural ordinary differential equations. *arXiv preprint arXiv:1806.07366* (2018)
8. Deschaud, J.E., Goulette, F.: Point cloud non local denoising using local surface descriptor similarity. *IAPRS* **38**(3A), 109–114 (2010)
9. Digne, J., De Franchis, C.: The bilateral filter for point clouds. *Image Processing On Line* **7**, 278–287 (2017)
10. Dinh, L., Krueger, D., Bengio, Y.: Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516* (2014)
11. Dinh, L., Sohl-Dickstein, J., Bengio, S.: Density estimation using real NVP. In: *5th International Conference on Learning Representations, ICLR 2017* (2017)
12. Duan, C., Chen, S., Kovacevic, J.: 3d point cloud denoising via deep neural network based local surface estimation. In: *2019 IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2019 - Proceedings*. pp. 8553–8557 (May 2019)
13. Fleishman, S., Cohen-Or, D., Silva, C.T.: Robust moving least-squares fitting with sharp features. *ACM transactions on graphics (TOG)* **24**(3), 544–552 (2005)
14. Gao, X., Hu, W., Guo, Z.: Graph-based point cloud denoising. In: *2018 IEEE Fourth International Conference on Multimedia Big Data (BigMM)*. pp. 1–6 (2018)
15. Grathwohl, W., Chen, R.T.Q., Bettencourt, J., Sutskever, I., Duvenaud, D.: Ffjord: Free-form continuous dynamics for scalable reversible generative models. *International Conference on Learning Representations* (2019)
16. Guerrero, P., Kleiman, Y., Ovsjanikov, M., Mitra, N.J.: PCPNet: Learning local shape properties from raw point clouds. *Computer Graphics Forum* **37**(2), 75–85 (2018). <https://doi.org/10.1111/cgf.13343>
17. Hermosilla, P., Ritschel, T., Ropinski, T.: Total denoising: Unsupervised learning of 3d point cloud cleaning. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 52–60 (2019)
18. Hu, W., Gao, X., Cheung, G., Guo, Z.: Feature graph learning for 3d point cloud denoising. *IEEE Transactions on Signal Processing* **68**, 2841–2856 (2020)
19. Huang, H., Li, D., Zhang, H., Ascher, U., Cohen-Or, D.: Consolidation of unorganized point clouds for surface reconstruction. *ACM transactions on graphics (TOG)* **28**(5), 1–7 (2009)

20. Huang, H., Wu, S., Gong, M., Cohen-Or, D., Ascher, U., Zhang, H.: Edge-aware point set resampling. *ACM transactions on graphics (TOG)* **32**(1), 1–12 (2013)
21. Huhle, B., Schairer, T., Jenke, P., Straßer, W.: Robust non-local denoising of colored depth data. In: 2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops. pp. 1–7 (2008)
22. Jenke, P., Wand, M., Bokeloh, M., Schilling, A., Straßer, W.: Bayesian point cloud reconstruction. In: *Computer Graphics Forum*. vol. 25, pp. 379–388 (2006)
23. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: *European Conference on Computer Vision*. pp. 694–711 (2016)
24. Kalogerakis, E., Nowrouzezahrai, D., Simari, P., Singh, K.: Extracting lines of curvature from noisy point clouds. *Computer-Aided Design* **41**, 282–292 (2009)
25. Kingma, D.P., Dhariwal, P.: Glow: Generative flow with invertible 1x1 convolutions. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 31. Curran Associates, Inc. (2018)
26. Klokov, R., Boyer, E., Verbeek, J.: Discrete point flow networks for efficient point cloud generation. In: *Proceedings of the 16th European Conference on Computer Vision (ECCV)* (2020)
27. Kumar, M., Babaeizadeh, M., Erhan, D., Finn, C., Levine, S., Dinh, L., Kingma, D.: Videoflow: A flow-based generative model for video. *arXiv preprint arXiv:1903.01434* **2**(5) (2019)
28. Lipman, Y., Cohen-Or, D., Levin, D., Tal-Ezer, H.: Parameterization-free projection for geometry reconstruction. *ACM Transactions on Graphics (TOG)* **26**(3), 22–es (2007)
29. Liu, Y., Anwar, S., Qin, Z., Ji, P., Caldwell, S., Gedeon, T.: Disentangling noise from images: A flow-based image denoising neural network. *arXiv preprint arXiv:2105.04746* (2021)
30. Lu, X., Wu, S., Chen, H., Yeung, S.K., Chen, W., Zwicker, M.: Gpf: Gmm-inspired feature-preserving point set filtering. *IEEE transactions on visualization and computer graphics* **24**(8), 2315–2326 (2017)
