

PCT: LARGE-SCALE 3D POINT CLOUD REPRESENTATIONS VIA GRAPH INCEPTION NETWORKS WITH APPLICATIONS TO AUTONOMOUS DRIVING

Siheng Chen, Sufeng Niu, Tian Lan, Baoan Liu

ABSTRACT

We present a novel graph-neural-network-based system to effectively represent large-scale 3D point clouds with the applications to autonomous driving. Many previous works studied the representations of 3D point clouds based on two approaches, voxelization, which causes discretization errors and learning, which is hard to capture huge variations in large-scale scenarios. In this work, we combine voxelization and learning: we discretize the 3D space into voxels and propose novel graph inception networks to represent 3D points in each voxel. This combination makes the system avoid discretization errors and work for large-scale scenarios. The entire system for large-scale 3D point clouds acts like the blocked discrete cosine transform for 2D images; we thus call it the point cloud neural transform (PCT). We further apply the proposed PCT to represent real-time LiDAR sweeps produced by self-driving cars and the PCT with graph inception networks significantly outperforms its competitors.

Index Terms— 3D point cloud representations, graph deep neural networks, autonomous driving

1. INTRODUCTION

With the growth of 3D sensing technologies, one can now use a large number of 3D points to precisely represent objects' surfaces and surrounding environments. We call those 3D points a *3D point cloud*; it has a growing impact on various applications, including autonomous driving, virtual reality and scanning of historical artifacts [1]. In this paper, we consider the setting of autonomous driving. A self-driving car could use multiple sensors to observe the world, such as LiDARs, cameras and RADARs [2]. Among those, LiDARs produce two types of 3D point clouds, real-time LiDAR sweeps and high-precision maps. Both include accurate range information, which are critical to perception and localization systems. We consider both types of point clouds *large-scale point clouds* because they contain a large number of 3D points and record outdoor environments.

To fully exploit real-time LiDAR sweeps and high-precision maps, we need advanced techniques to handle a series of challenges, including 3D point cloud compression, 3D localization and 3D object detection. A common task shared in those challenges is *3D point cloud representations*;

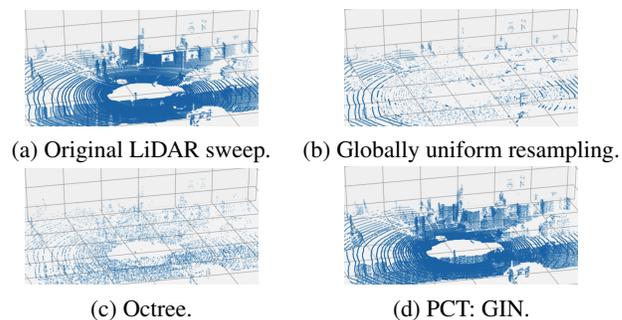


Fig. 1: Comparison between an original sweep and its reconstructions. All reconstructions use 2.78% of original data.

that is, representing a 3D point cloud in a compact format, such that it is easy to conduct subsequent processing procedures. For 1D time-series, the basic representation is the Fourier transform; for 2D images, it is the discrete cosine transform and 2D wavelet transform [3]; for 3D point clouds, Octree partitions the 3D space adaptively and has been an effective representation tool [4]; however, Octree represents a 3D point cloud only in the 3D spatial domain and does not fully exploit shapes formed by 3D points.

To represent large-scale 3D point clouds, we propose a novel graph-neural-network-based system, called the *point cloud neural transform* (PCT). The PCT includes two phases: voxelization, which adopts the standard Octree-like partition and splits a large-scale space into a series of small-scale spaces (voxels), and voxel-level encoding, which adopts graph neural networks to capture the complicated underlying distributions of 3D points within each voxel. In the phase of voxel-level encoding, we propose novel graph inception networks, which transform 3D points in a voxel to a low-dimensional feature vector and extend to the translation-invariant graph convolution [5] to the 3D space. Because of the two-phase design, PCT can be considered as the 3D counterpart of the windowed Fourier transform for 1D time series and the blocked discrete cosine transform for 2D images. We further apply the PCT to represent real-time LiDAR sweeps collected by self-driving cars and compare it with the standard Octree-based representations. As a general tool, the PCT can be potentially used to 3D compression, 3D object detection and many others in autonomous driving. The

main contributions of the paper include: (i) we propose the PCT to represent large-scale 3D point clouds; (ii) we propose novel graph inception networks to implement the voxel-level encoding; (iii) the proposed PCT is applied to represent real-time LiDAR sweeps produced by self-driving cars and outperforms its competitors.

