



**HAL**  
open science

## **Denoising Diffusion Probabilistic Model for Point Cloud Compression at Low Bit-Rates**

Gabriele Spadaro, Alberto Presta, Jhony Heriberto Giraldo Zuluaga, Marco Grangetto, Wei Hu, Giuseppe Valenzise, Attilio Fiandrotti, Enzo Tartaglione

► **To cite this version:**

Gabriele Spadaro, Alberto Presta, Jhony Heriberto Giraldo Zuluaga, Marco Grangetto, Wei Hu, et al.. Denoising Diffusion Probabilistic Model for Point Cloud Compression at Low Bit-Rates. IEEE International Conference on Multimedia and Expo (ICME 2025), Jun 2025, Nantes, France. ⟨hal-05160415⟩

**HAL Id: hal-05160415**

**<https://centralesupelec.hal.science/hal-05160415v1>**

Submitted on 14 Jul 2025

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

# Denoising Diffusion Probabilistic Model for Point Cloud Compression at Low Bit-Rates

Gabriele Spadaro<sup>1,2</sup>, Alberto Presta<sup>2</sup>, Jhony H. Giraldo<sup>1</sup>, Marco Grangetto<sup>2</sup>, Wei Hu<sup>3</sup>,  
Giuseppe Valenzise<sup>4</sup>, Attilio Fiandrotti<sup>1,2</sup>, Enzo Tartaglione<sup>1</sup>

<sup>1</sup>LTCI, Télécom Paris, Institut Polytechnique de Paris, France <sup>2</sup>University of Turin, Italy

<sup>3</sup>Wangxuan Institute of Computer Technology, Peking University

<sup>4</sup>Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des Signaux et Systèmes

**Abstract**—Efficient compression of low-bit-rate point clouds is critical for bandwidth-constrained applications. However, existing techniques mainly focus on high-fidelity reconstruction, requiring many bits for compression. This paper proposes a “Denoising Diffusion Probabilistic Model” (DDPM) architecture for point cloud compression (DDPM-PCC) at low bit-rates. A PointNet encoder produces the condition vector for the generation, which is then quantized via a learnable vector quantizer. This configuration allows to achieve a low bitrates while preserving quality. Experiments on ShapeNet and ModelNet40 show improved rate-distortion at low rates compared to standardized and state-of-the-art approaches. We publicly released the code at <https://github.com/EIDOSLAB/DDPM-PCC>.

**Index Terms**—Geometry Point Cloud Compression, Denoising Diffusion Probabilistic Models, Learnable Vector Quantizer.

## I. INTRODUCTION

Recent advances in 3D imaging have driven interest in Point Clouds Compression (PCC) technologies. Namely, the complexity of storing and delivering point cloud geometry poses significant challenges due to their spatial complexity. For this reason, traditional codecs based on effective data structure have been proposed. The MPEG Geometry-based Point Cloud Compression (G-PCC) [1] standard, for instance, aims at achieving high-quality compression of static, highly detailed point clouds, leveraging grid-based and spatial subdivision techniques. Similarly, Google released DRACO [2], an open-source library for point cloud compression.

The recent success of neural models in image and video compression [3], [4] prompted learning-based approaches for PCC. In [5], the authors introduced a pair of specific analysis-synthesis convolutional transforms for PCC compression. In [6] the authors proposed a PointNet-based [7] encoder, while [8] adopted a PointNet++ [9] module. More recently, attention-based point cloud deconvolution modules have been proposed to improve reconstruction by addressing localized artifacts from linear interpolation [8]. Li *et al.* [10], instead, introduced an optimal transport system to learn the distribution transformation from original to reconstructed point clouds, while He *et al.* [11] highlighted that maintaining local density is crucial for achieving faithful reconstruction, proposing a novel deep learning-based autoencoder designed to maintain this local property. Moreover, hybrid solutions were also proposed to synergistically employ deep learning and octree-based representations [12]–[14], or following a voxel-based

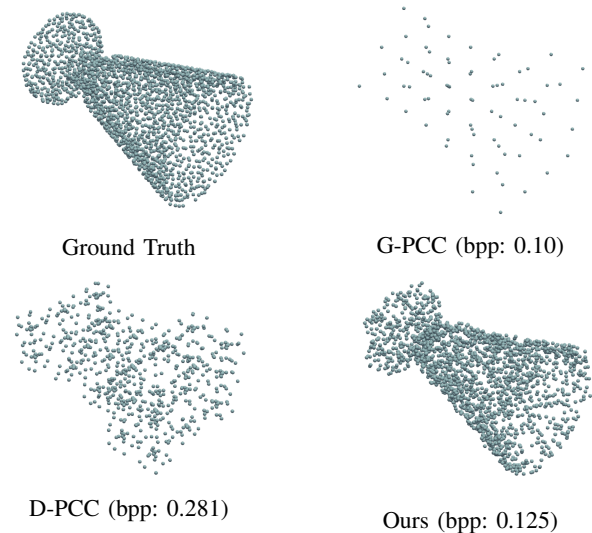


Fig. 1: Low bit-rates PCC and compression artifacts.

approach [5], [15]–[17], where a point cloud is represented as a structured grid of volumetric cells (voxels). All the above methods target high fidelity reconstruction, struggling to preserve quality at bandwidths practically available on real-world channels (as visualized in Fig. 1), motivating the present work.