31. Lugmayr, A., Danelljan, M., Van Gool, L., Timofte, R.: Srflo: Learning the super-resolution space with normalizing flow. In: *European Conference on Computer Vision*. pp. 715–732. Springer (2020)
32. Luo, S., Hu, W.: Differentiable manifold reconstruction for point cloud denoising. In: *Proceedings of the 28th ACM International Conference on Multimedia*. pp. 1330–1338 (2020)
33. Luo, S., Hu, W.: Score-based point cloud denoising. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 4583–4592 (October 2021)
34. Mattei, E., Castrodad, A.: Point cloud denoising via moving rpca. In: *Computer Graphics Forum*. vol. 36, pp. 123–137. Wiley Online Library (2017)
35. Narváez, E.A.L., Narváez, N.E.L.: Point cloud denoising using robust principal component analysis. In: *Proceedings of the First International Conference on Computer Graphics Theory and Applications (GRAPP)*. pp. 51–58 (2006)
36. Pauly, M., Gross, M.: Spectral processing of point-sampled geometry. In: *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. pp. 379–386 (2001)
37. Pistilli, F., Fracastoro, G., Valsesia, D., Magli, E.: Learning graph-convolutional representations for point cloud denoising. In: *The European Conference on Computer Vision (ECCV)* (2020)

38. Postels, J., Liu, M., Spezialetti, R., Van Gool, L., Tombari, F.: Go with the flows: Mixtures of normalizing flows for point cloud generation and reconstruction. arXiv preprint arXiv:2106.03135 (2021)
39. Prenger, R., Valle, R., Catanzaro, B.: Waveglow: A flow-based generative network for speech synthesis. In: ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 3617–3621 (2019)
40. Pumarola, A., Popov, S., Moreno-Noguer, F., Ferrari, V.: C-flow: Conditional generative flow models for images and 3d point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7949–7958 (2020)
41. Rakotosaona, M.J., La Barbera, V., Guerrero, P., Mitra, N.J., Ovsjanikov, M.: Pointcleannet: Learning to denoise and remove outliers from dense point clouds. In: Computer Graphics Forum. vol. 39, pp. 185–203. Wiley Online Library (2020)
42. Schall, O., Belyaev, A., Seidel, H.P.: Robust filtering of noisy scattered point data. In: Proceedings Eurographics/IEEE VGTC Symposium Point-Based Graphics, 2005. pp. 71–144 (2005)
43. Schoenenberger, Y., Paratte, J., Vandergheynst, P.: Graph-based denoising for time-varying point clouds. In: 2015 3DTV-Conference: The True Vision-Capture, Transmission and Display of 3D Video (3DTV-CON). pp. 1–4 (2015)
44. Serna, A., Marcotegui, B., Goulette, F., Deschaud, J.E.: Paris-rue-madame database: a 3d mobile laser scanner dataset for benchmarking urban detection, segmentation and classification methods. In: 4th International Conference on Pattern Recognition, Applications and Methods ICPRAM 2014 (2014)
45. Sun, Y., Schaefer, S., Wang, W.: Denoising point sets via l_0 minimization. Computer Aided Geometric Design **35**, 2–15 (2015)
46. Wang, R.f., Chen, W.z., Zhang, S.y., Zhang, Y., Ye, X.z.: Similarity-based denoising of point-sampled surfaces. Journal of Zhejiang University-Science A **9**(6), 807–815 (2008)
47. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1912–1920 (2015)
48. Yang, G., Huang, X., Hao, Z., Liu, M.Y., Belongie, S., Hariharan, B.: Pointflow: 3d point cloud generation with continuous normalizing flows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4541–4550 (2019)
49. Yu, L., Li, X., Fu, C.W., Cohen-Or, D., Heng, P.A.: Pu-net: Point cloud upsampling network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2790–2799 (2018)
50. Zeng, J., Cheung, G., Ng, M., Pang, J., Yang, C.: 3d point cloud denoising using graph laplacian regularization of a low dimensional manifold model. IEEE Transactions on Image Processing **29**, 3474–3489 (2019)
51. Zhang, D., Lu, X., Qin, H., He, Y.: Pointfilter: Point cloud filtering via encoder-decoder modeling. IEEE Transactions on Visualization and Computer Graphics (2020)
52. Zheng, Q., Sharf, A., Wan, G., Li, Y., Mitra, N.J., Cohen-Or, D., Chen, B.: Non-local scan consolidation for 3d urban scenes. ACM transactions on graphics (TOG) **29**(4), 94 (2010)