Related works. 3D point cloud processing has become important in various 3D imaging and autonomous systems. The topic broadly includes compression [6, 7, 8], denoising [9, 10], surface reconstruction [11], feature extraction [12], localization [13], 3D object detection [14] and many others. Here we consider 3D point cloud representations, which mainly based on three approaches, including resampling, voxelization and learning. Resampling represents a 3D point cloud by selecting a subset of 3D points [15]. In many applications, resampling can enhance key geometry information and make subsequent processing both cheaper and more accurate [16]. Voxelization represents a 3D point cloud by partitioning the 3D space into a series of voxels and use the corresponding voxel center as the proxy for each 3D point, such as regular voxels [17] and Octree [18]; however, they all suffer from discretization errors. Learning-based models can also be trained to capture the underlying distribution of 3D points. For example, [19] proposed a probabilistic generative model to model the distribution of 3D point clouds; however, such model is inefficient in inferring parameters; [20] proposed a deep autoencoder that directly handles 3D point clouds; [21, 22] introduced a 2D lattice to help decoding.

2. METHODOLOGY

Similarly to many standard representation problems, the overall goal is to use a low-dimensional feature vector to represent a large-scale 3D point cloud; however, a large-scale 3D point cloud has its own challenges: (i) variations. 3D points captured in an outdoor environment have huge variations, while the available training data are limited. The representations need to learn rich variations from limited amount of 3D point clouds; (ii) irregularity. 3D points are irregularly and sparsely scattered in the 3D space. The representations need to go beyond the regular lattices and capture irregular and nonuniform distributions of 3D points; (iii) invariances and equivalences. The representations need to promote basic geometric properties, including permutation-invariance, translation-invariance, scale-equivalent and rotation-equivalence. We propose the point cloud neural transform (PCT) to handle these challenges.

Point cloud neural transform. To handle large variations, we propose a two-stage framework. In the first stage, we partition the 3D space into a series of voxels; in the second stage, we encode 3D points in each voxel to a few codes; see the entire system in Figure 2. The intuitions are (i) since the representations in each individual voxel is much more constraint, we are able to specifically and effectively learn local

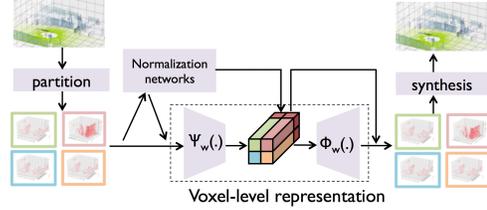


Fig. 2: Proposed point cloud neural transform (PCT).

shapes and patterns; and (ii) since each 3D point cloud produces multiple voxels, we are able to increase the number of training data and potentially train strong voxel-level representations. This two-stage representation is similar to the windowed Fourier transform for 1D time series and the blocked discrete cosine transform for 2D images. To handle irregularity, we propose graph inception networks, which use a graph to capture the underlying distribution of 3D points. We treat 3D points as nodes and connects each point to its neighboring points to formulate a spatial graph. The edge weights reflect irregular and nonuniform distribution of 3D points. The graph inception networks also ensure permutation and translation-invariances. To handle equivalences, we propose normalization networks to ensure scale and rotation-equivalences.