When working at lower bit-rates, a recent yet effective approach to perform compression is towards generative models like diffusion models [18]. Seminal works like [19] exploit a conditional diffusion model for image compression. Here, the reverse diffusion process is conditioned using a latent representation extracted via analysis transform. In [20], a stable diffusion model was conditioned on text embeddings extracted from a pre-trained image captioning model and transmitted to the decoder. Here, a pre-trained latent diffusion model (LDM) was used to encode the image into a latent representation and then quantized via a hyperprior autoencoder. Although effective, these approaches are designed for images with a fixed spatial structure, which is fundamentally different from the irregular and unstructured nature of point clouds. To the best of our knowledge, no attempts have been made to employ such methods specifically for point cloud data.

In this work, we propose a Denoising Diffusion Probabilistic Model architecture for PCC (DDPM-PCC). Specifically, we

use PointNet [7] to map the initial point cloud to a latent representation for initial compression, followed by vector quantization. The resulting quantized representation is then passed to the decoder, which conditions a diffusion process to produce the final reconstruction.

Our contributions can be summarized as follows:

- To the best of our knowledge, we are the first to employ a diffusion-based generative model for geometry point cloud compression, detailed in Sec. III-A.
- We adopted a compression solution based on a learnable vector quantizer to compress the latent space (Sec. III-B).
- Our experimental results, presented in Sec. IV-B, show that our method improves the encoding efficiency at a low bit rate with respect to other state-of-the-art approaches.

## II. BACKGROUND ON DDPM

Denoising Diffusion Probabilistic Models (DDPM) are a class of generative methods adopting a diffusion process to generate data. This process is defined as a Markov chain and allows the generation of high-fidelity samples of complex distributions. In a nutshell, given a sample from the target distribution  $\mathbf{x}^{(0)}$ , the *diffusion* process gradually adds noise to  $\mathbf{x}^{(0)}$ . The *reverse* process, indeed, involves a neural network to remove noise for generating a sample. More formally, the diffusion process (or forward process) is defined as

$$q(\mathbf{x}^{(1:T)}|\mathbf{x}^{(0)}) = \prod_{t=1}^T q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}), \quad (1)$$

where

$$q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) = \mathcal{N}\left(\mathbf{x}^{(t)}; \sqrt{1 - \beta_t}\mathbf{x}^{(t-1)}, \beta_t\mathbf{I}\right) \quad (2)$$

having  $\mathbf{x}^{(t)}$  sampled from a Gaussian distribution with mean  $\sqrt{1 - \beta_t}\mathbf{x}^{(t-1)}$  and variance  $\beta_t\mathbf{I}$ , where  $\beta_t$  is the variance scheduler and  $\mathbf{I}$  is the identity matrix. For generating a latent sample  $\mathbf{x}^{(t)}$ , this process would require to perform  $t - 1$  steps in the Markov chain. To avoid this, following [18], we can express  $\mathbf{x}^{(t)}$  directly in terms of  $\mathbf{x}^{(0)}$ :

$$q(\mathbf{x}^{(t)}|\mathbf{x}^{(0)}) = \mathcal{N}\left(\mathbf{x}^{(t)}; \sqrt{\bar{\alpha}_t}\mathbf{x}^{(0)}, (1 - \bar{\alpha}_t)\mathbf{I}\right), \quad (3)$$

where  $\alpha_t = 1 - \beta_t$  and  $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ .

For large  $T$ , the forward process converges to an isotropic Gaussian distribution centered in zero. The reverse process samples from this distribution and, adopting a neural network parameterized with  $\theta$ , gradually denoises the Gaussian noise into a real sample from the target distribution. Specifically, this process is defined as:

$$p_\theta(\mathbf{x}^{(0:T)}) = p(\mathbf{x}^{(T)}) \prod_{t=1}^T p_\theta(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}), \quad (4)$$

where

$$p_\theta(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) = \mathcal{N}(\mathbf{x}^{(t-1)}; \mu_\theta(\mathbf{x}^{(t)}, t), \beta_t\mathbf{I}), \quad (5)$$

and  $\mu_\theta$  is a network that predicts the mean of the Gaussian distribution for  $\mathbf{x}^{(t-1)}$ .