Mathematically, let $P \in \mathbb{R}^{N \times 3}$ be the matrix representation of this 3D point cloud, whose i th row $\mathbf{p}_i = [x_i, y_i, z_i] \in \mathbb{R}^3$ represents the 3D coordinate of the i th point. The overall procedures of the PCT are

$$\{P^{(i)}\}_{i=1}^M = \text{partition}(P), \quad (1a)$$

$$\mathbf{c}^{(i)} = \Psi_w(P^{(i)}) \in \mathbb{R}^\ell, \quad (1b)$$

$$\widehat{P}^{(i)} = \Phi_w(\mathbf{c}^{(i)}) \quad (1c)$$

$$\widehat{P} = \text{synthesis}\left(\{\widehat{P}^{(i)}\}_{i=1}^M\right). \quad (1d)$$

In (1a), we partition a 3D space into a series of nonoverlapping voxels based on the spatial structure; correspondingly, a large-scale 3D point cloud is partitioned into a series of small-scale 3D point clouds, where $P^{(i)}$ represents the 3D point cloud in the i th voxel. In (1b), we encode the 3D points in each voxel to a low-dimensional feature vector. In (1c), we decode a low-dimensional feature vector back to the 3D coordinates. In (1d), we concatenate the 3D coordinates in all the voxels and reconstruct a 3D point cloud.

Partition. We simply partition the 3D space into equally-spaced nonoverlapping voxels from each of three dimensions. Let each voxel is of size H, W, D along the X, Y, Z axes respectively. The (h, w, d) th voxel represents a 3D space,

$$\mathcal{V}_{h,w,d} = \{(x, y, z) \mid \begin{aligned} (h-1)H &\leq x < hH, \\ (w-1)W &\leq y < wW, \\ (d-1)D &\leq z < dD. \end{aligned}\} \quad (2)$$

The points inside $\mathcal{V}_{h,w,d}$ form a point cloud $P_{h,w,d}$. The partition (2) contributes to the implement of (1a) and (1d). The PCT is also compatible to multiscale voxels, such as Octree.

Graph inception networks as voxel-level encoder. We use the graph inception networks (GIN) to implement the voxel-level encoding (1b). The operations and weights of GIN are shared across all the voxels. To handle the irregularity, we introduce a spatial graph to capture the distribution of 3D points. For the simplicity, here we consider a K -nearest-neighbor graph, where each 3D points connects to its K closest 3D points. We denote the neighboring set for of i -th point as $\mathcal{N}_i^{(K)}$. We then use an adjacency matrix $A^{(K)} \in \mathbb{R}^{n \times n}$ to reflect the pairwise connectivity, where n is the number of 3D points in the corresponding voxel. The edge weight between two points \mathbf{p}_i and \mathbf{p}_j is

$$A_{i,j}^{(K)} = \begin{cases} e^{-\|\mathbf{p}_i - \mathbf{p}_j\|_2^2}, & j \in \mathcal{N}_i^{(K)}; \\ 0, & \text{otherwise.} \end{cases}$$

A key issue raised by a graph-based approach is how to choose the number of neighbors K . Especially, the point density could vary a lot in various voxels. A fixed K could either capture limited information or irrelevant information. To solve this issue, we adopt an inception-like structure, where we construct a series of K -nearest-neighbor graphs with multiple K values. This K value is equivalent to the kernel size of the classical 2D convolution: a larger K indicates a larger reception field. We thus consider a graph inception convolution to extract features from 3D points. Let $\mathbf{x}_i \in \mathbb{R}^\ell$ be the features of the i -th point (the initial feature is $\mathbf{p}_i \in \mathbb{R}^3$). The response is

$$\begin{aligned} \mathbf{x}'_i &= g_w \left(\left[\mathbf{y}_i^{(k)} \right]_{k \in \mathcal{K}} \right) \quad (3) \\ &= g_w \left(\left[\sum_{j \in \mathcal{N}_i^{(k)}} h_w \left(A_{i,j}^{(k)}, \mathbf{x}_j - \mathbf{x}_i \right) \right]_{k \in \mathcal{K}} \right) \in \mathbb{R}^{\ell'}, \quad (4) \end{aligned}$$

where $h_w(\cdot)$ is a standard multilayer perceptron with parameters w , \mathcal{K} is a set of K values, $[\cdot]_{k \in \mathcal{K}}$ denotes the concatenation and $g_w(\cdot)$ is the inception network that combine responses from multiple graph convolutions. The edge weight $A_{i,j}^{(K)}$ and the difference $\mathbf{x}_j - \mathbf{x}_i$ reflects the relative difference of two points in the original 3D space and the feature space, respectively. The corresponding matrix representation is

$$X' = \text{conv}_{\text{gin}} \left(\left[A^{(k)} \right]_{k \in \mathcal{K}}, X \right) \in \mathbb{R}^{n \times \ell'}, \quad (5)$$

where $X \in \mathbb{R}^{n \times \ell}$ is the input feature matrix with the i th row vector \mathbf{x}_i in (3) and X' is the output feature matrix with the i th row vector \mathbf{x}'_i . Since we only consider relative differences, the graph convolution is translation-invariant; that is,

$$\text{conv}_{\text{gin}}(X + \mathbf{1}_n \mathbf{a}^T) = \text{conv}_{\text{gin}}(X),$$