Ho *et al.* [18] proposed a tractable way for training the reverse process. In particular, they minimized the Mean Squared Error between the noise added in the forward process and the noise predicted by the model:

$$\begin{aligned} \mathcal{L}_{\text{Diff}} &= \mathbb{E}_{t, \mathbf{x}^{(0)}, \epsilon} [\|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}\mathbf{x}^{(0)} + \sqrt{1 - \bar{\alpha}_t}\epsilon, t)\|_2^2] \\ &= \mathbb{E}_{t, \mathbf{x}^{(0)}, \epsilon} [\|\epsilon - \epsilon_\theta(\mathbf{x}^{(t)}, t)\|_2^2], \end{aligned} \quad (6)$$

where  $t \sim U(1, T)$ ,  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$  is the noise and  $\epsilon_\theta(\cdot)$  is the trainable network that predicts the added noise.

**Point Cloud Generation.** Lou *et al.* [21] proposed a DDPM-based model for 3D Point Cloud Generation. More specifically, they considered a point cloud  $\mathbf{X}^{(0)} = \{\mathbf{x}_i^{(0)}\}_{i=1}^P$  as a set of  $P$  points in an evolving system. Each of these points  $\mathbf{x}_i$  can then be treated as independently sampled from a point distribution. For this reason, the forward diffusion process and the reverse diffusion one (detailed in (1) and (4) respectively) are directly applied independently to these points.

In our work, we show how such modeling can be leveraged towards the distinct problem of compressing point clouds at very low bit rates.

## III. PROPOSED METHOD

In the previous section, we provided some background knowledge regarding Denoising Diffusion Probabilistic Models (DDPM). Here, we explain how this class of generative models can be employed as an autoencoder trained for a compression task. Next, we discuss the compression strategy employed to quantize the latent space and finally, we present the overall proposed framework.

### A. Diffusion-based Point Cloud Compression Model

To perform compression, our generative diffusion model needs to be conditioned during the reverse process to generate a specific sample. Moreover, each sample from the data distribution is represented as point cloud  $\mathbf{X}^{(0)} \in \mathbb{R}^{P \times 3}$ .

Leveraging the point cloud auto-encoder proposed by [21] we generate the original point cloud from a random Gaussian distribution in a point-to-point way. Specifically, to obtain certain sample  $\mathbf{X}^{(0)}$ , the generation is guided with a latent representation of the sample  $\mathbf{z} = E_\varphi(\mathbf{X}^{(0)})$ , where in our case  $E_\varphi(\cdot)$  is a PointNet-based [7] encoder having parameters  $\varphi$ . Next, we adopt the reverse diffusion process as a decoder, conditioned on the latent representation  $\mathbf{z}$ . To do this, we can define  $p_\theta(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}, \mathbf{z})$  in (4) and rewrite the Loss function (6) training the auto-encoder to minimize:

$$\mathcal{L}_{\text{Diff}} = \mathbb{E}_{t, \mathbf{X}^{(0)}, \epsilon} [\|\epsilon - \epsilon_\theta(\mathbf{x}^{(t)}, t, E_\varphi(\mathbf{X}^{(0)}))\|_2^2]. \quad (7)$$

At this point, by appropriately quantizing and then compressing the latent representation  $\mathbf{z}$ , we obtain a model able to compress and then generate the input point cloud. In the next section, we describe how we compress this latent code.

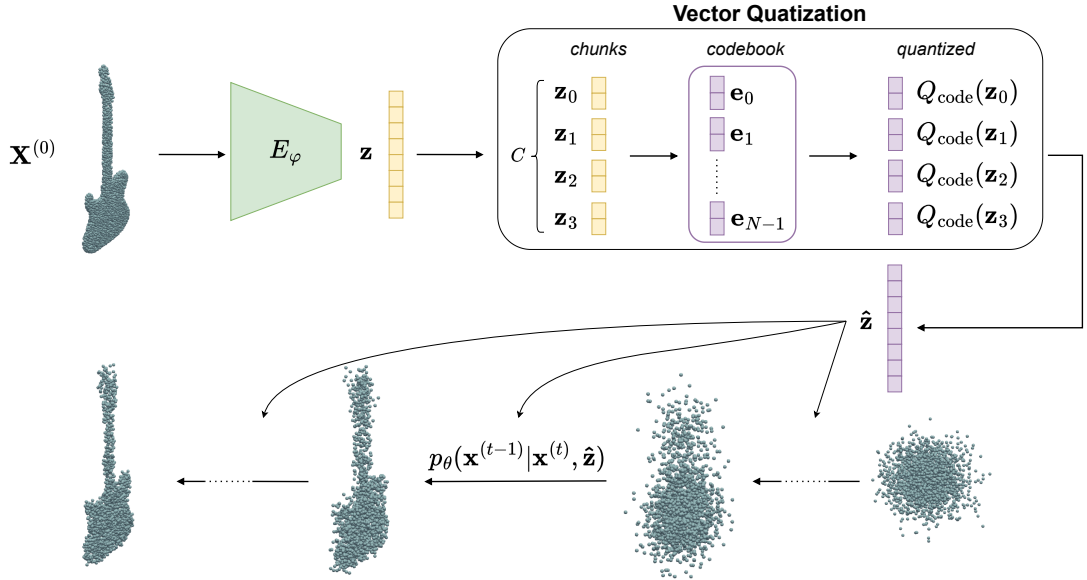


Fig. 2: Overview of the proposed method. Here for simplicity,  $\mathbf{x}^{(t-1)}$  represents a single point, while the diffusion process shows the final reconstructed point cloud (conditioned by the quantized latent representation  $\hat{\mathbf{z}}$ ). Here  $\mathbf{z}$  is split into  $C = 4$  chunks, making  $\hat{\mathbf{z}}$  representable with only 4 codebook indices.