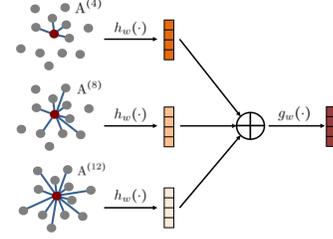


Fig. 3: Proposed graph inception convolution: $\text{conv}_{\text{gin}}(\cdot)$.

holds for arbitrary $\mathbf{a} \in \mathbb{R}^\ell$, where $\mathbf{1}_n \in \mathbb{R}^n$ is a all-one vector. We call (5) *graph inception convolution*; see Figure 3.

After a series of graph inception convolution, we obtain deep point-wise features. To produce voxel-level features, we average the features across all the points. Let $X \in \mathbb{R}^{n \times \ell}$ be the final point-wise feature matrix, $\mathbf{c} = \text{agg}_{\text{mean}}(X) = X^T \mathbf{1}_n / n \in \mathbb{R}^\ell$. Since we aggregate along the point dimension, the final voxel-level features are permutation invariant. Comparing to the max-aggregation [20], the mean-aggregation allows all points contributes to the code, which preserves richer information for reconstruction.

The voxel-level encoder $\Psi_w(\cdot)$ is thus a combination of graph inception convolutions and a mean aggregation. For example, the encoder of the i th voxel with a single layer graph inception convolution is

$$\mathbf{c}^{(i)} = \Psi_w(P^{(i)}) = \text{agg}_{\text{mean}} \left(\rho \left(\text{conv}_{\text{gin}} \left(\left[A^{(k)} \right]_{k \in \mathcal{K}}, P^{(i)} \right) \right) \right),$$

where $\rho(\cdot)$ is the nonlinear activation, such as ReLU. In our experiments, we use 3-layer graph inception convolutions.

Graph inception convolution is a 3D convolution. Here we consider graph inception convolution from a different perspective. We can represent each 3D point as a delta function in the 3D space. A 3D point cloud is then a train of delta functions; that is,

$$\mathbf{s}(\mathbf{p}) = \sum_j \delta(\mathbf{p} - \mathbf{p}_j),$$

where $\mathbf{p}, \mathbf{p}_j \in \mathbb{R}^3$. Let a 3D convolution be $h(\cdot) : \mathbb{R}^3 \rightarrow \mathbb{R}$. The response is then

$$\begin{aligned} y(\mathbf{p}) &= \int_{\tau} h(\mathbf{p} - \tau) \mathbf{s}(\tau) d\tau \\ &= \int_{\tau} h(\mathbf{p} - \tau) \sum_j \delta(\tau - \mathbf{p}_j) d\tau \\ &= \sum_j h(\mathbf{p} - \mathbf{p}_j). \end{aligned}$$

This is equivalent to the graph inception convolution in (3). In other words, we aim to learn a kernel function $h(\cdot)$ through neural networks. The number of neighbors in the graph reflects the size of the reception field of a kernel function. This

kernel function is learnt in the 3D continuous space and operates in the graph domain. The nature of a 3D convolution indicates the properties of translation invariance and weight sharing in the 3D space; it also ensures that the same kernel function works for various graph topologies. Once the neighboring points form a same shape pattern, we will construct the same local graph topology and obtain the same response.

Fully connected layers as voxel-level decoder. To design a voxel-level decoder $\Phi_w(\cdot)$ in (1c), we cannot use graph-based approaches, because the only information that decoder can access is the code given by the encoder and the graph information no longer exists. We consider two approaches to design a voxel-level decoder. The first approach is based on fully-connected layers, which use more trainable parameters and work better in practice [20]; the second approach is based on FoldingNet [21, 22], which considers that points are warped from a 2D map.