### B. Encoding the Latent Representation

The input point cloud  $\mathbf{X}^{(0)}$  is encoded into the latent representation vector  $\mathbf{z} \in \mathbb{R}^d$ , where  $d$  is the hidden dimension. Inspired by VQ-VAE [22], we adopt a vector quantization strategy based on a learnable codebook to efficiently compress our latent code. Specifically, to control the rate, we divide  $\mathbf{z}$  into  $C$  chunks:  $\mathbf{z} = \{\mathbf{z}_0, \dots, \mathbf{z}_{C-1}\}$ , having the same dimension  $\mathbf{z}_i \in \mathbb{R}^{d/C}$ . We choose  $C$  as an integer multiple of  $d$ . Then, the  $i$ -th chunk  $\mathbf{z}_i$  is mapped to a vector in the learned codebook minimizing the vector quantization loss:

$$\mathcal{L}_{\text{VQ}} = \mathbb{E}_{\mathbf{z}_i} [\|\text{sg}(\mathbf{z}_i) - \mathbf{e}_q\|_2^2 + \|\text{sg}(\mathbf{e}_q) - \mathbf{z}_i\|_2^2], \quad (8)$$

where  $\text{sg}(\cdot)$  stands for the stopgradient operator that is defined having zero partial derivatives, while  $\mathbf{e}_q$  is the closest codebook entry to  $\mathbf{z}_i$ , defined as  $\mathbf{e}_q = Q_{\text{code}}(\mathbf{z}_i)$ . This loss ensures that the codebook entries match the encoder's output and, at the same time, forces the encoder to produce latent vectors close to the codebook entries.

The quantized version of the latent code  $\hat{\mathbf{z}}$  is given by mapping each of its chunks with the closest entries in the codebook  $\hat{\mathbf{z}} = \{Q_{\text{code}}(\mathbf{z}_0), \dots, Q_{\text{code}}(\mathbf{z}_{C-1})\}$ .

At this point, the rate of  $\hat{\mathbf{z}}$  can be easily defined in terms of  $C$  and the codebook dimension. In particular,  $\hat{\mathbf{z}}$  can be represented by specifying the indexes of  $C$  entries in the codebook. The number of bits per index depends on the size of the codebook. For instance, if the codebook has a size of  $N = 256$ , each index requires 8 bits for representation. Therefore, the final rate can be computed as  $C \times \log_2 N$ .

### C. Overall Framework

In our framework, the conditional diffusion model and the learnable vector quantizer are jointly trained minimizing:

$$\mathcal{L} = \mathcal{L}_{\text{diff}} + \mathcal{L}_{\text{VQ}}. \quad (9)$$

Notably, the  $\mathcal{L}_{\text{diff}}$  adopted here differs from (6) since  $E_\varphi(\mathbf{X}^{(0)}) = \mathbf{z}$  is replaced with  $\hat{\mathbf{z}}$ . This joint optimization enables better integration of the quantization process with the generative model.

Fig. 2 shows an illustrative overview of our compression pipeline once the model is trained. Here, the point cloud to be compressed is first passed through the encoder  $E_\varphi(\cdot)$ , producing a latent representation  $\mathbf{z}$ . This vector is then divided into  $C$  chunks, and each chunk is quantized using the learned codebook to produce  $\hat{\mathbf{z}}$ . The latter is used to condition the reverse denoising process of the diffusion model to reconstruct the input point cloud, which thus functions as a true decoder.

As discussed in Sec. III-B, the final rate depends on the number of chunks  $C$  and the codebook size  $N$ . For simplicity, we set  $N = 128$  constant. Then, to change the coding rate, we can use different  $C$  values from 1 up to  $d$  (where  $\mathbf{z} \in \mathbb{R}^d$ ). Intuitively, with  $C = d$ , each value in  $\mathbf{z}$  is independently quantized, leading smaller quantization error. On the contrary, with  $C = 1$ , the entire vector is quantized with a single index, resulting in only  $N$  possible reconstructions. It is important to notice that each codebook has to be constructed based on the chunk dimension:  $d/C$ . For this reason, given  $C$  and thus the target rate, the corresponding codebook is trained with the entire model and then used for coding at that rate.