Normalization networks. To handle equivalences, the voxel-level representation also needs to capture the scale and rotation variances. Here we use normalization networks to explicitly learn the scale and rotation. Before we feed the points to the voxel-level encoder, we use multilayer perceptions to learn the 3×1 scale vector and the 3×3 rotation matrix. We keep those geometric information to the code and apply to the reconstructed points after the decoder. In this way, we promote scale and rotation equivalences: the 3D point cloud in each voxel is distributed in the unit space and has a similar orientation. The voxel-level encoder and decoder can thus focus on learning local shapes.

Training implementations. The encoder $\Phi_w(\cdot)$ and decoder $\Psi_w(\cdot)$ is implemented by using deep-neural-networks with trainable weights. To train the networks, we consider the following optimization problem:

$$O(\ell) = \underset{\Psi_w(\cdot), \Phi_w(\cdot)}{\text{minimize}} \sum_{i=1}^M \text{CD}(\mathbf{P}^{(i)}, \widehat{\mathbf{P}}^{(i)}),$$

$$\text{subject to} \quad \mathbf{c} = \Psi_w(\mathbf{P}), \widehat{\mathbf{P}} = \Phi_w(\mathbf{c}),$$

$$\dim(\mathbf{c}) \leq \ell,$$

where $\text{CD}(\mathbf{P}, \widehat{\mathbf{P}}) = \sum_{j=1}^N \min_{i \in \{1, 2, \dots, N\}} \|\widehat{\mathbf{p}}_j - \mathbf{p}_i\|_2^2 / N + \sum_{i=1}^N \min_{j \in \{1, 2, \dots, N\}} \|\mathbf{p}_i - \widehat{\mathbf{p}}_j\|_2^2 / N$ is the Chamfer distance [20]. An effective representation should be $\ell M \ll 3N$, $O(\ell) \rightarrow 0$.

3. EXPERIMENTAL RESULTS

Dataset. We validate the proposed PCT in a standard autonomous-driving dataset, KITTI [2], which has been recorded from a moving platform while driving in and around Karlsruhe. Real-time LiDAR sweeps are collected by a Velodyne HDL-64E rotating 3D laser scanner, with 10 Hz, 64 beams, 0.09 degree angular resolution, around 1.3 million points/second, 360 horizontal, 26.8 vertical field of view.

Experimental Setup. For each real-time sweep in KITTI, we partition the space into voxels with the size of $1 \times 1 \times$

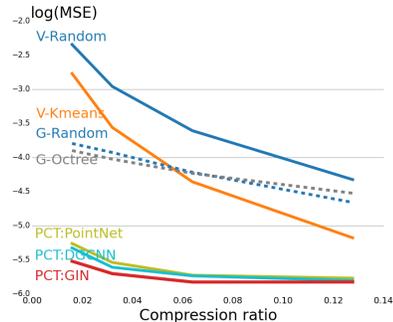


Fig. 4: Mean square errors as a function of the compression ratio in the dataset of KITTI.

10 meter³. We train 400 LiDAR sweeps and test 100 LiDAR sweeps. We select training sweeps and testing sweeps from separate logsets to avoid data snooping. For GIN, $\mathcal{K} = \{1, 4, 8, 16\}$. To evaluate the performance, we compare the reconstruction based on the codes to the original LiDAR sweep.

Results. We validate the proposed PCT from two aspects. In the ablation study, we add each component at a time to validate the effectiveness of each component; in the ratio-distortion analysis, we compare the PCT with other competitors. Due to the limited space, we only show the quantitative results for KITTI. Table 1 shows the ablation study in the dataset of KITTI. For the encoder, we consider three components: graph inception convolution (inception), mean-aggregation (mean-agg), normalization (norm). When the graph inception convolution is not checked, we consider a single $K = 1$; when the mean-aggregation is not checked, we consider the maximum-aggregation; when the normalization is not checked, we do not use normalization networks. For the decoder, we consider two components: fully-connected layers and FoldingNet. We use five metrics to evaluate the performance: earth-mover distance (EMD), Chamfer distance (CD), mean square error (MSE), mean and variance [20]. For MSE, we consider the difference between each original 3D point and its closest correspondence in the reconstruction. For all the metrics, lower values indicate better results. We see that each component of graph inception convolution improves the reconstruction performance; fully-connected layers consistently outperform FoldingNet.