## IV. EXPERIMENTS

In this section, we detail the experimental setup and the results obtained with our compression method. Lastly, we

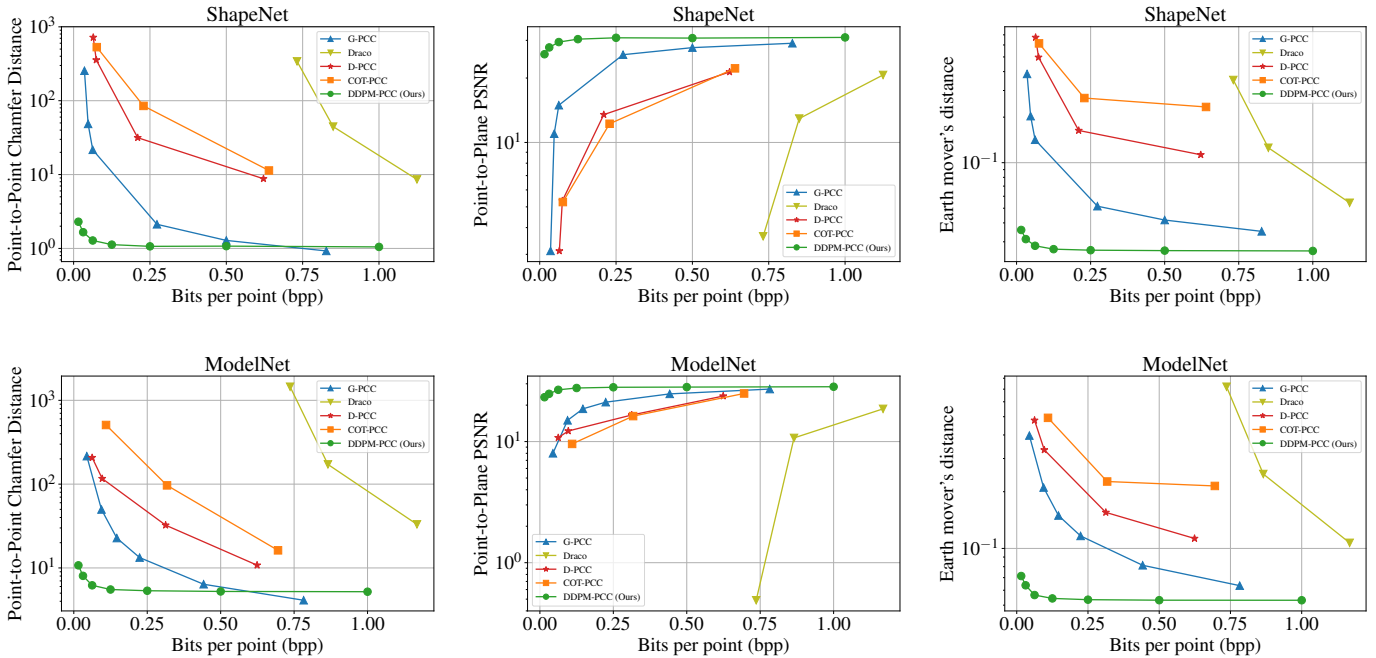


Fig. 3: Quantitative results on ShapeNet (first row) and ModelNet40 (second row). The y-axes are plotted in logarithmic scale. Our method achieves the best geometry reconstruction at a very low bit rate according to all quality metrics.

present two ablation studies on key design choices: the compression method and the encoder architecture.

#### A. Experimental Setup

**Datasets.** We conduct our experiments on ShapeNet [23] and ModelNet40 [24]. ShapeNet [23] contains a total of 51,127 shapes, divided in 55 categories. This dataset was randomly split to obtain 43,433 samples for training and 7,694 samples for testing. We sample 2,048 points from each shape to acquire the point clouds and we normalize each of them to zero mean and unit variance. ModelNet40 [24] is a dataset containing 12,308 meshes CAD model from 40 categories. From this dataset, we adopted 9,843 samples for training and the remaining 2,468 for testing. Also in such case, we sample 2,048 points from each mesh and we normalized the coordinates in the same manner.

**Training Details.** Following the same training procedure adopted in [21], we trained our compression model for around 1,000,000 steps, using a batch of 128 samples for each step. We used Adam as optimizer [25] with an initial learning rate of  $1 \times 10^{-4}$  and a final learning rate of  $1 \times 10^{-5}$ , decreased linearly from step  $200k$  to  $400k$ . Regarding the model, we choose a dimension  $d = 256$  for the encoder’s output. As implemented in [21], we used  $T = 200$  denoising steps and a linear variance schedule from  $\beta_1 = 0.0001$  to  $\beta_T = 0.05$ .

As detailed in Sec. III-B, we obtained a pool of models to compress the input point cloud at different bit-rates by using chunks numbers  $C$  belonging to  $\{4, 8, 16, 32, 64, 128, 256\}$  (lowering the  $C$  lowers the rate). For each of these  $C$  values, a specific codebook is jointly trained with the autoencoder to quantize chunks of dimension  $d/C$ . For simplicity, we fix the

codebook size  $N = 128$ .