Figure 4 shows the mean square error as a function of compression ratios. We vary the code length in each voxel as 9, 18, 36, 72, leading to various compression ratios. The x -axis is the compression ratio; the y -axis is the logarithm-scale mean square error. We consider give comparison methods. Uniform resampling randomly selects a few 3D points and use this subset to represent the overall 3D point clouds. We consider resampling based on either the entire 3D spatial space (G-Random) or the voxels (V-Random). For voxel-based resampling, we select the same number of 3D points in

Encoder: GIN			Decoder		Metrics				
Inception	Mean-Agg	Norm	FC	FoldingNet	EMD	CD	Mean	Variance	MSE
			✓		2.947×10^1	2.689×10^{-2}	4.531×10^{-2}	1.832×10^{-3}	3.933×10^{-3}
✓			✓		2.552×10^1	2.342×10^{-2}	4.220×10^{-2}	1.842×10^{-3}	3.668×10^{-3}
✓				✓	2.707×10^1	2.489×10^{-2}	4.354×10^{-2}	1.930×10^{-3}	3.871×10^{-3}
✓	✓		✓		2.429×10^1	2.296×10^{-2}	4.164×10^{-2}	1.716×10^{-3}	3.489×10^{-3}
✓	✓			✓	2.606×10^1	2.464×10^{-2}	4.321×10^{-2}	1.792×10^{-3}	3.699×10^{-3}
✓	✓	✓	✓		2.237×10^1	1.88×10^{-2}	3.982×10^{-2}	1.706×10^{-3}	3.335×10^{-3}
✓	✓	✓		✓	2.328×10^1	1.936×10^{-2}	4.051×10^{-2}	1.784×10^{-3}	3.468×10^{-3}

Table 1: Ablation study in the dataset of KITTI. The code length is 18, corresponding to the compression ratio 3.19%.

each voxel to promote spatial uniformity; kmeans-based representation (V-Kmeans) adaptively selects cluster centers in each voxel and we select the same number of 3D points in each voxel. Kmeans is computationally expensive and cannot afford global optimization. These three approaches preserve information based on the spatial domain. We also consider two learning-based approaches to implement the voxel-level encoding in the framework of PCT. PointNet encodes the 3D points in each voxel into a few features by using deep neural networks [20]; dynamic graph convolutional neural networks (DGCNN) are the extension of PointNet by introducing a graph structure [23]. Compared to DGCNN, GIN ensures translation invariance and adopts inception structures, mean-aggregation and the normalization networks. We ensure all methods preserves the same number of data from the original sweeps. We see that PCT with GIN outperforms its competitors. The advantage of the PCT is to transform the 3D points from the spatial domain to a feature domain, which is similar to the mechanism of the classical Fourier transform.

4. CONCLUSIONS

We propose the PCT to provide compact representations for large-scale 3D point clouds. The PCT includes two phases: 3D partition and voxel-level representations, which makes it acts like the blocked discrete cosine transform for 2D images. We propose GIN to improve voxel-level representations. The proposed PCT is applied to represent real-time LiDAR sweeps and significantly outperforms its competitors.