**Evaluation.** We compare our method with both non-learning-based techniques: G-PCC (release-v23.0) [1], Google Draco [2]; and learning-based models: D-PCC [11], COT-PCC [10]. For these latter models, we re-trained them on the same datasets as our method, using  $\lambda$  values in the Rate-Distortion Loss function to achieve a low bit-rate compression.

Following [10], [11], we evaluate the accuracy in geometry reconstruction in terms of symmetric point-to-point Chamfer Distance (CD) and point-to-plane PSNR. We also compare the methods by leveraging the Earth Mover’s distance (EMD), while the compression rate is measured in Bits per Point (bpp).

#### B. Rate-Distortion Comparison

**Quantitative Results.** We first provide quantitative results by comparing our method with previous state-of-the-art approaches. In Fig. 3 we show the Chamfer Distance (CD), the PSNR, and Earth Mover’s distance (EMD) of all methods against Bits per point (bpp). To highlight differences among the methods, the y-axes are plotted using a logarithmic scale.

The obtained results show that our method achieves superior performance compared to the other analyzed approaches on both datasets and according to all quality metrics. In particular, the advantages of our method are particularly pronounced at low bit rates, where it provides significantly better reconstruction quality. Draco [2], instead, requires a large number of bits to achieve acceptable reconstruction. Although G-PCC [1] shows good capabilities in preserving the overall geometry, it suffers from a loss of finer details due to the voxelization process. Finally, the learnable methods [10], [11] do not converge to satisfactory results for the compression