5. REFERENCES

- [1] R. B. Rusu and S. Cousins, “3D is here: Point cloud library (PCL),” in *Proc. IEEE Int. Conf. Robot. Autom.*, Shanghai, May 2011.
- [2] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? the kitti vision benchmark suite,” in *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recogn.*, Providence, RI, June 2012.
- [3] M. Vetterli, J. Kovačević, and V. K. Goyal, *Foundations of Signal Processing*, Cambridge University Press, Cambridge, 2014, <http://foundationsofsignalprocessing.org>.
- [4] R. Schnabel and R. Klein, “Octree-based point-cloud compression,” in *SPBG’06 Proceedings of the 3rd Eurographics / IEEE VGTC conference on Point-Based Graphics*, Boston, MA, July 2006.
- [5] S. Niu, S. Chen, H. Guo, C. Targonski, M. C. Smith, and J. Kovačević, “Generalized value iteration networks: Life beyond lattices,” in *AAAI*, Feb. 2018.
- [6] S. Schwarz, M. Preda, V. Baroncini, M. Budagavi, P. César, P. A. Chou, R. A. Cohen, M. Krivokuca, S. Lasserre, Z. Li, J. Llach, K. Mammou, R. Mekuria, O. Nakagami, E. Siahaan, A. J. Tabatabai, A. M. Tourapis, and V. Zakharchenko, “Emerging MPEG standards for point cloud compression,” *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 9, no. 1, pp. 133–148, 2019.
- [7] D. Thanou, P. A. Chou, and P. Frossard, “Graph-based compression of dynamic 3D point cloud sequences,” *IEEE Trans. Image Process.*, vol. 25, no. 4, pp. 1765–1778, Feb. 2016.
- [8] A. Anis, P. A. Chou, and A. Ortega, “Compression of dynamic 3D point clouds using subdivisional meshes and graph wavelet transforms,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Shanghai, Mar. 2016, pp. 6360–6364.
- [9] C. Duan, S. Chen, and J. Kovačević, “Weighted multi-projection: 3d point cloud denoising with estimated tangent planes,” in *Proc. IEEE Glob. Conf. Signal Information Process.*, Anaheim, CA, Nov. 2018.
- [10] J. Zeng, G. Cheung, M. Ng, and C. Yang J. Pang, “3d point cloud denoising using graph laplacian regularization of a low dimensional manifold model,” *arXiv preprint arXiv:1803.07252*, 2018.

- [11] B. F. Gregorski, B. Hamann, and K. I. Joy, "Reconstruction of B-spline surfaces from scattered data points," in *Proc. Comput. Graphics Int'l*, Geneva, June 2000, pp. 163–170.
- [12] C. Feng, Y. Taguchi, and V. Kamat, "Fast plane extraction in organized point clouds using agglomerative hierarchical clustering," in *Proc. IEEE Int. Conf. Robot. Autom.*, Hong Kong, May 2014, pp. 6218–6225.
- [13] R. Spangenberg, D. Goehring, and R. Rojas, "Pole-based localization for autonomous vehicles in urban scenarios," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Dec. 2016.
- [14] M. Liang, B. Yang, S. Wang, and R. Urtasun, "Deep continuous fusion for multi-sensor 3d object detection," in *The European Conference on Computer Vision (ECCV)*, Sept. 2018.
- [15] S. Chen, D. Tian, C. Feng, A. Vetro, and J. Kovačević, "Fast resampling of 3d point clouds via graphs," *IEEE Trans. Signal Process.*, vol. 66, no. 3, pp. 666–681, Feb. 2018.
- [16] N. Gelfand, L. Ikemoto, S. Rusinkiewicz, and M. Levoy, "Geometrically stable sampling for the ICP algorithm," in *Fourth International Conference on 3D Digital Imaging and Modeling (3DIM)*, Oct. 2003.
- [17] C. T. Loop, C. Zhang, and Z. Zhang, "Real-time high-resolution sparse voxelization with application to image-based modeling," in *Proc. High-Perform. Graphics*, Anaheim, CA, July 2013, pp. 73–80.
- [18] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: An efficient probabilistic 3D mapping framework based on octrees," *Autonom. Robots*, pp. 189–206, Apr. 2013.
- [19] B. Eckart, K. Kim, A. Troccoli, A. Kelly, and J. Kautz, "Accelerated generative models for 3D point cloud data," in *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recogn.*, Las Vegas, NV, June 2016.
- [20] P. Achlioptas, O. Diamanti, I. Mitliagkas, and L. Guibas, "Representation learning and adversarial generation of 3d point clouds," arXiv:1707.02392., June 2017.
- [21] Y. Yang, C. Feng, Y. Shen, and D. Tian, "Foldingnet: Point cloud auto-encoder via deep grid deformation," in *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recogn.*, June 2018.
- [22] T. Groueix, M. Fisher, V. Kim, B. Russell, and M. Aubry, "Atlasnet: A papier-mâché approach to learning 3d surface generation," in *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recogn.*, June 2018.
- [23] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *arXiv preprint arXiv:1801.07829*, 2018.