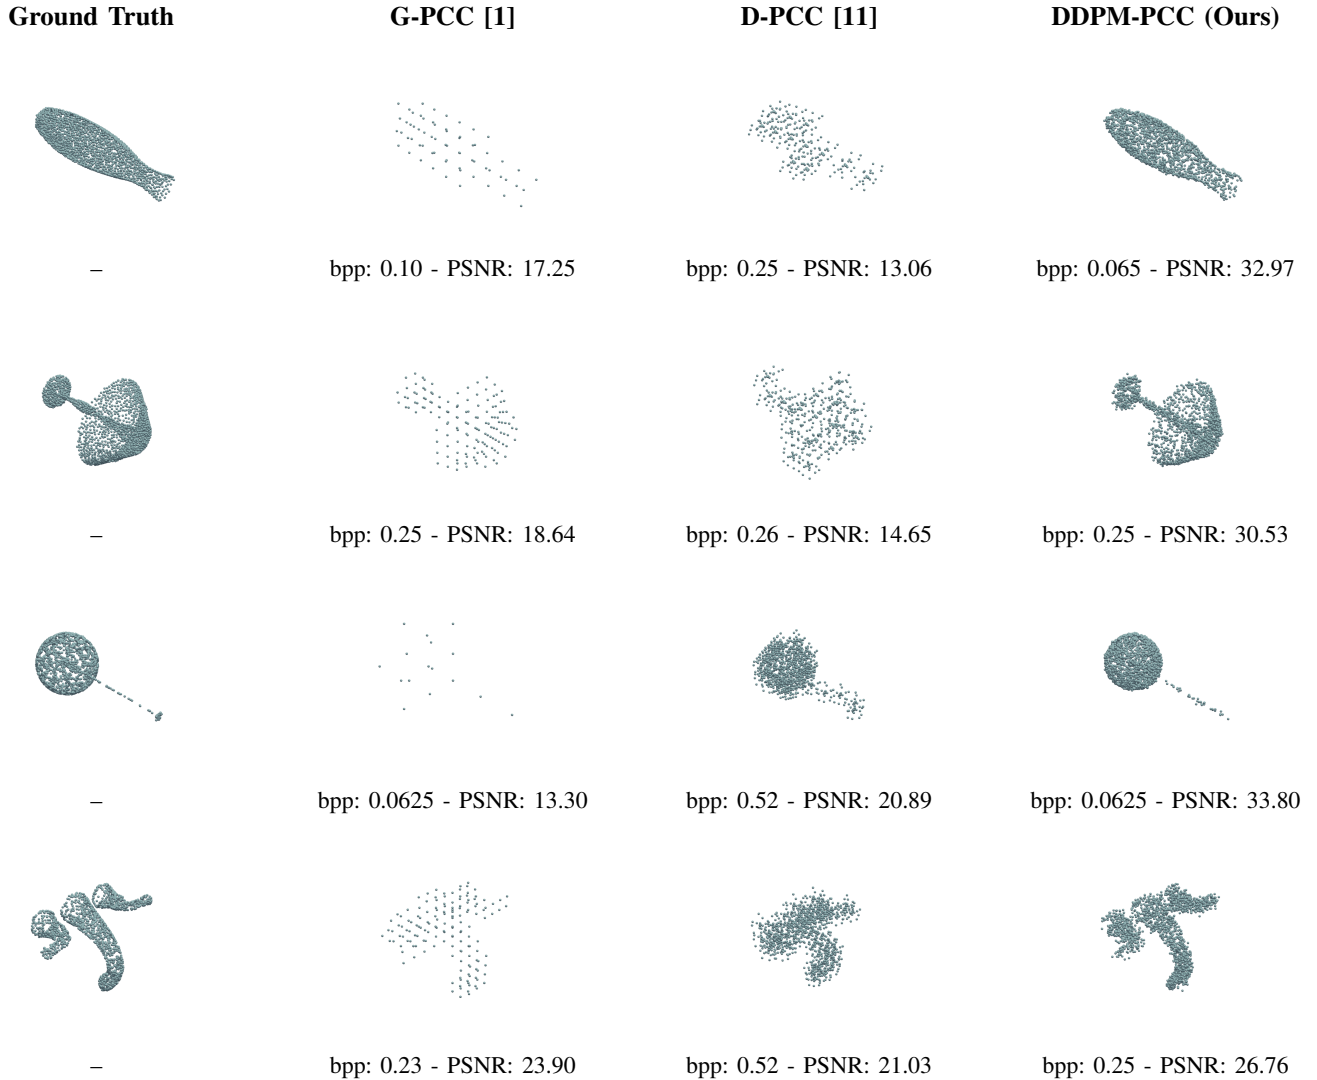


Fig. 4: Qualitative results on ModelNet (first two rows) and ShapeNet (last two rows). We compare the reconstruction results of different compression methods providing the bpp and the PSNR values for each analyzed sample.

levels analyzed, as they were designed to operate at higher bit-rates (typically in the range of 1 to 6 bpp).

Despite these advantages, our method is limited by the performance of the generative model. In fact, even by increasing the coding rate, the results obtained cannot outperform the quality of reconstruction achieved by the generative model.

**Qualitative Results.** To clearly show the reconstruction ability of our compression model, even at a very low bit-rate (around 0.065 bpp), in Fig. 4 we report some qualitative results at various bit-rates. From these reconstructions, we can notice the limitations encountered by previous techniques in such a context. Here, G-PCC [1] shows the effect of the voxelization and quantization needed to compress at these rates. This leads to the loss of many local details during reconstruction. D-PCC [11], indeed, consistently requires more bits for low-quality reconstructions. On the other hand, our model can achieve satisfactory PSNR results even for very low bpp in all reconstructions analyzed.

### C. Ablation Studies

**Compression Strategy.** In Fig. 5 we compare our chosen quantization technique, which exploits learnable vector quantizer (VQ), with a more traditional end-to-end compression approach [3], where the latent space is modeled with a fully factorized (FF) distribution estimated by an auxiliary neural network. In the latter, the compression is achieved by balancing (with a  $\lambda$  parameter) the reconstruction and the rate of the latent representations in the well-known rate-distortion (RD) loss. From this figure we notice that our method achieves better compression trade-offs, obtaining closer results to the model that does not implement any type of compression. Moreover, our approach does not involve the use of a  $\lambda$  parameter to adjust the importance of reconstruction versus rate. Indeed, as reported in Fig. 5, by changing the value of  $C$  we can easily tailor the model to a specific coding rate.

**Encoder Architecture.** Finally, we replace PointNet with PointNet++ [9] concerning the encoder architecture for the

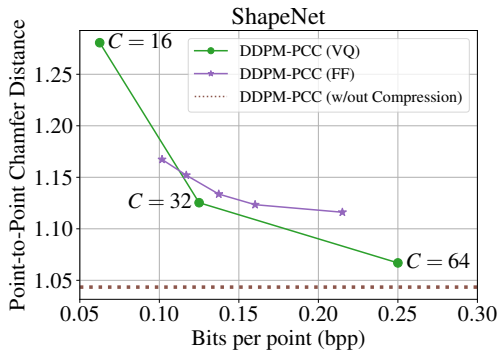


Fig. 5: Comparison between a learnable vector quantize approach (VQ) and a fully factorize-based strategy (FF) to encode the latent representation. The dashed line represents the lower bound in Chamfer Distance: the generative model. Different values of  $C$  allow for different RD trade-offs

Encoder	Chamfer-Distance ( $\downarrow$ )	Enc. (s)	Dec. (s)
PointNet [7]	1.043	$0.001 \pm 0.001$	$0.254 \pm 0.0102$
PointNet++ [9]	<b>0.953</b>	$0.134 \pm 0.004$	$0.259 \pm 0.012$

TABLE I: Ablation study concerning the encoder architecture. Encodings and Decodings times are in seconds and evaluated on NVIDIA GeForce RTX 3090 GPU.

generative model (*i.e.*, without quantization). In Tab I we notice that PointNet++ [9], since it leverages the geometry of the input point cloud, can result in a richer latent representation. This increases the performance of the model in terms of CD, but at the cost of increasing coding time by  $\times 100$ . PointNet [7], indeed, considers each point independently. This results in a less informative representation but obtained in a much faster time.

## V. CONCLUSIONS & LIMITATIONS

In this work we proposed DDPM-PCC, a simple yet effective model to compress the point cloud geometry at very low bit rates. Our compression model consists of a PointNet [7] encoder used to generate a latent representation. This vector is then quantized with a learnable quantizer strategy and used to condition the diffusion reverse process to generate the input point cloud. The compression rate is then controlled during the quantization, dividing the latent code in a certain number of chunks and quantizing each of these sub-vectors independently. This allows us to compress the input at very low bit rates. Despite its simplicity, the generation quality is upper-bounded by the performance of the generative model. A possible solution to increase the fidelity of the reconstruction at higher bit rates could be to provide side information regarding the geometry of the input point cloud to guide the generation.

**Acknowledgements.** This research was partially funded by Hi!PARIS Center on Data Analytics and Artificial Intelligence. This project was provided with computer and storage resources by GENCI at IDRIS thanks to the grant 2024-AD011015338 on the supercomputer Jean Zay.

## REFERENCES

- [1] D Graziosi, O Nakagami, S Kuma, A Zaghetto, T Suzuki, and A Tabatabai, "An overview of ongoing point cloud compression standardization activities: video-based (v-pcc) and geometry-based (g-pcc)," *APSIPA Transactions on Signal and Information Processing*, 2020.
- [2] Frank Galligan, Michael Hemmer, Ondrej Stava, Fan Zhang, and Jamieson Brettle, "Google/draco: a library for compressing and decompressing 3d geometric meshes and point clouds," <https://github.com/google/draco>, 2018.
- [3] Johannes Ballé et al., "End-to-end optimized image compression," in *ICLR*, 2017.
- [4] Jinming Liu, Heming Sun, and Jiro Katto, "Learned image compression with mixed transformer-cnn architectures," in *CVPR*, 2023.
- [5] Maurice Quach, Giuseppe Valenzise, and Frederic Dufaux, "Learning convolutional transforms for lossy point cloud geometry compression," in *ICIP*. IEEE, 2019.
- [6] Wei Yan, Shan Liu, Thomas H Li, Zhu Li, Ge Li, et al., "Deep autoencoder-based lossy geometry compression for point clouds," *arXiv preprint arXiv:1905.03691*, 2019.
- [7] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *CVPR*, 2017.
- [8] Wenxiang Shen, BaoYe Zhang, Hao Xu, XiaoHan Li, and Jun Wu, "Multi-space point geometry compression with progressive relation-aware transformer," *IEEE TMM*, 2024.
- [9] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," *NeurIPS*, 2017.
- [10] Zezeng Li, Weimin Wang, Ziliang Wang, and Na Lei, "Point cloud compression via constrained optimal transport," in *ICME*, 2024.
- [11] Yun He, Xinlin Ren, Danhang Tang, Yinda Zhang, Xiangyang Xue, and Yanwei Fu, "Density-preserving deep point cloud compression," in *CVPR*, 2022.
- [12] Zizheng Que, Guo Lu, and Dong Xu, "Voxelcontext-net: An octree based framework for point cloud compression," in *CVPR*, 2021.
- [13] Chunyang Fu, Ge Li, Rui Song, Wei Gao, and Shan Liu, "Octattention: Octree-based large-scale contexts model for point cloud compression," in *AAAI*, 2022.
- [14] Vicente Morell, Sergio Orts, Miguel Cazorla, and Jose Garcia-Rodriguez, "Geometric 3d point cloud compression," *Pattern Recognition Letters*, 2014.
- [15] Maurice Quach, Giuseppe Valenzise, and Frederic Dufaux, "Improved deep point cloud geometry compression," in *IEEE MMS-W*. IEEE, 2020.
- [16] Jianqiang Wang, Hao Zhu, Haojie Liu, and Zhan Ma, "Lossy point cloud geometry compression via end-to-end learning," *IEEE TCSVT*, 2021.
- [17] Jianqiang Wang, Dandan Ding, Zhu Li, and Zhan Ma, "Multiscale point cloud geometry compression," in *DCC*, 2021.
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel, "Denoising diffusion probabilistic models," in *NeurIPS*, 2020.
- [19] Ruihan Yang and Stephan Mandt, "Lossy image compression with conditional diffusion models," *NeurIPS*, 2023.
- [20] Marlene Careil, Matthew J Muckley, Jakob Verbeek, and Stéphane Lathuilière, "Towards image compression with perfect realism at ultra-low bitrates," in *ICLR*, 2023.
- [21] Shitong Luo and Wei Hu, "Diffusion probabilistic models for 3d point cloud generation," in *CVPR*, 2021.
- [22] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu, "Neural discrete representation learning," in *NeurIPS*, 2017.
- [23] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu, "ShapeNet: An Information-Rich 3D Model Repository," Tech. Rep. arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015.
- [24] Zhirong Wu, Shuran Song, Aditya Khosla, Xiaoou Tang, and Jianxiong Xiao, "3d shapenets for 2.5d object recognition and next-best-view prediction," in *CVPR*, 2015.
- [25] Diederik P Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2015.